From the holdings of the Department of Special Collections , Princeton University Library, Princeton, NJ

These documents are either in the public domain, or can only be used for educational and research purposes (Fair use) as per U.S. Copyright law (text below). By accessing this file, all users agree that their use falls in the public domain or within fair use as defined by the copyrightlaw. We do not charge any permission or usage fees for the publication of images of material inour collections, including those provided by Princeton University Library, via our website, catalog, or directly from staff.For more information regarding copyright, credit, and citation please review the information onthe Copyright, Credit and Citation Guidelines page of our website.
https://library.princeton.edu/special-collections/policies/copyright-credit-and-citation-guidelines. If you have any questions, including how to order higher-quality images from the collection, contact Special Collections Public Services via our Ask Us! Form. https://library.princeton.edu/special-collections/ask-us

---

Doctoral Dissertations, 2011-2026 Mathematics

2022

Invariants from Equivariant Transversality in Symplectic Topology and Some Results on the Rouquier Dimension of Wrapped Fukaya Categories

Bai, Shaoyun

Princeton, NJ : Princeton University

---

http://arks.princeton.edu/ark:/88435/dsp01jw827f828

U.S. Copyright law: The copyright law of the United States (Title 17, §108, United States Code) governs the making of photocopies or other reproductions of copyrighted material. Under certain conditions specified in the law, libraries and archives are authorized to furnish a photocopy or other reproduction. One of these specified conditions is that the photocopy or reproduction is not to be used for any purpose other than private study, scholarship, or research. If a user makes a request for, or later uses, a photocopy or reproduction for purposes in excess of fair use that user may be liable for copyright infringement. This institution reserves the right to refuse to accept a copying order if, in its judgment, fulfillment of the order would involve a violation of copyright law. https://www.copyright.gov/title17/title17.pdf

Invariants from equivariant transversality in symplectic topology and some results on the Rouquier dimension of wrapped Fukaya categories

Shaoyun Bai

A Dissertation
Presented to the Faculty
of Princeton University
in Candidacy for the Degree
of Doctor of Philosophy

Recommended for Acceptance
by the Department of
Mathematics
Advisor: Professor John Vincent Pardon

May 2022

© Copyright by Shaoyun Bai, 2022.

All Rights Reserved

Abstract

In this thesis, we construct several invariants in low-dimensional topology and symplectic topology, including a symplectic definition of generalized Casson invariants, an extension of Taubes' Gromov invariants to symplectic Calabi–Yau 3-folds, and $\mathbb{Z}$-valued genus 0 Gromov–Witten type invariants for general compact symplectic manifolds, based on solutions to various equivariant transversality problems in symplectic topology. In a different direction, we also study the Rouquier dimension of wrapped Fukaya categories of Liouville manifolds to obtain applications in algebraic geometry and symplectic geometry.

Acknowledgements

First of all, I would like to thank my advisor John Pardon for his constant support in the past four and half years. This thesis would not come into the current shape without the freedom I got from John's hands-off advising style or his patience in listening to my two-hour lectures on his blackboard every other week. Moreover, as is manifest in the body part, the fundamental tools developed by John are crucial for crystallizing many ideas in my work.

Secondly, I am very grateful to my collaborators Laurent Côté, Mohan Swaminathan, Guangbo Xu, and Boyu Zhang for teaching me different aspects of symplectic topology and gauge theory through our collaborations. I have benefited much from their insights and mastery of various techniques.

Thirdly, I would like to mention that my limited understanding of mathematical research and ideas for future directions are largely based on studying the works of and the conversations with the following mathematicians: most notably, Mohammed Abouzaid, Mark McLean, Paul Seidel, and Chris Wendl; additionally, Daniel Alvarez-Gavela, Sheel Ganatra, Helmut Hofer, Eleny Ionel, Yusuf Baris Kartal, Cheuk-Yu Mak, Ciprian Manolescu, Kyler Siegel, Ivan Smith, and Zhengyi Zhou.

Fourthly, I warmly thank my fellow graduate students in Fine Hall for making my years at Princeton a far more interesting experience, including Eric Chen, Weibo Fu, Shilin Lai, Yangyang Li, Thomas Masssoni, Lue Pan, Rohil Prasad, Lingfu Zhang, and Ziquan Zhuang.

Finally, I wish to thank Fanying Chen for joining the journey near the end.

To my parents.

Contents

Abstract iii

Acknowledgements iv

List of Tables xi

List of Figures xii

1 Introduction **1**

2 A symplectic formula of generalized Casson invariants **4**

2.1 Introduction 4

2.2 Recap of 3–dimensional gauge theory 9

2.2.1 Some linear algebra 10

2.2.2 Some 3–dimensional gauge theory 12

2.3 Extended moduli spaces 19

2.4 Equivariant transversality of Lagrangian submanifolds 26

2.4.1 Transversality 26

2.4.2 An equivariant Maslov index 33

2.5 The gauge-symplectic correspondence 39

2.5.1 Perturbations 39

2.5.2 Spectral flow comparison 44

2.5.3 An adiabatic limit result 54

2.5.4 From Maslov index to spectral flow and back 61

3 Bifurcations of embedded curves and towards an extension of Taubes' Gromov invariant to Calabi–Yau 3-folds **66**

3.1 Introduction 66

3.1.1 Statement of the main Theorem 67

3.1.2 Ingredients of the proof 70

3.1.3 Relations with other work 74

3.1.4 Technical remarks concerning transversality arguments 78

3.1.5 Structure of the chapter 80

3.2 Review of embedded curves and super-rigidity 81

3.2.1 Cauchy–Riemann operators and Galois covers 82

3.2.2 Super-rigidity 87

3.3 A necessary condition for bifurcations 91

3.3.1 Outline of the argument 94

3.3.2 Preparation 95

3.3.3 Thickened moduli space of covers 98

3.3.4 Thickened moduli space of stable maps 101

3.3.5 Concluding the proof of Theorem 3.3.1 104

3.4 Bifurcation analysis 105

3.4.1 Simplifications 112

3.4.2 Preparation 114

3.4.3 Implicit function theorem 120

3.4.4 Analysis of the Kuranishi map 123

3.5 An extension of Taubes' Gromov invariant to dimension 6 138

3.5.1 Linear wall crossing 139

3.5.2 Proof of Theorem 3.1.5 156

3.6 Application to the Gopakumar–Vafa invariants 164

3.6.1 Obstruction bundles and their Euler classes 164

3.6.2 Comparision to Gopakumar–Vafa invariants 174

4 On the Rouquier dimension of wrapped Fukaya categories and a conjecture of Orlov 176

4.1 Introduction 176

4.1.1 Overview 176

4.1.2 An application of symplectic topology to algebraic geometry . 177

4.1.3 Quantitative problems in symplectic topology 183

4.2 Category theory 191

4.2.1 Notation and conventions 191

4.2.2 Triangulated categories 192

4.2.3 A_∞ categories 194

4.2.4 The Rouquier dimension 197

4.2.5 Homotopy colimits 199

4.2.6 The Rouquier dimension of some triangulated categories . . . 200

4.3 Wrapped Fukaya categories 203

4.3.1 Liouville manifolds and related objects 203

4.3.2 Invariants of Liouville manifolds 205

4.3.3 Structural facts about wrapped Fukaya categories 207

4.3.4 The Rouquier dimension of wrapped Fukaya categories 210

4.4 Symplectic flexibility and Orlov's conjecture 212

4.4.1 Stratifications and Goresky's triangulation theorem 214

4.4.2 Arboreal singularities 216

4.4.3 Upper bounds on the Rouquier dimensions 221

4.4.4 Applications to Orlov's conjecture 226

4.4.5 A digression on non-commutative motives 230

4.5 Central actions on the wrapped Fukaya category 232

4.5.1 Central actions on triangulated categories 232

4.5.2 Central actions and the Rouquier dimension 235
4.5.3 Central actions on Fukaya categories 242
4.5.4 Computations via string topology 244
4.6 Applications to symplectic geometry 250
4.6.1 Embedding monotonicity 250
4.6.2 Lefschetz fibrations 252
4.6.3 Rigidity of skeleta 255
4.7 Appendix 256

5 An integral Euler cycle in normally complex orbifolds and $\mathbb{Z}$-valued Gromov–Witten type invariants 264

5.1 Introduction 264
5.1.1 Statement of main results 265
5.1.2 A brief discussion of the proof 271
5.1.3 Relations with other work 276
5.1.4 Plan of the chapter 278
5.1.5 Acknowledgements 278
5.2 A simple proof of Abouzaid–McLean–Smith's splitting theorem 279
5.3 Normally complex orbifolds and bundles 284
5.3.1 Orbifolds and orbifold vector bundles 284
5.3.2 The isotropy prestratification 287
5.3.3 Riemannian metric 289
5.3.4 Connections 298
5.3.5 Normal complex structures 300
5.3.6 Normally complex sections 302
5.4 Whitney stratifications on the variety Z 308
5.4.1 The canonical Whitney stratification 309
5.4.2 Regularity of the isotropy-free part 315

5.4.3 Change of degrees 318
5.4.4 Change of groups 324
5.5 Pseudocycles in Thom–Mather stratified spaces 331
5.5.1 Pseudocycles 334
5.5.2 A technical lemma 335
5.6 Construction of the integral Euler cycle 343
5.6.1 Strong transversality 344
5.6.2 The integral Euler cycles 347
5.6.3 Existence of strongly transverse perturbations 348
5.6.4 The pseudocycle property 354
5.6.5 Invariance of the Euler cycles 358
5.7 Stable complex derived orbifold bordism 359
5.7.1 Orbispaces 360
5.7.2 Stable complex structures 364
5.7.3 Derived orbifolds and bordism groups 366
5.7.4 Integral homology classes from stable complex derived orbifold bordism 369
5.7.5 Naturality of the $\mathcal{FOP}$ map 372
5.8 Whitney stratifications 374
5.8.1 Prestratifications and stratifications 375
5.8.2 Whitney stratifications 376
5.8.3 Invariance of minimal Whitney stratifications 378
5.8.4 Existence results in the complex algebraic case 381
5.8.5 Relative minimal Whitney stratifications 385
5.8.6 Relative to a prestratification 387

List of Tables

List of Figures

3.1 A schematic picture of the moduli space (when n is odd). 132

4.1 The simply-connected compact Lie groups 249

5.1 The fibration $\tilde{P}$. 280

Chapter 1

Introduction

The modern study of rigidity aspects of symplectic manifolds started from Gromov's seminal work [Gro85]. Following Gromov, counting pseudo-holomorphic curves and related objects usually require conducting "Fredholm differential topology" in infinite dimensions. The first theme of this thesis is to develop enumerative differential geometry in the presence of *symmetries*. Various results based on solutions to variants of equivariant transversality problems in symplectic topology are presented in Chapter 2-4. The second theme of this thesis is concerned with *Floer theory* of Lagrangian submanifolds. More precisely, we use *Fukaya categories*, which could be viewed as "non-commutative spectra" of symplectic manifolds and assemble the information of Lagrangian Floer homology of many Lagrangians, to study Weinstein manifolds. Chapter 5 showcases how to use a noncommutative invariant associated with the wrapped Fukaya categories, known as the Rouquier dimension, to derive interesting geometric applications.

We present a brief introduction in this chapter. The reader should refer to the introduction section of each chapter to get a comprehensive overview.

The second Chapter, which is posted as arXiv:2102.03665 , is entitled "A symplectic formula of generalized Casson invariants". It contains results that could be

viewed a decategorified version of the Atiyah–Floer conjecture. Suppose Y is an integer homology 3-sphere, Taubes [Tau90] proved that the number of irreducible critical orbits of the perturbed Chern–Simons functional on Y, counted with signs, is equal to the algebraic intersection number of two character varieties associated with Heegaard splittings when the structure group is SU(2). Taubes' result established a relationship between gauge theory and the Casson invariant. This chapter proves an analogous identification result for SU(n) generalized Casson invariants. As a special case, we show that the SU(3) Casson invariant of Boden–Herald [BH98] can be equivalently calculated by taking an appropriate intersection number of Lagrangian submanifolds.

The third Chapter, which is a joint work with Mohan Swaminathan and appears as arXiv:2106.01206, is entitled "Bifurcations of embedded curves and towards an extension of Taubes' Gromov invariant to Calabi–Yau 3-folds". Here we define an integer-valued virtual count of embedded pseudo-holomorphic curves of two times a primitive homology class and arbitrary genus in symplectic Calabi–Yau 3-folds, which can be viewed as an extension of Taubes' Gromov invariant [Tau96a]. The construction depends on a detailed study of bifurcations of moduli spaces of embedded pseudo-holomorphic curves which is partially motivated by Wendl's recent solution [Wen21] of Bryan–Pandharipande's super-rigidity conjecture.

The fourth chapter, which represents a joint work with Laurent Côté and could be found online as arXiv:2110.10663, is entitled "On the Rouquier dimension of wrapped Fukaya categories and a conjecture of Orlov". In this chapter, we study the Rouquier dimension [Rou08] of wrapped Fukaya categories of Liouville manifolds and pairs, and apply this invariant to various problems in algebraic and symplectic geometry. On the algebro-geometric side, we introduce a new method based on symplectic flexibility and mirror symmetry to bound the Rouquier dimension of derived categories of coherent sheaves on certain complex algebraic varieties and stacks. These bounds are sharp in dimension at most 3. As a result, we resolve a well-known conjecture

of Orlov [Orl09b] for a large class of new examples, including all toric 3-folds and certain log Calabi–Yau surfaces. On the symplectic side, we study various quantitative questions such as: (1) given a Weinstein manifold, what is the minimal number of intersection points between the skeleton and its image under a generic compactly-supported Hamiltonian diffeomorphism? (2) what is the minimal number of critical points of a Lefschetz fibration on a Liouville manifold with Weinstein fibers? We give lower bounds for these quantities which are to our knowledge the first to go beyond the basic flexible/rigid dichotomy.

The fifth chapter, which consists of a joint work with Guangbo Xu and appears as arXiv:2201.02688, is entitled "An integral Euler cycle in normally complex orbifolds and $\mathbb{Z}$-valued Gromov–Witten type invariants." In this chapter, we define an integral Euler cycle for a vector bundle E over an effective orbifold X for which (E, X) is (stably) normally complex. The transversality is achieved by using Fukaya–Ono's [FO97] "normally polynomial perturbations" and Brett Parker's [Par13] generalization to "normally complex perturbations." Two immediate applications in symplectic topology are the definition of integer-valued genus-zero Gromov–Witten type invariants for general compact symplectic manifolds using the global Kuranishi chart constructed by Abouzaid–McLean–Smith [AMS21], and an alternative proof of the cohomological splitting theorem of Abouzaid–McLean–Smith for Hamiltonian fibrations over S^2 with integer coefficients.

Chapter 2

A symplectic formula of generalized Casson invariants

2.1 Introduction

In 1985, Casson introduced an invariant for integer homology 3-spheres (see [AM90, Mar88]) by considering the SU(2)–representations of the fundamental group. Suppose Y is an integer homology 3-sphere with a Heegaard splitting $Y = H_1 \cup_\Sigma H_2$, then the inclusions of Σ in H_1 and H_2 induce embeddings of the SU(2)–character varieties of $\pi_1(H_1)$ and $\pi_1(H_2)$ into the SU(2)–character variety of $\pi_1(\Sigma)$. The character variety of Y is given by the intersection of the character varieties of M_1 and M_2 as a set. After taking a sufficiently small generic perturbation on the embeddings, one obtains a well-defined intersection number of the character varieties of $\pi_1(H_1)$ and $\pi_1(H_2)$. The intersection number is always even (see, for example, [Sav11, Corollary 17.6]), and the Casson invariant is defined to be half of this intersection number.

On the other hand, let P be the trivial SU(2) bundle over Y with a fixed trivialization, and let $\mathcal{C}(Y)$ be the affine space of smooth connections on P. Let θ be the flat connection associated with the given trivialization of P. For $B \in \mathcal{C}(Y)$, suppose

$B = \theta + b$ where $b \in \Omega^1(Y, \mathfrak{su}(2))$, then the *Chern-Simons functional*

$$\mathrm{CS} : \mathcal{C}(Y) \to \mathbb{R}$$

is defined by

$$\mathrm{CS}(B) = \frac{1}{2} \int_Y \mathrm{Tr}\bigg(db \wedge b + \frac{2}{3} b \wedge b \wedge b \bigg).$$

After fixing a Riemannian metric on Y, the formal gradient of CS at $B \in \mathcal{C}(Y)$ is equal to $*F_B$, where $*$ is the Hodge star operator of the given metric; therefore, the critical points of CS are given by flat connections, and the set of critical orbits of CS has a one-to-one correspondence with the conjugacy classes of SU(2)–representations of $\pi_1(Y)$. Taubes [Tau90] proved that after adding a generic perturbation to the Chern-Simons functional, the signed count of its critical orbits is equal to two times the Casson invariant.

A natural question about the Casson invariant is whether it can be extended to more general groups and more general 3-manifolds. This question has been studied intensively since the introduction of the theory, and many constructions are given using either intersection of character varieties arisen from Heegaard splittings or the Chern-Simons functionals. The Heegaard-splitting approach has been studied by Boyer-Nicas [BN90], Walker [Wal92a], Cappell-Lee-Miller [CLM90], and Curtis [Cur94]. More recently, Abouzaid-Manolescu [AM17] studied the intersection of $\mathrm{SL}_2(\mathbb{C})$ character varieties arisen from Heegaard splittings and obtained a sheaf-theoretic 3-manifold invariant. On the gauge theory side, Boden-Herald [BH98] defined an SU(3) Casson invariant for integer homology spheres by introducing real-valued correction terms on the reducible critical orbits. Variations of the SU(3) Casson invariant have also been defined by Boden-Herald-Kirk [BHK01] and Cappell-Lee-Miller [CLM02]. Together with Zhang, the author [BZ20] generalized the construction of Boden-Herald and defined perturbative SU(n)–Casson invariants for integer homology spheres. The

construction of [BZ20] is reviewed in Section 2.2.

This chapter studies the relationship of the gauge-theoretic generalization of Casson invariants in [BZ20] with Heegaard splittings. We obtain an identification result analogous to Taubes' theorem [Tau90]: the quantities entering into the definition of the SU(n)–Casson invariant could all be expressed in terms of quantities constructed using finite-dimensional symplectic geometry. As a special case, we show that the SU(3) Casson invariant of Boden-Herald [BH98] is equal to an appropriate equivariant intersection number of character varieties associated with Heegaard splittings.

Now we describe the contents of this chapter in more detail. Let $Y = H_1 \cup_\Sigma H_2$ be a closed, oriented 3-manifold with a Heegaard splitting, and let G be a compact Lie group with Lie algebra $\mathfrak{g}$. In general, the G–character variety of $\pi_1(\Sigma)$ have singularities. Their appearance put forth difficulties on constructing perturbations and studying intersection theory from the differential-geometric approach. To get around this difficulty, we make use of the *extended moduli space of flat connections* $\mathcal{M}^{\mathfrak{g}}(\Sigma')$ on Σ' introduced by Jeffrey [Jef94] and Huebschmann [Hue95] to study an equivalent question using equivariant geometry, where Σ' be the Riemann surface obtained from Σ be removing a disc. Roughly speaking, there is an open subset $\hat{\mathcal{M}}^{\mathfrak{g}}(\Sigma') \subset \mathcal{M}^{\mathfrak{g}}(\Sigma')$ which is a smooth symplectic manifold and there is a Hamiltonian G–action on $\hat{\mathcal{M}}^{\mathfrak{g}}(\Sigma')$ whose moment map reduction $\mu^{-1}(0)/G$ is naturally identified with the G–character variety of $\pi_1(\Sigma)$. The G–character varieties of $\pi_1(H_1)$ and $\pi_1(H_2)$ could be lifted to two smooth Lagrangian submanifolds L_1, L_2 of $\hat{\mathcal{M}}^{\mathfrak{g}}(\Sigma')$ which are invariant under the G–action. They are also contained in $\mu^{-1}(0)$. Details of the definition and properties of the extended moduli space will be reviewed in Section 2.3. Note that the extended moduli spaces were used in the work of Manolescu-Woodward [MW12] to define a symplectic version of instanton Floer homology.

We are then interested in studying the G–equivariant intersection theory of L_1 and L_2. In Section 2.4, we introduce a transversality result on Hamiltonian pertur-

bations that allows us to perturb L_1 and L_2 equivariantly so that their intersection is non-degenerate (see Definition 2.4.1). Actually, in Section 2.5.1, we introduce the notion of *a compatible pair of* a holonomy perturbation $f_{\mathbf{q}}$ on the 3–manifold Y and a G–equivariant Hamiltonian perturbation Φ_H on $\hat{\mathcal{M}}^{\mathfrak{g}}(\Sigma')$ such that there is a one-to-one correspondence between gauge orbits of $f_{\mathbf{q}}$–perturbed flat connections on Y and G–orbits of $\Phi_H(L_1) \cap L_2$. Moreover, the non-degeneracy of connection is equivalent to the non-degeneracy of Lagrangian intersection for such a compatible pair. Therefore, we can use the equivariant transversality result on the symplectic side to construct preferred holonomy perturbation on Y to simplify the calculation of generalized Casson invariants.

Let p, q be two intersection points of L_1 and L_2 after perturbation, and choose a disk D such that $\partial D = \gamma_1 \cup \gamma_2$, where γ_1 and γ_2 are two arcs connecting p and q on L_1 and L_2 respectively. Suppose $H \subset G$ is a closed subgroup and for every $z \in D$ we assume that the stabilizer of z contains H as a subgroup. For such data, we define an *equivariant Maslov index* $\mu^H(D)$ in Section 2.4.2 which takes value in the representation ring of H. On the other hand, the disk D defines a 1-parameter family $B_t(D)$ of (not necessarily) flat G connections on Y connecting the perturbed flat connections on Y induced by p and q. In [BZ20], the authors defined the *equivariant spectral flow* for a family of self-adjoint operators. The main identification result of this chapter is the following theorem, where the more precise statement is given in Theorem 2.5.19.

Theorem 2.1.1. *The equivariant Maslov index $\mu^H(D)$ is equal to the equivariant spectral flow associated with the 1-parameter family of connections $B_t(D)$ on Y.*

The proof of Theorem 2.1.1 is presented in Section 2.5. It follows from a combination of Nicolaescu's arguments in [Nic95], which allow us to identify the H–equivariant spectral flow of $B_t(D)$ with certain infinite-dimensional Maslov index (Theorem 2.5.8), Dostoglou-Salamon's adiabatic limit result in [DS94] which reduces

the infinite-dimensional Maslov index to some finite-dimensional Maslov index (Theorem 2.5.17), and some observations regarding the relation between tangent space of the extended moduli space $\hat{\mathcal{M}}^{\mathfrak{g}}(\Sigma')$ and twisted harmonic forms.

After a further identification between the symplectic areas and Chern-Simons invariants coming into the definitions of SU(3)–Casson invariants on symplectic and gauge-theoretic sides, we prove

Theorem 2.1.2. *Let $Y = H_1 \cup_\Sigma H_2$ be an integer homology 3–sphere equipped with a Heegaard splitting such that the genus of Σ is at least 3. Suppose Φ_H is a sufficiently small G–equivariant Hamiltonian perturbation on $\hat{\mathcal{M}}^{\mathfrak{g}}(\Sigma')$ such that $\Phi_H(L_1)$ intersect L_2 non-degenerately. Let $[\tilde{\theta}] \in \Phi_H(L_1) \cap L_2$ be the intersection coming from perturbing the product connection $[\theta] \in L_1 \cap L_2$. Given any $p \in \Phi_H(L_1) \cap L_2$, let $D(p)$ be a disc in $\hat{\mathcal{M}}^{\mathfrak{g}}(\Sigma')$ as above which connects p and $[\tilde{\theta}]$. For any p, choose a point $\hat{p}$ in $L_1 \cap L_2$ close to p and let $D(\hat{p})$ be a disc connecting $\hat{p}$ and $[\theta]$ which is close to $D(p)$. Then the Boden-Herald* SU(3)*–Casson invariant is equal to*

$$\begin{aligned}\lambda_{\mathrm{SU}(3)}(Y) = &\sum_{[p] \in (\Phi_H(L_1) \cap L_2)^{irr}} (-1)^{\mu(D(p))} \\ &- \sum_{[p] \in (\Phi_H(L_1) \cap L_2)^{red}} (-1)^{\mu_t(D(p))} (\mu_n(D(p)) - \frac{\omega(D(\hat{p}))}{2\pi^2} + 1).\end{aligned} \tag{2.1.1}$$

In the above formula, $(\Phi_H(L_1) \cap L_2)^{irr}$ is the set of G–orbits of $(\Phi_H(L_1) \cap L_2)$ whose stabilizer is isomorphic to the center of SU(3) *and $(\Phi_H(L_1) \cap L_2)^{red}$ is the set of G–orbits of $(\Phi_H(L_1) \cap L_2)$ whose stabilizer is isomorphic to* U(1) *which corresponds to perturbed* SU(2)*–flat connections on Y. The index $\mu(D(p))$ is the equivariant Maslov index with respect to the trivial group, while the indices $\mu_t(D(p))$ and $\mu_n(D(p))$ are the trivial and nontrivial component of the* U(1)*–equivariant Maslov index $\mu^{\mathrm{U}(1)}(D(p))$ respectively. The symbol $\omega(D(\hat{p}))$ is the symplectic area of the disc $D(\hat{p})$ with respect to the symplectic form on $\hat{\mathcal{M}}^{\mathfrak{g}}(\Sigma')$.*

The reader might have realized that this is parallel to the definition of Casson-Walker invariant [Wal92a, equation (2.2)] for rational homology spheres.

We hope that the symplectic approach used in this chapter could help understand structural aspects of the generalized Casson invariants. Indeed, in [Wal92a], Walker proved that the SU(2)–Casson-Walker invariant is combinatorial in nature based on studying various isotopies of character varieties relevant to Dehn surgeries on 3–manifolds. The combinatorial approach was generalized by Lescop in [Les96] to define invariants for all 3–manifolds. The symplectic geometry provides much flexibility, though the use of Chern-Simons type invariants in our definition might put obstructions to the understanding. However, given a knot $K \subset Y$, let $Y_{1/k}(K)$ be the 3–manifold obtained by doing Dehn surgery of Y along K with slope $1/k$, it is reasonable to expect that the SU(n)–Casson invariants of $Y_{1/k}(K)$ might satisfy the asymptotic growing rate k^{n-1} as $k \to \infty$ for a fixed n. When $n = 2$, this is the classical surgery formula for Casson-Walker invariants, see [Wal92a, Chapter 3]; when $n = 3$, such speculation is supported by the discussion in [BHK05, Section 6]. We will not discuss this topic further in this chapter.

Acknowledgements

The author would like to thank his advisor John Pardon for constant support and encouragement. Special thanks are given to Boyu Zhang for numerous discussions and suggestions.

2.2 Recap of 3–dimensional gauge theory

In this section, we recall some definitions and results from [BZ20] which will be used in this chapter. Interested readers should consult *loc.cit.* for full details.

2.2.1 Some linear algebra

Let G be a connected compact Lie group. Write the representation ring of G as $\mathcal{R}(G)$ and let $\mathcal{R}(G) \otimes \mathbb{R}$ be its natural $\mathbb{R}$–extension.

Definition 2.2.1. *Define $\mathcal{R}_G$ to be the set of conjugation equivalence classes of the triple (H, V, ρ), where H is a closed subgroup of G, and $\rho : H \to \mathrm{Hom}(V, V)$ is a finite-dimensional $\mathbb{R}$–linear representation of H. We say that (H, V, ρ) is* conjugation equivalent *to (H', V', ρ'), if there exists $g \in G$ and an isomorphism $\varphi : V \to V'$, such that*

$$H' = gHg^{-1},$$

and

$$\rho'(ghg^{-1}) = \varphi \circ \rho(h) \circ \varphi^{-1}.$$

Given a closed subgroup $H \subset G$, let $\mathcal{R}_G([H])$ be the subset of $\mathcal{R}_G$ consisting of elements represented by representations of H. Any $\sigma \in \mathcal{R}_G$ is called irreducible if it is represented by an irreducible representation.

Definition 2.2.2. *Let H be a closed subgroup of G. Define*

$$i_G^H : \mathcal{R}_H \to \mathcal{R}_G$$

be the tautological map by identifying subgroups of H as subgroups of G.

Given a compact Lie group H, if V is an irreducible $\mathbb{R}$-linear H–representation, recall that V is said to be of type $\mathbb{R}, \mathbb{C}$ or $\mathbb{H}$, if $\hom_H(V, V) \cong \mathbb{R}, \mathbb{C}$ or $\mathbb{H}$ respectively.

Definition 2.2.3. *Suppose V is an orthogonal irreducible H–representation of type $\mathbb{K} \in \{\mathbb{R}, \mathbb{C}, \mathbb{H}\}$. Given $r \in \mathbb{Z}_{>0}$, define $d_V(r)$ to be the dimension of self-adjoint maps*

on V^r intertwining with the H action. In other words,

$$d_V(r) = \begin{cases} \frac{1}{2}r(r+1) & \text{if } \mathbb{K} = \mathbb{R}, \\ r^2 & \text{if } \mathbb{K} = \mathbb{C}, \\ 2r^2 - r & \text{if } \mathbb{K} = \mathbb{H}. \end{cases}$$

Definition 2.2.4. *Given $\sigma \in \mathcal{R}_G$ with representative given by (H, V, ρ) such that V is an orthogonal H–representation, suppose the isotypic decomposition of V as an H–representation is given by*

$$V \cong V_1^{\oplus a_1} \oplus \cdots \oplus V_m^{\oplus a_m}.$$

Define the dimension function *of σ, denoted by $d(\sigma)$, to be the following quantity:*

$$d(\sigma) = \sum_{i=1}^{m} d_V(a_i).$$

This definition is independent of the choice of the representative of the equivalence class σ. According to Schur's lemma, $d(\sigma)$ is equal to the dimension of H–equivariant self-adjoint linear endomorphisms, denoted by $\mathrm{Sym}_H(V)$, of the representation V. The following lemma is [BZ20, Lemma 3.18].

Lemma 2.2.5. *Suppose $\sigma \in \mathcal{R}_G([H])$ is represented by the triple (H, V, ρ) such that V is an orthogonal H–representation. Let $\mathrm{Sym}_{H,\sigma}(V) = \{s \in \mathrm{Sym}_H(V) | \sigma \cong (H, \ker(s), \rho|_{\ker(s)})\}$. Then $\mathrm{Sym}_{H,\sigma}(V)$ is a smooth submanifold of $\mathrm{Sym}_H(V)$ of codimension $d(\sigma)$. If $K \subset \mathrm{Sym}_H(V)$ is a linear space, let $\Pi : V \to \ker(s)$ be the orthogonal projection, then $s + K$ is transverse to $\mathrm{Sym}_{H,\sigma}(V)$ at s if and only if the map taking $l \in K$ to $(x \mapsto \Pi(l(x))) \in \mathrm{Sym}_H(\ker(s))$ is surjective.*

2.2.2 Some 3–dimensional gauge theory

Let Y be a smooth, oriented, closed 3–manifold equipped with a Riemannian metric. Suppose G is a compact, connected and simply-connected Lie group and let $\mathfrak{g}$ be its Lie algebra. Let $P = Y \times G$ be the product principal G–bundle over Y. We use θ to represent the trivial product connection on P. Fix an integer $k \geq 2$. Let $\mathcal{C}(Y)$ be the affine space of L^2_k–connections over P, which is an affine space modeled on $L^2_k(T^*Y \otimes \mathfrak{g})$. Let $\mathcal{G}_Y$ be the L^2_{k+1}–gauge group of P. Then $\mathcal{G}_Y$ acts smoothly on $\mathcal{C}(Y)$ via gauge transformation. Given $B = \theta + b \in \mathcal{C}(Y)$, where $b \in L^2_k(T^*Y \otimes \mathfrak{g})$, let $\mathrm{Stab}(B)$ and $\mathrm{Orb}(B)$ be the stablizer group and orbit of B under the $\mathcal{G}_Y$–action respectively. For $H \subset G$, let $\mathcal{C}(Y)^H$ be the subset of $\mathcal{C}(Y)$ consisting of connections whose stabilizer contains a subgroup of G which is conjugate to H.

Recall the Chern-Simons functional CS on $\mathcal{C}(Y)$ is an $\mathbb{R}$–valued smooth function

$$\mathrm{CS}(\theta + b) := \frac{1}{2}\int_Y d_\theta b \wedge b + \frac{1}{3}\int_Y [b \wedge b] \wedge b. \tag{2.2.1}$$

Using the Riemannian metric on Y, we have

$$(\nabla \mathrm{CS})(B) = *F_B \text{ for all } B \in \mathcal{C}(Y),$$

where $F_B = d_\theta b + \frac{1}{2}[b, b]$ is the curvature of B and $*$ is the Hodge star operator. In particular, critical points of CS correspond to flat connections over P.

Next we discuss about a slight generalization of the holonomy perturbation [Don87], [Flo88]. We follow the notation of [KM11]. This form of holonomy perturbation will be used later in Section 2.5. Regard the circle S^1 as $\mathbb{R}/\mathbb{Z}$. We identify the open unit disc in the plane D^2 with the open cube $(-1, 1)^2$.

Definition 2.2.6. *A* cylinder datum *is a tuple* $(q_1, \cdots, q_m, \beta, h_s)$ *with* $m \in \mathbb{Z}_{>0}$ *that satisfies the following conditions.*

1. $q_i : S^1 \times D^2 \cong S^1 \times (-1,1)^2 \to Y$ *is a smooth immersion for* $i = 1, \cdots, m$;

2. *there exists* $\epsilon > 0$ *such that* $q_1, \ldots, q_m$ *coincide on* $(-\epsilon, \epsilon) \times D^2$,

3. β *is a non-negative, compactly supported smooth function on* $(-1,1)$, *such that*

$$\int_{-1}^{1} \beta = 1;$$

4. *For any* $-1 \leq s \leq 1$, $h_s : G^m \to \mathbb{R}$ *is a smooth function that is invariant under the diagonal action of* G *by conjugations and the family* $\{h_s\}_{-1 \leq s \leq 1}$ *depends smoothly on* s.

Given $B \in \mathcal{C}(Y)$ and cylinder datum $\mathbf{q} = (q_1, \cdots, q_m, \beta, h_s)$, we can define the *holonomy map* of B by

$$\begin{aligned} \mathrm{Hol}_{\mathbf{q}}(B) : D^2 &\to G^m \\ z &\mapsto \big(\mathrm{Hol}_{q_1,z}(B), \cdots, \mathrm{Hol}_{q_m,z}(B)\big), \end{aligned}$$

where $\mathrm{Hol}_{q_i,z}(B)$ is the element in G obtained by the holonomy of B along the loop $q_i(S^1 \times \{z\})$ for $1 \leq i \leq m$ using the trivialization $P_{\{0\}\times\{z\}} \cong G$. Then the *cylinder function* associated to $\mathbf{q}$ is defined to be the map

$$\begin{aligned} f_{\mathbf{q}} : \mathcal{C}(Y) &\to \mathbb{R} \\ B &\mapsto \int_{-1}^{1} \int_{-1}^{1} \beta(\lambda) h_s(\mathrm{Hol}_{\mathbf{q}}(B)) d\lambda ds, \end{aligned}$$

where we use the coordinate $z = (\lambda, s)$ under the identification $D^2 \cong (-1,1)^2$. When $h_s = \chi(s)h$ such that $h : G^m \to \mathbb{R}$ is a smooth function invariant under the diagonal action of G by conjugations and $\chi : (-1,1) \to \mathbb{R}$ is a compactly supported smooth function with $\int_{-1}^{1} \chi(s) ds = 1$, the 2–form $\beta(\lambda)\chi(s) d\lambda \wedge ds$ defines a non-negative bump 2–form over D^2 with total integration 1 and $f_{\mathbf{q}}$ is the usual cylinder function

as in [KM11]. These cylinder functions are smooth over $\mathcal{C}(Y)$ and have well-defined formal gradients. Then one can follow the procedure in [KM11, Definition 3.6] to construct the Banach space of holonomy perturbations, which is denoted by $\mathcal{P}$. Given any $\pi \in \mathcal{P}$, write the associated function over $\mathcal{C}(Y)$ as f_π and let V_π be the formal gradient of f_π. Let DV_π be the derivative of V_π.

Definition 2.2.7. *A connection B is called π–flat if*

$$*F_B + V_\pi(B) = 0.$$

In other words, B is a critical point of the functional $\mathrm{CS} + f_\pi$.

Given a pair $B \in \mathcal{C}(Y)$ and $\pi \in \mathcal{P}$, we can define a self-adjoint Fredholm operator with index zero which has discrete real spectrum

$$K_{B,\pi} : L^2_k(\mathfrak{g}) \oplus L^2_k(T^*Y \otimes \mathfrak{g}) \to L^2_{k-1}(\mathfrak{g}) \oplus L^2_{k-1}(T^*Y \otimes \mathfrak{g})$$

by

$$K_{B,\pi}(\xi, b) := (d^*_B b, d_B \xi + *d_B b + DV_\pi(B)(b)). \tag{2.2.2}$$

When B is π–flat, define $\mathrm{Hess}_{B,\pi}$ to be the operator

$$\mathrm{Hess}_{B,\pi} : \ker d^*_B \cap L^2_k(T^*Y \otimes \mathfrak{g}) \to \ker d^*_B \cap L^2_{k-1}(T^*Y \otimes \mathfrak{g})$$

given by

$$\mathrm{Hess}_{B,\pi} := *d_B + DV_\pi(B).$$

This is a self-adjoint Fredholm operator with index zero. Note that both $K_{B,\pi}$ and $\mathrm{Hess}_{B,\pi}$ are $\mathrm{Stab}(B)$–equivariant operators.

Definition 2.2.8. *A π–flat connection B is called* non-degenerate *if* $\mathrm{Hess}_{B,\pi}$ *is an isomorphism. A perturbation $\pi \in \mathcal{P}$ is called* non-degenerate *if all the critical points*

of $\mathrm{CS}+f_\pi$ *are non-degenerate.*

Because the classical holonomy perturbations all arise as special cases of our formulation of holonomy perturbation, the transversality results proved in [BZ20, Section 4] all carry over due to the abundance of perturbations. The following proposition is [BZ20, Corollary 4.20, 4.21].

Proposition 2.2.9. *The set* $\mathcal{P}^{reg} \subset \mathcal{P}$ *of non-degenerate holonomy perturbations is of Baire second category.*

For any pair $\pi_0, \pi_1 \in \mathcal{P}^{reg}$*, one can find a generic smooth path* $\pi_t : [0,1] \to \mathcal{P}$ *from* π_0 *to* π_1*, such that there are only countably many* t *where* π_t *is degenerate. Moreover, for every such* t *there is exact one degenerate critical orbit* $\mathrm{Orb}(B)$ *of* $\mathrm{CS}+f_{\pi_t}$*, and the kernel of* Hess_{B,π_t} *is an irreducible representation of* $\mathrm{Stab}(B)$*.*

Next we recall definitions and results which lead to the definition of $\mathrm{SU}(n)$–Casson invariants using gauge theory.

Let $\mathcal{V}$ be a Hilbert space and $\mathcal{D} \subset \mathcal{V}$ is a space whose image in $\mathcal{V}$ is dense and the inclusion map is a compact operator. Let $D_t : \mathcal{D} \to \mathcal{V}, t \in [0,1]$ be a 1–parameter family of self-adjoint operators whose spectra are real and discrete. Let H be a connected compact Lie group and $\mathcal{V}$ admits a linear H–action which preserves $\mathcal{D}$. Suppose that for all $0 < t < 1$, the operator D_t is equivariant with respect to the H–action. Given any W which is a finite-dimensional irreducible H–representation, we have an induced family of self-adjoint operators on $\mathrm{Hom}_H(W, \mathcal{V})$ and let the spectral flow of this family be n_W. We use the convention that if D_0 or D_1 has nontrivial kernel, the spectral flow of the family $\{D_t\}_{0\le t\le 1}$ is calculated by the spectral flow of $\{D_t + \epsilon \cdot \mathrm{id}\}_{0\le t\le 1}$ where $\epsilon > 0$ is sufficiently small. Let $\mathcal{R}^{irr}(H)$ be the set of finite-dimensional irreducible H–representations.

Definition 2.2.10. *Define the* H–equivariant spectral flow *of* $\{D_t\}_{0\le t\le 1}$*, denoted by*

$Sf_H(D_t)$ to be the following element in the representation ring $\mathcal{R}(H)$:

$$\sum_{W\in\mathcal{R}^{irr}(H)} n_W \cdot [W]. \tag{2.2.3}$$

Let $\pi \in \mathcal{P}$ be a non-degenerate perturbation and suppose B is a π–flat connection. Suppose H is a closed subgroup of G and the group $\mathrm{Stab}(B) = H$.

Definition 2.2.11. *Define $Sf_H(B,\pi) \in \mathcal{R}(H)$ to be the H–equivariant spectral flow given by the homotopy from $K_{B,\pi}$ to $K_{\theta,0}$ constructed from the linear homotopy from (B,π) to $(\theta,0)$.*

Because of the homotopy invariance of the spectral flow, $Sf_H(B,\pi)$ could be calculated by any path of H–equivariant operators connecting $K_{B,\pi}$ and $K_{\theta,0}$.

Let $n \geq 2$ be an integer. Suppose the Lie group G is given by $\mathrm{SU}(n)$ and $\mathbb{C}^n$ is the standard representation of $\mathrm{SU}(n)$. We further assume that Y is an integer homology sphere. Let Σ_n be the set of tuples of pairs of integers

$$((n_1,m_1),\ldots,(n_r,m_r)), r \geq 1$$

such that $n = \sum_{i=1}^r n_i m_i$. We also assume that $n_1 \leq \cdots \leq n_r$ and m_i's are in non-decreasing order if the corresponding n_i's are the same. By the discussion in [BZ20, Section 5], given any connection $B \in \mathcal{C}(Y)$, there exists $\sigma \in \Sigma_n$, such that the vector bundle

$$E = P \times_{\mathrm{SU}(n)} \mathbb{C}^n$$

could be decomposed into

$$E \cong E(n_1)^{\oplus m_1} \oplus \cdots \oplus E(n_r)^{\oplus m_r} \tag{2.2.4}$$

satisfying

1. $E(n_i)$ is a trivial $\mathbb{C}$–vector bundle over Y of rank n_i for $1 \leq i \leq r$;

2. The connection B preserves this decomposition and is decomposed into direct sums of irreducible unitary connections on each component $E(n_i)$.

Such a connection B is called *of type* σ. Fix a $\sigma \in \Sigma_n$, connections with above decomposition define a subset $\mathcal{C}_\sigma(Y) \subset \mathcal{C}(Y)$. All elements in $\mathcal{C}_\sigma(Y)$ have stabilizers conjugate to a subgroup $H_\sigma \subset \mathrm{SU}(n)$. Then $\mathcal{C}_\sigma(Y) \subset \mathcal{C}(Y)^{H_\sigma}$ and the linear path from $B \in \mathcal{C}_\sigma(Y)$ to θ stays in $\mathcal{C}(Y)^{H_\sigma}$.

Suppose $g : E \to E$ is a gauge transformation that is decomposed into the direct sum of gauge transformations $g_i : E(n_i) \to E(n_i)$ using (2.2.4). If $B \in \mathcal{C}_\sigma(Y)$, one can find $\tau_1, \ldots, \tau_r \in \mathcal{R}(H_\sigma)$ such that the H_σ–equivariant spectral flow from $K_{B,\pi}$ to $K_{g(B),\pi}$ is given by

$$\sum_{n_i \geq 2} \deg(g_i) \cdot \tau_i$$

which is independent of B and $\pi \in \mathcal{P}$, where $\deg(g_i)$ is the degree of the map $H_3(Y;\mathbb{Z}) \to H_3(\mathrm{U}(n);\mathbb{Z})$ induced by g_i. If B is furthermore a flat connection and let B_i be the $E(n_i)$–component of B, define

$$\mathrm{CS}_\sigma(B) = \sum_{n_i \geq 2} \frac{\mathrm{CS}(B_i)}{4\pi^2 n_i} \cdot \tau_i \in \mathcal{R}(H_\sigma) \otimes \mathbb{R}. \tag{2.2.5}$$

Now we are ready to recall the definition of the $\mathrm{SU}(n)$–Casson invariant. Let $\mathcal{C}_F(Y) \subset \mathcal{C}(Y)$ be the subspace consisting of flat connection on P. Let $\mathcal{U} \subset \mathcal{C}(Y)$ be an open subset containing $\mathcal{C}_F(Y)$ such that the inclusion

$$\mathcal{C}_F(Y) \cap \mathcal{C}(Y)^{H_\sigma} \hookrightarrow \mathcal{U} \cap \mathcal{C}(Y)^{H_\sigma}$$

induces a one-to-one correspondence on connected components. By Uhlenbeck compactness, there exists $r_0 > 0$ such that as long as $\pi \in \mathcal{P}$ satisfies $\|\pi\| < r_0$, the space of π–flat connections is contained in U.

Definition 2.2.12. *Suppose $\|\pi\| < r_0$ is a non-degenerate perturbation and let B be a π–perturbed flat connection of type $\sigma \in \Sigma_n$. Take $g \in \mathcal{G}_Y$ such that $g(B) \in \mathcal{C}(Y)^{H_\sigma}$ and let $g(\hat{B})$ be a genuine flat connection lying in the same connected component of $\mathcal{U} \cap \mathcal{C}(Y)^{H_\sigma}$ as $g(B)$. Define* $\mathrm{ind}(B, \pi)$ *to be*

$$Sf_{H_\sigma}(g(B), \pi) - [\ker d_{g(B)}] - \mathrm{CS}_\sigma(g(\hat{B})) \in \mathcal{R}(H_\sigma) \otimes \mathbb{R}, \tag{2.2.6}$$

where $[\ker d_{g(B)}]$ *is the H_σ–representation associated to the Lie algebra of* $\mathrm{Stab}(B)$.

The following is claimed in [BZ20, Section 5].

Theorem 2.2.13. *There exists a set $\widetilde{\mathcal{R}}_{\mathrm{SU}(n)}$, a well-defined map $\mathcal{R}(H_\sigma) \otimes \mathbb{R} \to \widetilde{\mathcal{R}}_{\mathrm{SU}(n)}$ and a $\mathbb{Z}$–submodule $\widetilde{\mathrm{Bif}}_{\mathrm{SU}(n)} \subset \mathbb{Z}\widetilde{\mathcal{R}}_{\mathrm{SU}(n)}$ with the following significance. Let* $[\mathrm{ind}(B, \pi)]$ *be the image of* $\mathrm{ind}(B, \pi)$ *in $\widetilde{\mathcal{R}}_{\mathrm{SU}(n)}$. If $\|\pi\| < r_0$ is a non-degenerate perturbation, define*

$$\mathrm{ind}_\pi := \sum_{\mathrm{Orb}(B) \ \textit{is } \pi-\textit{ flat}} [\mathrm{ind}(B, \pi)] \in \mathbb{Z}\widetilde{\mathcal{R}}_{\mathrm{SU}(n)}.$$

Then for any pair of non-degenerate $\pi_1, \pi_2 \in \mathcal{P}$ such that $\|\pi_i\| < r_0$ for $i = 1, 2$, we have

$$\mathrm{ind}_{\pi_1} - \mathrm{ind}_{\pi_2} \in \widetilde{\mathrm{Bif}}_{\mathrm{SU}(n)}.$$

Therefore the image of ind_π in the $\mathbb{Z}$–module $\mathbb{Z}\widetilde{\mathcal{R}}_{\mathrm{SU}(n)}/\widetilde{\mathrm{Bif}}_{\mathrm{SU}(n)}$ is a topological invariant of the integer homology sphere Y based on the above claim. One can then use this quantity to construct Casson-type invariants using different normalizations. In particular, the $\mathbb{R}$–valued SU(3)–Casson invariant of Boden-Herald [BH98] could be recovered as a special case of the above construction.

However, it seems to the author that the wall-crossing terms coming from the values of Chern-Simons' functional were not treated properly in [BZ20]. The revision of [BZ20] would carry out an alternative definition using reference connections different from θ. Nevertheless, it still makes sense to talk about the equivariant index

$\mathrm{ind}(B,\pi)$ regardless of the definition of Casson invariants. Our goal in this chapter, namely Theorem 2.1.1, is to present $\mathrm{ind}(B,\pi)$ using geometric quantities constructed from finite-dimensional symplectic geometry when Y is equipped with a Heegaard splitting and the perturbation π has certain adapted behavior. Its application to establish Theorem 2.1.2 is independent of [BZ20, Section 5].

2.3 Extended moduli spaces

In this section, we recall the construction of the extended moduli space in [Jef94] and discuss about some of its properties. Most importantly, we will see that the essential information of the tangent spaces of the extended moduli space is encoded in the corresponding space of twisted harmonic forms. We also recall the construction of Lagrangian submanifolds in extended moduli spaces from [MW12] which will be used in this chapter.

Suppose Σ is a closed Riemann surface with genus $h \geq 2$ equipped with a Riemannian metric. Let $p \in \Sigma$ be a base point. Define $\Sigma' = \Sigma \backslash D_p^2$ to be the Riemann surface with one boundary component obtained by removing a disc D_p^2 around p. Suppose G is a simply-connected compact Lie group and $\mathfrak{g}$ is its Lie algebra. We identify $\mathfrak{g}$ with its dual $\mathfrak{g}^*$ using the Killing form. Let P be the product principal G–bundles over Σ. We continue to use θ to represent the trivial product connection.

Fix an integer $k \geq 2$.Let $\mathcal{A}(\Sigma)$ and $\mathcal{A}(\Sigma')$ be the affine space of L_k^2–connections over P and $P|_{\Sigma'}$ respectively. Denote the L_{k+1}^2–gauge groups of P and $P|_{\Sigma'}$ by $\mathcal{G}_\Sigma$ and $\mathcal{G}_{\Sigma'}$ respectively. Let $\mathcal{M}(\Sigma) = \mathcal{A}_F(\Sigma)/\mathcal{G}_\Sigma$ be the moduli space of flat connections over P, where $\mathcal{A}_F(\Sigma)$ is the space of flat connections over P. We also need the group $\mathcal{G}_{\Sigma'}^c$ consisting of elements in $\mathcal{G}_{\Sigma'}$ which are identity near the boundary of Σ'.

Definition 2.3.1. *Let s be the coordinate on $S^1 \cong \partial\Sigma'$. Define*

$$\mathcal{A}_F^{\mathfrak{g}}(\Sigma') := \{A \in \mathcal{A}(\Sigma') \big| F_A = 0, A = \theta + \xi ds \text{ near the boundary for some } \xi \in \mathfrak{g}\}.$$

Then the extended moduli space $\mathcal{M}^{\mathfrak{g}}(\Sigma')$ *is defined to be the quotient*

$$\mathcal{M}^{\mathfrak{g}}(\Sigma') := \mathcal{A}_F^{\mathfrak{g}}(\Sigma')/\mathcal{G}_{\Sigma'}^c.$$

$\mathcal{M}^{\mathfrak{g}}(\Sigma')$ has a finite-dimensional description similar to character varieties. Let us choose $\{\alpha_i, \beta_i\}_{1 \le i \le h}$ to be a set of simple loops on Σ with common base point at some $p' \in \partial\Sigma' \subset \Sigma' \subset \Sigma$ which gives the standard presentation of $\pi_1(\Sigma)$. Choose a simple loop γ in Σ' based at p' wrapping around the boundary once then we have $\Pi_{i=1}^h [\alpha_i, \beta_i] = \gamma$. Using the holonomy map, one can show that

$$\mathcal{M}^{\mathfrak{g}}(\Sigma') = \{A_i, B_i \in G, \forall 1 \le i \le h, \xi \in \mathfrak{g} \big| \Pi_{i=1}^h [A_i, B_i] = \exp(\xi)\}. \tag{2.3.1}$$

We will use both the gauge-theoretic definition and the representation-theoretic definition in the sequel. Note that $\mathcal{M}^{\mathfrak{g}}(\Sigma')$ admits a G–action: given $g \in G$, it acts on $\mathcal{M}^{\mathfrak{g}}(\Sigma')$ by

$$((A_i, B_i)_{1 \le i \le h}, \xi) \mapsto ((gA_ig^{-1}, gB_ig^{-1})_{1 \le i \le h}, \mathrm{Ad}_g(\xi)),$$

where Ad_g is the adjoint action of g on $\mathfrak{g}$. Use the Killing form on $\mathfrak{g}$ to define a norm on $\mathfrak{g}$, the next proposition is proved in [Jef94].

Proposition 2.3.2. *1. The open G–invariant subset $\hat{\mathcal{M}}^{\mathfrak{g}}(\Sigma') \subset \mathcal{M}^{\mathfrak{g}}(\Sigma')$ defined by $|\xi| < \delta$ for some sufficiently small $\delta > 0$ is actually a smooth manifold;*

2. There exists a symplectic form ω on $\hat{\mathcal{M}}^{\mathfrak{g}}(\Sigma')$ and the G–action defines a smooth Hamiltonian action with respect to ω. The moment map $\mu : \hat{\mathcal{M}}^{\mathfrak{g}}(\Sigma') \to \mathfrak{g}^*$ of

this G–action is given by

$$((A_i, B_i)_{1\leq i\leq h}, \xi) \mapsto \xi;$$

3. *The moment map reduction $\mu^{-1}(0)/G$ is naturally identified with $\mathcal{M}(\Sigma)$;*

4. *If $[A] \in \mathcal{M}(\Sigma)$ is represented by some $A \in \mathcal{A}_F^{\mathfrak{g}}(\Sigma')$, then the stabilizer group of A under the $\mathcal{G}_{\Sigma'}$–action, written as* $\mathrm{Stab}(A)$*, is isomorphic to* $\mathrm{Stab}([A])$*, the stabilizer of $[A] \in \mathcal{M}(\Sigma)$ under the G–action.*

Making δ smaller if necessary, we can assume that the exponential map $\exp : \mathfrak{g} \to G$ is a diffeomorphism onto its image over $\{\xi \in \mathfrak{g} \big| |\xi| < \delta\}$. Then the following statement follows from the equation (2.3.1).

Lemma 2.3.3. *The map $\hat{\mathcal{M}}^{\mathfrak{g}}(\Sigma') \to G^{2h}$ defined by mapping $((A_i, B_i)_{1\leq i\leq h}, \xi)$ to the G^{2h}–component is a G–equivariant embedding, where G^{2h} is equipped with the diagonal conjugation action.* □

Introduce the notation

$$\mathcal{A}_F^0(\Sigma) := \{A \in \mathcal{A}_F(\Sigma) \big| A|_{D_p^2} \text{ is given by the product connection}\}.$$

Then $\mathcal{A}_F^0(\Sigma)$ could be viewed as a subspace of $\mathcal{A}_F^{\mathfrak{g}}(\Sigma')$ corresponding to $\xi = 0$. Choose $A \in \mathcal{A}_F^0(\Sigma)$. Denote by $[A] \in \hat{\mathcal{M}}^{\mathfrak{g}}(\Sigma')$ the image of A in $\hat{\mathcal{M}}^{\mathfrak{g}}(\Sigma')$, we now give an explicit description of the tangent space $T_{[A]}\hat{\mathcal{M}}^{\mathfrak{g}}(\Sigma')$.

Let $\Omega^i(\Sigma') \otimes \mathfrak{g}$ (resp. $\Omega^i(\Sigma) \otimes \mathfrak{g}$) be the space of $\mathfrak{g}$–valued i–forms on Σ' (resp. Σ) for $i = 0, 1, 2$. Also write the subspace of $\Omega^i(\Sigma') \otimes \mathfrak{g}$ consisting of elements with compact support as $\Omega_c^i(\Sigma') \otimes \mathfrak{g}$. Recall that s is the coordinate on $\partial\Sigma'$. Define

$$\Omega^{1,\mathfrak{g}}(\Sigma') = \{a \in \Omega^1(\Sigma') \otimes \mathfrak{g} \big| a = \xi ds \text{ near the boundary for some } \xi \in \mathfrak{g}\ \}.$$

We have the elliptic complex

$$\Omega_c^0(\Sigma') \otimes \mathfrak{g} \xrightarrow{d_A} \Omega^{1,\mathfrak{g}}(\Sigma') \xrightarrow{d_A} \Omega_c^2(\Sigma') \otimes \mathfrak{g}. \tag{2.3.2}$$

Then the tangent space $T_{[A]}\hat{\mathcal{M}}^{\mathfrak{g}}(\Sigma')$ is identified with the first cohomology group of (2.3.2), written as $\tilde{H}_A^{1,\mathfrak{g}}$. The bilinear form

$$\omega_{\Sigma'}(a_1, a_2) = \int_{\Sigma'} \mathrm{Tr}(a_1 \wedge a_2)$$

over $\Omega^{1,\mathfrak{g}}(\Sigma')$ defines the symplectic form on $\hat{\mathcal{M}}^{\mathfrak{g}}(\Sigma')$ under this identification.

On the other hand, view $A \in \mathcal{A}_F(\Sigma)$, we have the complex

$$\Omega^0(\Sigma) \otimes \mathfrak{g} \xrightarrow{d_A} \Omega^1(\Sigma) \otimes \mathfrak{g} \xrightarrow{d_A} \Omega^2(\Sigma) \otimes \mathfrak{g}. \tag{2.3.3}$$

The cohomolgy groups H_A^i for $i = 0, 1, 2$ could be identified with the space of twisted harmonic i–forms. Similarly, the bilinear form

$$\omega_{\Sigma}(a_1, a_2) = \int_{\Sigma} \mathrm{Tr}(a_1 \wedge a_2)$$

over $\Omega^1(\Sigma) \otimes \mathfrak{g}$ is non-degenerate and descends down to a non-degenerate pairing on H_A^1.

We need to introduce some auxiliary objects to understand the relation between $\tilde{H}_A^{1,\mathfrak{g}}$ and H_A^1 as symplectic vector spaces. Let $\overline{\Sigma} \cong \Sigma \backslash \{p\}$ be the Riemann surface obtained from Σ' by attaching a cylindrical end $S^1 \times [0, +\infty)$ along $\partial\Sigma'$. We still use P to represent the product principal G–bundle over $\overline{\Sigma}$. Choose $\epsilon > 0$ to be a positive real number. Define $L_k^{2,\epsilon}(\Omega^i(\overline{\Sigma}) \otimes \mathfrak{g})$ the completion of the space of $\mathfrak{g}$–valued i–forms

over Σ' under the norm

$$\|a\|_{L^{2,\epsilon}_k} = (\int_{\overline{\Sigma}} e^{\epsilon\chi}(|a|^2 + |\nabla a|^2 + \cdots + |\nabla^k a|^2))^{\frac{1}{2}},$$

where χ is a smooth function on $\overline{\Sigma}$ which is equal to 1 over $S^1 \times [0, +\infty)$ and ∇ is some background connection. This is the norm associated to the inner product

$$\langle a_1, a_2\rangle_{k,\epsilon} = \int_{\overline{\Sigma}} e^{\epsilon\chi}(\langle a_1, a_2\rangle + \langle \nabla a_1, \nabla a_2\rangle + \cdots + \langle \nabla^k a_1, \nabla^k a_2\rangle).$$

Note that the bilinear form on $L^{2,\epsilon}_k(\Omega^i(\overline{\Sigma}) \otimes \mathfrak{g})$ given by

$$\omega_{\overline{\Sigma}}(a_1, a_2) = \int_{\overline{\Sigma}} \mathrm{Tr}(a_1 \wedge a_2)$$

is non-degenerate, which could be proved using the Hodge $*$–operator. Given $A \in \mathcal{A}^0_F(\Sigma)$, it naturally defines a connection on $P|_{\overline{\Sigma}}$, namely extending A by the product connection on the cylindrical end. Therefore, we can define the operator

$$d_A : L^{2,\epsilon}_k(\Omega^i(\overline{\Sigma}) \otimes \mathfrak{g}) \to L^{2,\epsilon}_{k-1}(\Omega^{i+1}(\overline{\Sigma}) \otimes \mathfrak{g})$$

using the connection A the the L^2–formal adjoint operator of d_A

$$d^*_A : L^{2,\epsilon}_k(\Omega^i(\overline{\Sigma}) \otimes \mathfrak{g}) \to L^{2,\epsilon}_{k-1}(\Omega^{i-1}(\overline{\Sigma}) \otimes \mathfrak{g}).$$

Lemma 2.3.4. *For $\epsilon > 0$ sufficiently small, the kernel of the map*

$$d^*_A \oplus d_A : L^{2,\epsilon}_k(\Omega^1(\overline{\Sigma}) \otimes \mathfrak{g}) \to L^{2,\epsilon}_{k-1}(\Omega^0(\overline{\Sigma}) \otimes \mathfrak{g}) \oplus L^{2,\epsilon}_{k-1}(\Omega^2(\overline{\Sigma}) \otimes \mathfrak{g})$$

is naturally isomorphic to H^1_A and this isomorphism respects the symplectic structure.

Proof. The map $\ker(d_A \oplus d^*_A) \to H^1_A$ is given by the natural inclusion $L^{2,\epsilon}_k(\Omega^1(\overline{\Sigma}) \otimes \mathfrak{g}) \to$

$L^2_k(\Omega^1(\Sigma))$, after identifying $\overline{\Sigma}$ with $\Sigma\backslash\{p\}$. This map is the same as the composition $H^1_A(\Sigma', \partial\Sigma') \cong H^1_c(\Sigma', d_A) \to H^1_A$, which is an isomorphism by counting dimensions. This map obviously respects the symplectic structure. □

Lemma 2.3.5. *There exists a map $P_0 : \ker(d_A \oplus d_A^*) \to \tilde{H}_A^{1,\mathfrak{g}}$ which defines a linear symplectic embedding. The symplectic orthogonal complement of $\mathrm{Im}(P_0)$ is isomorphic to $(\mathfrak{g}/H^0_A)\oplus(\mathfrak{g}/H^0_A)^*$, where H^0_A is identified with the Lie algebra of the stabilizer group of $A \in \mathcal{A}_\Sigma$. Moreover, the induced symplectic structure on $(\mathfrak{g}/H^0_A)\oplus(\mathfrak{g}/H^0_A)^*$ from $\omega_{\Sigma'}$ is the same as the canonical linear symplectic form on $T^*(\mathfrak{g}/H^0_A) \cong (\mathfrak{g}/H^0_A)\oplus(\mathfrak{g}/H^0_A)^*$.*

Proof. The proof is essentially the same as the proof of [Jef94, Proposition 3.9]. In *loc.cit*, Jeffrey constructed a family of linear symplectic embeddings $P_r : \ker(d_A \oplus d_A^*) \to \tilde{H}_A^{1,\mathfrak{g}}$ parametrized by $r \in [0, +\infty)$ and the P_0 here is the specialization at $r = 0$. Although [Jef94, Proposition 3.9] was stated for $G = \mathrm{SU}(2)$, the same arguments work for general Lie groups. □

Using these two lemmas, and note that all the isomorphisms commute with the $\mathrm{Stab}(A)$–action, we conclude

Corollary 2.3.6. *Suppose $[A] \in \hat{\mathcal{M}}^{\mathfrak{g}}(\Sigma')$ is represented by some $A \in \mathcal{A}^0_F(\Sigma)$, then there is a $\mathrm{Stab}(A)$–equivariant symplectic isomorphism*

$$T_{[[A]]}\hat{\mathcal{M}}^{\mathfrak{g}}(\Sigma') \cong H^1_A \oplus T^*(\mathfrak{g}/H^0_A), \tag{2.3.4}$$

where $T^*(\mathfrak{g}/H^0_A)$ is equipped with the canonical symplectic form and trivial $\mathrm{Stab}(A)$–action. □

Remark 2.3.7. The above corollary could be viewed as a consequence of the symplectic slice theorem which appeared in the proof of [SL91, Theorem 2.1]. However, we cannot apply the slice theorem directly because we need to identify the tangent bundle of

the normal slice explicitly as twisted harmonic 1–forms. This will be important in our calculations of spectral flows later.

Now suppose Y' is a smooth, oriented, 3–manifold with boundary $\partial Y' = \Sigma$. Choose the base point $p \in \Sigma$ as above then p also define base points in both Y'. Assume that $\pi_1(\Sigma) \to \pi_1(Y')$ induced by the inclusion defines a surjective map. Denote by P the trivial G–bundle on both Y' and Σ. Let $\mathcal{A}^{\mathfrak{g}}_F(\Sigma'|Y') \subset \mathcal{A}^{\mathfrak{g}}_F(\Sigma')$ be the space consisting of elements in $\mathcal{A}^{\mathfrak{g}}_F(\Sigma')$ which extend to flat connection on P over Y'. Such space is preserved by the action of $\mathcal{G}^c_{\Sigma'}$. Furthermore, let $\mathcal{A}_F(Y')$ be the space of flat connections on $P|_{Y'}$ and let $\mathcal{G}_0(Y')$ be the based gauge group consisting of gauge transformations $g \in \mathcal{G}_{Y'}$ such that $g(p) = \mathrm{id}$. Then the map

$$\mathcal{A}_F(Y')/\mathcal{G}_0(Y') \to \mathcal{A}^{\mathfrak{g}}_F(\Sigma'|Y')/\mathcal{G}^c_{\Sigma'}$$

is a diffeomorphism. By [MW12, Lemma 5.1], the embedding

$$\mathcal{A}_F(Y')/\mathcal{G}_0(Y') \cong \mathcal{A}^{\mathfrak{g}}_F(\Sigma'|Y')/\mathcal{G}^c_{\Sigma'} \hookrightarrow \mathcal{A}^{\mathfrak{g}}_F(\Sigma')/\mathcal{G}^c_{\Sigma'} = \mathcal{M}^{\mathfrak{g}}(\Sigma')$$

defines a Lagrangian submanifold $L_{Y'}$ in $\mathcal{M}^{\mathfrak{g}}(\Sigma')$ with respect to the symplectic form in Proposition 2.3.2. The manifold $L_{Y'}$ diffeomorphic to G^h and it is preserved by the G–action on $\mathcal{M}^{\mathfrak{g}}(\Sigma')$. Additionally, this Lagrangian is contained in the 0–level set of the moment map $\mu^{-1}(0)$. As a result, $L_{Y'}$ is contained in the smooth locus $\hat{\mathcal{M}}^{\mathfrak{g}}(\Sigma')$.

One could obtain many G–invariant Lagrangians in $\hat{\mathcal{M}}^{\mathfrak{g}}(\Sigma')$ by the construction above. We will study the equivariant intersection theory of a pair of such Lagrangians and define some version of equivariant intersection number which calculates the generalized Casson invariants.

2.4 Equivariant transversality of Lagrangian submanifolds

This section develops an equivariant transversality result for Lagrangian submanifolds inside symplectic manifolds equipped with Lie group actions, and defines the notion *equivariant Maslov index* which will be used in the calculation of the generalized Casson invariants.

Suppose (M, ω) is a (not necessarily closed) smooth symplectic manifold of dimension $2n$, where ω is the symplectic form . Let G be a connected, compact Lie group that acts on (M, ω) such that $\forall g \in G$, $g^*\omega = \omega$. Let L_1, L_2 be two oriented Lagrangian submanifolds in M which are invariant under the G–action.

For $p \in M$, we use $\mathrm{Orb}(p)$ to denote the orbit of p under the action of G. Note that $\mathrm{Orb}(p)$ could be identified with $G/\operatorname{Stab}(p)$, where $\operatorname{Stab}(p)$ is the stabilizer group of p under the G–action. In particular, $\mathrm{Orb}(p)$ is a smooth manifold. A subset X of M is called *G–invariant* if $G(X) = X$; and similarly, a subbundle E of $TM|_X$ where X is G–invariant is called *G–invariant* if $G(E) = E$.

Definition 2.4.1. *We say that L_1 and L_2 intersect* non-degenerately *at p, if for $p \in L_1 \cap L_2$, we have $T_pL_1 \cap T_pL_2 = T_p \operatorname{Orb}(p)$. L_1 and L_2 are said to intersect* non-degenerately *if they intersect* non-degenerately *at each of their intersection point.*

Remark 2.4.2. Using the terminology in [Poz99], the above definition is equivalent to saying that L_1 and L_2 *have clean intersection along* $\mathrm{Orb}(p)$ *for each* $p \in L_1 \cap L_2$. Note that if L_1 and L_2 are compact and intersect non-degenerately, they only intersect along finitely many G–orbits.

2.4.1 Transversality

Suppose $(M, \omega), L_1, L_2$ are as above, we prove that L_1 and L_2 intersect non-degenerately after applying a generic G–equivariant Hamiltonian perturbation on L_1. We also

characterize the failure of equivariant transversality when varying the Hamiltonian perturbations in generic 1–paramter families.

We start by fixing some notations from symplectic geometry. Let $C^\infty_G(M)$ be the space of compactly supported smooth G–invariant time-independent Hamiltonians on the symplectic manifold M equipped with the C^∞–topology.

Definition 2.4.3. *Let $C^\infty_0\Big([0,1], C^\infty_G(M)\Big)$ be the space of time-dependent smooth G–invariant Hamiltonians which vanish near $s = 0$ and $s = 1$, where s is the coordinate on $[0,1]$. Take $\{H^0_s, H^1_s, \cdots\}$ which form a countable dense subset of $C^\infty_0\Big([0,1], C^\infty_G(M)\Big)$ and let*

$$N_m := \sup\{\|H^0_s\|_{C^m}, \cdots, \|H^m_s\|_{C^m}\}.$$

Define $\mathcal{H}$ to be the completion of the subspace of $C^\infty_0\Big([0,1], C^\infty_G(M)\Big)$ consisting of functions of form

$$\sum_{m\geq 0} a_m H^m_s$$

such that $\sum_{m\geq 0} N_m|a_m| < +\infty$ under the norm $\|\sum_{m\geq 0} a_m H^m_s\| := \sum_{m\geq 0} N_m|a_m|$.

Given any $H_s \in C^\infty_0\Big([0,1], C^\infty_G(M)\Big)$, the Hamiltonian vector field X_{H_s} of H_s is determined by

$$dH_s = \omega(\cdot, X_{H_s}).$$

Let Φ_{H_s} be the time 1–flow generated by integrating X_{H_s}. Then for any G–invariant Lagrangian $L \subset M$, $\Phi_{H_s}(L)$ is also a G–invariant Lagrangian.

Recall that for every manifold S, there is a canonical symplectic form $d\lambda$ on T^*S (see, for example, [MS17, p. 90]). When S admits a smooth G–action, $d\lambda$ is preserved by the induced G–action on T^*S. We need the following equivariant local version of the Lagrangian neighborhood theorem:

Lemma 2.4.4. *Suppose L is a G–invariant Lagrangian submanifold of (M, ω), and suppose* $\mathrm{Orb}(p) \subset L$. *Let $E \subset TM|_{\mathrm{Orb}(p)}$ be a G–invariant Lagrangian subbundle that is transverse to $TL|_{\mathrm{Orb}(p)}$. Then there exists a G–invariant open subset $U \subset M$ with the following property: if we write $S = U \cap L$, then there is a G–equivariant symplectomorphism φ from an open neighborhood of the zero section of T^*S to U, such that*

1. *φ equals the identity map on S,*
2. *the tangent map of φ on S sends the fibers of T^*S to the fibers of E on* $\mathrm{Orb}(p)$.

Proof. Because E and $TL|_{\mathrm{Orb}(p)}$ are transverse to each other fiberwisely as Lagrangian subbundles of $TM|_{\mathrm{Orb}(p)}$, the map $\beta_E : E \to T^*L|_{\mathrm{Orb}(p)}$ given by $v \mapsto \omega(v, \cdot)$ is an isomorphism. Using the identification

$$TM|_{\mathrm{Orb}(p)} \cong E \oplus TL|_{\mathrm{Orb}(p)} \cong T^*L|_{\mathrm{Orb}(p)} \oplus TL|_{\mathrm{Orb}(p)},$$

we can construct a fiberwise G–invariant almost complex structure J_p on $TM|_{\mathrm{Orb}(p)}$ which is compatible with the fiberwise symplectic form on $TM|_{\mathrm{Orb}(p)}$ such that $J_p(TL|_{\mathrm{Orb}(p)}) = E$.

Extend J_p to a G–equivariant compatible almost complex structure J on M and let $g_J = \omega(\cdot, J\cdot)$ be the compatible Riemannian metric. After choosing sufficiently small G–invariant neighborhood of $\mathrm{Orb}(p) \in T^*L$ so that exponential map of g_J is well defined, the argument of [MS17, Theorem 3.4.13] carries over. Our special choice of the almost complex structure guarantees the second condition holds. □

Lemma 2.4.5. *Let S be a G–manifold and suppose that U is a sufficiently small G–invariant neighborhood of* $\mathrm{Orb}(p) \subset S$. *Then any G–invariant Lagrangian $L \subset T^*S$ which is transverse to the Lagrangian foliation on T^*U induced by cotangent fibers is locally presented as the graph of an exact G-equivariant 1–form over U.*

Proof. Because L is transverse to the Lagrangian foliation on T^*U induced by cotangent fibers, it is locally represented by the graph of some 1–form θ_L over U. The Lagrangian condition implies that θ_L is closed.

Using the slice theorem, without loss of generality we can assume that $U = G \times_{\mathrm{Stab}(p)} U'$ where U' is a contractile G–invariant neighborhood of the origin in $T_pS/T_p \operatorname{Orb}(p)$. Because the quotient of U by G is contractible, the cohomology class represented by the equivariant form θ_L is 0. Using the de Rham complex of G–equivariant differential forms on U, we conclude that θ_L is exact. □

Now we are ready to establish the equivariant transversality result. Let $A \subset M$ be a compact subset such that $L_1 \cap L_2$ is included in the interior of A. Let $\mathcal{H}_0 \subset \mathcal{H}$ be a sufficiently small open ball centered at zero, such that for every $H_s \in \mathcal{H}_0$, the map Φ_{H_s} is defined on A, and $\Phi_{H_s}(L_1) \cap L_2$ is included in the interior of A.

Suppose $\sigma \in \mathcal{R}_G$. Define $\mathcal{M}_\sigma \subset \mathcal{H}_0 \times (A \cap L_2)$ to be the set of points (H_s, p) such that

1. $p \in \Phi_{H_s}(L_1) \cap L_2$,
2. The linear space $\big(T_p\Phi_{H_s}(L_1) \cap T_p(L_2)\big)/T_p \operatorname{Orb}(p)$ as a $\mathrm{Stab}(p)$–representation represents σ.

Proposition 2.4.6. *The space $\mathcal{M}_\sigma$ is a Banach manifold. The projection of M_σ to $\mathcal{H}_0$, written as π_σ, is Fredholm with index $-d(\sigma)$.*

Proof. Let $(H_s, p) \in \mathcal{M}_\sigma$. Let $Q_p \subset TM|_{\mathrm{Orb}(p)}$ be a G–invariant Lagrangian subbundle such that Q_p is transverse to $T\Phi_{H_s}(L_1)|_{\mathrm{Orb}(p)}$ and $TL_2|_{\mathrm{Orb}(p)}$ along $\operatorname{Orb}(p)$. Let B_2 be a small G–invariant open neighborhood of p in L_2. Then by Lemma 2.4.4, there exists be a G–invariant open neighborhood B of p in M, such that

1. There exists a G–equivariant symplectomorphism to a neighborhood of the zero section of T^*B_2,

2. $T_p^* B_2$ is tangent to Q_p at p under the above diffeomorphism.

After further shrinking B_2 and B if necessary, by Lemma 2.4.5, $\Phi_{H_s}(L_1) \cap B$ is the graph of the differential of a G–invariant function on $\overline{B_2}$. This function is uniquely determined if we require that its integration over $\overline{B_2}$ is 0. Therefore we have a map from a sufficiently small neighborhood $N(H_s)$ of H_s to $C_G^\infty(\overline{B_2})$

$$\pi : N(H_s) \to C_G^\infty(\overline{B_2})$$

such that for each $H'_s \in N(H)$, we have $\Phi_{H'_s}(L_1) \cap \overline{B}$ is the graph of $d\pi(H'_s)$, and $\int_{\overline{B_2}} \pi(H'_s) = 0$.

Claim: The map π is a submersion at H_s onto the subspace

$$\big\{ g \in C_G^\infty(\overline{B_2}) | \int_{\overline{B_2}} g = 0 \big\}.$$

For each $g \in C_G^\infty(\overline{B_2})$ such that $\int_{\overline{B_2}} g = 0$, let $\hat{g} \in C_G^\infty(M)$ be a smooth, compactly supported function on M, such that $\hat{g}|_B$ is given by the pull-back of g from B_2 to $T^* B_2$. Let $\chi : [0,1] \to \mathbb{R}^{\geq 0}$ be a smooth, compactly supported function such that $\int \chi = 1$. For $s \in [0,1]$, let $\Phi_{H_s}(s)$ be the integration of X_{H_s} from s to 1. Define

$$\hat{H}_s = \chi(s) \cdot \big(\hat{g} \circ \Phi_{H_s}(s) \big).$$

Then

$$\frac{d}{dv} \pi(H + v\hat{H}) = g.$$

Therefore the claim is proved.

Let D be a slice of B_2 under the G–action through p. Let D^0 be the fixed point subset of D under the $\mathrm{Stab}(p)$–action. Then for any $(\Phi_{H'_s}, q) \in \mathcal{M}_\sigma$ such that $H'_s \in N(H_s)$ and $q \in B$, we have $q \in D^0$ where D^0 is viewed as a subset of the zero section of the Weinstein neighborhood. Choose a G–invariant Riemannian metric on

B_2 and let S_p be the orthogonal complement of $T_p \operatorname{Orb}(p)$ in $T_p B_2$ and let S_p^0 be the fixed point subset of S_p under the $\operatorname{Stab}(p)$–action. Using parallel transports, we can identify TD with S_p. Now we can define a map

$$\psi : N(H_s) \times D^0 \to (S_p^0)^* \times \operatorname{Sym}_{\operatorname{Stab}(p)}(S_p)$$

$$(H_s', q) \mapsto (\nabla_q(\pi(H_s'))|_{D^0}, \operatorname{Hess}(\pi(H_s'))|_{S_p}),$$

where ∇ and Hess stand for the gradient and Hessian respectively. It is easy to see that

$$\mathcal{M}_\sigma = \psi^{-1}(\{0\} \times \operatorname{Sym}_{\operatorname{Stab}(p),\sigma}(S_p)).$$

Using the abundance result [BZ20, Lemma 2.25] and the result established in the claim, we conclude that the differential of ψ at (H_s, p) is surjective. Therefore $\mathcal{M}_\sigma$ is a Banach manifold near (H_s, p). Because the inclusion

$$\psi^{-1}(\{0\} \times \operatorname{Sym}_{\operatorname{Stab}(p),\sigma}(S_p)) \hookrightarrow N(H_s) \times D^0$$

is a Fredholm map of index $-\dim(D^0) - d(\sigma)$, the projection map from $\mathcal{M}_\sigma$ to $\mathcal{H}$ is Fredholm of index $-d(\sigma)$. The proposition is proved because of the separability of $\mathcal{H}_0 \times (A \cap L_2)$. □

Corollary 2.4.7. *Let L_1 and L_2 be two G–invariant Lagrangians in (M, ω) such that their intersection is contained in a compact subset of M. Let $\mathcal{H}_0$ be a small neighborhood of 0 in $\mathcal{H}$. Then for a generic perturbation $H_s \in \mathcal{H}_0$, the pair $\Phi_{H_s}(L_1)$ and L_2 intersect non-degenerately.*

Proof. Using Proposition 2.4.6, because $d(\sigma) > 0$ as long as σ is not given by the 0–representation, $\mathcal{H}_\sigma := \pi_\sigma(\mathcal{M}_\sigma)$ is a C^∞–subvariety of $\mathcal{H}^0$ of positive codimension.

As a result,

$$\bigcup_{\sigma \in \mathcal{R}_G} \mathcal{H}_\sigma$$

is a meager subset of $\mathcal{H}_0$. Any Hamiltonian perturbation constructed from elements in $\mathcal{H}_0 \backslash \bigcup_{\sigma \in \mathcal{R}_G} \mathcal{H}_\sigma$ establishes the transversality result. □

Now we show that for a generic 1–paramter family $\{H_{s,t}\}_{0 \leq t \leq 1}$, the intersection $\Phi_{H_{s,t}}(L_1) \cap L_2$ has at most one degenerate intersection orbit which corresponds to an irreducible representation for each t.

Choose $\mathcal{H}_0$ and A as before. Suppose $\sigma_1, \sigma_2 \in \mathcal{R}_G$. Define $\mathcal{M}_{\sigma_1,\sigma_2} \subset \mathcal{H}_0 \times (A \cap L_2) \times (A \cap L_2)$ to be the set of points (H_s, p_1, p_2) such that

1. $p_1, p_2 \in \Phi_{H_s}(L_1) \cap L_2$,
2. $\mathrm{Orb}(p_1) \cap \mathrm{Orb}(p_2) = \emptyset$,
3. The spaces $\big(T_{p_i}\Phi_{H_s}(L_1) \cap T_{p_i}(L_2)\big)/T_{p_i}\,\mathrm{Orb}(p_i)$ as $\mathrm{Stab}(p_i)$–representations represent σ_i for $i = 1, 2$.

Let $\mathcal{H}_{\sigma_1,\sigma_2}$ be the projection of $\mathcal{M}_{\sigma_1,\sigma_2}$ onto $\mathcal{H}_0$. The following statement follows from the arguments in the proof of Proposition 2.4.6 and [BZ20, Lemma 3.23], thus the proof is omitted.

Proposition 2.4.8. *$\mathcal{H}_{\sigma_1,\sigma_2}$ is a C^∞–subvariety of $\mathcal{H}_0$ of codimension at least $d(\sigma_1) + d(\sigma_2)$.* □

Note that $d(\sigma) = 1$ if and only if σ is irreducible. Therefore the following corollary follows immediately from Proposition 2.4.6 and Proposition 2.4.8 by an application of the Sard-Smale theorem.

Corollary 2.4.9. *Let $\{H_{s,t}\}_{0 \leq s \leq 1}$ be a generic path in $\mathcal{H}_0$ connecting $H_{s,0}, H_{s,1} \in \mathcal{H}_0 \backslash \bigcup_{\sigma \in \mathcal{R}_G} \mathcal{H}_\sigma$ which intersect all $\mathcal{H}_\sigma$ and $\mathcal{H}_{\sigma_1,\sigma_2}$ transversely. Then there are at most countably many t such that $\Phi_{H_{s,t}}(L_1)$ do not intersect L_2 non-degenerately. For every*

such $t \in (0,1)$*, there is exactly one intersection orbit* $\mathrm{Orb}(p_t)$ *such that* $\big(T_{p_t}\Phi_{H_{s,t}}(L_1) \cap T_{p_t}(L_2)\big)/T_{p_t}\,\mathrm{Orb}(p_t)$ *is an irreducible* $\mathrm{Stab}(p_t)$*–representation, and* $\Phi_{H_{s,t}}(L_1)$ *intersect* L_2 *non-degenerately along any other orbits.* □

2.4.2 An equivariant Maslov index

In this subsection, we give the definition of equivariant Maslov index which will appear in the index formula later on.

We recall the definition of the classical Maslov index. Let (V_0, ω_0) be a symplectic vector space. For $a, b \in \mathbb{R}$ such that $a < b$, let $\mathcal{P}_{[a,b]}(V_0)$ be the space of continuous and piecewise smooth maps

$$q : [a,b] \to \{\text{paris of Lagrangian subspaces in } V_0\}.$$

Then there exists a map

$$\mu_{V_0} : \mathcal{P}_{[a,b]}(V_0) \to \mathbb{Z}$$

which is called the *Maslov index* uniquely characterized by the following properties (see [CLM94]):

1. (Affine scale invariance) Suppose $k > 0$ and $l \geq 0$ and let $\Gamma : [a,b] \to [ka+l, kb+l]$ be the affine linear map given by $t \mapsto kt + l$. Then for $q \in \mathcal{P}_{[ka+l,kb+l]}(V_0)$, we have
$$\mu_{V_0}(q) = \mu_{V_0}(q \circ \Gamma);$$

2. (Homotopy invariance) Given $q_0, q_1 \in \mathcal{P}_{[a,b]}(V_0)$ such that $q_0(a) = q_1(a)$ and $q_0(b) = q_1(b)$, if there is a homotopy between q_0 and q_1 in $\mathcal{P}_{[a,b]}(V_0)$ relative to end points, then
$$\mu_{V_0}(q_0) = \mu_{V_0}(q_1);$$

3. (Additivity under concatenation of paths) If $a < c < b$ and $q \in \mathcal{P}_{[a,b]}(V_0)$, then the following additive property is satisfied:

$$\mu_{V_0}(q) = \mu_{V_0}(q|_{[a,c]}) + \mu_{V_0}(q|_{[c,b]});$$

4. (Symplectic additivity) Let (V_0', ω_0') be another symplectic vector space and suppose $q \in \mathcal{P}_{[a,b]}(V_0), q' \in \mathcal{P}_{[a,b]}(V_0')$. Then

$$\mu_{V_0 \oplus V_0'}(q \oplus q') = \mu_{V_0}(q) + \mu_{V_0'}(q');$$

5. (Independence of symplectic framing) Suppose $\phi_t : [a,b] \to \mathrm{Sp}(V_0)$ is a 1–parameter family of symplectic matrices. Denote by $q(t) = (L_1(t), L_2(t))$ for some $q \in \mathcal{P}_{[a,b]}(V_0)$ and define $\phi_*(q)(t) = (\phi_t(L_1(t)), \phi_t(L_2(t)))$. Then

$$\mu_{V_0}(\phi_*(q)(t)) = \mu_{V_0}(q(t));$$

6. (Normalization) Equip $\mathbb{R}^2_{x,y} \cong \mathbb{C}_{z=x+iy}$ with the symplectic form $dx \wedge dy = \frac{i}{2} dz \wedge d\bar{z}$. Consider the path of Lagrangians $p(t)$ in $\mathcal{P}_{[-\frac{\pi}{4}, \frac{\pi}{4}]}(\mathbb{R}^2)$ given by

$$t \mapsto (\mathbb{R}\{1\}, \mathbb{R}\{e^{it}\}).$$

Then

$$\mu_{\mathbb{R}^2}(p|_{[-\frac{\pi}{4}, \frac{\pi}{4}]}) = 1,\ \mu_{\mathbb{R}^2}(p|_{[-\frac{\pi}{4}, 0]}) = 0,\ \mu_{\mathbb{R}^2}(p|_{[-\frac{\pi}{4}, 0]}) = 1.$$

There are many ways to construct the map μ_{V_0} explicitly. For our purpose later, we use the characterization of μ_{V_0} as the spectral flow of certain self-adjoint operators, see [CLM94, Section 7]. Let $J_0 \in GL(V_0)$ be a complex structure which is compatible with ω_0. Given $L_1, L_2 \subset V_0$ which are linear Lagrangian subspaces, let $L_1^2([0,1]; L_1, L_2)$

be the completion of smooth maps $f : [0,1] \to V$ with $f(0) \in L_1, f(1) \in L_2$ under L^2_1–norm induced by the metric $\omega(\cdot, J_0\cdot)$. Let $L^2([0,1]; V_0)$ be the space of L^2–maps from $[0,1]$ to V_0. Then the operator

$$D_{L_1,L_2} : L^2_1([0,1]; L_1, L_2) \to L^2([0,1]; V_0)$$

$$f \mapsto -J_0 \frac{df}{dt}$$

is a real self-adjoint unbounded Fredholm operator with discrete spectrum. The kernel of D_{L_1,L_2} is the same as constant functions taking value in $L_1 \cap L_2$.

If we take a path of pairs of Lagrangian subspaces $q(t) = (L_1(t), L_2(t))_{t\in[a,b]}$ in V_0, we have a family of self-adjoint operators $\{D_{L_1(t),L_2(t)}\}_{a\le t\le b}$. Let $\epsilon > 0$ be small enough so that $D_{L_1(a),L_2(a)}$ and $D_{L_1(b),L_2(b)}$ do not have any other eigenvalues between $-\epsilon$ and ϵ except for 0.

Proposition 2.4.10. *The spectral flow of* $\{D_{L_1(t),L_2(t)} - \epsilon\,\mathrm{id}\}_{a\le t\le b}$ *is independent of the choice of* J_0 *and equal to* $\mu_{V_0}(p)$. □

Now suppose V_0 is a H–representation where H is a connected compact Lie group, and the H–action preserves the symplectic form ω_0. Let

$$V_0 \cong (V_0^1)^{\oplus a_1} \oplus \cdots \oplus (V_0^m)^{\oplus a_m}$$

be the isotypic decomposition of V_0. Then each isotypic piece inherits a symplectic structure and the above decomposition is compatible with the symplectic structure. If $q(t) = (L_1(t), L_2(t))_{t\in[a,b]}$ is a path such that both $L_1(t)$ and $L_2(t)$ are preserved under the H–action, then

$$(V_0^1)^{\oplus a_1} \cap L_i(t) \oplus \cdots \oplus (V_0^m)^{\oplus a_m} \cap L_i(t)$$

gives the isotypic decomposition of $L_i(t)$ for $i = 1, 2$ and $t \in [a, b]$. Write $q_k(t)$ as the induced path on the isotypic piece $(V_0^k)^{\oplus a_k}$ for $1 \leq k \leq m$.

Definition 2.4.11. *The* equivariant Maslov index $\mu^H(q)$ *of the path* $q(t)$ *as above is the element*

$$\sum_{i=1}^{m} \frac{1}{\dim_{\mathbb{R}} V_i} \cdot \mu_{(V_0^i)^{\oplus a_i}}(q_i) \cdot [V_i] \in \mathcal{R}(H). \tag{2.4.1}$$

Remark 2.4.12. The above quantity is well-defined because the V_i–isotypic piece of $L_1(t) \cap L_2(t)$ would change the Maslov index by a multiple of $\dim_{\mathbb{R}} V_i$. Alternatively, we can define $\mu^H(q)$ using the equivariant spectral flow of the H–equivariant operator $-J_0 \frac{d}{dt}$ where J_0 is an H–equivariant complex structure on V_0 compatible with ω_0, based on Definition 2.2.10 and Proposition 2.4.10.

Now we go back to discuss about global symplectic geometry. Recall that (M, ω) is a symplectic manifold with a symplectic G–action and L_1 and L_2 are two G–invariant Lagrangians of M. Suppose H is a closed subgroup of G. Let $\overline{\mathbb{D}} \subset \mathbb{C}_z = \mathbb{R}^2_{x,y}$ be the closed unit disc in the plane. Suppose we have a map $u : \overline{\mathbb{D}} \to M$ which is smooth except at ± 1 satisfying:

1. $u(\partial\overline{\mathbb{D}} \cap \{y \leq 0\}) \subset L_1$ and $u(\partial\overline{\mathbb{D}} \cap \{y \geq 0\}) \subset L_2$;

2. For any $z \in \overline{\mathbb{D}}$, the Lie group $\mathrm{Stab}(u(z))$ contains H as a subgroup.

Choose $p = u(-1)$ as the base point. Then the symplectic vector bundle u^*TM is isomorphic to the product bundle $T_pM \times \overline{\mathbb{D}}$ respecting the fiberwise symplectic H–action under a trivialization Ψ. Using Ψ,

$$q_u = (u^*|_{\partial\overline{\mathbb{D}} \cap \{y \leq 0\}} TL_1, u^*|_{\partial\overline{\mathbb{D}} \cap \{y \geq 0\}} TL_2)$$

defines a path of pair of H–invariant Lagrangian subspaces in T_pM, after identifying $\partial\overline{\mathbb{D}} \cap \{y \leq 0\}$ and $\partial\overline{\mathbb{D}} \cap \{y \geq 0\}$ with the interval $[-1, 1]$ by projecting to the x–coordinate. Recall that $\mathfrak{h}$ is the Lie algebra of H. Because $H \subset \mathrm{Stab}(u(z))$ for

all $z \in \overline{\mathbb{D}}$, the symplectic slice theorem implies that $T_{u(z)}M$ contains the symplectic subspace $T^*(\mathfrak{g}/\mathfrak{h})$, i.e. the tangent space of the cotangent bundle of $\mathrm{Orb}(u(z))$. The tangent spaces of $T_{u(z)}L_i, z \in \partial\mathbb{D}$ contain a copy of $\mathfrak{g}/\mathfrak{h}$ as a result of the slice theorem. Define the symplectic vector space

$$\overline{T_pM} = T_pM/T^*(\mathfrak{g}/\mathfrak{h})$$

with the induced symplectic structure and the path of pairs of linear Lagrangian subspaces

$$\overline{q}_u = (u^*|_{\partial\overline{\mathbb{D}}\cap\{y\leq 0\}}TL_1/(\mathfrak{g}/\mathfrak{h}), u^*|_{\partial\overline{\mathbb{D}}\cap\{y\geq 0\}}TL_2/(\mathfrak{g}/\mathfrak{h}))$$

which lies inside $\overline{T_pM}$.

Definition 2.4.13. *Define the* H–equivariant Maslov index $\mu^H(u) \in \mathcal{R}(H)$ *of the map* u *as above to be*

$$\mu^H(u) := \mu^H(\overline{q}_u). \tag{2.4.2}$$

Using the properties of the Maslov index, it is easy to see that $\mu^H(u)$ is independent of the trivialization Ψ. It is also invariant under the homotopies of u which satisfy (1) and (2) above and keep both $u(-1)$ and $u(1)$ fixed.

Example 2.4.14. Suppose V is an orthogonal representation of G, and suppose (M, ω) is given by the cotangent bundle of V. Let f be a smooth G–invariant function on V with finitely many critical orbits. Furthermore, we assume that f is a G–Morse function, i.e. the Hessian of f is non-degenerate on $T_pV/T_p\,\mathrm{Orb}(p)$ if p is a critical point of f. Let L_1 be the graph of the zero section of $M = T^*V$, let L_2 be the graph of df. Then it is easy to see that L_1 and L_2 intersect non-degenerately. Let $p \in V$ be critical point of f such that $H \subset \mathrm{Stab}(p)$. Let

$$V/T\,\mathrm{Orb}(p) = (V_1)^{\oplus a_1} \oplus \cdots \oplus (V_m)^{\oplus a_m}$$

be the isotypic decomposition of V as an H–representation. Then for each $1 \leq i \leq m$, the Hessian of f induces an invertible self-adjoint H–equivariant linear map

$$(V_i)^{\oplus a_i} \mapsto (V_i)^{\oplus a_i}$$

whose negative eigenspace has dimension denoted by $\dim_{\mathbb{R}} V_i \cdot \mathrm{ind}(p, V_i)$. Define

$$\mathrm{ind}_H(p) = \sum_{i=1}^{m} \mathrm{ind}(p, V_i) \cdot [V_i] \in \mathcal{R}(H).$$

Now consider that we have a pair of critical points p_1 and p_2 of f which do not lie on the same G–orbit and they are both fixed under the action of H. Then we can find a path γ in V which is invariant under the H–action connecting p_1 and p_2. Note that the time-t Hamiltonian flow of the function $f : T^*V \to \mathbb{R}$ is given by

$$(v, v') \to (v, v' + t \cdot df),$$

where we extend $f : V \to \mathbb{R}$ to a smooth function over T^*V which is constant along each fiber. The trace of the image of γ for $t \in [0, 1]$ defines a map $u_\gamma : \overline{\mathbb{D}} \to T^*V$ which has a well-defined Maslov index. It could be shown that

$$\mu^H(u_\gamma) = \mathrm{ind}_H(p_2) - \mathrm{ind}_H(p_1)$$

which is a classical computation of Lagrangian Floer homology of cotangent bundles.

Example 2.4.14 shows that in this simple case, the non-degeneracy of critical points of a G–invariant smooth function is equivalent to the non-degeneracy of intersection between a closely-related pair of Lagrangians. Moreover, the equivariant Maslov index could be computed as a relative equivariant Morse index. We extend these basic correspondences in the next section to infinite dimensions which eventually allow

us to interpret the generalized Casson invariants using finite-dimensional symplectic constructions.

2.5 The gauge-symplectic correspondence

This section proves the most important technical results of this chapter. Given a three manifold Y equipped with a Heegaard splitting $Y = H_1 \cup_\Sigma H_2$, we show that there is a bijection between gauge equivalence classes of perturbed flat G–connections on Y and intersection G–orbits of Lagrangians in $\hat{\mathcal{M}}^{\mathfrak{g}}(\Sigma')$ induced by H_1 and H_2 after applying a compatible G–equivariant Hamiltonian isotopy. Moreover, the non-degeneracy condition from gauge-theoretic context is shown to be equivalent to the non-degeneracy condition of Lagrangian intersections. We then prove several technical results which eventually reduce the computation of spectral flows on Y to the computation of Maslov indices on $\hat{\mathcal{M}}^{\mathfrak{g}}(\Sigma')$.

2.5.1 Perturbations

Let Y be a closed, oriented, Riemannian 3-manifold with a splitting

$$Y = H_1 \cup_\Sigma \big([-2,2] \times \Sigma\big) \cup_\Sigma H_2,$$

where H_1 and H_2 are handlebodies and the genus of Σ is $h \geq 2$. Assume further that the metric on $[-2,2] \times \Sigma$ is cylindrical. Let $\{p\} \in \Sigma$, D^2_p and Σ' be the same as in Section 2.3. Throughout this section, the product principal G–bundle over any manifold will be denoted by P. We introduce a family of cylinder functions on Y supported on $[-2,2] \times \Sigma$, which are closely related to Hamiltonian perturbations on the extended moduli space $\hat{\mathcal{M}}^{\mathfrak{g}}(\Sigma')$.

Suppose

$$\gamma_j : (\mathbb{R}/\mathbb{Z}) \times (-1,1) \to \Sigma' \subset \Sigma, \qquad j = 1, \cdots, m$$

are smooth immersions, and suppose there exists $\epsilon > 0$ such that the γ_j's coincide on $(-\epsilon, \epsilon) \times (-1, 1)$. Let

$$h_s : G^m \to \mathbb{R}, \qquad s \in [-1, 1]$$

be a 1–parameter family of smooth functions on G^m invariant under diagonal conjugations by G, such that h_s vanishes when s is near -1 and 1. Let

$$\beta : (-1, 1) \to \mathbb{R}$$

be a smooth, compactly supported function that integrates to 1. Define the function $H_s : \mathcal{A}(\Sigma)$ or $\mathcal{A}(\Sigma') \to \mathbb{R}$ by

$$H_s(A) = \int_{-1}^{1} \beta(\lambda) h_s(\rho_\lambda(A))\, d\lambda \tag{2.5.1}$$

for $s \in [-1, 1]$. It is easy to see that H_s is invariant under the gauge group actions. Using the affine linear map $x \mapsto 2x - 1$ identifying $[0, 1]$ and $[-1, 1]$, we can use the coordinate $\hat{s}$ on $[0, 1]$ to get the family $\{H_{\hat{s}}\}_{0 \le \hat{s} \le 1}$ from the family $\{H_s\}_{-1 \le s \le 1}$. This fits in with the convention set up in Section 2.4. The Hamiltonian $H_{\hat{s}}$ descends to a G–invariant time-dependent Hamiltonian on the moduli space $\hat{\mathcal{M}}^{\mathfrak{g}}(\Sigma')$. The induced Hamiltonian flow preserves the set $\mu^{-1}(0) \subset \hat{\mathcal{M}}^{\mathfrak{g}}(\Sigma')$.

Lemma 2.5.1. *Every time-dependent smooth G–invariant Hamiltonian on $\hat{\mathcal{M}}^{\mathfrak{g}}(\Sigma')$ could be constructed in this way.*

Proof. Choose $\{\alpha_i, \beta_i\}_{1 \le i \le h}$ to be loops with common base point in Σ' such that they generate the fundamental group of Σ'. They could be thickened into immersions of annuli of the form described above. For any $[A] \in \hat{\mathcal{M}}^{\mathfrak{g}}(\Sigma')$, it is uniquely determined by its holonomies around $\alpha_i, \beta_i, 1 \le i \le h$. This is the embedding described in Lemma 2.3.3. Thus all the G–invariant functions on $\hat{\mathcal{M}}^{\mathfrak{g}}(\Sigma')$ all come from restricting conjugation invariant functions on G^{2h}. Given a smooth 1–parameter family of conjugation

invariant functions on G^{2h}, denoted by $h_{\hat{s}}$, we can construct the function $H_{\hat{s}}$ using $h_{\hat{s}}$, the annuli and the cut-off function β. It is easy to see that the function on $\hat{\mathcal{M}}^{\mathfrak{g}}(\Sigma')$ induced from $H_{\hat{s}}$ is the same as $h_{\hat{s}}$. This proves the lemma. □

On the other hand, the maps $\{\gamma_j\}_{1\leq j\leq m}$ could be used to define smooth immersions

$$q_j : (\mathbb{R}/\mathbb{Z}) \times (-1,1)^2 \to \Sigma' \times [-1,1] \qquad j = 1, \cdots, m$$

simply by requiring that $q_j|_{\Sigma'\times\{s\}} = \gamma_j$ for all $1 \leq j \leq m$. Note that the tuple

$$\mathbf{q} = (q_1, \cdots, q_m, \beta, h_s)$$

defines a cylinder datum on $\Sigma' \times [-1,1] \subset \Sigma \times [-1,1]$ as in Definition 2.2.6. Because h_s vanishes for s near ± 1, we see that $\mathbf{q}$ defines a cylinder datum on Y. Therefore we have the function

$$f_{\mathbf{q}} : \mathcal{C}_Y \to \mathbb{R}$$

which defines a holonomy perturbation.

Definition 2.5.2. *The Hamiltonian $H_{\hat{s}} : \hat{\mathcal{M}}^{\mathfrak{g}}(\Sigma') \to \mathbb{R}$ and the cylinder function $f_{\boldsymbol{q}} : \mathcal{C}_Y \to \mathbb{R}$ constructed as above are called a* compatible pair.

Suppose $A \in \mathcal{C}(Y)$ is in temporal gauge on $[-1,1]\times\Sigma$, i.e. $\iota(\frac{\partial}{\partial s})(A) = 0$ where s is the coordinate on $[-1,1]$. It is well-known that every connection is gauge-equivalent to a connection in temporal gauge on $[-1,1]\times\Sigma$ in this case. Let A_s be the restriction of A to the slice $\{s\} \times \Sigma$. Then $A \in \mathcal{C}(Y)$ is a critical point of the perturbed Chern-Simons functional $\mathrm{CS} + f_{\mathbf{q}}$ if and only if

$$A \text{ is flat on the complement of } [-1,1] \times \Sigma, \tag{2.5.2}$$

$$F_{A_s} = 0 \text{ for } s \in [-1,1], \tag{2.5.3}$$

$$\dot{A} - X_s(A) = 0 \text{ for } s \in [-1, 1], \tag{2.5.4}$$

where $X_s(A)$ is a smooth map $X_s : \mathcal{A}(\Sigma) \to \Omega^1(\Sigma) \otimes \mathfrak{g}$ such that

$$dH_s(A)a = \int_\Sigma \mathrm{Tr}(a \wedge X_s(A)). \tag{2.5.5}$$

Notice that $X_s(A) \in \Omega^1(\Sigma) \otimes \mathfrak{g}$ is supported in the union of $\mathrm{Im}(\gamma_j)$.

Notice that $H_1 \cup [-2, -1] \times \Sigma$ and $[1, 2] \times H_2$ are manifolds with boundary both given by Σ. Let $L_1, L_2 \subset \mu^{-1}(0)$ be the corresponding G–invariant Lagrangian submanifolds of $\hat{\mathcal{M}}^{\mathfrak{g}}(\Sigma')$ as described in Section 2.3. Recall that $\Phi_{H_{\hat{s}}}$ is the time 1–map of the Hamiltonian flow induced by $H_{\hat{s}}$. Now we are ready to establish the following set-theoretic correspondence.

Proposition 2.5.3. *Let $H_{\hat{s}}$ and $f_{\boldsymbol{q}}$ be a compatible pair. Then there is a one-to-one correspondence between the orbits of perturbed flat connections on Y with respect to $f_{\boldsymbol{q}}$ and the G–orbits of $\Phi_{H_{\hat{s}}}(L_1) \cap L_2$. Moreover, a perturbed flat connection is non-degenerate if and only if the corresponding intersection of $\Phi_{H_{\hat{s}}}(L_1)$ and L_2 is non-degenerate.*

Proof. Throughout the proof, we write $\Phi_{H_s} = \Phi_{H_{\hat{s}}}$ and work with the s–coordinate.

Suppose A_s with $-1 \leq s \leq 1$ is a smooth 1-parameter family in $\mathcal{A}(\Sigma)$ that satisfies equation (2.5.3) and equation (2.5.4). Because A_s is flat, without loss of generality we can assume that $A_s|_{D^2_p}$ is given by the product connection after applying a gauge transformation. Then the image $[A_s] \in \hat{\mathcal{M}}^{\mathfrak{g}}(\Sigma')$ defines a smooth path. Recall that H_s is a family of smooth Hamiltonian functions on $\hat{\mathcal{M}}^{\mathfrak{g}}(\Sigma')$. Then $[A_s]$ is a Hamiltonian chord of H_s by equation (2.5.5) and the description of the symplectic form $\omega_{\Sigma'}$. Now if A is an $f_{\mathbf{q}}$–perturbed flat connection, by equation (2.5.2) $[A_{-1}]$ (resp. $[A_1]$) actually lies in L_1 (resp. L_2). By equivariance we construct a G–orbit in $\Phi_{H_s}(L_1) \cap L_2$.

On the other hand, suppose $[A_s]$ is a Hamiltonian chord of H_s on $\hat{\mathcal{M}}^{\mathfrak{g}}(\Sigma') \cap \mu^{-1}(0)$. Let A_s be a lift of $[A_s]$ to $\mathcal{A}^0_F(\Sigma)$, namely a family of flat connections on Σ which are

trivial over the disc D^2_p, then we have

$$\int_\Sigma \mathrm{Tr}(\dot{A}_s - X_s(A_s), a) = 0$$

for any s and any $a \in \Omega^1(\Sigma) \otimes \mathfrak{g}$ such that $d_{A_s} a = 0$. Therefore there exists ϕ_s such that

$$\dot{A}_s - X_s(A_s) = d_{A_s}\phi_s.$$

Hence A_s solves the equations (2.5.3) and (2.5.4) after a 1-parameter family of gauge transformations. The Lagrangian boundary condition lifts to (2.5.2). This proves the first part of the statement.

For the second part, suppose $A \in \mathcal{C}(Y)$ is a degenerate perturbed flat connection, then there exists $v \in \ker d^*_A \backslash \{0\}$ such that $K_{A,f_\mathfrak{q}} v = 0$. We have $[A_{-1}] \in L_1$, $[A_1] \in L_2$, and $\Phi_{H_s}(\bar{A}_{-1}) = A_1$. Restricting v to $\{\pm 1\} \times \Sigma$ defines twisted harmonic 1–forms in $H^1_{A_{\pm 1}}$ denoted by $v_{\pm 1}$. Under the identification

$$T_{[A_{\pm 1}]}\hat{\mathcal{M}}^\mathfrak{g}(\Sigma') \cong H^1_{A_{\pm 1}} \oplus T^*(\mathfrak{g}/H^0_{A_{\pm 1}})$$

from Corollary 2.3.6, it is easy to see that v_{-1} and v_1 define tangent vectors in $T_{[A_1]}L_1$ and $T_{[A_1]}L_2$ which are orthogonal to the directions corresponding to the G–orbits respectively. By linearizing equation (2.5.4) and integrate along the s–direction, we see that $(\Phi_{H_s})_*(v_{-1}) = v_1$. Therefore the intersection of $\Phi_{H_s}(L_1)$ and L_2 is degenerate at $[A_1]$.

On the other hand, suppose the intersection of $\Phi_{H_s}(L_1)$ and L_2 is degenerate, then we obtain a Hamiltonian chord $[A_s]$ of H_s connecting L_1 and L_2 and there exists $v_{\pm 1} \in T_{[A_{\pm 1}]}L_{\pm 1}$ such that $v_{\pm 1} \notin T\,\mathrm{Orb}([A_{\pm 1}])$ and $(\Phi_{H_s})_*(v_{-1}) = v_{+1}$. Therefore, we can identify $v_{\pm 1}$ as elements in $H_{A_{\pm 1}}$. Lift $v_{\pm 1}$ to $w_{\pm 1} \in \Omega^1(\Sigma) \otimes \mathfrak{g}$, then $w_\pm$ extends to $w \in \Omega^1(Y) \otimes \mathfrak{g}$ such that $w \in \ker(*d_A + dV_{f_\mathfrak{q}}(A))$ and $w \notin \mathrm{Im}\, d_A$. Decompose w

as $w = d_A u + w'$ where $d_A^* w' = 0$, we have $w' \in \operatorname{Hess} K_{A,f_{\mathbf{q}}}$, therefore A is degenerate as a perturbed flat connection. □

Using Corollary 2.4.7, Lemma 2.5.1 and Proposition 2.5.3, we conclude

Proposition 2.5.4. *For a generic choice of small Hamiltonian $H_{\hat{s}} : \mathcal{A}(\Sigma) \to \mathbb{R}$ constructed as above, the intersection between G–equivariant Lagrangians $\Phi_{H_{\hat{s}}}(L_1)$ and L_2 is non-degenerate. Accordingly, the perturbation of the Chern-Simons functional* CS *on Y induced by $f_{\boldsymbol{q}}$ constructed as above is non-degenerate.* □

Remark 2.5.5. The above statement shows that we can obtain equivariant transversality of CS by a perturbation supported in the neck region $[-1, 1] \times \Sigma$ instead of using loops wiggling all over the 3–manifold Y. Of course, this would not be a surprise because the inclusion $\Sigma \hookrightarrow Y$ defines a surjection between fundamental groups so a generic perturbation of the character variety of Y should result form a generic perturbation of it as a subset in the character variety of Σ.

2.5.2 Spectral flow comparison

Let $Y = H_1 \cup_\Sigma \big([-2, 2] \times \Sigma\big) \cup_\Sigma H_2$ and P be the same as before. Suppose H is a closed subgroup of the structural group G. We choose a generic Hamiltonian and the corresponding compatible holonomy perturbation $f_{\mathbf{q}}$ as from Proposition 2.5.4. Suppose $B_0, B_1 \in \mathcal{C}(Y)$ are two critical points of $\mathrm{CS} + f_{\mathbf{q}}$ such that $H = \mathrm{Stab}(B_1) \subset \mathrm{Stab}(B_0)$ and let $A_{s,0}, A_{s,1}$ be the corresponding path in $\mathcal{A}(\Sigma)$ satisfying equations (2.5.3) and (2.5.4). Let us further assume that we actually have a smooth 2–parameter family

$$A_{s,t} \in \mathcal{A}(\Sigma), \quad (s, t) \in [-1, 1] \times [0, 1]$$

connecting the 1–parameter families $A_{s,0}$ and $A_{s,1}$ such that the following holds:

$$F_{A_{s,t}} = 0 \text{ for } (s, t) \in [-1, 1] \times [0, 1]; \tag{2.5.6}$$

$$\text{For all } t \in [0,1], A_{-1,t} \text{ and } A_{1,t} \text{ extend to flat connections on} \\ H_1 \cup [-2,-1] \times \Sigma \text{ and } [1,2] \times \Sigma \cup H_2 \text{ respectively.} \tag{2.5.7}$$

Using the above two conditions, we can construct a 1–parameter family of connections $\{B_t\}_{0\leq t\leq 1}$ on Y connecting B_0 and B_1 which are in temporal gauge on $[-1,1]\times\Sigma$ by extending $A_{s,t}$. We further assume that

$$H \subset \mathrm{Stab}(B_t) \text{ for all } t \in [0,1]. \tag{2.5.8}$$

Then we have a 1–parameter family of self-adjoint operators $K_{B_t,f_{\mathbf{q}}}$ and it has a well-defined H–equivariant spectral flow $Sf_H(K_{B_t,f_{\mathbf{q}}})$ as in Definition 2.2.10. By applying a homotopy and reparametrization, we can assume that for all $t \in [0,1]$, the connection B_t is in temporal gauge on $[-2,2]\times\Sigma$ which defines the family $\{A_{s,t}\}_{-2\leq s\leq 2} \subset \mathcal{A}_\Sigma$ and this family is constant (i.e. independent of s) over $[-2,1]\times\Sigma$ and $[1,2]\times\Sigma$.

On the other hand, we can consider the space

$$\Omega^1(\Sigma)\otimes\mathfrak{g} \oplus \Omega^0(\Sigma)\otimes\mathfrak{g} \oplus \Omega^2(\Sigma)\otimes\mathfrak{g}.$$

It comes with a natural symplectic form

$$\tilde{\omega}_\Sigma : ((a_1,\phi_1,\psi_1),(a_2,\phi_2,\psi_2)) \mapsto \int_\Sigma \mathrm{Tr}(a_1\wedge a_2) + \int_\Sigma \mathrm{Tr}(\phi_1\wedge\psi_2) - \int_\Sigma \mathrm{Tr}(\psi_1\wedge\phi_2)$$

and a compatible complex structure $\tilde{J}_\Sigma$ given by

$$(a,\phi,\psi) \mapsto (*a, -*\psi, *\phi).$$

Given $A \in \mathcal{A}_F(\Sigma)$, we define S_A to be the *twisted de Rham operator* on $\Omega^1(\Sigma)\otimes\mathfrak{g} \oplus$

$\Omega^0(\Sigma)\otimes\mathfrak{g}\oplus\Omega^2(\Sigma)\otimes\mathfrak{g}$ given by

$$(a,\phi,\psi)\mapsto(-*d_A\phi-d_A*\psi,*d_Aa,d_A*a).$$

This is a self-adjoint operator with respect to the natural L^2–norm. Now consider the space

$$L^2_k([-1,1];\Omega^1(\Sigma)\otimes\mathfrak{g}\oplus\Omega^0(\Sigma)\otimes\mathfrak{g}\oplus\Omega^2(\Sigma)\otimes\mathfrak{g})$$

which is the L^2_k–completion of the space of smooth maps from $[-1,1]$ to $\Omega^0(\Sigma)\otimes\mathfrak{g}\oplus\Omega^1(\Sigma)\otimes\mathfrak{g}\oplus\Omega^2(\Sigma)\otimes\mathfrak{g}$ using the cylindrical metric on $[-1,1]\times\Sigma$. Note that this space could be identified with

$$L^2_k(\Omega^0([-1,1]\times\Sigma)\otimes\mathfrak{g})\oplus L^2_k(\Omega^1([-1,1]\times\Sigma)\otimes\mathfrak{g})$$

on $[-1,1]\times\Sigma$ by mapping $(a,\phi,\psi)\in L^2_k([-1,1];\Omega^1(\Sigma)\otimes\mathfrak{g}\oplus\Omega^0(\Sigma)\otimes\mathfrak{g}\oplus\Omega^2(\Sigma)\otimes\mathfrak{g})$ to $(\phi,a+*\psi ds)$. Because the holonomy perturbation $f_{\mathbf{q}}$ is supported on $[-1,1]\times\Sigma$, for each $t\in[0,1]$, the operator $K_{B_t,f_{\mathbf{q}}}$ is equal to the twisted odd signature operator

$$\begin{pmatrix} *d_{B_t} & d_{B_t} \\ d^*_{B_t} & 0 \end{pmatrix} \tag{2.5.9}$$

on $(\Omega^0(Y)\oplus\Omega^1(Y))\otimes\mathfrak{g}$ in the complement of $[-1,1]\times\Sigma$. The restrictions of this operator on $[-2,-1]\times\Sigma$ and $[1,2]\times\Sigma$ are all of the form

$$*(ds\wedge\frac{\partial}{\partial s}+S_{A_{\mp1,t}})$$

under the identification described above. Let $\Lambda_{-1}(t)\subset L^2(\Omega^1(\Sigma)\otimes\mathfrak{g}\oplus\Omega^0(\Sigma)\otimes\mathfrak{g}\oplus\Omega^2(\Sigma)\otimes\mathfrak{g})$ be the subspace consisting of elements that extend to $L^2_{1/2}$ elements in $\ker(K_{B_t,f_{\mathbf{q}}})$ on $H_1\cup[-2,-1]\times\Sigma$ and let $\Lambda_1(t)\subset L^2(\Omega^1(\Sigma)\otimes\mathfrak{g}\oplus\Omega^0(\Sigma)\otimes\mathfrak{g}\oplus\Omega^2(\Sigma)\otimes\mathfrak{g})$

be the similar subspace constructed from $[1,2]\times\Sigma\cup H_2$. Then it is well-known that $\Lambda_{\pm 1}(t)$ defines a pair of *Lagrangian subspaces* in $L^2(\Omega^1(\Sigma)\otimes\mathfrak{g}\oplus\Omega^0(\Sigma)\otimes\mathfrak{g}\oplus\Omega^2(\Sigma)\otimes\mathfrak{g})$ with respect to the symplectic form $\tilde{\omega}_\Sigma$, see [Nic95, Section 2] for example. The paths $t\mapsto\Lambda_\pm(t)$ are continuously differentiable in the sense that the paths of orthogonal projection operators $P_{\Lambda_\pm(t)}$ onto $\Lambda_\pm(t)$ are continuously differentiable.

Definition 2.5.6. *Let $L^2_1([-1,1]\times\Sigma;\Lambda_{-1}(t),\Lambda_1(t))$ be the subspace of $L^2_1([-1,1];\Omega^1(\Sigma)\otimes\mathfrak{g}\oplus\Omega^0(\Sigma)\otimes\mathfrak{g}\oplus\Omega^2(\Sigma)\otimes\mathfrak{g})$ consisting of elements whose restriction to $\{\pm1\}\times\Sigma$ lies in the space $\Lambda_{\pm1}(t)$. Let $A_{s,t}\in\mathcal{A}(\Sigma)$ be the family described above. Define*

$$D(t):L^2_1([-1,1]\times\Sigma;\Lambda_{-1}(t),\Lambda_1(t))\to L^2([-1,1];\Omega^1(\Sigma)\otimes\mathfrak{g}\oplus\Omega^0(\Sigma)\otimes\mathfrak{g}\oplus\Omega^2(\Sigma)\otimes\mathfrak{g})$$

to be the operator given by

$$D(t)\begin{pmatrix}a\\ \phi\\ \psi\end{pmatrix}=\begin{pmatrix}*\dot{a}+d_{A_{s,t}}\phi+d^*_{A_{s,t}}\psi-*dX_s(A_{s,t})(a)\\ -*\dot{\psi}+d^*_{A_{s,t}}a\\ *\dot{\phi}+d_{A_{s,t}}a\end{pmatrix},$$

where the dot represents taking derivative along the s–direction.

Remark 2.5.7. The operator $D(t)$ could be written as $\tilde{J}_\Sigma(\frac{\partial}{\partial s}+S_{A_{s,t}}-dX_s(A_{s,t}))$. Using the Lagrangian boundary conditions $\Lambda_{\pm1}(t)$, it is easy to see that $D(t)$ defines a self-adjoint operator such that the inclusion map from its domain to the background Hilbert space is a compact operator. Meanwhile, $D(t)$ is exactly the restriction of the operator $K_{B_t,\mathbf{q}}$ on the region $[-1,1]\times\Sigma\subset Y$.

As a consequence, we obtain a 1–parameter family of self-adjoint operators $\{D(t)\}_{0\le t\le1}$. Although this family has varying domains $L^2_1([-1,1]\times\Sigma;\Lambda_{-1}(t),\Lambda_1(t))$, its spectral flow is nonetheless well-defined in this situation, see [SW08, Appendix A]. Note that all the discussions above are compatible with the $\mathrm{Stab}(B_t)$–action, therefore the H–equivariant spectral flow $Sf_H(D(t))$ of the family $\{D(t)\}_{0\le t\le1}$ is also well-defined.

Now we are ready to state the main result of this subsection.

Theorem 2.5.8. *Let* $\{K_{B_t,f_q}\}_{0\leq t\leq 1}$ *and* $\{D(t)\}_{0\leq t\leq 1}$ *be as above. Then*

$$Sf_H(K_{B_t,f_q}) = Sf_H(D(t)).$$

The proof of Theorem 2.5.8 is based on adapting the arguments in [Nic95]. The idea is to modify the path of self-adjoint operators to a preferred form to simplify the calculation of the spectral flow.

We start by reducing the calculation of the equivariant spectral flow to the calculation of the ordinary spectral flow. Identifying $(\Omega^0(Y)\oplus\Omega^1(Y))\otimes\mathfrak{g}$ with the space of smooth sections of the vector bundle $\mathfrak{g}\oplus T^*Y\otimes\mathfrak{g}$, where $\mathfrak{g}$ stands for the product $\mathfrak{g}$–bundle over Y, for any connection $B\in\mathcal{C}(Y)$, its stabilizer $\mathrm{Stab}(B)\subset G$ acts on $(\Omega^0(Y)\oplus\Omega^1(Y))\otimes\mathfrak{g}$ by the fiberwise adjoint G–action on the $\mathfrak{g}$ factor. Choose $H\subset G$ as above. Suppose the adjoint action of H on $\mathfrak{g}$ decomposes $\mathfrak{g}$ into isotypic pieces

$$\mathfrak{g}\cong V_1^{\oplus a_1}\oplus\cdots\oplus V_m^{\oplus a_m},$$

we define the vector bundle

$$\tilde{E}_i = (\Omega^0(Y)\oplus\Omega^1(Y))\otimes V_i^{\oplus a_i}$$

and the vector bundle

$$E_i = (\Omega^0(Y)\oplus\Omega^1(Y))\otimes\mathbb{K}_i^{a_i},$$

where $\mathbb{K}_i=\mathrm{Hom}_H(V_i,V_i)$. Then the global sections of E_i could be identified with

$$\mathrm{Hom}_H(V_i,(\Omega^0(Y)\oplus\Omega^1(Y))\otimes\mathfrak{g}).$$

Let $L_k^2(E_i)$ be the L_k^2–completion of the space of smooth sections of E_i over Y. Then

this encodes the information of the V_i-isotypic piece of the H–Banach space $L^2_k(\mathfrak{g} \oplus T^*Y \otimes \mathfrak{g})$. By the H–equivariance of the operator $K_{B_t,f_\mathfrak{q}}$, it induces an operator

$$K^i_{B_t,f_\mathfrak{q}} = L^2_k(E_i) \to L^2_{k-1}(E_i).$$

Correspondingly, we have

$$\mathrm{Hom}_H(V_i, L^2_k(\Omega^0([-1,1]\times\Sigma)\otimes\mathfrak{g}) \oplus L^2_k(\Omega^1([-1,1]\times\Sigma)\otimes\mathfrak{g})) \cong L^2_k(E_i|_{[-1,1]\times\Sigma})$$

and the spaces over the Riemann surface Σ given by

$$\mathrm{Hom}_H(V_i, L^2_k(\Omega^1(\Sigma)\otimes\mathfrak{g} \oplus \Omega^0(\Sigma)\otimes\mathfrak{g} \oplus \Omega^2(\Sigma)\otimes\mathfrak{g})) \cong L^2_k(E_i|_{\{0\}\times\Sigma}).$$

The symplectic form $\tilde{\omega}_\Sigma$ induces a symplectic form on $L^2(E_i|_{\{0\}\times\Sigma})$. By restricting $\Lambda_{\pm1}(t)$ to the V_i–isotypic piece, we obtain Lagrangian subspaces

$$\Lambda^i_{\pm1}(t) \subset L^2(E_i|_{\{0\}\times\Sigma}).$$

As a result, the H–equivariant operator $D(t)$ induces operators

$$D^i(t) : L^2_1(E_i|_{[-1,1]\times\Sigma}; \Lambda^i_{-1}(t), \Lambda^i_1(t)) \to L^2(E_i|_{[-1,1]\times\Sigma})$$

for all $i = 1, \ldots, m$. According to Definition 2.2.10, we have

Lemma 2.5.9. *$Sf_H(K_{B_t,f_\mathfrak{q}}) = Sf_H(D(t))$ if and only if $Sf(K^i_{B_t,f_\mathfrak{q}}) = Sf(D^i(t))$ for any $1 \le i \le m$.* □

From now on, we will focus on the operators $K^i_{B_t,f_\mathfrak{q}}$ and $D^i(t)$ for a fixed $i \in \mathbb{Z} \cap [1, m]$.

Definition 2.5.10. *An operator $D : L^2_k(E_i) \to L^2_{k-1}(E_i)$ is said to satisfy* unique

continuation *if for any* $v \in \ker(D)$ *such that* $v|_U = 0$ *where* $U \subset H_2 \subset Y$ *is a non-empty open set, we have* $v \equiv 0$.

Lemma 2.5.11. *For any* $t \in [0,1]$*, the operator* $K^i_{B_t,f_q}$ *satisfies unique continuation.*

Proof. Let $v \in \ker(K^i_{B_t,f_\mathbf{q}})$ and $v|_U = 0$ where U is as above. The operator $K^i_{B_t,f_\mathbf{q}}|_{[1,2]\times\Sigma\cup H_2}$ comes from the twisted signature operator therefore is a Dirac operator. By the unique continuation property of Dirac operators, we know that v is identically 0 over $[1,2]\times\Sigma \cup H_2$. Over the region $[-1,1]\times\Sigma$, the element v defines a curve $v(s)$ in $L^2(E_i|_{\{0\}\times\Sigma})$ satisfying the ODE

$$\tilde{J}_\Sigma(\frac{\partial}{\partial s} + S_{A_{s,t}} - dX_s(A_{s,t}))v(s) = 0.$$

Here we use the same notation to represent the induced operators on $L^2(E_i|_{[-1,1]\times\Sigma})$. Since the solution is uniquely determined by the initial value and $v|_{\{1\}\times\Sigma} = 0$, we know that $v|_{[-1,1]\times\Sigma} = 0$. The vanishing of v over $H_1 \cup [-2,-1]\times\Sigma$ comes again from the unique continuation property of the twisted signature operator on $H_1 \cup [-2,-1]\times\Sigma$ and the fact that we have already shown that $v_{\{-1\}\times\Sigma} = 0$. Therefore the lemma is proved. □

Lemma 2.5.12. *The operator* $K^i_{B_1,f_q} - K^i_{B_0,f_q}$ *is a bounded operator on* $L^2_k(E_i)$.

Proof. By definition, $K^i_{B_1,f_\mathbf{q}} - K^i_{B_0,f_\mathbf{q}}$ is the sum of a 0–th order differential operator depending on B_0, B_1 and the induced operator on $L^2_k(E_i)$ given by the difference

$$DV_{f_\mathbf{q}}(B_1) - DV_{f_\mathbf{q}}(B_0).$$

Use the definition and [KM11, Proposition 3.5], these 2 terms are both bounded by $C(\|B_0\|_{L^2_k} + \|B_1\|_{L^2_k})$ where $C > 0$ is some constant depending on Y, so the lemma is proved. □

Recall that for a family of self-adjoint operators $\{D_t\}_{0\leq t\leq 1}$, its *resonance set* $Z(D_t)$ consists of $t \in [0,1]$ such that $\ker(D_t) \neq \{0\}$. Suppose $t_0 \in Z(D_t)$ and let P_{t_0} be the orthogonal projection onto $\ker(D_{t_0})$. Then the *resonance matrix* $R(D_{t_0})$ at t_0 is defined to be the map $P_{t_0}\dot{D}_{t_0} : \ker(D_{t_0}) \to \ker(D_{t_0})$.

Definition 2.5.13. *A family of self-adjoint operators $\{D_t\}_{0\leq t\leq 1}$ is said to be* positive *if the resonance matrices are all positive definite. It is said to be* negative *if the family $\{-D_t\}_{0\leq t\leq 1}$ is positive.*

Using Kato's selection theorem, if the family $\{D_t\}_{0\leq t\leq 1}$ is positive or negative, the set $Z(D_t)$ is discrete.

Lemma 2.5.14. *The 1–parameter family $\{K^i_{B_t,f_\mathbf{q}}\}_{0\leq t\leq 1}$ defined over $L^2_1(E_i) \subset L^2(E_i)$ is homotopic relative end points to the concatenation of a positive 1–parameter family and a negative 1–parameter family such that each operator in these 2–families satisfies unique continuation.*

Proof. Using Lemma 2.5.12, we can choose some $C > 0$ such that the norm of $K^i_{B_1,f_\mathbf{q}} - K^i_{B_0,f_\mathbf{q}}$ over $L^2(E_i)$ is bounded above by $C - 1$. Consider the families

$$K^i_+(t) = K^i_{B_0,f_\mathbf{q}} + tC \cdot \mathrm{id}$$

$$K^i_-(t) = K^i_{B_0,f_\mathbf{q}} + C \cdot \mathrm{id} + t(K^i_{B_1,f_\mathbf{q}} - K^i_{B_0,f_\mathbf{q}} - C \cdot \mathrm{id})$$

for $0 \leq t \leq 1$. It is easy to see that $K^i_+(t)$ is positive and $K^i_-(t)$ is negative. The affine homotopy between $\{K^i_{B_t,f_\mathbf{q}}\}_{0\leq t\leq 1}$ and the concatenation of these 2 families gives the desired homotopy. For the unique continuation property, one proceeds exactly the same as in the proof of Lemma 2.5.11 by noticing that $K^i_\pm(t)$ could be described by Dirac operators on $H_1 \cup [-2,-1] \times \Sigma \cup [1,2] \times \Sigma \cup H_2$ and the kernel of $K^i_\pm(t)$ consists of solutions to an ODE on $[-1,1] \times \Sigma$. □

Given a positive or negative family $\{D_t\}_{0\leq t\leq 1}$, we wish to perturb it further so that for any $t_0 \in Z(D_t)$, the kernel of D_{t_0} becomes 1–dimensional. This is the transversality result established in [Nic95, Proposition 3.6]. The unique continuation condition stated above guarantees that the proof of [Nic95, Proposition 3.6] works without change to give the following statement.

Lemma 2.5.15. *For the two families $\{K^i_\pm(t)\}_{0\leq t\leq 1}$ constructed in Lemma 2.5.14, one can find*

$$\alpha_\pm : [0,1] \to C^\infty(\mathrm{Hom}(E_i, E_i))$$

which are arbitrarily small such that the following holds:

1. *$\alpha_\pm(t)$ is supported away from $[-2,2]\times\Sigma$;*
2. *The family $K^i_+(t)+\alpha_+(t)$ (resp. $K^i_-(t)+\alpha_-(t)$) remains to be positive (resp. negative) and for any $t_0 \in Z(K^i_+(t)+\alpha_+(t))$ (resp. $t_0 \in Z(K^i_-(t)+\alpha_-(t))$), the kernel of $K^i_+(t_0)+\alpha_+(t_0)$ (resp. $K^i_-(t_0)+\alpha_-(t_0)$) is actually 1–dimensional.* □

Proof of Theorem 2.5.8. Let us fix $\alpha_\pm(t)$ as in Lemma 2.5.15 and denote $\tilde{K}^i_\pm(t) = K^i_\pm(t)+\alpha_\pm(t)$. By construction, $\{\tilde{K}^i_\pm(t)\}_{0\leq t\leq 1}$ defines operators

$$\tilde{K}^i_\pm(t) : L^2_1(E_i) \to L^2(E_i)$$

for all $0 \leq t \leq 1$ which are Dirac operators when restricted to $H_1 \cup [-2,-1]\times\Sigma \cup [1,2]\times\Sigma \cup H_2$ and they are cylindrical over $[-2,-1]\times\Sigma \cup [1,2]\times\Sigma$. Therefore, they define Lagrangian subspaces

$$\tilde{\Lambda}^i_{-1}(t)_\pm \subset L^2(E_i|_{\{0\}\times\Sigma})$$

consisting of restriction of elements in $L^2_{1/2}(E_i|_{H_1\cup[-2,-1]\times\Sigma})$ lying in

$$\ker(\tilde{K}^i_\pm(t)|_{H_1\cup[-2,-1]\times\Sigma}).$$

Similarly one can construct

$$\tilde{\Lambda}_1^i(t)_\pm \subset L^2(E_i|_{\{0\}\times\Sigma})$$

by considering the restriction to $[1,2]\times\Sigma \cup H_2$. We can then construct the operators

$$\tilde{D}_\pm^i(t) : L_1^2(E_i|_{[-1,1]\times\Sigma}; \tilde{\Lambda}_{-1}^i(t)_\pm, \tilde{\Lambda}_1^i(t)_\pm) \to L^2(E_i|_{[-1,1]\times\Sigma})$$

similar to the construction of $D(t)$, which define self-adjoint operators using the Lagrangian boundary conditions.

Making $\alpha_\pm(t)$ smaller if necessary, the spectral flow of the family $\{K_{B_t,f_\mathbf{q}}\}_{0\le t\le 1}$ is the same as the sum of the spectral flows of $\{\tilde{K}_+^i(t)\}_{0\le t\le 1}$ and $\{\tilde{K}_-^i(t)\}_{0\le t\le 1}$. Similarly, by concatenating the Lagrangian boundary conditions, it is easy to see that the spectral flow of $\{D^i(t)\}_{0\le t\le 1}$ is equal to the sum of the spectral flows of $\{\tilde{D}_+^i(t)\}_{0\le t\le 1}$ and $\{\tilde{D}_-^i(t)\}_{0\le t\le 1}$. Using Lemma 2.5.9, it suffices to show that

$$Sf(\tilde{K}_+^i(t)) = Sf(\tilde{D}_+^i(t)) \text{ and } Sf(\tilde{K}_-^i(t)) = Sf(\tilde{D}_-^i(t)).$$

We prove the plus version of the above statement and the negative version holds using the same argument. From the definition, we see that

$$Z(\tilde{K}_+^i(t)) = Z(\tilde{D}_+^i(t))$$

therefore it suffices to show that the resonance matrices of $\{\tilde{K}_+^i(t)\}_{0\le t\le 1}$ and $\{\tilde{D}_+^i(t)\}_{0\le t\le 1}$ have the same sign at any $t_0 \in Z(\tilde{K}_+^i(t)) = Z(\tilde{D}_+^i(t))$. In other words, we just need to show that the resonance matrix of $\tilde{D}_+^i(t)$ at $t = t_0$ is positive definite. Let $v \in \ker(\tilde{D}_+^i(t_0)) - \{0\}$. Note that v could also be viewed as an element of $\ker(\tilde{K}_+^i(t_0))$. Let $\chi : [0,1] \to \mathbb{R}$ be a non-negative smooth function supported in $(0,1)$ which is

strictly positive over $[\frac{1}{4}, \frac{3}{4}]$. Then χ defines a function on $[0,1] \times \Sigma \subset Y$ by composing it with the projection to the s–coordinate. This further extends to a smooth function over Y in the obvious way. Then for $\epsilon > 0$ sufficiently small, the family

$$\{\tilde{K}^i_+(t_0) + (t - t_0)\chi \cdot \mathrm{id}\}_{t_0-\epsilon \leq t \leq t_0+\epsilon} \tag{2.5.10}$$

is C^0–close to $\{\tilde{K}^i_+(t)\}_{t_0-\epsilon \leq t \leq t_0+\epsilon}$ and the affine homotopy between them preserves the spectral flow. This is because $\langle v, \chi v \rangle_{L^2} > 0$ and the kernel is 1–dimensional: if $\langle v, \chi v \rangle_{L^2} = 0$, the restriction of v to the time slice $\{\frac{1}{2}\} \times \Sigma$ is 0 thus it vanishes on $[-1,1] \times \Sigma$ by uniqueness of solutions to ODE therefore vanishes on whole Y by unique continuation of Dirac operators. Accordingly, we can construct the induced affine homotopy from $\{\tilde{D}^i_+(t)\}_{t_0-\epsilon \leq t \leq t_0+\epsilon}$ to

$$\{\tilde{D}^i_+(t_0) + (t - t_0)\chi \cdot \mathrm{id}\}_{t_0-\epsilon \leq t \leq t_0+\epsilon} \tag{2.5.11}$$

preserving the spectral flow. Now notice that because χ is supported in the region $[0,1] \times \Sigma$ and the kernels of the new families at time t_0 are all spanned by v, we see that the resonance matrices of these 2 families (2.5.10) and (2.5.11) have the same value at time t_0. This finishes the proof. □

2.5.3 An adiabatic limit result

Let $\{A_{s,t}\}_{-1 \leq s \leq 1, 0 \leq t \leq 1} \subset A_F(\Sigma)$ be a smooth family of flat connections on Σ. Viewing Σ as the boundary of $H_1 \cup [-2,-1] \times \Sigma$ and $[1,2] \times \Sigma \cup H_2$, we further suppose that $A_{-1,t}$ extends to a flat connection $B_{-1,t}$ on $H_1 \cup [-2,-1] \times \Sigma$ and $A_{1,t}$ extends to a flat connection $B_{1,t}$ on $[1,2] \times \Sigma \cup H_2$ for all $t \in [0,1]$. Then we can define $\{\Lambda_{\pm 1}(t)\}_{0 \leq t \leq 1}$ be the same as in the previous subsection. Note that we do not require that $A_{s,0}$ or $A_{s,1}$ comes from the restriction of a $f_{\mathbf{q}}$–perturbed flat connection on Y as in the previous subsection. This setup gives us more flexibility for later calculations. Furthermore, we

assume that $\mathrm{Stab}(A_{s,t}) \cong H \subset G$ remains to be the same for all $(s,t) \in [-1,1]\times[0,1]$.

Given $A_{s,t} \in \mathcal{A}_F(\Sigma)$, let $\mathcal{H}_{A_{s,t}}$ be the kernel of the twisted de Rham operator

$$S_{A_{s,t}} : L^2(\Omega^1(\Sigma)\otimes\mathfrak{g}\oplus\Omega^0(\Sigma)\otimes\mathfrak{g}\oplus\Omega^2(\Sigma)\otimes\mathfrak{g}) \to L^2(\Omega^1(\Sigma)\otimes\mathfrak{g}\oplus\Omega^0(\Sigma)\otimes\mathfrak{g}\oplus\Omega^2(\Sigma)\otimes\mathfrak{g}).$$

In other words, $\mathcal{H}_{A_{s,t}}$ is given by the total space of twisted harmonic forms on Σ with respect to the flat connection $A_{s,t}$. This is a symplectic vector space using the restriction of the symplectic form $\tilde{\omega}_\Sigma$. Let

$$\pi_{A_{s,t}} : L^2(\Omega^1(\Sigma)\otimes\mathfrak{g}\oplus\Omega^0(\Sigma)\otimes\mathfrak{g}\oplus\Omega^2(\Sigma)\otimes\mathfrak{g}) \to \mathcal{H}_{A_{s,t}}$$

be the L^2–orthogonal projection. Define

$$\mathcal{L}_{\pm1}(t) := \mathcal{H}_{A_{\pm1,t}} \cap \Lambda_{\pm1}(t), \text{ for } 0 \le t \le 1.$$

Lemma 2.5.16. *$\mathcal{L}_i(t)$ is a Lagrangian subspace of $\mathcal{H}_{A_{i,t}}$ for $i = \pm1$.*

Proof. By definition, $\mathcal{L}_i(t)$ is isotropic. So we just need to prove that $\dim \mathcal{L}_i(t) \ge \frac{1}{2}\dim \mathcal{H}_{A_{i,t}}$. Notice that $(0, \ker d_{A_{i,t}}, 0) \subset \mathcal{L}_i$ because this comes from the Lie algebra of the stabilizer group and $\mathrm{Stab}(B_{i,t}) = \mathrm{Stab}(A_{i,t})$.

Let $m = \dim\ker d_{B_{-1,t}}$. Recall that h is the genus of the Riemann surface Σ. Fix a base point p_0 on $\Sigma \cong \partial H_1$, let $\gamma_1, \cdots, \gamma_h$ be h closed curved based on p_0 that generates the fundamental group of H_1. Let $\hat{\mathcal{C}}$ be the space of smooth connections B on $H_1 \cup [-2,-1]\times\Sigma$ such that $B(\partial_s) = 0$ on the boundary. Define

$$\mathrm{Hol} : \hat{\mathcal{C}} \to G^h$$

to be the map given by holonomies along γ_i. Let $U = (-\epsilon,\epsilon)^h$, let $\Phi : U \to \hat{\mathcal{C}}$ be a map such that $\Phi(0)$ equals $B_{-1,t}$, and that $\mathrm{Im}\, d(\mathrm{Hol}\circ\Phi)$ is surjective at zero. By

construction, $\operatorname{Im} d\Phi(0) \subset \Omega^1(H_1 \cup [-2,-1] \times \Sigma) \otimes \mathfrak{g}$ is a $h \cdot \dim G$ dimensional linear space. For every $u \in \operatorname{Im} d\Phi(0)$, there is a unique v such that $v \in \operatorname{Im} d_{B_{-1,t}}$, $*v = 0$ on $\{-1\} \times \Sigma$, and $u' := u - d_{B_{-1,t}} v \in \ker(d^*_{B_{-1,t}})$. The kernel of $\operatorname{Im} d\Phi(0)$ under the map $u \mapsto u'$ has dimension at most m by considering the holonomies, therefore the image of $\operatorname{Im} d\Phi(0)$ under the map $u \mapsto u'$ has dimension at least $(h-1) \cdot \dim G + m$. Therefore

$$\mathcal{L}_{-1}(t) \cap (\ker(d_{B_{-1,t}} + d^*_{B_{-1,t}}), 0, 0) \subset \mathcal{H}_{A_{-1,t}}$$

has dimension at least $(h-1) \cdot \dim G + m$, which proved that $\mathcal{L}_{-1}(t)$ is Lagrangian. The same argument works to prove $\mathcal{L}_i(t)$ is Lagrangian. □

By definition, $\dim \mathcal{H}_{A_{s,t}}$ remains to be the same for all $(s,t) \in [-1,1] \times [0,1]$. For a given $t \in [0,1]$, consider the spaces

$$L^2_k([-1,1]; \mathcal{H}_{A_{s,t}}) \subset L^2_k([-1,1]; \Omega^1(\Sigma) \otimes \mathfrak{g} \oplus \Omega^0(\Sigma) \otimes \mathfrak{g} \oplus \Omega^2(\Sigma) \otimes \mathfrak{g})$$

$$L^2_k([-1,1]; \mathcal{H}_{A_{s,t}}; \mathcal{L}_{-1}(t), \mathcal{L}_1(t)) \subset L^2_k([-1,1] \times \Sigma; \Lambda_{-1}(t), \Lambda_1(t))$$

consisting of elements (a, ϕ, ψ) whose value at each $s \in [-1,1]$ lies in $\mathcal{H}_{A_{s,t}}$. Then we can define a self-adjoint operator

$$D_0(t) : L^2_1([-1,1]; \mathcal{H}_{A_{s,t}}; \mathcal{L}_{-1}(t), \mathcal{L}_1(t)) \to L^2([-1,1]; \mathcal{H}_{A_{s,t}}) \tag{2.5.12}$$

given by

$$D_0(t) \begin{pmatrix} a \\ \phi \\ \psi \end{pmatrix} = \pi_{A_{s,t}} \begin{pmatrix} *\dot{a} - *dX_s(A_{s,t})(a) \\ - * \dot{\psi} \\ *\dot{\phi} \end{pmatrix}.$$

By construction, the family $\{D_0(t)\}_{0 \le t \le 1}$ has a well-defined H–equivariant spectral flow $Sf_H(D_0(t))$. The following is another comparison result of spectral flows.

Theorem 2.5.17. *Let $\{D(t)\}_{0\leq t\leq 1}$ be the family constructed in Section 2.5.2 and let $\{D_0(t)\}_{0\leq t\leq 1}$ be the family constructed above. Then*

$$Sf_H(D(t)) = Sf_H(D_0(t)).$$

The proof of Theorem 2.5.17 is exactly the same as the proof of [DS94, Theorem 6.1]. We will sketch a proof for completeness. Because we have used connections in temporal gauge, the covariant differentiation ∇_s in [DS94] is replaced by the ordinary differentiation ∂_s. Note that the calculation of equivariant spectral flow could be reduced to the calculation of ordinary spectral flow as in Section 2.5.2 so the H–equivariance will be suppressed in this subsection.

The idea is to use the 1–parameter family of operators

$$D_\epsilon(t) : L^2_1([-1,1]\times\Sigma;\Lambda_{-1}(t),\Lambda_1(t)) \to L^2([-1,1];\Omega^1(\Sigma)\otimes\mathfrak{g}\oplus\Omega^0(\Sigma)\otimes\mathfrak{g}\oplus\Omega^2(\Sigma)\otimes\mathfrak{g})$$

given by the formula

$$D_\epsilon(t)\begin{pmatrix} a \\ \phi \\ \psi \end{pmatrix} = \begin{pmatrix} *\dot{a} + d_{A_{s,t}}\phi + d^*_{A_{s,t}}\psi - *dX_s(A_{s,t})(a) \\ - * \dot{\psi} + \frac{1}{\epsilon^2} d^*_{A_{s,t}} a \\ *\dot{\phi} + \frac{1}{\epsilon^2} d_{A_{s,t}} a \end{pmatrix}$$

depending on the parameter $\epsilon > 0$, so that the spectral flow of $\{D(t)\}_{0\leq t\leq 1} = \{D_1(t)\}_{0\leq t\leq 1}$ could be reduced to the spectral flow of $\{D_0(t)\}_{0\leq t\leq 1}$ by passing to the adiabatic limit $\epsilon \to 0$.

Fix $t \in [0,1]$. Introduce norms

$$\|(a,\phi,\psi)\|^2_{0,\epsilon} = \|a\|^2 + \epsilon^2\|\phi\|^2 + \epsilon^2\|\psi\|^2$$

on the space $L^2([-1,1];\Omega^1(\Sigma)\otimes\mathfrak{g}\oplus\Omega^0(\Sigma)\otimes\mathfrak{g}\oplus\Omega^2(\Sigma)\otimes\mathfrak{g})$ and

$$\begin{aligned}\|(a,\phi,\psi)\|^2_{1,\epsilon} = \|a\|^2 + \|d_{A_{s,t}}a\|^2 + \|d_{A_{s,t}}*a\|^2 + \epsilon^2\|\dot{a}\|^2 \\ + \epsilon^2\|d_{A_{s,t}}\phi\|^2 + \epsilon^4\|\dot{\phi}\|^2 + \epsilon^2\|d^*_{A_{s,t}}\psi\|^2 + \epsilon^4\|\dot{\psi}\|^2\end{aligned}$$

on $L^2_1([-1,1];\Omega^1(\Sigma)\otimes\mathfrak{g}\oplus\Omega^0(\Sigma)\otimes\mathfrak{g}\oplus\Omega^2(\Sigma)\otimes\mathfrak{g})$ using the product metric on $[-1,1]\times\Sigma$.

Sketch of proof of Theorem 2.5.17. The proof breaks into five steps.

*Step 1-Elliptic estimate:*There is an $\epsilon_0 > 0$ such that for any $0 < \epsilon < \epsilon_0$, there exists a constant $C > 0$ independent of ϵ such that

$$\|u - \pi_{A_{s,t}}(u)\|_{1,\epsilon} \leq C\epsilon(\|D_\epsilon(t)u\|_{0,\epsilon} + \|\pi_{A_{s,t}}(u)\|_{L^2}) \tag{2.5.13}$$

for $u \in L^2_1([-1,1]\times\Sigma;\Lambda_{-1}(t),\Lambda_1(t))$.

This is [DS94, Lemma 7.3]. Note that although our 3–manifold $[-1,1]\times\Sigma$ has boundary, the Lagrangian boundary condition guarantees that the integration by parts formula entering the proof works without change. Meanwhile, our projection map $\pi_{A_{s,t}}$ involves the H^0 and H^2 part so the reducible connections would not affect the estimate.

Step 2-Convergence of resolvent set: For every $C_0 > 0$ there exists constants $\epsilon_0 > 0$ and $C > 0$ such that the following holds for all $t \in [0,1]$ and $|\lambda| \leq C_0$. If

$$\|u_0\|_{L^2} \leq C_0\|D_0(t)u_0 - \lambda u_0\|_{L^2}$$

for all $u_0 \in L^2_1([-1,1];\mathcal{H}_{A_{s,t}};\mathcal{L}_{-1}(t),\mathcal{L}_1(t))$, then

$$\|u\|_{0,\epsilon} \leq C\|D_\epsilon(t)u - \lambda u\|_{0,\epsilon}$$

for $0 < \epsilon < \epsilon_0$ and $u \in L^2_1([-1,1]\times\Sigma;\Lambda_{-1}(t),\Lambda_1(t))$.

This is [DS94, Lemma 7.4]. Note that the proof is based on a direct computation and an application of equation (2.5.13) therefore it works in our situation without change.

Step 3-Refined convergence of operators: Define

$$R_0 = \{(t, \lambda) \in [0, 1] \times \mathbb{C} \big| \lambda \text{ is not an eigenvalue of } D_0(t)\}$$

to be the resolvent set of the family $\{D_0(t)\}_{0 \leq t \leq 1}$. Denote by R_ϵ is resolvent set of the family $\{D_\epsilon(t)\}_{0 \leq t \leq 1}$. Then for every compact subset $K \subset R_0$, there exists a constant $\epsilon_0 > 0$ such that $K \subset R_\epsilon$ for $0 < \epsilon < \epsilon_0$ and

$$\|\pi_{A_{s,t}}((\lambda \cdot \mathrm{id} - D_\epsilon(t))^{-1} v) - (\lambda \cdot \mathrm{id} - D_0(t))^{-1} \pi_{A_{s,t}}\|_{L^2} \leq C\epsilon \|v\|_{0,\epsilon} \tag{2.5.14}$$

for $(t, \lambda) \in K$ and $v \in L^2([-1, 1]; \Omega^1(\Sigma) \otimes \mathfrak{g} \oplus \Omega^0(\Sigma) \otimes \mathfrak{g} \oplus \Omega^2(\Sigma) \otimes \mathfrak{g})$.

This is [DS94, Lemma 7.5].

Step 4-Multiplicity estimate: For all $C_0 > 0$ there exists a constant $k_0 > 0$ such that the following holds for all $t \in [0, 1]$ and $0 < \epsilon^2 + \delta^2 < k_0$. If λ_0 is an eigenvalue of $D_0(t)$ of multiplicity m_0 and

$$u_0 \perp \ker(\lambda_0 \cdot \mathrm{id} - D_0(t)) \Longrightarrow \|u_0\|_{L^2} \leq C_0 \|\lambda_0 u_0 - D_0(t) u_0\|_{L^2}$$

for all $u \in L^2_1([-1, 1]; \mathcal{H}_{A_{s,t}}; \mathcal{L}_{-1}(t), \mathcal{L}_1(t))$, then the multiplicity of all eigenvalues of λ of $D_\epsilon(t)$ with $|\lambda - \lambda_0| \leq \delta$ does not exceed m_0.

This follows from [DS94, Lemma 7.6].

Step 5-Concluding the proof. By Kato's selection theorem, we can find $\delta \in \mathbb{R}$ which is arbitrarily close to 0 such that the resonance set of the family of the operators $\{D_0(t) - \delta \cdot \mathrm{id}\}_{0 \leq t \leq 1}$, written as $Z(D_0(t) - \delta \cdot \mathrm{id})$ is finite and contained in $(0, 1)$. Moreover, we can guarantee that the resonance matrices of this family is non-degenerate.

Use $\sigma(D)$ to represent the spectrum of D. Let $Z(D_0(t)-\delta\cdot\mathrm{id}) = \{t_1,\ldots,t_N\}$. Recall that $R(D_0(t_i)-\delta\cdot\mathrm{id})$ for $i=1,\ldots,N$ is the resonance matrix. Define

$$m_i := \mathrm{sign}(R(D_0(t_i)-\delta\cdot\mathrm{id})), i=1,\ldots,N$$

to be the signatures of the resonance matrices. Then

$$Sf(\{D_0(t)\}_{0\le t\le 1}) = \sum_{i=1}^{N} m_i$$

and m_i could be computed as follows. Choose $\kappa > 0$ such that $\lambda = \delta$ is the only eigenvalue of $D_0(t_i)$ in the interval $[\delta-\kappa, \delta+\kappa]$. Now choose $\tau > 0$ such that $\delta \pm \kappa \notin \sigma(D_0(t))$ for all $t \in [t_i-\tau, t_i+\tau]$. Then

$$\begin{aligned} m_i = &\#\{\lambda\in\sigma(D_0(t_i-\tau))\big|\delta-\kappa<\lambda<\delta\} \\ &-\#\{\lambda\in\sigma(D_0(t_i+\tau))\big|\delta-\kappa<\lambda<\delta\}. \end{aligned}$$

By *Step 2*, there exists a constant $\epsilon_0 > 0$ such that $\delta \notin \sigma(D_\epsilon(t))$ for $t = t_i-\tau$ or $t_i+\tau$) and $\delta\pm\kappa \notin \sigma(D_\epsilon(t))$ for all $t \in (t_i-\tau, t_i+\tau)$ as long as $0<\epsilon<\epsilon_0$. Using *Step 3*, *Step 4* and the spectral projection operators, one can show that

$$\begin{aligned} m_i = &\#\{\lambda\in\sigma(D_\epsilon(t_i-\tau))\big|\delta-\kappa<\lambda<\delta\} \\ &-\#\{\lambda\in\sigma(D_\epsilon(t_i+\tau))\big|\delta-\kappa<\lambda<\delta\}. \end{aligned}$$

Therefore we conclude that

$$Sf(D_\epsilon(t)) = \sum_{i=1}^{N} m_i = Sf(D_0(t))$$

for $\epsilon > 0$ sufficiently small. By homotopy invariance, $Sf(D_\epsilon(t))$ is independent of

$\epsilon > 0$. This concludes the proof. For the full details of the proof, we refer the readers to [DS94, Section 7]. □

2.5.4 From Maslov index to spectral flow and back

Throughout this subsection, the structure group G is equal to $\mathrm{SU}(n)$ and Y is an integer homology 3–sphere. Furthermore, suppose that the genus of the Heegaard surface Σ is $h \geq 3$. Let $H_{\hat{s}}$ and $f_{\mathbf{q}}$ be a compatible pair of small non-degenerate perturbations. Suppose $B \in \mathcal{C}(Y)$ is a $f_{\mathbf{q}}$–perturbed flat connection of type $\sigma = ((n_1, m_1), \ldots, (n_r, m_r)) \in \Sigma_n$ (see the end of Section 2.2). Without loss of generality, we can assume that B is in temporal gauge over $[-2, 2] \times \Sigma$ and it is represented by a path $A_s : [-2, 2] \to \mathcal{A}^0_F(\Sigma)$ which is constant for $s \in [-2, -1] \cup [1, 2]$. After applying a gauge transform, the connection B is given by the direct sum of irreducible $\mathrm{SU}(n_i)$-connections $B^{(i)}$ on $E(n_i)$ where

$$P \times_{\mathrm{SU}(n)} \mathbb{C}^n \cong E(n_1)^{\oplus m_1} \oplus \cdots \oplus E(n_r)^{\oplus m_r}. \tag{2.5.15}$$

Let $A^{(i)}_s$ be the induced path on the $E(n_i)$–component. Then $A^{(i)}_{-1}$ extends to a flat connection $B^{(i)}_{-1}$ on $H_1 \cup [-2, -1] \times \Sigma$ and $A^{(i)}_1$ extends to a flat connection $B^{(i)}_1$ on $[1, 2] \times \Sigma \cup H_2$. Note that because Y is an integer homology sphere, the product connection θ is a non-degenerate critical point of the unperturbed functional CS. Because $f_{\mathbf{q}}$ is a small perturbation, there exists a $f_{\mathbf{q}}$–perturbed flat connection $\tilde{\theta}$ in temporal gauge over $[-2, 2] \times \Sigma$ with stabilizer G which lies in a contractible neighborhood of the product connection $\theta \in \mathcal{C}(Y)$. Denote by $\tilde{\theta}^{(i)}$ the $E(n_i)$–component of $\tilde{\theta}$ under the decomposition (2.5.15) and $\tilde{\theta}_s$ is the induced path of flat connections in $\mathcal{A}_F(\Sigma)$, which can be further assumed to lie in $\mathcal{A}^0_F(\Sigma)$. Then we can find a path of connections

$$A^{(i)}_{-1,t} : [0, 1] \to \mathcal{A}^0_F(\Sigma) \text{ with } A^{(i)}_{-1,0} = \tilde{\theta}^{(i)}_{-1} \text{ and } A^{(i)}_{-1,1} = A^{(i)}_{-1}$$

such that $A^{(i)}_{-1,t}$ extends to a flat connection $B^{(i)}_{-1,t}$ on $E(n_i)$ over $H_1 \cup [-2,-1] \times \Sigma$ for all $t \in [0,1]$ which is irreducible expect for $t = 0$. Similarly, we can find the family $A^{(i)}_{1,t}$ on $[1,2] \times \Sigma \cup H_2$ and $B^{(i)}_{1,t}$ with the same properties. Then the union of paths $\{A^{(i)}_{-1,t}\}_{0\le t\le 1}, \{\tilde{\theta}^{(i)}_s\}_{-1\le s\le 1}, \{A^{(i)}_{1,t}\}_{0\le t\le 1}$ and $\{A^{(i)}_s\}_{-1\le s\le 1}$ defines a loop inside the space of flat connections on $E(n_i)$, denoted by $\mathcal{A}^{(i)}_F(\Sigma)$. By our assumption that the genus of Σ is at least 3, [DU95, Corollary 2.7] shows that the subset of $\mathcal{A}^{(i)}_F(\Sigma)$ consisting of irreducible connections is simply-connected. Consequently, the loop constructed above could be extended to a smooth 2–parameter family

$$A^{(i)}_{s,t} : [-1,1] \times [0,1] \to \mathcal{A}^{(i)}_F(\Sigma)$$

such that $A^{(i)}_{s,t}$ is irreducible except for $(s,t) \in [-1,1] \times \{0\}$. Take the direct sum of these families for $1 \le i \le m$ and apply further gauge transformations if necessary, we obtain a 2–parameter family

$$A_{s,t} : [-1,1] \times [0,1] \to \mathcal{A}^0_F(\Sigma) \subset \mathcal{A}_F(\Sigma).$$

Let $\tilde{\theta} = B_0$ and $B = B_1$, we can use $A_{s,t}$ to construct a family $\{B_t\}_{0\le t\le 1}$ satisfying the assumptions in the beginning of Subsection 2.5.2. According to this construction, $\mathrm{Stab}(A_{s,t})$ remains invariant for $(s,t) \in [-1,1] \times (0,1]$.

Lemma 2.5.18. *Suppose the perturbations $H_{\hat{s}}$ and $f_{\boldsymbol{q}}$ are 0. Then for each $1 \le i \le r$, the symplectic area of $A^{(i)}_{s,t} : [-1,1] \times [0,1] \to \mathcal{A}^{(i)}_F(\Sigma)$ under the Atiyah-Bott symplectic form ω_Σ is one-half of the Chern-Simons invariant of $B^{(i)}$.*

Proof. According to our convention of Chern-Simons functional (2.2.1),

$$\begin{aligned}\mathrm{CS}(B^{(i)}) &= \mathrm{CS}(B^{(i)}|_{H_1\cup[-2,-1]\times\Sigma}) + \mathrm{CS}(B^{(i)}|_{[-1,1]\times\Sigma}) + \mathrm{CS}(B^{(i)}|_{[1,2]\times\Sigma\cup H_2}) \\ &= \int_{[0,1]\times[-1,1]\times\Sigma} \mathrm{Tr}(F_{B_t}\wedge F_{B_t}) \\ &= \int_{[0,1]\times[-1,1]\times\Sigma} \mathrm{Tr}((ds\wedge\partial_s A^{(i)}_{s,t} + dt\wedge\partial_t A^{(i)}_{s,t} + F_{A^{(i)}_{s,t}})^2) \\ &= 2\int_{[0,1]\times[-1,1]\times\Sigma} \mathrm{Tr}(\partial_s A^{(i)}_{s,t}\wedge\partial_t A^{(i)}_{s,t})dsdt.\end{aligned}$$

The second equality is an application of the Stokes' formula over the 4–manifolds $[0,1]\times[-1,1]\times\Sigma, [0,1]\times(H_1\cup[-2,-1]\times\Sigma), [0,1]\times([1,2]\times\Sigma\cup H_2)$ and notice that $B^{(i)}|_{H_1\cup[-2,-1]\times\Sigma}, B^{(i)}|_{[1,2]\times\Sigma\cup H_2}$ are connected to the product connection through flat connections. The above calculation proves the lemma. □

Recall that L_1 and L_2 are Lagrangians in the extended moduli space $\hat{\mathcal{M}}^{\mathfrak{g}}(\Sigma')$ constructed from handlebodies. The $f_{\mathbf{q}}$–perturbed flat connection B is associated with the G–orbit $[A_1]\in\Phi_{H_{\hat{s}}}(L_1)\cap L_2$ in $\hat{\mathcal{M}}^{\mathfrak{g}}(\Sigma')$ in Proposition 2.5.3. Let $\Phi_{H_{\hat{s}}}(\eta)$ be the time-η–flow of the Hamiltonian vector field $X_{H_{\hat{s}}}$. Then the composition

$$u(A_{s,t}) := \Phi_{H_{\hat{s}}}(1-\eta)([A_{2\eta-1,t}]) : (\eta,t)\in[0,1]\times[0,1]\to\hat{\mathcal{M}}^{\mathfrak{g}}(\Sigma')$$

maps $\{0\}\times[0,1]$ to $\Phi_{H_{\hat{s}}}(L_1)$ and maps $\{1\}\times[0,1]$ to L_2 and the stabilizer of every point in the image contains $H=\mathrm{Stab}(B)\cong\mathrm{Stab}([A_1])$ as a subgroup. Using Definition 2.4.13, $u(A_{s,t})$ has a well-defined H–equivariant Maslov index. Here is the precise statement of Theorem 2.1.1.

Theorem 2.5.19. $Sf_H(B,f_{\boldsymbol{q}}) - [\ker d_B] = \mu^H(u(A_{s,t}))$.

Proof. By definition, the equivariant Maslov index $\mu^H(u(A_{s,t}))$ could be computed using the equivariant spectral flow of the operator $-J\frac{\partial}{\partial s}$ on the space of L^2_1–sections

of the symplectic vector bundle

$$u(A_{s,t})^*(T\hat{\mathcal{M}}^{\mathfrak{g}}(\Sigma'))$$

with Lagrangian boundary conditions induced by $T\Phi_{H_{\hat{s}}}(L_1)$ and TL_2. Apply the fiberwise symplectomorphism induced by linearizing $\Phi_{H_{\hat{s}}}(1-\eta)$, use the identification (2.3.6) and use the fact that the intersection $[\tilde{\theta}]$ is non-degenerate, we can see that $\mu^H(u(A_{s,t}))$ is given by the H–equivariant spectral flow of the family

$$-D_0(t) : L_1^2([-1,1]; \mathcal{H}_{A_{s,t}} \cap H^1; \mathcal{L}_{-1}(t), \mathcal{L}_1(t)) \to L^2([-1,1]; \mathcal{H}_{A_{s,t}} \cap H^1)$$

from (2.5.12) where $t \in [\epsilon, 1]$ such that $\epsilon > 0$ is sufficiently small. The symbol $\cap H^1$ means that we require that the 0-form and 2–form components are 0. This is exactly where the shifting $[\ker d_B]$ comes in. By Theorem 2.5.17, we see that $\mu^H(u(A_{s,t}))$ is equal to the H–equivariant spectral flow of the family $\{-K_{B_t,f_{\mathbf{q}}}\}_{\epsilon\leq t\leq 1}$ up to a shifting by $[\ker d_B]$. By concatenating $\{K_{B_t,f_{\mathbf{q}}}\}_{0\leq t\leq 1}$ with a linear path between $K_{\tilde{\theta},f_{\mathbf{q}}}$ and $K_{\theta,0}$, using Theorem 2.5.8, we see that $\mu^H(u(A_{s,t}))$ exactly computes $Sf_H(B, f_{\mathbf{q}})$ (recall that this is defined using the linear path from $(B, f_{\mathbf{q}})$ to $(\theta, 0)$) by homotopy invariance. □

Therefore, we have the following immediate corollary from Definition 2.2.12, Lemma 2.5.18 and Theorem 2.5.19:

Corollary 2.5.20. $\mathrm{ind}(B, f_{\boldsymbol{q}}) = \mu^H(u(A_{s,t})) - \sum_{n_i\geq 2} \frac{\langle\omega, u(\hat{A}^{(i)}_{s,t})\rangle}{2\pi^2 n_i} \cdot \tau_i$, *where* $\langle\omega, u(\hat{A}^{(i)}_{s,t})\rangle$ *is the symplectic area of the family* $u(\hat{A}^{(i)}_{s,t})$ *induced from the* i*–th component of the decomposition of a genuine flat* $\hat{B}$ *near* B. □

Proof of Theorem 2.1.2. Let $H_{\hat{s}}$ and $f_{\mathbf{q}}$ be a compatible pair. For any G–orbits of $\Phi_{H_{\hat{s}}}(L_1) \cap L_2$, without loss of generality we can assume that the disc $D(p)$ comes from a family $A_{s,t}(p) \subset \mathcal{A}^0_F(\Sigma)$ as in this section. If p corresponds to an irreducible SU(3)–

connection $B(p)$, by Theorem 2.5.19 the spectral flow from $K_{B(p),f_{\mathbf{q}}}$ is equal to $\mu(D(p))$ by noticing that the group acts trivially here and $[\ker d_B] = 0$. If p corresponds to a reducible SU(3)–connection $B(p)$, it has stabilizer U(1) and $B(p)$ is gauge-equivalent to an SU(2)–connection. The equivariant Maslov index $\mu^H(u(A_{s,t}))$ is the linear combination of the trivial representation and the weight (-2)–representation of U(1) in the representation ring of U(1). The coefficient before each of them, is given by $\mu_t(D(p))$ and $\mu_n(D(p))$ respectively. Then the theorem follows from checking the formula (2.1.2) term-wisely with the formula in [BH98, Theorem 1]. Note that our definition of the equivariant spectral flow cancels out the factor $\frac{1}{2}$ in *loc.cit.* and our convention of spectral flow adds the term $[\ker d_B]$. Although the formula of Boden-Herald is written down using a holonomy perturbation which is not of the preferred form, the identification between (2.1.2) and $\lambda_{\mathrm{SU}(3)}$ results from the independence of $\lambda_{\mathrm{SU}(3)}$ on holonomy perturbations, established in [BZ20] and Section 2.2. Therefore the theorem is proved. □

Chapter 3

Bifurcations of embedded curves and towards an extension of Taubes' Gromov invariant to Calabi–Yau 3-folds

3.1 Introduction

Let (X, ω) be a symplectic Calabi–Yau 3-fold, i.e., it is a closed symplectic manifold of (real) dimension 6 with $c_1(TX, \omega) = 0 \in H^2(X, \mathbb{Z})$. After choosing an ω-compatible almost complex structure $J \in \mathcal{J}(X, \omega)$, the genus g Gromov–Witten invariant of (X, ω) of class $0 \neq A \in H_2(X, \mathbb{Z})$ with no marked points, denoted by $\mathrm{GW}_{A,g}(X)$, is defined by certain counts of stable J-holomorphic maps into X and it is independent of J. In general, $\mathrm{GW}_{A,g}(X) \in \mathbb{Q}$ because some J-holomorphic maps may possess non-trivial automorphism groups.

On the other hand, Taubes' Gromov invariants [Tau96b] define integer-valued invariants for symplectic 4-manifolds by counting *embedded* pseudo-holomorphic curves.

The construction heavily relies on intersection theory in dimension 4. The goal of this chapter is to generalize Taubes' Gromov invariants to define invariants of symplectic Calabi–Yau 3-folds by studying moduli spaces of embedded pseudo-holomorphic curves.

3.1.1 Statement of the main Theorem

We start by introducing some notions important for our discussion. The term *J-curve* throughout this chapter stands for a simple J-holomorphic map with smooth domain into an almost complex manifold (X, J). Any closed embedded 2-dimensional almost complex submanifold $C \subset X$, gives rise to a J-curve unique up to the reparametrization of the domain. In the following, unless otherwise specified, (X, ω) will denote a symplectic Calabi–Yau 3-fold.

Definition 3.1.1. *Define $\mathcal{W}_{\text{emb}} \subset \mathcal{J}(X, \omega)$ to be the subset of almost complex structures J for which there exists either a simple and non-embedded J-curve or two simple J-curves with distinct but intersecting images.*

It is a basic fact (see [DW20, Theorem 2.1.18] for example) that $\mathcal{W}_{\text{emb}}$ is of codimension 2 in $\mathcal{J}(X, \omega)$. Thus, along generic paths $\gamma \subset \mathcal{J}(X, \omega)$, we may assume that all simple curves are embedded and pairwise disjoint by appealing to the Sard–Smale theorem.

Definition 3.1.2 (Normal deformation operator)**.** *Let $J \in \mathcal{J}(X, \omega)$ and let $C \subset X$ be a smooth closed embedded J-curve, i.e. the tangent bundle of C, written as T_C, is invariant under the J–action. The normal bundle $N_C = T_X|_C / T_C$ is then naturally a complex vector bundle. The linearized Cauchy–Riemann operator on $T_X|_C$ descends to an operator*

$$D^N_{C,J} : \Omega^0(C, N_C) \to \Omega^{0,1}(C, N_C) \tag{3.1.1}$$

called the normal deformation operator *of the embedded curve C.*

Given an embedded J-curve C, the Fredholm index of the operator (3.1.1) is 0. When $\ker D^N_{C,J} = 0$ (as a consequence $\operatorname{coker} D^N_{C,J} = 0$), recall that we can define the sign of C as a point in the moduli space of J-curves as follows.

Definition 3.1.3 (Sign)**.** *Let $C \subset X$ be an embedded J-curve such that $D^N_{C,J}$ is an isomorphism. We then define the sign of the curve C, denoted by*

$$\operatorname{sgn}(C) \in \{-1, +1\} \tag{3.1.2}$$

as follows. We deform $D^N_{C,J}$ in a 1-parameter family to a $\mathbb{C}$-linear operator which is an isomorphism and define the sign of C according to the parity of the associated spectral flow. Alternatively, we consider the straight line homotopy $\{D_t\}_{0\leq t\leq 1}$ from $D_0 = D^N_{C,J}$ to its $\mathbb{C}$-linear part D_1 and trivialize the determinant line bundle

$$t \mapsto \det(D_t) = \det\ker D_t \otimes (\det \operatorname{coker}\ D_t)^* \tag{3.1.3}$$

and define the sign of C to be the sign of the resulting isomorphism between $\det(D_0) = \det(0) \otimes \det(0)^* = \mathbb{R}$ and $\det(D_1)$ with the complex orientation.

The notion of super-rigidity requires that any holomorphic multiple cover of any J-curve $C \subset X$ does not admit any non-trivial normal deformation, more precisely:

Definition 3.1.4 (Super-rigidity)**.** *Given an embedded J-curve $C \subset X$ as in Definition 3.1.2, we say that it is* super-rigid *if for all holomorphic (possibly branched) covers $\varphi : C' \to C$, the pullback operator*

$$\varphi^* D^N_{C,J} : \Omega^0(C', \varphi^* N_C) \to \Omega^{0,1}(C', \varphi^* N_C) \tag{3.1.4}$$

is injective. We say that a given almost complex structure $J \in \mathcal{J}(X, \omega)$ is super-rigid

if and only if $J \in \mathcal{J}(X,\omega) \setminus \mathcal{W}_{\text{emb}}$ *and every embedded* J*-curve is super-rigid.*

Recent work of Wendl [Wen21] shows that the set of super-rigid compatible almost complex structures is a Baire subset of $\mathcal{J}(X,\omega)$, verifying the generic super-rigidity conjecture of Bryan–Pandharipande [BP01]. An important consequence of super-rigidity is the following finiteness result. Suppose J is super-rigid, then given any class $A \in H_2(X,\mathbb{Z})$ and $g \in \mathbb{Z}_{\geq 0}$, there are only finitely many embedded J-curves of class A and genus g and they have disjoint images in X, see [Wen21, Corollary 1.3]. We remark that the work of Doan–Walpuski [DW21, Theorem 1.6] proves this finiteness result with only an *a priori* energy bound, without any assumption on the genus, though we do not need this stronger result in the present work.

Given a homology class $A \in H_2(X,\mathbb{Z})$, we say it is *primitive* if $\omega(A) > 0$ and we cannot write $A = kB$ for any $B \in H_2(X,\mathbb{Z})$ and integer $k \geq 2$. For a fixed primitive homology class A, the work of Zinger [Zin11] and Doan–Walpuski [DW19a] show that the signed count (c.f. Definition 3.1.3) of embedded J-curves of class A and genus $h \geq 0$ for $J \in \mathcal{J}(X,\omega) \setminus \mathcal{W}_{\text{emb}}$ defines an invariant of (X,ω) independent of J. If A is not primitive, although there are only finitely many embedded J-curves of class A and genus h for a super-rigid almost complex structure J, the direct signed count of all such J-curves may not define an invariant of (X,ω) because *bifurcations* could happen when varying the almost comnplex structures. The main result of this chapter is to define a virtual count of embedded J-curves of two times a primitive homology class for *all genera*. We adopt the notation Gr in order to be consistent with Taubes' pioneering work in dimension 4 [Tau96b].

Theorem 3.1.5. *Fix a primitive class* $A \in H_2(X,\mathbb{Z})$ *and an integer* $h \geq 0$*. For a super-rigid* $J \in \mathcal{J}(X,\omega)$*, define the* virtual count of embedded genus h curves of class

$2A$ to be the integer

$$\mathrm{Gr}_{2A,h}(X,\omega,J) = \sum_{C'\colon 2A,h} \mathrm{sgn}(C') + \sum_{g\le h} \sum_{C\colon A,g} \mathrm{sgn}(C)\cdot \mathrm{w}_{2,h}(D^N_{C,J}) \tag{3.1.5}$$

where the first sum is over embedded J-curves C' of genus h and class $2A$, the second sum is over all genera $0 \le g \le h$ and embedded J-curves C of genus g and class A, sgn *is as in Definition 3.1.3 and the integers* $\mathrm{w}_{2,h}$ *are as in Definition 3.5.17.*

Then, $\mathrm{Gr}_{2A,h}(X,\omega,J)$ *is independent of the choice of super-rigid J and*

$$\mathrm{Gr}_{2A,h}(X,\omega) := \mathrm{Gr}_{2A,h}(X,\omega,J) \tag{3.1.6}$$

defines a symplectic invariant of (X,ω).

In other words, we construct correction terms to the naive signed count of embedded J-curves of class $2A$ and genus h from embedded J-curves of class A and genus $g \le h$ to get a well-defined invariant. For a thorough discussion of the correction terms in (3.1.5), see Section 3.5.1.

Remark 3.1.6 (Symplectic deformation invariance). In fact, $\mathrm{Gr}_{2A,h}$ is a symplectic deformation invariant of (X,ω). This follows from the fact that all the results about super-rigidity continue to hold when we replace ω-compatible almost complex structures by ω-tame almost complex structures (see the first paragraph of [Wen21, Section 1.1] for a discussion of this).

3.1.2 Ingredients of the proof

The proof of Theorem 3.1.5 (carried out in Section 3.5.2) relies on a detailed study of bifurcations of moduli spaces of embedded J-curves when varying J in a generic 1-parameter family. To be more precise, although super-rigidity is a generic condition for $J \in \mathcal{J}(X,\omega)$, there are codimension 1 subsets of $\mathcal{J}(X,\omega)$ where super-rigidity

fails to hold. Given a path of almost complex structures $\gamma : [-1, 1] \to \mathcal{J}(X, \omega)$ such that both $\gamma(-1)$ and $\gamma(1)$ are super-rigid, we need to understand the bifurcations of moduli spaces of embedded pseudo-holomorphic curves when γ intersects these codimension 1 subsets.

The first key ingredient is a *necessary* condition for the occurrence of a bifurcation. For a more general and complete statement, see Theorem 3.3.1.

Theorem 3.1.7 (special case of Theorem 3.3.1)**.** *Given a sequence $t_k \to t'$ in $(-1, 1)$, define $J_k := \gamma(t_k)$ and $J := \gamma(t')$ so that $J_k \to J$ in the C^∞-topology. Suppose $h \in \mathbb{Z}_{\geq 0}$ and we have a sequence of genus h embedded J_k-holomorphic curves $\varphi_k : \Sigma'_k \to X$ converging to a stable J-holomorphic map $C' \xrightarrow{\varphi} C \subset X$ in the Gromov topology. Here C is an embedded J-curve of genus $g \leq h$ and $\varphi : C' \to C$ is a stable genus h map of degree k with target C.*

Then we either have $g = h, k = 1$ and $\varphi : C' \to C$ is an isomorphism, or the pullback map $\varphi^* : \ker D^N_{C,J} \to \ker \varphi^* D^N_{C,J}$ is only injective but not surjective.

Theorem 3.1.7 allows us to exclude the possibility of bifurcations from nodal curves (possibly with ghost components) into embedded curves once the path γ is chosen generically. When the homology class represented by $\varphi_k : \Sigma'_k \to X$ is primitive, the degree k is then necessarily equal to 1 and Theorem 3.1.7 implies that the limit φ must be an isomorphism. In particular, the stable map $C' \xrightarrow{\varphi} C \subset X$ does not have any ghost irreducible component. Combined with Ionel–Parker's construction of Kuranishi models for birth-death bifurcations of embedded J-curves [IP18, Theorem 6.2], one can prove the aforementioned result of Zinger and Doan–Walpuski regarding counting embedded J-curves representing primitive homology classes without appealing to Gromov–Witten theory as in [Zin11] or the delicate gluing analysis as from [DW19a], see also Remark 3.3.2. The full power of Theorem 3.1.7 is used in Section 3.5, which tells us ultimately we just need to worry about bifurcations from (branched) double covers with smooth domains in order to check the invariance of

$\mathrm{Gr}_{2A,h}(X,\omega)$.

The second key ingredient is a *sufficient* condition for bifurcations to occur. We construct Kuranishi models of moduli spaces of embedded curves based on a detailed analytical argument. The technical condition of "elementary wall type" includes all the cases in the construction of $\mathrm{Gr}_{2A,h}(X,\omega)$, and in fact is satisfied by a large class of examples, see Example 3.4.2.

Theorem 3.1.8 (see Theorem 3.4.7)**.** *Let $\{J_t\}_{t\in[-1,1]}$ be a generic path in $\mathcal{J}(X,\omega)$. Assume that there exists an embedded rigid J_0-curve $\Sigma \hookrightarrow X$ along with a d-fold genus h branched multiple cover $\varphi : \Sigma' \to \Sigma$ which has non-trivial normal deformations. If this cover determines an elementary wall type (see Definition 3.4.1), then* $\mathrm{Aut}(\varphi) \subset \mathbb{Z}/2\mathbb{Z}$ *and, moreover, the change in the signed count of embedded curves of genus h and class $d[\Sigma]$ near φ is given by $\pm 2/|\mathrm{Aut}(\varphi)|$.*

The notion of "generic path" deserves some further comments. The work of Wendl [Wen21], recalled in Section 3.2.2, specifies the codimension 1 subsets ("walls") of $\mathcal{J}(X,\omega)$ where a bifurcation could occur. Away from a subset of higher codimension, these subsets have manifold structures and we require that the path $\{J_t\}_{t\in[-1,1]}$ only intersects this manifold locus and the intersection is transverse. However, this notion of transversality is *not* sufficient for our discussion. In fact, we define another subset of the walls with positive codimension, called the *degeneracy locus*, in Definition 3.4.21 and the path $\{J_t\}_{t\in[-1,1]}$ is required to avoid the degeneracy locus as well. The condition of "elementary wall type" is needed in this step to guarantee the existence and abundance of such paths. This allows us to guarantee the non-vanishing of certain Taylor coefficients of the Kuranishi maps which leads to Theorem 3.1.8.

For instance, suppose we have a path $\{J_t\}_{t\in[-1,1]} \subset \mathcal{J}(X,\omega)$ and an embedded J_0-holomorphic torus $T \subset X$ with $\ker(D^N_{T,J_0}) = 0$, i.e. the moduli space of J_0-curves is cut out transversely at T. Suppose there exists an unbranched holomorphic double cover from another torus $\varphi : T' \to T$ such that the normal deformation operator

$\varphi^* D^N_{T,J_0}$ satisfies $\dim_{\mathbb{R}} \ker(\varphi^* D^N_{T,J_0}) = 1$. Using Kuranishi reduction, one can show that the parametric moduli space of J_t-holomorphic genus 1 maps in class $2[T]$ (near the J_0-holomorphic map $T' \xrightarrow{\varphi} T \subset X$) is locally modeled on

$$x \cdot (ax^2 + bt) = 0 \text{ for some } a, b \in \mathbb{R} \tag{3.1.7}$$

for $(x, t) \in \mathbb{R} \times [-1, 1]$ with a $\mathbb{Z}/2\mathbb{Z}$-action given by $(x, t) \to (-x, t)$. Wendl's result allows us to ensure that $b \neq 0$ by choosing the path $\{J_t\}_{t\in[-1,1]}$ generically. However, this is *not* sufficient for our application. We argue that one can choose the path generically so that we have $a \neq 0$ as well. As a consequence, the relevant moduli space here is indeed described by the "pitchfork bifurcation".

We remark that our arguments fill in a technical gap in Taubes' paper [Tau96b, Lemma 5.18] and the applicability of our result goes beyond the bifurcation coming from double unbranched covers of tori as in [Tau96b]. As far as we know, this is the first result studying equivariant bifurcations of pseudo-holomorphic curves in dimension 6.

The above two ingredients give a complete characterization of the moduli spaces of embedded pseudo-holomorphic curves of class $2A$ and any genus along a generic path of almost complex structures connecting two given super-rigid almost complex structures. The third key ingredient is the construction of the correction terms in (3.1.5), which are called *linear wall crossing invariants* in Definition 3.5.17. These correction terms are constructed by studying the spaces of $\mathbb{R}$-linear Cauchy–Riemann operators on the normal bundles of embedded J-curves in a symplectic Calabi–Yau 3-fold and they should be thought of as a version of "equivariant spectral flow". This algebraic invariant is used to compensate the change in the signed count of embedded curves of class $2A$ when a bifurcation happens and it serves as the last piece of the construction of $\mathrm{Gr}_{2A,h}(X, \omega)$.

Constructions of this kind probably date back to Walker's extension of the SU(2)-Casson invariant to rational homology 3-spheres [Wal92b] and similar ideas appear in other literature of gauge theory, including the SU(3)-Casson invariants for integer homology 3-spheres [BH98] and the Seiberg–Witten invariants for homology $S^1 \times S^3$ [MRS11]. Recent work of Eftekhary [Eft20] applies similar ideas to a counting problem of periodic orbits. However, in our case, even the definition of the wall crossing numbers requires some work (see Section 3.5.1).

3.1.3 Relations with other work

There are some works in the literature (besides Taubes' Gromov invariant) which have close relations with the current work.

Gopakumar–Vafa invariants

Given a symplectic Calabi–Yau 3-fold (X, ω), it was conjectured by Gopakumar–Vafa [GV98b] that the Gromov–Witten generating function of X satisfies the formula

$$\sum_{\substack{A\neq 0\\ g\geq 0}} \mathrm{GW}_{A,g}(X) t^{2g-2} q^A = \sum_{\substack{A\neq 0\\ h\geq 0}} n_{A,h}(X) \sum_{k=1}^{\infty} \frac{1}{k} \left(2\sin\left(\frac{kt}{2}\right)\right)^{2h-2} q^{kA}, \tag{3.1.8}$$

where the "BPS numbers" $n_{A,h}(X)$ are integers and $n_{A,h}(X) = 0$ for a fixed $A \in H_2(X, \mathbb{Z})$ if h is sufficiently large. The integrality the BPS numbers was proved by Ionel–Parker [IP18] and the finiteness has been established by the recent work of Doan–Ionel–Walpuski [DIW21].

However, Ionel–Parker's proof of $n_{A,h}(X) \in \mathbb{Z}$ does not illustrate the geometric meaning of these Gopakumar–Vafa invariants. Their proof proceeds as follows. [IP18,

Equation (8.1)] shows that the Gromov–Witten generating function satisfies

$$\sum_{A\neq 0, g\geq 0} \mathrm{GW}_{A,g}(X) t^{2g-2} q^A = \sum_{A\neq 0, g\geq 0} e_{A,g}(X) \cdot \mathrm{GW}_g^{\mathrm{elem}}(q^A, t) \tag{3.1.9}$$

where $\mathrm{GW}_g^{\mathrm{elem}}$ is some universal generating series and $e_{A,g}(X) \in \mathbb{Z}$ is regarded as a "*virtual count of embedded clusters in* X". They also show that $n_{A,h}(X)$ and $e_{A,g}(X)$ actually determine each other via the formula [IP18, Equation (8.9)]

$$n_{A,h}(X) = \sum_{\substack{A=dB \\ B\in\mathbb{Z}_{>0}}} \sum_{g=0}^{h} n_{d,h}(g) \cdot e_{B,g}(X) \tag{3.1.10}$$

where $n_{d,h}(g) \in \mathbb{Z}$ are certain universal coefficients determined by d, h, g. Therefore the integrality of $n_{A,h}(X)$ is a result of integrality of $e_{A,g}(X)$.

Conjecture 3.1.9. *The invariant* $\mathrm{Gr}_{2A,h}(X,\omega)$ *is equal to the virtual count of embedded clusters* $e_{2A,h}$ *from* (3.1.9) *for all* $h \geq 0$ *where* A *is a primitive homology class.*

We show that Conjecture 3.1.9 is true for $h = 0$, see Theorem 3.6.10. The proof relies on a wall-crossing result of virtual Euler numbers, which might be of independent interest.

Theorem 3.1.10 (see Theorem 3.6.5)**.** *Suppose* $N \to C$ *is a rank* 2 *complex vector bundle of degree* $2g-2$ *on a curve* C *of genus* g*. Let* $\mathcal{D} = \{D_t\}_{t\in[-1,1]}$ *be a generic* 1*-parameter family of Cauchy–Riemann operators on* N*. Fix integers* $d \geq 2$*,* $h \geq g$ *and consider the moduli stack* $\overline{\mathcal{M}}_h(C,d)$ *of stable genus* h *maps to* C *of degree* d*. Assume further that* $\mathcal{D}$ *is* (d,h)*-rigid (see Definition 3.6.2) except over a finite set corresponding to holomorphic maps with smooth domains*

$$\Delta \times \{0\} \subset \mathcal{M}_h(C,d) \times [-1,1]. \tag{3.1.11}$$

Then, the assignment

$$[f : C' \to C] \in \overline{\mathcal{M}}_h(C, d) \mapsto \text{coker } (f^* D_t) \tag{3.1.12}$$

gives a well-defined vector bundle on $\overline{\mathcal{M}}_h(C, d)$*, equipped with a natural orientation, for small* $t \neq 0$*. Denote the Euler number of this bundle by* $e_{d,h}(C, D_\pm)$ *for small* $\pm t > 0$*. Then, we have*

$$e_{d,h}(C, D_+) - e_{d,h}(C, D_-) = \sum_{p \in \Delta} \frac{2 \cdot \text{sgn}(\mathcal{D}, p)}{|\text{Aut}(p)|} \tag{3.1.13}$$

where $\text{sgn}(\mathcal{D}, p) \in \{-1, +1\}$ *are determined by the local behavior of* $\mathcal{D}$ *near* p*.*

The main obstacle to proving Conjecture 3.1.9 in general is to drop the assumption that Δ is a finite subset and corresponds to stable maps with smooth domains in the above statement.

When (X, ω) is further assumed to be a smooth projective variety, Maulik–Toda [MT18] propose a sheaf-theoretic definition of the Gopakumar–Vafa invariants, based on earlier works of Hosono–Saito–Takahashi [HST01] and Kiem–Li [KL12]. Despite the beauty of the definition, it is not known how to prove the deformation invariance of Maulik–Toda's ansatz and the proof of formula (3.1.8) using this definition seems to be out of reach in general. Moreover, the sheaf-theoretic definition relies crucially on the algebraicity of X, making it hard to generalize it to general symplectic Calabi–Yau 3-folds (equipped with not necessarily integrable almost complex structures). Conjecture 3.1.9 proposes a geometric definition of the Gopakumar–Vafa invariants of two times a primitive homology class via the formula (3.1.10) by counting embedded pseudo-holomorphic curves.

Remark 3.1.11. Conjecture 7.4.5 in Cox–Katz's book [CK99] conjectures that the genus 0 BPS numbers should come from virtual counts of embedded genus 0 J-

curves, with correction terms coming from genus 0 J-curves of lower degree. Later on, Katz proposes an algebro-geometric counterpart of this conjecture in [Kat08] using Donaldson–Thomas invariants. Our conjecture and result are consistent with these expectations.

Counting calibrated submanifolds

For a given symplectic manifold (X, ω) and an ω-compatible almost complex structure J, almost complex submanifolds, in particular embedded J-curves, define *calibrated submanifolds* of X. Another example of calibrated submanifolds are associative submanifolds of 7-dimensional Riemannian manifolds with holonomy group G_2.

There is a proposal from Doan–Walpuski [DW19b] which defines weights from counting solutions to gauge-theoretic equations on the associative submanifolds such that the weighted count of associatives defines an invariant. In this work, a symplectic definition of Pandharipande–Thomas invariant is discussed, based on a weighted count of embedded pseudo-holomorphic curves with weights again coming from counts of gauge-theoretic equations.

Our formula (3.1.5) could also be viewed as a weighted count of embedded pseudo-holomorphic curves which defines an invariant of (X, ω). This definition bypasses the difficulty coming from compactness issues of moduli spaces of generalized Seiberg–Witten equations which enters the proposal in [DW19b] and has a closer relation with the Gromov–Witten invariants.

Study of multi-valued sections

Given an embedded J-curve C in a symplectic Calabi–Yau 3-fold and a holomorphic (branched) cover $\varphi : C' \to C$, any non-zero element $s \in \ker(\varphi^* D^N_{C,J})$ defines a *multi-valued section* of the normal bundle N of C in X by associating any point $x \in C$ with the values of s over the set $\varphi^{-1}(x)$. Our bifurcation result (Theorem 3.1.8) can be

interpreted as a sufficient condition to perturb a (holomorphic) multi-valued section into a "nearby" embedded pseudo-holomorphic curve in X. Moreover, the discussion in Section 3.5.1 which is required for the definition of linear wall crossing numbers studies some generic behaviors of these multi-valued sections.

As speculated in [Don19], 2-valued harmonic 1-forms are important for special Lagrangian geometry. Our study of multi-valued sections might have some applicability in this context, see Remark 3.4.8 and Remark 3.5.13.

3.1.4 Technical remarks concerning transversality arguments

To carry out the numerous generic transversality arguments in this chapter (and in many other places in the theory of pseudo-holomorphic curves), one needs to appeal to the Sard–Smale theorem (for Fredholm maps of second countable *Banach* manifolds) after showing that certain universal moduli spaces are smooth (respectively, C^∞ subvarieties as in [Wen21, Appendix C]) by showing that the linearized operators associated to their defining equations are surjective (respectively, satisfy a lower bound on their rank).

Let $f : \mathcal{X} \to \mathcal{Y}$ be the map for which we wish to prove that the set $\mathcal{Y}_{\mathrm{reg}} \subset \mathcal{Y}$ of regular values is a Baire subset (with $\mathcal{Y}$ being a space of smooth objects and both $\mathcal{X}$, $\mathcal{Y}$ being second countable). In our applications, this will be the space of smooth almost complex structures on a manifold or Cauchy–Riemann operators on a fixed bundle or spaces of paths/homotopies in the aforementioned spaces or, in our most complicated situation, a space of embedded pseudo-holomorphic curves equipped with the data of a branched cover of the domain (Definition 3.4.21). The first step of the proof is always to establish *formal regularity*, i.e., surjectivity (or rank lower bounds) for the linearized operators of the defining equations of $\mathcal{X}$ over $\mathcal{Y}$ at the level of smooth objects (i.e. without taking any Banach space completions or restricting to any Banach subspace).

Now, for most applications (including all the applications in this chapter), one wishes to prove C^∞-genericity statements (e.g. "simple J-holomorphic maps are unobstructed for a C^∞-generic J"). However, the space $\mathcal{Y}$ does not form a Banach manifold (typically, it is just a Fréchet manifold) and thus, we cannot directly appeal to the Sard–Smale theorem. Fortunately, there is a well-known technique to overcome this difficulty (i.e. to upgrade formal regularity into genericity of regular values) which we briefly outline below. For the complete details and non-trivial examples of the application of this technique, the reader is referred to the very clear and detailed exposition in [Wen21] (specifically the proofs of Theorems 5.25 and D therein). Going from formal regularity to the statement that regular values are generic involves showing that showing that $\mathcal{Y}_{\text{reg}}$ is a countable intersection of dense open subsets of $\mathcal{Y}$.

For this, we use the popular trick (due to Taubes) of exhausting $\mathcal{X}$ by a sequence of subspaces $\mathcal{X}^1 \subset \mathcal{X}^2 \subset \cdots$ such that $f|_{\mathcal{X}^N} : \mathcal{X}^N \to \mathcal{Y}$ is proper for all N. We then define $\mathcal{Y}^N_{\text{reg}} \subset \mathcal{Y}$ to be the subset of points $y \in \mathcal{Y}$ for which every point in $f^{-1}(y) \cap \mathcal{X}^N$ is regular. Now, since regularity is an open property on $\mathcal{X}$ and $f|_{\mathcal{X}^N}$ is proper, it follows that $\mathcal{Y}^N_{\text{reg}}$ is open. As a consequence, $\mathcal{Y}_{\text{reg}} = \bigcap_{N \geq 1} \mathcal{Y}^N_{\text{reg}}$ is a countable intersection of open subsets of $\mathcal{Y}$. The exact construction of the subsets $\mathcal{X}^N \subset \mathcal{X}$ is straightforward but somewhat tedious and we refer the reader to the proofs of Theorems 5.2.5 and D in [Wen21] for details.

We are left to argue that each $\mathcal{Y}^N_{\text{reg}} \subset \mathcal{Y}$ is dense. For this, we fix any reference point $y_{\text{ref}} \in \mathcal{Y}$ introduce Floer's C_ϵ spaces centered at y_{ref}, which we denote by $\mathcal{Y}_\epsilon(y_{\text{ref}})$. These spaces are Banach manifolds, consist entirely of smooth objects close to y_{ref} but carry a topology depending on the parameter ϵ which is much finer than the C^∞ topology. For any $x_{\text{ref}} \in \mathcal{X}_{\text{ref}} = f^{-1}(y_{\text{ref}})$, formal regularity and the implicit function theorem imply that for all sufficiently small ϵ, the space $\mathcal{X}_\epsilon(y_{\text{ref}}) = f^{-1}(\mathcal{Y}_\epsilon(y_{\text{ref}}))$ is a Banach manifold (or C^∞ subvariety of expected codimension) near x_{ref}. By the

second countability of $\mathcal{X}_{\mathrm{ref}}$, we can choose a uniform ϵ which works for all $x_{\mathrm{ref}} \in \mathcal{X}_{\mathrm{ref}}$. Thus, we obtain an open neighborhood $\mathcal{X}_\epsilon(y_{\mathrm{ref}})^{\mathrm{reg}}$ of $\mathcal{X}_{\mathrm{ref}}$ in $\mathcal{X}_\epsilon(y_{\mathrm{ref}})$ such that the restriction of f as a map $\mathcal{X}_\epsilon(y_{\mathrm{ref}})^{\mathrm{reg}} \to \mathcal{Y}_\epsilon(y_{\mathrm{ref}})$ is a smooth map of Banach manifolds. Now, the Sard–Smale theorem shows that we can find a sequence $y_k \in \mathcal{Y}_\epsilon(y_{\mathrm{ref}})$ of regular values of this map converging to y_{ref} in $\mathcal{Y}$. For any $N \geq 1$, the properness of $\mathcal{X}^N$ over $\mathcal{Y}$ and the openness of $\mathcal{X}_\epsilon(y_{\mathrm{ref}})^{\mathrm{reg}} \subset \mathcal{X}_\epsilon(y_{\mathrm{ref}})$ now show that we must have $y_k \in \mathcal{Y}^N_{\mathrm{reg}}$ for all sufficiently large k (depending on N). Since y_{ref} was arbitrary, this shows that each $\mathcal{Y}^N_{\mathrm{reg}}$ is dense in $\mathcal{Y}$.

In view of the above, we adopt the convention of only proving formal regularity in all transversality proofs that arise in this chapter (and omitting the use of Floer's C_ϵ spaces and Taubes' trick to upgrade from formal regularity to genericity of regular values). Consistent with this, we use the notions of codimension of a subset of $\mathcal{Y}$ (or a C^∞ subvariety of $\mathcal{Y}$), the index of the map f etc in this formal sense as a shorthand for the more detailed proof scheme above. The corresponding notions for maps of Banach manifolds have standard definitions (a subset is of codimension $\geq k$ if it is covered by the images of countably many Fredholm maps with indices $\leq -k$; C^∞ subvarieties of codimension $\geq k$ are defined as in [Wen21, Appendix C] and we note that they are, in particular, subsets of codimension $\geq k$ by [Wen21, Proposition C.1]). The above formal notions will translate to the corresponding notions for Banach manifolds when we apply the implicit function theorem in the above proof scheme to a suitable C_ϵ space.

3.1.5 Structure of the chapter

We end the introduction by outlining the structure of this chapter. In Section 3.2, we recall the necessary background on (twisted) Cauchy–Riemann operators and superrigidity (following [Wen21]). Section 3.3 establishes a necessary condition for bifurcations to occur (Theorem 3.3.1) while Section 3.4 gives a sufficient condition (Theorem

3.4.7). In Section 3.5, we apply the results of the previous two sections to construct the extension of Taubes' Gromov invariant for symplectic Calabi–Yau 3-folds (counting curves of genus $h \geq 0$ lying in two times a primitive homology class) and prove the main result of this chapter (Theorem 3.1.5). Finally, in Section 3.6, we prove a result on the virtual Euler numbers of certain obstruction bundles (Theorem 3.6.5) which allows us to identify our invariant with the Gopakumar–Vafa invariant when the genus h is equal to 0 (Theorem 3.6.10).

Acknowledgements.

We would like to thank Chris Wendl for helpful correspondences and for pointing out an error in an earlier version of Lemma 3.4.10. We would also like to thank Eleny Ionel for a careful reading of an earlier version of this chapter and for several useful suggestions on improving the exposition. We are also grateful to Aleksander Doan, Sridhar Venkatesh and Thomas Walpuski for useful discussions. Finally, we would like to thank our advisor John Pardon for his constant support and encouragement.

Conventions.

Unless otherwise stated, all vector spaces, vector bundles and their determinants are defined over $\mathbb{R}$. Dimensions, Hom-spaces and tensor products without any subscripts should be understood to be over $\mathbb{R}$.

3.2 Review of embedded curves and super-rigidity

In this section, we will summarize the background on super-rigidity in a form relevant for our discussion in sections 3.4–3.5. The notation closely follows that of [Wen21, Section 2]. We will begin by recalling some basic facts about branched covers of curves and Cauchy–Riemann operators (in §3.2.1) and then move onto super-rigidity

(in §3.2.2).

3.2.1 Cauchy–Riemann operators and Galois covers

We will consistently use the word 'curve' to mean 'irreducible smooth projective algebraic curve over $\mathbb{C}$' (i.e. a closed connected smooth Riemann surface). A 'morphism' between curves will mean a 'non-constant holomorphic map'.

Definition 3.2.1 (Function field)**.** *Let Σ be a curve. Define $\mathbb{C}(\Sigma)$ to be the* function field *of Σ (i.e., the field of meromorphic functions on Σ). Given a degree $d \geq 1$ morphism of curves $\varphi : \Sigma' \to \Sigma$, we get the associated degree d field extension $\varphi^* : \mathbb{C}(\Sigma) \hookrightarrow \mathbb{C}(\Sigma')$. It is a basic fact from algebraic geometry that this is a contravariantly functorial one-to-one correspondence (i.e. an equivalence of categories) between curves and fields of transcendence degree 1 over $\mathbb{C}$.*

Definition 3.2.2 (Galois closure)**.** *Let $\varphi : \Sigma' \to \Sigma$ be a morphism of curves. We say that φ is* Galois *if and only if the associated field extension is Galois. In general, φ may not be Galois and we can form the Galois closure, a morphism $\tilde{\Sigma} \to \Sigma' \xrightarrow{\varphi} \Sigma$, by taking the Galois closure of the function field extension $\varphi^* : \mathbb{C}(\Sigma) \to \mathbb{C}(\Sigma')$. For a Galois cover $\tilde{\Sigma} \to \Sigma$, we define the Galois group $\mathrm{Gal}(\tilde{\Sigma}/\Sigma)$ to be the group of automorphisms of $\tilde{\Sigma}$ over Σ (or equivalently the Galois group of the associated function field extension).*

Given a morphism $\varphi : \Sigma' \to \Sigma$ of curves, we call its critical points as *ramification points* and its critical values as *branch points*. Away from the ramification points, φ restricts to a covering map with connected domain over the complement of the branched points, which is called *the associated covering space*. It is a basic fact that $\varphi : \Sigma' \to \Sigma$ is Galois if and only if the associated covering space is normal. Also, the associated covering space of the Galois closure of φ is the normal closure of the associated covering space of φ.

Definition 3.2.3 (Branched covers and Galois groups)**.** *Let the following data be given: integers $g \geq 0$ and $d \geq 2$ and* branching data $\boldsymbol{b} = (\boldsymbol{b}_1, \ldots, \boldsymbol{b}_r)$ *where each $\boldsymbol{b}_i = (b_i^1, \ldots, b_i^{q_i})$ is a tuple of positive integers adding up to d with $b_i^1 \geq \cdots \geq b_i^{q_i}$ and $b_i^1 \geq 2$. Define*

$$b = \sum_{i=1}^{r} \sum_{j=1}^{q_i} (b_i^j - 1) \tag{3.2.1}$$

and define h using the formula $2h-2 = b+d(2g-2)$. Assume that h is a non-negative integer. Now, let $\varphi : \Sigma' \to \Sigma$ be a morphism of curves of type $(g, \boldsymbol{b})$*, i.e., $h =$* genus(Σ')*, $g =$* genus(Σ)*, $d = \deg(\varphi)$ and φ has exactly r critical values $w_1, \ldots, w_r \in \Sigma$ with $\varphi^{-1}(w_i) = \{\zeta_i^1, \ldots, \zeta_i^{q_i}\}$ and φ is b_i^j-to-1 in a punctured neighborhood of each $\zeta_i^j \in \Sigma'$. We define the* generalized automorphism group *of φ to be the Galois group of the Galois closure of φ.*

Given a finite group G and a genus g curve Σ, define $\mathcal{M}_{\boldsymbol{b},G}^d(\Sigma)$ to be the orbifold of pairs $(\varphi : \Sigma' \to \Sigma, \iota)$, up to isomorphism, where φ is a morphism of type $(g, \boldsymbol{b})$ and ι is an isomorphism of G with the generalized automorphism group of φ.

We will next discuss Cauchy–Riemann operators on curves and their pullbacks along Galois covers. Denote by E a complex vector bundle over a curve Σ. Recall that an $\mathbb{R}$-*linear Cauchy–Riemann operator* is a first-order real linear partial differential operator $D : \Omega^0(\Sigma, E) \to \Omega^{0,1}(\Sigma, E)$ such that for any smooth function $f \in C^\infty(\Sigma, \mathbb{R})$ and section $s \in \Omega^0(\Sigma, E)$, we have

$$D(fs) = (\overline{\partial} f)s + fD(s), \tag{3.2.2}$$

where $\overline{\partial} f = \frac{1}{2}(df + i \cdot df \circ j)$ and j denotes the complex structure on Σ.

Definition 3.2.4 (Serre dual)**.** *Let $D : \Omega^0(\Sigma, E) \to \Omega^{0,1}(\Sigma, E)$ be an $\mathbb{R}$-linear*

Cauchy–Riemann operator. Define the Serre dual bundle *of E to be*

$$E^\dagger = \Lambda^{1,0} T_\Sigma^* \otimes_{\mathbb{C}} E^* \tag{3.2.3}$$

and the Serre dual operator $D^\dagger$ *to be the $\mathbb{R}$-linear Cauchy–Riemann operator on $E^\dagger$ determined by the following condition: for all $s \in \Omega^0(\Sigma, E)$ and $s^\dagger \in \Omega^0(\Sigma, E^\dagger)$,*

$$\int_\Sigma \langle Ds, s^\dagger \rangle + \langle s, D^\dagger s^\dagger \rangle = 0. \tag{3.2.4}$$

The pairing $\langle \cdot, \cdot \rangle$ followed by integration gives the Serre duality *isomorphisms*

$$\ker D^\dagger = (\text{coker } D)^* \tag{3.2.5}$$

$$\ker D = (\text{coker } D^\dagger)^*. \tag{3.2.6}$$

Definition 3.2.5 (Twisted operators)**.** *Let $D : \Omega^0(\Sigma, E) \to \Omega^{0,1}(\Sigma, E)$ be an $\mathbb{R}$-linear Cauchy–Riemann operator on the curve Σ and let $\psi : \tilde{\Sigma} \to \Sigma$ be a Galois cover with Galois group $G = \text{Gal}(\tilde{\Sigma}/\Sigma)$. We can then form the pullback operator*

$$\psi^* D : \Omega^0(\tilde{\Sigma}, \psi^* E) \to \Omega^{0,1}(\tilde{\Sigma}, \psi^* E) \tag{3.2.7}$$

which is G-equivariant with respect to the natural linear G-actions on the source and target. Let θ be a linear G-representation over $\mathbb{R}$. We define the twisted operator

$$D_{\psi,\theta} : \text{Hom}^G(\theta^*, \Omega^0(\tilde{\Sigma}, \psi^* E)) \to \text{Hom}^G(\theta^*, \Omega^{0,1}(\tilde{\Sigma}, \psi^* E)) \tag{3.2.8}$$

by taking the θ-isotypic piece of Equation (3.2.7)*. The operator $D_{\psi,\theta}$ is naturally linear over the endomorphism algebra $\mathbb{K}_\theta = \text{End}^G(\theta)$.*

Remark 3.2.6. This definition is easily seen to be equivalent to the twisted operators from [Wen21, §2.2 and §3.4] since taking G-invariant sections of $\theta \otimes \psi^* E$ on $\tilde{\Sigma}$ amounts

to taking sections of $(\tilde{\Sigma} \times V)/G$ on the orbifold $\tilde{\Sigma}/G$ (which has a natural map to Σ being an isomorphism except over the branch points).

Remark 3.2.7. Suppose $G \to G'$ is a quotient map between groups, corresponding to a factorization $\tilde{\Sigma} \xrightarrow{\varphi} \Sigma' \xrightarrow{\psi'} \Sigma$ of ψ as the composition of two Galois covers φ, ψ' and $G' = \mathrm{Gal}(\Sigma'/\Sigma)$. Also, suppose that θ' is a G'-representation and that θ is its pullback to G. Then, φ^* gives a natural map of complexes from $D_{\psi',\theta'}$ to $D_{\psi,\theta}$ and this map is easily seen to be a quasi-isomorphism (using the removable singularity theorem near the branch points of φ for D and $D^\dagger$). Similarly, if $H \subset G$ is any subgroup and $\tilde{\Sigma}' = \tilde{\Sigma}/H$ with the associated morphism $\psi' : \tilde{\Sigma}' \to \Sigma$, then the natural pullback map from $(\psi')^* D$ to $(\psi^* D)^H$ is also a quasi-isomorphism (i.e. an isomorphism on kernels and cokernels).

The following result, which appears in [DW20, Remark 1.4.15] in a slightly different form, will also be useful for the discussion below.

Lemma 3.2.8. *Let θ_i for $i = 1, \ldots, p$ be an enumeration of all the irreducible representations of G over $\mathbb{R}$ with $\mathbb{K}_i = \mathrm{End}^G(\theta_i)$. Then, the tautological map*

$$\bigoplus_i \theta_i^* \otimes_{\mathbb{K}_i} D_{\psi,\theta_i} \to \psi^* D \tag{3.2.9}$$

of G-equivariant two term complexes is a quasi-isomorphism.

Proof. This is because for any finite dimensional $\mathbb{R}[G]$-module V, the tautological map $\bigoplus_i \theta_i^* \otimes_{\mathbb{K}_i} \mathrm{Hom}^G(\theta_i^*, V) \to V$ is an isomorphism of $\mathbb{R}[G]$-modules by Maschke's theorem. We apply this to the (co)kernels of the operators in question. □

Given a family of curves and a corresponding family of Cauchy–Riemann operators, we introduce the following definition which measures the variation of the Cauchy–Riemann operators on a particular finite-dimensional vector space. It will be used later in the analysis of Kuranishi maps.

Let $\pi : \mathcal{C} \to \mathcal{B}$ be a smooth family of curves, i.e., π is a proper submersion between two manifolds $\mathcal{C}$ and $\mathcal{B}$ with fibers of $\mathbb{R}$-dimension 2 along with a smooth endomorphism j of $\ker(d\pi) = T_{\mathcal{C}/\mathcal{B}}$ satisfying $j^2 = -1$. Let $E \to \mathcal{C}$ be a $\mathbb{C}$-vector bundle and set

$$F = \Lambda^{0,1} T^*_{\mathcal{C}/\mathcal{B},j} \otimes_{\mathbb{C}} E. \tag{3.2.10}$$

Assume that we have a smooth family $\mathcal{D} = \{D_b\}_{b\in\mathcal{B}}$ of Cauchy–Riemann operators

$$D_b : \Omega^0(\mathcal{C}_b, E_b) \to \Omega^0(\mathcal{C}_b, F_b) \tag{3.2.11}$$

where $\mathcal{C}_b = \pi^{-1}(b)$, $E_b = E|_{\mathcal{C}_b}$ and $F_b = F|_{\mathcal{C}_b}$ for $b \in \mathcal{B}$.

Definition 3.2.9. *For any point $a \in \mathcal{B}$, define the* linearized map *at a associated to the family $\mathcal{D}$*

$$\mathcal{D}'|_a : T_a\mathcal{B} \to \mathrm{Hom}(\ker D_a, \mathrm{coker}\ D_a) \tag{3.2.12}$$

as follows. Given $\sigma \in \ker D_a$, extend it arbitrarily to a section $\tilde{\sigma} \in \Omega^0(\mathcal{C}, E)$. Then, the section $\tilde{\tau} = \mathcal{D}(\tilde{\sigma}) \in \Omega^0(\mathcal{C}, F)$ defined by $\tilde{\tau}|_{\mathcal{C}_b} = D_b(\tilde{\sigma}|_{\mathcal{C}_b})$ for $b \in \mathcal{B}$ vanishes along $\mathcal{C}_a$ and thus, has a well-defined normal derivative

$$\nabla\tilde{\tau}|_a \in \mathrm{Hom}(N_{\mathcal{C}_a/\mathcal{C}}, F_a) = \mathrm{Hom}(T_a\mathcal{B}, \Omega^0(\mathcal{C}_a, F_a)) \tag{3.2.13}$$

where we have used the identification $N_{\mathcal{C}_a/\mathcal{C}} = \pi^*(T_a\mathcal{B})$ in the last equality. Now, for any $v \in T_a\mathcal{B}$ we set $\mathcal{D}'|_a(v) \cdot \sigma \in \mathrm{coker}\ D_a$ to be the projection of the section $\nabla\tilde{\tau}|_a \cdot v \in \Omega^0(\mathcal{C}_a, F_a)$.

The operator (3.2.12) is well-defined because if $\sigma = 0$, then any choice of $\tilde{\sigma}$ has a

well-defined normal derivative

$$\nabla\tilde{\sigma}|_a \in \mathrm{Hom}(N_{\mathcal{C}_a/\mathcal{C}}, E_a) = \mathrm{Hom}(T_a\mathcal{B}, \Omega^0(\mathcal{C}_a, E_a)) \tag{3.2.14}$$

which satisfies $D_a(\nabla\tilde{\sigma}|_a \cdot v) = \nabla\tilde{\tau}|_a \cdot v$ for all $v \in T_a\mathcal{B}$. Note that if a finite group G acts compatibly on $\mathcal{B}$, $\mathcal{C}$, E, F, with the action on $\mathcal{B}$ being trivial, such that the family $\mathcal{D}$ is G-equivariant, then (3.2.12) factors through $\mathrm{Hom}^G(\ker D_a, \mathrm{coker}\ D_a)$.

Remark 3.2.10. It is easily seen from Definition 3.2.9 that we can define the linearized map $\mathcal{D}'|_a$ by first choosing suitable trivializations to regard $\mathcal{D}$ as a family of G-equivariant operators

$$b \mapsto D_b : \Omega^0(\mathcal{C}_a, E_a) \to \Omega^0(\mathcal{C}_a, F_a) \tag{3.2.15}$$

and setting $\mathcal{D}'|_a(v) \cdot \sigma$ to be the projection of $\left[\frac{d}{dt} D_{\gamma(t)}\sigma|_{t=0}\right]$ onto coker D_a, where γ is any smooth path in $\mathcal{B}$ with $\gamma(0) = a$ and $\dot{\gamma}(0) = v$.

3.2.2 Super-rigidity

For the remainder of this section, we assume that the manifold X is a symplectic Calabi–Yau 3-fold with symplectic form ω and, moreover, all almost complex structures under consideration will be assumed to be ω-compatible.

Fix $g, d, \mathbf{b}$ as in Definition 3.2.3 and a finite group G. Now, given $A \in H_2(X, \mathbb{Z})$, define the universal moduli space

$$\mathcal{M}^d_{\mathbf{b},G}(\mathcal{M}^*_g(X, A)) \to \mathcal{J}(X, \omega) \tag{3.2.16}$$

to consist of tuples $(J, \Sigma, \varphi : \Sigma' \to \Sigma, \iota)$ up to isomorphism where $J \in \mathcal{J}(X, \omega)$, $\Sigma \subset X$ is an embedded J-curve of genus g in class A and φ is a branched cover of type $(g, \mathbf{b})$ along with an isomorphism ι between G and its generalized automorphism

group (this is clearly a purely ineffective smooth Fréchet orbifold). Given any such tuple, we can then form the associated Galois closure $\psi : \tilde{\Sigma} \to \Sigma$ and consider the G-equivariant operator $\psi^* D^N_{\Sigma,J}$. Clearly its kernel and cokernel are finite dimensional real G-representations.

Definition 3.2.11 (Loci of failure of super-rigidity)**.** *Given two nonzero finite dimensional G-representations $\boldsymbol{k}$ and $\boldsymbol{c}$ over $\mathbb{R}$, define the universal moduli space*

$$\mathcal{M}^d_{\boldsymbol{b},G}(\mathcal{M}^*_g(X,A);\boldsymbol{k},\boldsymbol{c}) \to \mathcal{J}(X,\omega) \tag{3.2.17}$$

to be the subset of $\mathcal{M}^d_{\boldsymbol{b},G}(\mathcal{M}^*_g(X,A))$ consisting of those tuples $(J,\Sigma,\varphi : \Sigma' \to \Sigma,\iota)$, with associated Galois cover $\psi : \tilde{\Sigma} \to \Sigma$, for which we have isomorphisms $\ker \psi^* D^N_{\Sigma,J} \simeq \boldsymbol{k}$ *and* $\operatorname{coker} \psi^* D^N_{\Sigma,J} \simeq \boldsymbol{c}$ *of G-representations.*

We now state the version of [Wen21, Theorem D] that we will need for the rest of the article. Some immediate corollaries are also explained below.

Theorem 3.2.12 (Stratification)**.** *In the setting of Definition 3.2.11, away from an exceptional subset of infinite codimension of $\mathcal{J}(X,\omega)$, the subspace*

$$\mathcal{M}^d_{\boldsymbol{b},G}(\mathcal{M}^*_g(X,A);\boldsymbol{k},\boldsymbol{c}) \subset \mathcal{M}^d_{\boldsymbol{b},G}(\mathcal{M}^*_g(X,A)) \tag{3.2.18}$$

is a Fréchet sub-orbifold of codimension given by $\dim_{\mathbb{R}} \operatorname{Hom}^G(\mathbf{k},\mathbf{c})$.

Remark 3.2.13. In the statement of Theorem 3.2.12, we are using the fact that the Petri condition (see [Wen21, Section 5] for more details) holds on $\mathcal{J}(X,\omega)$ away from an exceptional subset which has infinite codimension. In practice, we can ignore this exceptional set as in the next corollary.

Corollary 3.2.14. *Let P be a second countable manifold and let $\mathcal{J}_P(X,\omega)$ denote the space of smooth ω-compatible almost complex structures parametrized by P with the*

C^∞ topology. Let $\mathcal{P} = \{J_s\}_{s\in P}$ be an element in $\mathcal{J}_P(X,\omega)$ for which the associated map $P \to \mathcal{J}(X,\omega)$ is transverse to the forgetful map (3.2.17) *over a closed submanifold $Q \subset P$. Define $\mathcal{J}(X,\omega)_{\mathcal{P},Q}$ to be the Fréchet manifold of $\mathcal{P}' \in \mathcal{J}_P(X,\omega)$ which coincide with $\mathcal{P}$ over $Q \subset P$. Then, the subset of $\mathcal{J}(X,\omega)_{\mathcal{P},Q}$ consisting of those $\mathcal{P}'$ for which the associated map $P \to \mathcal{J}(X,\omega)$ is transverse to* (3.2.17) *is a Baire set, and therefore, dense.*

Definition 3.2.15 (Wall)**.** *For $\delta = (g, d, \boldsymbol{b}, G, \boldsymbol{k}, \boldsymbol{c}, A)$ as in Definition 3.2.11 define*

$$\mathcal{W}^d_{\boldsymbol{b},G}(g, A, \boldsymbol{k}, \boldsymbol{c}) \subset \mathcal{J}(X,\omega) \tag{3.2.19}$$

to be the image of the forgetful map (3.2.17)*. Define the* wall $\mathcal{W} \subset \mathcal{J}(X,\omega)$ *to be the set of all J which are not super-rigid. Clearly, we have*

$$\mathcal{W} = \mathcal{W}_{\text{emb}} \cup \bigcup_{\delta} \mathcal{W}^d_{\mathbf{b},G}(g, A, \mathbf{k}, \mathbf{c}) \tag{3.2.20}$$

by passing to Galois covers in Definition 3.1.4. We refer to each term in the union in (3.2.20) *as a* stratum *of the wall $\mathcal{W}$.*

We will now determine a way to separate $\mathcal{W}$ into those strata which have codimension 1 in $\mathcal{J}(X,\omega)$ and those which have codimension ≥ 2. We do this essentially by rewording the proof of [Wen21, Theorem A] for X and following the remarks from [Wen21, §2.4].

Lemma 3.2.16 (Wall: codimension 1 strata)**.** *Away from a codimension 2 subset of $\mathcal{J}(X,\omega)$, the codimension 1 stratum of $\mathcal{W} \subset J(X,\omega)$ is a union of those strata $\mathcal{W}^d_{\boldsymbol{b},G}(g, A, \boldsymbol{k}, \boldsymbol{c})$ which are non-empty such that*

(i) *each $b^j_i \in \{1, 2\}$, i.e., each ramification point is simple and,*

(ii) *the G-representation $\boldsymbol{k}$ is faithful, irreducible and of real type, i.e., its equivariant endomorphism algebra* $\mathrm{End}^G(\mathbf{k})$ *is isomorphic to $\mathbb{R}$.*

Proof. Begin by noting that $\mathcal{W}_{\text{emb}}$ is of codimension 2 in $\mathcal{J}(X,\omega)$. Now, take J in $\mathcal{W} \setminus \mathcal{W}_{\text{emb}}$ and not in the exceptional set from Theorem 3.2.12. This means that we can find a minimum possible $d \geq 1$, an embedded J-curve $\Sigma \subset X$ and a d-fold branched cover $\varphi : \Sigma' \to \Sigma$ such that $\ker \varphi^* D^N_{\Sigma,J} \neq 0$. Let $\psi : \tilde{\Sigma} \to \Sigma$ be the Galois closure of φ and let $G = \text{Gal}(\tilde{\Sigma}/\Sigma)$ while $H = \text{Gal}(\tilde{\Sigma}/\Sigma')$. Let θ_i for $i = 1, \ldots, p$ denote the complete list of irreducible representations of G over $\mathbb{R}$. For each i, recall that we have the twisted operators D^N_{ψ,θ_i}, as in [Wen21, §2.2 and §3.4] and Definition 3.2.5, whose kernels (and similarly cokernels) are related to $\psi^* D^N_{\Sigma,J}$ by the following formulas

$$\ker D^N_{\psi,\theta_i} = \text{Hom}^G(\theta_i^*, \ker \psi^* D^N_{\Sigma,J}) \tag{3.2.21}$$

$$\ker \psi^* D^N_{\Sigma,J} = \bigoplus_i \theta_i^* \otimes_{\mathbb{K}_i} \ker D^N_{\psi,\theta_i} \tag{3.2.22}$$

where $\mathbb{K}_i = \text{End}^G(\theta_i)$, and (3.2.21), (3.2.22) are isomorphisms of $\mathbb{K}_i$-modules and G-representations respectively. Now, the assumption implies that we can find an $1 \leq j \leq p$ such that $(\theta_j^*)^H \neq 0$ and $\ker D^N_{\psi,\theta_j} \neq 0$. We contend that θ_j must be a faithful G-representation. If not, let $(1) \neq K \lhd G$ be such that θ_j descends to G/K. Note that $K \not\subset H$ (since H has trivial normal core) and thus, $D^N_{\Sigma,J}$ pulled back to $\tilde{\Sigma}/HK$ has a non-trivial kernel (contradicting the minimality of d). Now that we know that θ_i is faithful, we can use [Wen21, Lemma 2.10] to get

$$\text{ind}_{\mathbb{K}_j} D^N_{\psi,\theta_j} \leq -2r \tag{3.2.23}$$

where r is the number of distinct branch points of φ. Moreover, this estimate is strict in case $\mathbb{K}_j = \mathbb{R}$ unless all the ramification points of φ are simple (i.e. they're all locally 2-to-1). Now, set $t_i = \dim_{\mathbb{R}} \mathbb{K}_i$, $k_i = \dim_{\mathbb{K}_i} \ker D^N_{\psi,\theta_i}$ and $c_i = \dim_{\mathbb{K}_i} \text{coker } D^N_{\psi,\theta_i}$ and

note that the codimension of $\mathcal{W}^d_{\mathbf{b},G}(g, A, \mathbf{k}, \mathbf{c})$ near J is

$$\geq -2r + \sum_i t_i k_i c_i \geq -2r + t_j k_j (k_j + 2r) \geq 1 \tag{3.2.24}$$

using Theorem 3.2.12. Note that the shifting $-2r$ comes from the dimension of the moduli space of branched covers. (Here, $\mathbf{b}$ is constructed using φ, $A = [\Sigma]$, $g = \text{genus}(\Sigma)$, $\mathbf{k} = \ker \psi^* D^N_{\Sigma,J}$ and $\mathbf{c} = \text{coker } \psi^* D^N_{\Sigma,J}$.) Moreover, to obtain the equality in the last estimate, we must have $t_j = k_j = 1$, $k_i = 0$ for all $j \neq i$ (this last part follows since $\text{ind}_{\mathbb{K}_i} D^N_{\psi,\theta_i} \leq 0$ by [Wen21, Lemma 2.10]) and $\text{ind}_{\mathbb{K}_j} D^N_{\psi,\theta_j} = -2r$. [Wen21, Lemma 2.10] now implies that all the ramification points of φ must be locally 2-to-1. □

3.3 A necessary condition for bifurcations

In this section, we will describe a necessary condition for a sequence of simple pseudo-holomorphic curves (defined with respect to a sequence of almost complex structures) to converge to a multiple cover of an embedded curve.

We emphasize that the main result of this section (Theorem 3.3.1) is independent of [Wen21] *and holds for almost complex manifolds in any dimension.*

Fix an even dimensional manifold X^{2n} for the rest of this section. Let

$$\mathcal{F} : V \to \mathcal{J}(X) \tag{3.3.1}$$

be a smooth family of almost complex structures on X parametrized by a smooth finite dimensional manifold V (i.e. $\mathcal{F}$ defines a smooth section $\pi_X^* \text{End}(T_X)$ with $\mathcal{F}^2 = -\text{id}$ with $\pi_X : V \times X \to X$ being the projection). Given a smooth projective curve C, denote by $\overline{\mathcal{M}}_h(C, k)$ the moduli stack of genus h stable maps of degree k with target C.

Theorem 3.3.1 (Necessary condition for bifurcation)**.** *Let $x_\nu \to x$ be a convergent sequence in V, with $J_\nu := \mathcal{F}(x_\nu)$ for $\nu \geq 0$ and $J := \mathcal{F}(x)$. Let $h \geq 0$ be an integer and suppose that we have a sequence $(J_\nu, \varphi_\nu : \Sigma'_\nu \to X)$ of simple J_ν-curves of genus h converging, in the Gromov topology, to a stable map*

$$(J, C' \xrightarrow{\varphi} C \subset X) \tag{3.3.2}$$

with C being a smooth embedded J-curve of genus $g \leq h$ and $\varphi : C' \to C$ being an element of $\overline{\mathcal{M}}_h(C, k)$ for some integer $k \geq 1$. Then, exactly one of the following must be true.

(i) *We have $g = h$, $k = 1$ and $\varphi : C' \to C$ is an isomorphism.*

(ii) *The natural pullback map*

$$\varphi^* : \ker D^N_{C,J} \to \ker \varphi^* D^N_{C,J} \tag{3.3.3}$$

is injective but not surjective.

The alternative (ii) provides a necessary condition for bifurcations of the moduli spaces of embedded pseudo-holomorphic curves to occur. Before proceeding to the proof, we make a few remarks comparing this to similar statements in the literature and pointing out some natural extensions of this result.

Remark 3.3.2. (a) If we further suppose that there exists a sequence C_ν of embedded J_ν-curves of genus g converging to C (this is ensured, for example, if $D^N_{C,J}$ is surjective), then the rescaling argument used to prove [Wen21, Proposition B.1] yields the weaker conclusion $\ker \varphi^* D^N_{C,J} \neq 0$.

(b) Theorem 3.3.1 rules out the possibility of a sequence of simple curves of genus h converging in the Gromov topology to an embedded curve of genus $g < h$ with

some ghost components attached. In particular, if $A \in H_2(X, \mathbb{Z})$ is a primitive homology class in a 6-dimensional symplectic manifold (X, ω) with $c_1(A) = 0$, then for any path γ in $\mathcal{J}_{\mathrm{emb}}(X, \omega)$ (in the sense of [DW19a, Definition 2.36]), the moduli space $\mathcal{M}_h^{\mathrm{emb}}(X, \gamma, A)$ of embedded genus h curves of class A is compact. This can be used to recover [DW19a, Theorem 1.5(1)].

(c) If $\iota : C \to X$ is just a simple J-curve and not necessarily an embedding, then the 2-term complex

$$N_\iota^\bullet = [T_C \xrightarrow{d\iota} \iota^* T_X] \tag{3.3.4}$$

of vector bundles concentrated in degrees -1 and 0 (equipped with their respective Cauchy–Riemann operators) serves as a replacement for the normal bundle N_C equipped with the operator $D^N_{C,J}$. It should be possible to extend our proof of Theorem 3.3.1 to the case when ι is simple provided we replace $D^N_{C,J}$ by $N_\iota^\bullet$ in the conclusion.

(d) The restriction in Theorem 3.3.1 to finite dimensional families of almost complex structures suffices for all applications we are aware of but it should be possible to extend our proof to the infinite dimensional case as well. We would have to replace the appeal to Brouwer's Invariance of Domain (as in §3.3.5) by a more careful analysis of the differential of the inclusion from Lemma 3.3.7. To deal with the lack of differentiability of the moduli space near curves with nodal domains, one would have to use suitable relative tangent bundles associated to canonical "rel–C^∞ structures" on the moduli spaces as in [Swa21].

The remainder of this section is devoted to the proof of Theorem 3.3.1. The main idea of the argument is given in §3.3.1 and involves counting dimensions of moduli spaces. However, there are some issues related to transversality which we need to overcome before this outline can be turned into a rigorous proof. In §3.3.2–§3.3.4, we

show how to tackle these issues and complete the proof in §3.3.5.

3.3.1 Outline of the argument

Set $A = [C] \in H_2(X, \mathbb{Z})$ and let

$$\mathcal{M}_g^{\text{emb}}(X, \mathcal{F}, A) \tag{3.3.5}$$

be the moduli space of pairs (y, Σ), where $y \in V$ and $\Sigma \subset X$ is an embedded $\mathcal{F}(y)$-curve of genus g in class A. Also consider the moduli space

$$\overline{\mathcal{M}}_h(X, \mathcal{F}, kA) \tag{3.3.6}$$

consisting of pairs $(y, \varphi' : \Sigma' \to X)$, where $y \in V$ and φ' is a stable $\mathcal{F}(y)$-holomorphic map in X of genus h and class $k[C]$ and the moduli space

$$\overline{\mathcal{M}}_h(\mathcal{M}_g^{\text{emb}}(X, \mathcal{F}, A), k) \tag{3.3.7}$$

consisting of triples $(y, \Sigma, \psi : \Sigma' \to \Sigma)$, where (y, Σ) lies in (3.3.5) and (Σ', ψ) lies in $\overline{\mathcal{M}}_h(\Sigma, k)$. There is a natural inclusion of (3.3.7) inside (3.3.6), given by

$$(y, \Sigma, \psi : \Sigma' \to \Sigma) \mapsto (y, \Sigma' \xrightarrow{\psi} \Sigma \subset X), \tag{3.3.8}$$

and the proof of Theorem 3.3.1 will follow by showing that codimension of (3.3.7) inside (3.3.6) at the point $(x, C' \xrightarrow{\varphi} C \subset X)$ is given by the dimension of the cokernel of (3.3.3). However, both (3.3.6) and (3.3.7) are not smooth in general therefore the notion of codimension is not well-defined. To remedy the situation, we will need to thicken the moduli spaces (in the sense of [Par16, Definition 9.2.3]) compatibly so that (3.3.8) becomes an inclusion of manifolds, at least near the point (x, φ), whose

codimension can then be computed.

3.3.2 Preparation

Lemma 3.3.3 (Transversality for embedded curves)**.** *We can find a smooth manifold V' containing V as a submanifold and a smooth family $\mathcal{F}' : V' \to \mathcal{J}(X)$ extending $\mathcal{F}$ such that the moduli space*

$$\mathcal{M}_g^{\mathrm{emb}}(X, \mathcal{F}', A) \tag{3.3.9}$$

is a smooth manifold of the expected dimension $= \dim V' + \operatorname{ind} D^N_{C,J}$ *near* (x, C).

Proof. This is a variant of the argument in [MS12, Proposition 3.2.1]. Indeed, by the implicit function theorem, it suffices to arrange that the linearized map

$$T_x V' \to \operatorname{coker} D^N_{C,J} \tag{3.3.10}$$

$$v' \mapsto \tfrac{1}{2}(d\mathcal{F}'|_x \cdot v') \circ d\iota \circ j \tag{3.3.11}$$

is surjective, where ι denotes the inclusion $C \subset X$ and j denotes the restriction of J to T_C. The right side of (3.3.11) is an element of $\Omega^{0,1}(C, T_X|_C)$ which we are projecting to $\Omega^{0,1}(C, N_C)$ and then to coker $D^N_{C,J}$. Note that the restriction map

$$T_J\mathcal{J} = \overline{\operatorname{End}}_J(T_X) \to \overline{\operatorname{Hom}}_{\mathbb{C}}(T_C, N_C) = \Omega^{0,1}(C, N_C) \tag{3.3.12}$$

is surjective and thus, we may lift a (finite) basis of coker $D^N_{C,J}$ to $T_J\mathcal{J}$. It is now easy to define $V \subset V'$ and $\mathcal{F}'$ extending $\mathcal{F}$ such that (3.3.10) is surjective. □

Before proceeding further, we pick some auxiliary data.

1. Choose a Riemannian metric on X such that we can identify the unit disc bundle in N_C with a tubular neighborhood of $C \subset X$, with fiber over $p \in C$ denoted

by D_p (a codimension 2 submanifold with boundary of X).

2. Choose an integer $m \geq 0$ and finitely many points $p_1, \ldots, p_m \in C$ (which are not the images of critical points of φ or nodal points of C') such that

   (a) $(C, p_1, \ldots, p_m)$ is a stable curve and,

   (b) the evaluation map

   $$\ker D_C^N \to \bigoplus_{i=1}^{m} N_{C,p_i} \tag{3.3.13}$$

   is injective.

   Note that for all (y, Σ) in (3.3.5) close to (x, C), the curve Σ intersects each $D_i := D_{p_i}$ transversally at a unique point (and is disjoint from ∂D_i). The submanifolds $D_1, \ldots, D_m$ will serve as "stabilizing divisors" in the sense of [Par16, Lemma 9.2.8].

3. Choose a complex manifold $\mathcal{M}$ and a complex analytic family $\pi : \mathcal{C} \to \mathcal{M}$ of genus g curves with sections $\sigma_1, \ldots, \sigma_m$ (corresponding to m marked points which stabilize the fibers), equipped with a basepoint $* \in \mathcal{M}$ and an isomorphism

   $$i : (C, p_1, \ldots, p_m) \xrightarrow{\simeq} (\mathcal{C}_*, \sigma_1(*), \ldots, \sigma_m(*)) \tag{3.3.14}$$

   such that the resulting map $\mathcal{M} \to \mathcal{M}_{g,m}$ is étale at $*$ (i.e. provides a local orbifold chart near $*$). Here, $\mathcal{C}_s$ denotes the curve $\pi^{-1}(s)$ for $s \in \mathcal{M}$. The choice of the family $(\pi : \mathcal{C} \to \mathcal{M}, \sigma_1, \ldots, \sigma_m)$ is analogous to choosing a local Teichmüller slice at $(C, p_1, \ldots, p_m)$.

4. For $1 \leq i \leq m$, choose a numbering $\varphi^{-1}(p_i) = \{q_i^1, \ldots, q_i^k\}$. Note that each q_i^j must be a smooth point of C' with $d\varphi(q_i^j) \neq 0$. Moreover, $(C', \{q_i^j\})$ is a stable

curve. Choose a complex manifold $\mathcal{M}'$ and a (flat) complex analytic family $\pi' : \mathcal{C}' \to \mathcal{M}'$ of (possibly nodal) arithmetic genus h curves with sections τ_i^j (corresponding to km marked points which stabilize the fibers), equipped with a basepoint $\bullet \in \mathcal{M}'$ and an isomorphism

$$i' : (C', \{q_i^j\}) \xrightarrow{\simeq} (\mathcal{C}'_\bullet, \{\sigma_i^j(\bullet)\}) \tag{3.3.15}$$

such that the resulting map $\mathcal{M}' \to \overline{\mathcal{M}}_{h,km}$ is étale at $\bullet$ (i.e. provides a local orbifold chart near $\bullet$).

5. Let us denote by

$$\mathcal{C}_g^{\mathrm{emb}}(X, \mathcal{F}', A) \hookrightarrow \mathcal{M}_g^{\mathrm{emb}}(X, \mathcal{F}', A) \times X \tag{3.3.16}$$

the universal embedded curve. As a subset, it is given by the union of the subsets $\{(y, \Sigma)\} \times \Sigma$ over all $(y, \Sigma) \in \mathcal{M}_g^{\mathrm{emb}}(X, \mathcal{F}', A)$. The first projection of (3.3.16) defines a smooth family of curves (i.e. Riemann surfaces) over $\mathcal{M}_g^{\mathrm{emb}}(X, \mathcal{F}', A)$ and thus, there is a smooth classifying map

$$\begin{array}{ccc} \mathcal{C}_g^{\mathrm{emb}}(X, \mathcal{F}', A) & \xrightarrow{\Psi} & \mathcal{C} \\ \downarrow & & \downarrow \\ \mathcal{M}_g^{\mathrm{emb}}(X, \mathcal{F}', A) & \xrightarrow{\Phi} & \mathcal{M} \end{array} \tag{3.3.17}$$

defined near (x, C) which maps $(y, \Sigma) \in \mathcal{M}_g^{\mathrm{emb}}(X, \mathcal{F}', A)$ to the point of $\mathcal{M}$ corresponding to $(\Sigma, \Sigma \cap D_1, \ldots, \Sigma \cap D_m)$ and specializes to the isomorphism (3.3.14) over (x, C).

Remark 3.3.4. All the assertions for the remainder of the argument are to be understood to hold only when we are sufficiently close to the basepoints (x, C), $(x, C, \varphi :$

$C' \to C)$, $(x, C' \xrightarrow{\varphi} C \subset X)$, $*$ and $\bullet$. The size of these neighborhoods may shrink from one line to the next.

Lemma 3.3.5 (Auxiliary embedding)**.** *The map*

$$\mathcal{M}_g^{\mathrm{emb}}(X, \mathcal{F}', A) \to V' \times D_1 \times \cdots \times D_m \tag{3.3.18}$$

defined in a neighborhood of (x, C) *by*

$$(y, \Sigma) \mapsto (y, \Sigma \cap D_1, \cdots, \Sigma \cap D_m) \tag{3.3.19}$$

is a smooth embedding near (x, C)*.*

Proof. By the inverse function theorem, it will suffice to check that the differential

$$T_{(x,C)}\mathcal{M}_g^{\mathrm{emb}}(X, \mathcal{F}', A) \to T_x V' \oplus \bigoplus_{i=1}^{m} N_{C,p_i} \tag{3.3.20}$$

of (3.3.18) at (x, C) is injective. This follows from the exact sequence

$$0 \to \ker D_{C,J}^N \to T_{(x,C)}\mathcal{M}_g^{\mathrm{emb}}(X, \mathcal{F}', A) \to T_x V' \tag{3.3.21}$$

and the fact that (3.3.13) is injective. □

3.3.3 Thickened moduli space of covers

As an intermediate step to tackling (3.3.7), we first study the moduli space

$$\overline{\mathcal{M}}_h(\mathcal{M}, k) \tag{3.3.22}$$

which consists of pairs $(s, \psi : \Sigma' \to \mathcal{C}_s)$, where $s \in \mathcal{M}$ and (Σ', ψ) is in $\overline{\mathcal{M}}_h(\mathcal{C}_s, k)$. We immediately note that

$$\overline{\mathcal{M}}_h(\mathcal{M}_g^{\mathrm{emb}}(X, \mathcal{F}', A), k) = \mathcal{M}_g^{\mathrm{emb}}(X, \mathcal{F}', A) \times_{\mathcal{M}} \overline{\mathcal{M}}_h(\mathcal{M}, k) \tag{3.3.23}$$

where $\overline{\mathcal{M}}_h(\mathcal{M}, k) \to \mathcal{M}$ is the natural projection while $\mathcal{M}_g^{\mathrm{emb}}(X, \mathcal{F}', A) \to \mathcal{M}$ is the classifying map from (3.3.17). Now, let

$$D\bar{\partial}|_{\varphi} : C^\infty(C', \varphi^* T_C) \to C^\infty(\tilde{C}', \Omega^{0,1}_{\tilde{C}'} \otimes_{\mathbb{C}} \varphi^* T_C) \tag{3.3.24}$$

be the linearized Cauchy–Riemann operator associated to the holomorphic map $\varphi : C' \to C$, where $\tilde{C}'$ is the normalization of the curve C'. Choose E to be a finite dimensional vector space and

$$\lambda : E \to C^\infty(\mathcal{C}'^{\circ} \times \mathcal{C}, \Omega^{0,1}_{\mathcal{C}'^{\circ}/\mathcal{M}'} \boxtimes_{\mathbb{C}} T_{\mathcal{C}/\mathcal{M}}) \tag{3.3.25}$$

to be a linear map (where $\mathcal{C}'^{\circ} \subset \mathcal{C}'$ is the complement of the nodal points of the fibers of π') with support proper over $\mathcal{M}' \times \mathcal{M}$ such that its restriction along the map $C' \hookrightarrow \mathcal{C}' \times \mathcal{C}$ given by $i' \times (i \circ \varphi \circ i'^{-1})$ yields an isomorphism

$$E = \operatorname{coker} D\bar{\partial}|_{\varphi}. \tag{3.3.26}$$

Now introduce the moduli space

$$\overline{\mathcal{M}}_h(\mathcal{M}, k)_\lambda \tag{3.3.27}$$

consisting of tuples $(s \in \mathcal{M}, s' \in \mathcal{M}', \psi : \mathcal{C}'_{s'} \to \mathcal{C}_s, e \in E)$ satisfying the λ-thickened Cauchy–Riemann equation (in the sense of [Par16, Definition 9.2.3]) and incidence

conditions

$$\bar{\partial}\psi + \lambda(e)(\cdot, \psi(\cdot)) = 0 \tag{3.3.28}$$

$$\psi(\tau_i^j(s')) = \sigma_i(s) \quad \text{transverse} \tag{3.3.29}$$

for each $1 \leq i \leq m$ and $1 \leq j \leq k$. The word *transverse* in (3.3.29) indicates that $d\psi(\tau_i^j(s')) \neq 0$. It now follows from [Par16, Proposition 9.2.6 & Appendix B] that (3.3.27), near the point $(*, \bullet, i \circ \varphi \circ i'^{-1}, 0)$, is a topological manifold of the expected dimension with the projection to $\mathcal{M}$ being a topological submersion. More explicitly, the dimension of (3.3.27) relative to $\mathcal{M}$ is $=$ vdim $\overline{\mathcal{M}}_h(C,k) + \dim E$ where we have set

$$\text{vdim } \overline{\mathcal{M}}_h(C,k) = 2r \tag{3.3.30}$$

with $r = (2h-2) - k(2g-2)$ being the expected complex dimension of $\overline{\mathcal{M}}_h(C,k)$.

We can now define the λ-thickened moduli space

$$\overline{\mathcal{M}}_h(\mathcal{M}_g^{\text{emb}}(X, \mathcal{F}', A), k)_\lambda = \mathcal{M}_g^{\text{emb}}(X, \mathcal{F}', A) \times_{\mathcal{M}} \overline{\mathcal{M}}_h(\mathcal{M}, k)_\lambda \tag{3.3.31}$$

which is also a topological manifold, near the point $(x, C, \varphi : C' \to C, 0)$. Moreover, its natural projection to $\mathcal{M}_g^{\text{emb}}(X, \mathcal{F}', A)$ is a topological submersion of relative dimension $2r + \dim E$ near this point. The loci in (3.3.27) and (3.3.31) where the E-coordinate vanishes respectively form local orbispace charts for (3.3.22) and (3.3.7).

Pulling back (3.3.25) along (3.3.17), we obtain a thickening datum

$$E \to C^\infty(\mathcal{C}'^\circ \times \mathcal{C}_g^{\text{emb}}(X, \mathcal{F}', A), \Omega^{0,1}_{\mathcal{C}'^\circ/\mathcal{M}'} \otimes_{\mathbb{C}} T_{\mathcal{C}_g^{\text{emb}}(X,\mathcal{F}',A)/\mathcal{M}_g^{\text{emb}}(X,\mathcal{F}',A)}) \tag{3.3.32}$$

which we continue to denote by the same symbol λ.

3.3.4 Thickened moduli space of stable maps

We will now transfer the thickening datum (3.3.32) to a suitably compatible thickening datum

$$\mu : E \to C^\infty(\mathcal{C}'^\circ \times X \times V' \times D_1 \times \cdots \times D_m, \Omega^{0,1}_{\mathcal{C}'^\circ/\mathcal{M}'} \boxtimes_{\mathbb{C}} T_{\mathcal{X}'/V'}). \tag{3.3.33}$$

We explain the notation in (3.3.33). Firstly, $T_{\mathcal{X}'/V'}$ denotes the relative tangent bundle of the family $\mathcal{X}' = X \times V' \to V'$ of almost complex manifolds defined by $\mathcal{F}'$ (i.e. it is the vertical tangent bundle of $X \times V' \to V'$ with fiberwise almost complex structures specified by the family $\mathcal{F}'$). Secondly, the bundle $\Omega^{0,1}_{\mathcal{C}'^\circ/\mathcal{M}'} \boxtimes_{\mathbb{C}} T_{\mathcal{X}'/V'}$ is defined on $\mathcal{C}'^\circ \times (X \times V')$ and the right side of (3.3.33) is to be interpreted as the space of sections of this bundle pulled back along the coordinate projection $\mathcal{C}'^\circ \times X \times V' \times D_1 \times \cdots \times D_m \to \mathcal{C}'^\circ \times X \times V'$.

Combining the embeddings (3.3.16) and (3.3.18), we get compatible embeddings

$$\begin{array}{ccc} \mathcal{C}^{\mathrm{emb}}_g(X, \mathcal{F}', A) & \longrightarrow & X \times V' \times D_1 \times \cdots \times D_m \\ \downarrow & & \downarrow \\ \mathcal{M}^{\mathrm{emb}}_g(X, \mathcal{F}', A) & \longrightarrow & V' \times D_1 \times \cdots \times D_m \end{array} \tag{3.3.34}$$

which are fiberwise pseudo-holomorphic, i.e., their vertical differential induces fiberwise injective $\mathbb{C}$-linear map from $T_{\mathcal{C}^{\mathrm{emb}}_g(X,\mathcal{F}',A)/\mathcal{M}^{\mathrm{emb}}_g(X,\mathcal{F}',A)}$ to the pullback of $T_{\mathcal{X}'/V'}$. Using this embedding we can regard λ as a section of $\Omega^{0,1}_{\mathcal{C}'^\circ/\mathcal{M}} \boxtimes_{\mathbb{C}} T_{\mathcal{X}'/V'}$ defined over $\mathcal{C}'^\circ \times$ the image of $\mathcal{C}^{\mathrm{emb}}_g(X, \mathcal{F}', A)$ in $V' \times D_1 \times \cdots \times D_m$. Now, by using a partition of unity argument, we may extend (a suitable cutoff of) this vector bundle section to the whole of $\mathcal{C}'^\circ \times X \times V' \times D_1 \times \cdots \times D_m$. This yields a linear map μ as in (3.3.33) which agrees with λ near (x, C). Next, choose E' to be a finite dimensional vector

space and

$$\mu' : E' \to C^\infty(\mathcal{C}'^\circ \times X \times V', \Omega^{0,1}_{\mathcal{C}'^\circ/\mathcal{M}'} \boxtimes_{\mathbb{C}} T_{\mathcal{X}'/V'}) \tag{3.3.35}$$

to be a linear map inducing an isomorphism

$$E' = \frac{\text{coker } \varphi^* D^N_{C,J}}{\text{coker } D^N_{C,J}} \tag{3.3.36}$$

when restricted to $\{x\} \times \text{graph}(\varphi \circ i'^{-1})$.

Now, we may define the $(\mu \oplus \mu')$-thickened moduli space

$$\overline{\mathcal{M}}_h(X, \mathcal{F}', kA)_{\mu,\mu'} \tag{3.3.37}$$

to consist of all tuples $(y \in V', s' \in \mathcal{M}, f : \mathcal{C}'_{s'} \to X, e \in E, e' \in E')$ such that

$$f_*[\mathcal{C}'_{s'}] = kA \tag{3.3.38}$$

$$r_i^j := f(\tau_i^j(s')) \in D_i \quad \text{transverse intersection} \tag{3.3.39}$$

$$\bar\partial_{\mathcal{F}(y)} f + \mu(e)(\cdot, f(\cdot), y, r_1^1, \ldots, r_m^1) + \mu'(e')(\cdot, f(\cdot), y) = 0. \tag{3.3.40}$$

The words *transverse intersection* in (3.3.39) indicate that the image of $df(\sigma_i^j(s'))$ and $T_{r_i^j} D_i$ together span $T_{r_i^j} X$. Note that the expected dimension of (3.3.37) is

$$(n-3)(2-2h) + 2c_1(kA) + \dim(E \oplus E') + \dim V' \tag{3.3.41}$$

$$= (n-1)(2-2h) + 2k(c_1(A) + 2g - 2) + 2r + \dim(E \oplus E') + \dim V' \tag{3.3.42}$$

$$= \text{ind } \varphi^* D^N_C + \text{vdim } \overline{\mathcal{M}}_h(C, k) + \dim(E \oplus E') + \dim V' \tag{3.3.43}$$

where r is as in (3.3.30). Moreover, the locus in (3.3.37) where the $(E \oplus E')$-coordinates vanish gives a local orbispace chart near $(x, \bullet, \varphi \circ i'^{-1}, 0, 0)$ for $\overline{\mathcal{M}}_h(X, \mathcal{F}', kA)$.

Lemma 3.3.6. *The moduli space* $\overline{\mathcal{M}}_h(X, \mathcal{F}', kA)_{\mu,\mu'}$ *is a topological manifold of the expected dimension near* $(x, \bullet, \varphi \circ i'^{-1}, 0, 0)$.

Proof. By [Par16, Proposition 9.2.6 & Appendix B], it is enough to check that the relevant linearized operator at $(x, \bullet, \varphi \circ i'^{-1}, 0, 0)$ is surjective. Splitting the linearized Cauchy–Riemann operator $D\bar{\partial}_J|_\varphi$ on $\varphi^* T_X$ along the directions tangent and normal to C, we can argue for surjectivity as follows. The cokernel of the $\varphi^* T_C$ component of $D\bar{\partial}_J|_\varphi$ is annihilated by the linearization of the μ-term using isomorphism (3.3.26).

For the normal part, i.e. $\varphi^* D^N_{C,J}$, we argue as follows. Using the surjectivity of (3.3.10), we see that coker $D^N_{C,J}$ is annihilated by the variation $d\mathcal{F}|_x$ of almost complex structures parametrized by V'. The remaining part of coker $\varphi^* D^N_{C,J}$ is annihilated by the linearization of μ' using the isomorphism (3.3.36). □

Lemma 3.3.7 (Compatibility of thickenings)**.** *The assignment*

$$\begin{aligned} &\overline{\mathcal{M}}_h(\mathcal{M}_g^{emb}(X, \mathcal{F}', A), k)_\lambda \to \overline{\mathcal{M}}_h(X, \mathcal{F}', kA)_{\mu,\mu'} \\ &((y, \Sigma), (s, s', \psi : \mathcal{C}'_{s'} \to \mathcal{C}_s, e)) \mapsto (y, s', \Psi^{-1}_{(y,\Sigma)} \circ \psi : \mathcal{C}'_{s'} \to \Sigma \subset X, e, 0) \end{aligned} \tag{3.3.44}$$

defines a continuous injective map from a neighborhood of $((x, C), (*, \bullet, i \circ \varphi \circ i'^{-1}, 0))$ *to a neighborhood of* $(x, \bullet, \varphi \circ i'^{-1}, 0, 0)$. *Here* Ψ *is the classifying map from* (3.3.17).

Proof. It will suffice to show that the map is well-defined since continuity and injectivity are evident. So, let us assume that $((y, \Sigma), (s, s', \psi : \mathcal{C}'_{s'} \to \mathcal{C}_s, e))$ is a point in (3.3.31). Setting $\{\hat{p}_i\} = D_i \cap \Sigma$, this means that we have $\Phi(y, \Sigma) = s$, a biholomorphism

$$\Psi_{(y,\Sigma)} : (\Sigma, \hat{p}_1, \ldots, \hat{p}_m) \xrightarrow{\simeq} (\mathcal{C}_s, \sigma_1(s), \ldots, \sigma_m(s)) \tag{3.3.45}$$

of stable curves and that the equations (3.3.28) and (3.3.29) are satisfied. We then need to verify that the tuple $(y, s', \Psi^{-1}_{(y,\Sigma)} \circ \psi : \mathcal{C}'_{s'} \to \Sigma \subset X, e, 0)$ satisfies the equations

(3.3.38)–(3.3.40). If (y, Σ) is sufficiently close to (x, C), then we have $[\Sigma] = [C] = A$ and thus, (3.3.38) is satisfied. (3.3.29) has (3.3.39) as an immediate consequence. To deduce (3.3.40) from (3.3.28), it will suffice to check that

$$\Psi^*_{(y,\Sigma)}\left[\lambda(e)(\cdot, \psi(\cdot))\right] = \mu(e)(\cdot, \Psi^{-1}_{(y,\Sigma)} \circ \psi(\cdot), y, \hat{p}_1, \ldots, \hat{p}_m) \tag{3.3.46}$$

since $\hat{p}_i = \Psi^{-1}_{(y,\Sigma)}(\sigma_i(s)) = \Psi^{-1}_{(y,\Sigma)} \circ \psi(\tau^1_i(s'))$ for $1 \leq i \leq m$. By construction, μ agrees with λ via the embedding (3.3.34) in a neighborhood of (x, C) and (3.3.46) follows from this immediately. □

Corollary 3.3.8. *The map from Lemma 3.3.7 is an inclusion of codimension given by*

$$\dim \ker \varphi^* D^N_{C,J} - \dim \ker D^N_{C,J} \geq 0. \tag{3.3.47}$$

Proof. The dimension of (3.3.31) is given by

$$\dim V' + \text{ind } D^N_{C,J} + \text{vdim } \overline{\mathcal{M}}_h(C, k) + \dim E \tag{3.3.48}$$

while the dimension of (3.3.37) is given by

$$\dim V' + \text{ind } \varphi^* D^N_{C,J} + \text{vdim } \overline{\mathcal{M}}_h(C, k) + \dim E + \dim E' \tag{3.3.49}$$

and thus, from (3.3.36), we find that the codimension of (3.3.31) in (3.3.37) is precisely given by (3.3.47). □

3.3.5 Concluding the proof of Theorem 3.3.1

We will now prove Theorem 3.3.1 by contradiction. To this end, assume that (3.3.47) is zero. Then, the map from Lemma 3.3.7 is an injective continuous map between

topological manifolds of the same dimension and thus, by Brouwer's Invariance of Domain ([Hat02, Theorem 2B.3]), its image must cover an open neighborhood of $(x, \bullet, \varphi \circ i'^{-1}, 0, 0)$. Restricting this over $V \times 0 \times 0 \subset V' \times E \times E'$, we conclude that

$$\overline{\mathcal{M}}_h(\mathcal{M}_g^{\mathrm{emb}}(X, \mathcal{F}, A), k) = \overline{\mathcal{M}}_h(X, \mathcal{F}, kA) \tag{3.3.50}$$

in a neighborhood of $(x, \varphi : C' \to C \subset X)$. This is a contradiction to the existence of the sequence $(J_\nu, \varphi_\nu : \Sigma'_\nu \to X)$ unless $h = g$ and $k = 1$ (in which case φ must also be an isomorphism). □

3.4 Bifurcation analysis

We begin by describing the class of bifurcations that we will treat in this article. Given any $\delta = (g, d, \mathbf{b}, G, \mathbf{k}, \mathbf{c}, A)$ as in Definition 3.2.11, let r, h and $H \subset G$ be the number of branch points, domain genus and the (conjugacy class of) subgroup corresponding to the d-fold covers in question. We remind the reader that, as in §3.2.2, we will be assuming throughout this section that (X, ω) is a symplectic Calabi–Yau 3-fold.

Definition 3.4.1 (Elementary wall type)**.** *We say that δ is an* elementary wall type *if it satisfies conditions* (1)–(2) *below.*

(1) *$d \geq 2$ and either $r = 0$, i.e. unbranched covers, or for each $1 \leq i \leq r$, we have $b_i^1 = 2$ and $b_i^j = 1$ for $1 < j \leq q_i$.*

(2) *The G-representation $\boldsymbol{k}$ is faithful, irreducible and of real type. Moreover, if $r = 0$, then we have $\dim_{\mathbb{R}} \boldsymbol{k}^H \leq 1$.*

In this case, the subset $\mathcal{W}^d_{\boldsymbol{b},G}(g, A, \boldsymbol{k}, \boldsymbol{c}) \subset \mathcal{J}(X, \omega)$ from Definition 3.2.15 is called an elementary wall. *If we also have $\boldsymbol{k}^H = 0$, then we say that the elementary wall type δ is* trivial.

Note that condition (1) in Definition 3.4.1 is just a strengthening of condition (i) in Lemma 3.2.16 while condition (2) of Definition 3.4.1 is a strengthening of condition (ii) of Lemma 3.2.16. Using [Wen21, Appendix B] and Lemma 3.2.8, the non-triviality of δ is a necessary condition for a bifurcation to occur (see Lemma 3.4.9 below for more details).

Example 3.4.2. *Definition 3.4.1 covers many cases, including the following.*

(1) *Any degree 2 morphism $\Sigma' \to \Sigma$ of curves gives an elementary wall type δ, with $G = \mathbb{Z}/2\mathbb{Z}$ and $\boldsymbol{k}$ being the unique one-dimensional non-trivial representation of $\mathbb{Z}/2\mathbb{Z}$. A result analogous to Theorem 3.4.7 below in the case of unbranched double covers of tori in dimension 4 appears in* [Tau96b].

(2) *For any $d \geq 2$, consider a degree d morphism $\Sigma' \to \Sigma$ with distinct simple branch points with generalized automorphism group being the whole symmetric group S_d. Such covers always exist provided*

$$2h - 2 - d(2g-2) = r \geq 2d - 2 \tag{3.4.1}$$

where $h = \text{genus}(\Sigma')$ and $g = \text{genus}(\Sigma)$ and they give rise to elementary wall types with $G = S_d$ and $\boldsymbol{k}$ being the subrepresentation $\{(x_1, \ldots, x_d) : \sum_i x_i = 0\}$ of the tautological representation of S_d on $\mathbb{R}^d$. The subgroup $H \subset G$ in this case is given by S_{d-1}.

(3) *For any $n \geq 1$, take $d = 2n+1$ and consider a degree d unbranched cover $\Sigma' \to \Sigma$ of a genus $g \geq 2$ curve given by the permutation representation*

$$\rho : \pi_1(\Sigma_g) = \langle a_1, b_1, \ldots, a_g, b_g : [a_1, b_1] \cdots [a_g, b_g] = 1 \rangle \to S_{2n+1} \tag{3.4.2}$$

$$\rho(a_1) = \rho(b_1) = (1\,2)(3\,4)\cdots(2n-1\,2n) \tag{3.4.3}$$

$$\rho(a_2) = \rho(b_2) = (2\,3)(4\,5)\cdots(2n\,2n+1) \tag{3.4.4}$$

and $\rho(a_i) = \rho(b_i) = 1$ *for* $i > 2$. *Note that the generalized automorphism group* $\operatorname{im}\rho = G \subset S_{2n+1}$ *of this cover is isomorphic to the dihedral group*

$$D_{2n+1} = \langle \tau, \theta : \tau^2 = \theta^{2n+1} = 1, \tau\theta\tau^{-1} = \theta^{-1} \rangle \tag{3.4.5}$$

with $\tau = (1\,2)(3\,4)\cdots(2n-1\,2n)$ *and* $\tau\theta = (2\,3)(4\,5)\cdots(2n\,2n+1)$. *The subgroup* H *is given by* $\langle\tau\rangle \subset D_{2n+1}$. *The standard* 2*-dimensional representation* $\boldsymbol{k}$ *which realizes* D_{2n+1} *as the isometry group of a regular* $(2n+1)$*-gon in* $\mathbb{R}^2$ *gives rise to a non-trivial elementary wall type.*

(4) *Let* $n \geq 2$ *be an integer and let* $d = 2n$. *For integers* $g, s \geq 1$, *define a degree* d *branched cover* $\Sigma' \to \Sigma$ *of a genus* g *curve with* $2s$ *distinct simple branch points* $\{q_1, \ldots, q_{2s}\} \subset \Sigma$ *by the permutation representation*

$$\rho : \pi_1(\Sigma_g \setminus \{q_1, \ldots, q_{2s}\}) \tag{3.4.6}$$

$$= \langle a_1, b_1, \ldots, a_g, b_g, c_1, \ldots, c_{2s} : [a_1, b_1]\cdots[a_g, b_g]c_1\cdots c_{2s} = 1\rangle \to S_d \tag{3.4.7}$$

$$\rho(a_i) = \rho(b_i) = (1\,2\,\cdots\,n)(n+1\,\cdots\,2n), \quad \rho(c_j) = (1\ n+1) \tag{3.4.8}$$

for all $1 \leq i \leq g$ *and* $1 \leq j \leq 2s$. *Note that the generalized automorphism group* $\operatorname{im}\rho = G \subset S_d$ *is isomorphic to* $(\mathbb{Z}/2\mathbb{Z})^n \rtimes \mathbb{Z}/n\mathbb{Z}$ *via*

$$((0, 0, \ldots, 0), 1) \mapsto (1\,2\,\cdots\,n)(n+1\,\cdots\,2n) \tag{3.4.9}$$

$$((1, 0, \ldots, 0), 0) \mapsto (1\ n+1) \tag{3.4.10}$$

and the subgroup $H \subset G$ *is given by* $(\mathbb{Z}/2\mathbb{Z})^{n-1} \times \{0\}$. *The group* G *has a representation* $\boldsymbol{k} = \mathbb{R}^n$, *where* $\mathbb{Z}/n\mathbb{Z}$ *acts by cyclically permuting the coordinates while* $(\alpha_1, \ldots, \alpha_n) \in (\mathbb{Z}/2\mathbb{Z})^n$ *acts by the matrix* $\operatorname{diag}[(-1)^{\alpha_1}, \ldots, (-1)^{\alpha_n}]$, *and this gives rise to a non-trivial elementary wall type.*

Let $J_+, J_- \in \mathcal{J}(X,\omega)$ be super-rigid almost complex structures (which will remain fixed for the rest of this section). Define $\mathcal{P}(J_-, J_+)$ to be the space of smooth paths $\gamma : [-1,1] \to \mathcal{J}(X,\omega)$ with $\gamma(\pm 1) = J_\pm$. Before giving the statement of our main result on bifurcations, we need some preliminaries which we now turn to.

Lemma 3.4.3 (Generic transversality to wall strata)**.** *There is a Baire subset*

$$\mathcal{B}' \subset \mathcal{P}(J_-, J_+) \tag{3.4.11}$$

such that any $\gamma \in \mathcal{B}'$ has the following property. The image $\gamma([-1,1]) \subset \mathcal{J}(X,\omega)$ avoids $\mathcal{W}_{\mathrm{emb}}$ and γ is transverse to each of the maps (3.2.17)*. In particular, whenever the Fredholm index of* (3.2.17) *is ≤ -2, the image $\gamma([-1,1])$ avoids the image $\mathcal{W}^d_{\boldsymbol{b},G}(g, A, \boldsymbol{k}, \boldsymbol{c})$ of* (3.2.17)*.*

Proof. Note that $J_\pm$ are not contained in the wall $\mathcal{W}$ (see Definition 3.2.15). Applying the Sard–Smale theorem (countably many times) to suitable universal moduli spaces over $\mathcal{P}(J_-, J_+)$, we obtain a Baire subset $\mathcal{B}' \subset \mathcal{P}(J_-, J_+)$ consisting of paths γ which are transverse to $\mathcal{W}_{\mathrm{emb}}$ as well as to each of the maps (3.2.17). Since $\mathcal{W}_{\mathrm{emb}} \subset \mathcal{J}(X,\omega)$ is a subset of codimension ≥ 2 (as are those $\mathcal{W}^d_{\mathbf{b},G}(g, A, \mathbf{k}, \mathbf{c})$ for which (3.2.17) has Fredholm index ≤ -2), transversality implies that $\gamma([-1,1])$ avoids all of these. □

Let $\gamma \in \mathcal{B}'$ be a path from J_- to J_+. Suppose that we have $t_0 \in (-1,1)$ and $p = (\gamma(t_0), \Sigma, \varphi : \Sigma' \to \Sigma, \iota) \in \mathcal{M}^d_{\mathbf{b},G}(\mathcal{M}^*_g(X,A); \mathbf{k}, \mathbf{c})$. Let $\psi : \tilde{\Sigma} \to \Sigma$ be a Galois closure of φ with $H \subset G$ being such that $\Sigma' = \tilde{\Sigma}/H$. Introduce the abbreviation

$$D = D^N_{\Sigma, \gamma(t_0)}. \tag{3.4.12}$$

Lemma 3.4.3 guarantees that γ and (3.2.17) are transverse at (t_0, p), which yields a

canonical isomorphism

$$T_{t_0}\gamma \xrightarrow{\simeq} \mathrm{Hom}(\ker \varphi^* D, \mathrm{coker}\ \varphi^* D)/T_\varphi \mathcal{M}^d_{\mathbf{b},G}(\Sigma) \tag{3.4.13}$$

between the tangent space of γ at t_0 and the normal space of $\mathcal{W}^d_{\mathbf{b},G}(g, A, \mathbf{k}, \mathbf{c}) \subset \mathcal{J}(X, \omega)$ at $\gamma(t_0)$, where we have used Lemmas 3.2.8 and 3.2.16 to make the natural identification

$$\mathrm{Hom}^G(\ker \psi^* D, \mathrm{coker}\ \psi^* D) = \mathrm{Hom}(\ker \varphi^* D, \mathrm{coker}\ \varphi^* D) \tag{3.4.14}$$

given by restriction to the H-fixed subspace. Orient the two sides of (3.4.13) by identifying $T_{t_0}\gamma$ with $\mathbb{R}$ in the obvious way, using the complex structure on the orbifold $\mathcal{M}^d_{\mathbf{b},G}(\Sigma)$ and using the standard orientation on the determinant line $\det(\varphi^* D)$, which is obtained via the spectral flow when we deform $\varphi^* D$ to a $\mathbb{C}$-linear Cauchy–Riemann operator.

Definition 3.4.4 (Sign of wall crossing)**.** *With notation as above, we call the pair* (t_0, p) *a* wall crossing *along* γ*. The* sign *of this wall crossing is defined to be the sign of the isomorphism* (3.4.13) *and is denoted by*

$$\mathrm{sgn}(\gamma, t_0; p) \in \{-1, +1\}. \tag{3.4.15}$$

We call the wall crossing positive *or* negative *in accordance with its sign. When* $\delta = (g, d, \boldsymbol{b}, G, \boldsymbol{k}, \boldsymbol{c}, A)$ *is an elementary wall type, we say that the wall crossing is* elementary*.*

If $\delta = (g, d, \mathbf{b}, G, \mathbf{k}, \mathbf{c}, A)$ is an elementary wall type for which (3.2.17) has Fredholm index ≤ -1, then we have a natural codimension ≥ 1 subspace

$$\mathcal{M}^d_{\mathbf{b},G}(\mathcal{M}^*_g(X, A); \mathbf{k}, \mathbf{c})_\spadesuit \subset \mathcal{M}^d_{\mathbf{b},G}(\mathcal{M}^*_g(X, A); \mathbf{k}, \mathbf{c}) \tag{3.4.16}$$

consisting of points which are "degenerate". (See Definition 3.4.21 and Lemma 3.4.24 for details.) Denote its image under (3.2.17) by $\mathcal{W}^d_{\mathbf{b},G}(g,A,\mathbf{k},\mathbf{c})_\spadesuit$ and note that it is a subset of codimension ≥ 2 in $\mathcal{J}(X,\omega)$.

Lemma 3.4.5 (Generic non-degeneracy)**.** *There is a Baire subset*

$$\mathcal{B}'' \subset \mathcal{P}(J_-,J_+) \tag{3.4.17}$$

with the property that every $\gamma \in \mathcal{B}''$ has image disjoint from $\mathcal{W}^d_{\boldsymbol{b},G}(g,A,\boldsymbol{k},\boldsymbol{c})_\spadesuit$ for each elementary wall type $\delta = (g,d,\boldsymbol{b},G,\boldsymbol{k},\boldsymbol{c},A)$ for which the map (3.2.17) *has Fredholm index ≤ -1.*

Proof. This follows by an application of the Sard–Smale theorem (countably many times) to suitable universal moduli spaces over $\mathcal{P}(J_-,J_+)$. The key observation is that the subsets $\mathcal{W}^d_{\mathbf{b},G}(g,A,\mathbf{k},\mathbf{c})_\spadesuit$ have codimension ≥ 2 in $\mathcal{J}(X,\omega)$. □

Remark 3.4.6. The subsets $\mathcal{B}'$ and $\mathcal{B}''$ can both be chosen to be invariant under smooth reparametrizations of paths.

Theorem 3.4.7 (Elementary bifurcations)**.** *Let $\gamma \in \mathcal{P}(J_-,J_+)$ be a path lying in the Baire subset $\mathcal{B}' \cap \mathcal{B}''$. Consider a point*

$$p = (J,\Sigma,\varphi:\Sigma'\to\Sigma,\iota) \in \mathcal{M}^d_{\boldsymbol{b},G}(\mathcal{M}^*_g(X,A);\boldsymbol{k},\boldsymbol{c}) \tag{3.4.18}$$

such that $(0,p)$ is an elementary wall crossing.

We can then find an open neighborhood $\mathcal{O}$ of $\varphi:\Sigma'\to X$ in the Gromov topology and an $\epsilon>0$ such that the parametric moduli space

$$((-\epsilon,\epsilon)\times\mathcal{O}) \cap \mathcal{M}^*_h(X,\gamma,dA) \to [-1,1] \tag{3.4.19}$$

of smooth embedded *$\gamma(t)$-curves, with $|t|<\epsilon$, of genus h and class dA, lying in $\mathcal{O}$*

has the following properties. The fiber of (3.4.19) *over* $t = 0$ *is empty and* (3.4.19) *is a smooth oriented* 1*-dimensional manifold, with the projection to* $(-\epsilon, 0) \cup (0, \epsilon)$ *being a proper submersion. Let* $m_\pm$ *and* $n_\pm$ *denote the unsigned and signed number, respectively, of points in the fiber of* (3.4.19) *over* t *for* $\pm t > 0$.

(i) *We have* $m_\pm = 0$ *if* δ *is trivial, or if* $r = 0$ *and* $H \subset G$ *is not a maximal subgroup such that* $\boldsymbol{k}^H \neq 0$.

(ii) *If we're not in either of the cases from* (i)*, then we have* $\Gamma = \mathrm{Aut}(\varphi) \subset \mathbb{Z}/2\mathbb{Z}$*, the estimate* $|\Gamma| \cdot |m_\pm| \leq 2$ *and the identity*

$$n_+ - n_- = -\frac{2 \cdot \mathrm{sgn}(\gamma, 0; p)}{|\Gamma|}. \tag{3.4.20}$$

Moreover, if $(t_i, \Sigma_i' \subset X)$ *is any sequence of points in* (3.4.19) *with* $t_i \to 0$*, then the Gromov limit of this sequence exists and is given by* $\varphi : \Sigma' \to X$.

Remark 3.4.8. Note that for $J = \gamma(0)$, an element $\sigma \in \ker(\varphi^* D^N_{\Sigma,J})$ can be viewed as a d-valued holomorphic section of the bundle N_Σ over the curve Σ with singular points given by the branched points of the cover $\varphi : \Sigma' \to \Sigma$. Theorem 3.4.7 asserts that by varying J along a generic path γ, we can obtain a 1-parameter family of solutions to the non-linear Cauchy–Riemann equations in X which give rise to embedded $\gamma(t)$-holomorphic curves of class $d[\Sigma]$. We expect the proof of Theorem 3.4.7 (which is carried out in the next section via Kuranishi models) to be useful in other similar situations. For instance, it is speculated in [Don19, Section 1] that a 2-valued harmonic 1-form defined over a special Lagrangian submanifold M in a Calabi–Yau manifold (X, ω, Ω) should give rise to a 1-parameter family of special Lagrangian immersions $\iota_t : \tilde{M} \to X$, where $\tilde{M}$ is the branched double cover of M branching along the singular set of the given harmonic 1-form. The similarity is quite transparent, though the analysis needed to carry out the program in [Don19] is still under development.

The rest of this section is devoted to the proof of Theorem 3.4.7. To orient the reader, we start with a brief plan of the argument. In §3.4.1, we deal with the trivial case of the theorem and focus our attention on the non-trivial case for the remainder. In §3.4.2 and §3.4.3, we use the implicit function theorem to construct a Kuranishi model for the moduli space of interest. The detailed analysis of the Kuranishi model, which is the heart of the proof, is carried out in §3.4.4. The proof is completed in §3.4.4 and §3.4.4 (modulo a crucial technical result whose proof is carried out in §3.4.4).

Throughout the argument, we will be freely using the notation from the statement of Theorem 3.4.7. In particular, the super-rigid almost complex structures $J_\pm$, the elementary wall type $\delta = (g, d, \mathbf{b}, G, \mathbf{k}, \mathbf{c}, A)$, the path $\gamma \in \mathcal{B}' \cap \mathcal{B}'' \subset \mathcal{P}(J_-, J_+)$ and the point $p = (J, \Sigma, \varphi : \Sigma' \to \Sigma, \iota)$ with $J = \gamma(0)$ will remain fixed for the remainder of this section. We will be abbreviating $D^N_{\Sigma,J}$ to D and will sometimes denote the almost complex structure $\gamma(t)$ by the symbol J_t.

3.4.1 Simplifications

Let r, h and $H \subset G$ be the number of branch points, domain genus and the (conjugacy class of) subgroup corresponding to the d-fold cover φ (as in Definition 3.4.1). We will study what happens to the moduli space

$$\mathcal{M}^*_h(X, \gamma, dA) \to [-1, 1] \tag{3.4.21}$$

of embedded genus h curves of class dA (lying in a neighborhood of $\varphi : \Sigma' \to X$ in the Gromov topology) as we go from $t < 0$ to $t > 0$.

Write $\mathbf{k} = \theta^*$ for some irreducible faithful representation θ of G of real type. At the point $p = (J, \Sigma, \varphi : \Sigma' \to \Sigma, \iota) \in \mathcal{M}^d_{\mathbf{b},G}(\mathcal{M}^*_g(X, A); \mathbf{k}, \mathbf{c})$, with associated Galois

closure $\psi : \tilde{\Sigma} \to \Sigma$, we have $\dim_{\mathbb{R}} \ker D^N_{\psi,\theta} = 1$ and

$$\ker \psi^* D = \theta^* \otimes \ker D^N_{\psi,\theta}. \tag{3.4.22}$$

Since $d = [G : H] \geq 2$, it follows that G is non-trivial (note that we study bifurcation analysis only in the multiply covered case since [IP18, Section 6] already treats the corresponding analysis in the somewhere injective case). Taking G-invariants in (3.4.22) implies that $\ker D = 0$. There is therefore a unique embedded $\gamma(t)$-curve Σ_t close to Σ for each $|t|$ small enough (obtained using the inverse function theorem).

Lemma 3.4.9 (No bifurcations for trivial δ)**.** *If δ is a trivial elementary wall type, then we can find $\epsilon > 0$ and a neighborhood $\mathcal{O}$ of $\varphi : \Sigma' \to X$ in the Gromov topology such that* (3.4.21) *has no points contained in $(-\epsilon, \epsilon) \times \mathcal{O}$.*

Proof. Assume to the contrary that we can find a sequence $(t_i, \Sigma'_i \subset X)$ of embedded $\gamma(t_i)$-curves of genus h and class dA converging in the Gromov topology to $(0, \varphi : \Sigma' \to X)$. Now, exactly as in the proof of [Wen21, Proposition B.1], we may suitably rescale Σ'_i in the normal direction to the embedded $\gamma(t_i)$-curve Σ_{t_i} of genus g and class A to obtain a nonzero element in $\ker \varphi^* D = (\ker \psi^* D)^H$ in the Gromov limit as $i \to \infty$. Now, (3.4.22) shows that $\mathbf{k}^H \neq 0$, which implies that δ is non-trivial. □

Lemma 3.4.9 completes the proof of Theorem 3.4.7 in the case when δ is trivial. Thus, for the rest of this section, we shall focus exclusively on the case when δ is non-trivial, i.e., $(\theta^*)^H \neq 0$. We record below an important consequence of δ's non-triviality.

Lemma 3.4.10. *The non-triviality of δ implies that $p = (J, \Sigma, \varphi : \Sigma' \to \Sigma, \iota) \in \mathcal{M}^d_{\boldsymbol{b},G}(\mathcal{M}^*_g(X, A); \boldsymbol{k}, \boldsymbol{c})$ satisfies the following.*

(i) *We have* $\dim_{\mathbb{R}} (\theta^*)^H = 1$.

(ii) *For each irreducible representation ρ of G over $\mathbb{R}$ with $\rho \not\cong \theta$, we have either* $\mathrm{ind}_{\mathbb{R}} D^N_{\psi,\rho} = 0$ *or* $(\rho^*)^H = 0$.

As a consequence, we have the identifications

$$\ker \varphi^* D = (\theta^*)^H \otimes \ker D^N_{\psi,\theta} \tag{3.4.23}$$

$$\text{coker } \varphi^* D = (\theta^*)^H \otimes \text{coker } D^N_{\psi,\theta}. \tag{3.4.24}$$

Proof. We use Lemma 3.2.8 to deduce the quasi-isomorphism

$$\bigoplus_i (\theta_i^*)^H \otimes_{\mathbb{K}_i} D^N_{\psi,\theta_i} \to \varphi^* D, \tag{3.4.25}$$

which implies that we have the identity

$$\mathrm{ind}_{\mathbb{R}}(\varphi^* D) = \sum_i \dim_{\mathbb{R}} (\theta^*)^H \cdot \mathrm{ind}_{\mathbb{K}_i} D^N_{\psi,\theta_i}. \tag{3.4.26}$$

Now, a simple computation using the Riemann–Roch formula (combined with the Riemann–Hurwitz formula for φ and condition (1) of Definition 3.4.1) shows that the left side of (3.4.26) is $-2r$. On the other hand, using [Wen21, Lemma 2.10], we have $\dim_{\mathbb{R}}(\theta^*)^H \cdot \mathrm{ind}_{\mathbb{R}}(D^N_{\psi,\theta}) \leq -2r$ while each summand on the right side of (3.4.26) is ≤ 0. Thus, we must have equality in all these inequalities. If $r > 0$, this immediately gives us (i) and (ii). If $r = 0$, then we get (ii) as above and (i) follows from Definition 3.4.1 and the non-triviality of δ. Finally, (3.4.23) and (3.4.24) are formal consequences of statements (i) and (ii). □

3.4.2 Preparation

Since $\gamma \in \mathcal{B}'$, we know that it is transverse to the map (3.2.17) at $t = 0$ and $p = (J, \Sigma, \varphi : \Sigma' \to \Sigma, \iota)$. Let us clarify the meaning of this transversality statement a little further. Choose $\epsilon > 0$ small enough so that for $|t| < \epsilon$ the inverse function

theorem may be applied to produce an embedded $\gamma(t)$-curve Σ_t (of genus g and class A) close to Σ. Define

$$\mathcal{M}^d_{\mathbf{b},G}(\Sigma,\gamma) \to (-\epsilon,\epsilon) \tag{3.4.27}$$

to be the smooth orbifold consisting of pairs $(t, \varphi'' : \Sigma'' \to \Sigma_t)$ such that $t \in (-\epsilon,\epsilon)$ and φ'' belongs to $\mathcal{M}^d_{\mathbf{b},G}(\Sigma_t)$ as in Definition 3.2.3 (equipped with the obvious projection $(t,\varphi'') \mapsto t$). Being a smooth orbifold, it has a well-defined tangent space at each point (which may be regarded as representations of the isotropy group at that point). It is easy to see that (3.4.27) is a smooth locally trivial bundle over $(-\epsilon,\epsilon)$. Thus, we have a natural short exact sequence

$$0 \to T_\varphi \mathcal{M}^d_{\mathbf{b},G}(\Sigma) \to T_{(0,\varphi)}\mathcal{M}^d_{\mathbf{b},G}(\Sigma,\gamma) \to \mathbb{R}_t \to 0 \tag{3.4.28}$$

relating the tangent spaces of the fiber, the total space and the base. Transversality of γ and (3.2.17) at $(0,p)$ then says that the natural map

$$T_{(0,\varphi)}\mathcal{M}^d_{\mathbf{b},G}(\Sigma,\gamma) \xrightarrow{\simeq} \mathrm{Hom}^G(\ker \psi^* D, \mathrm{coker}\ \psi^* D) \tag{3.4.29}$$

is an isomorphism, where $\psi : \tilde{\Sigma} \to \Sigma$ is the Galois closure of φ. Here (3.4.29) is the linearized map, in the sense of Definition 3.2.9, associated to the family of G-equivariant Cauchy–Riemann operators on $\mathcal{M}^d_{\mathbf{b},G}(\Sigma,\gamma)$ given by

$$p' = (t, \varphi'' : \Sigma'' \to \Sigma_t, \iota_t) \mapsto \psi^*_{p'}(D^N_{\Sigma_t,\gamma(t)}) \tag{3.4.30}$$

where $p' \mapsto (\psi_{p'} : \tilde{\Sigma}'' \to \Sigma'' \to \Sigma_t)$ is a continuous family of Galois closures extending ψ at p.

Next we will study the stable map moduli space $\mathcal{M}_h(X,\gamma,dA)$ locally near $(0,\varphi)$

as an intermediate step to understanding (3.4.19). The strategy is to use the implicit function theorem to first produce a Kuranishi model and then explicitly analyze the resulting Kuranishi map. Before proceeding to this, let us establish the notations for the domain and target of the Kuranishi map (which are respectively the kernel and cokernel of the corresponding linearized operator).

Domain of the Kuranishi model

To describe the domain explicitly, let us denote the inclusion of Σ in X by ι and the linearized Cauchy–Riemann operator on ι^*T_X (resp. φ^*T_X) by $\tilde{D}$ (resp. $\tilde{D}_\varphi = \varphi^*\tilde{D}$). Moreover, write $K = \dot{\gamma}(0)$. The domain of the Kuranishi map is then an open neighborhood of 0 in the vector space

$$T^{\mathrm{Zar}}_{(0,\varphi)}\mathcal{M}_h(X,\gamma,dA) \tag{3.4.31}$$

which consists of "cocycle" triples (ξ, y, t), where $\xi \in \Omega^0(\Sigma', \varphi^*T_X)$, $y \in \Omega^{0,1}(\Sigma', T_{\Sigma'})$ and $t \in \mathbb{R}$ are such that

$$\tilde{D}_\varphi\xi + \tfrac{1}{2}(J \circ d\varphi \circ y + tK \circ d\varphi \circ j_{\Sigma'}) = 0 \tag{3.4.32}$$

modulo the "coboundary" triples $(\frac{1}{2}d\varphi(V), j_{\Sigma'} \cdot \bar{\partial}V, 0)$ where $V \in \Omega^0(\Sigma', T_{\Sigma'})$ and $\bar{\partial}$ stands for the Cauchy-Riemann operator on $\Omega^0(\Sigma', T_{\Sigma'})$. This vector space contains

$$T^{\mathrm{Zar}}_{\varphi}\mathcal{M}_h(X,J,dA) \tag{3.4.33}$$

which is the subspace of $T^{\mathrm{Zar}}_{(0,\varphi)}\mathcal{M}_h(X,\gamma,dA)$ defined by $t = 0$.

Similarly, we can consider the parametrized moduli space of d–fold genus h branched covers of $\{\Sigma_t\}_{|t|<\epsilon}$, denoted by $\mathcal{M}_h(\Sigma,\gamma,d)$ and its (Zariski) tangent space $T_{(0,\varphi)}\mathcal{M}_h(\Sigma,\gamma,d)$ at $(0,\varphi)$ by considering the vector bundle φ^*T_Σ, instead of φ^*T_X, and the correspond-

ing cohomology group. The specialization of the space $\mathcal{M}_h(\Sigma, \gamma, d)$ at $t = 0$ is denoted by $\mathcal{M}_h(\Sigma, d)$ and its tangent space at φ, written as $T_\varphi \mathcal{M}_h(\Sigma, d)$, is constructed accordingly.

These vector spaces fit into the following commutative diagram with exact rows and columns (we recall again that $D = D^N_{\Sigma,J}$).

$$\begin{array}{ccccccccc}
 & & 0 & & 0 & & & & \\
 & & \downarrow & & \downarrow & & & & \\
0 & \longrightarrow & T_\varphi \mathcal{M}_h(\Sigma, d) & \longrightarrow & T_{(0,\varphi)} \mathcal{M}_h(\Sigma, \gamma, d) & \longrightarrow & \mathbb{R}_t & \longrightarrow & 0 \\
 & & \downarrow & & \downarrow & & \| & & \\
0 & \longrightarrow & T^{\mathrm{Zar}}_\varphi \mathcal{M}_h(X, J, dA) & \longrightarrow & T^{\mathrm{Zar}}_{(0,\varphi)} \mathcal{M}_h(X, \gamma, dA) & \longrightarrow & \mathbb{R}_t & \longrightarrow & 0 \\
 & & \downarrow & & \downarrow & & & & \\
 & & \ker(\varphi^* D) & = & \ker(\varphi^* D) & & & & \\
 & & \downarrow & & \downarrow & & & & \\
 & & 0 & & 0 & & & &
\end{array} \tag{3.4.34}$$

The exactness of the first row is obvious and the exactness of the second row follows from construction. The exactness of the left and middle columns follow by using the long exact sequence associated to $0 \to \varphi^* T_\Sigma \to \varphi^* T_X \to \varphi^* N_\Sigma \to 0$. The exactness of the right column now follows formally.

In view of condition (1) of Definition 3.4.1, the natural map from the short exact sequence (3.4.28) to the upper row short exact sequence in this diagram is an isomorphism. In fact, the following diagram

$$\begin{array}{ccc}
T_{(0,\varphi)} \mathcal{M}^d_{\mathbf{b},G}(\Sigma, \gamma) & \xrightarrow{\simeq} & \mathrm{Hom}^G(\ker \psi^* D, \mathrm{coker}\ \psi^* D) \\
\downarrow \simeq & & \downarrow \simeq \\
T_{(0,\varphi)} \mathcal{M}_h(\Sigma, \gamma, d) & \xrightarrow{\mathbf{w}_\varphi} & \mathrm{Hom}(\ker \varphi^* D, \mathrm{coker}\ \varphi^* D)
\end{array} \tag{3.4.35}$$

is commutative and, because of Lemma 3.4.10, all its arrows are isomorphisms. Indeed, the isomorphism given by the left vertical arrow is induced by the natural map $\mathcal{M}^d_{\mathbf{b},G}(\Sigma,\gamma) \to \mathcal{M}_h(\Sigma,\gamma,d)$, the right vertical arrow comes by restricting to the H-fixed subspace while the horizontal arrows come by considering the families of operators (3.4.30) and $(t,\varphi'' : \Sigma'' \to \Sigma_t) \mapsto \varphi''^*(D^N_{\Sigma_t,\gamma(t)})$ on $\mathcal{M}^d_{\mathbf{b},G}(\Sigma,\gamma,d)$ and $\mathcal{M}_h(\Sigma,\gamma,d)$ respectively. For future reference, we will label the bottom horizontal arrow of this commutative diagram as $\mathbf{w}_\varphi$.

Observe that both sides of the isomorphism $\mathbf{w}_\varphi$ have natural orientations. Indeed, we have the following.

1. $\mathcal{M}_h(\Sigma,\gamma,d)$ is fibered over the oriented interval $(-\epsilon,\epsilon)$ with complex fibers of complex dimension r. Recall that $r = (2h-2) - d(2g-2)$ is the number of branch points of φ.

2. An orientation of the vector space

$$\operatorname{Hom}(\ker \varphi^* D, \operatorname{coker}\, \varphi^* D) = (\ker \varphi^* D)^* \otimes \operatorname{coker}\, \varphi^* D \tag{3.4.36}$$

   is the same as an orientation of its determinant line

$$(\ker \varphi^* D)^{\otimes 1+2r} \otimes \det(\operatorname{coker}\, \varphi^* D) \tag{3.4.37}$$

   where we recall that $\dim \ker \varphi^* D = 1$ and $\dim \operatorname{coker}\, \varphi^* D = 1 + 2r$. Since the tensor square of a 1-dimensional space is canonically oriented, it follows that we simply need an orientation of the determinant line

$$\det(\varphi^* D) = \det(\ker \varphi^* D) \otimes \det(\operatorname{coker}\, \varphi^* D)^* \tag{3.4.38}$$

   of the Cauchy–Riemann operator $\varphi^* D$. This can be defined in the usual manner by homotoping D to a $\mathbb{C}$-linear Cauchy–Riemann operator and using the

spectral flow.

Thus, the isomorphism $\mathbf{w}_\varphi$ has a well-defined sign. Notice that the isomorphism (3.4.13) for $t_0 = 0$, whose sign is referenced in Theorem 3.4.7, can be derived as follows from $\mathbf{w}_\varphi$ by passing to the quotient $\mathbb{R}_t$ of $T_{(0,\varphi)}\mathcal{M}_h(\Sigma, \gamma, d)$, using the short exact sequence

$$0 \to T_\varphi \mathcal{M}_h(\Sigma, d) \to T_{(0,\varphi)}\mathcal{M}_h(\Sigma, \gamma, d) \to \mathbb{R}_t \to 0 \tag{3.4.39}$$

and making the identification $T_{(0,\varphi)}\mathcal{M}^d_{\mathbf{b},G}(\Sigma, \gamma) \xrightarrow{\simeq} T_{(0,\varphi)}\mathcal{M}_h(\Sigma, \gamma, d)$. In particular, we have

$$\mathrm{sgn}(\mathbf{w}_\varphi) = \mathrm{sgn}(\gamma, 0; p). \tag{3.4.40}$$

Without loss of generality, we will assume for the rest of this section that $(0, p)$ is a positive wall-crossing, i.e., $\mathrm{sgn}(\mathbf{w}_\varphi) = +1$. This assumption can be achieved by reversing the direction of γ if it is not already satisfied.

Target of the Kuranishi model

The target of the Kuranishi map is the space $E_{(0,\varphi)}\mathcal{M}_h(X, \gamma, dA)$ of all $\eta \in \Omega^{0,1}(\Sigma', \varphi^* T_X)$ modulo all elements of the kind that appear on the left side of (3.4.32). In fact, in view of the surjectivity of $T^{\mathrm{Zar}}_{(0,\varphi)}\mathcal{M}_h(X, \gamma, dA) \to \mathbb{R}_t$, we can restrict attention to just elements as in the left side of (3.4.32) for which $t = 0$. Using the long exact sequence associated to the short exact sequence $0 \to \varphi^* T_\Sigma \to \varphi^* T_X \to \varphi^* N_\Sigma \to 0$, it now follows that we have a natural identification

$$E_{(0,\varphi)}\mathcal{M}_h(X, \gamma, dA) = \mathrm{coker}\ \varphi^* D. \tag{3.4.41}$$

3.4.3 Implicit function theorem

We will now explain how to construct the Kuranishi model for $\mathcal{M}_h(X, \gamma, dA)$ near $(0, \varphi)$ using the implicit function theorem on an appropriate map $\Phi : \mathcal{X} \to \mathcal{Y}$ of Banach spaces (defined on a neighborhood of 0). Let $m \geq 2$ be an integer. Define $\mathcal{Y}$ to be the space of sections of $\Lambda^{0,1}T^*_{\Sigma'} \otimes_{\mathbb{C}} \varphi^* T_X$ of class L^2_m (i.e. m derivatives in L^2). Note that $\mathcal{Y}$ has a natural action of

$$\Gamma = \mathrm{Aut}(\varphi) = N_G(H)/H \tag{3.4.42}$$

We define $\mathcal{X} = \mathcal{X}_{\mathrm{vf}} \oplus \mathcal{X}_{\mathrm{acs}} \oplus \mathbb{R}_t$, where the summands are defined as follows. Choose a Γ-invariant set of points $M \subset \Sigma'$ where φ is unramified such that Σ' has no automorphisms fixing M. For each $[z] \in M/\Gamma$, choose a complement $D_{[z]} \subset T_{\varphi(z)}X$ for $d\varphi_z(T_z\Sigma')$. Define $\mathcal{X}_{\mathrm{vf}}$ to be the kernel of the Γ-equivariant evaluation map

$$L^2_{m+1}(\Sigma', \varphi^* T_X) \to \bigoplus_{z \in M} (T_{\varphi(z)}X)/D_{[z]}. \tag{3.4.43}$$

Define $\mathcal{X}_{\mathrm{acs}}$ to be a (finite dimensional) Γ-invariant complement of the map

$$\Omega^0_M(\Sigma', T_{\Sigma'}) \xrightarrow{\bar{\partial}} \Omega^{0,1}(\Sigma', T_{\Sigma'}) \tag{3.4.44}$$

consisting of sections compactly supported in $\Sigma' \setminus M$. Here, $\Omega^0_M(\Sigma', T_{\Sigma'})$ is the space of vector fields on Σ' which vanish on M.

To define the map $\Phi : \mathcal{X} \to \mathcal{Y}$, we need to fix the following choices.

1. A metric on X with its associated exponential map exp.

2. A J-linear connection ∇ on T_X.

3. A smooth 1-parameter family of complex vector bundle isomorphisms

$$I_{X,t} : (T_X, J) \xrightarrow{\simeq} (T_X, J_t) \tag{3.4.45}$$

with $I_{X,t} = \mathrm{id} + \frac{1}{2} J A_{X,t}$, where $A_{X,0} = 0$ and $A_{X,t}$ is J-antilinear for all $t \in \mathbb{R}$ with $|t| < \epsilon$ (after shrinking ϵ if necessary). Explicitly, we can define the endomorphism $A_{X,t}$ by the formula

$$A_{X,t} = -2J(\mathrm{id} + JJ_t)(\mathrm{id} - JJ_t)^{-1}. \tag{3.4.46}$$

Note that we have $\dot{\gamma}(0) = \frac{d}{dt}|_{t=0} J_t = \frac{d}{dt}|_{t=0} A_t = K$.

4. A smooth family $y \mapsto j_y$ of almost complex structures on Σ' parametrized by $y \in \mathcal{X}_{\mathrm{acs}}$ with $\frac{d}{d\tau}|_{\tau=0} j_{\tau y} = y$ and $j_0 = j_{\Sigma'}$ being the given almost complex structure on Σ'. The concrete choice we shall work with is given by declaring the following to be an isomorphism of complex vector bundles

$$I_{\Sigma',y} = \mathrm{id} + \frac{1}{2} j_{\Sigma'} y : (T_{\Sigma'}, j_{\Sigma'}) \to (T_{\Sigma'}, j_y) \tag{3.4.47}$$

for $y \in \mathcal{X}_{\mathrm{acs}}$ with small $|y|$.

We can now define the map $\Phi : \mathcal{X} \to \mathcal{Y}$ for small $(\xi, y, t) \in \mathcal{X}$ as

$$\Phi(\xi, y, t) = PT_\nabla \left[I_{X,t}^{-1} \circ (\bar{\partial}_{J_t, j_y} \exp_\varphi \xi) \circ I_{\Sigma',y} \right]. \tag{3.4.48}$$

Note that Φ is a smooth map of Banach manifolds (see e.g. the discussion in [MS12, page 43] or [Par16, Section B.5]). Here, the symbol $\exp_\varphi \xi$ denotes the map $z \mapsto \exp_{\varphi(z)} \xi(z)$ of regularity L^2_{m+1} (i.e. $m+1$ derivatives in L^2), the symbol $\bar{\partial}_{J_t, j_y}$ denotes

the non-linear Cauchy–Riemann operator

$$u \mapsto \tfrac{1}{2}(du + J_t \circ du \circ j_y) \tag{3.4.49}$$

and PT_∇ denotes the backward parallel transport along the paths $\tau \mapsto \exp_{\varphi(z)} \tau\xi(z)$ induced by the connection ∇. Clearly, Φ is Γ-equivariant, $\Phi(0) = 0$ and it is immediate that a neighborhood of $(0, \varphi) \in \mathcal{M}_h(X, \gamma, dA)$ is identified homeomorphically with a neighborhood of $[0] \in \Phi^{-1}(0)/\Gamma$.

To apply the implicit function theorem to Φ, we first note that the linearization $\tilde{D} = D\Phi(0) : \mathcal{X} \to \mathcal{Y}$ is given by the following formula (compare (3.4.32) from §3.4.2)

$$(\xi, y, t) \mapsto \tilde{D}_\varphi \xi + \tfrac{1}{2}(J \circ d\varphi \circ y + tK \circ d\varphi \circ j_{\Sigma'}) \tag{3.4.50}$$

where recall (from §3.4.2) that $\tilde{D}_\varphi$ denotes the linearized Cauchy–Riemann operator on $\varphi^* T_X$ (defined with $j_{\Sigma'}$ and J held fixed). We can now readily see that we have natural identifications

$$\ker \tilde{D} = T^{\mathrm{Zar}}_{(0,\varphi)} \mathcal{M}_h(X, \gamma, dA) \tag{3.4.51}$$

$$\operatorname{coker} \tilde{D} = E_{(0,\varphi)} \mathcal{M}_h(X, \gamma, dA). \tag{3.4.52}$$

Split the maps $\ker \tilde{D} \to \mathcal{X}$ and $\mathcal{Y} \to \operatorname{coker} \tilde{D}$ by choosing a Γ-invariant complement $\mathcal{X}' \subset \mathcal{X}$ for $\ker \tilde{D}$ and a Γ-invariant lift $E \subset \mathcal{Y}$ of $\operatorname{coker} \tilde{D}$.

Lemma 3.4.11. *There are Γ-invariant neighborhoods $0 \in U \subset \ker \tilde{D}$ and $0 \in V \subset \mathcal{X}'$ such that $\Phi^{-1}(E) \cap (U \times V)$ is given by the graph of a Γ-equivariant smooth function $\Psi : U \to \mathcal{X}'$ satisfying $\Psi(0) = 0$ and $D\Psi(0) = 0$.*

Proof. Apply the implicit function theorem to the Γ-equivariant map Φ. □

Definition 3.4.12 (Kuranishi map)**.** *Define the* Γ*-equivariant map* $F : U \to E$ *by*

$$F(\kappa) = \Phi(\kappa + \Psi(\kappa)). \tag{3.4.53}$$

This map F *yields the germ at* 0 *of a* Γ*-equivariant smooth map from* $\ker \tilde{D}$ *to* $\operatorname{coker} \tilde{D}$ *or equivalently*

$$T^{\mathrm{Zar}}_{(0,\varphi)}\mathcal{M}_h(X,\gamma,dA) \to E_{(0,\varphi)}\mathcal{M}_h(X,\gamma,dA) \tag{3.4.54}$$

which we call a Kuranishi map *for* $(0,\varphi) \in \mathcal{M}_h(X,\gamma,dA)$.

It is clear from Lemma 3.4.11 that there is a neighborhood $\mathcal{U}'$ of $[0] \in F^{-1}(0)/\Gamma$ and a neighborhood $\mathcal{U}''$ of $(0,\varphi) \in \mathcal{M}_h(X,\gamma,dA)$ with an obvious homeomorphism

$$\Xi : \mathcal{U}' \xrightarrow{\simeq} \mathcal{U}'' \tag{3.4.55}$$

mapping $[0]$ to $(0,\varphi)$. We may assume, after shrinking $\mathcal{U}'$, $\mathcal{U}''$ and ϵ, that we have

$$\mathcal{U}' = ((-\epsilon,\epsilon) \times \mathcal{O}) \cap \mathcal{M}_h(X,\gamma,dA) \tag{3.4.56}$$

for a neighborhood $\mathcal{O}$ of $\varphi : \Sigma' \to X$ in the Gromov topology (see the discussion following [Par16, Definition 9.1.3] for an explicit description of a basis of neighborhoods).

3.4.4 Analysis of the Kuranishi map

All functions in this subsection are to be understood to be germs near 0 of smooth functions which vanish at 0.

Definition 3.4.13. *Given two finite dimensional vector spaces* W_1, W_2*, two functions*

$F', F'' : W_1 \to W_2$ and an integer $l \geq 1$, we write

$$F' \equiv_l F'' \tag{3.4.57}$$

to mean that the i^{th} derivatives of $F' - F''$ vanish at 0 for all $0 \leq i \leq l$. Clearly, $\equiv_l$ is an equivalence relation.

We will now analyze the Kuranishi map F from Definition 3.4.12. From Lemma 3.4.11, it follows that $F(0) = 0$ and $DF(0) = 0$. We fix some notations first. Let

$$\varsigma : \ker \tilde{D} \to \ker(\varphi^* D) \tag{3.4.58}$$

$$t : \ker \tilde{D} \to \mathbb{R}_t \tag{3.4.59}$$

be the natural projections (see the exact commutative diagram in §3.4.2 for more details, recalling that $T^{\mathrm{Zar}}_{(0,\varphi)}\mathcal{M}_h(X, \gamma, dA) = \ker \tilde{D}$). Fix a splitting

$$\ker \tilde{D} = T_{(0,\varphi)}\mathcal{M}_h(\Sigma, \gamma, d) \oplus \ker \varphi^* D = \mathbb{R}_t \oplus T_\varphi \mathcal{M}_h(\Sigma, d) \oplus \ker \varphi^* D \tag{3.4.60}$$

compatible with the inclusions $T_\varphi \mathcal{M}_h(\Sigma, d) \subset T_{(0,\varphi)}\mathcal{M}_h(\Sigma, \gamma, d) \subset \ker \tilde{D}$ and the projections (3.4.58), (3.4.59).

Observe that $F^{-1}(0) \subset \ker \tilde{D} = T^{\mathrm{Zar}}_{(0,\varphi)}\mathcal{M}_h(X, \gamma, dA)$ contains a Γ-invariant submanifold Z (passing through 0) corresponding to $\mathcal{U}' \cap \mathcal{M}_h(\Sigma, \gamma, d)$ under Ξ with

$$T_0 Z = T_\varphi \mathcal{M}_h(\Sigma, \gamma, d) \subset \ker \tilde{D}. \tag{3.4.61}$$

Choosing $\varsigma' : \ker \tilde{D} \to \ker(\varphi^* D) \simeq \mathbb{R}$ to be a smooth function such that, near 0, $Z = (\varsigma')^{-1}(0)$ is a regular level set, we can then find a smooth function

$$F_1 : \ker \tilde{D} \to (\ker \varphi^* D)^* \otimes \operatorname{coker} \tilde{D} = \operatorname{Hom}(\ker \varphi^* D, \operatorname{coker} \varphi^* D) \tag{3.4.62}$$

such that $F = \langle \varsigma', F_1 \rangle$, where

$$\langle , \rangle : \ker \varphi^* D \otimes \mathrm{Hom}(\ker \varphi^* D, \mathrm{coker}\ \varphi^* D) \to \mathrm{coker}\ \varphi^* D \tag{3.4.63}$$

is the natural pairing. In view of (3.4.61), Z is given by the graph of a function $f : T_0 Z \to \ker(\varphi^* D)$ with $f(0) = 0$ and $Df(0) = 0$. As a result, we may take $\varsigma' = \varsigma - f$ and, in particular, we have the relation

$$\varsigma' \equiv_1 \varsigma. \tag{3.4.64}$$

Since $DF(0) = 0$, it follows that $F_1(0) = 0$. Let F_1' denote the function obtained by restricting F_1 to the subspace $T_0 Z \subset \ker \tilde{D}$. We then have the following.

Lemma 3.4.14. *The relation*

$$F_1' \equiv_1 \boldsymbol{w}_\varphi \tag{3.4.65}$$

holds, where $\boldsymbol{w}_\varphi$ is the isomorphism introduced in diagram (3.4.35).

Proof. Given any $b \in T_0 Z$, we need to show that

$$\mathbf{w}_\varphi(b) \cdot \sigma = \left(\tfrac{d}{d\lambda} F_1(b_\lambda) \cdot \tilde{\sigma} \right) |_{\lambda=0} \tag{3.4.66}$$

where $\tilde{\sigma} \in \ker \tilde{D}$ is any element (with σ being its image in $\ker \varphi^* D$) and $\lambda \mapsto b_\lambda$ is a path in Z with $b_\lambda(0) = 0$ and $\dot{b}_\lambda(0) = b$. We will make use of the fact that $\Psi(0) = 0$ and $D\Psi(0) = 0$ (Lemma 3.4.11). Recall that $\mathbf{w}_\varphi$ comes from linearizing the Cauchy-Riemann operator. Accordingly, we have

$$\mathbf{w}_\varphi(b) \cdot \sigma \equiv \left(\tfrac{d}{d\lambda} D\Phi(b_\lambda + \Psi(b_\lambda)) \cdot \tilde{\sigma} \right) |_{\lambda=0} \pmod{\mathrm{im}\ \tilde{D}} \tag{3.4.67}$$

$$\equiv (D^2\Phi)(0) \cdot b \cdot \tilde{\sigma} \pmod{\mathrm{im}\ \tilde{D}} \tag{3.4.68}$$

where $D^2\Phi = D(D\Phi)$ denotes the second derivative of Φ. On the other hand,

$$F_1(b_\lambda)\cdot\tilde{\sigma} = \tfrac{\partial}{\partial\mu}F(b_\lambda+\mu\tilde{\sigma})|_{\mu=0} \tag{3.4.69}$$

for each λ, by the definition of F_1. We now compute

$$\begin{aligned}\left(\tfrac{d}{d\lambda}F_1(b_\lambda)\cdot\tilde{\sigma}\right)|_{\lambda=0} &= \tfrac{\partial^2}{\partial\lambda\,\partial\mu}F(b_\lambda+\mu\tilde{\sigma})|_{\lambda=\mu=0} &(3.4.70)\\ &\equiv \tfrac{\partial^2}{\partial\lambda\,\partial\mu}\Phi(b_\lambda+\mu\tilde{\sigma}+\Psi(b_\lambda+\mu\tilde{\sigma}))|_{\lambda=\mu=0} \pmod{\text{im }\tilde{D}} &(3.4.71)\\ &\equiv (D^2\Phi)(0)\cdot b\cdot\tilde{\sigma} \pmod{\text{im }\tilde{D}} &(3.4.72)\end{aligned}$$

which proves (3.4.66). □

Corollary 3.4.15. *$Z' = F_1^{-1}(0) \subset \ker\tilde{D}$ is a Γ-invariant submanifold passing through 0 and its natural projection to $\ker\varphi^*D$ is a local diffeomorphism at 0. In other words, Z' is given by the graph of a smooth function $g : \ker\varphi^*D \to T_{(0,\varphi)}\mathcal{M}_h(\Sigma,\gamma,d)$ such that $g(0)=0$ and g is Γ-equivariant. In particular, the submanifolds Z and Z' are transverse at $0\in\ker\tilde{D}$.*

Proof. Apply the implicit function theorem to the Γ-equivariant map F_1, noting that $DF_1(0)|_{T_0Z} = \mathbf{w}_\varphi$ is an isomorphism. □

To continue with the bifurcation analysis, we need to understand the natural projection $Z'\subset\ker\tilde{D}\to\mathbb{R}_t$, which equivalent to understanding the smooth function

$$\ker\varphi^*D \xrightarrow{g} T_{(0,\varphi)}\mathcal{M}_h(\Sigma,\gamma,d) \xrightarrow{t} \mathbb{R}_t. \tag{3.4.73}$$

This will be more convenient to do in explicit coordinates which we now introduce. Recall that $\dim_{\mathbb{R}}\ker\varphi^*D = 1$ and fix an element $0\neq\sigma\in\ker\varphi^*D$ for the rest of this section. Using σ, we will sometimes regard the map ς from (3.4.58) as an $\mathbb{R}$-valued linear functional on $\ker\tilde{D}$ (with σ corresponding to 1). Finally, let $\mathbf{z} : T_\varphi\mathcal{M}_h(\Sigma,d)\to$

$\mathbb{C}^r$ be a $\mathbb{C}$-linear isomorphism. Thus, the splitting (3.4.60) now gives us explicit coordinates

$$\ker \tilde{D} = T_{(0,\varphi)}\mathcal{M}_h(X,\gamma,dA) \xrightarrow{(t,\mathbf{z},\varsigma)} \mathbb{R}_t \oplus \mathbb{C}^r_{\mathbf{z}} \oplus \mathbb{R}_\varsigma. \tag{3.4.74}$$

We also obtain (non-canonical) splittings $\ker \tilde{D} = T_{(0,\varphi)}\mathcal{M}_h(\Sigma,\gamma,d) \oplus \ker \varphi^* D$ and $T_{(0,\varphi)}\mathcal{M}_h(\Sigma,\gamma,d) = \mathbb{R}_t \oplus T_\varphi \mathcal{M}_h(\Sigma,d)$ which we will use without further comment below.

Orient the 1-dimensional space $\ker \varphi^* D$ using σ. Then the natural orientation on $\mathrm{Hom}(\ker \varphi^* D, \operatorname{coker} \varphi^* D)$ gives rise to an orientation on $\operatorname{coker} \varphi^* D$. Using the isomorphism $\mathbf{w}_\varphi$, the element σ and the coordinates (3.4.74), we can identify

$$\operatorname{coker} \varphi^* D = \mathbb{R}_t \oplus \mathbb{C}^r_{\mathbf{z}}. \tag{3.4.75}$$

This isomorphism preserves orientation since $\mathrm{sgn}(\mathbf{w}_\varphi) = +1$. With these choices of coordinates, we can write the Kuranishi map as

$$F(t,\mathbf{z},\varsigma) = (\mathrm{id} + A(t,\mathbf{z},\varsigma)) \cdot (\varsigma - f(t,\mathbf{z})) \cdot ((t,\mathbf{z}) - g(\varsigma)) \tag{3.4.76}$$

with $f \equiv_1 0$ and $A(\cdot)$ is a $(1+2r)\times(1+2r)$ real matrix valued function with $A(0) = 0$. The zeros of F are exactly the points of Z (zeros of the factor $\varsigma - f(t,\mathbf{z})$) and the points of Z' (zeros of the factor $(t,\mathbf{z}) - g(\varsigma)$). Clearly, the points in Z represent multiple covers and we are left to determine whether or not the points in $Z' \setminus \{0\}$ represent *embedded* pseudo-holomorphic curves of genus h and class dA. This depends on the section $\sigma \in \ker \varphi^* D$ as we explain below.

We begin by recalling that the normal bundle N_Σ admits an almost complex structure $\hat{J}^0_\infty$ (see [Wen21, Equation (B.1)] for the definition, also compare Lemma 3.5.9) such that the projection $N_\Sigma \to \Sigma$ is pseudo-holomorphic and the element $\sigma \in$

$\ker \varphi^* D$ can be regarded as a $\hat{J}^0_\infty$-holomorphic map $\hat{\sigma} : \Sigma' \to N_\Sigma$ (see [Wen21, Lemma B.4] for a proof).

Lemma 3.4.16. *$\hat{\sigma} : \Sigma' \to N_\Sigma$ is a multiply covered map if and only if there is a subgroup $H \subsetneq H' \subset G$ with $(\theta^*)^{H'} = (\theta^*)^H \neq 0$. Moreover, in this case,*

(1) *φ is unramified, i.e., $r = 0$ and,*

(2) *if σ descends to some intermediate cover $\Sigma' \to C \to \Sigma$, then so does any element of* coker $\varphi^* D$.

Proof. Let $\tilde{\Sigma} \to \Sigma' = \tilde{\Sigma}/H \xrightarrow{\varphi} \Sigma$ be a Galois closure. If there exists a subgroup $H \subsetneq H' \subset G$ with $(\theta^*)^{H'} = (\theta^*)^H \neq 0$, then it follows from Lemma 3.4.10 that $\hat{\sigma}$ (or equivalently σ) descends to $\tilde{\Sigma}/H'$ and thus, $\hat{\sigma}$ is a multiply covered map.

Conversely, suppose $\hat{\sigma}$ is multiply covered. Then, it descends to $\tilde{\Sigma}/H' \xrightarrow{\zeta} \Sigma$ for some subgroup $H \subsetneq H' \subset G$. This implies that $(\theta^*)^{H'} \neq 0$ and thus, $(\theta^*)^{H'} = (\theta^*)^H$ since the latter is 1-dimensional by Lemma 3.4.10. It follows that the (co)kernels of $\varphi^* D$ and $\zeta^* D$ are identified, using formulas (3.4.23) and (3.4.24). This proves the statement about the descent of elements in coker $\varphi^* D$. It also follows that the indices of $\varphi^* D$ and $\zeta^* D$ are the same, i.e., the number n_φ of zeros of $d\varphi$ is the same as the number n_ζ of zeros of $d\zeta$ (when counted with multiplicity). This cannot happen with $[H' : H] > 1$, unless φ is unramified. □

We now split our discussion into two cases depending on whether $\hat{\sigma}$ is a simple or multiply covered $\hat{J}^0_\infty$-holomorphic map.

Case 1 ($\hat{\sigma}$ is multiply covered)

By Lemma 3.4.16 we can now assume that φ is unramified and that both $\ker \varphi^* D$ and coker $\varphi^* D$ are pulled back from some intermediate cover $\Sigma' = \tilde{\Sigma}/H \to C = \tilde{\Sigma}/H' \xrightarrow{\zeta} \Sigma$ with ζ being of degree d' and C being of genus h'. We can compare the Kuranishi

models for $(0,\varphi)$ and $(0,\zeta)$ and we find that they are compatible under the natural pullback maps

$$T^{\mathrm{Zar}}_{(0,\zeta)}\mathcal{M}_{h'}(X,\gamma,d'A) \to T^{\mathrm{Zar}}_{(0,\varphi)}\mathcal{M}_h(X,\gamma,dA) \tag{3.4.77}$$

$$E_{(0,\zeta)}\mathcal{M}_{h'}(X,\gamma,d'A) \to E_{(0,\varphi)}\mathcal{M}_h(X,\gamma,dA) \tag{3.4.78}$$

both of which are isomorphisms. Thus, every curve close to $(0,\varphi) \in \mathcal{M}_h(X,\gamma,dA)$ is a (d/d')-fold cover of a curve close to $(0,\zeta) \in \mathcal{M}_{h'}(X,\gamma,d'A)$. Thus, the space $\mathcal{M}^*_h(X,\gamma,dA)$ is empty near $(0,\varphi) \in \mathcal{M}_h(X,\gamma,dA)$. This completes the proof of Theorem 3.4.7 in this case.

Case 2 ($\hat{\sigma}$ is simple)

Lemma 3.4.17. $\Gamma = \mathrm{Aut}(\varphi) \subset \mathbb{Z}/2\mathbb{Z}$.

Proof. It suffices to show that Γ acts faithfully on the 1-dimensional space $(\theta^*)^H$. Suppose to the contrary that there is an element $h' \in N_G(H) \setminus H$ which acts trivially on $(\theta^*)^H$. Letting H' be the subgroup of $N_G(H)$ generated by H and h', we find that $(\theta^*)^{H'} = (\theta^*)^H \neq 0$. By Lemma 3.4.16, we see that this is a contradiction to the simplicity of the $\hat{J}^0_\infty$-holomorphic map $\hat{\sigma}$. □

Lemma 3.4.18 (Nearby embedded curves)**.** *All points in $Z' \setminus \{0\}$ which are sufficiently close to $0 \in \ker \tilde{D}$ are mapped into $\mathcal{M}^*_h(X,\gamma,dA)$ by the map Ξ from* (3.4.55)*.*

Proof. Since $\gamma \in \mathcal{B}'$, it avoids $\mathcal{W}_{\mathrm{emb}}$ and it is therefore enough to show that the curves in question are simple. We argue by contradiction. To this end, suppose $(t_n, \varphi'_n : \Sigma'_n \to X)$ is a sequence in $\mathcal{M}_h(X,\gamma,dA)$ converging to $(0,\varphi)$ such that for each $n \geq 1$, we have the following properties.

1. $(t_n, \varphi'_n : \Sigma'_n \to X)$ is the image under Ξ of a point $(t_n, \varphi_n, \sigma_n) \in \ker \tilde{D}$ lying in $Z' \setminus \{0\}$. Here, we are using the splitting (3.4.60).

2. φ_n' is a multiply covered J_{t_n}-holomorphic map.

Z and Z' are transverse at 0 by Corollary 3.4.15. Thus, (1) implies that, for $n \gg 1$, the φ_n' are not covers of the embedded $\gamma(t_n)$-curves $\Sigma_{t_n} \subset X$ defined in the discussion following (3.4.22) of genus g and class A. Arguing as in the proof of [Wen21, Theorem B.1], we conclude that by rescaling the φ_n' (and the almost complex structures J_{t_n}) suitably in the normal directions to Σ_{t_n} (and passing to a subsequence), we obtain a Gromov limit $\hat{\sigma}' : \Sigma' \to N_\Sigma$ which is $\hat{J}^0_\infty$-holomorphic and corresponds (under [Wen21, Lemma B.4]) to an element $0 \neq \sigma' \in \ker(\varphi^* D)$. Since $\dim_{\mathbb{R}} \ker \varphi^* D = 1$ and $\hat{\sigma}$ is simple, it follows that $\hat{\sigma}'$ is also simple. Since being a simple map is an open property, it follows that the rescalings of the φ_n' in the convergent subsequence are also simple maps for $n \gg 1$. As a result, for $n \gg 1$, the maps $\varphi_n' : \Sigma_n' \to X$ themselves are also simple J_{t_n}-holomorphic maps and this contradicts (2). □

Lemma 3.4.19. *The map* (3.4.73) *has a non-trivial Taylor expansion at the origin. More precisely, there exists an integer* $1 \leq n \leq d$ *and a real number* $c \neq 0$ *with*

$$(t \circ g)(\varsigma) =_n c\varsigma^n. \tag{3.4.79}$$

Proof. In view of the expression (3.4.76) for F, the assertion we wish to prove is equivalent to saying that $\chi(\varsigma) = (t \circ F)(0, \mathbf{0}, \varsigma)$ has a nonzero Taylor expansion

$$\chi(\varsigma) \equiv_{n+1} -c\varsigma^{n+1} \tag{3.4.80}$$

for some integer $1 \leq n \leq d$ and real number $c \neq 0$. Examining Definition 3.4.21, we see that this amounts to saying that the point $p = (J, \Sigma, \varphi : \Sigma' \to \Sigma, \iota)$ does not lie on the degeneracy locus $\mathcal{M}^d_{\mathbf{b},G}(\mathcal{M}^*_g(X, A); \mathbf{k}, \mathbf{c})_{\spadesuit}$. But this is guaranteed by the fact that γ belongs to the Baire set $\mathcal{B}''$ (see Lemma 3.4.5). □

Using Lemma 3.4.19, we will now determine the signed count of the zeros of F in

Z' for small $t \neq 0$. We divide into cases based on the sign of c and the parity of n. (Clearly, if $\Gamma = \mathbb{Z}/2\mathbb{Z}$, then n must be even.)

1. *$c \neq 0$, n odd.* The Taylor expansion (3.4.79) shows that for each small $t_0 \neq 0$ we have a unique ς_0 such that $t_0 = (t \circ g)(\varsigma_0)$. Moreover, since $\mathbf{z}$ is a complex coordinate, the sign of this zero can be determined by looking at the unique nonzero solution $\varsigma = \varsigma_0 = (c^{-1}t_0)^{1/n}$ of the simplified equation

$$\varsigma(t_0 - c\varsigma^n) = 0. \tag{3.4.81}$$

   We compute of the derivative at this point

$$\tfrac{d}{d\varsigma}\,\varsigma(t_0 - c\varsigma^n)|_{\varsigma=\varsigma_0} = t_0 - (n+1)c\varsigma_0^n = -nt_0 \tag{3.4.82}$$

   and find that the sign of the zero is $+1$ for $t < 0$ and -1 for $t > 0$.

2. *$c > 0$, n even.* The Taylor expansion (3.4.79) shows that for small $t_0 > 0$, there are exactly two distinct solutions $\varsigma_0 \sim \pm(c^{-1}t_0)^{1/n}$ to $t_0 = (t \circ g)(\varsigma_0)$ while there are no solutions for small $t_0 < 0$. Again, we can compute the signs of these zeros, by considering the model equation (3.4.81), and we find that both of these zeros have the sign -1.

3. *$c < 0$, n even.* Arguing exactly as in the previous case, we find that there are exactly two zeros $\varsigma_0 \sim \pm(c^{-1}t_0)^{1/n}$ for small $t_0 < 0$ (both with sign $+1$) and no zeros for small $t_0 > 0$.

Taking the Γ-symmetry into account and shrinking $\mathcal{O}$ and ϵ if necessary, we find that in all cases the signed number of embedded curves in a sufficiently small Gromov neighborhood $\mathcal{O}$ of $\varphi : \Sigma' \to \Sigma$ reduces by the amount $2/|\Gamma| \in \mathbb{Z}$ when γ goes from $-\epsilon < t < 0$ to $0 < t < \epsilon$. Except for the proof of Lemma 3.4.19, we have now completed the proof of Theorem 3.4.7. $\square$

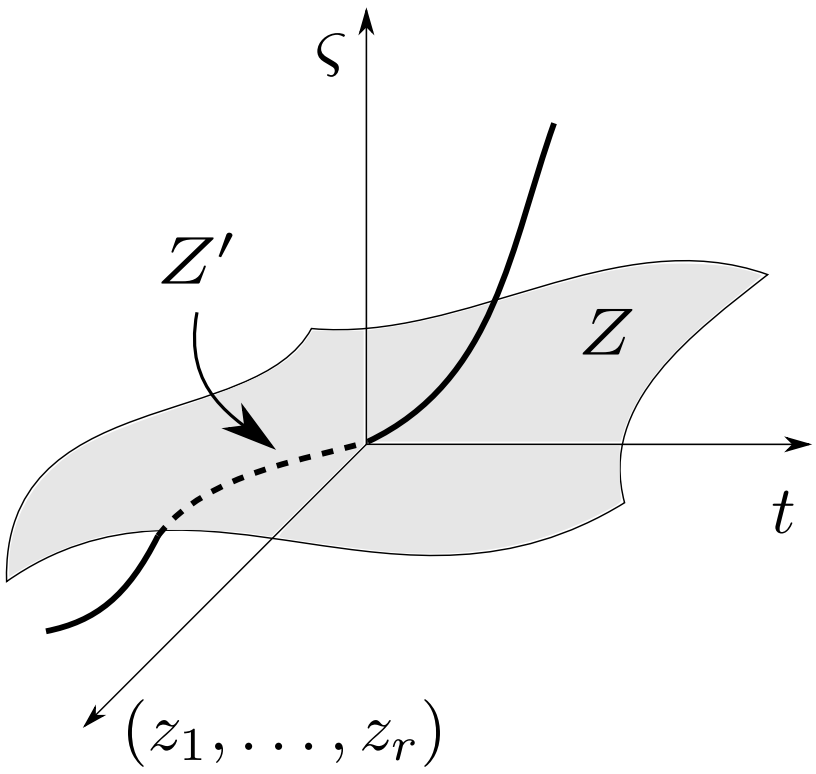


Figure 3.1: A schematic picture of the moduli space (when n is odd).

Non-degeneracy of Taylor expansion

For any point

$$\hat{p} = (\hat{J}, C, \hat{\varphi} : C' \to C, \hat{\imath}) \in \mathcal{M}^d_{\mathbf{b},G}(\mathcal{M}^*_g(X, A); \mathbf{k}, \mathbf{c}), \tag{3.4.83}$$

the moduli space $\mathcal{M}_h(X, \hat{J}, dA)$ of $\hat{J}$-holomorphic maps in class dA with domains of genus h (up to isomorphism) can be described locally near $\hat{\varphi}$, in the sense of the discussion surrounding (3.4.55), using a Kuranishi map

$$F_{\hat{p}} : T^{\mathrm{Zar}}_{\hat{\varphi}} \mathcal{M}_h(X, \hat{J}, dA) \to E_{\hat{\varphi}} \mathcal{M}_h(X, \hat{J}, dA) \tag{3.4.84}$$

with the domain and target of (3.4.84) being defined following §3.4.2 and §3.4.2. In particular, we have a short exact sequence and isomorphism

$$0 \to T_{\hat{\varphi}}\mathcal{M}_h(C,d) \to T^{\mathrm{Zar}}_{\hat{\varphi}}\mathcal{M}_h(X,\hat{J},dA) \xrightarrow{\hat{\varsigma}} \ker \hat{\varphi}^*\hat{D} \to 0 \tag{3.4.85}$$

$$E_{\hat{\varphi}}\mathcal{M}_h(X,\hat{J},dA) = \operatorname{coker}\ \hat{\varphi}^*\hat{D} \tag{3.4.86}$$

where $\hat{D} = D^N_{C,\hat{J}}$ is the normal deformation operator of embedded $\hat{J}$-curve $C \subset X$. Exactly as in (3.4.35), we have a natural map

$$T_{\hat{\varphi}}\mathcal{M}_h(C,d) \to \operatorname{Hom}(\ker \hat{\varphi}^*\hat{D}, \operatorname{coker}\ \hat{\varphi}^*\hat{D}). \tag{3.4.87}$$

Denote the image of the associated map

$$T_{\hat{\varphi}}\mathcal{M}_h(C,d) \otimes \ker \hat{\varphi}^*\hat{D} \to \operatorname{coker}\ \hat{\varphi}^*\hat{D} \tag{3.4.88}$$

by $W_{\hat{\varphi}}$. Note that the injectivity of (3.4.87) is equivalent to that of (3.4.88) which is in turn equivalent to saying that $W_{\hat{\varphi}}$ is a codimension 1 subspace of coker $\hat{\varphi}^*\hat{D}$.

Now, suppose that after choosing a splitting of (3.4.85), the Taylor expansion of the function $\hat{\varsigma} \mapsto F_{\hat{p}}(\mathbf{0},\hat{\varsigma}) \pmod{W_{\hat{\varphi}}}$ about $\hat{\varsigma} = 0$ is such that all terms of order $< n$ in $\hat{\varsigma}$ vanish (for some integer $n > 1$). Then, this is in fact independent of the choice of splitting (and even the particular choice of Kuranishi model) and in this case, we can view the order n term in the Taylor expansion invariantly as a linear map

$$T^{(n)}_{\hat{p}} : (\ker \hat{\varphi}^*\hat{D})^{\otimes n} \to (\operatorname{coker}\ \hat{\varphi}^*\hat{D})/W_{\hat{\varphi}}. \tag{3.4.89}$$

Remark 3.4.20. It is conceivable that we have $T^{(n)}_{\hat{p}} = 0$ for all $n > 1$.

Definition 3.4.21 (Degeneracy locus)**.** *We define the* degeneracy locus

$$\mathcal{M}^d_{\boldsymbol{b},G}(\mathcal{M}^*_g(X,A);\boldsymbol{k},\boldsymbol{c})_{\spadesuit} \subset \mathcal{M}^d_{\boldsymbol{b},G}(\mathcal{M}^*_g(X,A);\boldsymbol{k},\boldsymbol{c}) \tag{3.4.90}$$

to be the subspace consisting of all points $\hat{p} = (\hat{J}, C, \hat{\varphi} : C' \to C, \hat{\imath})$ *such that*

(1) $W_{\hat{\varphi}}$ *is a subspace of codimension* 1 *and,*

(2) *we have* $T^{(n)}_{\hat{p}} = 0$ *for each* $2 \leq n \leq d+1$.

Note that condition (1) in Definition 3.4.21 is open (as it is an injectivity condition) while condition (2) is closed (as it is a vanishing condition). As a result, the degeneracy locus is a locally closed subspace of $\mathcal{M}^d_{\mathbf{b},G}(\mathcal{M}^*_g(X,A);\mathbf{k},\mathbf{c})$.

Lemma 3.4.22. *Let* $\hat{p} = (\hat{J}, C, \hat{\varphi} : C' \to C, \hat{\imath})$ *be a point of* $\mathcal{M}^d_{\boldsymbol{b},G}(\mathcal{M}^*_g(X,A);\boldsymbol{k},\boldsymbol{c})$ *with* $\hat{D} = D^N_{C,\hat{J}}$ *being the normal deformation operator of* C*. Then, for any*

$$0 \neq \xi \in \ker \hat{\varphi}^* \hat{D} \tag{3.4.91}$$

$$0 \neq \xi' \in \ker \hat{\varphi}^* \hat{D}^\dagger = (\operatorname{coker}\, \hat{\varphi}^* \hat{D})^* \tag{3.4.92}$$

we can find an integer $1 \leq n \leq d$ *and a point* $q \in C$ *over which* $\hat{\varphi}$ *is unramified and the element*

$$\sum_{\tilde{q} \in \varphi^{-1}(q)} (\xi^{\otimes n+1} \otimes \xi')(\tilde{q}) \in \operatorname{Sym}^{n+1}_{\mathbb{R}} N_C \otimes_{\mathbb{R}} N_C^\dagger \tag{3.4.93}$$

is nonzero.

Proof. Choose a Galois closure $\tilde{C} \to \tilde{C}/H = C' \to C$. Enlarge the subgroup $H \subset G$ to a maximal subgroup H' such that we have $\mathbf{k}^{H'} = \mathbf{k}^H \neq 0$. Applying Lemma 3.4.16, we see that $\ker \hat{\varphi}^* \hat{D}$ and $\operatorname{coker} \hat{\varphi}^* \hat{D}$ both descend to the intermediate cover $\hat{\varphi}' : \tilde{C}/H' \to C$. The element defined by the formula (3.4.93) with $\hat{\varphi}'$ in place of

$\hat{\varphi}$ differs from (3.4.93) just by a factor of $[H' : H] \geq 1$. Thus, there is no loss of generality in assuming that $H' = H$.

Using this assumption, it follows that for all points $q \in C$, outside of a closed and at most countable set, $\xi(\hat{\varphi}^{-1}(q)) \subset N_{C,q}$ has d distinct nonzero elements. Choose such a point q. By perturbing it slightly if necessary, we may also assume that $\xi'(\tilde{q}) \neq 0$ for some $\tilde{q} \in \hat{\varphi}^{-1}(q)$. Enumerate the pre-images of q under $\hat{\varphi}$ as $\tilde{q}_1, \ldots, \tilde{q}_d \in C'$. Choose linear functionals

$$\Lambda : N_{C,q} \to \mathbb{R} \tag{3.4.94}$$

$$\Lambda' : N_{C,q}^{\dagger} \to \mathbb{R} \tag{3.4.95}$$

such that the d numbers $x_i = \Lambda(\xi(\tilde{q}_i))$ are *all distinct and nonzero* and the d numbers $y_j = \Lambda'(\xi'(\tilde{q}_j))$ are *not all zero*. This implies that the product

$$\begin{bmatrix} x_1^2 & x_2^2 & \cdots & x_d^2 \\ x_1^3 & x_2^3 & \cdots & x_d^3 \\ \vdots & \vdots & \ddots & \vdots \\ x_1^{d+1} & x_2^{d+1} & \cdots & x_d^{d+1} \end{bmatrix} \begin{bmatrix} y_1 \\ y_2 \\ \vdots \\ y_d \end{bmatrix} \tag{3.4.96}$$

cannot be zero, since the matrix is invertible (its determinant is $\prod_{i=1}^{d} x_i^2$ times the Vandermonde determinant associated to the d distinct numbers $x_1, \ldots, x_d$) and the column vector is nonzero. Take n to be the smallest number such that the n^{th} entry of this matrix product is nonzero. It then follows that

$$\left\langle \mathrm{Sym}_{\mathbb{R}}^{n+1}\Lambda \otimes_{\mathbb{R}} \Lambda', \sum_{\tilde{q} \in \varphi^{-1}(q)} (\xi^{\otimes n+1} \otimes \xi')(\tilde{q}) \right\rangle = \sum_{i=1}^{d} x_i^{n+1} y_i \neq 0. \tag{3.4.97}$$

which shows that this n suffices. □

Lemma 3.4.23. *Let $\hat{p}$, ξ, ξ' and $1 \leq n \leq d$ be as in Lemma 3.4.22. We can then*

find an element $\widehat{K} \in T_{\hat{\jmath}}\mathcal{J}(X,\omega) \subset \overline{\mathrm{End}}_{\hat{\jmath}}(T_X)$ *such that it and its first* n *derivatives vanish along* C *and its* $(n+1)^{\mathrm{st}}$ *derivative, in the directions normal to* C*, projects to a smooth* $\mathbb{R}$*-linear bundle homomorphism*

$$\widehat{K}_{n+1} : \mathrm{Sym}^{n+1}_{\mathbb{R}} N_C \to \overline{\mathrm{Hom}}_{\hat{\jmath}}(T_C, N_C) = \Lambda^{0,1}T_C^* \otimes_{\mathbb{C}} N_C \tag{3.4.98}$$

satisfying

$$\int_{C'} \langle \hat{\varphi}^*(i\widehat{K}_{n+1}), \xi^{\otimes n+1} \otimes \xi' \rangle \neq 0 \tag{3.4.99}$$

where i *denotes multiplication by* $\sqrt{-1}$ *on the fibers of* N_C*.*

Proof. It suffices to construct a bundle homomorphism L of the type (3.4.98) satisfying (3.4.99) with L in place of $\widehat{K}_{n+1}$. Indeed, in this case, we can then construct $\widehat{K} \in T_{\hat{\jmath}}\mathcal{J}(X,\omega)$ with $L = \widehat{K}_{n+1}$ (using a tubular neighborhood of $C \subset X$ and the argument from [Wen21, Lemma 5.26] to ensure ω-compatibility). Choose a smooth bundle homomorphism

$$L' : \mathrm{Sym}^{n+1}_{\mathbb{R}} N_C \to \Lambda^{0,1}T_C^* \otimes_{\mathbb{C}} N_C \tag{3.4.100}$$

which has positive pairing with (3.4.93) at the point $q \in C$ yielded by Lemma 3.4.22. Multiplying L' by a non-negative cutoff function on C supported near q, we can assume that its pairing with (3.4.93) is ≥ 0 at all points of C and > 0 at q. Taking $L = iL'$ completes the construction. We refer the reader to [MS12, Section 3.2] for comparison and further details. □

We can now state our main technical result.

Lemma 3.4.24. *The inclusion* (3.4.90) *has codimension* ≥ 1*.*

Proof. Near any point $\mathcal{M}^d_{\mathbf{b},G}(\mathcal{M}^*_g(X,A);\mathbf{k},\mathbf{c})_{\spadesuit}$, we can choosing auxiliary data (i.e.,

$\{D_{[z]}\}_{[z]\in M/\Gamma}$, (1), (2) and (4) as in §3.4.11) and define Kuranishi models $F_{\hat{p}}$ as in (3.4.84) for all points $\hat{p} \in \mathcal{M}^d_{\mathbf{b},G}(\mathcal{M}^*_g(X,A);\mathbf{k},\mathbf{c})$ in this neighborhood. Thus, in this neighborhood, the inclusion (3.4.90) is defined by the equations

$$T^{(2)}_{\hat{p}} = \cdots = T^{(d+1)}_{\hat{p}} = 0. \tag{3.4.101}$$

Given any point $\hat{p}$ satisfying (3.4.101), we show that the linearization of this system of d equations has rank ≥ 1. This will suffice to show that (3.4.90) is the inclusion of a C^∞ subvariety of codimension ≥ 1 (in the sense of [Wen21, Appendix C]).

Having fixed a point $\hat{p} \in \mathcal{M}^d_{\mathbf{b},G}(\mathcal{M}^*_g(X,A);\mathbf{k},\mathbf{c})_\spadesuit$ satisfying (3.4.101), pick a section $0 \neq \xi \in \ker \hat{\varphi}^* \hat{D}$. By condition (1) of Definition 3.4.21, we can also pick a section, unique up to scaling, $0 \neq \xi' \in \ker \hat{\varphi}^* \hat{D}^\dagger = (\operatorname{coker} \hat{\varphi}^* \hat{D})^*$ which annihilates $W_{\hat{\varphi}}$. Lemmas 3.4.22 and 3.4.23 then produce $1 \leq n \leq d$ and $\widehat{K}$ satisfying (3.4.93). We can now find a smooth path $\hat{\gamma} : [-1,1] \to \mathcal{J}(X,\omega)$ such with the following properties.

1. $\hat{\gamma}(0) = \hat{J}$ and $\frac{d}{ds}|_{s=0}\hat{\gamma}(s) = \widehat{K}$.

2. The path $\hat{\gamma}$ fixes the n-jet of $\hat{J}$ along C. More precisely, $|\hat{\gamma}(s) - \hat{J}| = O(|\nu_C|^{n+1})$ for all s, where $\nu_C = \operatorname{dist}(C,\cdot)$ denotes the distance to C induced by the given Riemannian metric.

3. $\hat{p}(s) = (\hat{\gamma}(s), C, \hat{\varphi} : C' \to C, \hat{\imath}) \in \mathcal{M}^d_{\mathbf{b},G}(\mathcal{M}^*_g(X,A);\mathbf{k},\mathbf{c})$ for all s.

Note that (3) actually follows from (1) and (2). Also observe that (2) implies that $T^{(i)}_{\hat{p}(s)} \equiv 0$ for $2 \leq i \leq n$. Thus, we have a well-defined element

$$DT^{(n+1)}|_{\hat{p}} \cdot \widehat{K} = \tfrac{d}{ds} T^{(n+1)}_{\hat{p}(s)}|_{s=0} \tag{3.4.102}$$

which can be computed by covariantly differentiating the non-linear Cauchy–Riemann operator and then pairing the result with $\xi' \in (\operatorname{coker} (\hat{\varphi}^* \hat{D})/W_{\hat{\varphi}})^*$. The computation

in Lemma 3.4.25 now shows that $DT^{(n+1)}|_{\hat{p}} \cdot \widehat{K} \neq 0$ and thus, we have proved that the rank in question is ≥ 1. □

Lemma 3.4.25 (Variation of Taylor coefficient)**.** *Using the notation from the proof Lemma 3.4.24, we have the following identity:*

$$\langle (DT^{(n+1)}|_{\hat{p}} \cdot \widehat{K}) \cdot \xi^{\otimes n+1}, \xi' \rangle = -\frac{1}{2} \int_{C'} \langle \hat{\varphi}^*(i\widehat{K}_{n+1}), \xi^{\otimes n+1} \otimes \xi' \rangle. \tag{3.4.103}$$

Proof. To get (3.4.103), we inspect the expression for the non-linear Cauchy–Riemann operator

$$(J', \varphi', j') \mapsto \bar{\partial}_{J',j'}(\varphi') = \tfrac{1}{2}(d\varphi' + J'(\varphi') \cdot d\varphi' \cdot j') \tag{3.4.104}$$

and differentiate it at $(\hat{J}, \hat{\varphi}, j_{C'})$ once with respect to the J' (in the direction $\widehat{K}$) and covariantly differentiate it $n+1$ times with respect to (φ', j') (in the direction $(\tilde{\xi}_{\mathrm{vf}}, \tilde{\xi}_{\mathrm{acs}}) \in \ker D(\bar{\partial}_{\hat{J}})|_{(\hat{\varphi}, j_{C'})}$, which is a lift of $\xi \in \ker \hat{\varphi}^* \hat{D}$) and use the fact that $\widehat{K}(\hat{\varphi}) \cdot d\hat{\varphi} \cdot j_{C'} = -\hat{J}(\hat{\varphi}) \cdot \widehat{K}(\hat{\varphi}) \cdot d\hat{\varphi}$. By assumption, $\widehat{K}$ and its first n derivatives vanish along C and thus, the only term which can survive the pairing with ξ' is the one from (3.4.103). □

3.5 An extension of Taubes' Gromov invariant to dimension 6

In this section, we will apply the results of the last two sections to prove Theorem 3.1.5. This invariant can be viewed as a partial extension of Taubes' Gromov invariant (defined in [Tau96b]) to symplectic Calabi–Yau 3-folds.

3.5.1 Linear wall crossing

The purpose of this subsection is to introduce the linear wall-crossing invariant (Definition 3.5.17), which will serve as a "correction term" in the definition of virtual count of embedded curves of class $2A$ in equation (3.1.5).

Fix a smooth closed oriented 2-manifold Σ of genus g and a $\mathbb{C}$-vector bundle $\pi : N \to \Sigma$ of rank 2 such that $c_1(N) \cdot [\Sigma] = 2g - 2$. Such pairs (Σ, N) arise by considering the normal bundle of any embedded pseudo-holomorphic curve of genus g in a symplectic Calabi–Yau 3-fold.

Definition 3.5.1 (Space of operators)**.** *Let $\mathcal{J}(\Sigma)$ denote the space of almost complex structures on Σ inducing the given orientation on Σ. For $\mathbb{K} \in \{\mathbb{R}, \mathbb{C}\}$, define the space*

$$\mathcal{CR}_{\mathbb{K}}(N) \tag{3.5.1}$$

to consist of pairs (j, D), where $j \in \mathcal{J}(\Sigma)$ and D is a $\mathbb{K}$-linear Cauchy–Riemann operator on N with respect to j, i.e.,

$$D : \Omega^0(\Sigma, N) \to \Omega^{0,1}_j(\Sigma, N) \tag{3.5.2}$$

is a linear first order differential operator such that for all $s \in \Omega^0(\Sigma, N)$ and $f \in C^\infty(\Sigma, \mathbb{K})$, we have $D(fs) = (\bar\partial_j f)s + fDs$. This is a bundle of affine spaces over the contractible space $\mathcal{J}(\Sigma)$ and therefore is also contractible. Every $(j, D) \in \mathcal{CR}_{\mathbb{K}}(N)$ satisfies $\operatorname{ind} D = 0$.

Definition 3.5.2 (Ordinary Brill–Noether loci)**.** *For $\mathbb{K} \in \{\mathbb{R}, \mathbb{C}\}$ and any integer $k \geq 1$, define the subspace*

$$\mathcal{W}^k_{\mathbb{K}}(N) \subset \mathcal{CR}_{\mathbb{K}}(N) \tag{3.5.3}$$

to consist of those pairs (j, D) *for which* $\dim_{\mathbb{K}} \ker D = k$. *Also define*

$$\mathcal{W}_{\mathbb{K}}(N) = \bigcup_{k \geq 1} \mathcal{W}^k_{\mathbb{K}}(N). \tag{3.5.4}$$

Finally, also define $\mathcal{W}'(N) \subset \mathcal{CR}_{\mathbb{R}}(N)$ *to be the set of pairs* (j, D) *such that there exists a section* $0 \neq \sigma \in \ker D$ *and a point* $z \in \Sigma$ *such that* $\sigma(z) = 0$.

Lemma 3.5.3. *We have the following codimension estimates on* $\mathcal{W}^k_{\mathbb{K}}(N)$ *and* $\mathcal{W}'(N)$.

(i) $\mathcal{W}^1_{\mathbb{R}}(N)$ *is a submanifold of* $\mathcal{CR}_{\mathbb{R}}(N)$ *of codimension* 1.

(ii) *For each* $k \geq 2$, $\mathcal{W}^k_{\mathbb{R}}(N)$ *is of codimension* ≥ 4 *in* $\mathcal{CR}_{\mathbb{R}}(N)$.

(iii) *For each* $k \geq 1$, $\mathcal{W}^k_{\mathbb{C}}(N)$ *is of codimension* ≥ 2 *in* $\mathcal{CR}_{\mathbb{C}}(N)$.

Proof. Given any element $(j, D) \in \mathcal{W}^k_{\mathbb{K}}(N)$, there exists (using [DW20, Lemma 1.1.4] for example) a smooth function

$$\mathcal{CR}_{\mathbb{K}}(N) \supset \mathcal{U} \xrightarrow{\mathcal{S}} \mathrm{Hom}_{\mathbb{K}}(\ker D, \mathrm{coker}\ D) \tag{3.5.5}$$

defined in a neighborhood $\mathcal{U}$ of (j, D) such that $\mathcal{U} \cap \mathcal{W}^k_{\mathbb{K}}(N) = \mathcal{S}^{-1}(0)$. Moreover, its linearization $d\mathcal{S}|_{(j,D)}$ restricted to $\Omega^{0,1}_j(\Sigma, \mathrm{Hom}_{\mathbb{K}}(N, N))$, i.e. the space of variations of D with j held fixed, is given by

$$A \mapsto (\kappa \mapsto A\kappa \pmod{\mathrm{im}\ D}) \tag{3.5.6}$$

where $\kappa \in \ker D$. The statements (i)–(iii) will therefore follow by proving suitable rank estimates on this map. By duality, it suffices to prove rank estimates for the "Petri map" (recall that $D^\dagger$ is the Serre dual of D as from Definition 3.2.4)

$$\ker D \otimes_{\mathbb{K}} \ker D^\dagger \to \Omega^0(C, N \otimes_{\mathbb{K}} N^\dagger) \tag{3.5.7}$$

$$s \otimes_{\mathbb{K}} s' \mapsto (z \mapsto s(z) \otimes_{\mathbb{K}} s'(z)). \tag{3.5.8}$$

By taking a $0 \neq s \in \ker D$ and $0 \neq s' \in \ker D^\dagger$ and using unique continuation for solutions of Cauchy–Riemann equations, it follows that the Petri map always has $\mathbb{K}$-rank ≥ 1. This proves (i) and (iii). Statement (ii) follows by using [DW20, Remark 1.1.15] (which is itself based on [Eft16, Lemma 4.4]) to conclude that the rank of the Petri map is $\geq 2k$ whenever $k \geq 2$ and $\mathbb{K} = \mathbb{R}$.

□

Lemma 3.5.4. *$\mathcal{W}'(N)$ is of codimension ≥ 3 in $\mathcal{CR}_{\mathbb{R}}(N)$.*

Proof. Since $\mathcal{W}'(N) \subset \mathcal{W}_{\mathbb{R}}(N)$, Lemma 3.5.3(ii) implies that it is enough to show that $\mathcal{W}'(N) \cap \mathcal{W}^1_{\mathbb{R}}(N)$ is of codimension ≥ 3 in $\mathcal{CR}_{\mathbb{R}}(N)$. To prove this, we argue as follows. Define

$$\widehat{\mathcal{W}'}(N) \subset \mathcal{CR}_{\mathbb{R}}(N) \times \Sigma \tag{3.5.9}$$

to consist of (j, D, z) such that $(j, D) \in \mathcal{W}^1_{\mathbb{R}}(N)$ and the elements of $\ker D$ vanish at z. It will be enough to show that $\widehat{\mathcal{W}'}(N)$ is a submanifold of codimension 4 in $\mathcal{W}^1_{\mathbb{R}}(N) \times \Sigma$. Fix any point $(j, D, z) \in \widehat{\mathcal{W}'}(N)$. Choose metrics on Σ, N – this yields an L^2 inner product on the space of sections of N. Fix a section $\sigma \in \ker D$ with $\|\sigma\|_{L^2} = 1$ (note that σ is unique up to multiplication by ± 1). Then, for every (j', D') in a sufficiently small neighborhood $\mathcal{U}$ of (j, D) in $\mathcal{W}^1_{\mathbb{R}}(N)$, there is a unique element $\sigma_{(j',D')} \in \ker D'$ such that its L^2 inner product with σ is 1. We now have the smooth map

$$\mathcal{U} \times \Sigma \xrightarrow{\mathcal{E}} N \tag{3.5.10}$$

$$(j', D', z') \mapsto \sigma_{(j',D')}(z') \tag{3.5.11}$$

which is compatible with the projection to Σ and satisfies $\widehat{\mathcal{W}'}(N) \cap (\mathcal{U} \times \Sigma) = \mathcal{E}^{-1}(0)$. It will now suffice to show that the linearization

$$d\mathcal{E}|_{(j,D,z)} : T_{(j,D)}\mathcal{W}^1_{\mathbb{R}}(N) \oplus T_z\Sigma \to N_z \tag{3.5.12}$$

is surjective. We will show the stronger statement that variations in D with j and z held constant suffice to surject onto N_z. First, we need to find an explicit expression for this linearization. To this end, let $\Pi : \Omega^{0,1}_j(\Sigma, N) \to \text{coker } D$ be the natural projection and let $Q : \text{im } D \to (\ker D)^\perp$ be the right inverse of D which we extend to $\Omega^{0,1}_j(\Sigma, N)$ by declaring $Q|_{(\text{im } D)^\perp} \equiv 0$.

Let $A \in \Omega^{0,1}_j(\Sigma, \text{Hom}_{\mathbb{R}}(N, N))$ be such that $A \in T_{(j,D)}\mathcal{W}^1_{\mathbb{R}}(N)$ (from the proof of Lemma 3.5.3(i), this means that $\Pi A\sigma = 0$). Then, there exists a smooth path $t \mapsto D_t \in \mathcal{W}^1_{\mathbb{R}}(N)$ such that $\dot{D}_0 = A$. Denoting $\sigma_t = \sigma_{(j,D_t)} = \sigma + \xi_t$ with $\xi_t \perp \sigma$ and differentiating the identity $D_t\sigma_t \equiv 0$, we find

$$0 = A\sigma + D\dot{\xi}_0 \tag{3.5.13}$$

which gives $\dot{\xi}_0 = -QA\sigma$. To summarize, we need to show that the map

$$L : \Omega^{0,1}_j(\Sigma, \text{Hom}_{\mathbb{R}}(N, N)) \to N_z \oplus \text{coker } D \tag{3.5.14}$$

$$L(A) := (-(QA\sigma)(z), \Pi A\sigma) \tag{3.5.15}$$

is surjective.

Now, if we let $D_z : \Omega^{0,1}_z(\Sigma, N) \to \Omega^{0,1}_j(\Sigma, N)$ be the restriction of D to the space of sections of N vanishing at z and $\Pi' : \Omega^{0,1}_j(\Sigma, N) \to \text{coker } D_z$ be the corresponding projection, we see that we have an injective map

$$\text{coker } D_z \to N_z \oplus \text{coker } D \tag{3.5.16}$$

$$\Pi'\eta \mapsto (-(Q\eta)(z), \Pi\eta) \tag{3.5.17}$$

which is an isomorphism by counting dimensions. Under this isomorphism, the map L is simply $A \mapsto \Pi' A\sigma$.

Suppose for the sake of contradiction that L is not surjective. Then, there is a nonzero distributional section τ of $N^\dagger$ such that $\langle A\sigma, \tau\rangle = 0$ for all smooth A (i.e., $\sigma \otimes \tau \equiv 0$ as a distributional section of $N \otimes_{\mathbb{R}} N^\dagger$) and $\tau|_{\operatorname{im} D_z} \equiv 0$. The second condition implies that $D^\dagger\tau$, viewed as a functional on $\Omega^0(\Sigma, N)$, vanishes on the subspace $\Omega^0_z(\Sigma, N)$. As a result, there is a linear functional $\Lambda_\tau : N_z \to \mathbb{R}$ such that for any $\sigma \in \Omega^0(\Sigma, N)$, we have $\langle D^\dagger\tau, \sigma\rangle = \Lambda_\tau(\sigma(z))$. Thus, $D^\dagger\tau$ is of class $(L^p_1)^*$ for any $2 < p < \infty$ (where L^p_1 is the space of L^p functions with one derivative in L^p and we are using the Sobolev embedding $L^p_1 \subset C^0$). By elliptic regularity, τ must be of class $(L^p)^* = L^q$ (where $\frac{1}{p} + \frac{1}{q} = 1$) and smooth in the complement of z (since it satisfies $D^\dagger\tau = 0$ here). But now $\sigma \otimes \tau \equiv 0$ with $\sigma \not\equiv 0$ and $\tau \not\equiv 0$ on the complement of z contradicts the fact that σ and τ satisfy unique continuation on the complement of z. □

Now, let $h \geq g$ be an integer such that $r := 2h - 2 - 2(2g-2) \geq 0$. Then, for any $j \in \mathcal{J}(\Sigma)$, the complex orbifold $\mathcal{M}_h(\Sigma, j, 2)$ of genus h holomorphic double covers of (Σ, j) branched at r points is non-empty. In fact, the underlying coarse moduli spaces are canonically identified for all j (since the isomorphism class of a cover is determined by purely topological data once j is specified). Call this common coarse moduli space $M_h(\Sigma, 2)$; it is naturally a smooth manifold of dimension $2r$.

Now, given $(j, D, [\varphi : \Sigma' \to \Sigma]) \in \mathcal{CR}_{\mathbb{K}}(N) \times M_h(\Sigma, 2)$, the operator $D_\varphi = \varphi^* D$ on $\varphi^* N$ admits a splitting $D_\varphi = D^+_\varphi \oplus D^-_\varphi$ into a direct sum of Cauchy–Riemann operators on isotypic subspaces corresponding to the trivial and sign representations of $\mathbb{Z}/2\mathbb{Z} = \operatorname{Aut}(\varphi)$. We have $\operatorname{ind} D^+_\varphi = \operatorname{ind} D = 0$ and $\operatorname{ind} D^-_\varphi = \operatorname{ind} D_\varphi = -2r$.

Definition 3.5.5 (Equivariant Brill–Noether loci)**.** *Given $\mathbb{K} \in \{\mathbb{R}, \mathbb{C}\}$, an integer*

$k \geq 1$ and integers h and r as above, define the subset

$$\mathcal{W}^k_{2,h,\mathbb{K}}(N) \subset \mathcal{CR}_{\mathbb{K}}(N) \times M_h(\Sigma, 2) \tag{3.5.18}$$

to consist of those $(j, D, [\varphi : \Sigma' \to \Sigma])$ for which $\dim_{\mathbb{K}} \ker D^-_\varphi = k$. Also define

$$\mathcal{W}_{2,h,\mathbb{K}}(N) = \bigcup_{k \geq 1} \mathcal{W}^k_{2,h,\mathbb{K}}(N). \tag{3.5.19}$$

Lemma 3.5.6. *We have the following codimension estimates on $\mathcal{W}^k_{2,h,\mathbb{K}}(N)$.*

(i) *$\mathcal{W}^1_{2,h,\mathbb{R}}$ is a submanifold of $\mathcal{CR}_{\mathbb{R}}(N) \times M_h(\Sigma, 2)$ of codimension $2r + 1$.*

(ii) *For each $k \geq 2$, $\mathcal{W}^k_{2,h,\mathbb{R}}$ has codimension $\geq 4r + 4$ in $\mathcal{CR}_{\mathbb{R}}(N) \times M_h(\Sigma, 2)$.*

(iii) *For each $k \geq 1$, $\mathcal{W}^k_{2,h,\mathbb{C}}$ has codimension $\geq 2r + 2$ in $\mathcal{CR}_{\mathbb{C}}(N) \times M_h(\Sigma, 2)$.*

Proof. The proof is a slight variation of the argument from Lemma 3.5.3. As in that argument, given any point $(j, D, [\varphi : \Sigma' \to \Sigma]) \in \mathcal{W}^k_{2,h,\mathbb{K}}$, we can find a smooth function

$$\mathcal{CR}_{\mathbb{K}}(N) \times M_h(\Sigma, 2) \supset \mathcal{U} \xrightarrow{\mathcal{S}} \mathrm{Hom}_{\mathbb{K}}(\ker D^-_\varphi, \mathrm{coker}\ D^-_\varphi) \tag{3.5.20}$$

defined on a neighborhood $\mathcal{U}$ of $(j, D, [\varphi : \Sigma' \to \Sigma])$ such that $\mathcal{U} \cap \mathcal{W}^k_{2,h,\mathbb{K}} = \mathcal{S}^{-1}(0)$. Moreover, the linearization $d\mathcal{S}|_{(j,D,[\varphi])}$ restricted to $\Omega^{0,1}_j(\Sigma, \mathrm{Hom}_{\mathbb{K}}(N, N))$, i.e. the space of variations of D with j and $[\varphi]$ held fixed, is given by

$$A \mapsto (\kappa \mapsto (\varphi^* A)\kappa \pmod{\mathrm{im}\ D^-_\varphi}) \tag{3.5.21}$$

where $\kappa \in \ker D^-_\varphi$. Thus, exactly as in Lemma 3.5.3, the proofs of (i)–(iii) will follow from suitable rank estimates on this map. By duality, it is enough to prove rank

estimates for the "Petri map"

$$\ker D_\varphi^- \otimes_{\mathbb{K}} \ker(D^\dagger)_\varphi^- \to \Omega^0(\Sigma', \varphi^* N \otimes_{\mathbb{K}} \varphi^* N^\dagger)^{\mathbb{Z}/2\mathbb{Z}} \tag{3.5.22}$$

$$\sigma \otimes_{\mathbb{K}} \sigma' \mapsto (\zeta \mapsto \sigma(\zeta) \otimes_{\mathbb{K}} \sigma'(\zeta)) \tag{3.5.23}$$

where $(\cdot)^{\mathbb{Z}/2\mathbb{Z}}$ denotes the $\mathbb{Z}/2\mathbb{Z} = \mathrm{Aut}(\varphi)$-invariant part. The key observation is that $\ker D_\varphi^- \otimes_{\mathbb{K}} \ker(D^\dagger)_\varphi^-$ already consists of $(\mathbb{Z}/2\mathbb{Z})$-invariant elements – indeed, $\mathbb{Z}/2\mathbb{Z}$ acts by the sign representation on both $\ker D_\varphi^-$ and $\ker(D^\dagger)_\varphi^-$ – and so, the map (3.5.22) is well-defined. Thus, we are reduced again to the argument of Lemma 3.5.3(i)–(iii). More precisely, we can show by unique continuation that the Petri map has $\mathbb{R}$-rank $= 2r+1$ in case of (i) and $\mathbb{C}$-rank $\geq r+1$ in case of (ii). For statement (iii), we invoke [DW20, Remark 1.1.15] (or [DW20, Remark 1.3.12]) to conclude that the rank of the Petri map is $\geq 2(2+2r)$ in case of (ii). □

Remark 3.5.7. The injectivity of Petri maps in the proofs of both Lemma 3.5.3 and Lemma 3.5.6 follows without appealing to Wendl's work on generic injectivity of Petri maps [Wen21] because the degrees of the corresponding covers are small enough.

Definition 3.5.8 (Co-orientation on $\mathcal{W}^1_{2,h,\mathbb{R}}(N)$)**.** *Given a tuple $(j, D, [\varphi : \Sigma' \to \Sigma])$ lying in $\mathcal{W}^1_{2,h,\mathbb{R}}(N)$, its normal space in $\mathcal{CR}_{\mathbb{R}}(N) \times M_h(\Sigma, 2)$ is canonically identified with*

$$\mathrm{Hom}(\ker D_\varphi^-, \mathrm{coker}\ D_\varphi^-). \tag{3.5.24}$$

We define an orientation on (3.5.24) *in the usual fashion, i.e., we deform D in a* 1*-parameter family to a $\mathbb{C}$-linear operator and use the spectral flow. This assignment of orientations to the normal spaces of $\mathcal{W}^1_{2,h,\mathbb{R}}(N) \subset \mathcal{CR}_{\mathbb{R}}(N)$ is called the natural co-orientation on $\mathcal{W}^1_{2,h,\mathbb{R}}(N)$.*

In order to make the next definition, we will need the following well-known fact

which is implicit in [Wen21, Appendix B] and [DW21, Lemma 3.6].

Lemma 3.5.9. *Given a complex curve C (not necessarily compact) and a complex vector bundle $\pi : N \to C$ on it, there is a natural bijection $D \mapsto J_D$ between*

(i) *$\mathbb{R}$-linear Cauchy–Riemann operators $D : \Omega^0(C, N) \to \Omega^{0,1}(C, N)$ and,*

(ii) *almost complex structures J on the total space of N which*

(a) *restrict to the standard complex structure on each fiber of π,*

(b) *restrict to the given complex structure on C along the zero section,*

(c) *are invariant under the $\mathbb{R}^\times$ action $\lambda \cdot (z, v) = (z, \lambda v)$ on N and,*

(d) *such that the projection π is J-holomorphic.*

Under this bijection, if C' is another complex curve (not necessarily compact), then a smooth map $\hat{s} : C' \to N$ is J_D-holomorphic if and only if $\varphi := \pi \circ \hat{s} : C' \to C$ is holomorphic and the corresponding section $s \in \Omega^0(C', \varphi^* N)$ given by

$$z \in C' \mapsto \hat{s}(z) \in N_{\varphi(z)} \tag{3.5.25}$$

satisfies $(\varphi^* D)s = 0$.

Proof. It is enough to consider the case of the trivial bundle $U \times \mathbb{C}^r \to U$ where $U \subset \mathbb{C}$ is an open subset with complex coordinate z and $\mathbb{C}^r$ is endowed with the complex coordinates $w = (w_1, \ldots, w_r)$. Let J be any almost complex structure on $U \times \mathbb{C}^r$ which satisfies conditions (a)–(d). Using (a), (b) and (d), we find that the bundle $\Omega^{1,0}_J$ near $U \times 0 \subset U \times \mathbb{C}^r$ has a local frame consisting of dz and

$$dw_i + A_i(z, w)d\bar{z} \tag{3.5.26}$$

for $1 \leq i \leq r$ where the A_i are smooth functions vanishing when $w_i = 0$ for all $1 \leq i \leq r$. Now, (c) implies that the A_i are actually $\mathbb{R}$-linear in w, i.e., they are of

the form

$$A_i(z,w) = \hat{A}_i(z)[w] \tag{3.5.27}$$

where $\hat{A} = (\hat{A}_i)_{i=1}^r : U \to \mathrm{End}_{\mathbb{R}}(\mathbb{C}^r)$ is a smooth function. Moreover, the expression (3.5.26) for a frame of $\Omega_J^{1,0}$ is valid on the whole of $U \times \mathbb{C}^r$. To this almost complex structure, we can associate the $\mathbb{R}$-linear Cauchy–Riemann operator $D = \bar{\partial} + \hat{A}[\cdot]$ on the bundle $U \times \mathbb{C}^r \to U$. Conversely, any $\mathbb{R}$-linear Cauchy–Riemann operator D on $U \times \mathbb{C}^r$ can be written as $\bar{\partial} + \hat{A}[\cdot]$ for some smooth function $\hat{A} : U \to \mathrm{End}_{\mathbb{R}}(\mathbb{C}^r)$ and thus, defines an almost complex structure J_D on N, satisfying (a)–(d), by taking the $(1,0)$ forms to be spanned by dz along with $dw_i + \hat{A}_i(z)[w]d\bar{z}$ for $1 \leq i \leq r$.

It is easy to see that this gives a bijective correspondence $D \mapsto J_D$ and that the bijection is preserved under the natural action of the group $C^\infty(U, GL_r(\mathbb{C}))$ on both sides. Finally, for any open subset $U' \subset \mathbb{C}$ and smooth map $\hat{s} : U' \to U \times \mathbb{C}^r$ given by $\zeta \mapsto (\varphi(\zeta), s(\zeta))$ is J_D-holomorphic if and only if we have

$$\bar{\partial}\varphi = 0, \quad \bar{\partial}s + \hat{A}(\varphi)[s]\bar{\partial}\bar{\varphi} = 0 \tag{3.5.28}$$

as desired. Now, the lemma extends to the general case by naturality. □

Let us now observe that for any $(j, D, [\varphi : \Sigma' \to \Sigma]) \in \mathcal{W}^1_{2,h,\mathbb{R}}(N)$, all non-zero elements of $\ker D_\varphi^-$ are non-zero scalar multiples of each other. Thus, for any $0 \neq \sigma \in \ker D_\varphi^-$, the corresponding J_D-holomorphic map $\hat{\sigma} : \Sigma' \to N$ is immersed/injective if and only if the same property holds for $a \cdot \sigma$ where $0 \neq a \in \mathbb{R}$. Therefore, we can make sense of the locus within $\mathcal{W}^1_{2,h,\mathbb{R}}$ where this immersion/injectivity property fails. The discussion which follows constructs Baire sets which avoid this locus (i.e. the immersion and injectivity property hold). These Baire sets will be used in the definition of linear wall crossing invariant (Definition 3.5.17).

Definition 3.5.10. *Define the subsets*

$$\mathcal{W}^1_{2,h,\mathbb{R}}(N)_{\mathrm{imm}}, \tag{3.5.29}$$

$$\mathcal{W}^1_{2,h,\mathbb{R}}(N)_{\mathrm{inj}} \tag{3.5.30}$$

of $\mathcal{W}^1_{2,h,\mathbb{R}}(N)$ *to consist of exactly those* $(j, D, [\varphi : \Sigma' \to \Sigma])$ *for which the* J_D*-holomorphic map* $\hat{\sigma} : \Sigma' \to N$ *corresponding to a nonzero* $\sigma \in \ker D^-_\varphi$ *fails to be immersed or injective respectively. Also define*

$$\mathcal{W}^1_{2,h,\mathbb{R}}(N)_{\mathrm{emb}} = \mathcal{W}^1_{2,h,\mathbb{R}}(N)_{\mathrm{imm}} \cup \mathcal{W}^1_{2,h,\mathbb{R}}(N)_{\mathrm{inj}}. \tag{3.5.31}$$

Lemma 3.5.11. *We have the following codimension estimates.*

(i) $\mathcal{W}^1_{2,h,\mathbb{R}}(N)_{\mathrm{imm}}$ *has codimension* $\geq 2r + 5$ *in* $\mathcal{CR}_{\mathbb{R}}(N) \times M_h(\Sigma, 2)$.

(ii) $\mathcal{W}^1_{2,h,\mathbb{R}}(N)_{\mathrm{inj}}$ *has codimension* $\geq 2r + 3$ *in* $\mathcal{CR}_{\mathbb{R}}(N) \times M_h(\Sigma, 2)$.

In particular, $\mathcal{W}^1_{2,h,\mathbb{R}}(N)_{\mathrm{emb}}$ *has codimension* $\geq 2r + 3$ *in* $\mathcal{CR}_{\mathbb{R}}(N) \times M_h(\Sigma, 2)$.

Proof. The proof is just a combination of the ideas from Lemmas 3.5.4 and 3.5.6. More precisely, the proof of (ii) exactly mirrors the proof of Lemma 3.5.4 while (i) requires a slightly more involved argument.

(i) Consider the finite cover $\tilde{\mathcal{C}}_{\mathrm{imm}} \to \mathcal{CR}_{\mathbb{R}}(N) \times M_h(\Sigma, 2)$ which consists of tuples $(j, D, [\varphi : \Sigma' \to \Sigma])$ along with the choice of a point $\zeta \in \mathrm{Crit}\ \varphi$. Within $\tilde{\mathcal{C}}_{\mathrm{imm}}$, we can consider the subset $\tilde{\mathcal{W}}^1_{2,h,\mathbb{R}}(N)_{\mathrm{imm}}$ consisting of $(j, D, [\varphi : \Sigma' \to \Sigma], \zeta)$ such that $(j, D, [\varphi : \Sigma' \to \Sigma]) \in \mathcal{W}^1_{2,h,\mathbb{R}}(N)$ and $d\hat{\sigma}|_\zeta = 0$ for all $0 \neq \sigma \in \ker D^-_\varphi$ (where $\hat{\sigma}$ is the J_D-holomorphic map corresponding to σ).

It will suffice to show that $\tilde{\mathcal{W}}^1_{2,h,\mathbb{R}}(N)_{\mathrm{imm}}$ is a submanifold of $\tilde{\mathcal{C}}_{\mathrm{imm}}$ of codimension $2r + 5$. To this end, fix a point $(j, D, [\varphi : \Sigma' \to \Sigma], \zeta) \in \tilde{\mathcal{W}}^1_{2,h,\mathbb{R}}(N)_{\mathrm{imm}}$.

Arguing exactly as in the proof of Lemma 3.5.4 (with Lemma 3.5.6 playing a role analogous to Lemma 3.5.3 in that proof), we are reduced to showing that

$$L : \Omega_j^{0,1}(\Sigma, \mathrm{Hom}_{\mathbb{R}}(N,N)) \to \mathrm{coker}\ (D_\varphi^-)_\zeta \tag{3.5.32}$$

$$L(A) := \Pi'((\varphi^* A)\sigma) \tag{3.5.33}$$

is surjective. Here, σ is some nonzero element of the 1-dimensional vector space $\ker D_\varphi^-$, the operator $(D_\varphi^-)_\zeta$ is the restriction of $D_\varphi^- : \Omega^0(\Sigma', \varphi^* N)^- \to \Omega_j^{0,1}(\Sigma', \varphi^* N)^-$ to the subspace of sections s for whose derivative at ζ (well-defined and independent of choice of connection on N since $\zeta \in \mathrm{Crit}\ \varphi$) is $\mathbb{C}$-anti-linear and Π' is the projection from $\Omega_j^{0,1}(\Sigma', \varphi^* N)^-$ to $\mathrm{coker}\ (D_\varphi^-)_\zeta$.

Assume, for the sake of contradiction, that L is not surjective. Then, arguing as in the proof of Lemma 3.5.4, we deduce the existence of a nonzero distributional section τ of $(\varphi^* N)^\dagger$ which

(a) is $\mathbb{Z}/2\mathbb{Z} = \mathrm{Aut}(\varphi)$-anti-invariant,

(b) vanishes on the image of $(D_\varphi^-)_\zeta$ and,

(c) is such that $\sigma \otimes \tau \equiv 0$ as a distributional section of $\varphi^* N \otimes_{\mathbb{R}} (\varphi^* N)^\dagger$ (for this last condition, we are using the observation that $\sigma \otimes \tau$ is invariant under $\mathbb{Z}/2\mathbb{Z}$ – see the parallel argument in the proof of Lemma 3.5.6).

Now, note that (a) and (b) imply that $(D_\varphi)^\dagger \tau$, viewed as a functional on $\Omega^0(\Sigma', \varphi^* N)$, factors through the projection

$$\Omega^0(\Sigma', \varphi^* N) \to \mathrm{Hom}_{\mathbb{C}}(T_\zeta \Sigma', N_{\varphi(\zeta)}) \tag{3.5.34}$$

$$s \mapsto (\varphi^* \nabla)^{1,0} s|_\zeta \tag{3.5.35}$$

and thus, it is of class $(L_2^p)^*$ for any $2 < p < \infty$ (where L_2^p is the space of L^p

functions with two derivatives in L^p and we are using the Sobolev embedding $L^p_2 \subset C^1$). By elliptic regularity it now follows that τ is of class $(L^p_1)^*$ and smooth in the complement of ζ (since it satisfies $(D_\varphi)^\dagger\tau = 0$ here). Now, if support of τ is not just the point ζ, then by unique continuation for σ and τ (on the complement of ζ), we must have $\sigma \otimes \tau \not\equiv 0$. Thus, we must have supp $\tau = \{\zeta\}$. It follows that there is an integer $m \geq 0$ such that τ applied to any element of $\Omega^{0,1}_j(\Sigma', \varphi^*N)$ depends only on the m-jet of this element at ζ. Since τ is known to be of class $(L^p_1)^*$ and $L^p_1 \not\subset C^1$, it follows that $m = 0$ and as a result τ acts on the space $\Omega^{0,1}_j(\Sigma', \varphi^*N)$ via a linear functional

$$M : (\Lambda^{0,1}T^*_{\Sigma',j} \otimes_{\mathbb{C}} \varphi^*N)_\zeta \to \mathbb{R}. \tag{3.5.36}$$

But now, it is easy to produce a smooth section s of φ^*N such that

$$-\langle (D_\varphi)^\dagger\tau, s\rangle = \langle \tau, D_\varphi s\rangle = M(D_\varphi s|_\zeta) = M((\varphi^*\nabla)^{0,1}s|_\zeta) \neq 0 \tag{3.5.37}$$

and $(\varphi^*\nabla)^{1,0}s|_\zeta = 0$. This contradicts the fact that $(D_\varphi)^\dagger\tau$ factors through the map (3.5.34).

(ii) The proof is similar to (i) but easier and so, we only indicate the setup. Define $\tilde{\mathcal{C}}_{\text{inj}}$ to be the subset of $\mathcal{CR}_{\mathbb{R}}(N) \times M_h(\Sigma, 2) \times \Sigma$ consisting of tuples $(j, D, [\varphi : \Sigma' \to \Sigma], z)$ such that $z \notin \varphi(\text{Crit } \varphi)$. We can then define the subset $\tilde{\mathcal{W}}^1_{2,h,\mathbb{R}}(N)_{\text{inj}} \subset \tilde{\mathcal{C}}_{\text{inj}}$ to consist of those elements $(j, D, [\varphi : \Sigma' \to \Sigma], z)$ for which $(j, D, [\varphi]) \in \mathcal{W}^1_{2,h,\mathbb{R}}(N)$ and $\sigma|_{\varphi^{-1}(z)} \equiv 0$ for all $0 \neq \sigma \in \ker D^-_\varphi$. It can now be shown as in (i) that $\tilde{\mathcal{W}}^1_{2,h,\mathbb{R}}(N)_{\text{inj}} \subset \tilde{\mathcal{C}}_{\text{inj}}$ is a submanifold of codimension $2r + 5$. Noting that $\mathcal{W}^1_{2,h,\mathbb{R}}(N)_{\text{inj}}$ is the image of $\tilde{\mathcal{W}}^1_{2,h,\mathbb{R}}(N)_{\text{inj}}$ under the obvious projection gives the desired result.

The statement concerning $\mathcal{W}^1_{2,h,\mathbb{R}}(N)_{\text{emb}}$ follows immediately from (i) and (ii). $\square$

Remark 3.5.12. The codimension estimates from Lemma 3.5.11 assert that the holomorphic curve associated to the kernel of a generic element in $\mathcal{W}^1_{2,h,\mathbb{R}}(N)$ is embedded. The injectivity of the relevant Petri map is not hard to establish because our discussion just involves $(\mathbb{Z}/2\mathbb{Z})$-equivariant Cauchy–Riemann operators from 2-fold covers. In fact, it is possible to extend the statements in Lemma 3.5.11 to elementary walls as in Definition 3.4.1. The fact that $\dim_{\mathbb{R}} \mathbf{k}^H = 1$, which follows from the definition of elementary wall type, is necessary in order to define the holomorphic sections uniquely up to rescaling. However, one needs to adapt the proof of genetric injectivity of Petri maps from [Wen21] to this context.

Remark 3.5.13. Any element $0 \neq \sigma \in \ker D^-_\varphi$ defines a *multi-valued section* of the bundle $N \to \Sigma$. The conclusion of Lemma 3.5.11 can therefore be phrased as the "generic embeddedness" of multi-valued sections. As pointed out in Remark 3.4.8, it may be possible to use similar arguments to establish similar generic behavior for 2-valued harmonic forms relevant to special Lagrangian geometry.

The next definition introduces the notion of "2-rigid" operators, paths between them and fixed endpoint homotopies between such paths. Such operators arise naturally as normal deformation operators of embedded pseudo-holomorphic curves with respect to any super-rigid almost complex structure.

Definition 3.5.14 (2-rigidity)**.** (a) *We say that $\alpha = (j, D) \in \mathcal{CR}_{\mathbb{R}}(N)$ is* 2-pre-rigid *if it misses the images of $\mathcal{W}'(N), \mathcal{W}^k_{2,h,\mathbb{R}}(N)$ and $\mathcal{W}^1_{2,h,\mathbb{R}}(N)_{\mathrm{emb}}$ in $\mathcal{CR}_{\mathbb{R}}(N)$ for all $h \geq g$ and $k \geq 2$. We say that α is 2-rigid if, in addition, it also misses the image of the map*

$$\mathcal{W}^1_{2,h,\mathbb{R}}(N) \to \mathcal{CR}_{\mathbb{R}}(N) \tag{3.5.38}$$

for all $h \geq g$. Denote the set of 2-rigid elements by $\mathcal{CR}_{\mathbb{R}}(N)_2$ and its intersection with $\mathcal{CR}_{\mathbb{C}}(N) \setminus \mathcal{W}_{\mathbb{C}}(N)$ by $\mathcal{CR}_{\mathbb{C}}(N)_2$.

(b) *Given* $\alpha_\pm \in \mathcal{CR}_\mathbb{R}(N)_2$*, define* $\mathcal{P}(\alpha_-, \alpha_+)$ *to be the space of smooth paths in* $\mathcal{CR}_\mathbb{R}(N)$ *from* α_- *to* α_+*. Define*

$$\mathcal{P}^{\pitchfork}(\alpha_-, \alpha_+) \subset \mathcal{P}(\alpha_-, \alpha_+) \tag{3.5.39}$$

to be the set of paths which consist entirely of 2*-pre-rigid elements and are transverse to* (3.5.38)*.*

(c) *Given* $\gamma_\pm \in \mathcal{P}^{\pitchfork}(\alpha_-, \alpha_+)$*, define* $\mathcal{P}(\gamma_-, \gamma_+)$ *to be the space of smooth paths in* $\mathcal{P}(\alpha_-, \alpha_+)$ *from* γ_- *to* γ_+*, i.e., smooth homotopies from* γ_- *to* γ_+*. Define*

$$\mathcal{P}^{\pitchfork}(\gamma_-, \gamma_+) \subset \mathcal{P}(\gamma_-, \gamma_+) \tag{3.5.40}$$

to be the set of homotopies which consist entirely of 2*-pre-rigid elements and are transverse to* (3.5.38)*.*

For the convenience of the reader, we now summarize the consequences of Lemmas 3.5.3, 3.5.6 and 3.5.11 which are relevant for the next lemma: $\mathcal{W}'(N)$ has codimension ≥ 3 in $\mathcal{CR}_\mathbb{R}(N)$, the image of $\mathcal{W}^k_{2,h,\mathbb{R}}(N)$ in $\mathcal{CR}_\mathbb{R}(N)$ has codimension ≥ 4 for $k \geq 2$ and the image of $\mathcal{W}^1_{2,h,\mathbb{R}}$ has codimension ≥ 3 in $\mathcal{CR}_\mathbb{R}(N)$. Moreover, both $\mathcal{W}^k_\mathbb{C}(N)$ and the image of $\mathcal{W}^k_{2,h,\mathbb{C}}$ in $\mathcal{CR}_\mathbb{C}(N)$ have codimension ≥ 2 for all $k \geq 1$.

Lemma 3.5.15. *We have the following genericity statements.*

(i) $\mathcal{CR}_\mathbb{R}(N)_2 \subset \mathcal{CR}_\mathbb{R}(N)$ *is a Baire subset.*

(ii) $\mathcal{CR}_\mathbb{C}(N)_2 \subset \mathcal{CR}_\mathbb{C}(N)$ *is a path connected Baire subset.*

(iii) *For* $\alpha_\pm \in \mathcal{CR}_\mathbb{R}(N)_2$*,* $\mathcal{P}^{\pitchfork}(\alpha_-, \alpha_+) \subset \mathcal{P}(\alpha_-, \alpha_+)$ *is a Baire subset.*

(iv) *For* $\gamma_\pm \in P^{\pitchfork}(\alpha_-, \alpha_+)$*,* $\mathcal{P}^{\pitchfork}(\gamma_-, \gamma_+) \subset \mathcal{P}(\gamma_-, \gamma_+)$ *is a Baire subset.*

Proof. This is a straightforward consequence of the Sard–Smale theorem combined with the codimension formulas/estimates from Lemmas 3.5.3, 3.5.6 and 3.5.11. □

Lemma 3.5.16. *Fix $\alpha_\pm \in \mathcal{CR}_{\mathbb{R}}(N)_2$, $\gamma_\pm \in \mathcal{P}^{\pitchfork}(\alpha_-, \alpha_+)$ and $\tilde{\gamma} \in \mathcal{P}^{\pitchfork}(\gamma_-, \gamma_+)$. Consider $\tilde{\gamma}$ as smooth map*

$$[-1,1] \times [-1,1] \to \mathcal{CR}_{\mathbb{R}}(N) \tag{3.5.41}$$

such that $\tilde{\gamma}(\cdot, \pm 1) \equiv \gamma_\pm(\cdot)$ and $\tilde{\gamma}(\pm 1, \cdot) \equiv \alpha_\pm$. Then, for any fixed $h \geq g$, the fiber product of (3.5.41) *and* (3.5.38) *is a naturally oriented smooth* 1*-manifold with boundary which is proper over $[-1,1]\times[-1,1]$ and has boundary exactly over $[-1,1]\times \pm 1$.*

Proof. The fact that the fiber product in question is a manifold with boundary (with boundary only over $[-1,1] \times \pm 1$) is obvious from the transversality assumption in the definition of $\mathcal{P}^{\pitchfork}(\gamma_-, \gamma_+)$. The orientation comes the standard orientation on $[-1,1] \times [-1,1]$ and the natural co-orientation from Definition 3.5.8. Thus, what remains to be proved is properness over $[-1,1]\times[-1,1]$ and this will be a consequence of the fact that all elements in the image of $\tilde{\gamma}$ are 2-pre-rigid. To establish properness, let $p_n \to p$ be a convergent sequence in $[-1,1] \times [-1,1]$ and let $\tilde{\alpha}_n = (j_n, D_n, [\varphi_n : \Sigma'_n \to \Sigma]) \in \mathcal{W}^1_{2,h,\mathbb{R}}(N)$ be a sequence such that we have $\tilde{\gamma}(p_n) = \alpha_n := (j_n, D_n)$ for all $n \geq 1$, with limit $\tilde{\gamma}(p) = \alpha := (j, D)$. We will show that $\tilde{\alpha}_n$ has a subsequence converging to an element $\tilde{\alpha} \in \mathcal{W}^1_{2,h,\mathbb{R}}(N)$. In other words, we will use the assumption $\tilde{\gamma} \in \mathcal{P}^{\pitchfork}(\gamma_-, \gamma_+)$ to argue that there is a limiting $(\mathbb{Z}/2\mathbb{Z})$-anti-invariant multi-valued section associated to $\alpha = (j, D)$ which is defined over a branched cover of Σ with smooth domain.

Fix a Hermitian inner product on N and let ω_0 be a symplectic form on the unit ball $B_1(N)$ taming the almost complex structure J_D (this exists by [DW21, Lemma 3.6]). Since $\alpha_n \to \alpha$, it follows that ω_0 tames J_{D_n} also for $n \gg 1$.

Let $\hat{\sigma}_n : \Sigma'_n \to N$ be the simple J_{D_n}-holomorphic map representing the class $2[\Sigma]$ corresponding to a section $\sigma_n \in \ker(D_n)^-_{\varphi_n}$ with $\|\sigma_n\|_{L^\infty} = \frac{1}{2}$. By Gromov

compactness, we can pass to a subsequence to assume that we have a limiting J_D-holomorphic stable map $\hat{\sigma} : \Sigma' \to N$ representing class $2[\Sigma]$ with Σ' being of arithmetic genus h. Note that the image of each $\hat{\sigma}_n$ (and therefore of $\hat{\sigma}$) is invariant under the automorphism $\mathcal{I} : (z, v) \mapsto (z, -v)$ of N. Define $\varphi = \pi \circ \hat{\sigma} : \Sigma' \to \Sigma$ and note that the corresponding section $\sigma \in \ker \varphi^* D$ satisfies $\|\sigma\|_{L^\infty} = \frac{1}{2}$. Each effective component of Σ' (i.e. an irreducible component on which $\hat{\sigma}$ is non-constant) must map with positive degree onto Σ via φ. Thus, there are at most two effective components in Σ'. We now consider two cases.

(a) (There are two effective components $C_1, C_2 \subset \Sigma$.) Since $\varphi|_{C_i} : C_i \to \Sigma$ are both of degree 1, they must be isomorphisms. Thus, $\hat{s}_1 := \hat{\sigma}|_{C_1}$ and $\hat{s}_2 := \hat{\sigma}|_{C_2}$ give us two sections $s_1, s_2 \in \ker D$ at least one of which must be nonzero. If $\hat{s}_1 \not\equiv 0$, then since the image of $\hat{\sigma}$ is connected invariant under $\mathcal{I}$, it follows that s_1 must vanish at a point $z \in C$. This contradicts the 2-pre-rigidity of $\alpha = (j, D)$.

(b) (There is a unique effective component $C \subset \Sigma'$.) Let $\psi : \tilde{C} \to C$ be the normalization. Since the image of $\hat{\sigma} \circ \varphi : \tilde{C} \to N$ is invariant under $\mathcal{I}$, it follows that $0 \neq \psi^* \sigma \in \ker D^-_{\varphi \circ \psi}$. Since $\alpha = (j, D)$ is 2-pre-rigid, it follows that $\dim \ker D^-_{\varphi \circ \psi} = 1$ and the J_D-holomorphic map $\hat{\sigma} \circ \psi$ is an embedding because α necessarily lies in the complement of $\mathcal{W}^1_{2,h,\mathbb{R}}(N)_{\mathrm{emb}}$. This implies that ψ is an isomorphism, C is smooth and $\hat{\sigma}|_C$ is an embedding representing class $2[\Sigma]$. Since $\hat{\sigma} : \Sigma' \to N$ is the Gromov limit of the simple (in fact, embedded by 2-pre-rigidity) J_{D_n}-curves $\hat{\sigma}_n : \Sigma'_n \to N$, Remark 3.3.2(b) shows that Σ' cannot have any ghost components. Thus, Σ' is smooth and coincides with C. Moreover, $0 \neq \sigma \in \ker D^-_\varphi$ and $\dim \ker D^-_\varphi = 1$.

We have now shown that Σ' is necessarily smooth. We conclude that the limit $\tilde{\alpha} = (j, D, [\varphi : \Sigma' \to \Sigma])$ lies in $\mathcal{W}^1_{2,h,\mathbb{R}}(N)$ as desired. $\square$

Definition 3.5.17 (Linear wall crossing invariant)**.** *Let $\alpha \in \mathcal{CR}_{\mathbb{R}}(N)_2$ be 2-rigid and let $h \geq g$ be given. We then define the invariant*

$$\mathrm{w}_{2,h}(\alpha) \in \mathbb{Z} \tag{3.5.42}$$

as follows. Set $r = (2h-2) - 2(2g-2)$. If $r < 0$, we define $\mathrm{w}_{2,h}(\alpha) = 0$.

If $r \geq 0$, choose a point $\beta \in \mathcal{CR}_{\mathbb{C}}(N)_2$ and choose a path $\gamma \in \mathcal{P}^{\mathrm{th}}(\alpha, \beta)$. Then, using Lemma 3.5.16 and the natural co-orientation from Definition 3.5.8, the fiber product of γ with (3.5.38) *is a compact oriented 0-manifold and we define $\mathrm{w}_{2,h}(\alpha)$ to be the signed number points in this fiber product.*

The following lemma shows that $\mathrm{w}_{2,h}(\alpha)$ does not depend on the auxiliary data therefore it is well-defined.

Lemma 3.5.18. *$\mathrm{w}_{2,h}(\alpha)$ in Definition 3.5.17 is independent of the choice of β and γ.*

Proof. It is enough to consider the case $r \geq 0$. Let us temporarily denote the quantity from Definition 3.5.17 as $\mathrm{w}_{2,h}(\alpha, \beta, \gamma)$ to emphasize the dependence on the choices β and γ.

(i) Fix the choice of β and consider $\gamma, \gamma' \in \mathcal{P}^{\mathrm{th}}(\alpha, \beta)$. Choose $\tilde{\gamma} \in \mathcal{P}^{\mathrm{th}}(\gamma, \gamma')$ and apply Lemma 3.5.16 (combined with the fact that an oriented 1-manifold with boundary has 0 boundary points when counted with sign) to get

$$\mathrm{w}_{2,h}(\alpha, \beta, \gamma) = \mathrm{w}_{2,h}(\alpha, \beta, \gamma'). \tag{3.5.43}$$

Thus, we may write $\mathrm{w}_{2,h}(\alpha, \beta, \gamma) = \mathrm{w}_{2,h}(\alpha, \beta)$.

(ii) Now, suppose $\beta, \beta' \in \mathcal{CR}_{\mathbb{C}}(N)_2$ are given. Choose a path δ in $\mathcal{CR}_{\mathbb{C}}(N)_2$ from β to β' – such a path exists by Lemma 3.5.15 and it is, moreover, automatically

in $\mathcal{P}^{\pitchfork}(\beta,\beta')$. Now, given any $\gamma \in \mathcal{P}^{\pitchfork}(\alpha,\beta)$, we can set $\gamma' \in \mathcal{P}^{\pitchfork}(\alpha,\beta')$ to be the concatenation of γ and δ. Since there are no intersections of δ with (3.5.38), it follows that

$$\mathrm{w}_{2,h}(\alpha,\beta) = \mathrm{w}_{2,h}(\alpha,\beta,\gamma) = \mathrm{w}_{2,h}(\alpha,\beta',\gamma') = \mathrm{w}_{2,h}(\alpha,\beta'). \tag{3.5.44}$$

Combining (i) and (ii), we conclude that $\mathrm{w}_{2,h}(\alpha,\beta,\gamma)$ is independent of the choices of β and γ. □

3.5.2 Proof of Theorem 3.1.5

The proof of Theorem 3.1.5 is an application of the arguments of §3.5.1 combined with Theorems 3.3.1 and 3.4.7.

We fix the primitive class $A \in H_2(X,\mathbb{Z})$ for the whole argument. Let $J_\pm \in \mathcal{J}(X,\omega)$ be super-rigid almost complex structures. We need to show that for all $h \geq 0$ we have

$$\mathrm{Gr}_{2A,h}(X,\omega,J_-) = \mathrm{Gr}_{2A,h}(X,\omega,J_+) \tag{3.5.45}$$

where $\mathrm{Gr}_{2A,h}(X,\omega,J_\pm)$ is the right side of (3.1.5) for $J = J_\pm$. We will accomplish this by studying the bifurcations that occur along a suitably chosen path $\gamma : [-1,1] \to \mathcal{P}(J_-,J_+)$. Before this, we need some preparatory facts analogous to Lemmas 3.5.4, 3.5.11, 3.5.15 and 3.5.16.

Definition 3.5.19. *Let $h \geq g \geq 0$ be integers.*

(i) *Define $\mathcal{J}'_{A,g} \subset \mathcal{J}(X,\omega)$ to be the set of all ω-compatible almost complex structures J on X for which there exists an embedded J-curve $C \subset X$ of genus g in class A, a point $z \in C$ and $0 \neq \sigma \in \ker D^N_{C,J}$ with $\sigma(z) = 0$.*

(ii) *Define* $\mathcal{W}_{2,h}(\mathcal{J}_{A,g}) \subset \mathcal{J}(X,\omega)$ *to be the set of all* ω*-compatible almost complex structures* J *on* X *for which there exists an embedded* J*-curve* $C \subset X$ *of genus* g *in class* A*, a holomorphic branched double cover* $\varphi : C' \to C$ *with* C' *being smooth of genus* h *and* $0 \neq \sigma \in \ker(D^N_{C,J})^-_\varphi$ *such that the corresponding* $J_{D^N_{C,J}}$*-holomorphic map* $\hat{\sigma} : C' \to N_C$ *defined using* (3.5.25) *is not an embedding.*

Lemma 3.5.20 (Analogue of Lemmas 3.5.4 and 3.5.11)**.** *We have the following codimension estimates.*

(i) $\mathcal{J}'_{A,g} \subset \mathcal{J}(X,\omega)$ *has codimension* ≥ 3 *for all* $g \geq 0$*.*

(ii) $\mathcal{W}_{2,h}(\mathcal{J}_{A,g}) \subset \mathcal{J}(X,\omega)$ *has codimension* ≥ 3 *for all* $h \geq g \geq 0$*.*

Proof. We first define the relevant universal moduli spaces and reduce the problem of estimating the codimension to proving a rank lower bound for the linearization of the defining equations of these universal moduli spaces. The proofs of the rank estimates from Lemmas 3.5.4 and 3.5.11 (which prove the analogous results for Cauchy–Riemann operators) apply verbatim once we combine them with [Wen21, Lemma 5.26] which shows that deformations in the space of Cauchy–Riemann operators can be realized by deformations in the space of ω-compatible almost complex structures. □

Lemma 3.5.21 (Analogue of Lemma 3.5.15)**.** *The subset* $\mathcal{B}^* \subset \mathcal{P}(J_-, J_+)$ *of paths* $\gamma : [-1,1] \to \mathcal{J}(X,\omega)$ *with* $\gamma(\pm 1) = J_\pm$ *which are disjoint from* $\mathcal{J}'_{A,g}$ *and* $\mathcal{W}_{2,h}(\mathcal{J}_{A,g})$ *for all* $A \in H_2(X,\mathbb{Z})$ *and* $h \geq g \geq 0$ *is a Baire subset.*

Proof. This is a direct consequence of a Sard–Smale argument applied to the codimension estimates from Lemma 3.5.20. □

For the statement of the next lemma, recall the Baire set $\mathcal{B}'$ from Lemma 3.4.3.

Lemma 3.5.22 (Analogue of Lemma 3.5.16)**.** *Fix a path* $\gamma \in \mathcal{B}^* \cap \mathcal{B}' \subset \mathcal{P}(J_-, J_+)$ *and integers* $0 \leq g \leq h$*.*

(i) *Consider all tuples of the form* $(t, \Sigma, \varphi : \Sigma' \to \Sigma)$ *where* $t \in [-1,1]$, $\Sigma \subset X$ *is an embedded* $\gamma(t)$*-curve of genus* g *in class* A *and* $(\Sigma', \varphi) \in \overline{\mathcal{M}}_h(\Sigma, 2)$ *with* Σ' *being smooth. Among these, the tuples satisfying*

$$\dim \ker D^N_{\Sigma,\gamma(t)} = 0, \quad \dim \ker \varphi^* D^N_{\Sigma,\gamma(t)} = 1 \tag{3.5.46}$$

form a finite set.

(ii) *Let* $t_n \to t$ *in* $[-1,1]$ *be a convergent sequence and suppose* $\Sigma'_n \subset X$ *is a sequence of embedded* $\gamma(t_n)$*-curves of genus* h *in class* $2A$ *which converge to a* $\gamma(t)$*-holomorphic stable map* $\varphi : \Sigma' \to \Sigma \subset X$ *in the Gromov topology with* Σ *being an embedded* $\gamma(t)$*-curve of genus* g *in class* $2A/d$ *for some* $d \in \{1,2\}$ *and* $(\Sigma', \varphi) \in \overline{\mathcal{M}}_h(\Sigma, d)$.

Then, Σ' *is necessarily smooth. If* $d = 1$*, then* φ *is an isomorphism. If* $d = 2$*, then the operator* $D^N_{\Sigma,\gamma(t)}$ *is an isomorphism and we have*

$$\dim \ker \varphi^* D^N_{\Sigma,\gamma(t)} = 1. \tag{3.5.47}$$

Proof. The key ingredients of the proof are Theorem 3.3.1 and Remark 3.3.2(b).

(i) Since $\gamma \in \mathcal{B}'$, it follows that this set is discrete. It will therefore suffice to prove compactness. To this end, suppose that $(t_n, \Sigma_n, \varphi_n : \Sigma'_n \to \Sigma_n)$ is a sequence of points in this set with $t_n \to t$. By Remark 3.3.2(b), the sequence Σ_n has, after passing to a subsequence, a Gromov limit Σ which is an embedded $\gamma(t)$-curve of the same genus g. Now, by identifying the normal bundles $N_{\Sigma_n} \to \Sigma_n$ for $n \gg 1$ with the fixed smooth vector bundle on $N_\Sigma \to \Sigma$, we can repeat the argument of Lemma 3.5.16 to conclude (the condition $\gamma \in \mathcal{B}' \cap \mathcal{B}^*$ is the analogue of the 2-pre-rigidity condition used in that proof).

(ii) First, suppose Σ itself is of class $2A$. Then, Remark 3.3.2(b) shows that $d = 1$,

φ is an isomorphism and, in particular, that Σ' is smooth. For the remainder suppose that Σ is of class A and $d = 2$. Theorem 3.3.1 now yields a section $s \in \ker \varphi^* D^N_{\Sigma,\gamma(t)}$ which is not pulled back from Σ. Let $\hat{s} : \Sigma' \to N_\Sigma$ be the associated $J_{D^N_{\Sigma,\gamma(t)}}$-holomorphic map.

Since $(\Sigma', \varphi) \in \overline{\mathcal{M}}_h(\Sigma, 2)$, it follows that Σ' has at most two effective components (i.e. irreducible components of Σ' on which φ is non-constant). We now split into two cases.

(a) (There are two effective components $C_1, C_2 \subset \Sigma'$.) Since $\varphi|_{C_i}$ has degree 1, it must be an isomorphism for $i = 1, 2$. Thus, $s|_{C_i}$ yields a section $s_i \in \ker D^N_{\Sigma,\gamma(t)}$ for $i = 1, 2$. Since s is not pulled back from Σ, it follows that $s_1 \not\equiv s_2$. Since C_1 can be connected to C_2 within Σ' by a (possibly empty) chain of ghost components, it also follows that there must be a point $z \in \Sigma$ such that $s_1(z) = s_2(z)$. The section $s_1 - s_2 \in \ker D^N_{\Sigma,\gamma(t)}$ now shows that $\gamma(t) \in \mathcal{J}'_{A,g}$. This contradicts $\gamma \in \mathcal{B}^*$.

(b) (There is a unique effective component $C \subset \Sigma$.) Let $\psi : \tilde{C} \to C$ be the normalization and set $\varphi' = \varphi \circ \psi : \tilde{C} \to \Sigma$. Since $\gamma(t) \notin \mathcal{W}_{2,h}(\mathcal{J}_{A,g'})$ for $g' = \text{genus}(\tilde{C})$, we deduce that $\hat{s} \circ \psi : \tilde{C} \to N_\Sigma$ is an embedding. This implies that ψ must be injective and therefore an isomorphism. Thus, C is smooth and $\hat{s}|_C$ is an embedding. Since $\ker(D^N_{\Sigma,\gamma(t)})^-_{\varphi'} \neq 0$, it follows from $\gamma \in \mathcal{B}'$ that this vector space must in fact be 1-dimensional and that $\ker D^N_{\Sigma,\gamma(t)}$ is an isomorphism. Using the inverse function theorem, we can now obtain a sequence $\Sigma_n \subset X$ of $\gamma(t_n)$-curves of genus g in class A converging to Σ. Arguing as in [Wen21, Appendix B], we can now identify tubular neighborhoods of Σ_n for $n \gg 1$ with the normal bundle of Σ equipped with the pullback of the almost complex structure $\gamma(t_n)$. Rescaling the images of the embedded curves Σ'_n in N_Σ to have distance $\frac{1}{2}$ from the zero section gives

a new sequence of embedded curves which are holomorphic with respect to the rescaled versions of $\gamma(t_n)$. The rescalings of the $\gamma(t_n)$ converge to $J_{D^N_{\Sigma,\gamma(t)}}$ and moreover, all these almost complex structures are tamed by a symplectic form for $n \gg 1$ ([DW21, Lemma 3.6]). Therefore, after passing to a subsequence, we can extract a Gromov limit given by a $J_{D^N_{\Sigma,\gamma(t)}}$-holomorphic map $\hat{s}' : \Sigma' \to N_\Sigma$ whose image doesn't lie in the zero section. Since C is the only effective component of $\hat{s}'$ (since a holomorphic mapping into any normal fiber must be constant), it follows that $\hat{s}'$ a scalar multiple of $\hat{s}$. We have therefore exhibited $\hat{s}$ as a limit of embedded pseudo-holomorphic curves of genus h and class $2[\Sigma]$. Recalling that $\hat{s}|_C$ is an embedding and applying Remark 3.3.2(b) shows that we must have $C = \Sigma'$.

We have now succeeded in showing that Σ' is smooth. The remaining facts concerning the operator $D^N_{\Sigma,\gamma(t)}$ have already been proved during the course of (b) above.

□

Remark 3.5.23. Although the proof of Lemma 3.5.22 is formally parallel to the proof of Lemma 3.5.16, it is crucial to note that we use Theorem 3.3.1 to construct the multisection s over Σ because the rescaling argument from [Wen21, Appendix B] or Remark 3.3.2(a) cannot be used unless we already know that $D^N_{\Sigma,\gamma(t)}$ is an isomorphism.

We can now complete the proof of Theorem 3.1.5. Choose a path γ from the Baire subset $\mathcal{B}^* \cap \mathcal{B}' \cap \mathcal{B}'' \subset \mathcal{P}(J_-, J_+)$. (Recall the definition of the Baire set $\mathcal{B}''$ from Lemma 3.4.5.) Define the moduli space

$$\mathcal{M}^*_h(X, \gamma, 2A) \to [-1, 1] \tag{3.5.48}$$

consisting of pairs (t, Σ'), where $t \in [-1, 1]$ and $\Sigma' \subset X$ is an embedded $\gamma(t)$-curve of

genus h in class $2A$. For each $0 \leq g \leq h$, also consider the moduli spaces

$$\mathcal{M}_g(X, \gamma, A) \to [-1, 1] \tag{3.5.49}$$

consisting of pairs (t, Σ), where $t \in [-1, 1]$ and $\Sigma \subset X$ is an embedded $\gamma(t)$-curve curve of genus g in class A. Both $\mathcal{M}_h^*(X, \gamma, 2A)$ and $\mathcal{M}_g(X, \gamma, A)$ are oriented 1-manifolds with boundary over $t = \pm 1$ since $\gamma \in \mathcal{B}'$. Moreover, $\mathcal{M}_g(X, \gamma, A)$ is compact by Remark 3.3.2(b). Since $\gamma \in \mathcal{B}' \cap \mathcal{B}^*$, Lemma 3.5.22 shows that the set of points in $(t, \Sigma, \varphi : \Sigma' \to \Sigma)$ with $(t, \Sigma) \in \mathcal{M}_g(X, \gamma, A)$ for some $g \leq h$, $(\Sigma', \varphi) \in \mathcal{M}_h(\Sigma, 2)$ with Σ' smooth and

$$\ker D^N_{\Sigma, \gamma(t)} = 0, \quad \ker \varphi^* D^N_{\Sigma, \gamma(t)} = 1 \tag{3.5.50}$$

forms a finite set. Moreover, γ is transverse to the map

$$\mathcal{M}^2_{\mathbf{b}^r, \mathbb{Z}/2\mathbb{Z}}(\mathcal{M}_0^*(X, A); \theta, \theta^{\oplus 1+2r}) \to \mathcal{J}(X, \omega) \tag{3.5.51}$$

at any such point $(t, \Sigma, \varphi : \Sigma' \to \Sigma)$, where we have set $r = 2h - 2 - 2(2g - 2) \geq 0$, $\mathbf{b}^r = ((2), \ldots, (2))$ and $\theta =$ the sign representation of $\mathbb{Z}/2\mathbb{Z}$.

Enumerate the times t for which there exists $\varphi : \Sigma' \to \Sigma$ satisfying (3.5.50) as

$$-1 < t_1 < \cdots < t_k < 1 \tag{3.5.52}$$

and, for each $1 \leq i \leq k$, enumerate the intersections with the maps (3.5.51) as

$$p_{i,j} = (t_i, \Sigma_{i,j} \subset X, \varphi_{i,j} : \Sigma'_{i,j} \to \Sigma_{i,j}) \tag{3.5.53}$$

for $1 \leq j \leq \ell_i$, where ℓ_i is a positive integer depending on i. Theorem 3.4.7 now provides numbers $\epsilon_i > 0$ and neighborhoods $\mathcal{O}_{i,j}$ of $\varphi_{i,j} : \Sigma'_{i,j} \to \Sigma_{i,j} \subset X$ and by

shrinking these further we can assume that the products

$$\mathcal{U}_{i,j} = (t_i - \epsilon_i, t_i + \epsilon_i) \times \mathcal{O}_{i,j} \tag{3.5.54}$$

are all pairwise disjoint. We now make a series of claims.

1. The $\mathbb{Z}$-valued function

$$(t, \Sigma) \mapsto \mathrm{w}_{2,h}(D^N_{\Sigma,\gamma(t)}) \tag{3.5.55}$$

defined on $\mathcal{M}_g(X, \gamma, A)$ away from $t = t_1, \ldots, t_k$ is locally constant.

Proof. Indeed, fix any point (t, Σ) such that $t \neq t_i$ for all $1 \leq i \leq k$. Then, $D^N_{\Sigma,\gamma(t)}$ is 2-rigid and the same is true for all nearby points (t', Σ'). Thus, any nearby (t', Σ') can be connected to (t, Σ) by a path of curves with 2-rigid normal deformation operators and so, $\mathrm{w}_{2,h}(D^N_{\Sigma,\gamma(t)}) = \mathrm{w}_{2,h}(D^N_{\Sigma',\gamma(t')})$. □

2. The moduli space $\mathcal{M}^*_h(X, \gamma, 2A) \setminus \bigcup_{i,j} \mathcal{U}_{i,j}$ is compact. Moreover, the restriction of (3.5.48) over the complement of $\{t_i\}_{1 \leq i \leq k}$ is proper.

Proof. Both assertions follow directly from Lemma 3.5.22. □

3. The expression (3.1.5) is well-defined with $J = \gamma(t)$ provided $t \in [-1, 1] \setminus \{t_1, \ldots, t_k\}$ is a regular value of (3.5.48) and each of the maps (3.5.49). If $t_i < t' \leq t'' < t_{i+1}$ are two such regular values (for some $1 \leq i < k$), then we have the equality

$$\mathrm{Gr}_{2A,h}(X, \omega, \gamma(t')) = \mathrm{Gr}_{2A,h}(X, \omega, \gamma(t'')). \tag{3.5.56}$$

Proof. Claim (2) shows that the first sum in the expression (3.1.5) is well-defined for such a t. The remaining sums in (3.1.5) are well-defined by Claim

(1). For regular values $t' < t''$ in (t_i, t_{i+1}), we can now use the cobordisms given by restricting (3.5.48) and (3.5.49) to $[t', t'']$ and the local constancy statement from Claim (1) to deduce (3.5.56). □

4. In the previous claim, if we instead have $t_i - \epsilon_i < t' < t_i < t'' < t_i + \epsilon_i$ for some $1 \leq i \leq k$, then (3.5.56) still holds.

Proof. By Claim (3), we can assume t', t'' are arbitrarily close to t_i. For small $|t - t_i|$ and $\Sigma \in \{\Sigma_{i,j} : 1 \leq j \leq \ell_i\}$, define $\Sigma(t)$ to be the unique embedded $\gamma(t)$-curve close to Σ (obtained by noting that $\ker D^N_{\Sigma,\gamma(t_i)} = 0$ and applying the inverse function theorem). For such a Σ, define the set $I_\Sigma(i) = \{1 \leq j \leq \ell_i : \Sigma = \Sigma_{i,j}\}$, and note that we have

$$\sum_{j \in I_\Sigma(i)} \operatorname{sgn}(\gamma, t_i; p_{i,j}) = \operatorname{sgn}(D^N_{\Sigma,\gamma(t_i)}) \cdot (\mathrm{w}_{2,h}(D^N_{\Sigma(t''),\gamma(t'')}) - \mathrm{w}_{2,h}(D^N_{\Sigma(t'),\gamma(t')})) \tag{3.5.57}$$

where we check orientations by using the identification

$$\det(D^N_{\Sigma,\gamma(t_i)}) \otimes \det(D^N_{\Sigma,\gamma(t_i)})^-_{\varphi_{i,j}} = \det(\varphi^*_{i,j} D^N_{\Sigma,\gamma(t_i)}) \tag{3.5.58}$$

for each $j \in I_\Sigma(i)$. The quantities $\operatorname{sgn}(\gamma, t_i; p_{i,j})$ are as in Definition 3.4.4. Now recall that the neighborhoods $\mathcal{U}_{i,j}$ were constructed by applying Theorem 3.4.7 to each point $p_{i,j} \in \mathcal{U}_{i,j}$ and note that the isotropy group Γ is $\mathbb{Z}/2\mathbb{Z}$ in all cases. Restrict (3.5.48) to $[t', t'']$ and note that this gives a compact 1-manifold with boundary over the complement of $\bigcup_{i,j} \mathcal{U}_{i,j}$ by Claim (2). Thus, as in Claim (3), this portion contributes equally to both sides of (3.5.56).

For the remaining contribution (coming from $\bigcup_{i,j} \mathcal{U}_{i,j}$), we see from Theorem 3.4.7 that the first sum in the definition of $\mathrm{Gr}_{2A,h}$ changes by the negative of the

left side of (3.5.57) as we go from t' to t'', while the second sum in the definition of $\mathrm{Gr}_{2A,h}$ changes by the right side of (3.5.57) when we go from t' to t''. These two changes cancel each other out and so, we again have (3.5.56) as desired. $\square$

Combining claims (3) and (4) and noting that ± 1 are regular values of (3.5.48) as well as each of the maps (3.5.49) completes the proof of Theorem 3.1.5.

3.6 Application to the Gopakumar–Vafa invariants

In this section, we will show that the invariant from Theorem 3.1.5 coincides with the Gopakumar–Vafa invariant for genus 0. The comparison is carried out in §3.6.2 and is an application of the discussion of obstruction bundles in §3.6.1. The main result of §3.6.1 (Theorem 3.6.5) is stated in greater generality than is needed for the application in §3.6.2 as it may be of independent interest.

3.6.1 Obstruction bundles and their Euler classes

Let C be a smooth curve of genus g and let $h \geq g$ and $d \geq 2$ be integers such that the integer

$$r = (2h-2) - d(2g-2) \tag{3.6.1}$$

is non-negative. The moduli stack $\overline{\mathcal{M}}_h(C,d)$, consisting of stable holomorphic maps $f : C' \to C$ with the arithmetic genus of C' being h and $f_*[C'] = d[C]$, has complex virtual dimension given by r. Moreover, the non-empty open substack $\mathcal{M}_h(C,d) \subset \overline{\mathcal{M}}_h(C,d)$, consisting of those $[f : C' \to C]$ for which the domain C' is smooth, is a complex orbifold of complex dimension r.

Definition 3.6.1 (Virtual fundamental class)**.** *Define the class*

$$[\overline{\mathcal{M}}_h(C,d)]^{\mathrm{vir}} \in \check{H}_{2r}(\overline{\mathcal{M}}_h(C,d),\mathbb{Q}) \tag{3.6.2}$$

to be the virtual fundamental class of this moduli space as in [Par16, Section 9]*. Note that* $\overline{\mathcal{M}}_h(C,d)$ *is algebraic and so, we do not actually need to distinguish between its Čech and singular (co)homologies.*

Now, let N be a rank 2 complex vector bundle on C with $\deg(N) = 2g-2$. We are interested in $\mathbb{R}$-linear Cauchy–Riemann operators on N because N naturally arises as the normal bundle of J-curves of genus g in a symplectic Calabi–Yau 3-fold. Then, for any $\mathbb{R}$-linear Cauchy–Riemann operator D on N, we have

$$\mathrm{ind}(D) = 2\deg(N) + (2-2g)\cdot\mathrm{rank}(N) = 0 \tag{3.6.3}$$

by the Riemann–Roch theorem. For the Cauchy–Riemann operator D and a point $[f : C' \to C] \in \overline{\mathcal{M}}_h(C,d)$, we can consider the vector spaces

$$\ker(f^*D)' := \mathrm{coker}\ (\ker D \xrightarrow{f^*} \ker(f^*D)) \tag{3.6.4}$$

$$\mathrm{coker}\ (f^*D)' := \mathrm{coker}\ (\mathrm{coker}\ D \xrightarrow{f^*} \mathrm{coker}\ (f^*D)). \tag{3.6.5}$$

Note that these spaces can also be defined as subspaces of $\ker(f^*D)$ and coker (f^*D), when C' is smooth, using Remark 3.2.7 and Lemma 3.2.8.

Definition 3.6.2 (Cokernel/obstruction bundle)**.** *The operator* D *is said to be* (d,h)-rigid *if and only if for all* $[f : C' \to C] \in \overline{\mathcal{M}}_h(C,d)$*, we have*

$$\ker(f^*D)' = 0. \tag{3.6.6}$$

In this case, the assignment

$$[f : C' \to C] \mapsto (\mathcal{N}_{C,D}^{d,h})_f := \text{coker } (f^*D)' \tag{3.6.7}$$

defines a rank $2r$ *vector bundle* $\mathcal{N}_{C,D}^{d,h} \to \overline{\mathcal{M}}_h(C,d)$*, which is called the* cokernel/obstruction bundle*.*

Next, define the canonically oriented *determinant line* of $(f^*D)'$ to be

$$\det(f^*D)' = \det \ker(f^*D)' \otimes (\det \text{coker } (f^*D)')^* \tag{3.6.8}$$

$$= \det(f^*D) \otimes \det(D)^* \tag{3.6.9}$$

where $\det(D)$ and $\det(f^*D)$ are the determinant lines of the Cauchy–Riemann operators D and f^*D, which are canonically oriented by considering the spectral flow of a homotopy between D and a $\mathbb{C}$-linear Cauchy–Riemann operator. If D is (d,h)–rigid, this gives rise to an orientation of the corresponding cokernel bundle. Accordingly, we have a well-defined *virtual Euler number* of the cokernel bundle

$$e_{d,h}(C,D) = \langle [\overline{\mathcal{M}}_h(C,d)]^{\text{vir}}, e(\mathcal{N}_{C,D}^{d,h}) \rangle \in \mathbb{Q}. \tag{3.6.10}$$

More generally, suppose P is a topological space and $\mathcal{D} = \{D_p\}_{p \in P}$ is a family of real Cauchy–Riemann operators

$$D_p : \Omega^0(C,N) \to \Omega^{0,1}(C,N) \tag{3.6.11}$$

parametrized by P. The family $\mathcal{D} = \{D_p\}_{p \in P}$ is called *continuous* if the coefficients of $\{D_p\}_{p \in P}$ (and all their derivatives in the C direction) are continuous as functions on $P \times C$ once we identify the space of real Cauchy–Riemann operators with $\Omega^{0,1}(C, \hom_{\mathbb{R}}(N,N))$ by choosing a reference operator. The notion of smooth family

$\mathcal{D} = \{D_p\}_{p\in P}$ is defined similarly when P is a smooth manifold. It is easy to see that this notion of continuity/smoothness is well-defined. Let $\hat{\Delta} \subset \overline{\mathcal{M}}_h(C,d) \times P$ be the closed substack over which (3.6.6) fails, with $\mathcal{D}$ in place of D. Then, the assignment (3.6.7), with $\mathcal{D}$ in place of D, gives a well-defined oriented rank $2r$ vector bundle

$$\mathcal{N}^{d,h}_{C,\mathcal{D}} \to (\overline{\mathcal{M}}_h(C,d) \times P) \setminus \hat{\Delta} \tag{3.6.12}$$

which has an Euler class $e(\mathcal{N}^{d,h}_{C,\mathcal{D}}) \in \check{H}^{2r}((\overline{\mathcal{M}}_h(C,d) \times P) \setminus \hat{\Delta}, \mathbb{Q})$.

Definition 3.6.3 (Simple wall crossing)**.** *Let $\mathcal{D} = \{D_t\}_{t\in[-1,1]}$ be a smooth family of real Cauchy–Riemann operators on N. We say that $\mathcal{D}$ is a* simple wall crossing *of type (d,h) at $t = 0$ if $\ker D_0 = 0$ and there is a finite set $\Delta \subset \mathcal{M}_h(C,d)$ such that*

(i) *Equation* (3.6.6) *holds on $(\overline{\mathcal{M}}_h(C,d) \times [-1,1]) \setminus (\Delta \times 0)$ and the vector bundle $\mathcal{N}^{d,h}_{C,\mathcal{D}}$ is well-defined here;*

(ii) *for each $[f : C' \to C] \in \Delta$, the linearized map*

$$T_f\mathcal{M}_h(C,d) \oplus \mathbb{R}_t \xrightarrow{\simeq} \mathrm{Hom}(\ker(f^*D)', \mathrm{coker}\ (f^*D)') \tag{3.6.13}$$

defined in Definition 3.2.9 is an isomorphism.

Remark 3.6.4. Note that the set Δ is assumed to be a subset of the locus $\mathcal{M}_h(C,d)$ parametrizing holomorphic maps with smooth domains. This is consistent with the Equation (3.1.11) in Theorem 3.1.10. This assumption is sufficient for the application in Section 3.6.2.

By counting dimensions and indices, it follows that we have

$$\dim_{\mathbb{R}} \ker(f^*D)' = 1 \tag{3.6.14}$$

$$\dim_{\mathbb{R}} \mathrm{coker}\ (f^*D)' = 1 + 2r \tag{3.6.15}$$

over the finite set Δ. Both sides of (3.6.13) are canonically oriented by combining the discussion surrounding equation (3.4.36) with the orientation on $\det(f^*D)'$. Thus, to each $[f : C' \to C] \in \Delta$, we can assign a well-defined number

$$\mathrm{sgn}(\mathcal{D}, [f : C' \to C]) \in \{+1, -1\}. \tag{3.6.16}$$

which is the sign of the isomorphism (3.6.13).

We want to compute the change in the virtual Euler number $e_{d,h}(C, D_t)$ as we go from $t < 0$ to $t > 0$ for a simple wall crossing $\mathcal{D}$ in terms of isotropy groups of the points in Δ and the signs $\{\mathrm{sgn}(\mathcal{D}, p)\}_{p\in\Delta}$. Our main result is the following.

Theorem 3.6.5 (Change in Euler number)**.** *In the situation of Definition 3.6.3,*

$$[-1, 1] \setminus \{0\} \to \mathbb{Q} \tag{3.6.17}$$

$$t \mapsto e_{d,h}(C, D_t) \tag{3.6.18}$$

is a locally constant function. Let $e_{d,h}(C, D_\pm)$ denote the value $e_{d,h}(C, D_t)$ for $\pm t > 0$. Then, we have the identity

$$e_{d,h}(C, D_+) - e_{d,h}(C, D_-) = \sum_{p\in\Delta} \frac{2 \cdot \mathrm{sgn}(\mathcal{D}, p)}{|\mathrm{Aut}(p)|}. \tag{3.6.19}$$

The local constancy assertion is obvious. The remaining assertion will be proved in the next two lemmas. We continue to use the notation from Definition 3.6.3.

Lemma 3.6.6 (Local to global)**.** *For each point $p \in \Delta$, choose an open neighborhood*

$$V_p \Subset \mathcal{M}_h(C, d) \times (-1, 1) \tag{3.6.20}$$

whose closure has smooth orbifold boundary ∂V_p. Endowing ∂V_p with the boundary orientation from V_p and assuming that the closures of $\{V_p\}_{p\in\Delta}$ are pairwise disjoint,

we have the identity

$$e_{d,h}(C, D_+) - e_{d,h}(C, D_-) = \sum_{p \in \Delta} \langle [\partial V_p], e(\mathcal{N}_{C,\mathcal{D}}^{d,h}) \rangle \tag{3.6.21}$$

where $[\partial V_p]$ *denotes the standard orbifold fundamental class for each* $p \in \Delta$.

Proof. Introduce the notation $\mathcal{M}_I := \overline{\mathcal{M}}_h(C, d) \times [-1, 1]$. It will be enough to show that the identity

$$[\overline{\mathcal{M}}_h(C, d)]^{\text{vir}} \times \{+1\} - [\overline{\mathcal{M}}_h(C, d)]^{\text{vir}} \times \{-1\} = \sum_{p \in \Delta} [\partial V_p] \tag{3.6.22}$$

holds in $H_*(\mathcal{M}_I \setminus \bigcup_{p \in \Delta} V_p; \mathbb{Q})$. This will follow from [Par16, Lemmas 5.2.4, 5.2.5 and 5.2.6] applied to the space

$$\mathcal{M}_I \setminus \bigcup_{p \in \Delta} V_p \tag{3.6.23}$$

provided we show that it has an implicit atlas with boundary which, when restricted to the boundary, gives (a) the usual atlas (from section [Par16, Section 9]) on $[\overline{\mathcal{M}}_h(C, d)]^{\text{vir}} \times \{\pm 1\}$ and (b) the standard orbifold atlas on each V_p. Constructing such an atlas is straightforward. Indeed, choose open neighborhoods

$$V_p \Subset W_p \Subset W_p' \Subset \mathcal{M}_h(C, d) \times (-1, 1) \tag{3.6.24}$$

with $\{W_p'\}_{p \in \Delta}$ still being pairwise disjoint. Endow each $W_p' \setminus V_p$ with the obvious orbifold-with-boundary atlas (with trivial obstruction spaces) as in [Par16, Section 2.1.2] and endow the open subset

$$\mathcal{M}_I \setminus \bigcup_{p \in \Delta} \overline{W}_p \tag{3.6.25}$$

with the restriction of the product atlas (see [Par16, Definition 6.3.1]) on $\mathcal{M}_I$. These $1+|\Delta|$ open sets cover the moduli space in question. Over the open overlaps $W_p'\backslash\overline{W}_p \subset \mathcal{M}_h(C,d)\times(-1,1)$ – note that there are no triple overlaps – both implicit atlases are equivalent (in the sense of [Par16, Remark 2.1.5]). This is because the standard orbifold atlas on $\mathcal{M}_h(C,d)$ is equivalent to the restriction of the implicit atlas on $\overline{\mathcal{M}}_h(C,d)$. Thus, we can combine this cover into an implicit atlas on the space (3.6.23) with the desired properties. □

Lemma 3.6.7 (Local contribution)**.** *In the situation of Lemma 3.6.6, the neighborhoods V_p can be chosen such that their closures are pairwise disjoint and we have*

$$\langle [\partial V_p], e(\mathcal{N}_{C,\mathcal{D}}^{d,h})\rangle = \frac{2\cdot \operatorname{sgn}(\mathcal{D},p)}{|\operatorname{Aut}(p)|} \tag{3.6.26}$$

for each $p\in\Delta$.

Proof. Fix $p=[f:C'\to C]\in\Delta\subset\mathcal{M}_h(C,d)$. Since $\mathcal{M}_h(C,d)$ is a smooth orbifold, we can identify a neighborhood U_p of p diffeomorphically with the orbifold $[\mathbb{V}/\Gamma]$, where $\Gamma=\operatorname{Aut}(C,f)$ is the isotropy group of p with its natural complex representation on $\mathbb{V}=T_f\mathcal{M}_h(C,d)$ and the point p is identified with $[0]$. Without loss of generality, we can assume that this neighborhood meets no other points of Δ.

Assume, for the sake of concreteness, that $\operatorname{sgn}(\mathcal{D},p)=+1$. This can be achieved by reparametrizing the interval $[-1,1]$ in reverse if necessary. The local chart $[\mathbb{V}/\Gamma]$ gives a map $\mathbb{V}\xrightarrow{\nu}\mathcal{M}_h(C,d)$ which, in turn, yields the following Γ-equivariant commutative diagram (with the square being a fiber square).

$$\begin{array}{ccccc} C' & \overset{f}{\overset{\curvearrowright}{\hookrightarrow}} & \mathcal{C}' & \xrightarrow{\text{ev}} & C \\ \downarrow & & \downarrow{\scriptstyle\pi} & & \\ * & \xrightarrow{0} & \mathbb{V} & & \end{array}$$

Here, π is a Γ-equivariant smooth family of curves and ev is a Γ-invariant smooth map which is holomorphic on each fibre of π. The maps π and ev together with the Γ-action

realize the open embedding $[\mathbb{V}/\Gamma] \hookrightarrow \mathcal{M}_h(C,d)$. We can then define $V_p = [B_{\mathbb{V}}/\Gamma]$, where $B_{\mathbb{V}}$ is the ball bounded by $S_{\mathbb{V}}$, and conclude by Lemma 3.6.8 and noting that $\langle [S_{\mathbb{V}}], e(TS_{\mathbb{V}})\rangle = \chi(S_{\mathbb{V}}) = 2$. It is useful to note that, for this last argument to work, we *do not* require the isomorphism (3.6.27) to be Γ-equivariant. □

Lemma 3.6.8. *Use the same notations as in Lemma 3.6.7. For any sufficiently small sphere $S_{\mathbb{V}} \subset \mathbb{V} \times [-1,1]$ centered at 0 (defined with respect to the orthogonal direct sum of a Γ-invariant Hermitian inner product on $\mathbb{V}$ and the standard inner product on $\mathbb{R}$), we have an isomorphism*

$$(\tilde{\nu}^* \mathcal{N}^{d,h}_{C,\mathcal{D}})|_{S_{\mathbb{V}}} \simeq TS_{\mathbb{V}} \tag{3.6.27}$$

of oriented vector bundles on $S_{\mathbb{V}}$, where $\tilde{\nu} : \mathbb{V} \times [-1,1] \to \mathcal{M}_h(C,d) \times [-1,1]$ is given by $\tilde{\nu} = \nu \times id$.

Proof. Let us first make the following auxiliary choices.

1. A Hermitian metric $(\cdot,\cdot)$ and a compatible $\mathbb{C}$-linear connection ∇ on N.

2. A Riemannian metric (and therefore, an area form) on C' whose conformal class coincides with the one on C' induced by j.

3. A smooth trivialization $\mathcal{C}' = C' \times \mathbb{V}$ of the bundle π. Let $\tilde{f} : C' \times \mathbb{V} \to C$ denote the map induced by ev, and set $f_v := \tilde{f}(\cdot, v) : C' \to C$ for $v \in \mathbb{V}$. Also, let $\tilde{j} = (j_v)_{v \in V}$ denote the induced family of almost complex structures on C'.

4. A smooth family of isomorphisms

$$I_v : (T_{C'}, j_0) \xrightarrow{\simeq} (T_{C'}, j_v). \tag{3.6.28}$$

We now have the following smooth family of Cauchy–Riemann operators on N parametrized

by $(v,t) \in \mathbb{V} \times [-1,1]$

$$D_{v,t} : L^2_1(C', f^*N) \to L^2(C', \Lambda^{0,1} T^*_{C',j_0} \otimes_{\mathbb{C}} f^*N) \tag{3.6.29}$$

defined by pulling back the operator D_t on f_v^*N to an operator on f^*N by using the ∇-parallel transport along the paths $\tau \in [0,1] \mapsto f_{\tau v}(z) \in C$ for each $z \in C'$ and the isomorphism I_v^{-1}. Obviously, we have $D_{0,0} = f^*D_0$ on smooth sections. Since $\ker D_0 = \text{coker } D_0 = 0$ (and this will continue to hold for small t), we have an obvious isomorphism

$$\text{coker } D_{v,t} \simeq (\tilde{\nu}^* \mathcal{N}^{d,h}_{C,\mathcal{D}})_{v,t} \tag{3.6.30}$$

for small $(v,t) \neq (0,0)$. Orient $A = \ker(f^*D_0)'$ and $B = \text{coker } (f^*D_0)'$ so that the induced orientation on $\det(f^*D_0)'$ is the canonical one from Definition 3.6.2. Define the bounded linear map

$$R : \text{im } D_{0,0} \to (\ker D_{0,0})^{\perp} \tag{3.6.31}$$

to be the inverse of $D_{0,0}|_{(\ker D_{0,0})^{\perp}}$, with $(\ker D_{0,0})^{\perp} \subset L^2_1$ denoting the L^2 orthogonal complement of $\ker D_{0,0}$. Let $\Pi_{\text{im}} : L^2 \to \text{im } D_{0,0}$ be the orthogonal projection and let Π_{coker} be the quotient projection onto the cokernel. For small (v,t), the quantity $\|1 - \Pi_{\text{im}} D_{v,t} R\|$ is small on $\text{im } D_{0,0}$ and so, we can define the map

$$R_{v,t} = R(\Pi_{\text{im}} D_{v,t} R)^{-1} : \text{im } D_{0,0} \to (\ker D_{0,0})^{\perp} \tag{3.6.32}$$

which is the inverse of $\Pi_{\text{im}} D_{v,t}|_{(\ker D_{0,0})^{\perp}}$. It is immediate that $\ker(\Pi_{\text{im}} D_{v,t})$ is given by the graph of the map $-R_{v,t}\Pi_{\text{im}} D_{v,t} : \ker D_{0,0} \to (\ker D_{0,0})^{\perp}$. Define the smooth

family of linear maps

$$L_{v,t} : \ker D_{0,0} \to \text{coker } D_{0,0} \tag{3.6.33}$$

$$L_{v,t} = \Pi_{\text{coker}} D_{v,t}(1 - R_{v,t}\Pi_{\text{im}} D_{v,t}) \tag{3.6.34}$$

for small (v, t). We now observe that that the natural maps $\ker L_{v,t} \xrightarrow{\simeq} \ker D_{v,t}$ and coker $L_{v,t} \xrightarrow{\simeq}$ coker $D_{v,t}$ are isomorphisms. It follows (using the identifications $A = \ker D_{0,0}$ and $B =$ coker $D_{0,0}$) that the oriented vector bundle $\tilde{\nu}^*\mathcal{N}^{d,h}_{C,\mathcal{D}}$ is given, on a punctured neighborhood of $(0, 0)$, by the assignment

$$(v, t) \mapsto \text{coker } (L_{v,t} : A \to B) \tag{3.6.35}$$

with the bundle orientation derived from the orientations already chosen on A and B. Note that $L_{0,0} = 0$ and that the first order Taylor expansion of $L_{v,t}$ about $(0, 0)$ is given by the positively oriented linear isomorphism

$$\mathbb{V} \oplus \mathbb{R}_t \xrightarrow{\simeq} \text{Hom}(A, B) \tag{3.6.36}$$

from (3.6.13). Fixing a unitary Γ-invariant metric on $\mathbb{V}$ and restricting to any sufficiently small sphere $S_{\mathbb{V}} \subset \mathbb{V} \times [-1, 1]$ centred at the origin, we can discard (i.e. homotope away) the higher order terms from the Taylor expansion of $L_{v,t}$. It follows that the fiber of the vector bundle at a point $(v_0, t_0) \in S_{\mathbb{V}}$ is identified with the quotient of the oriented vector space $\mathbb{V} \oplus \mathbb{R}_t$ by the oriented 1-dimensional space spanned by (v_0, t_0). But this is simply the standard description of the (oriented) tangent bundle of the sphere $S_{\mathbb{V}} \subset \mathbb{V} \oplus \mathbb{R}_t$, once we replace "quotient" by "orthogonal complement" and (3.6.27) now follows. □

Note that Lemma 3.6.8 holds for any d and h. It is ultimately a consequence of the isomorphism (3.6.13) via the finite dimensional reduction argument presented

above.

Remark 3.6.9. A coordinate invariant way of summarizing the above proof is the following. The bundle $\mathcal{N}^{d,h}_{C,\mathcal{D}}$, initially defined on $(\overline{\mathcal{M}}_h(C,d)\times[-1,1])\setminus(\Delta\times 0)$, extends to the real oriented blow-up of $\overline{\mathcal{M}}_h(C,d)\times[-1,1]$ along the smooth points $\Delta\times 0$ as a vector bundle $\overline{\mathcal{N}}^{d,h}_{C,\mathcal{D}}$. Moreover, for each point $p\in\Delta$, we have a natural vector bundle isomorphism with orientation $\mathrm{sgn}(\mathcal{D},p)$

$$\overline{\mathcal{N}}^{d,h}_{C,\mathcal{D}}|_{\mathbb{S}_p} = T\mathbb{S}_p \tag{3.6.37}$$

over $\mathbb{S}_p$, the exceptional boundary component created by blowing up at p. Now, a direct cobordism argument as in Lemma 3.6.6, applied to the real oriented blow-up $\mathrm{Bl}^+_{\Delta\times 0}(\overline{\mathcal{M}}_h(C,d)\times[-1,1])$ gives (3.6.19).

3.6.2 Comparision to Gopakumar–Vafa invariants

In this subsection, we identify Taubes' Gromov invariant in dimension 6 (Equation (3.1.5)) with the Gopakumar–Vafa invariants in the case of genus 0 and class $2A$ for $A\in H_2(X,\mathbb{Z})$ primitive. We omit the underlying Calabi–Yau 3-fold (X,ω) from the notation.

Theorem 3.6.10. *The invariant* $\mathrm{Gr}_{2A,0}$ *satisfies the Gopakumar–Vafa formula. More precisely, define, for super-rigid* J*,*

$$\mathrm{GW}_{A,0} := \mathrm{Gr}_{A,0} := \sum_{C\colon A,0}\mathrm{sgn}(D^N_{C,J}) \tag{3.6.38}$$

$$\mathrm{GW}_{2A,0} := \sum_{C'\colon 2A,0}\mathrm{sgn}(D^N_{C',J}) + \sum_{C\colon A,0}\mathrm{sgn}(D^N_{C,J})e_{2,0}(C,D^N_{C,J}) \tag{3.6.39}$$

where $e_{2,0}$ *denotes the virtual Euler number as defined in* (3.6.10) *and the summations follow the notation of* (3.1.5)*. Then, we have the following relation between* GW *and*

Gr.

$$\mathrm{GW}_{2A,0} = \mathrm{Gr}_{2A,0} + \frac{1}{2^3} \cdot \mathrm{Gr}_{A,0}. \tag{3.6.40}$$

Proof. From the definitions, it is clear that it will suffice to prove that

$$e_{2,0}(C, D^N_{C,J}) - \frac{1}{2^3} = \mathrm{w}_{2,0}(D^N_{C,J}). \tag{3.6.41}$$

Let $\bar{\partial}$ denote an operator on N_C which is conjugate to the standard operator on $\mathcal{O}(-1)^{\oplus 2} \to \mathbb{P}^1$. The Aspinwall–Morrison formula (see, for example [Voi96]) shows that we have the identity

$$e_{2,0}(C, \bar{\partial}) = \frac{1}{2^3}. \tag{3.6.42}$$

Thus, we are left to prove that the difference of Euler numbers is given by the linear wall crossing invariant $\mathrm{w}_{2,0}$. But this is immediate from the Definition 3.5.17 and Theorem 3.6.5 once we choose a path $\gamma \in \mathcal{P}^{\pitchfork}(D^N_{C,J}, \bar{\partial})$ in $\mathcal{CR}_{\mathbb{R}}(N_C)$ to compute $\mathrm{w}_{2,0}$. □

Chapter 4

On the Rouquier dimension of wrapped Fukaya categories and a conjecture of Orlov

4.1 Introduction

4.1.1 Overview

In his seminal work [Rou08], Rouquier introduced a notion of dimension for triangulated categories. Here is a sketch of the definition: if $\mathcal{T}$ is a triangulated category and G is a (split-)generator, then G is said to generate $\mathcal{T}$ in time t if every object of $\mathcal{T}$ can be built from G by taking finite direct sums, shifts, summands, and *at most* t cones. The Rouquier dimension of $\mathcal{T}$ is by definition the minimal generation time over all generators of $\mathcal{T}$.

In symplectic topology, an important class of triangulated categories are derived Fukaya categories in their various guises. We will be mainly interested in wrapped Fukaya categories of Liouville manifolds. More precisely, given a Liouville manifold (X, λ), its wrapped Fukaya category $\mathcal{W}(X)$ is an A_∞ category and as such, it canon-

ically determines a triangulated category $H^0(\operatorname{Perf} \mathcal{W}(X))$. The Rouquier dimension of this triangulated category is a numerical invariant of (X, λ) which will be the focus of this chapter.

Our applications will be twofold: first, to a problem in algebraic geometry; second, to quantitative questions in symplectic topology. We discuss each of these two classes of applications in turn in the remainder of this introduction.

4.1.2 An application of symplectic topology to algebraic geometry

Motivation

In classical algebraic geometry, one is typically interested in spaces (varieties, schemes, stacks, etc.) which are locally modeled on the spectrum of a commutative ring. According to the philosophy of non-commutative geometry, it is often profitable to regard A_∞ categories as "non-commutative" spaces. One can then attempt to generalize classical algebraic geometry to the setting of A_∞ category theory.

For example, there are standard notions smoothness, properness, dimension, etc. in classical algebraic geometry. It turns out that there also exist non-commutative analogs of some of these notions. Namely: an A_∞ category is defined to be smooth if its diagonal bimodule is perfect; it is defined to be proper if its diagonal bimodule is proper. The dimension of an A_∞ category is simply its Rouquier dimension.[1]

If Y is an ordinary algebraic variety, then one can associate to it the A_∞ category $D^b \operatorname{Coh}(Y)$. One typically regards $D^b \operatorname{Coh}(Y)$ as the non-commutative partner of the classical object Y. It is then natural to ask: do the non-commutative notions of smoothness, properness, dimension coincide with their classical counterparts?

As one might hope, it is known that $D^b \operatorname{Coh}(Y)$ is smooth and proper if and only

[1]Note that there are other reasonable notions of dimension discussed e.g. in [EL21], although the Rouquier dimension seems the most established in the literature.

if Y is smooth and proper in the classical sense. Concerning dimension, there is the following well-known conjecture or Orlov:

Conjecture 4.1.1 (Orlov; see Conjecture 10 in [Orl09b]). Let Y be a smooth quasi-projective scheme of dimension n. Then

$$\operatorname{Rdim} D^b \operatorname{Coh}(Y) = n. \tag{4.1.1}$$

The assumption that Y is smooth is necessary; see 4.2.3. Orlov's conjecture is in general wide open. It has been only established for the following rather restrictive classes of examples, thanks to the combined work of many authors (to the best of our knowledge, this list is exhaustive).

1. Smooth affine varieties, projective spaces and smooth projective quadrics [Rou08, Prop. 7.18, Ex. 7.7, 7.8];

2. All smooth projective curves [Orl09b, Thm. 6];

3. The following schemes defined over an algebraically-closed field of characteristic 0: del Pezzo surfaces with $\operatorname{rk} \operatorname{Pic}(Y) \leq 7$, Fano threefolds of type V_5 and V_{22}, toric surfaces with nef anti-canonical divisor, toric Deligne–Mumford stacks defined over $\mathbb{C}$ of dimension no more than 2 or Picard number no more than 2, Hirzebruch surfaces [BF12];

4. The product of varieties from the previous item [Yan16];

5. The product of two Fermat elliptic curves with the Fermat $K3$ surface [BFK14, Thm. 1.6];

6. Smooth toric Fano varieties of dimension 3 or 4 [BDM19, Prop. 2.15, 2.16];

7. Certain blow-ups of the projective space [Pir19, Thm. 4.1];

8. Some weak del Pezzo surfaces [EXZ21, Cor. 7.5];

We remark that Rouquier proved that

$$n \leq \operatorname{Rdim} D^b \operatorname{Coh}(Y) \leq 2n, \tag{4.1.2}$$

so the content of Orlov's conjecture lies in improving the upper bound. See 4.2.6 for further discussion and references.

Results

We introduce new methods from the flexible side of symplectic topology to study Orlov's conjecture. More precisely, we prove:

Theorem 4.1.1. Suppose that Y is a variety over $\mathbb{C}$ of complex dimension n which admits a homological mirror given by a polarizable Weinstein pair. Then $\operatorname{Rdim} D^b \operatorname{Coh}(Y) = n$ if $n \leq 3$. For $n \geq 4$, we have $\operatorname{Rdim} D^b \operatorname{Coh}(Y) \leq 2n - 3$.

(A Weinstein pair (X, A) is said to be polarizable if X admits a global Lagrangian plane field ξ.) 4.1.1 implies the following new cases of Orlov's conjecture.

Corollary 4.1.2. 4.1.1 is true for all 3-dimensional toric varieties over $\mathbb{C}$ and the log Calabi–Yau surfaces in 4.4.4.

In fact, 4.1.1 shows that (4.1.1) holds for a large class of (possibly singular) schemes and stacks over $\mathbb{C}$, see 4.4.4 for the list. It also gives a new proof of some of the known cases of Orlov's conjecture listed above. Finally, 4.1.1 improves Rouquier's upper bound (4.1.2) for $\operatorname{Rdim} D^b \operatorname{Coh}(-)$ for many examples in dimensions $n \geq 4$ (including all smooth toric varieties and DM stacks).

Methods: symplectic flexibility

Although Orlov's conjecture is a problem in algebraic geometry, our proof of 4.1.1 is entirely symplectic. As mentioned in (4.1.2), it is already known that $\operatorname{Rdim} D^b \operatorname{Coh}(Y) \geq$

n. Hence, we only need to prove the claimed upper bounds. To do this, we apply homological mirror symmetry. This transforms 4.1.1 into a statement about the Rouquier dimension of wrapped Fukaya categories of Weinstein pairs; namely, it is enough to prove that the Rouquier dimension of a (polarizable) Weinstein pair of real dimension $2n$ is bounded above by n when $n \leq 3$ and by $2n-3$ when $n \geq 4$. This is now purely a problem in symplectic topology which can be approached using symplectic tools.

According to the recent work of Ganatra–Pardon–Shende [GPS18b], wrapped Fukaya categories of Liouville manifolds satisfy "Weinstein sectorial descent". Essentially, this means that if we (suitably) cover a Liouville manifold X by Weinstein sectors $\{X_\alpha\}_{\alpha \in I}$, then $\mathcal{W}(X)$ can be expressed as the homotopy colimit of the wrapped Fukaya categories of the form $\mathcal{W}(\cap_{\alpha \in J} X_\alpha)$ (where J ranges over all non-empty subsets of I and the homotopy colimit is indexed by the obvious diagram of inclusions).

It is not hard to show that the Rouquier dimension of a homotopy colimit can be bounded from above by the Rouquier dimension of its individual pieces. This bound depends on (a) the "depth" of the diagram (i.e. the longest chain of arrows) and (b) the Rouquier dimension of the individual pieces; see 4.2.12 for a precise statement. This suggests a possible strategy for upper bounding the Rouquier dimension of a Weinstein manifold: construct a sectorial cover of "small" depth, such that each piece in the cover has "small" Rouquier dimension.

Let us first implement this strategy in the special case where $X = T^*M$, for M a closed n-manifold. We construct a sectorial cover as follows. First, triangulate M. Then set $X_\alpha := T^* \operatorname{star}(v_\alpha)$ where v_α ranges over the vertices of the triangulation. Observe that the intersection of any collection of stars in a triangulation is either empty, or is the star of some higher dimensional cell. From this, it follows that the cover has depth $n+1$. It also follows (after appropriate smoothing) that $\bigcap_{\alpha_1,\ldots,\alpha_n} T^* \operatorname{star}(v_{\alpha_i}) \simeq T^* D^n$, and hence the Rouquier dimension of the overlaps is

zero. Plugging these numbers into 4.2.12, one finds that $\operatorname{Rdim} \mathcal{W}(T^*M) \leq n$. This upper bound is sharp in general.[2]

We would like to extend this argument to an arbitrary Weinstein manifold X. Unfortunately, the naive generalization *does not work*. If X is an arbitrary Weinstein manifold, its skeleton may be a highly singular object. Even assuming the skeleton can be triangulated (which is not automatic), it is completely unclear how to control the Rouquier dimension of the resulting sectorial cover. It is also unclear how to convert a cover of the skeleton into a sectorial cover of its thickening (in the cotangent bundle case, there is a canonical Morse–Bott form which allows one to easily perform this conversion).

To overcome these difficulties, we will take advantage of recent advances from the flexible side of symplectic geometry. Namely, the arborealization theorem of Álvarez-Gavela–Eliashberg–Nadler [AGEN21, AGEN19, AGEN20] provides an h-principle for simplifying the singularities of the skeleton. The upshot is that, under rather mild topological assumptions, one can homotope a Weinstein manifold so that the singularities of its skeleton are of "arboreal" type. Such homotopies do not affect the wrapped Fukaya category.

Arboreal singularities are indexed by finite rooted trees. Morally, if $\mathcal{S}_\sigma$ is a Liouville sector obtained by thickening an arboreal singularity σ, one expects that $\mathcal{W}(\mathcal{S}_\sigma)$ is Morita equivalent to $\operatorname{Rep}(T_\sigma)$, the derived representation category of the corresponding tree T_σ. Moreover, it is known that $\operatorname{Rdim} \operatorname{Rep}(T_\sigma) = 0$ if T_σ is of Dynkin type (i.e. the underlying graph is a Dynkin diagram) and $\operatorname{Rdim} \operatorname{Rep}(T_\sigma) = 1$ otherwise. This suggests that sectorial thickenings of arboreal singularities are a good substitute for T^*D^n,

Arboreal skeleta admit a canonical stratification by "singularity-order". We prove in 4.4.2 that this stratification is Whitney. By a well-known theorem of Goresky, any

[2]For example, we will see that $\operatorname{Rdim} \mathcal{W}(T^*\mathbb{T}^n) = n$.

Whitney stratification can be refined to a (possibly non-Whitney!) triangulation. This suggests the following updated strategy: given a (polarizable) Weinstein manifold X of dimension $2n$, apply a homotopy to make the skeleton arboreal. It therefore admits a triangulation. Now construct a sectorial cover by taking arboreal thickenings of the stars of the vertices, like for cotangent bundles.

Unfortunately, there are significant obstructions to making this strategy rigorous. To begin with, one would need to develop a good theory of arboreal sectors (which should be closed under intersections), as well as a procedure for constructing a "sectorial thickening" of an arboreal singularity. Moreover, one would need a method for construct arboreal sectors from stars of (possibly non-Whitney) triangulations. Due to this difficulties, it is unclear to us whether the above strategy can feasibly be implemented in the framework of Fukaya categories.

Instead, we work in the framework of microlocal sheaf theory. The details of this approach are carried out in 4.4. After reviewing some basic properties of microlocal sheaves, we study the stratified topology of arboreal skeleta, and use this to construct a suitable triangulation. We then carry through a microlocal sheaf analog of the above covering argument to obtain an upper bound on the Rouquier dimension of the category of microlocal sheaves. Finally, we apply the main result of [GPS18a] to pass back to wrapped Fukaya categories.

Remark 4.1.1 (A digression on motives). As a byproduct of our method, we obtain some additional conclusions which are unrelated to the Rouquier dimension but may be of independent interest. For example, if (X, A) is an polarizable Weinstein pair such that $\mathcal{W}(X, A)$ is proper (it is a fortiori smooth), then we find that:

- $HH_*(\mathcal{W}(X, A)), K_*(\mathcal{W}(X, A))$ are concentrated in degree zero;
- Perf $\mathcal{W}(X, A)$ satisfies the *non-commutative Weil conjecture* introduced by Tabuada (following proposals of Kontsevich) in [Tab18].

Both of these statements follow from a more general result about non-commutative motives, which is itself an easy consequence of the arborealization method described above. We refer the reader to 4.4.5 for more.

4.1.3 Quantitative problems in symplectic topology

Motivation

As motivation for the geometrically-inclined symplectic topologist, we state some quantitative questions of rather classical flavor.

Question 4.1.2. Let (X, λ) be a Weinstein[3] manifold with skeleton $\mathfrak{c}_X$. If $\phi : X \to X$ is a generic compactly-supported Hamiltonian diffeomorphism, what is the minimal possible number of intersection points of $\phi(\mathfrak{c}_X)$ with $\mathfrak{c}_X$?

If $(X, \lambda) = (T^*M, pdq)$ for a closed manifold M, then $\mathfrak{c}_X = 0_M$ and 4.1.2 is classical. As is now well-known, it follows from the existence of Floer cohomology (or finite dimensional analogs) that $|\mathfrak{c}_X \cap \phi(\mathfrak{c}_X)|$ is bounded from below by the sum of the Betti numbers of M.

In general, however, the skeleton of a Weinstein manifold is a highly singular object for which there is no available notion of Floer cohomology (and it seems rather unlikely that such a notion exists). Note that *even for the cotangent bundle of a closed manifold*, equipped with a Weinstein structure homotopic to the standard Morse–Bott one, the skeleton is typically singular.

To the best of our knowledge, the only systematic result in the literature addressing 4.1.2 is the lower bound $|\mathfrak{c}_X \cap \phi(\mathfrak{c}_X)| \geq 1$, which holds whenever the Rabinowitz–Floer homology $RFH(X)$ is non-zero [CF09]. The nonvanishing of $RFH(-)$ for Weinstein manifolds is now known to be equivalent to the nonvanishing of $SH(-)$ and $\mathcal{W}(-)$; see [Rit13, Sec. 1.6].

[3]Recall that a Liouville manifold is said to be *Weinstein* if the Liouville vector field is gradient-like with respect to a proper Morse function.

Question 4.1.3. Let X be a Liouville manifold, considered up to homotopy. What is the minimal possible number of critical points of a Lefschetz fibration $f : X \to \mathbb{C}$ with Weinstein fibers?

Let us denote this number by $\mathrm{Lef}_w(X) \in \mathbb{N} \cup \{\infty\}$. According to work of Giroux–Pardon [GP17], any Weinstein manifold admits a Lefschetz fibration with Weinstein fibers, so $\mathrm{Lef}_w(X) < \infty$ iff X is Weinstein (up to homotopy). It is also clear that $\mathrm{Lef}_w(X) = 0$ iff X is subcritical. For topological reasons, one has the lower bound $\mathrm{Lef}_w(X) \geq \mathrm{rk}\, H_n(X;\mathbb{Z})$ when X has real dimension $2n$. These are the only lower bounds in the literature that we are aware of.

The real analog of $\mathrm{Lef}_w(X)$ was considered by Abouzaid–Seidel [AS10b] under the name "complexity". By definition, the complexity of a Weinstein manifold X, which we denote by $\mathrm{WMor}(X)$, is the minimal number of critical points of a Weinstein Morse function on X, where X is considered up to Weinstein homotopy.[4] Surprisingly, Lazarev [Laz20] showed that for Weinstein manifolds having non-zero middle homology, one has $\mathrm{WMor}(X) = \mathrm{Mor}(X)$, where $\mathrm{Mor}(X)$ denotes the minimal number of critical points of a *smooth* Morse function on X. Does a similar phenomenon hold for $\mathrm{Lef}_w(X)$?

Results

We give partial answers to the above questions. As we will explain, our results appear to be the first to go beyond the basic flexible/rigid dichotomy.

We first consider 4.1.2. To this end, suppose that (X, λ) be a Weinstein manifold with properly embedded cocores (this is a generic condition) and nonzero wrapped Fukaya category. Let $R \subseteq SH^0(X)$ be a subalgebra which is of finite type over a field k, and suppose that there exists a (split-)generating Lagrangian $K \in \mathcal{W}(X)$ such

[4] A *Weinstein Morse* function on (X, λ) is a proper Morse function with respect to which the Liouville vector field is gradient-like. Here λ is a particular choice of Liouville structure on X of the given homotopy class.

that $HW^\bullet(K,K)$ is a Noetherian R-module, where R acts via the closed-open string map.

Theorem 4.1.4 (= 4.6.6 + 4.5.1). Under the above assumptions, if $\phi : X \to X$ is a generic compactly-supported Hamiltonian symplectomorphism, then:

$$|\mathfrak{c}_X \cap \phi(\mathfrak{c}_X)| \geq \dim_R HW^\bullet(K,K) + 1. \tag{4.1.3}$$

Here $\dim_R(-)$ denotes the *Krull* dimension of an R-module. Both $\mathcal{W}(X)$ and $SH^\bullet(X)$ are defined over k and are, say, $\mathbb{Z}/2$-graded. Here are some examples illustrating applications of 4.1.4.

Example 4.1.2. For (T^*S^n, λ), where $n \geq 2$ and λ is homotopic to the standard Morse–Bott form pdq, we will see that $\operatorname{Rdim} \mathcal{W}(T^*S^n) = 1$ (say with $\mathbb{Z}/2$-gradings and $\mathbb{Q}$-coefficients) and hence $|\mathfrak{c}_X \cap \phi(\mathfrak{c}_X)| \geq 2$. This lower bound is of course sharp.

Example 4.1.3. Let G be a compact, simply-connected Lie group. Suppose that (T^*G, λ) is a Weinstein structure which is homotopic to (T^*G, pdq). Then we will see that $\operatorname{Rdim} \mathcal{W}(T^*G) \geq \operatorname{rank} G$ and hence $|\mathfrak{c}_X \cap \phi(\mathfrak{c}_X)| \geq \operatorname{rank} G + 1$.

Example 4.1.4. Let us consider the case where $2c_1(X) = 0$ so that both $SH^\bullet(X)$ and $\mathcal{W}(X)$ are $\mathbb{Z}$-graded. Following [Pom21], an object $L_0 \in \mathcal{W}(X)$ is called a *homological section* if the restriction of the closed-open map $\mathcal{CO}$ to the degree 0 part of $SH^\bullet(X)$

$$\mathcal{CO} : SH^0(X) \to HW^\bullet(L_0, L_0) \tag{4.1.4}$$

defines an isomorphism of algebras. If $\operatorname{Spec}(SH^0(X))$ is smooth, [Pom21, Cor. 1.4] tells us L_0 is a split-generator. Then we can conclude

$$|\mathfrak{c}_X \cap \phi(\mathfrak{c}_X)| \geq \operatorname{Rdim} \mathcal{W}(X) + 1 \geq \dim SH^0(X) + 1. \tag{4.1.5}$$

In practice, a large class of such Weinstein manifolds arise from the complement of

the anti-canonical divisor of log Calabi–Yau varieties.

We now turn to 4.1.3. We have the following result:

Theorem 4.1.5 (= 4.6.5 + 4.5.1). If $R \subseteq SH^0(X)$ and $K \in \mathcal{W}(X)$ are as in 4.1.4, then

$$\mathrm{Lef}_w(X) \geq \dim_R HW^\bullet(K,K) + 1. \tag{4.1.6}$$

Example 4.1.5. We have $\mathrm{Lef}_w(T^*S^n) = 2$ for $n \geq 2$. If G is a simply-connected compact Lie group, then $\mathrm{Lef}_w(T^*G) \geq \operatorname{rank} G + 1$. If X admits a homological section as in 4.1.4, then $\mathrm{Lef}_w(X) \geq \dim SH^0(X) + 1$.

Example 4.1.6. If $\mathcal{W}(X)$ is $\mathbb{Z}$-graded and $L \in \mathcal{W}(X)$ is diffeomorphic to a n-torus, then L is a *point-like object*. By work of Elagin–Lunts [EL18], the existence of such an L implies that $\operatorname{Rdim} \mathcal{W}(X) \geq n$ (see 4.3.8). Hence $\mathrm{Lef}_w(X) \geq n + 1$.

As an application of 4.1.5, we prove that the natural complex analog of Lazarev's previously mentioned result "$\mathrm{WMor}(X) = \mathrm{Mor}(X)$" *fails*. This is a consequence of the following corollary.

Corollary 4.1.3 (= 4.6.1). There exists a Liouville manifold $T^*S^3_{\mathrm{exotic}} = (T^*S^3, \lambda)$ which is formally isotopic to (T^*S^3, pdq), but where

$$\mathrm{Lef}_w(T^*S^3) \neq \mathrm{Lef}_w(T^*S^3_{\mathrm{exotic}}). \tag{4.1.7}$$

The term "formally isotopic" should be interpreted in the context of the h-principle: it means that pdq and λ are indistinguishable from the perspective of differential topology. 4.1.3 provides the first example of a pair of formally isotopic Weinstein manifolds with non-zero middle homology which can be distinguished by $\mathrm{Lef}_w(X)$.[5] In particular, 4.1.3 implies that, in contrast to $\mathrm{WMor}(-)$, the number $\mathrm{Lef}_w(-)$ is a truly symplectic invariant.

[5]For Weinstein manifolds with vanishing middle homology, $\mathrm{Lef}_w(-)$ distinguishes exotic examples for trivial reasons: for example, if $X = \mathbb{C}^n$ and X' is an exotic $\mathbb{C}^n$, then $\mathrm{Lef}_w(X) = 0$ but $\mathrm{Lef}_w(X') > 0$ since otherwise X' would be subcritical and hence standard.

The proof of 4.1.3 uses a construction of Eliashberg–Ganatra–Lazarev [EGL20], who exhibited an exotic $T^*S^3_{\text{exotic}}$ which contains a regular Lagrangian 3-torus. We will see that this implies that $\operatorname{Rdim} \mathcal{W}(T^*S^3_{\text{exotic}}) = 3$, and hence $\operatorname{Lef}_w(T^*S^3_{\text{exotic}}) \geq 4$. In contrast, it is known that T^*S^3 admits a Lefschetz fibration with 2 critical points.

Methods: bounds on the Rouquier dimension coming from geometry

It turns out that the Rouquier dimension of wrapped Fukaya categories provides a lower bound for the quantities we wish to study.

Proposition 4.1.6 (= 4.6.6). Let (X, λ) be a Weinstein manifold with properly embedded cocores (this is a generic condition) and nonzero wrapped Fukaya category. If $\phi : X \to X$ is a generic compactly-supported Hamiltonian symplectomorphism, then:

$$|\mathfrak{c}_X \cap \phi(\mathfrak{c}_X)| \geq \operatorname{Rdim} \mathcal{W}(X) + 1. \tag{4.1.8}$$

(The result we prove in 4.6.6 is stated slightly differently and is more general.) Let us sketch the proof of 4.1.6. To begin with, it can be shown that the Rouquier dimension of any A_∞ category is bounded above by the length of the shortest resolution of the diagonal bimodule by Yoneda bimodules (minus 1). This is a purely algebraic fact which has nothing to do with Fukaya categories.

We now set $r = |\mathfrak{c}_X \cap \phi(\mathfrak{c}_X)|$. Our goal is to construct a resolution of the diagonal bimodule $\Delta_{\mathcal{W}(X)}$ of length r. We do this by the following geometric argument. First, observe that the intersection points of $\mathfrak{c}_X$ with $\phi(\mathfrak{c}_X)$ are in bijection with the intersection points of the skeleton of $(\overline{X} \times X, -\lambda \oplus \phi_*\lambda)$ with the diagonal Lagrangian. Since $\mathfrak{c}_X \pitchfork \phi(\mathfrak{c}_X)$, the diagonal intersects the skeleton transversally in r points. By applying deep work of Chantraine–Dimitroglou-Rizell–Ghiggini–Golovko [CRGG17] or Ganatra–Pardon–Shende [GPS18b], one can then construct a resolution of the diagonal by r cocores of $\overline{X} \times X$. Morally, the idea is that as one scales the diagonal

Lagrangian by the Liouville flow, it converges to cocores at each point where it intersects the skeleton. The upshot is that we have resolved the diagonal Lagrangian by r cocores in $\overline{X} \times X$, which are just products of cocores in X. Finally, we can turn this resolution of the diagonal Lagrangian into a resolution of the diagonal bimodule. This last step is carried out in the Appendix, following the methods of [Gan12] and [GPS18b].

There is also an upper bound on the Rouquier dimension coming from Lefschetz fibrations.

Proposition 4.1.7 (= 4.6.5). Let X be a Liouville manifold with nonzero wrapped Fukaya category. Then we have

$$\mathrm{Lef}_w(X) \geq \mathrm{Rdim}\, \mathcal{W}(X) + 1. \tag{4.1.9}$$

The proof of 4.1.7 is straightforward and (at least implicitly) well-known to experts: the Fukaya–Seidel category associated to a Lefschetz fibration admits a semi-orthogonal decomposition of length equal to the number of critical points of the fibration, which immediately implies the corresponding upper bound for the Rouquier dimension.

Methods: commutative ring actions on wrapped Fukaya categories

In order to obtain conclusions about quantitative symplectic topology, we need a method for obtaining lower bounds on the Rouquier dimension of wrapped Fukaya categories which can then be combined with the upper bounds stated in Propositions 4.1.6 and 4.1.7. These lower bounds will ultimately arise from considering a canonical action of the Hochschild cohomology ring on the Fukaya category. To explain this, it is natural to start the discussion in a more general setting.

Let $\mathcal{C}$ be a pre-triangulated A_∞ category and let $\mathcal{T} = H^0(\mathcal{C})$ be its homotopy

category. Let $Z^\bullet(\mathcal{T})$ be the (graded-commutative) ring of graded natural transformations of the identity on $\mathcal{T}$. Given a graded-commutative ring R, a *central action* of R on the triangulated category $\mathcal{T}$ is by definition a morphism of graded-commutative rings $R \to Z^\bullet(\mathcal{T})$. Concretely, this is the data of a morphism $R \to \hom^\bullet_{H(\mathcal{C})}(K, K)$ for each object $K \in \mathcal{C}$ such that the induced left and right module structures on $\hom^\bullet_{H(\mathcal{C})}(K, L)$ coincide up to sign.

The upshot is that the morphism spaces of objects of $\mathcal{T}$ define modules over R. One can then study the triangulated category $\mathcal{T}$ using tools from (graded-)commutative algebra. This perspective leads to a very rich theory; see e.g. [BIK08].

The workhorse result for lower-bounding the Rouquier dimension of wrapped Fukaya categories in this chapter is the following:

Theorem 4.1.8 (= 4.5.1). Let $\mathcal{T}$ be a $\mathbb{Z}/m$-graded triangulated category ($1 \leq m \leq \infty$) and let R be a finite type k-algebra acting centrally on $\mathcal{T}$. Suppose that there exists a (split-)generator $G \in \mathcal{T}$ such that $\mathrm{Hom}^\bullet_{\mathcal{T}}(G, G)$ is a Noetherian R-module. Then $\mathrm{Rdim}\, \mathcal{T} \geq \dim_R \mathrm{Hom}^\bullet_{\mathcal{T}}(G, G)$.

4.5.1 can be viewed as an "affine" analog of a beautiful theorem of Bergh–Iyengar–Krause–Oppermann [BIKO10, Thm. 4.2]. In fact, (as we will see) one could directly apply [BIKO10, Thm. 4.2] and still obtain useful lower bounds for certain examples. However, for our purposes, [BIKO10, Thm. 4.2] suffers from two limitations. First of all, it is essentially never sharp (there is a pesky -1 which comes, ultimately, from working in the "projective" rather than "affine" setting). Secondly, it requires rather strong hypotheses on R and $\mathrm{Hom}^\bullet_{\mathcal{T}}(G, G)$ which do not hold in many cases of interest to us (for instance, 4.1.4 and 4.5.3 would fail).[6] In contrast, 4.5.1 is sharp for many examples (such as cotangent bundles of n-spheres and n-tori), and more widely applicable.

[6] There is a typo in the definition of the projective dimension in [BIKO10, Sec. 2] which propagates. As a result, one needs to add an assumption in the statement of [BIKO10, Thm. 4.2], the most natural (to us) being that $R^0 = k$.

Our general approach to proving 4.5.1 via *Koszul objects* follows [BIKO10]; however, the input from graded-commutative algebra is replaced by new arguments from (ungraded) commutative algebra. We crucially use the hypothesis that R is of finite type. This hypothesis is rather severe from the representation theoretic perspective taken in [BIKO10]. However, from the symplectic perspective which we adopt in this chapter, this assumption is rather mild.

Of course, 4.1.8 is useless without a good supply of central actions. Happily, for any pre-triangulated A_∞ category $\mathcal{C}$, there is a canonical central action of the Hochschild cohomology ring $HH^\bullet(\mathcal{C})$ on $\mathcal{T} = H^0(\mathcal{C})$ called the *characteristic morphism*. Given any object $K \in \mathcal{C}$, this is just the map $HH^\bullet(\mathcal{C}) \to \hom^\bullet_{H(\mathcal{C})}(K, K)$ that projects a Hochschild cochain onto its length zero part. In general, the characteristic morphism is neither injective nor surjective.

In the case which is relevant to us, namely when $\mathcal{T} = H^0(\operatorname{Perf} \mathcal{W}(X))$ and X is Weinstein, the characteristic morphism is known to admit a concrete description in terms of the familiar *closed-open map*. To explain this a little, recall that the closed-open map defines a ring morphism $SH^\bullet(X) \to HW^\bullet(L, L)$, for any object $L \in \mathcal{W}(X)$. On the other hand, if X is Weinstein, there is an isomorphism of graded-commutative rings $HH^\bullet(\mathcal{W}(X)) = SH^\bullet(X)$ [CRGG17]. Now the point is simply that the obvious diagram formed by these maps commutes. As a result, we obtain the following corollary which is the desired lower bound.

Corollary 4.1.4 (= 4.5.3). Suppose that $SH^0(X)$ admits a subring R of finite-type over k, and suppose there exists a (split-)generator $K \in \mathcal{W}(X)$ such that $HW^\bullet(K, K)$ is Noetherian over R, where R acts via the closed-open map. Then

$$\operatorname{Rdim} \mathcal{W}(X) \geq \dim_R HW^\bullet(K, K). \tag{4.1.10}$$

4.1.4 works with any choice of gradings, but $SH^\bullet(X)$ will typically be $\mathbb{Z}$-graded

or $\mathbb{Z}/2$-graded, depending on the application. The flexibility to choose different gradings is convenient in practice, since of course $SH^0(X)$ will admit different subrings depending on the choice of grading.

Acknowledgements

We are indebted to Daniel Álvarez-Gavela for answering countless questions about arborealization, Vivek Shende for help with microlocal sheaves and homological mirror symmetry of toric varieties, Srikanth Iyengar for discussions and comments related to 4.5.2, Y. Baris Kartal for telling us the proof of 4.6.5 and for useful discussions related to 4.4.5, and John Pardon for many helpful conversations about various aspects of the chapter. In addition, we would like to thank Denis Auroux, Petter Bergh, Georgios Dimitroglou Rizell, Sheel Ganatra, Wahei Hara, Oleg Lazarev, David Nadler, Amnon Neeman and Dmitri Orlov for answering our questions at various points in this project.

4.2 Category theory

4.2.1 Notation and conventions

Unless otherwise indicated, k denotes a field.

Let $\mathcal{T}$ be a category. We will routinely abuse notation by writing $K \in \mathcal{T}$ to mean that K is an object of $\mathcal{T}$. Similarly, we write $\mathcal{I} \subseteq \mathcal{T}$ to mean that $\mathcal{I}$ is a full subcategory of $\mathcal{T}$. Given two full subcategories $\mathcal{I}, \mathcal{J}$ of $\mathcal{T}$, we let $\mathcal{I} \cup \mathcal{J}$ be the full subcategory with objects the union of the objects in $\mathcal{I}$ and those in $\mathcal{J}$.

A full subcategory $\mathcal{I} \subseteq \mathcal{T}$ is said to be *strictly full* if is closed under isomorphisms. We will sometimes implicitly identify a set of objects with the corresponding full subcategory (resp. identify a set of isomorphism classes of objects with the corresponding strictly full subcategory).

Given an object $K \in \mathcal{T}$, a *summand* of K is a triple (Z, k, r) consists of an object

$Z \in \mathcal{T}$ along with maps $r \in \hom(Z,K), k \in \hom(K,Z)$ such that $k \circ r = \mathrm{id}$ and $r \circ k$ is an idempotent. An idempotent $e \in \hom(K,K)$ is said to *split* if $e = r \circ k$ for some summand (Z,k,r). A category is said to be *idempotent complete* (or Karoubi complete) if all idempotents split.

4.2.2 Triangulated categories

A triangulated category $\mathcal{T} = (\mathcal{T}, \Sigma, \Delta)$ is an additive category $\mathcal{T}$ equipped with an autoequivalence $\Sigma : \mathcal{T} \to \mathcal{T}$, called the *shift functor*, and a class of distinguished triangles Δ. This data is required to satisfy certain axioms. A standard reference for triangulated categories is [Nee01].

We say that a triangulated category $\mathcal{T}$ is $\mathbb{Z}/m$-graded (for $m \in \mathbb{N}_+ \cup \{\infty\}$ and $\mathbb{Z}/\infty := \mathbb{Z}$) if $\Sigma^m = \mathrm{id}$. If $\mathcal{T}$ is $\mathbb{Z}/m$-graded, we define $\mathrm{Hom}^\bullet_{\mathcal{T}}(X,Y) := \oplus_{i=1}^m \mathrm{Hom}_{\mathcal{T}}(X, \Sigma^i Y)$.

Given full subcategories $\mathcal{I}_0, \mathcal{I}_1$ of $\mathcal{T}$, let $\mathcal{I}_0 * \mathcal{I}_1$ be the strictly full subcategory which is uniquely characterized by the following property: an object $K \in \mathcal{T}$ is contained in $\mathcal{I}_0 * \mathcal{I}_1$ if there exists a distinguished triangle $K_0 \to K \to K_1 \to$ with $K_0 \in \mathcal{I}_0$ and $K_1 \in \mathcal{I}_1$. It is a straightforward consequence of the octahedral axiom for triangulated categories that the operation $(\mathcal{I}_0, \mathcal{I}_1) \mapsto \mathcal{I}_0 * \mathcal{I}_1$ is associative (see e.g. [BBD82, Lem. 1.3.10]).

Given $\mathcal{I} \subseteq \mathcal{T}$ a full subcategory, let $[\mathcal{I}]$ be the smallest strictly full subcategory of $\mathcal{T}$ containing $\mathcal{I}$ and closed under finite direct sums and shifts. Let $\langle \mathcal{I} \rangle$ be the smallest strictly full subcategory of $\mathcal{T}$ closed under finite direct sums, shifts and taking summands. Let $\mathrm{ads}(\mathcal{I})$ be the smallest strictly full subcategory of $\mathcal{T}$ containing $\mathcal{I}$ and closed under *arbitrary* direct sums and shifts.

We inductively define $[\mathcal{I}]_n = [\mathcal{I}]_{n-1} * [\mathcal{I}]$ and $\langle \mathcal{I} \rangle_n = \langle \langle \mathcal{I} \rangle_{n-1} * \langle \mathcal{I} \rangle \rangle$. We set $[\mathcal{I}]_\infty := \cup_{n \in \mathbb{N}} [\mathcal{I}]_n$ and $\langle \mathcal{I} \rangle_\infty = \cup_{n \in \mathbb{N}} \langle \mathcal{I} \rangle_n$.

Finally, we let $\mathcal{I}_{sd}$ be the smallest strictly full subcategory containing all sum-

mands of $\mathcal{I}$ (this is often called the "Karoubi completion"). A full subcategory of a triangulated category is called *thick* if it is closed under taking summands.

Remark 4.2.1. Some sources (e.g. [EL21]) adopt a different convention in which the indices in $\langle - \rangle_n$ and $[-]_n$ are shifted down by 1.

Lemma 4.2.1. The following properties hold for all $n \in \mathbb{N} \cup \{\infty\}$ and any $\mathcal{I} \subseteq \mathcal{T}$.

1. $[\mathcal{I}] = [[\mathcal{I}]]$, $\mathcal{I}_{sd} = (\mathcal{I}_{sd})_{sd}$ and $[\mathcal{I}_{sd}] \subseteq [\mathcal{I}]_{sd}$
2. $\langle \mathcal{I} \rangle_n = ([\mathcal{I}]_n)_{sd}$
3. $[[\mathcal{I}]_n] = [\mathcal{I}]_n$ and $\langle \langle \mathcal{I} \rangle_n \rangle = \langle \mathcal{I} \rangle_n$
4. Given objects $G, H \in \mathcal{T}$, we have $\langle G \rangle_n \subseteq \langle G \oplus H \rangle_n$ and $\langle H \rangle_n \subseteq \langle G \oplus H \rangle_n$.

Proof. (1) is obvious. (2) can be established by following the argument of [Sei08b, Lem. 4.1] (and is stated without proof in [Rou08, Rmk. 3.1]). (3) is straightforward using (2). (4) is obvious. □

Lemma 4.2.2. Let $\mathcal{T}$ be a triangulated category. Let $\mathcal{I}, \mathcal{J}$ be thick full triangulated subcategories such that $\mathcal{T} = \langle \mathcal{I} \cup \mathcal{J} \rangle_\infty$ and $\hom(K, L) = 0$ for all $K \in \mathcal{J}$ and $L \in \mathcal{I}$. Then $\mathcal{T} = \langle \mathcal{J} * \mathcal{I} \rangle$.

Proof. The proof is standard; see e.g. [Bon89, Lem. 3.1]. Since $\langle \mathcal{J} * \mathcal{I} \rangle$ contains $\mathcal{I}$ and $\mathcal{J}$, it is enough to verify that $\langle \mathcal{J} * \mathcal{I} \rangle$ is closed under cones. This is a straightforward check using [BBD82, Prop. 1.1.11]. □

Lemma 4.2.3. Let $\mathcal{T}$ be a triangulated category. Let $\mathcal{I}_1, \ldots, \mathcal{I}_r$ be full triangulated subcategories which are mutually orthogonal, in the sense that $\hom(K, L) = 0$ for all $K \in \mathcal{I}_k$ and $L \in \mathcal{I}_l$ if $k \neq l$. Then $\langle \mathcal{I}_1 \cup \cdots \cup \mathcal{I}_r \rangle_n = \langle \langle \mathcal{I}_1 \rangle_n \cup \cdots \cup \langle \mathcal{I}_r \rangle_n \rangle$ for all $n \in \mathbb{N} \cup \{\infty\}$. □

The following lemma is well-known; see e.g. [Rou08, Cor. 3.14].

Lemma 4.2.4. Let $\mathcal{T}$ be a triangulated category and let $\mathcal{T}^c$ denote the full triangulated subcategory of compact objects. Then given $\mathcal{I} \subseteq \mathcal{T}^c$, we have $\mathcal{T}^c \cap \langle \mathrm{ads}(\mathcal{I}) \rangle_d = \langle \mathcal{I} \rangle_d$. □

4.2.3 A_∞ categories

Basic definitions

Unless otherwise indicated, we follow the conventions of [Sei08b] and [Sei] when discussing A_∞ categories. One exception is that all A_∞ categories in this chapter are assumed to be *strictly* unital. We will work mostly with $\mathbb{Z}/2$-graded or $\mathbb{Z}$-graded A_∞ categories, which are always assumed to be defined over a field k. However, the categorical reasoning in this chapter applies with purely superficial changes to the setting of $\mathbb{Z}/n$-graded A_∞ categories for any $n \in \mathbb{N}_+$ (the case $n = 1$ corresponds to an ungraded A_∞ category).

Given A_∞ categories $\mathcal{C}_1$ and $\mathcal{C}_2$, an $(\mathcal{C}_1, \mathcal{C}_2)$-bimodule is an A_∞ functor $\mathcal{C}_1^{\mathrm{op}} \times \mathcal{C}_2 \to \mathrm{Ch}\, k$, where $\mathrm{Ch}\, k$ is the dg category of chain complexes of k-vector spaces. Note that $(\mathcal{C}_1, \mathcal{C}_2)$-bimodules form a dg category, which is denoted by $(\mathcal{C}_1, \mathcal{C}_2) - \mathrm{mod}$. As usual, two such bimodules are said to be quasi-isomorphic if they are isomorphic in the homotopy category $H^0((\mathcal{C}_1, \mathcal{C}_2) - \mathrm{mod})$. Let $\Delta_{\mathcal{C}}$ be the diagonal bimodule over an A_∞ category $\mathcal{C}$. A right (resp. left) module over $\mathcal{C}$ is a $(k \times \mathcal{C})$-bimodule (resp. a $(\mathcal{C} \times k)$-bimodule). The right (resp. left) modules over $\mathcal{C}$ form a dg category $\mathrm{Mod}\, \mathcal{C}$ (resp. $\mathrm{Mod}\, \mathcal{C}^{\mathrm{op}}$).

Let $\{\mathcal{C}_i\}_{i=1,2,3}$ be A_∞ categories. Given a $(\mathcal{C}_1, \mathcal{C}_2)$-bimodule $\mathcal{P}$ and a $(\mathcal{C}_2, \mathcal{C}_3)$-bimodule $\mathcal{Q}$, we can form their convolution (bimodule tensor product) $\mathcal{P} \otimes_{\mathcal{C}_2} \mathcal{Q}$. Using the same notations, the convolution has the following properties:

- it is strictly associative;
- we have quasi-isomorphisms between bimodules $\Delta_{\mathcal{C}_1} \otimes_{\mathcal{C}_1} \mathcal{P} \to \mathcal{P}$ and $\mathcal{P} \otimes_{\mathcal{C}_2} \Delta_{\mathcal{C}_2} \to$

$\mathcal{P}$.

- the convolution $- \otimes_{\mathcal{C}_2} \mathcal{Q}$ induces a dg functor $(\mathcal{C}_1, \mathcal{C}_2) - \mathrm{mod}$ to $(\mathcal{C}_1, \mathcal{C}_3) - \mathrm{mod}$. Similarly for $\mathcal{Q} \otimes_{\mathcal{C}_3} -$.

There is a canonical Yoneda embedding $\mathcal{C} \to \mathrm{Mod}\,\mathcal{C}$. Let $\mathrm{Tw}\,\mathcal{C} \subseteq \mathrm{Mod}\,\mathcal{C}$ be the closure of the image of the Yoneda embedding under taking mapping cones. Let $\mathrm{Perf}\,\mathcal{C} \subseteq \mathrm{Mod}\,\mathcal{C}$ be the closure of the image of the Yoneda embedding under taking mapping cones and summands. Note that we have

$$\mathrm{Perf}\,\mathcal{C} = (\mathrm{Mod}\,\mathcal{C})^c, \tag{4.2.1}$$

where $(-)^c$ means taking compact objects; see [GPS18a, (A.2) and Prop. A.3].

An A_∞ category $\mathcal{C}$ is said to be *pre-triangulated* if the canonical embedding $\mathcal{C} \hookrightarrow \mathrm{Tw}\,\mathcal{C}$ is a quasi-equivalence. The homotopy category $H^0(\mathcal{C})$ of a pre-triangulated A_∞ category is a triangulated category.

Given an A_∞ category $\mathcal{C}$ and a set of objects $\mathcal{A}$, we can form the *quotient* A_∞ category $\mathcal{C}/\mathcal{A}$ (see [LO06, Dri04]), which comes equipped with a canonical map $q : \mathcal{C} \to \mathcal{C}/\mathcal{A}$. Suppose $\mathcal{B}$ is another A_∞ category. Given a functor $\mathcal{C} \to \mathcal{B}$ which sends $\mathcal{A}$ to acyclic objects, there is an induced functor $\mathcal{C}/\mathcal{A} \to \mathcal{B}$ (see [GPS18a, Sec. A.7]). We note that the canonical map $(\mathrm{Tw}\,\mathcal{C})/\mathcal{A} \to \mathrm{Tw}(\mathcal{C}/\mathcal{A})$ is a quasi-equivalence, and the canonical map $\mathrm{Perf}\,\mathcal{C}/\mathcal{A} \to \mathrm{Perf}(\mathcal{C}/\mathcal{A})$ is a Morita equivalence (it is not in general a quasi-equivalence).

Equivalent constructions of the derived category of a (dg) k-algebra

In the literature, one encounters various definitions of the derived category of a (dg) k-algebra. We therefore collect some standard facts which will be implicitly assumed in the sequel.

Lemma 4.2.5. Let $\mathcal{A}$ be a dg k-algebra. The following constructions of the derived category $D(\mathcal{A})$ produce equivalent triangulated categories.

1. Consider the quotient (in the sense of Drinfeld [Dri04]) of the dg category of dg $\mathcal{A}$-modules by the dg subcategory of acyclic dg $\mathcal{A}$-modules. Then pass to $H^0(-)$.

2. View $\mathcal{A}$ as an A_∞ algebra and consider $\operatorname{Mod}\mathcal{A}$, the A_∞ category of A_∞ $\mathcal{A}$-modules. Now pass to $H^0(-)$.

□

Lemma 4.2.6. Let A be a k-algebra. The following constructions of the derived category $D(A)$ produce equivalent triangulated categories.

1. Consider the abelian category of modules over A and let $K(A)$ be the homotopy category. Then take the quotient (in the sense of Verdier) of $K(A)$ by the subcategory of acyclic complexes.

2. View A as a dg algebra concentrated in degree zero with the trivial differential. Now apply either one of the equivalent constructions of 4.2.5.

□

In contrast to 4.2.6, if A is a graded k algebra, it is in general *not* the case that the derived category of the abelian category of graded A-modules coincides with the derived category of dg A-modules.

Remark 4.2.2. Let A be a k-algebra. Viewing A as a dg k-algebra, we can consider the A_∞ category $\operatorname{Perf} A$. Then $H^0(\operatorname{Perf} A) \subseteq H^0(\operatorname{Mod} A) = D(A)$ is equivalent to the smallest idempotent complete triangulated subcategory of $D(A)$ containing A (where A is viewed as a module over itself). It is shown in [Sta21, Prop. 07LT] that the objects of this latter category are precisely the perfect complexes, i.e. the complexes of A-modules quasi-isomorphic to a bounded complex of finite projective A-modules.

4.2.4 The Rouquier dimension

In his remarkable paper [Rou08], Rouquier introduced the following notion of dimension for triangulated categories.

Definition 4.2.7. The *Rouquier dimension* of a triangulated category $\mathcal{T}$ is the smallest $n \in \mathbb{N} \cup \{\infty\}$ such that there exists an object $G \in \mathcal{T}$ with $\mathcal{T} = \langle G \rangle_{n+1}$. We denote the Rouquier dimension of $\mathcal{T}$ by $\operatorname{Rdim} \mathcal{T}$.

The Rouquier dimension of an A_∞ category $\mathcal{C}$ is defined to be the Rouquier dimension of $H^0(\operatorname{Perf} \mathcal{C})$ and is denoted by $\operatorname{Rdim} \mathcal{C}$.

We now collect some standard properties of the Rouquier dimension which will be needed in the sequel. The following lemma is a useful inequality for the Rouquier dimension of semi-orthogonal decompositions.

Lemma 4.2.8. Let $\mathcal{T}$ be a triangulated category. Let $\mathcal{I}_1, \ldots, \mathcal{I}_m$ be thick full triangulated subcategories satisfying the following conditions:

(i) for any $j > i$, we have $\hom(K, L) = 0$ for all $K \in \mathcal{I}_j$ and $L \in \mathcal{I}_i$;

(ii) $\langle \mathcal{I}_1 \cup \cdots \cup \mathcal{I}_m \rangle_\infty = \mathcal{T}$.

Then

$$\operatorname{Rdim} \mathcal{T} \leq \sum_{l=1}^{m} (\operatorname{Rdim} \mathcal{I}_l + 1) - 1. \tag{4.2.2}$$

(4.2.2) is in general far from sharp, as can be seen by taking for instance $\mathcal{I}_2, \ldots, \mathcal{I}_m = 0$.

Proof. Note that $\mathcal{T} = \langle \mathcal{I}_1 \cup \cdots \cup \mathcal{I}_m \rangle_\infty = \langle \langle \mathcal{I}_1 \cup \cdots \cup \mathcal{I}_{m-1} \rangle_\infty \cup \mathcal{I}_m \rangle_\infty$. Hence it is sufficient to treat the case where $m = 2$.

Suppose $\operatorname{Rdim} \mathcal{I}_1 = n_1$ and $\operatorname{Rdim} \mathcal{I}_2 = n_2$. This means that there exist objects $G \in \mathcal{I}_1$ and $H \in \mathcal{I}_2$ such that $\mathcal{I}_1 = \langle G \rangle_{n_1+1}$ and $\mathcal{I}_2 = \langle H \rangle_{n_2+1}$.

We now have, by 4.2.2, $\mathcal{T} = \langle\langle H\rangle_{n_2+1} * \langle G\rangle_{n_1+1}\rangle \subseteq \langle\langle G\oplus H\rangle_{n_2+1} * \langle G\oplus H\rangle_{n_1+1}\rangle = \langle([G\oplus H]_{n_2+1})_{sd} * ([G\oplus H]_{n_1+1})_{sd}\rangle \subseteq \langle([G\oplus H]_{n_2+1} * [G\oplus H]_{n_1+1})_{sd}\rangle = \langle([G\oplus H]_{n_1+n_2+2})_{sd}\rangle = \langle\langle G\oplus H\rangle_{n_1+n_2+2}\rangle$. □

The following statements show that the Rouquier dimension is non-increasing under quotients of categories.

Lemma 4.2.9 (Lem. 3.4 in [Rou08]). Let $\mathcal{T}$ and $\mathcal{T}'$ be triangulated categories. If there exists an exact functor $\mathcal{T} \to \mathcal{T}'$ with dense image (i.e. every object in $\mathcal{T}'$ is isomorphic to a summand of an object in the image), then $\operatorname{Rdim}\mathcal{T}' \leq \operatorname{Rdim}\mathcal{T}$. □

Corollary 4.2.1. Let $\mathcal{C}$ be an A_∞ category and let $\mathcal{A}$ denote a set of objects. Then $\operatorname{Rdim}\mathcal{C}/\mathcal{A} \leq \operatorname{Rdim}\mathcal{C}$. □

When a triangulated category comes from an A_∞ category, the following lemma shows that the Rouquier dimension is bounded above by the length of the shortest resolution of the diagonal bimodule.

Lemma 4.2.10. Let $\mathcal{C}$ be an A_∞ category and let $\Delta_{\mathcal{C}}$ denote the diagonal bimodule. Suppose there exists a resolution of length $\ell \in \mathbb{N}_+$

$$0 = \mathcal{M}_0 \to \mathcal{M}_1 \to \cdots \to \mathcal{M}_\ell \tag{4.2.3}$$

of $(\mathcal{C},\mathcal{C})$-bimodules where

- $\Delta_{\mathcal{C}}$ is a summand of $\mathcal{M}_\ell$;
- there exists perfect modules $\mathcal{P}_i^r \in \operatorname{Perf}\mathcal{C}$ and $\mathcal{P}_i^l \in \operatorname{Perf}\mathcal{C}^{\mathrm{op}}$ such that $\operatorname{cone}(\mathcal{M}_{i-1} \to \mathcal{M}_i) \in [\mathcal{P}_i^l \otimes \mathcal{P}_i^r]$ in $H^0((\mathcal{C},\mathcal{C}) - \mathrm{mod})$.

Then $\operatorname{Rdim}\mathcal{C} \leq \ell - 1$.

Proof. Given $Q \in \operatorname{Perf}\mathcal{C}$, we convolve $\mathcal{Q}$ with the above resolution (i.e. we apply the functor $\mathcal{Q} \otimes_{\mathcal{C}} -$). Note that $\mathcal{Q} \otimes_{\mathcal{C}} \mathcal{P}_i^l \otimes \mathcal{P}_i^r \simeq H^\bullet(\mathcal{Q} \otimes_{\mathcal{C}} \mathcal{P}_i^l) \otimes \mathcal{P}_i^r$. Writing $C_i :=$

$H^\bullet(\mathcal{Q} \otimes_{\mathcal{C}} \mathcal{P}_i^l)$, it follows that $\mathcal{Q} \in \langle \oplus_{i=1}^{\ell}(C_i \otimes \mathcal{P}_i^r)\rangle_\ell \subseteq \langle \mathrm{ads}(\mathcal{G})\rangle_\ell$, where $\mathcal{G} = \oplus_{i=1}^{\ell}\mathcal{P}_i^r$. Since $\mathcal{Q}$ and $\mathcal{G}$ are compact objects in $\operatorname{Mod}\mathcal{C}$ (see (4.2.1)), this implies that $\mathcal{Q} \in \langle \mathcal{G}\rangle_\ell$ by 4.2.4. □

4.2.5 Homotopy colimits

Given a diagram of A_∞ categories $\{\mathcal{C}_\sigma\}_{\sigma\in\Sigma}$ indexed by a poset Σ, one can form the *Grothendieck construction* (or semiorthogonal gluing)

$$\operatorname{Groth}_{\sigma\in\Sigma}\mathcal{C}_\sigma. \tag{4.2.4}$$

This is an A_∞ category with $\operatorname{Ob}(\operatorname{Groth}_{\sigma\in\Sigma}\mathcal{C}_\sigma) = \bigsqcup_{\sigma\in\Sigma}\operatorname{Ob}\mathcal{C}_\sigma$. It satisfies the following important property: if $K \in \mathcal{C}_\sigma$ and $L \in \mathcal{C}_{\sigma'}$ are objects, then the space of morphism from K to L in $\operatorname{Groth}_{\sigma\in\Sigma}\mathcal{C}_\sigma$ is zero unless $\sigma \leq \sigma'$.

The category $\operatorname{Groth}_{\sigma\in\Sigma}\mathcal{C}_\sigma$ admits a distinguished collection of morphism A_Σ called *adjacency morphisms*. The *homotopy colimit* of the diagram $\{\mathcal{C}_\sigma\}_{\sigma\in\Sigma}$ is defined to be the localization of the Grothendieck construction at the adjacency morphisms:

$$\operatorname{hocolim}_{\sigma\in\Sigma} := (\operatorname{Groth}_{\sigma\in\Sigma}\mathcal{C}_\sigma)\,[A_\Sigma]^{-1}. \tag{4.2.5}$$

We refer the reader to [GPS18b, A.4] for details on the Grothendieck construction and homotopy colimit. We will not need any properties of these constructions beyond those mentioned above.

We now collect some lemmas about the Rouquier dimension of homotopy colimits. We begin with a small piece of notation: given a finite poset Σ and an element $\sigma \in \Sigma$, we let $h(\sigma) \in \mathbb{N}_+$ be the length of the longest chain $\sigma_1 < \cdots < \sigma_{h(\sigma)} = \sigma$ terminating at σ. Note that $h(\sigma) = h(\sigma')$ implies that σ and σ' are either incomparable or equal. Let D be the maximum of $h(\sigma)$ for all $\sigma \in \Sigma$.

Let us now consider the triangulated categories $\mathcal{T} = H^0(\operatorname{Perf}\operatorname{Groth}_{\sigma\in\Sigma}\mathcal{C}_\sigma)$ and

$\mathcal{T}' = H^0(\operatorname{Perf}\operatorname{hocolim}_{\sigma\in\Sigma}\mathcal{C}_\sigma)$. We write $\mathcal{I}_\ell = \langle\bigcup_{h(\sigma)=\ell}\mathcal{C}_\sigma\rangle \subseteq \mathcal{T}$.

Lemma 4.2.11. $\operatorname{Rdim}\mathcal{I}_\ell = \max_{h(\sigma)=\ell}\operatorname{Rdim}\mathcal{C}_\sigma$.

Proof. Follows from 4.2.3 and the fact that $\mathcal{C}_\sigma, \mathcal{C}_{\sigma'}$ are orthogonal whenever $\sigma \neq \sigma'$. □

Lemma 4.2.12. We have $\operatorname{Rdim}\mathcal{T}' \leq \operatorname{Rdim}\mathcal{T} \leq \sum_{\ell=1}^{D}(\operatorname{Rdim}\mathcal{I}_\ell + 1) - 1$.

Proof. The first inequality follows from 4.2.1. For the second one, observe that $\langle\bigcup_{\ell=1}^{L}\mathcal{I}_\ell\rangle_\infty = \mathcal{T}$. As noted above, we have $\hom_{\mathcal{T}}(K, L) = 0$ if $K \in \mathcal{I}_j$ and $L \in \mathcal{I}_i$ for $j > i$. The claim thus follows from 4.2.8. □

4.2.6 The Rouquier dimension of some triangulated categories

We collect some computations from the literature to be used later. For many other interesting computations, we refer to [Ela20].

Coherent sheaves

We let $D^b\operatorname{Coh}(Y)$ denote the bounded derived category of coherent sheaves over a scheme Y, which is always assumed to be separated and of finite type over a field.

The following facts are due to Rouquier [Rou08].

- If Y is a separated scheme of finite type over a perfect field, then $\operatorname{Rdim} D^b\operatorname{Coh}(Y)$ is finite; [Rou08, Thm. 7.38].
- If Y is a smooth quasi-projective scheme over a field, then $\operatorname{Rdim} D^b\operatorname{Coh}(Y) \leq 2\dim(Y)$; [Rou08, Prop. 7.9].
- If Y is a reduced separated scheme of finite type over a field, then $\operatorname{Rdim} D^b\operatorname{Coh}(Y) \geq \dim(Y)$; [Rou08, Prop 7.16].

- If Y is a smooth affine scheme of finite type over a field, then $\operatorname{Rdim} D^b \operatorname{Coh}(Y) = \dim(Y)$; [Rou08, Thm. 7.17].

It follows from the above facts that $\operatorname{Rdim} D^b \operatorname{Coh}(Y) = \dim Y$ if Y is affine. Orlov conjectured in [Orl09b, Conj. 10] that this equality holds for any smooth quasi-projective variety. We will return to this conjecture in 4.4.4.

If Y is a variety which is not assumed to be smooth, then the Rouquier dimension $D^b \operatorname{Coh}(Y)$ can be bigger than its Krull dimension. The following example, which was communicated to us by W. Hara, shows that there exist projective curves whose which have arbitrarily large Rouquier dimension.

Example 4.2.3. Fix a positive integer m. Wantanabe [Wat73, Thm. 1] constructed a k-algebra $(A, \mathfrak{m})$ which is a complete intersection Noetherian local domain of Krull dimension 1 and with $\dim_{R/\mathfrak{m}} \mathfrak{m}/\mathfrak{m}^2 = m$. Moreover, by construction, A is the localization of a finite type k-algebra B at a maximal ideal.

Let $Y_0 = \operatorname{Spec} B$ and let $Y \subseteq \mathbb{P}^n$ be the projective compactification of Y_0. We then have $\operatorname{Rdim} D^b \operatorname{Coh}(Y) \geq \operatorname{Rdim} D^b \operatorname{Coh}(Y_0) \geq \operatorname{Rdim} D^b Coh(R)$, where the second inequality is by [AT15, Lem. 4.2]. We now appeal to [BIKO10, Cor. 5.10], which implies that $\operatorname{Rdim} D^b Coh(R) \geq m - \dim Y - 1 = m - 2$.

Ballard–Favero gave examples of singular varieties whose Rouquier dimension equals their Krull dimension (see [BF12, Cor. 3.3]). We will prove in 4.4.4 that such an equality holds for all toric boundary divisors of smooth quasiprojective DM stacks over the complex numbers of dimension at most 3 and some other singular varieties.

Representation categories of quivers

In this chapter, a *quiver* Γ is a directed graph which is connected, finite, and acyclic (i.e. having no loops and no cycles).

Given a quiver Γ and a field k, we can consider the *path algebra* $k[\Gamma]$. As a vector space, $k[\Gamma]$ is generated by words of composable arrows; in particular, it is finite dimensional as a k-vector space since Γ is acyclic. The product of two words is given by concatenation when this makes sense and is defined to be zero otherwise. The path algebra of a quiver is in general non-commutative. Unless otherwise indicated, we always view $k[\Gamma]$ as an *ungraded* algebra (although it admits a natural grading by path length).

For a quiver Γ, we will be interested in the A_∞ category $\operatorname{Perf} k[\Gamma]$, as well as in the associated triangulated category $H^0(\operatorname{Perf} k[\Gamma])$. As noted in 4.2.2, the later category is equivalent to the smallest idempotent complete triangulated subcategory of $D(k[\Gamma])$ containing $k[\Gamma]$.

A quiver is said to be of Dynkin ADE type if its underlying graph is a Dynkin diagram of type A_n, D_n, E_6, E_7, E_8.

The following proposition is a special case of [Ela20, Prop. 4.3, 4.4] and is ultimately an application of Gabriel's theorem [DW17, Chap. 4]. (We warn the reader that the notation $\operatorname{Perf}(A)$ in [Ela20] denotes the triangulated category which we have been denoting by $H^0(\operatorname{Perf} A)$.)

Proposition 4.2.13. Let Γ be a quiver.

- If Γ is of Dynkin ADE type, then $\operatorname{Rdim}(\operatorname{Perf} k[\Gamma]) = 0$;
- if Γ is not of Dynkin ADE type, then $\operatorname{Rdim}(\operatorname{Perf} k[\Gamma]) = 1$.

4.3 Wrapped Fukaya categories

4.3.1 Liouville manifolds and related objects

Basic definitions

A Liouville vector field on a symplectic manifold (X, ω) is a vector field V which satisfies $\mathcal{L}_V \omega = \omega$. Dually, a Liouville 1-form λ on a symplectic manifold (X, ω) is a primitive for ω.

A *Liouville cobordism* is an exact symplectic manifold (X_0, λ_0) with boundary $\partial X_0 = \partial_+ X_0 \sqcup \partial_- X_0$, such that the Liouville vector field is outward pointing along $\partial_+ X_0$ and inward pointing along $\partial_- X_0$. If $\partial_- X_0 = \emptyset$, such a cobordism is called a *Liouville domain*.

A *Liouville manifold* is an exact symplectic manifold (X, λ) which is modeled near infinity on the symplectization of a closed contact manifold $(Y, \xi = \ker \lambda_Y)$. More formally, this means that there exists a proper embedding

$$e : (SY, \lambda_Y) \hookrightarrow (X, \lambda) \tag{4.3.1}$$

which covers the complement of a compact set and satisfies $e^* \lambda = \lambda_Y$. Here $(SY, \lambda_Y) = \{\sigma \in T^*Y \mid \sigma|_\xi = 0\}$ denotes the symplectization of (Y, ξ), where λ_Y is the (restriction of) the tautological 1-form on T^*Y. A Liouville manifold can equivalently be viewed as the result of completing a Liouville domain by gluing on the positive symplectization of its boundary. The contact manifold Y is written as $\partial_\infty X$ and is called the *ideal contact boundary* of X.

A *Liouville sector* is a Liouville manifold with boundary which is modeled near infinity on the positive symplectization of a contact manifold with convex boundary (see [GPS20, Def. 2.4] for the precise definition). A Liouville manifold is in particular a Liouville sector. One example of a Liouville sector which is not a Liouville manifold

is the cotangent bundle of a compact manifold with boundary. Given a Liouville sector (X, λ), a *stop* is an arbitrary closed subset $\mathfrak{f} \subseteq (\partial_\infty X)^\circ$. Such a pair $(X, \mathfrak{f})$ is often called a stopped Liouville sector.

A Liouville manifold (X, λ) is said to be *Weinstein* if the Liouville vector field V is gradient-like with respect to a proper Morse function ϕ. Note that according to this definition, being Weinstein is a *property* of Liouville manifolds rather than an extra structure. A Liouville cobordism is said to be Weinstein if its completion is Weinstein; a Liouville sector is said to be Weinstein if its convex completion [GPS20, Sec. 2.7] is Weinstein. A *Liouville pair* consists of a pair (X, A) such that X is a Liouville manifold and A (possibly empty) is a Liouville hypersurface embedded inside the ideal contact boundary $\partial_\infty X$, i.e. the induced contact form $\lambda_{\partial_\infty X}$ from the Liouville 1–form λ restricts to the Liouville 1–form on A. The pair (X, A) is called a *Weinstein pair* of both X and A are Weinstein. A *deformation* of Liouville manifolds/sectors/cobordisms is a one parameter family of such objects.

Skeleta

The *skeleton* $\mathfrak{c}_X \subset X$ of a Liouville manifold or sector (X, λ) is the set of points which do not escape to infinity under the positive Liouville flow. (More formally: $p \in \mathfrak{c}_X \subset X$ iff there exists a compact subset $K \subset X$ containing p and such that p does leave K under the positive Liouville flow). The skeleton is sometimes called the *core* in the literature.

The (relative) skeleton $\mathfrak{c}_{X,\mathfrak{f}}$ of a stopped Liouville manifold $(X, \mathfrak{f})$ is defined to be the set of points which do not escape to $\partial_\infty X - \mathfrak{f}$ under the positive Liouville flow. If (X, A) is a Liouville pair, we similarly define the (relative) skeleton $\mathfrak{c}_{X,A}$ to be the set of points which do not escape to $\partial_\infty X - \mathfrak{c}_A$ under the positive Liouville flow.

The skeleton of an arbitrary Liouville manifold always has measure zero. If (X, λ) is Weinstein, then a choice of a gradient-like proper Morse function $\phi : X \to \mathbb{R}$

defines a handlebody decomposition of X. The skeleton is then precisely the union of the stable manifolds associated to this decomposition. In particular, the skeleton is stratified by isotropic submanifolds. In contrast, the skeleton of an arbitrary Liouville manifold can be much larger (see e.g. McDuff's examples in [McD91]).

4.3.2 Invariants of Liouville manifolds

The *wrapped Fukaya category* $\mathcal{W}(X)$ of a Liouville manifold (X, λ) is an A_∞ category. The objects of this category are exact, cylindrical Lagrangian submanifolds. The A_∞ operations are constructed by counting pseudoholomorphic curves. The wrapped Fukaya category was first constructed by Abouzaid–Seidel [AS10c] and has become a fundamental invariant in the study of Liouville manifolds.

More generally, one can also consider the wrapped Fukaya category $\mathcal{W}(X, \mathfrak{f})$ of a stopped Liouville sector. The relevant theory was developed by Ganatra–Pardon–Shende (see [GPS18b, Sec. 2] for an overview), following earlier work of Sylvan [Syl19a].

The wrapped Fukaya category is invariant up to quasi-equivalence under deformations of the underlying Liouville manifold/sector. More generally, the wrapped Fukaya category of a stopped Liouville manifold is invariant under deformations which preserve the contactomorphism type of the complement of the stop [GPS18b, Thm. 1.4].

The *symplectic cohomology* $SH^\bullet(X)$ of a Liouville manifold (X, λ) is a graded-commutative ring (meaning that $ab = (-1)^{|a||b|}ba$). The generators can be interpreted as orbits for a Hamiltonian on X which grows suitably fast at infinity, and the differential counts pseudoholomorphic curves. Our grading conventions for symplectic cohomology are uniquely determined by stipulating that the differential increases grading and that the unit lives in degree zero. They match those of e.g. [Sei08a] and [Gan12].

Symplectic cohomology was originally introduced by Floer and Hofer [FH94] and subsequently recast by Viterbo [Vit99]. We refer to [Sei08a] for a beautiful survey which also matches our sign and grading conventions.

Remark 4.3.1 (Orientations/gradings). Note that one must fix additional orientations/grading data in order for the above invariants to be $\mathbb{Z}$-graded and defined over a field of characteristic zero. We generally avoid discussing signs/gradings in this chapter (see instead [GPS18a, Sec. 5.3]).

Remark 4.3.2 (Technical setups). There are multiple technical setups for Floer theoretic invariants on Liouville manifolds in the literature. In particular, there are (at least) three distinct versions of the wrapped Fukaya category of a Liouville manifold. The first is the so-called "linear" version of [AS10c]. The second is the so-called "quadratic" version, which first appeared in [Abo10]. The third is the so-called "localization" version which was first implemented in the literature by Ganatra–Pardon–Shende in [GPS20]. Some standard references (such as [AS10a] and [Gan12]) also construct Floer theoretic invariants using almost-complex structures of "rescaled contact type", which is in contrast to the previously cited references which only use almost-complex structures of contact type.

The fact that all of these setups yield isomorphic theories is a folklore result which is routinely assumed in the literature. However, to the best of our knowledge, detailed proofs for most of these equivalences have yet to appear (see however [Syl19b, Prop. 2.6] concerning the comparison between the quadratic and localization versions). In this chapter, the reader is free to assume throughout that the wrapped Fukaya category is defined according to the localization setup of [GPS20] (this is the only setup available for sectors). However, we will quote facts from [AS10a] and [Gan12] in 4.5.4, and we will sometimes also appeal to certain mirror symmetry results from the literature which are proved using other setups.

4.3.3 Structural facts about wrapped Fukaya categories

We collect some structural properties of wrapped Fukaya categories which will be used in the sequel. Unless otherwise indicated, the theory summarized here is due to Ganatra–Pardon–Shende [GPS20, GPS18b].

Definition 4.3.1 (Def. 1.6 in [GPS18b])**.** A closed subset $\mathfrak{c}$ of a symplectic manifold W is said to be *mostly Lagrangian* if it admits a decomposition $\mathfrak{c} = \mathfrak{c}^{\text{crit}} \cup \mathfrak{c}^{subcrit}$, where $\mathfrak{c}^{crit}$ is a Lagrangian submanifold (in general non-compact) and $\mathfrak{c}^{\text{subcrit}}$ is closed and covered by the smooth image of a second countable manifold of dimension at most $\frac{1}{2}\dim W - 1$. We similarly define the notion of a *mostly Legendrian* subset of a contact manifold.

If $\mathfrak{c}_X \subset X$ is mostly Lagrangian, a *generalized cocore* associated to a connected component of $\mathfrak{c}_X^{\text{crit}}$ is an exact (cylindrical at infinity) Lagrangian which intersects this component transversally in a single point. There is a closely related notion for stops: suppose $\mathfrak{c} \subseteq \partial_\infty X$ is a stop and let $p \in \mathfrak{c}^{\text{crit}}$. Then there is an associated Lagrangian submanifold $L_p \subseteq X$ called the *linking disk*, which is well-defined up to isotopy and depends only on the connected component of $\mathfrak{c}^{\text{crit}}$ containing p.

Example 4.3.3. If the stable manifolds of (X, λ) are properly embedded (this condition holds after a generic perturbation of λ), then the skeleton is mostly Lagrangian. The cocores of the critical handles are in particular generalized cocores. Similarly, the skeleton of a Weinstein hypersurface $V \subseteq \partial_\infty X$ is mostly Legendrian after possibly perturbing V.

If X is a Liouville sector and $\mathfrak{c} \subseteq \partial_\infty X$ is a stop, then there is a natural functor (often called "pushforward" or "stop-removal")

$$\mathcal{W}(X, \mathfrak{c}) \to \mathcal{W}(X). \tag{4.3.2}$$

This functor is characterized as follows when $\mathfrak{c}$ is mostly Legendrian.

Theorem 4.3.2 (see Thm. 1.16 in [GPS18b]). If $\mathfrak{c} \subseteq \partial_\infty X$ is mostly Legendrian, then (4.3.2) induces a quasi-equivalence

$$\mathcal{W}(X, \mathfrak{c})/\mathcal{D} \xrightarrow{\sim} \mathcal{W}(X), \tag{4.3.3}$$

where $\mathcal{D}$ is the collection of linking disks associated to $\mathfrak{c}^{\mathrm{crit}}$.

It is a fundamental fact due independently to Chaintraine–Dimitroglou Rizell–Gigghini–Golovko [CRGG17] and Ganatra–Pardon–Shende [GPS18b, Thm. 1.10] that the wrapped Fukaya category of a Weinstein manifold is generated by cocores. We will need a slightly more refined statement, which follows from essentially the same proof. This is the content of 4.3.3 below.

To set the stage, let (X, λ) be a Liouville manifold. Suppose that the skeleton $\mathfrak{c}_X = \mathfrak{c}^{\mathrm{crit}} \cup \mathfrak{c}^{\mathrm{subcrit}}$ is mostly Lagrangian, and that every connected component of $\mathfrak{c}^{\mathrm{crit}}$ admits a generalized cocore. Let L be an object of $\mathcal{W}(X)$ such that the underlying Lagrangian satisfies the following conditions:

- $L \cap \mathfrak{c}_V \subseteq \mathfrak{c}_V^{\mathrm{crit}}$;
- L intersects $\mathfrak{c}_V^{\mathrm{crit}}$ transversally.

Theorem 4.3.3. With the above assumptions, suppose that $|L \cap \mathfrak{c}_V| = r \in \mathbb{N}$. Then there exists a resolution

$$0 = K_0 \to K_1 \to \cdots \to K_{r-1} \to K_r = L \tag{4.3.4}$$

where each $\mathrm{cone}(K_{i-1} \to K_i)$ is isomorphic in $H^0(\mathcal{W}(X))$ to a generalized cocore.

Proof. The proof is essentially the same as that of [GPS18b, Thm. 1.11]. After possibly modifying L outside a compact set, we may assume that $\lambda|_L$ has a compactly-supported primitive [GPS18b, Lem. 6.1]. Let $f : X \to \mathbb{R}$ be an extension of this primitive to X.

Consider the functors

$$\mathcal{W}(X) \hookrightarrow \mathcal{W}(X \times \mathbb{C}, \mathfrak{c}_X \times \{\pm\infty\}) \xrightarrow{\sim} \mathcal{W}(X \times \mathbb{C}, \mathfrak{c}_X \times \{+\infty\} \cup \mathfrak{c}_X \times \{e^{4\pi i/3} \cdot \infty\}), \quad (4.3.5)$$

where the first arrow is the Künneth embedding [GPS18b, Thm. 1.5] and the second arrow is induced by the isotopy of stops $\mathfrak{c}_X \times \{-\infty\} \rightsquigarrow \mathfrak{c}_X \times \{e^{\pi i \cdot (1+s)}\infty\})$ for $s \in [0, 1/3]$, which is a quasi-equivalence by [GPS18b, Thm. 1.4].

On objects, (4.3.5) sends $L \mapsto L\tilde{\times} i\mathbb{R}$, which is the *cylindrization* of the product Lagrangian $L \times i\mathbb{R}$ as defined in [GPS18b, Sec. 6.2]. The cylindrization is well-defined up to a contractible choice of Lagrangian isotopies. It will be convenient to construct a nice cylindrization. To this end, let $C \gg 1$ and $0 < \epsilon \ll 1$ be constants which will be fixed later. Let $\phi(r, \theta) = \phi(r) : \mathbb{C} \to \mathbb{R}$ be a function with the following properties:

- the derivative of ϕ in the radial direction is non-negative;
- $\phi(r) = 0$ for $r \in [0, C - \epsilon]$ and $\phi(r) = 1$ if $r \geq C$.

To define the cylindrization, we consider the Liouville 1-form $\lambda_X + \lambda_{\mathbb{C}} - d(\phi f)$, where $\lambda_{\mathbb{C}} = r^2 d\theta$. The associated Liouville vector field is

$$(Z_X + \phi X_f) + (\frac{1}{2} r\partial_r + f X_\phi). \quad (4.3.6)$$

Assuming C is large enough, it was shown in [GPS18b, Sec. 6.2] that there is a unique Hamiltonian isotopy Φ which is the identity on a Liouville subdomain of $X \times \mathbb{C}$ and sends $Z_X + \frac{1}{2} r\partial_r$ (i.e. the Liouville vector field of $\lambda_X + \lambda_{\mathbb{C}}$) to (4.3.6).

For $t \in [0, 2/3]$, let $L_t := L \times e^{\pi i \cdot (t+1/2)}\mathbb{R}$. We define $\tilde{L}_t := \Phi(L_t)$ to be the cylindrization.

Consider now the contact manifold $\mathcal{O} = X \times \{r = C\}$, with contact form $\lambda_X + C^2 d\theta$. This defines coordinates on (a region of) $(\partial_\infty(X \times \mathbb{C}), \lambda_X + \lambda_{\mathbb{C}})$. We set $\Lambda_t = \mathcal{O} \cap L_t$ and $\tilde{\Lambda}_t = \mathcal{O} \cap \tilde{L}_t$. By construction, $\tilde{\Lambda}_t \subseteq \mathcal{O}$ is Legendrian. Moreover, by taking

ϵ small enough, we can ensure that $\tilde{L}_t$ is arbitrarily close to the Legendrian lift of L_t in any C^k norm. In particular, the projection of $\tilde{L}_t$ onto $X \times \{r = C, \theta = \pi \cdot (t+1/2)\}$ is arbitrarily close to Λ_t.

We now consider the isotopy $\tilde{L}_t$ for $t \in [0, 2/3]$. The intersection points of $\tilde{L}_t$ with the stop $\mathfrak{c}_X \times \{+\infty\}$ are in correspondence with the intersection points of Λ_t with $\mathfrak{c}_X \times \{\infty\}$, and hence with the intersection points $L \cap \mathfrak{c}_X$. By assumption, these intersections are all transverse and contained in $\mathfrak{c}_X^{\mathrm{crit}}$.

The rest of the argument is identical to the proof of [GPS18b, Thm. 1.11] and is therefore only sketched. We may assume (e.g. by a small perturbation of the stop) that the intersections happens at discrete times $0 < t_1 < \cdots < t_r < 2/3$. At each intersection point t_i, the wrapping exact triangle [GPS18b, Thm. 1.9] gives $L_{t_i+\delta} \to L_{t_i-\delta} \to D_i \xrightarrow{[1]}$, where D_i is a linking disk. One can verify that $L_{t_r+\delta}$ is the zero object (for example, because it is contained in a Liouville sector deformation equivalent to $X \times \mathbb{C}_{\mathrm{Im} \geq 0}$ whose wrapped Fukaya category thus vanishes). Finally, it can be shown that the linking disks correspond to generalized cocores under (4.3.5). □

Remark 4.3.4. 4.3.3 also follows from [CRGG17, Prop. 9.3] (in a slightly less general form). Note that in the statement of [CRGG17, Prop. 9.3], the authors implicitly assume that the Liouville structure has already been put in a "nice" form (see the discussion in [CRGG17, Sec. 2.2, 9.1]). However, these deformations can always be assumed to preserve the skeleton and the cocores set-wise so do not cause any problems.

4.3.4 The Rouquier dimension of wrapped Fukaya categories

We now arrive at the main topic of this chapter.

Definition 4.3.4. Let (X, λ) be a Liouville manifold. We let

$$\mathrm{Rdim}\, \mathcal{W}(X) \in \mathbb{N} \cup \{\infty\} \tag{4.3.7}$$

denote the Rouquier dimension of the wrapped Fukaya category $\mathcal{W}(X)$, as defined in 4.2.7.

More generally, if $(X, \mathfrak{f})$ is a stopped Liouville sector, we let $\operatorname{Rdim} \mathcal{W}(X, \mathfrak{f})$ denote the Rouquier dimension of its wrapped Fukaya category.

Remark 4.3.5 (Orientations/gradings). Of course, the category $\mathcal{W}(X)$ depends on auxiliary orientations/grading data, so $\operatorname{Rdim} \mathcal{W}(X)$ also depends on this data. Thus $\operatorname{Rdim} \mathcal{W}(-)$ should be viewed as an invariant of Liouville manifolds *equipped with orientation/grading data*.

To obtain an honest invariant of Liouville manifolds defined in maximum generality, one could simply choose to work exclusively with $\mathbb{Z}/2$-coefficients and $\mathbb{Z}/2$-gradings. On the other hand, for the purpose of this chapter, we are mostly interested in the Rouquier dimension as a tool for symplectic and algebro-geometric applications. Thus we will typically restrict our discussion to whichever choice of coefficients/gradings is most convenient for computations or applications.

In the following examples, we collect some preliminary facts about the Rouquier dimension of wrapped Fukaya categories.

Example 4.3.6. If X is Weinstein, then it can be shown using either [CRGG17, Thm. 1.1] and [GPS18b, Thm. 1.10] that the diagonal $\Delta \subseteq \mathcal{W}(\overline{X} \times X, -\lambda \oplus \lambda)$ is generated by product-cocores. It then follows by combining 4.7.1 in the appendix with 4.2.10 that

$$\operatorname{Rdim} \mathcal{W}(X) < \infty. \tag{4.3.8}$$

We do not know if there exists a Liouville manifold whose wrapped Fukaya category has infinite Rouquier dimension.

Example 4.3.7. If $\mathcal{W}(X) = 0$, then $\operatorname{Rdim} \mathcal{W}(X) = 0$. This holds of course for subcritical Weinstein manifolds. More generally, the same holds for *flexible* Weinstein manifold, namely those Weinstein manifolds which admit a handlebody presentation

where all critical handles are attached along loose Legendrians.

Example 4.3.8 (Point-like objects). Suppose that $\mathcal{W}(X)$ is equipped with $\mathbb{Z}$-grading and admits a point-like object (e.g. an exact Lagrangian torus). Then it follows from [EL18, Thm. 5.2] that $\operatorname{Rdim} \mathcal{W}(X) \geq \frac{1}{2} \dim_{\mathbb{R}} X$. (For an A_∞ category $\mathcal{C}$, we say that $L \in \mathcal{C}$ is *point-like* if (i) $\mathcal{C}$-module $\mathcal{Y}_L^r$ is proper and; (ii) $\hom(L, L) = \Lambda^\bullet V$ as A_∞ algebras, where the right hand side is the exterior algebra (viewed as an A_∞ algebra) of an r-dimensional graded vector space V which is supported in degree 1. Strictly speaking, [EL18, Thm. 5.2] applies only to proper categories. However, one can verify that the proof only uses the weaker assumption (i).)

Example 4.3.9 (Homological mirror symmetry). If X happens to be mirror to an algebraic variety $X^\vee$, then we can immediately transfer the computations of $\operatorname{Rdim} D^b \operatorname{Coh}(X^\vee)$ discussed in 4.2.6 into computations of $\operatorname{Rdim} \mathcal{W}(X)$. A partial list of homological mirror symmetry results for Weinstein manifolds and sectors is [HK20, LP20, LU18, Gam20, GS17, Pom21]. We leave it to the reader to verify which results from 4.2.6 apply in each of these instances.

4.4 Symplectic flexibility and Orlov's conjecture

Before delving into the technical details, let us illustrate the ideas of this section by considering the example $(X, \lambda) = (T^*M, pdq)$, where M is an n-dimensional closed manifold. Choose a smooth triangulation (see the discussion in Section 4.4.1 for a precise definition) $\mathcal{S}$ of M and denote by $\{v_\alpha\}_{\alpha \in I}$ the set of 0-simplices of $\mathcal{S}$. Denote by $\operatorname{star}(v_\alpha)$ the star of the vertex v_α and let $\overline{\operatorname{star}(v_\alpha)}$ be its closure. Then $\overline{\operatorname{star}(v_\alpha)}$ is an embedded polyhedron. We choose an inward smoothing of each $\overline{\operatorname{star}(v_\alpha)}$ to obtain a smooth manifold with boundary M_α such that the boundaries $\{\partial M_\alpha\}_{\alpha \in I}$ intersect transversely. By [GPS18b, Ex. 1.25], the collection of cotangent bundles $\{T^*M_\alpha\}_{\alpha \in I}$ forms a *sectorial cover* (see [GPS18b, Def. 1.24]) of the Weinstein mani-

fold (T^*M, pdq). Then [GPS18b, Thm. 1.27] asserts a pre-triangulated equivalence

$$\operatorname{hocolim}_{\emptyset \neq J \subset I} \mathcal{W}(\bigcap_{\alpha \in J} T^*M_\alpha) \xrightarrow{\sim} \mathcal{W}(T^*M). \tag{4.4.1}$$

The sectorial cover $\{T^*M_\alpha\}_{\alpha \in I}$ admits the following special properties:

1. it has depth at most $n+1$ if the inward smoothing is small enough; i.e. if $\bigcap_{\alpha \in J} T^*M_\alpha \neq \emptyset$, then $|J| \leq n+1$;
2. each local piece $\bigcap_{\alpha \in J} T^*M_\alpha$ is Weinstein homotopic to the cotangent bundle of a closed n-dimensional ball $\overline{D}^n$.

Note that $\mathcal{W}(T^*\overline{D}^n)$ is equivalent to $\operatorname{Perf}(k[A_1])$ where A_1 is the quiver with only one vertex so $\operatorname{Rdim} \mathcal{W}(T^*\overline{D}^n) = 0$ by Proposition 4.2.13. Then Lemma 4.2.12 and 4.2.11 and the fact that this cover has depth at most $n+1$ implies $\operatorname{Rdim} \mathcal{W}(T^*M) \leq n$ because of (4.4.1).

We wish to apply the above ideas to general Weinstein manifolds or sectors with arboreal skeleta (see Section 4.4.2). For purely technical reasons, it is more convenient to implement this argument using microlocal sheaf theory. We therefore study the *Nadler–Shende microlocal sheaf* (see Section 4.4.3) as an alternative model of the wrapped Fukaya category, and appeal to [GPS18a, Thm. 1.4] to translate our results back into the framework of wrapped Fukaya categories. The Nadler–Shende sheaf just depends on the topology of the skeleta and satisfies a local-to-global descent formula similar to (4.4.1). Then we can cover the skeleta by the stars of topological triangulations and conclude our result Theorem 4.4.13 as above.

We will use the notations $\operatorname{skel}(X)$ and $\mathfrak{c}_X$ interchangeably throughout this section to denote the skeleton of (X, λ).

4.4.1 Stratifications and Goresky's triangulation theorem

A *stratification* $\mathcal{S}$ of a topological space X is a locally finite decomposition $\{X_\sigma\}_{\sigma\in\mathcal{S}}$ of X into disjoint, locally closed subsets (called strata[7]) which satisfies the frontier condition: if $\overline{X_\alpha}\cap X_\beta\neq\emptyset$, then $X_\beta\subseteq X_\alpha$. The pair $(X,\mathcal{S})$ is said to be a *stratified space*. Note that $\mathcal{S}$ is naturally a partially ordered set, where we set $X_\alpha\leq X_\beta$ if and only if $X_\alpha\subseteq\overline{X_\beta}$.

Given a stratum X_α, the (open) star of X_α is defined to be the union of strata from $\mathcal{S}$ whose closures contain X_α and is denoted by $\mathrm{star}(\alpha)$. The *link* of X_α is denoted by $\mathrm{link}(\alpha)$ and defined as the closure of $\mathrm{star}(\alpha)$ minus its interior.

Given a fixed topological space X, a stratification $\mathcal{S}'$ of X is said to be a *refinement* of some coarser stratification $\mathcal{S}$ of X if every stratum in $\mathcal{S}'$ is contained in a stratum in $\mathcal{S}$ (equivalently, every stratum in $\mathcal{S}$ is a union of strata in $\mathcal{S}'$).

If X is a topological subspace of a manifold (without loss of generality $X\subseteq\mathbb{R}^N$), then a stratification $\mathcal{S}$ of X is said to be C^μ (for $\mu\in\mathbb{Z}_{\geq 0}\cup\{\infty\}$) if each stratum is a locally closed C^μ submanifold.

Definition 4.4.1. A C^1 stratification $\mathcal{S}$ of $X\subset\mathbb{R}^N$ is called *Whitney* if the following condition is satisfied.

Whitney (b): given any pair of strata X_α and X_β for $\alpha,\beta\in\mathcal{S}$ with $X_\alpha<X_\beta$, suppose that $x_i\in X_\beta$ and $y_i\in X_\alpha$ are a sequences of points both converging to $y\in X_\alpha$. Suppose also that $T_{x_i}X_\beta$ converges to a subspace $T\subseteq T_y\mathbb{R}^N$, and the secant lines $\overline{x_iy_i}$ also converge to a line $\ell\subseteq T_y\mathbb{R}^N$. Then $\ell\subseteq T$.

Next we review the notion of a triangulation. A k-dimensional *simplex* is a subset of $\mathbb{R}^d$ (for some $d>0$) which is the convex hull of some finite set of independent points $p_0,\ldots,p_k$. The *faces* of a simplex are defined in the obvious way. A *simplicial*

[7]The strata are required to be connected under some conventions. We find it more convenient to also allow disconnected strata – however, this is just a matter of terminology which does not affect the substances of the arguments.

complex (see [Mun16, Def. 7.1]) is a set $\mathcal{K}$ of simplices in $\mathbb{R}^d$ (for some fixed $d > 0$) which satisfies the following conditions:

1. every face of a simplex $\sigma \in \mathcal{K}$ is also in $\mathcal{K}$;

2. the non-empty intersection of any two simplices $\sigma_1, \sigma_2 \in \mathcal{K}$ is a face of both σ_1 and σ_2;

3. local finiteness (i.e. each point of $|\mathcal{K}|$ is contained in finitely many simplices, where $|\mathcal{K}|$ is the topological space formed by the union of the simplices from $\mathcal{K}$).

Note that the space $|\mathcal{K}|$ admits a stratification whose strata are given by interiors of simplices from $\mathcal{K}$ and we denote this stratification by $\mathcal{S}_{\mathcal{K}}$. Each stratum of $\mathcal{S}_{\mathcal{K}}$ is naturally a smooth submanifold of $\mathbb{R}^d$ and it is easy to see that $\mathcal{S}_{\mathcal{K}}$ defines a Whitney stratification.

The notion of a simplicial subcomplex is defined in the obvious way.

Definition 4.4.2. Given a smooth manifold M, a *smooth triangulation* of M consists of the following data:

1. a simplicial complex $\mathcal{K}$ and a simplicial subcomplex $\mathcal{L} \subset \mathcal{K}$;

2. a homeomorphism $f : |\mathcal{K}| \setminus |\mathcal{L}| \to X$ such that for each simplex $\sigma \in \mathcal{K}$ and $x \in \sigma \setminus |\mathcal{L}|$, there is a neighborhood U of x in the plane containing σ and a smooth embedding $\tilde{f} : U \to X$ extending $f|_{U \cap \sigma}$.

Definition 4.4.3. For a Whitney stratified space $(X, \mathcal{S})$ with smooth strata, a triangulation of $(X, \mathcal{S})$ is given by a simplicial complex $\mathcal{K}$ and a homeomorphism $f : |\mathcal{K}| \to X$ such that for each stratum X_α, the set $f^{-1}(X_\alpha)$ is the underlying topological space of some simplicial subcomplex of $\mathcal{K}$ and $f : f^{-1}(X_\alpha) \to X_\alpha$ is a smooth triangulation.

Given a triangulation of $(X, \mathcal{S})$ as in Definition 4.4.3, the stratification $\mathcal{S}_{\mathcal{K}}$ on $|\mathcal{K}|$ defines a (topological) stratification of X and we call a stratification of this from *triangulation*. It is easy to see that this stratification is a refinement of the stratification $\mathcal{S}$. The following theorem of Goresky asserts that it is always possible to refine a Whitney stratification into a triangulation.

Theorem 4.4.4 (Goresky [Gor78b]). Every Whitney stratified space $(X, \mathcal{S})$ with smooth strata admits a triangulation.

Note that Theorem 4.4.4 does not prove that the stratification induced by the triangulation is Whitney. It is an open question whether the theorem can be strengthened in that way.

For the purpose of our application, we introduce some more notions. Given two stratified topological spaces $(X, \mathcal{S})$ and $(X', \mathcal{S}')$, a *stratified homeomorphism* $(X, \mathcal{S}) \to (X', \mathcal{S}')$ is a homeomorphism $X \to X'$ which sends strata to strata and defines a bijection between strata.

Definition 4.4.5. For a stratified topological space $(X, \mathcal{S})$, a *topological triangulation* of $(X, \mathcal{S})$ is defined by a simplicial complex $\mathcal{K}$ and a homeomorphism $f : |\mathcal{K}| \to X$ such that $\mathcal{S}$ admits a refinement $\mathcal{S}'$ with the property that $f : (|\mathcal{K}|, \mathcal{S}_{\mathcal{K}}) \to (X, \mathcal{S}')$ defines a stratified homeomorphism.

Then Theorem 4.4.4 implies the following result which suffices for our application.

Corollary 4.4.1. Any Whitney stratified space with smooth strata admits a topological triangulation. □

4.4.2 Arboreal singularities

We review the necessary inputs from the arborealization program [AGEN21, AGEN19, AGEN20] and develop some topological aspects of arboreal Lagrangians needed for our applications.

Arborealization of polarized Weinstein manifolds

A *signed rooted tree* $\mathfrak{T} = (T, \rho, \epsilon)$ consists of an acyclic finite graph T, a distinguished vertex ρ which is called the *root*, and a decoration ϵ of edges which are not adjacent to ρ by signs $\pm$. Let $n(\mathfrak{T})$ be the number of vertices of T which are not the root. To any such $\mathfrak{T}$, [AGEN21, Def. 2.19] associates a closed Lagrangian subset

$$L_{\mathfrak{T}} \subset T^*\mathbb{R}^{n(\mathfrak{T})} \tag{4.4.2}$$

stratified by isotropic submanifolds, where $T^*\mathbb{R}^{n(\mathfrak{T})}$ is endowed with the standard symplectic structure.

Definition 4.4.6 (Def. 3.1 in [AGEN21])**.** A subset L in a symplectic manifold (M^{2n}, ω) is called a *(boundaryless) arboreal Lagrangian* if, for any $p \in L$, there is a neighborhood U containing p and a symplectic embedding

$$(U, U \cap L) \hookrightarrow (T^*\mathbb{R}^{n(\mathfrak{T})} \times T^*\mathbb{R}^{n-n(\mathfrak{T})}, L_{\mathfrak{T}} \times \mathbb{R}^{n-n(\mathfrak{T})}), \tag{4.4.3}$$

for some rooted tree $\mathfrak{T}$. Here $\mathbb{R}^{n-n(\mathfrak{T})} \subset T^*\mathbb{R}^{n-n(\mathfrak{T})}$ is the 0-section. We say that L is an *arboreal Lagrangian with boundary* (or just an *arboreal Lagrangian*) if for any $q \in L$, there is a neighborhood U containing q and a symplectic embedding modeled either on (4.4.3) or

$$(U, U \cap L) \hookrightarrow (T^*\mathbb{R}^{n(\mathfrak{T})} \times T^*\mathbb{R}^{n-n(\mathfrak{T})-1} \times T^*\mathbb{R}_{\geq 0}, L_{\mathcal{T}} \times \mathbb{R}^{n-n(\mathfrak{T})-1} \times \mathbb{R}_{\geq 0}). \tag{4.4.4}$$

The following theorem is due to Álvarez-Gavela–Eliashberg–Nadler (see Thm. 1.5 and Rmk. 1.6(iii) in [AGEN20]).

Theorem 4.4.7 (Arborealization). Let (X, λ) be a Weinstein manifold which has a polarization. Let $A \subseteq \partial_\infty X$ be a Weinstein hypersurface. Then there exists the following data:

1. A family of Liouville forms $(\lambda_t)_{t\in[0,1]}$ on X with $\lambda_0 = \lambda$ such that (X, λ_t) and $(A, \lambda_t|_A)$ are both Weinstein homotopies, and such that $\mathrm{skel}((X, \lambda_1), A)$ is arboreal with smooth boundary and possibly with "corner-type" singularities (i.e. locally modeled on $(0_{\mathbb{R}^n} \cup \mathit{conormal}(\{x_n = 0, x_{n-1} \geq 0\})) \subseteq T^*\mathbb{R}^n$);

2. A Weinstein hypersurface $B \subseteq \partial_\infty X$ such that $\mathrm{skel}(B) \setminus \mathrm{skel}(A, \lambda_1)$ is a smooth Legendrian with finitely many connected components, and such that $\mathrm{skel}((X, \lambda_1), B)$ is arboreal. (To be clear, we do not allow any "corner-type" singularities here).

Remark 4.4.1. When $A = \emptyset$, the skeleton $\mathrm{skel}(X, \lambda_1) = \mathrm{skel}((X, \lambda_1), A)$ from (1) of 4.4.7 can be made into an arboreal Lagrangian with *smooth boundary*.

Remark 4.4.2. If X has real dimension at most 4, then both conclusions of 4.4.7 hold without assuming the existence of a global Lagrangian plane field. (This is essentially a consequence of work of Starkston [Sta18], although strictly speaking the result there is only stated under the assumption that $A = \emptyset$, and produces an arboreal skeleton with "corners".)

The singularity-order stratification

We introduce a stratification on arboreal Lagrangians according to the local behavior of the singularities.

Definition 4.4.8. Suppose $L \subset (M^{2n}, \omega)$ is an arboreal Lagrangian. Given a rooted tree $\mathfrak{T}$ with $n(\mathfrak{T}) \leq n$, define $L(\mathfrak{T})$ to be the subset of L such that for any $p \in L(\mathfrak{T})$, there is a neighborhood U containing p and a symplectic embedding modeled on (4.4.3). For any rooted tree $\mathfrak{T}$ with $n(\mathfrak{T}) \leq n-1$, define $L'(\mathfrak{T})$ to be the subset of L such that for any $q \in L'(\mathfrak{T})$, there is a neighborhood U containing q and a symplectic embedding modeled on (4.4.4).

Then the following decomposition, which is denoted by $\mathcal{S}_{\text{arb}}$,

$$L = \bigcup_{\mathfrak{T}, n(\mathfrak{T}) \leq n} L(\mathfrak{T}) \cup \bigcup_{\mathfrak{T}, n(\mathfrak{T}) \leq n-1} L'(\mathfrak{T}) \tag{4.4.5}$$

defines a C^∞ stratification of L called *the singularity-order stratification.*

Proposition 4.4.9. For any arboreal Lagrangian $L \subset (M, \omega)$, the singularity-order stratification $\mathcal{S}_{\text{arb}}$ on L is Whitney.

The rest of this subsection is devoted to the proof of Proposition 4.4.9.

Suppose we have a possibly singular Legendrian $\Lambda \subset S^*\mathbb{R}^n$ endowed with a C^∞ Whitney stratification $\mathcal{S}_\Lambda$ by isotropic submanifolds. When Λ is connected, assume that (i) the front projection restricted to each stratum from $\mathcal{S}_\Lambda$ defines an embedding into the 0-section $\mathbb{R}^n \subset T^*\mathbb{R}^n$; (ii) different strata have disjoint images under the projection π. If Λ is disconnected, we require that (i) each component satisfies the previous conditions and; (ii) given connected components $\Lambda_1, \ldots, \Lambda_k$ together with strata $Y_1 \subset \Lambda_1, \ldots, Y_k \subset \Lambda_k$, the restriction of the front projection $\pi : S^*\mathbb{R}^n \to \mathbb{R}^n$ along $Y_1 \cup \cdots \cup Y_k$ is self-transverse.

Then we can consider the singular Lagrangian $\text{Cone}(\Lambda) \cup \mathbb{R}^n \subset T^*\mathbb{R}^n$ (where $\text{Cone}(\Lambda)$ is the Liouville cone and we write $\mathbb{R}^n \equiv 0_{\mathbb{R}^n}$). The Lagrangian $\text{Cone}(\Lambda) \cup \mathbb{R}^n$ admits a natural stratification $\mathcal{S}_{\text{Cone}(\Lambda) \cup \mathbb{R}^n}$ as follows. Given any singular isotropic Λ', note that $\pi(\Lambda') \subset \text{Cone}(\Lambda')$ and define $C^\circ(\Lambda') := \text{Cone}(\Lambda') \setminus \pi(\Lambda')$. Then the union of all $C^\circ(Y)$ for $Y \in \mathcal{S}_\Lambda$ defines a decomposition of $C^\circ(\Lambda)$ and we declare each $C^\circ(Y)$ to be a stratum from $\mathcal{S}_{\text{Cone}(\Lambda) \cup \mathbb{R}^n}$. For the 0-section $\mathbb{R}^n$, define

$$(\text{Cone}(\Lambda) \cup \mathbb{R}^n)_{Y_1, \ldots, Y_k} = \bigcap_{i=1}^k \pi(Y_i) \setminus \left(\bigcup_{Y'} \pi(Y') \cap \bigcap_{i=1}^k \pi(Y_i) \right) \tag{4.4.6}$$

where $Y_1 \subset \Lambda_1, \ldots, Y_k \subset \Lambda_k$ are strata of distinct components $\Lambda_1, \ldots, \Lambda_k$ and Y' is any stratum of Λ different from $Y_1, \ldots, Y_k$. Finally, define $(\mathbb{R}^n)_\emptyset$ to be the complement of

the union of all $(\mathrm{Cone}(\Lambda)\cup\mathbb{R}^n)_{Y_1,\ldots,Y_k}$ in the 0-section. Then we obtain a decomposition

$$\mathbb{R}^n = (\mathbb{R}^n)_\emptyset \cup \bigcup_{Y_1,\ldots,Y_k} (\mathrm{Cone}(\Lambda)\cup\mathbb{R}^n)_{Y_1,\ldots,Y_k} \tag{4.4.7}$$

and the decomposition $\mathcal{S}_{\mathrm{Cone}(\Lambda)\cup\mathbb{R}^n}$ is simply

$$\mathrm{Cone}(\Lambda)\cup\mathbb{R}^n = \bigcup_Y C^\circ(Y) \cup (\mathbb{R}^n)_\emptyset \cup \bigcup_{Y_1,\ldots,Y_k} (\mathrm{Cone}(\Lambda)\cup\mathbb{R}^n)_{Y_1,\ldots,Y_k}. \tag{4.4.8}$$

$\mathcal{S}_{\mathrm{Cone}(\Lambda)\cup\mathbb{R}^n}$ indeed defines a stratification as the frontier condition holds.

Lemma 4.4.10. The stratification $\mathcal{S}_{\mathrm{Cone}(\Lambda)\cup\mathbb{R}^n}$ is Whitney.

Proof. It suffices to check condition *Whitney (b)* for any pair of strata X_α and X_β with $X_\alpha < X_\beta$ from $\mathcal{S}_{\mathrm{Cone}(\Lambda)\cup\mathbb{R}^n}$.

If both X_α and X_β are strata of the form $C^\circ(Y)$, *Whitney (b)* holds because $\mathcal{S}_\Lambda$ is a Whitney stratification.

If $X_\alpha \subset \mathbb{R}^n$ and X_β is of the form $C^\circ(Y)$, then X_α is necessarily of the form $(\mathrm{Cone}(\Lambda)\cup\mathbb{R}^n)_{Y_1,\ldots,Y_k}$ such that there exists Y_i with $Y_i \subset \overline{Y}$. *Whitney (b)* holds by observing that the canonical stratification on $[0,\infty)$ by $\{\{0\},(0,\infty)\}$ is Whitney.

If both X_α and X_β are contained in the 0-section $\mathbb{R}^n$, there are two cases. When Λ only has one connected component, the projection π is assumed to be a topological embedding along Λ and is a smooth immersion when restricted to each stratum. Then *Whitney (b)* holds because it holds for the stratification $\mathcal{S}_\Lambda$. If Λ has multiple connected components, we make the following observation. Given two smooth submanifolds N_1 and N_2 of a smooth manifold M, if N_1 and N_2 intersect transversely, then the decompositions $N_1 \cup N_2 = (N_1 \setminus N_2) \cup (N_1 \cap N_2) \cup (N_2 \setminus N_1)$ and $M = (M \setminus N_1 \cup N_2) \cup (N_1 \setminus N_2) \cup (N_1 \cap N_2) \cup (N_2 \setminus N_1)$ are all Whitney stratifications. The obvious generalizations for multiple submanifolds are easy to verify. Then *Whitney (b)* for the case when Λ has multiple connected components just follows

from this observation, the discussion for one single component and the transversality assumption of the front projection π for strata from different connected components. (Note that it is possible that the front projections of components of Λ are disjoint from each other, but *Whitney (b)* holds nevertheless similar to the discussion of Λ with only 1 component.) □

Proof of Proposition 4.4.9. This is because of Lemma 4.4.10 and the axiomatic characterization of arboreal Lagrangians [AGEN21, Def. 1.1]. Note that the stratification constructed in Lemma 4.4.10 coincides with the singularity-order stratification using the inductive description of arboreal Lagrangians from [AGEN21, Def. 1.1]. □

4.4.3 Upper bounds on the Rouquier dimensions

Now we combine microlocal sheaf theory computations with the main result of [GPS18a] to obtain upper bounds for the Rouquier dimension of wrapped Fukaya categories.

Microlocal sheaves

Given a topological space X, one can consider sheaves/cosheaves valued in a symmetric monoidal stable ∞-category on X. For foundational material on sheaves of categories we refer to [ST16, Sec. 3.1], [NS20, Rmk. 5.1] and the references therein.

If $(X, \mathcal{S})$ is a stratified space, then a sheaf of (stable ∞-)categories on X is said to be *constructible* with respect to $\mathcal{S}$ if its restriction to each stratum is locally constant. Thus constructibility with respect to $\mathcal{S}$ is a *property*. Under (rather weak) hypotheses on $(X, \mathcal{S})$, it is a theorem that the full subcategory of $\mathcal{S}$-constructible sheaves is equivalent to a certain representation category of the "infinity exit-path category"; see [ST16, Sec. 4.2.1].

We record the following standard lemma:

Lemma 4.4.11. Let $\mathcal{F}$ be a sheaf of (stable ∞-)categories on a simplicial complex $|\mathcal{K}|$

which is constructible with respect to the canonical stratification $\mathcal{S}_{\mathcal{K}}$. Given $\sigma \in \mathcal{K}$, we have $\mathcal{F}(\mathrm{star}(\sigma)) = \mathcal{F}_c$, for any $c \in |\sigma|$.

Proof. There is a homeomorphism $\mathrm{star}(s) = [0,1) \times \mathrm{link}(\sigma)/(\{0\} \times \mathrm{link}(\sigma))$ which is compatible with the stratification. So p admits a neighborhood base by open sets $U_t := [0,t) \times \mathrm{link}(\sigma)/(\{0\} \times \mathrm{link}(\sigma))$. The restriction maps are isomorphisms since $\mathcal{F}$ is constructible: indeed, the stratification is independent of t and the restriction maps are isomorphisms on each stratum. See also [ST16, Sec. 4.2.1]. □

Suppose now that (X, λ) is a Weinstein manifold of dimension $2n$, equipped with a polarization $\xi \subset TX$ (or more generally, with "Maslov data" [NS20, Sec. 10]). Given any conic subset $\Lambda \subseteq X$, there is a sheaf of stable ∞-categories $\mu sh_\Lambda(-)$ on Λ, called the *Nadler–Shende sheaf*, defined in [She17, NS20].[8] For the purposes of this chapter, we will always assume that μsh takes values in the dg derived category of $\mathbb{Z}$ or $\mathbb{Z}/2$-graded k-modules.

We emphasize that although the *construction* of μsh_Λ uses symplectic topology (i.e. it uses the fact that Λ is a conical subset of a Liouville manifold), μsh_Λ itself is not a symplectic object: it is just a sheaf of categories on the topological space Λ.

Here are some facts about the Nadler–Shende sheaf for Weinstein pairs with arboreal skeleton. Denote by $\mu sh^c_\Lambda(-)$ the assignment $U \mapsto \mu sh_\Lambda(U)^c$, where $\mu sh_\Lambda(U)^c$ is the full subcategory of compact objects of $\mu sh_\Lambda(U)$.

Proposition 4.4.12. Suppose that (X, A) is a Weinstein pair (A could be empty), X is endowed with a polarization ξ, and $\mathrm{skel}(X, A) = \mathfrak{c}_{X,A} \subseteq X$ is an arboreal Lagrangian.

1. The sheaf $\mu sh_{\mathfrak{c}_{X,A}}$ is *constructible* with respect to the singularity-order stratification.

[8] This sheaf of categories is often called the *Kashiwara–Schapira stack* in the literature. We find this terminology misleading since $\mu sh_\Lambda(-)$ is not a stack and was not, to our knowledge, defined by Kashiwara–Schapira (at least not in its current form).

2. Suppose $p \in \mathfrak{c}_{X,A} \subset X$ has a neighborhood U containing p and a symplectic embedding modeled on (4.4.3). Let T be the underlying tree of $\mathfrak{T}$ and we endow T with the natural quiver structure $\overrightarrow{T}$ with arrows pointing towards the root. Then the stalk of $\mu sh^c_{\mathfrak{c}_{X,A}}(-)$ at p satisfies

$$\mu sh^c_{\mathfrak{c}_X}(-)_p \simeq \mathrm{Perf}(k[\overrightarrow{T}]). \tag{4.4.9}$$

In particular, the stalks are independent of the signs on the edges of $\mathfrak{T}$.

3. The category $\mu sh_{\mathfrak{c}_{X,A}}(U)$ is compactly-generated *for any* open set $U \subseteq \mathfrak{c}_{X,A}$. Hence $\mu sh^c_{\mathfrak{c}_{X,A}}(-)$ forms a cosheaf of categories on $\mathfrak{c}_{X,A}$. In particular, if $\{U_\sigma\}_{\sigma\in\Sigma}$ is an open cover of Λ closed under intersections and indexed by a finite poset Σ, then

$$\mu sh^c_{\mathfrak{c}_{X,A}}(\mathfrak{c}_{X,A}) \simeq \mathrm{hocolim}_{\sigma\in\Sigma}\, \mu sh^c_{\mathfrak{c}_{X,A}}(U_\sigma). \tag{4.4.10}$$

Proof. (1) follows from combining 4.4.6 with [NS20, Lem. 5.3]. (2) was proved by [Nad17, Thm. 4.4] in a somewhat different language since the definition of μsh^c_Λ was not in the literature (cf. also [She17, Thm. 4] for our more modern formulation.) For (3), combine 4.4.9 and [NS20, Rmk. 1.6, 1.7]. (More precisely, for any open set U, it follows from 4.4.9 that $U \cap \mathfrak{c}_{X,A}$ is Whitney stratified by isotropics. Hence it follows from [NS20, Rmk. 1.6, 1.7] that $\mu sh_{\mathfrak{c}_{X,A}}(U)$ is compactly-generated by microstalks, which are corepresentable.) Note that this property holds with much weaker assumptions on the skeleton. □

It turns out these properties of microlocal sheaves give rather strong control on the Rouquier dimensions of the sheaf of global sections, after using the topological results from Section 4.4.2.

Theorem 4.4.13. Suppose that (X^{2n}, A) is a Weinstein pair (A could be empty), X is endowed with a polarization ξ, and $\mathrm{skel}(X, A) = \mathfrak{c}_{X,A} \subseteq X$ is an arboreal Lagrangian.

Then

$$\operatorname{Rdim} \mu sh^c_{\mathfrak{c}_{X,A}}(\mathfrak{c}_{X,A}) \leq n \text{ if } n \leq 3 \text{ and} \tag{4.4.11}$$

$$\operatorname{Rdim} \mu sh^c_{\mathfrak{c}_{X,A}}(\mathfrak{c}_{X,A}) \leq 2n-3 \text{ for } n \geq 4. \tag{4.4.12}$$

Proof. We prove the theorem assuming $A \neq \emptyset$ and the case $A = \emptyset$ follows with a similar (but simpler) proof. Note that $\mathfrak{c}_{X,A}$ is the mapping cone $\mathfrak{c}_A \times [0,\infty)$ glued to $\mathfrak{c}_X$ along $\mathfrak{c}_A \times \{0\}$ using the Liouville flow. Choose $N > 0$ and consider the "truncation" $\Lambda_N := \mathfrak{c}_X \cup \mathfrak{c}_A \times [0,N]$. Note that Λ_N is an arboreal Lagrangian with boundary. By Proposition 4.4.9, the singularity-order stratification $\mathcal{S}_{\text{arb}}$ on Λ_N defines a Whitney stratification. Note that $\mathcal{S}_{\text{arb}}$ restricts to a Whitney stratification $\mathcal{S}^{\partial}_{\text{arb}}$ along the "boundary" $\mathfrak{c}_A \times \{N\} \subset \Lambda_N$.

Using Corollary 4.4.1, we can find a topological triangulation $\mathcal{S}_{\mathcal{K}}$ of $(\Lambda_N, \mathcal{S}_{\text{arb}})$ which restricts to a topological triangulation $\mathcal{S}^{\partial}_{\mathcal{K}}$ of $(\mathfrak{c}_A = \mathfrak{c}_A \times \{N\}, \mathcal{S}^{\partial}_{\text{arb}})$. In fact, we can extend the stratification $\mathcal{S}_{\mathcal{K}}$ induced by $\mathcal{K}$ to a stratification $\tilde{\mathcal{S}}_{\mathcal{K}}$ of the non-compact space $\mathfrak{c}_{X,A}$ simply by introducing strata of the form $X_\alpha \times (N,\infty)$ where X_α is a stratum from the triangulation $\mathcal{S}^{\partial}_{\mathcal{K}}$.

Then the stars of the stratification $\tilde{\mathcal{S}}_{\mathcal{K}}$ on $\mathfrak{c}_{X,A}$ define an open cover $\{U_\sigma\}_\sigma$. This cover has the following nice properties.

1. This cover has depth at most $n+1$: given any chain of inclusions $U_{\sigma_1} \subsetneq \cdots \subsetneq U_{\sigma_k}$, we have $k \leq n+1$.

2. Given any $\operatorname{star}(X_\alpha)$ for a stratum X_α from $\tilde{\mathcal{S}}_{\mathcal{K}}$, the category $\mu sh^c_{\mathfrak{c}_{X,A}}(X_\alpha)$ is equivalent to $\operatorname{Perf}(k[\overrightarrow{T}])$ for some quiver $\overrightarrow{T}$ with at most $n+1$ vertices. Indeed, this follows from 4.4.12(2) and Lemma 4.4.11.

The theorem is then a consequence of the local-to-global formula (4.4.10), 4.2.12 and 4.2.11. Indeed, note that any quiver $\overrightarrow{T}$ with at most 4 vertices must be of Dynkin ADE type so $\operatorname{Rdim} \operatorname{Perf}(k[\overrightarrow{T}]) = 0$ by Proposition 4.2.13. This proves (4.4.11). For

$n \geq 4$, note that for any chain of inclusions $U_{\sigma_1} \subsetneq \cdots \subsetneq U_{\sigma_k}$, there are least 4 of the members satisfying $\mu sh^c_{\mathfrak{c}_{X,A}}(U_{\sigma_i}) = 0$ and (4.4.12) is proved again by using Proposition 4.2.13. □

Corollary 4.4.2. Suppose (X^{2n}, A) is a Weinstein pair and (X, λ) is equipped with a polarization. Then

$$\operatorname{Rdim} \mathcal{W}(X, A) \leq n \text{ if } n \leq 3 \text{ and} \tag{4.4.13}$$

$$\operatorname{Rdim} \mathcal{W}(X, A) \leq 2n - 3 \text{ for } n \geq 4. \tag{4.4.14}$$

Proof. By [GPS18a, Cor. 7.28], we can assume that the Weinstein manifold (X, λ) and Weinstein hypersurface A are real analytic and the relative core $\mathfrak{c}_{X,A}$ is subanalytic singular isotropic. Then [GPS18a, Thm. 1.4] asserts the equivalence

$$\operatorname{Perf} \mathcal{W}(X, A) \simeq \mu sh^c_{\mathfrak{c}_{X,A}}(\mathfrak{c}_{X,A}). \tag{4.4.15}$$

Because X is endowed with a polarization, we can find a Weinstein homotopy $\{\lambda_t\}_{0 \leq t \leq 1}$ and a Weinstein hypersurface B with $\lambda_0 = \lambda$ such that for the Weinstein structure λ_1 we have $\operatorname{skel}(B) \setminus \operatorname{skel}(A)$ being a smooth Legendrian and $\mathfrak{c}(X, B) = \operatorname{skel}((X, \lambda_1), B)$ is arboreal by Theorem 4.4.7. Then using Lemma 4.2.1 and the stop removal result [NS20, Thm. 1.5 (2)], we see that for the Weinstein structure λ_1,

$$\operatorname{Rdim} \mu sh^c_{\mathfrak{c}_{X,A}}(\mathfrak{c}_{X,A}) \leq \operatorname{Rdim} \mu sh^c_{\mathfrak{c}_{X,B}}(\mathfrak{c}_{X,B}) \tag{4.4.16}$$

because $\mu sh^c_{\mathfrak{c}_{X,A}}(\mathfrak{c}_{X,A})$ is a quotient of $\mu sh^c_{\mathfrak{c}_{X,B}}(\mathfrak{c}_{X,B})$. The corollary is thus proved by Theorem 4.4.13 the invariance of $\mu sh^c_{\mathfrak{c}_{X,A}}(\mathfrak{c}_{X,A})$ under Weinstein homotopy [NS20, Thm. 1.9(2), Rmk. 8.16]. □

4.4.4 Applications to Orlov's conjecture

Now we can apply the results obtained above to resolve more cases of Orlov's conjecture 4.1.1.

Proof of Theorem 4.1.1. Given a quasi-projective variety Y over $\mathbb{C}$, the assumption on homological mirror symmetry identifies $D^b \operatorname{Coh}(Y)$ with the derived (partially) wrapped Fukaya category of some Weinstein pair satisfying the assumptions of Corollary 4.4.2. Therefore we have $\operatorname{Rdim} D^b \operatorname{Coh}(Y) \leq \dim_{\mathbb{C}} Y$. On the other hand, [Rou08, Prop 7.17] tells us that $\operatorname{Rdim} D^b \operatorname{Coh}(Y) \geq \dim_{\mathbb{C}} Y$. Therefore the Orlov conjecture holds. □

Remark 4.4.3. When $n \geq 4$ under the same assumptions, we obtain the bound

$$\operatorname{Rdim} D^b \operatorname{Coh}(Y) \leq 2n - 3. \tag{4.4.17}$$

This improves the general upper bound $\operatorname{Rdim} D^b \operatorname{Coh}(Y) \leq 2n$ as in [Rou08, Prop. 7.9].

Next we list the new instances of Orlov's conjecture covered by Theorem 4.1.1.

Example 4.4.4 (Log Calabi–Yau surfaces). In dimension 2, we can consider non-toric examples arising in the work of Keating [Kea18], as we now explain.

Let $Y_{p,q,r}$ be the complex algebraic variety constructed as follows. We consider the toric anti-canonical divisor of $D \subseteq \mathbb{P}^2$, which is a wheel of three $\mathbb{P}^1$. Pick p copies of a point lying on the first component, q copies of a point lying on the second component and r copies lying on the third. Assume these points are colinear, lying in the interior of each irreducible component, and that $1/p + 1/q + 1/r < 1$. Now blow up $\mathbb{CP}^2$ at these points. (By definition, the blowup along a point $x \in D$ with multiplicity $m > 1$ means that we first blow up at x, then we blow up at the the intersection of the exceptional divisor with the strict transform of D, and so forth m times).

According to the main result of [Kea18], there are equivalences of categories

$$D^b \operatorname{Coh}(Y_{p,q,r}) \cong \operatorname{Perf} \mathcal{W}(\mathcal{T}_{p,q,r}, V), \quad D^b \operatorname{Coh}(Y_{p,q,r} - D) \cong \operatorname{Perf} \mathcal{W}(\mathcal{T}_{p,q,r}) \tag{4.4.18}$$

where

- $\mathcal{T}_{p,q,r}$ is a Weinstein 4-manifold (a certain Milnor fiber in $\mathbb{C}^3$);
- $V \subseteq \partial_\infty \mathcal{T}_{p,q,r}$ is a Weinstein hypersurface;
- (by abuse of notation) D denotes the strict transform of the toric-anticanonical divisor on $\mathbb{P}^2$.

In particular, $\mathcal{T}_{p,q,r}$ is a complete intersection, so it admits a polarization. Therefore, we see Conjecture 4.1.1 holds for $Y_{p,q,r}$ and $Y_{p,q,r} - D$.

Example 4.4.5 (Three-dimensional toric varieties). Homological mirror symmetry for toric varieties (on the B-side) is proved by Fang–Liu–Treumann–Zaslow in [FLTZ12]. We follow the more modern fomulation from [GPS18a, Cor. 6.16]. For any toric variety $\mathbf{T}$ of dimension n defined over $\mathbb{C}$, there exists a singular Legendrian $\partial_\infty \Lambda_{\mathbf{T}} \subset \partial_\infty T^* T^n$ such that

$$\operatorname{Coh}(\mathbf{T}) \cong \mathcal{W}(T^* T^n, \partial_\infty \Lambda_{\mathbf{T}}). \tag{4.4.19}$$

According to the work of Zhou [Zho20], which generalizes the arguments in [GS17], the singular Legendrian $\partial_\infty \Lambda_{\mathbf{T}}$ arises as the skeleton of a Weinstein hypersurface $F_{\mathbf{T}}$. Note that the ambient Weinstein manifold $(\mathbb{C}^*)^n = T^* T^n$ comes with a natural polarization therefore Theorem 4.1.1 applies. We conclude that Conjecture 4.1.1 is true for all 3-dimensional toric varieties, using the fact that $\mathcal{W}(T^* T^n, \partial_\infty \Lambda_{\mathbf{T}})$ is equivalent to $\mathcal{W}(T^* T^n, F_{\mathbf{T}})$.

Using similar techniques, one can also obtain upper bounds for the Rouquier dimension of derived categories of coherent sheaves on singular algebraic varieties.

Although Orlov's conjecture 4.1.1 is formulated for smooth quasi-projective varieties, its extension to possibly singular varieties/stacks holds for the following examples.

Example 4.4.6 (Toric boundary divisors). We recall the following result of homological mirror symmetry by Gammage–Shende [GS17, Thm. 1.0.1]. Denote by $\partial\mathbf{T}_\Sigma$ the toric boundary divisor of a smooth toric stack $\mathbf{T}$ over $\mathbb{C}$. Then there exists a Laurent polynomial $W_{\mathbf{T}} : (\mathbb{C}^*)^n \to \mathbb{C}$ and a natural Weinstein structure on the generic fiber $F_{W_{\mathbf{T}}}$ of $W_{\mathbf{T}}$ such that

$$\mathrm{Coh}(\partial\mathbf{T}_\Sigma) \cong \mathcal{W}(F_{W_{\mathbf{T}}}).$$

Note that $F_{W_{\mathbf{T}}}$ is a hypersurface in the affine variety $(\mathbb{C}^*)^n$ so it admits a polarization. When $\dim_{\mathbb{C}}(\partial\mathbf{T}_\Sigma) \leq 3$, Theorem 4.1.1 tells us

$$\mathrm{Rdim}\, D^b\,\mathrm{Coh}(\partial\mathbf{T}_\Sigma) = \dim_{\mathbb{C}}(\partial\mathbf{T}_\Sigma). \tag{4.4.20}$$

Example 4.4.7 (The mirror of pair of pants). Recall that the $(n-1)$–dimensional pair of pants $\mathcal{P}_{n-1}$ is defined to be the zero set of the polynomial

$$f(x_1, \ldots, x_n) = x_1 + \cdots + x_n + 1$$

defined over $(\mathbb{C}^*)^n$. This description of $\mathcal{P}_{n-1}$ endows it a Weinstein structure and a polarization. The mirror symmetry result [LP20, Cor. 1.1.2] tells us

$$D^b\,\mathrm{Coh}(\{x_1x_2\cdots x_n = 0\}) \cong \mathcal{W}(\mathcal{P}_{n-1}).$$

Then the arguments in the proof of Theorem 4.1.1 for $n \leq 4$ show that

$$\mathrm{Rdim}\, D^b\,\mathrm{Coh}(\{x_1x_2\cdots x_n = 0\}) = \dim_{\mathbb{C}}\{x_1x_2\cdots x_n = 0\}. \tag{4.4.21}$$

Example 4.4.8 (Mirrors of Milnor fibers of weighted homogeneous polynomials). A

weighted homogeneous polynomial $\mathbf{w} \in \mathbb{C}[x_1, \ldots, x_n]$ with an isolated singularity at the origin is called *invertible* if there exists a matrix $(a_{ij})_{i,j=1}^n$ with integer-valued entries and non-zero determinant such that

$$\mathbf{w} = \sum_{i=1}^{n} \prod_{j=1}^{n} x_j^{a_{ij}}.$$

The corresponding weight system $(d_1, \ldots, d_n, h)$ is determined uniquely by requiring

$$w(t^{d_1} x_1, \ldots, t^{d_n} x_n) = t^h w(x_1, \ldots, x_n)$$

for all $t \in \mathbb{G}_m$ and $\gcd(d_1, \ldots, d_n, h) = 1$. The *transpose* of $\mathbf{w}$ is defined to be

$$\check{\mathbf{w}} = \sum_{i=1}^{n} \prod_{j=1}^{n} x_j^{a_{ji}}$$

and let $(\check{d}_1, \ldots, \check{d}_n, \check{h})$ be the weight system of $\check{\mathbf{w}}$. Denote by $\check{V}_{\check{\mathbf{w}}}$ (the Liouville completion of) the Milnor fiber of $\check{\mathbf{w}}$, which is a Weinstein manifold coming with a global Lagrangian plane field. The symmetry group of $\mathbf{w}$ is defined by

$$\Gamma_{\mathbf{w}} := \{(t_1, \ldots, t_n) \in (\mathbb{G}_m)^n | t_1^{a_{11}} \cdots t_n^{a_{1n}} = \cdots = t_1^{a_{n1}} \cdots t_n^{a_{nn}}\}.$$

It is conjectured in [LU18, Conj. 1.4] and proved in [LU20, Thm. 1.2] and [Gam20, Thm. 1.6] that

$$\mathrm{mf}(\mathbb{A}^{n+1}, \mathbf{w} + x_0 \cdots x_n, \Gamma_{\mathbf{w}}) \cong \mathcal{W}(\check{V}_{\check{\mathbf{w}}}),$$

where $\mathrm{mf}(\mathbb{A}^{n+1}, \mathbf{w} + x_0 \cdots x_n, \Gamma_{\mathbf{w}})$ is the (dg) category of equivariant matrix factorizations. If $\check{h} - \sum_{i=1}^n \check{d}_i$ is positive, a theorem of Orlov [Orl09a, Thm. 3.11] identifies $\mathrm{mf}(\mathbb{A}^{n+1}, \mathbf{w} + x_0 \cdots x_n, \Gamma_{\mathbf{w}})$ with $\mathrm{Coh}(\mathcal{Z}_{\mathbf{w}})$. Here $\mathcal{Z}_{\mathbf{w}}$ is the stack

$$[(\operatorname{Spec} \mathbb{C}[x_0, \ldots, x_n]/\mathbf{w} + x_0 x_1 \cdots x_n)/\Gamma_{\mathbf{w}}].$$

For $n \leq 4$, Theorem 4.1.1 applies so we conclude that

$$\operatorname{Rdim} D^b \operatorname{Coh}(\mathcal{Z}_{\mathbf{w}}) = \dim_{\mathbb{C}} \mathcal{Z}_{\mathbf{w}}. \tag{4.4.22}$$

4.4.5 A digression on non-commutative motives

A functor from the category of small A_∞ categories over a field k to some additive category $\mathcal{D}$ is called an *additive invariant* if:

(i) it sends Morita equivalences to isomorphisms;

(ii) it sends semiorthogonal decompositions to direct sums.

. Well-known examples of additive invariants include Hochschild homology, algebraic K-theory, cyclic homology, topological Hochschild homology, etc. We recommend [Tab15, Sec. 2.2] for details and many more examples. As explained in [Tab15, Sec. 2.3], all additive invariants factor through a *universal* additive invariant, which takes values in the additive category $\operatorname{Hmo}_0(k)$ of non-commutative Chow motives; the objects in this category are small dg/A_∞ categories, morphisms are obtained by taking the Grothendieck group of a certain category of bimodules.

As a byproduct of the argument of the previous section, we obtain the following (to us) rather surprising result:

Corollary 4.4.3. Let (X, A) be a polarizable Weinstein pair such that $\mathcal{W}(X, A)$ is proper (it is a fortiori smooth). Let $U(-)$ be the universal additive invariant. Then $U(\mathcal{W}(X, A))$ is a summand of $\oplus_{\sigma\in\Sigma} U(\operatorname{Perf} k[T_\sigma])$, where Σ is a finite set and T_σ is a tree quiver.

Proof. The proof of 4.4.13 and 4.4.2 furnishes a localization map $p : \mathcal{G} := \operatorname{Groth}_{\sigma\in\Sigma} \operatorname{Perf}(k[T_\sigma]) \to \mathcal{W}(X, A)$, where Σ is a finite poset. Since p is a localization, the pullback p^* on modules is fully faithful. Moreover, since $\operatorname{Perf} \mathcal{W}(X, A)$ and $\mathcal{G}$ are smooth and proper, p^*

takes $\operatorname{Perf}\mathcal{W}(X,A)$ into $\operatorname{Perf}\mathcal{G} \subset \operatorname{Mod}\mathcal{G}$. This implies that there is a semi-orthogonal decomposition $\operatorname{Perf}\mathcal{G} = \operatorname{Perf}\langle \operatorname{Perf}\mathcal{W}(X,A), \mathcal{B}\rangle$, where $\mathcal{B}$ is right orthogonal to p^* and $\mathcal{B}$ is smooth and proper; cf. [Tab16, Sec. 4]. Hence $U(\operatorname{Perf}\mathcal{W}(X,A))) \oplus U(\mathcal{B}) = U(\mathcal{G}) = \oplus_{\sigma\in\Sigma} U(\operatorname{Perf} k[T_\sigma])$. □

Here are some consequences.

Corollary 4.4.4. Under the assumptions of 4.4.3, $K_*(\mathcal{W}(X,A))$ and $HH_*(\mathcal{W}(X,A))$ are concentrated in degree 0 and are in fact isomophic. (If one works instead with $\mathbb{Z}$-coefficients, one similarly learns that $HH_*(\mathcal{W}(X,A))$ and $K_*(\mathcal{W}(X,A))$ are free abelian groups.)

Proof. It is known that the trace $K_*(\operatorname{Perf} k[T_\sigma]) \to HH_*(\operatorname{Perf} k[T_\sigma])$ is an isomorphism and that both sides are concentrated in degree zero (see e.g. [ST16, Sec. 5.3]). The claim follows from 4.4.3, since $K_*(-), HH_*(-)$ factor through U. □

Example 4.4.9. As a sanity check, suppose that $k = \mathbb{C}$ and let $(T^*T^n, \partial_\infty \Lambda_{\mathbf{T}})$ be mirror to a smooth toric variety X as in 4.4.5. By the HKR theorem, $HH_*(D^b\operatorname{Coh}(X)) = \oplus_p H^{p,p-*}(X)$. Hence we learn from 4.4.4 that $h^{p,q} = 0$ whenever $p \neq q$. This is indeed true: it is a consequence of the general fact that the cohomology of a smooth toric variety coincides with its Chow ring.

If we assume that $\mathcal{W}(X,A)$ is defined over a perfect field, then $U(\operatorname{Perf} kT_\sigma) = U(k) \oplus \cdots \oplus U(k)$ ($|T_\sigma|$ times); see [Tab15, Sec. 2.4.8]. Hence:

Corollary 4.4.5. Assume k is a perfect field and let $V(-)$ be *any* additive invariant (such as Hochschild homology, K-theory, cyclic homology, topological Hochschild homology, topological cyclic homology, etc.), Under the assumptions of 4.4.3, $V(\mathcal{W}(X,A))$ is a direct summand of $V(k)^m$, for some $m \geq 1$. □

Here is another consequence of 4.4.3: Tabuada, building on proposals of Kontsevich, introduced non-commutative analogs of various classical conjectures in arithmetic geometry; see [Tab18]. By combing 4.4.3 and [Tab18, Cor. 12.2], we learn:

Corollary 4.4.6. Assume k is a perfect field. If (X, A) is a polarizable Weinstein pair such that $\mathcal{W}(X, A)$ is proper, then $\mathcal{W}(X, A)$ satisfies the non-commutative Weil conjecture and the strong non-commutative Tate conjecture. □

4.5 Central actions on the wrapped Fukaya category

The main goal of this section is to prove 4.5.3, which gives a lower bound on the Rouquier dimension of certain wrapped Fukaya categories. We do this by considering a central action of $HH^\bullet(\mathcal{W}(X))$ on $\mathcal{W}(X)$ and adapting methods of Bergh–Iyengar–Krause–Opperman [BIKO10] to deduce lower bounds on the Rouquier dimension. The crucial point is that this central action can be understood geometrically in terms of the closed-open string map.

4.5.1 Central actions on triangulated categories

Let $\mathcal{T} = (\mathcal{T}, \Sigma, \Delta)$ be a triangulated category. Let $Z^d(\mathcal{T})$ be the set of natural transformations $\eta : \mathrm{id} \to \Sigma^d$ such that $\eta \circ \Sigma = (-1)^d \Sigma \circ \eta$. Observe that $Z^\bullet(\mathcal{T}) = \oplus_{d\in\mathbb{Z}} Z^d(\mathcal{T})$ inherits the structure of a graded-commutative ring by composition. We call $Z^\bullet(\mathcal{T})$ the *graded center* of $\mathcal{T}$.

We will be mainly interested in the case where $\mathcal{T} = H^0(\mathcal{C})$ for $\mathcal{C}$ being a pre-triangulated A_∞ category. In that case, the graded center of $H^0(\mathcal{C})$ is precisely the set of graded natural transformations of the identity in $H^\bullet(\mathcal{C})$; cf. [Low08, Sec. 4.2]. Concretely, the elements of the graded center of $H^0(\mathcal{C})$ are products

$$\zeta = (\zeta_K)_{K\in\mathcal{C}} \in \prod_{K\in\mathcal{C}} H^\bullet(K, K) \tag{4.5.1}$$

subject to the naturality condition that the following diagram commutes

$$\begin{array}{ccc} H^\bullet(K,L) & \xrightarrow{\zeta_L \otimes 1} & H^\bullet(L,L) \otimes H^\bullet(K,L) \\ \downarrow{\scriptstyle 1\otimes\zeta_K} & & \downarrow \\ H^\bullet(K,L) \otimes H^\bullet(K,K) & \longrightarrow & H^\bullet(K,L). \end{array} \tag{4.5.2}$$

Given a graded-commutative ring R, a *central action* of R on the triangulated category $\mathcal{T}$ is a morphism of graded-commutative rings $R \to Z^\bullet(\mathcal{T})$.

Any triangulated category $\mathcal{T}$ which comes from a pre-triangulated A_∞ category $\mathcal{C}$ (i.e. which is of the form $\mathcal{T} = H^0(\mathcal{C})$) admits a canonical central action

$$\chi_{\mathcal{C}} : HH^\bullet(\mathcal{C},\mathcal{C}) \to Z^\bullet(\mathcal{T}), \tag{4.5.3}$$

where $HH^\bullet(\mathcal{C},\mathcal{C})$ is the Hoschschild cohomology of $\mathcal{C}$. In general, this map is neither injective nor surjective. (Recall that the Hochschild cohomology of an A_∞ category is naturally a graded-commutative ring. We refer to [RS17, Sec. 2.13] for a very detailed overview of Hochschild cohomology of A_∞ categories.)

The morphism (4.5.3) which realizes this canonical action is called the *characteristic morphism*. To describe the characteristic morphism concretely, let us recall the bar model for the Hochschild cochain complex.

$$CC^k(\mathcal{C},\mathcal{C}) = \prod_{\substack{l\geq 0 \\ K_0,\ldots,K_l\in\mathcal{C}}} \mathrm{Hom}_{k-l}(\hom_{\mathcal{C}}(X_{l-1},X_l) \otimes \cdots \otimes \hom_{\mathcal{C}}(X_0,X_1), \hom_{\mathcal{C}}(X_0,X_l)) \tag{4.5.4}$$

where $\mathrm{Hom}_s(-,-)$ is the space of degree s morphisms of graded vector spaces.

The differential applied to an element ϕ of degree d is

$$\delta(\varphi)(a_l,\dots,a_1) = \sum_{r,s}(-1)^{\clubsuit_1}\mu_{\mathcal{C}}^{l-r+s}(a_l,\dots,a_{r+1},\varphi^{r-s+1}(a_r,\dots,a_s),a_{s-1},\dots,a_1) \tag{4.5.5}$$

$$-(-1)^d\sum_{r,s}(-1)^{\clubsuit_2}\varphi^{l-r+s}(a_l,\dots,a_{r+1},\mu_{\mathcal{C}}^{r-s+1}(a_r,\dots,a_s),a_{s-1},\dots,a_1), \tag{4.5.6}$$

where $\clubsuit_1 := d(|a_{s-1}|+\cdots+|a_1|-(s-1))$ and $\clubsuit_2 = (|a_{s-1}|+\cdots+|a_1|-(s-1))$.

Given any object $K\in\mathcal{C}$, there is a natural projection map

$$\pi_K : CC^k(\mathcal{C},\mathcal{C}) \to \mathrm{Hom}_k(k,\mathrm{hom}_{\mathcal{C}}(K,K)) \equiv \mathrm{hom}_{\mathcal{C}}(K,K) \tag{4.5.7}$$

$$\varphi \mapsto \varphi_K^0 \equiv \varphi_K^0(1). \tag{4.5.8}$$

It can be verified (cf. [RS17, Sec. 2.14]) that the projection map is a morphism of rings. The characteristic morphism is then precisely

$$\chi_{\mathcal{C}}(\varphi) = (\varphi_K^0(1))_{K\in\mathcal{C}}. \tag{4.5.9}$$

The fact that this is a central action follows from the definition of the Hochschild differential. Indeed, given a closed element $x\in\mathrm{hom}(K,L)$, we have $0=\delta(\varphi)(x) = \mu_{\mathcal{C}}^2(\varphi_L^0(1),x)-\mu_{\mathcal{C}}^2(x,\varphi_K^0(1))$. This matches (4.5.2).

Remark 4.5.1. The Hochschild cohomology of an A_∞ category is canonically equivalent to the space of natural transformations of the identity functor, i.e. the "derived center" of the category. The characteristic morphism thus realizes a map from the derived center of an A_∞ category to the center of its derived category. For more on this perspective, see [NS17] and the references therein.

4.5.2 Central actions and the Rouquier dimension

The goal of this section is to prove the following theorem. We refer to 4.1.3 for motivation and a discussion of the proof.

Theorem 4.5.1. Let $\mathcal{T}$ be a $\mathbb{Z}/m$-graded triangulated category ($1 \leq m \leq \infty$) and let R be a finite type k-algebra acting centrally on $\mathcal{T}$. Suppose that there exists a (split-)generator $G \in \mathcal{T}$ such that $\operatorname{Hom}^\bullet_{\mathcal{T}}(G,G)$ is a Noetherian R-module. Then $\operatorname{Rdim} \mathcal{T} \geq \dim_R \operatorname{Hom}^\bullet_{\mathcal{T}}(G,G)$.

Recollections from commutative algebra

Given a module M over a commutative ring R, the *annihilator* is the ideal $\operatorname{Ann}_R(M) := \{r \in R \mid r \cdot m = 0 \text{ for all } m \in M\} \subseteq R$.

The *Krull dimension* of a commutative ring R is defined as the length of the longest strictly increasing chain of prime ideals $\mathfrak{p}_0 \subsetneq \cdots \subsetneq \mathfrak{p}_l$ and is denoted by $\dim(R)$. Given a finitely generated module M over R, one defines the Krull dimension of M to be $\dim_R(M) := \dim(R/\operatorname{Ann}_R(M))$. Finally, the Krull dimension of a (Zariski) closed subset of $\operatorname{Spec} R$ is the length of the longest strictly increasing chain of closed subsets. When the ring R is clear from the context, we usually write the Krull dimension as $\dim(M) := \dim_R(M)$.

The *support* of a module M over R is the set $\operatorname{Supp}_R(M) := \{\mathfrak{p} \in \operatorname{Spec} R \mid M_\mathfrak{p} \neq 0\}$. The support enjoys the following standard properties:

Fact 4.5.2 (Properties of the support). Let R be a commutative ring.

1. If M is a finitely generated R module, then $\operatorname{Supp}_R(M) = V(\operatorname{Ann}_R(M)) := \{\mathfrak{p} \in \operatorname{Spec} R \mid \operatorname{Ann}_R(M) \subseteq \mathfrak{p}\}$. In particular, $\dim_R(M) = \dim(\operatorname{Supp}_R(M))$.

2. If $0 \to M_1 \to M \to M_2 \to 0$ is an exact sequence of R-modules, then $\operatorname{Supp}_R(M) = \operatorname{Supp}_R(M_1) \cup \operatorname{Supp}_R(M_2)$.

3. If M, N are finitely generated R-modules, then $\operatorname{Supp}_R(M \otimes_R N) = \operatorname{Supp}_R(M) \cap \operatorname{Supp}_R(N)$.

4. If M is a finitely generated R-module and $I \subseteq R$ is an ideal, then $\operatorname{Supp}_{R/I}(M/I) = V(I) \cap \operatorname{Supp}_R(M)$.

We also need to review the notion of regular sequences.

Definition 4.5.3. Let R be a commutative ring and let M be an R-module. A sequence $x_1, \ldots, x_n$ is *M-regular* (or just *regular*) if for each $k \leq n$, we have that x_k is not a zero-divisor in $M/(x_1, \ldots, x_{k-1})M$, and that $(x_1, \ldots, x_k)M \neq M$. Given an ideal $I \subseteq R$, we let $\operatorname{depth}_I M$ denote the length of the longest M-regular sequence contained in I. (If R is local and $I \subseteq R$ is the maximal ideal, we usually just write $\operatorname{depth}_I M$ as $\operatorname{depth} M$).

Fact 4.5.4 (Regular sequences). Let R be a commutative ring and let M be a module.

1. Suppose R is Noetherian and M is finitely generated. If $\mathfrak{m} \in \operatorname{Spec} R$ is maximal, then
$$\operatorname{depth}_{\mathfrak{m}} M = \operatorname{depth}_{\mathfrak{m}R_{\mathfrak{m}}} M_{\mathfrak{m}}. \tag{4.5.10}$$

2. Let $x_1, \ldots, x_n$ be elements of R and let and $a_1, \ldots, a_n > 0$ be integers. Then $x_1, \ldots, x_n$ is an M-regular sequence if and only if $x_1^{a_1}, \ldots, x_n^{a_n}$ is an M-regular sequence.

3. Suppose that $x_1, \ldots, x_n$ is an M-regular sequence. Suppose that there exists $\mathfrak{p} \in \operatorname{Supp}_R(M)$ such that $x_i \in \mathfrak{p}$. Then $x_1, \ldots, x_n$, viewed as a sequence of elements of $R_{\mathfrak{p}}$, is an $M_{\mathfrak{p}}$ regular sequence.

Proof. (1) follows from [Aea17, Cor. 1.27]. (2) is [Sta21, Tag 0AUH]. For (3), note that since multiplication by x_i is injective on $M/(x_1, \ldots, x_{i-1})M$, the same is true after localizing (recall that localization is exact). Finally, it follows from Nakayama's lemma that $M_{\mathfrak{p}}/(x_1, \ldots, x_n)M_{\mathfrak{p}} = 0$ iff $M_{\mathfrak{p}} = 0$. □

Lemma 4.5.5. Let R be a finite type algebra over a field k. Suppose that R is equidimensional (meaning that all its minimal primes have the same dimension). Given a minimal prime $\mathfrak{p} \subset R$, if $f \in R \setminus \mathfrak{p}$, then $\dim(R) = \dim(R_f)$.

Proof. Observe that $\dim(R_f) = \dim(R_f/\mathfrak{p})$. We have a canonical isomorphism $R_f/\mathfrak{p} = (R/\mathfrak{p})_f$. Finally, since $R/\mathfrak{p}$ and $(R/\mathfrak{p})_f$ are integral domains and finite type algebras over k, it follows from [Vak17, Thm. 11.2.1] that $\dim(R) = \dim(R/\mathfrak{p}) = \mathrm{tr.deg}\, K(R/\mathfrak{p}) = \mathrm{tr.deg}\, K((R/\mathfrak{p})_f) = \dim((R/\mathfrak{p})_f) = \dim(R_f)$. □

Note that the hypothesis that R is of finite type is essential. (For example, $R = k[x]_{(x)}$ has Krull dimension 1. However, x is not contained in the unique minimal prime (0) and $R_x = k(x)$ has dimension 0.)

Corollary 4.5.1. Let R be a finite type algebra over a field k and let M be a finitely generated R-module. Let $\mathfrak{p} \in \mathrm{Ass}_R(M)$ be an associated prime with the property that $\dim(M/\mathfrak{p}) = \dim(M)$. If $f \in R \setminus \mathfrak{p}$ and f is contained in every other associated prime of M, then:

- $\mathrm{Ass}_{R_f}(M_f)$ has exactly one associated prime (hence $\mathrm{Supp}_{R_f}(M_f)$ has exactly one minimal prime).
- $\dim(M) = \dim(M_f)$.

Proof. It follows from [Sta21, Tag 05BZ] that $\mathfrak{p}$ is the unique associated prime of M_f. Hence [Sta21, Tag 02CE] implies that it is also the unique minimal prime. This is the first claim.

We have that $\dim(M_f) = \dim(M_f/\mathfrak{p}) = \dim((M/\mathfrak{p})_f)$. It is enough to show that $\dim(M/\mathfrak{p}) = \dim((M/\mathfrak{p})_f)$. It follows from 4.5.2 that $\mathrm{Supp}(M/\mathfrak{p}) = V(\mathfrak{p}) = \mathrm{Spec}(R/\mathfrak{p})$, while the support of $(M/\mathfrak{p})_f$ is $V(\mathfrak{p}) \cap D(f) = \mathrm{Spec}((R/\mathfrak{p})_f)$ (where $D(f) := \{\mathfrak{q} \in \mathrm{Spec}\, R \mid f \notin \mathfrak{q}\}$). Hence the second claim follows from 4.5.5 applied to $R/\mathfrak{p}$. □

Proposition 4.5.6. Let R be a finite type algebra over a field k and let M be a finite R-module of Krull dimension $c \geq 0$. Then there exists an M-regular sequence of length c.

Proof. Choose a minimal associated prime $\mathfrak{p} \in \mathrm{Ass}(M)$ such that $M/\mathfrak{p}$ has the same dimension as M (see [Sta21, Tag 02CE]). Since M is finite, $\mathrm{Ass}_R(M)$ is finite [Sta21, Tag 00LC]. Choose an element $f_1 \in R$ such that $f_1 \notin \mathfrak{p}$ but $f_1 \in \mathfrak{q}$ for all $\mathfrak{q} \in \mathrm{Ass}(M) \setminus \mathfrak{p}$. According to 4.5.1, $\dim(M_{f_1}) = \dim(M) = c$.

Now choose a tower of primes $\mathfrak{p}R_{f_1} = \mathfrak{p}_0 \subseteq ... \subseteq \mathfrak{p}_c$ in R_{f_1} of maximal length. Pick $r_1 \in \mathfrak{p}_1 \setminus \mathfrak{p}_0$. Then r_1 is not a zero-divisor of M_{f_1} (indeed, by [Sta21, Tag 00LD] $\mathfrak{p}_0$ contains all zero divisors). According to 4.5.2(4), the dimension of the quotient module is at least $c - 1$.

We now replace M with M_{f_1}/r_1 and R with R_{f_1}/r_1, and repeat the same argument. After c steps, we end up with $(\dots(M_{f_1}/r_1)_{f_2}\dots/r_{c-1})_{f_c}/r_c$. This is the same as $M_{f_1\dots f_c}/(r_1, \dots, r_c)$. (Note that this module has dimension zero, so it is not the zero module!) So $r_1, \dots, r_c$ is a regular sequence for $M_{f_1\dots f_c}$.

Since $M_{f_1\dots f_c}$ is not zero, it has a maximal ideal $\mathfrak{m}$ in its support. We now localize at $\mathfrak{m}$. Observe first that $(M_{f_1\dots f_c})_{\mathfrak{m}} = M_{\mathfrak{m}}$ (indeed, it follows from 4.5.2(3) that $\mathfrak{m} \in D(f_1 \dots f_m)$, so $f_1 \dots f_m$ is a unit in $R_{\mathfrak{m}}$). Moreover, it follows from 4.5.4(3) that $r_1, \dots, r_c$ is still a regular sequence for the $R_{\mathfrak{m}}$ module $M_{\mathfrak{m}}$. The conclusion now follows from 4.5.4(1). □

Remark 4.5.2. In fact, more is true: with the notation of the proof of 4.5.6, one can apply [Kap70, Sec. 3.1, Ex. 23] to show that there exists an M-regular sequence generating $\mathfrak{m}$ which is *stable under permutations*. We thank Srikanth Iyengar for this reference.

Support varieties

We begin with some standard facts about Noetherian modules. First, recall that a module over a Noetherian ring is Noetherian iff it is finitely generated. Also, Noetherian modules are well behaved under exact sequences, meaning that if $0 \to M_1 \to M \to M_2 \to 0$ is an exact sequence of R-modules, then M is Noetherian iff M_1, M_2 are Noetherian.

Lemma 4.5.7. Assume that R is Noetherian and acts centrally on $\mathcal{T}$. Given any object $X \in \mathcal{T}$, the subcategory $\{Y \in \mathcal{T} \mid \operatorname{Hom}^\bullet_{\mathcal{T}}(X, Y) \text{ is Noetherian over } R\} \subseteq \mathcal{T}$ is a thick triangulated subcategory. Similarly for $\{Y \in \mathcal{T} \mid \operatorname{Hom}^\bullet_{\mathcal{T}}(Y, X) \text{ is Noetherian over } R\} \subseteq \mathcal{T}$.

Proof. We only treat the first case as the second one is similar. Consider an exact triangle $Y_1 \to Q \to Y_2 \to$ where $\operatorname{Hom}^\bullet_{\mathcal{T}}(X, Y_i)$ is Noetherian over R. Then we have an exact sequence $0 \to M \to \operatorname{Hom}^\bullet_{\mathcal{T}}(X, Q) \to N \to 0$ where $M = \operatorname{Hom}^\bullet_{\mathcal{T}}(X, Y_1)/\ker(\operatorname{Hom}^\bullet_{\mathcal{T}}(X, Y_1) \to \operatorname{Hom}^\bullet_{\mathcal{T}}(X, Q))$ and $N = \operatorname{im}(\operatorname{Hom}^\bullet_{\mathcal{T}}(X, Q) \to \operatorname{Hom}^\bullet_{\mathcal{T}}(X, Y_2))$. Observe that M, N are Noetherian over R. Hence so is the middle term.

Suppose that $\operatorname{Hom}^\bullet_{\mathcal{T}}(X, Y_1) \hookrightarrow \operatorname{Hom}^\bullet_{\mathcal{T}}(X, Y_2)$, e.g. if Y_1 is a direct summand of Y_2. If $\operatorname{Hom}^\bullet_{\mathcal{T}}(X, Y_2)$ is Noetherian over R, it follows that $\operatorname{Hom}^\bullet_{\mathcal{T}}(X, Y_1)$ is too. □

We also need the following lemma, which can be deduced from 4.5.2.

Lemma 4.5.8. Let R be a commutative Noetherian ring which acts centrally on $\mathcal{T}$. Fix an object $X \in \mathcal{T}$. Then for any subset $\mathcal{U} \subseteq \operatorname{Spec} R$, the subcategories $\{Y \in \mathcal{T} \mid \operatorname{Supp}_R \operatorname{Hom}^\bullet_{\mathcal{T}}(X, Y) \in \mathcal{U}\}$ and $\{Y \in \mathcal{T} \mid \operatorname{Supp}_R \operatorname{Hom}^\bullet_{\mathcal{T}}(Y, X) \in \mathcal{U}\}$ are thick triangulated subcategories. □

As a corollary of 4.5.7 and 4.5.8, we get:

Corollary 4.5.2. Let R be a commutative Noetherian ring which acts centrally on $\mathcal{T}$. Suppose that $G \in \mathcal{T}$ is a generator and let $X, Y \in \mathcal{T}$ be arbitrary objects.

- If $\mathrm{Hom}^\bullet_{\mathcal{T}}(G,G)$ is a Noetherian R-module, then $\mathrm{Hom}^\bullet_{\mathcal{T}}(X,Y)$ is also Noetherian.
- We have $\mathrm{Supp}_R(\mathrm{Hom}^\bullet_{\mathcal{T}}(X,Y)) \subseteq \mathrm{Supp}_R(\mathrm{Hom}^\bullet_{\mathcal{T}}(G,G))$. In particular, all generators have the same support.

□

Proof of 4.5.1

A standard (and possibly the only) tool for obtaining lower bounds on the Rouquier dimension of a triangulated category is the so-called "ghost lemma".

Lemma 4.5.9 (Ghost lemma; see Lem. 4.2 in [EL18]). Let $\mathcal{T}$ be a triangulated category and let F, G be objects. Suppose there exist morphisms

$$K_c \xrightarrow{\theta_c} K_{c-1} \xrightarrow{\theta_{c-1}} \dots \xrightarrow{\theta_1} K_0 \tag{4.5.11}$$

such that the following conditions hold:

1. $\mathrm{Hom}^\bullet_{\mathcal{T}}(G, \theta_i) = 0$ for each $i = 1, \dots, c$;
2. $\mathrm{Hom}^\bullet_{\mathcal{T}}(F, \theta_1 \dots \theta_c) \neq 0$.

Then one has $F \notin \langle G \rangle_c$.

Following [BIKO10], we recall the notion of a Koszul object. Suppose that R is a commutative ring acting centrally on a triangulated category $\mathcal{T}$. Given $X \in \mathcal{T}$ and $r \in R$, we say that $X/\!\!/r$ is a *Koszul object* for r if it fits into the exact triangle

$$X \xrightarrow{r} X \to X/\!\!/r \to \Sigma X. \tag{4.5.12}$$

Note that $X/\!\!/r$ is well-defined up to isomorphism.

By splicing together exact triangles, we get

$$X \xrightarrow{r} X \to X/\!\!/r \to \Sigma X \xrightarrow{-r} \Sigma X. \tag{4.5.13}$$

More generally, given a sequence of elements $r_1, \dots, r_n \in R$, we inductively define $X/\!\!/\boldsymbol{r} := X/\!\!/(r_1, \dots, r_n)$ as the unique object fitting into the triangle

$$X/\!\!/(r_1, \dots, r_{n-1}) \xrightarrow{r_n} X/\!\!/(r_1, \dots, r_{n-1}) \to X/\!\!/(r_1, \dots, r_n) \to \Sigma(X/\!\!/(r_1, \dots, r_{n-1})). \tag{4.5.14}$$

Lemma 4.5.10. Suppose that $a_1, \dots, a_n \in R$ is an M-regular sequence, where $M = \mathrm{Hom}^\bullet_{\mathcal{T}}(X, X)$. Then there are natural isomorphisms of R-modules:

1. $\mathrm{Hom}^\bullet_{\mathcal{T}}(X/\!\!/(a_1, \dots, a_n), X) \simeq M[-n]/(r_1, \dots, r_n)$;
2. $\mathrm{Hom}^\bullet_{\mathcal{T}}(X, X/\!\!/(a_1, \dots, a_n)) \simeq M/(r_1, \dots, r_n)$

Proof. We only treat (i) because (ii) is similar. Suppose first $n = 1$. Apply the functor $\mathrm{Hom}^\bullet_{\mathcal{T}}(-, X)$ to (4.5.13) (with $r = a_1$) gives an exact sequence of R-modules

$$M[-1] \xrightarrow{-a_1} M[-1] \to \mathrm{Hom}^\bullet_{\mathcal{T}}(X/\!\!/a_1, X) \to M \xrightarrow{a_1} M, \tag{4.5.15}$$

from which the desired claim follows. In general, suppose that the claim holds for all regular sequences of length at most $n - 1$. Consider the spliced triangle

$$\begin{aligned} X/\!\!/(a_1, \dots, a_{n-1}) \xrightarrow{a_n} X/\!\!/(a_1, \dots, a_{n-1}) \to X/\!\!/(a_1, \dots, a_n) \\ \to \Sigma(X/\!\!/(a_1, \dots, a_{n-1})) \to \Sigma(X/\!\!/(a_1, \dots, a_{n-1})). \end{aligned} \tag{4.5.16}$$

and apply the same argument. □

We will also need the following lemma.

Lemma 4.5.11 (cf. Lem. 5.11(1) in [BIK08] and Lem. 3.1 in [BIKO10]). Let $n \geq 1$ be an integer and set $s = 2^n$. For any sequence of elements $r_1, \dots, r_n \in R$, we have $r_i^s \cdot \mathrm{Hom}^\bullet_{\mathcal{T}}(X/\!\!/\boldsymbol{r}, -) = 0 = \mathrm{Hom}^\bullet_{\mathcal{T}}(-, X/\!\!/\boldsymbol{r}) \cdot r_i^s$. □

Proof of 4.5.1. By assumption, $N := \mathrm{Hom}^\bullet_{\mathcal{T}}(G, G)$ is a Noetherian R-module. Let $n = \dim_R(N)$. If G' is another generator, then it follows from 4.5.2 that $\mathrm{Hom}^\bullet_{\mathcal{T}}(G', G')$

also a Noetherian R-module of Krull dimension n. Hence, we can restrict our attention to G. Our goal is to construct an object which is not contained in $\langle G \rangle_n$.

By 4.5.6, there exists an N-regular sequence $r_1, \dots, r_n$. Set $s = 2^n$ and note that $r_1^s, \dots, r_n^s$ is still a regular sequence (4.5.4). For $l \leq n$, consider the exact triangles (obtained by rotating (4.5.14) and setting $X = G$)

$$\Sigma^{-1}(G/\!\!/(r_1^s, \dots, r_l^s)) \xrightarrow{-w[-1]} G/\!\!/(r_1^s, \dots, r_{l-1}^s) \xrightarrow{r_l^s} G/\!\!/(r_1^s, \dots, r_{l-1}^s) \to G/\!\!/(r_1^s, \dots, r_l^s), \tag{4.5.17}$$

where w is the connecting map. Suspending $l-1$ more times, we get

$$\begin{aligned} \Sigma^{-l}(G/\!\!/(r_1^s, \dots, r_l^s)) &\xrightarrow{\theta_l} \Sigma^{-(l-1)}(G/\!\!/(r_1^s, \dots, r_{l-1}^s)) \\ \xrightarrow{\pm r_l^s} \Sigma^{-(l-1)}(G/\!\!/(r_1^s, \dots, r_{l-1}^s)) &\to \Sigma^{-(l-1)}(G/\!\!/(r_1^s, \dots, r_l^s)). \end{aligned} \tag{4.5.18}$$

(note that each application of Σ^{-1} has the effect of changing the sign of $\cdot r_l^s$; here θ_l is defined as the morphism obtained by applying $\Sigma^{-(l-1)}(-)$ to $-w[-1]$).

Applying the functor $\mathrm{Hom}^\bullet_{\mathcal{T}}(G, -)$, it follows from 4.5.10(2) that we have a natural isomorphism $\mathrm{Hom}^\bullet_{\mathcal{T}}(G, \Sigma^{-(l-1)}(G/\!\!/r_1^s, \dots, r_{l-1}^s)) \simeq N/(r_1^s, \dots, r_{l-1}^s)[-(l-1)]$. By regularity, $\pm r_l^s$ acts injectively, and it follows that $\mathrm{Hom}^\bullet_{\mathcal{T}}(G, \theta_l) = 0$ by (4.5.18).

Next, it follows from 4.5.11 that $\pm r_l^s$ acts by zero on $\mathrm{Hom}(G/\!\!/\boldsymbol{r}, \Sigma^{-(l-1)}(G/\!\!/(r_1^s, \dots, r_{l-1}^s))$, which implies that $\mathrm{Hom}^\bullet_{\mathcal{T}}(G/\!\!/\boldsymbol{r}, \theta_l)$ is surjective by (4.5.18). But 4.5.10(1) implies that $\mathrm{Hom}^\bullet_{\mathcal{T}}(G/\!\!/\boldsymbol{r}, G) \simeq N[-n]/(r_1, \dots, r_n)$. By regularity, $N[-n]/(r_1, \dots, r_n) \neq 0$. Since this module is the target of the surjective map $\mathrm{Hom}^\bullet_{\mathcal{T}}(G, \theta_1 \dots \theta_n)$, this map is nonzero. We conclude by 4.5.9 (with $G/\!\!/\boldsymbol{r}$ in place of F) that $G/\!\!/\boldsymbol{r} \notin \langle G \rangle_n$. □

4.5.3 Central actions on Fukaya categories

We now specialize the discussion of the previous sections to wrapped Fukaya categories. For simplicity, we assume that the wrapped Fukaya category is defined with field coefficients and $\mathbb{Z}/2$-gradings.

Given a Liouville manifold (X, λ), we wish to consider the canonical central action of $HH^\bullet(\mathcal{W}(X), \mathcal{W}(X))$ on the triangulated category $H^0(\operatorname{Perf} \mathcal{W}(X))$ furnished by the characteristic morphism. The following lemma states that this characteristic morphism admits a purely symplectic topological interpretation, in terms of the familiar closed-open string map which is constructed for instance in [Gan12, Sec. 5.4].

Lemma 4.5.12. For any object $K \in \mathcal{W}(X)$, the following diagram of unital rings commutes

$$\begin{array}{ccccc} SH^\bullet(X) & \xrightarrow{\mathcal{CO}} & HH^\bullet(\mathcal{W}(X), \mathcal{W}(X)) & \xleftarrow{\simeq} & HH^\bullet(\operatorname{Perf} \mathcal{W}(X), \operatorname{Perf} \mathcal{W}(X)) \\ & \searrow^{\mathcal{CO}(-)^0_K(1)} & \downarrow & \swarrow & \\ & & HW^\bullet(K, K) & & \end{array} \tag{4.5.19}$$

Proof. The fact that $\mathcal{CO}$ is a ring morphism is proved in [Gan12, Prop. 5.3], and the left triangle is then tautologically commutative. The top right arrow is induced by the canonical (Yoneda) embedding $\mathcal{W}(X) \to \operatorname{Perf} \mathcal{W}(X)$, and the left triangle is also tautologically commutative; cf. [RS17, Lem. 2.11]. The fact that the top right arrow is an isomorphism of graded commutative rings is a manifestation of the familiar Morita invariance of Hochschild cohomology (see e.g. [Sei, Rmk. 9.5]). □

Corollary 4.5.3. Let $R \subseteq SH^0(X)$ be a subring which is of finite type as an algebra over k. Suppose that there exists a (split-)generator $K \in \mathcal{W}(X)$ such that $\operatorname{Hom}^\bullet_{\mathcal{T}}(K, K)$ is a Noetherian R-module, where R acts via the closed-open map. Then

$$\operatorname{Rdim} \mathcal{W}(X) \geq \dim_R \operatorname{Hom}^\bullet_{\mathcal{T}}(K, K). \tag{4.5.20}$$

Proof. Combine 4.5.1 and 4.5.12. □

We will see in the next section how to perform concrete computations via string topology. For now, let us record the following rather trivial example.

Example 4.5.3. If $X = T^*\mathbb{T}^n$ and F is a cotangent fiber, then we have $\mathcal{CO} : SH^\bullet(X) = \Lambda^\bullet[v_1, \ldots, v_n] \otimes k[x_1^\pm, \ldots, x_n^\pm] \to k[x_1^\pm, \ldots, x_n^\pm] = HW^\bullet(F, F)$, where $|v_i| = 1, |x_i| = 0$ and v are in degree 1 and $\mathcal{CO}(1 \otimes x_i) = x_i$. Hence (taking $R = k[x_1^\pm, \ldots, x_n^\pm]$, we find that $\mathcal{W}(T^*\mathbb{T}^n) \geq n$, where $\mathcal{W}(-)$ is defined with k-coefficients).

For a related (but significantly more complicated) example, see 4.1.4.

4.5.4 Computations via string topology

We now explain how to obtain concrete lower bounds for the Rouquier dimension of certain cotangent bundles via string topology. We will always assume in this section that the wrapped Fukaya category is defined with $\mathbb{Q}$ coefficients and $\mathbb{Z}/2$-gradings (so in particular, all Liouville manifolds carry the appropriate orientations/grading data). Similarly, all singular homology groups considered in this subsection are assumed to have coefficients in $\mathbb{Q}$ and all graded objects (such as homology/cohomology grousp) are assumed to be $\mathbb{Z}/2$-graded (although the gradings all come from a $\mathbb{Z}$-grading, so the reader is welcome to work with $\mathbb{Z}$-gradings if they prefer).

Given a closed, oriented manifold M of dimension n, the (shifted) homology of its free loop space $H_{\bullet+n}(\mathcal{L}M)$ carries the structure of a graded-commutative algebra (in fact a BV algebra). The product is called the *Chas–Sullivan product* and was introduced by Chas and Sullivan in [CS99]. At an intuitive level, this product is constructed as follows: choose cycles $\sigma_p \in C_p(\mathcal{L}M)$ and $\sigma_q \in C_q(\mathcal{L}M)$, and consider the intersection $\mathrm{ev}(\sigma_p) \cap \mathrm{ev}(\sigma_q)$, where $\mathrm{ev} : \mathcal{L}M \to M$ is the evaluation map sending a loop $\gamma : S^1 = \mathbb{R}/\mathbb{Z} \to M$ to $\gamma(1) \in M$. Along this intersection, we can concatenate loops coming from σ_p with loops from σ_q. This defines a $p+q-n$ chain in $\mathcal{L}M$, which is precisely the desired product of σ_p and σ_q.

Next, it will be useful to recall the construction of the *umkehr* (or "wrong-way") map in differential topology. Let $\mathcal{M}$ be a smooth manifold and let $\mathcal{N} \subseteq \mathcal{M}$ be a closed, co-oriented submanifold of codimension d. Let $\tau_{\mathcal{M}/\mathcal{N}} \to \mathcal{N}$ be the normal

bundle of $\mathcal{N} \subseteq \mathcal{M}$ and let $\mathcal{U} \subseteq \mathcal{M}$ be a tubular neighborhood. We have a sequence of maps

$$H_\bullet(\mathcal{M}) \to H_\bullet(\mathcal{M}, \mathcal{M} - \mathcal{N}) \to H_\bullet(\mathcal{U}, \mathcal{U} - \mathcal{N}) \to H_\bullet(\tau_{\mathcal{M}/\mathcal{N}}, \tau_{\mathcal{M}/\mathcal{N}} - 0_{\mathcal{N}}) \to H_{\bullet - d}(\mathcal{N}), \tag{4.5.21}$$

where we are considering singular homology. (The first map is the natural projection, the second one is obtained by excision and the fourth one is the Thom isomorphism.) The composition (4.5.21) is called the *umkehr* map associated to the embedding $\mathcal{N} \subseteq \mathcal{M}$. Note that it is not necessary for $\mathcal{M}, \mathcal{N}$ to be finite dimensional in this construction. It works as long as $\mathcal{N} \subseteq \mathcal{M}$ has finite codimension and admits a tubular neighborhood.

We will in fact mostly be interested in the umkehr map associated to the embedding $\Omega_p M \hookrightarrow \mathcal{L}M$, where M is a closed manifold of dimension $n < \infty$ with a basepoint $p \in M$.

The associated umkehr map

$$\mathcal{I} : H_{\bullet + n}(\mathcal{L}M) \to H_\bullet(\Omega_p M) \tag{4.5.22}$$

is called the *intersection morphism*. It was also first introduced (under a slightly different guise) by Chas and Sullivan [CS99] and plays an important role in string topology. The construction of the intersection morphism via the umkehr map is found in [FTVP07, Sec. 9].

The connection to symplectic topology arises via the work of Abbondandolo–Schwarz [AS10a].

Theorem 4.5.13 (Abbondandolo–Schwarz, [AS10a]). Let M be a closed manifold of dimension n. Let $F_p \subseteq T^*M$ be a cotangent fiber. Then the following diagram

commutes

$$\begin{array}{ccc} H_{-\bullet+n}(\mathcal{L}M) & \xrightarrow{\simeq} & SH^\bullet(T^*M) \\ \downarrow{\scriptstyle \mathcal{I}} & & \downarrow{\scriptstyle c\mathcal{O}(-)^0_{F_p}(1)} \\ H_{-\bullet}(\Omega_p M) & \xrightarrow{\simeq} & HW^\bullet(F_p, F_p) \end{array} \tag{4.5.23}$$

Proof. All of this is contained in [AS10a]. However, since the notation and terminology in *loc. cit.* differs from ours, we give a brief overview for the reader.

Given a manifold M, the wrapped Floer cohomology of a cotangent fiber is denoted by $HF^\Omega(T^*M)$ in [AS10a] and called "Floer homology for Hamiltonian orbits with Dirichlet boundary conditions". The intersection morphism that we denote by $\mathcal{I}$ is denoted by $i_!$. The authors describe how to realize this map in Morse theory in [AS10a, Sec. 2.2], which also depends on material from the appendix. The closed-open map is denoted by $I_!$ and constructed in [AS10a, Sec. 3.5]. The top horizontal arrow in (4.5.23) corresponds to [AS10a, (3)], while the bottom horizontal arrow is [AS10a, Thm. B]. The fact that (4.5.23) commutes amounts to the assertion that the diagram [AS10a, (7)] has a Floer theoretic counterpart; as is often the case in Floer theory, the proof consists in showing that the two ways of traveling around the diagram correspond to the two ends of a cobordism of moduli spaces (see [AS10a, Thm. 4.11]). □

We now discuss some concrete computations. We begin by recording the following lemma.

Lemma 4.5.14. Let M and N be closed oriented manifolds of dimension m and n respectively. Then the following diagram of vector spaces commutes, where the vertical arrow is the usual Eilenberg–Zilber map:

$$\begin{array}{ccc} H_{\bullet+m}(\mathcal{L}M) \otimes H_{\bullet+n}(\mathcal{L}N) & \xrightarrow{\mathcal{I}\otimes\mathcal{I}} & H_\bullet(\Omega_* M) \otimes H_\bullet(\Omega_* N) \\ \downarrow{\scriptstyle \simeq} & & \downarrow{\scriptstyle \simeq} \\ H_{\bullet+m+n}(\mathcal{L}(M \times N)) & \xrightarrow{\mathcal{I}} & H_\bullet(\Omega_*(M \times N)) \end{array} \tag{4.5.24}$$

The vertical arrows are in fact isomorphisms of algebras with respect to the Chas–Sullivan and Pontryagin products.

Proof. There are canonical splittings of the based loop and free loop spaces which are compatible with the inclusions (of the former into the latter). It follows from the Künneth formula and the definition of the umkehr map that the umkehr map takes products of spaces to tensor products. Finally, it is apparent from the above description of the Chas–Sullivan product that it is compatible with the Künneth isomorphism, which means that the vertical left arrow is indeed an isomorphism of algebras. Similar considerations apply for the Pontryagin product. □

The following theorem is extremely useful, as it essentially reduces the computation of the intersection morphism $\mathcal{I}$ to rational homotopy theory.

Theorem 4.5.15 (Félix–Thomas–Vigué-Poirrier; see Thm. F. in [FTVP07]). Let M be a simply-connected closed oriented manifold of dimension n with basepoint p. Then there exists a commutative diagram of morphisms of graded-commutative algebras

$$\begin{array}{ccc} HH^\bullet(C^\bullet(M), C^\bullet(M)) & \xleftarrow[\simeq]{} & H_{\bullet+n}(\mathcal{L}M) \\ \downarrow{\scriptstyle HH^\bullet(C^\bullet(M),\epsilon)} & & \downarrow{\scriptstyle \mathcal{I}} \\ HH^\bullet(C^\bullet(M), \mathbb{Q}) & \xleftarrow[\simeq]{} & H_\bullet(\Omega_p M), \end{array} \tag{4.5.25}$$

where $\epsilon : C^\bullet(M) \to \mathbb{Q}$ is the augmentation induced by the inclusion $p \hookrightarrow M$.

Given a graded vector space $V = \{V_i\}_{i\in\mathbb{Z}}$, let $T(V) := \oplus_{i\geq 0} V^{\otimes i}$ be the *tensor algebra* on V. Let $I \subseteq T(V)$ be the ideal generated by elements of the form $x \otimes y - (-1)^{|x||y|} y \otimes x$ with $x, y \in V$. Then $\Lambda V := T(V)/I$ be the *free commutative graded algebra* on V. See [FHT01, Sec. 3(a)]. If V is generated by a single element v, we write $T(V), \Lambda(V)$ as $T(v), \Lambda v$ respectively.

Let us now record some useful facts drawn from [FTVP04, Sec. 6.1]. First of all, $HH^\bullet(C^\bullet(S^n), \mathbb{Q}) = T(v)$ with $|v| = n-1$. If $n > 1$ is odd, then $HH^\bullet(C^\bullet(S^n), C^\bullet(S^n)) \simeq$

$\Lambda u \otimes T(v)$, where $|u| = -n$ and $|v| = n-1$. We have $\mathcal{I} = \epsilon \otimes 1 : \Lambda u \otimes T(v) \to T(v)$, where we implicitly identify both columns of (4.5.25) under the horizontal isomorphisms.

If $n > 1$ is even, then

$$HH^\bullet(C^\bullet(S^n), C^\bullet(S^n)) \simeq \Lambda b \otimes k[a,c]/(2ac, a^2, ab) \tag{4.5.26}$$

with $|a| = -n, |b| = -1$ and $|c| = 2n-2$. We have $\mathcal{I}(a) = \mathcal{I}(b) = 0$ and $\mathcal{I}(c) = v^2$.

Example 4.5.4 (Products of odd dimensional spheres). Let $M = S^{n_1} \times \ldots \times S^{n_k}$ be a product of odd dimensional spheres with $n_i \geq 3$. Fix a basepoint $p \in M$. If $k = 1$, then it follows from the computations above that there is an inclusion $\mathbb{Q}[x] \hookrightarrow H_{\bullet+n_1}(\mathcal{L}S^{n_1})$ sending $x \mapsto 1 \otimes v$, such that the composition $\mathcal{Q}[x] \to H_{\bullet+n_1}(\mathcal{L}S^{n_1}) \to H_\bullet(\Omega S^{n_1})$ is an isomorphism.

For $k > 1$, similar considerations along with 4.5.14 furnish maps

$$\mathbb{Q}[x_1, \ldots, x_k] \hookrightarrow H_{\bullet+n_1}(\mathcal{L}M) \to H_\bullet(\Omega M) \tag{4.5.27}$$

where $|x_i| = n_i - 1$ and the composition is an isomorphism.

By combining 4.5.13 and 4.5.3, it follows that

$$\operatorname{Rdim} \mathcal{W}(T^*M) \geq k. \tag{4.5.28}$$

Example 4.5.5 (Compact Lie groups). Let G be a simply-connected compact Lie group. Then G has a Sullivan minimal model of the form $(\Lambda V, d = 0)$, where $V = \{V^i\}_{i \in \mathbb{N}}$ is a graded $\mathbb{Q}$-vector space supported in odd degrees and satisfying $V^1 = 0$. Thus G has the rational homotopy type of a product of odd dimensional spheres.

The dimension of V is precisely the rank of the group G [FOT08, Thm. 3.33]. It

follows from 4.5.4 that we have

$$\operatorname{Rdim} \mathcal{W}(T^*G) \geq \operatorname{rank} G. \tag{4.5.29}$$

For the reader's convenience, we list the simply-connected simple compact Lie groups in 4.1. It follows from the classification theorem that all other simply-connected compact Lie groups are obtained from these by taking products. For $i = 6, 7, 8$, we write $\tilde{E}_i$ to denote the universal cover of the exceptional Lie group E_i.

Simply-connected compact Lie group	Real dimension	Rank	$\operatorname{Rdim} \mathcal{W}(T^*G)$
$\mathrm{Sp}(n), n \geq 1$	$n(2n+1)$	n	$\geq n$
$\mathrm{SU}(n), n \geq 3$	$n^2 - 1$	$n-1$	$\geq n-1$
$\mathrm{Spin}(n), n \geq 7$	$n(n-1)/2$	$\lfloor n/2 \rfloor$	$\geq \lfloor n/2 \rfloor$
G_2	14	2	≥ 2
F_4	52	4	≥ 4
$\tilde{E}_6$	156	6	≥ 6
$\tilde{E}_7$	266	7	≥ 7
$\tilde{E}_8$	496	8	≥ 8

Figure 4.1: The simply-connected compact Lie groups

Example 4.5.6 (Products of arbitrary spheres). We again assume $M = S^{n_1} \times \ldots \times S^{n_k}$. Suppose first that $k = 1$ and $n_1 > 1$ is even. Then there is a natural inclusion $\mathbb{Q}[x] \hookrightarrow \Lambda b \otimes \mathbb{Q}[a, c]/(2ac, a^2, ab)$ with $x \mapsto c$. Composing with the intersection morphism gives the map $\mathbb{Q}[c] \to \mathbb{Q}[v]$ sending $c \mapsto v^2$. More generally, we can apply 4.5.14, to get a map

$$\mathbb{Q}[x_1, \ldots, x_k] \to H_{\bullet + n_1}(\mathcal{L}M) \to \mathbb{Q}[v_1, \ldots, v_k], \tag{4.5.30}$$

where

- $|x_i| = n_i - 1$ and $x_i \mapsto v_i$ if n_i is odd;
- $|x_i| = 2n_i - 2$ and $x_i \mapsto v_i^2$ if n_i is even.

Observe that $\mathbb{Q}[v_1, \ldots, v_k]$ is a Noetherian module over $\mathbb{Q}[x_1, \ldots, x_k]$ (indeed, it is finitely generated as a module over a Noetherian ring). By combining 4.5.13 and 4.5.3, it follows that

$$\operatorname{Rdim} \mathcal{W}(T^*M) \geq k. \tag{4.5.31}$$

Example 4.5.7 (Complex projective spaces). Let us also briefly mention the case of complex projective spaces. The cohomology of $\Omega\mathbb{CP}^n$ is computed in [See14, Thm. 1.3]. It contains a free subalgebra $\mathbb{Q}[y]$ where $|y| = 2n$. On the other hand, as explained in e.g. [See14, Sec. 3], we have $\Omega\mathbb{CP}^n = S^1 \times \Omega S^{2n+1}$ and hence $H_\bullet(\Omega\mathbb{CP}^n) = H_\bullet(S^1) \otimes H_\bullet(\Omega S^{2n+1})$. It is shown in [BG17, Ex. 5.7] that the characteristic morphism hits the class $1 \otimes v$ of degree $2n$ in $H_\bullet(\Omega\mathbb{CP}^n)$. For degree reasons, the composition of $\mathbb{Q}[y] \to H_{\bullet+2n}(\mathcal{L}\mathbb{CP}^n) \to H_\bullet(\Omega\mathbb{CP}^n)$ must send y to $1 \otimes v$.

By a similar argument to 4.5.4, we find that if $M = M_1 \times \ldots \times M_k$ where each M_i is a sphere of dimension at least 2 or a complex projective space of arbitrary dimension, then

$$\operatorname{Rdim} \mathcal{W}(M) \geq k. \tag{4.5.32}$$

4.6 Applications to symplectic geometry

We now describe how the Rouquier dimension of wrapped Fukaya categories is related to certain quantities and operations of geometric interest.

4.6.1 Embedding monotonicity

An important fact about the Rouquier dimension of Weinstein manifolds is that it is monotonic under Liouville embeddings which are Weinstein up to deformation. To explain this precisely, we begin with the following definition.

Definition 4.6.1. A *Liouville embedding* of Liouville domains is a codimension 0

smooth embedding

$$i : (X_0^{\text{in}}, \lambda_0^{\text{in}}) \hookrightarrow (X_0, \lambda_0) \tag{4.6.1}$$

where $i^*\lambda_0 = e^\rho \lambda_0^{\text{in}} + df$ for some $\rho \in \mathbb{R}$ and $f : X_0^{\text{in}} \to \mathbb{R}$.

We refer to [GS20, Sec. 2.1.3] for a detailed discussion of Liouville embeddings. As explained in [GS20, Lem. 2.2], any Liouville embedding as in 4.6.1 extends to a family of Liouville embeddings $i^t : (X_0^{\text{in}}, \lambda_0^{\text{in}}) \hookrightarrow (X_0, \lambda_0^t)$ such that $i = i^0$ and i^1 is a strict Liouville embedding, i.e. $(i^1)^*\lambda_0^1 = \lambda_0^{\text{in}}$.

Proposition 4.6.2. Let $i : (X_0^{\text{in}}, \lambda^{\text{in}}) \subseteq (X_0, \lambda)$ be a Liouville embedding and suppose that $(X_0^{\text{in}}, \lambda^{\text{in}})$ and $(X_0 - i(X_0^{\text{in}}), \lambda)$ are Weinstein up to deformation. Then

$$\operatorname{Rdim} \mathcal{W}(X^{\text{in}}) \leq \operatorname{Rdim} \mathcal{W}(X) \tag{4.6.2}$$

where X^{in}, X are the completions of X_0^{in}, X_0 respectively.

Proof. Since the wrapped Fukaya category is invariant under deformation, we may assume that the embedding is strict, i.e. $i^*\lambda = \lambda^{\text{in}}$. We may therefore appeal to [GPS18b, Prop. 8.15], which constructs a Viterbo restriction functor

$$\operatorname{Tw} \mathcal{W}(X) \to \operatorname{Tw} \mathcal{W}(X^{\text{in}}). \tag{4.6.3}$$

As explained in [GPS18b, Prop. 8.15], the hypothesis that the X_0^{in} and $X_0 - i(X_0^{\text{in}})$ are Weinstein implies that there is in fact a quasi-equivalence $\operatorname{Tw} \mathcal{W}(X)/\mathcal{D} \to \operatorname{Tw} \mathcal{W}(X^{\text{in}})$, where $\mathcal{D}$ is a set of objects. The claim now follows from 4.2.1 upon passing to $\operatorname{Perf}(-)$. □

In the statement of 4.6.2, note that the orientation/grading data on $\mathcal{W}(X^{\text{in}})$ is implicitly induced from the orientation/grading data on X.

It is natural to ask whether one can drop the hypothesis that $X_0 - i(X_0^{\text{in}})$ is Weinstein 4.6.2. We note that Sylvan has proved [Syl19b] that the Viterbo restriction

functor in this setting functor is a *homological epimorphism*, although it is not clear if this helps.

By combining various upper and lower bounds for $\operatorname{Rdim} \mathcal{W}(-)$ discussed in the introduction and below, it is easy to write down examples of putative embeddings $X_0^{\text{in}} \hookrightarrow X_0$ which are obstructed by $\operatorname{Rdim} \mathcal{W}(-)$. For instance:

Example 4.6.1. Let X be a Weinstein manifold of real dimension $2(n^2-1)$ admitting a Lefschetz fibration with Weinstein fibers having $< n-2$ critical points. Then X does not contain $T^*SU(n)$ as a Weinstein subdomain.

There is of course nothing special about 4.6.1, and there is a plethora of similar examples. It seems unlikely to the authors that examples such as these could be handled using invariants already in the literature. From this perspective, $\operatorname{Rdim} \mathcal{W}(-)$ is a useful addition to the symplectic topologist's toolkit of embedding obstructions. On the other hand, $\operatorname{Rdim} \mathcal{W}(-)$ has many blind spots. The most obvious defect of $\operatorname{Rdim} \mathcal{W}(-)$ as an embedding obstruction is that it outputs a very limited range of values: indeed, it seems plausible that $\operatorname{Rdim} \mathcal{W}(X^{2n}) \in [0, n]$ (see 4.4.2 for a somewhat weaker statement).

4.6.2 Lefschetz fibrations

We now discuss upper bounds for the Rouquier dimension of wrapped Fukaya categories in terms of the number of critical points of a Lefschetz fibration.

Definition 4.6.3 (see Sec. 2 in [MS10])**.** Given a Liouville manifold X, a *Lefschetz fibration* is a smooth map $f : X \to \mathbb{C}$ having Morse-type critical points and such that $\ker df$ is symplectic away from the critical points. This data is required to satisfy some additional technical conditions at infinity detailed in [MS10, Sec. 2]. A Lefschetz fibration is said to have *Weinstein fibers* if the general fiber is a Weinstein manifold up to deformation.

There are various essentially equivalent notions of a Lefschetz fibration in the literature. We have adopted the setup of [MS10, Sec. 2] for consistency with [GPS20, Sec. 1.1].

Definition 4.6.4. Given a Liouville manifold X (considered up to Liouville homotopy), we let

$$\mathrm{Lef}_w(X) \in \mathbb{N} \cup \{\infty\} \tag{4.6.4}$$

denote the minimal number of critical points of a Lefschetz fibration on X with Weinstein fibers. We set $\mathrm{Lef}_w(X) = \infty$ if no such fibration exists.

It was proved by Giroux–Pardon [GP17, Thm. 1.10] using quantitative transversality techniques that any Weinstein manifold admits a Lefschetz fibration with Weinstein fibers. Conversely, if a Liouville manifold admits a Lefschetz fibration with Weinstein fibers, then it is Weinstein (this is tautological if one uses [GP17, Def. 1.9]). Hence $\mathrm{Lef}_w(X) < \infty$ iff X admits a Weinstein structure.

Concerning lower bounds, there is a purely topological bound $\mathrm{rk}\, H_n(X;\mathbb{Z}) \leq \mathrm{Lef}_w(X)$, where X has real dimension $2n$. It is also immediate that $\mathrm{Lef}_w(X) = 0$ iff X is subcritical, meaning that it splits as a product of Liouville manifolds $(X, \lambda) = (F_0 \times \mathbb{C}, \lambda_0 \times \lambda_{\mathbb{C}})$.

The following proposition, which was explained to us by Y. Baris Kartal, states that the Rouquier dimension gives another lower bound.

Proposition 4.6.5. Let X be a Liouville manifold with nonzero wrapped Fukaya category. Then

$$\mathrm{Lef}_w(X) \geq \mathrm{Rdim}\, \mathcal{W}(X) + 1. \tag{4.6.5}$$

Proof. We may assume without loss of generality that $0 < \mathrm{Lef}_w(X) = k < \infty$. Fix a Lefschetz fibration $f : X \to \mathbb{C}$ with k critical points and choose a set of thimbles $\{L_1, \ldots, L_k\}$. These form an exceptional collection for the partially wrapped Fukaya category $\mathcal{W}(X, f^{-1}(\infty))$; see [GPS20, Ex. 1.4]. According to [GPS18b, Cor. 1.14], this

exceptional collection is *full*, meaning that it generates $\mathcal{W}(X, f^{-1}(\{\infty\}))$. It follows by 4.2.8 that $\operatorname{Rdim} \mathcal{W}(X, f^{-1}(\{\infty\})) \leq k-1$.

It follows from [GPS18b, Cor. 2.11] that $\mathcal{W}(X, f^{-1}(\{\infty\})) = \mathcal{W}(X, \mathfrak{c})$, where $\mathfrak{c}$ is the skeleton of $f^{-1}(\infty)$. The assumption that $f^{-1}(\{\infty\})$ is Weinstein implies that $\mathfrak{c}$ is mostly Legendrian [GPS18b, see Def. 1.6]. It then follows from 4.3.2, that the natural map $\mathcal{W}(X, f^{-1}(\{\infty\})) \to \mathcal{W}(X)$ is a quotient of A_∞ categories. 4.2.1 thus implies that $\operatorname{Rdim} \mathcal{W}(X) \leq k-1$. □

We do not expect (4.6.5) to be sharp except in very special cases.

Example 4.6.2. The affine quadric $\{z_1^2 + \cdots + z_n^2 + z_{n+1}^2 = 1\}$ with the Liouville structure $\sum y_i dx_i$ is isomorphic to T^*S^n as a Liouville manifold. The projection $(z_1, \ldots, z_{n+1}) \mapsto z_{n+1}$ is a Lefschetz fibration with two critical points. Hence $\operatorname{Rdim} \mathcal{W}(T^*S^n) \leq 1$. By 4.5.6, this is an equality if $\mathcal{W}(T^*S^n)$ is defined with $\mathbb{Z}/2$-gradings and $\mathbb{Q}$-coefficients.

Although the definition of $\operatorname{Lef}_w(X)$ only makes sense on Liouville manifolds, it is not a priori clear to what extent this invariant depends on the Liouville structure (recall Lazarev's result "$\operatorname{WMor}(X) = \operatorname{Mor}(X)$" discussed in the introduction, which implies that $\operatorname{WMor}(X)$ does not depend on the Liouville structure). The following corollary shows that the natural complex analog of "$\operatorname{WMor}(X) = \operatorname{Mor}(X)$" is false.

Corollary 4.6.1. There exists a Weinstein manifold $T^*S^3_{\text{exotic}}$ which is formally isotopic to T^*S^3, but such that $\operatorname{Lef}_w(T^*S^3) \neq \operatorname{Lef}_w(T^*S^3_{\text{exotic}})$.

Proof. [EGL20, Thm. 4.7] produces a Weinstein manifold $T^*S^3_{\text{exotic}}$ which contains a *flexible* Lagrangins $L \simeq \mathbb{T}^3$. This means by definition that L is *regular* (see [EGL20, Def. 3.1]), and hence the cobordism $T^*S^3_{\text{exotic}} - T^*_{\leq \epsilon} L$ is Weinstein. Now apply the monotonicity of Rdim under Weinstein embeddings. It follows that $\operatorname{Lef}_w(T^*S^3_{\text{exotic}}) \geq \operatorname{Rdim} \mathcal{W}(T^*S^3_{\text{exotic}}) + 1 \geq \operatorname{Rdim} \mathcal{W}(T^*\mathbb{T}^3) + 1 \geq 4$ by 4.5.3. On the other hand, we saw in 4.6.2 that $\operatorname{Lef}_w(T^*S^3) \leq 2$. □

4.6.3 Rigidity of skeleta

Finally, we prove the quantitative symplectic rigidity result for the skeleton of a Weinstein manifold which was advertised in the introduction.

Proposition 4.6.6. Let (X, λ) be a Liouville manifold with nonzero wrapped Fukaya category. Suppose that the skeleton $\mathfrak{c}_X$ is mostly Lagrangian and that every component of $\mathfrak{c}_X^{\text{crit}}$ admits a generalized cocore.

Let $\phi : X \to X$ be a Hamiltonian symplectomorphism. If $\mathfrak{c}_X \cap \phi(\mathfrak{c}_X) = \mathfrak{c}_X^{\text{crit}} \cap \phi(\mathfrak{c}_X^{\text{crit}})$ and these intersection points are transverse (a generic condition on ϕ), then

$$|\mathfrak{c}_X \cap \phi(\mathfrak{c}_X)| \geq \operatorname{Rdim} \mathcal{W}(X) + 1. \tag{4.6.6}$$

Proof. We consider the Liouville manifold $(\overline{X} \times X, -\lambda \oplus \phi_* \lambda)$ which has skeleton $\mathfrak{c}_X \times \phi(\mathfrak{c}_X)$. Let $\Delta \subseteq \overline{X} \times X$ be the diagonal. By assumption, Δ intersects $\mathfrak{c}_X \times \phi(\mathfrak{c}_X)$ at finitely many points. Suppose there are $r \geq 0$ such points.

Consider the diffeomorphism

$$(\overline{X} \times X, -\lambda \oplus \phi_* \lambda) \to (\overline{X} \times X, -\lambda \oplus \lambda) \tag{4.6.7}$$

$$(x_1, x_2) \mapsto (x_1, \phi^{-1}(x_2)), \tag{4.6.8}$$

which sends Δ to the Lagrangian $L = \{(x, \phi^{-1}(x)\}$.

We now apply 4.3.3. This produces a resolution of L of length $r \geq 0$ by generalized cocores of $\mathfrak{c}_{\overline{X} \times X} = \mathfrak{c}_X \times \mathfrak{c}_X$. Observe that one can choose these generalized cocores to be products of generalized cocores of $\mathfrak{c}_X$.

Since L is isotopic to Δ through a compactly-supported Lagrangian isotopy, they define isomorphic objects in $\mathcal{W}(\overline{X} \times X)$. Applying 4.7.1, we obtain a resolution of the diagonal bimodule of length r by Yoneda bimodules. If $r = 0$, then we find that the diagonal bimodule is the zero bimodule, which contradicts our assumption that

$\mathcal{W}(X)$ is not the zero category. Hence $r > 0$, and the conclusion now follows from 4.2.10. □

Remark 4.6.3. The assumption that $\mathfrak{c}_X^{\mathrm{crit}}$ and $\phi(\mathfrak{c}_X^{\mathrm{crit}})$ intersect transversally is used in the proof of 4.6.6 when appealing to 4.3.3. However, we do not know whether this assumption is necessary for the conclusion of the theorem to hold.

The proof of 4.6.6 can actually be adapted to derive lower bounds for $|\mathfrak{c}_X \cap \phi(\mathfrak{c}_X)|$ which do not involve the Rouquier dimension of $\mathcal{W}(X)$.

For example, suppose that $\mathcal{W}(X)$ admits a nonzero proper module. Then we claim $|\mathfrak{c}_X \cap \phi(\mathfrak{c}_X)| > 1$. Indeed, if $r = 0$, then $\mathcal{W}(X) = 0$. If $r = 1$, then the proof of 4.6.6 implies that the diagonal bimodule is isomorphic to $\mathcal{Y}_K^l \otimes_k \mathcal{Y}_L^r$ for some objects K, L of $\mathcal{W}(X)$. Convolving with $\mathcal{P}$, we find that $H^\bullet(P, P) = H^\bullet(K, K) \otimes H^\bullet(\mathcal{Y}_L^r, P)$. Hence $H^\bullet(K, K)$ is finite-dimensional. On the other hand, the convolution argument of 4.2.10 proves that K split-generates $\mathcal{W}(X)$. Thus $\mathcal{W}(X)$ is proper, which contradicts the main result of [Gan21].

Similarly, one can show that $|\mathfrak{c}_X \cap \phi(\mathfrak{c}_X)|$ is strictly greater than the dimension of the image of $K_0(\mathrm{Prop}\,\mathcal{C})$ in $K_0(\mathrm{Perf}\,\mathcal{C})$.

Example 4.6.4. If X contains an orientable closed exact Lagrangian, then this defines a proper module over $\mathcal{W}(X)$ (where we work with $\mathbb{Z}/2$ gradings and $\mathbb{Z}/2$ coefficients). Hence $|\mathfrak{c}_X \cap \phi(\mathfrak{c}_X)| \geq 2$.

4.7 Appendix

Let (X, λ) be a Liouville manifold and let $\mathcal{W}(X)$ be its wrapped Fukaya category. Let $\mathcal{W}(\overline{X} \times X)$ denote the wrapped Fukaya category of the product $(\overline{X} \times X, -\lambda \oplus \lambda)$.

If K and L are cylindrical Lagrangians in X, their product $K \times L$ is not a cylindrical Lagrangian in $\overline{X} \times X$ in general. However, following [GPS18b, Section 6.2], the product can be "cylindrized". This results in a cylindrical Lagrangian $K \tilde{\times} L$ which

behaves like $K \times L$ in the Floer-theoretic sense; see Lemma 4.7.6. If K, L are invariant under the Liouville flow (e.g. they could be cocores), then $K \times L = K \tilde{\times} L$.

Let $\mathcal{W}_{\mathrm{cyl}}(\overline{X} \times X) \subset \mathcal{W}(\overline{X} \times X)$ denote the smallest full, replete subcategory of $\mathcal{W}(\overline{X} \times X)$ containing the diagonal $\Delta \subset \overline{X} \times X$ and all cylindrized products (i.e. all objects of the form $K \tilde{\times} L$).

The purpose of this appendix is to prove the following theorem.

Theorem 4.7.1. There exists a fully faithful A_∞ functor

$$\mathcal{F} : \mathcal{W}_{\mathrm{cyl}}(\overline{X} \times X) \to (\mathcal{W}(X), \mathcal{W}(X)) - \mathrm{mod} \tag{4.7.1}$$

which satisfies the following properties:

- the image of the diagonal $\Delta \subseteq \overline{X} \times X$ is the diagonal bimodule;
- the image of any product Lagrangian $K \times L$ is the Yoneda bimodule $\mathcal{Y}^l_K \otimes_k \mathcal{Y}^r_L$.

Corollary 4.7.1. If X is Weinstein, then $\mathcal{F}$ extends to a fully-faithful functor $\mathrm{Tw}\, \mathcal{W}(\overline{X} \times X) \to (\mathcal{W}(X), \mathcal{W}(X)) - \mathrm{mod}\,$. □

The proof of 4.7.1 will occupy the remainder of this appendix.

Definition 4.7.2 (Wrapping category of (X, λ))**.** Define the *wrapping category* $((X, \lambda) \rightsquigarrow -)^+$ of (X, λ) as follows. The objects of $((X, \lambda) \rightsquigarrow -)^+$ are Hamiltonian isotopies of X generated by Hamiltonians $\{H_t\}_{0 \le t \le 1}$ such that H_t is linear at infinity. Given two objects $\{H_t\}_{0 \le t \le 1}$ and $\{H'_t\}_{0 \le t \le 1}$ in $((X, \lambda) \rightsquigarrow -)^+$, denote by ϕ_{H_t} and $\phi_{H'_t}$ respectively the Hamiltonian diffeomorphisms generated by them. Morphisms from $\{H_t\}_{0 \le t \le 1}$ to $\{H'_t\}_{0 \le t \le 1}$ are given by Hamiltonian diffeomorphisms satisfying $\phi_{H'_t} = \phi_{H_t} \sharp \phi_{\tilde{H}_t}$ and the generating Hamiltonians $\{\tilde{H}_t\}_{0 \le t \le 1}$ are *positive* and linear at infinity. Here $\phi_{H_t} \sharp \phi_{\tilde{H}_t}$ means the Hamiltonian diffeomorphism generated by $H_t \sharp \tilde{H}_t(x) = H_t(x) + \tilde{H}_t((\Phi^t_{H_t})^{-1}(x))$ and $\Phi^t_{H_t}$ is the time-t flow of the Hamiltonian vector field of H_t.

Recall that a category $\mathcal{C}$ is *filtered* if

1. $\mathcal{C}$ is non-empty;
2. for every pair of objects $x, y \in \mathcal{C}$, there exists an object $z \in \mathcal{C}$ and morphisms $x \to z, y \to z$;
3. for every pair of morphisms $f, g : x \to y$, there exists an object $z \in \mathcal{C}$ and a morphism $h : y \to z$ such that $h \circ f = h \circ g$.

A filtered category $\mathcal{C}$ is has *countable cofinality* if there exists a cofinal functor $\mathbb{Z}_{\geq 0} \to \mathcal{C}$.

Lemma 4.7.3. The category $((X, \lambda) \rightsquigarrow -)^+$ is filtered and has countable cofinality.

Proof. The proof that the category is filtered proceeds exactly as [GPS20, Lem. 3.27], except replacing the Lagrangian isotopies by Hamiltonian isotopies. That $((X, \lambda) \rightsquigarrow -)^+$ has countable cofinality can be proved by choosing a countable family of linear Hamiltonians with slope going to infinity. □

Lemma 4.7.4. Consider the diagonal $\Delta \subset \overline{X} \times X$ as a cylindrical Lagrangian. Let $(\Delta \rightsquigarrow -)^+$ be the positive wrapping category of Δ in the sense of [GPS20, Sec. 3.4]. Denote by π_1 and π_2 the projection from $\overline{X} \times X$ to the first and second factors respectively. Then the functor from $((X, \lambda) \rightsquigarrow -)^+$ to $(\Delta \rightsquigarrow -)^+$ by taking a generating Hamiltonian H_t to the Lagrangian isotopy of Δ which is generated by $\pi_1^* H_{1-t/2} + \pi_2^* H_{t/2}$ is cofinal.

Proof. This holds by inspecting the definitions. □

We introduce a category $\widetilde{\mathcal{O}}^{\mathrm{prod}}_{\overline{X} \times X}$ following [GPS18b, Sec. 6.6]. Fix a countable collection of Lagrangians $\tilde{I}^{\mathrm{prod}}_{\overline{X} \times X}$ in $\overline{X} \times X$ consisting of Δ and product Lagrangians $K \times L \subset \overline{X} \times X$ which represent every pair of isotopy classes. For the diagonal Δ, choose a countable cofinal family $\Delta = \Delta^{(0)} \rightsquigarrow \Delta^{(1)} \rightsquigarrow \cdots$ in $(\Delta \rightsquigarrow -)^+$ induced

from $((X,\lambda) \rightsquigarrow -)^+$ as in Lemma 4.7.4. Also choose cofinal wrappings $K = K^{(0)} \rightsquigarrow K^{(1)} \rightsquigarrow \cdots$ and $L = L^{(0)} \rightsquigarrow L^{(1)} \rightsquigarrow \cdots$. Then the category $\widetilde{\mathcal{O}}^{\mathrm{prod}}_{\overline{X}\times X}$ has objects given by $\mathbb{Z}_{\geq 0} \times \tilde{I}^{\mathrm{prod}}_{\overline{X}\times X}$, which are either of the form $\Delta^{(i)}$ or $K^{(i)} \times L^{(i)}$. The objects of $\widetilde{\mathcal{O}}^{\mathrm{prod}}_{\overline{X}\times X}$ admits a total order induced from $\mathbb{Z}_{\geq 0}$.

Note that by choosing the wrappings generically, we can assume that for any pair of Lagrangians in $\widetilde{\mathcal{O}}^{\mathrm{prod}}_{\overline{X}\times X}$, they intersect transversally and are disjoint at infinity with distance uniformly bounded from below. Using generic product almost complex structures, as they induce uniformly bounded geometry, one can define A_∞ operations on $\widetilde{\mathcal{O}}^{\mathrm{prod}}_{\overline{X}\times X}$ by counting suitable holomorphic discs based on the compactness arguments in [GPS20, Prop. 3.19]. Of course, the morphism spaces are given by the Floer cochain complexes. The wrappings define a collection of continuation elements [GPS20, Def. 3.25] C_{prod} inside the Floer cohomology groups $HF^0(K^{(i+1)} \times L^{(i+1)}, K^{(i)} \times L^{(i)})$ and $HF^0(\Delta^{(i+1)}, \Delta^{(i)})$. Then we define an A_∞ category by localizing $\widetilde{\mathcal{O}}^{\mathrm{prod}}_{\overline{X}\times X}$ along C_{prod}:

$$\widetilde{\mathcal{W}}^{\mathrm{prod}}(\overline{X} \times X) := \widetilde{\mathcal{O}}^{\mathrm{prod}}_{\overline{X}\times X}[C^{-1}_{\mathrm{prod}}]. \tag{4.7.2}$$

By construction and [GPS20, Lem. 3.37], the following isomorphisms hold in $\widetilde{\mathcal{W}}^{\mathrm{prod}}(\overline{X} \times X)$ (this is parallel to the computations in [Gan12, Sec. 8.1] under the "quadratic" setup):

$$\hom^\bullet(\Delta, \Delta) \cong SH^\bullet(X), \tag{4.7.3}$$

$$\hom^\bullet(K \times L, \Delta) \cong HW^\bullet(L, K), \hom^\bullet(\Delta, K \times L) \cong HW^\bullet(K, L), \tag{4.7.4}$$

$$\hom^\bullet(K \times L, K' \times L') \cong HW^\bullet(K, K') \otimes HW^\bullet(L, L'). \tag{4.7.5}$$

Proposition 4.7.5. There exists a fully faithful A_∞ functor $\mathcal{F}^{\mathrm{prod}} : \widetilde{\mathcal{W}}^{\mathrm{prod}}(\overline{X} \times X) \to (\mathcal{W}(X), \mathcal{W}(X)) - \mathrm{mod}$ which satisfies the following properties:

- the image of the diagonal $\Delta \subseteq \overline{X} \times X$ is the diagonal bimodule;
- the image of any product Lagrangian $K \times L$ is the Yoneda bimodule $\mathcal{Y}^l_K \otimes_k \mathcal{Y}^r_L$.

Proof. Recall that the wrapped Fukaya category $\mathcal{W}(X)$ is defined as a localization $\mathcal{O}_X[C^{-1}]$. Here $\mathcal{O}_X$ is an A_∞ category whose objects are given by $\mathbb{Z}_{\geq 0} \times I_X$, where I_X is a countable collection of cylindrical Lagrangians in X representing all isotopy classes and for each Lagrangian $L \in I_X$ we have a cofinal wrapping sequence $L = L^{(0)} \rightsquigarrow L^{(1)} \rightsquigarrow \cdots$. The set of morphisms C consists of continuation elements in $HF^0(L^{(i+1)}, L^{(i)})$. For details, see [GPS18b, Sec. 3.5].

Without loss of generality, we can assume that the product Lagrangians in the set $\tilde{I}^{\mathrm{prod}}_{\overline{X}\times X}$ in $\overline{X} \times X$ are all given by product of Lagrangians from I_X. Then we can define an A_∞ functor

$$\widetilde{\mathcal{O}}^{\mathrm{prod}}_{\overline{X}\times X} \to (\mathcal{O}_X, \mathcal{O}_X) - \mathrm{mod} \tag{4.7.6}$$

by considering quilted holomorphic maps as in [Gan12, Sec. 9]. More precisely, for any Lagrangian $\mathcal{L} \in \widetilde{\mathcal{O}}^{\mathrm{prod}}_{\overline{X}\times X}$, it defines an A_∞ bimodule $\mathcal{M}_{\mathcal{L}}$ over $\mathcal{O}_X$ whose morphism space associates to a pair of Lagrangians $K, L \in \mathcal{O}_X$ the Floer cochain complex

$$\hom^\bullet_{\widetilde{\mathcal{O}}^{\mathrm{prod}}_{\overline{X}\times X}}(K \times L, \mathcal{L}). \tag{4.7.7}$$

The A_∞ operations are defined by counting quilted holomorphic maps into $\overline{X} \times X$ with seam along the Lagrangian $\mathcal{L}$, see [GPS18b, Figure 19] without inserting marked point along the seam on the left picture and replacing X^-, Y both by X. The almost complex structures are of product type and are chosen generically. Once again, the compactness arguments from [GPS20, Prop. 3.19] and [GPS18b, Sec. 6.6] guarantee that the relevant moduli spaces which define $\mathcal{M}_{\mathcal{L}}$ and verify the A_∞ relations are compact. To complete the definition of the A_∞ functor $\widetilde{\mathcal{O}}^{\mathrm{prod}}_{\overline{X}\times X} \to (\mathcal{O}_X, \mathcal{O}_X) - \mathrm{mod}$, it suffices to consider quilted holomorphic maps with marked points inserted along the seam, which are exactly the left picture of [GPS18b, Figure 19] after replacing both X^- and Y by X. Note that we allow Lagrangians $\Delta^{(i)}$ as markings on the seam, but they do not impose further difficulties in the construction. By definition, this functor

takes the (strict) diagonal Lagrangian Δ to the diagonal bimodule $\Delta_{\mathcal{O}_X}$ and takes the product Lagrangian $K \times L$ to the Yoneda bimodule $\mathcal{Y}^l_K \otimes_k \mathcal{Y}^r_L$ over the category $\mathcal{O}_X$.

The localization functor $\mathcal{O}_X \to \mathcal{W}(X)$ induces an A_∞ functor

$$(\mathcal{O}_X, \mathcal{O}_X) - \mathrm{mod} \to (\mathcal{W}(X), \mathcal{W}(X)) - \mathrm{mod} \tag{4.7.8}$$

and its composition with the functor defined in the previous paragraph defines an A_∞ functor

$$\widetilde{\mathcal{O}}^{\mathrm{prod}}_{\overline{X}\times X} \to (\mathcal{W}(X), \mathcal{W}(X)) - \mathrm{mod}\,. \tag{4.7.9}$$

Using the definition of continuation elements in C_{prod}, we see that the above functor descends to an A_∞ functor

$$\mathcal{F}^{\mathrm{prod}} : \widetilde{\mathcal{W}}^{\mathrm{prod}}(\overline{X} \times X) \to (\mathcal{W}(X), \mathcal{W}(X)) - \mathrm{mod} \tag{4.7.10}$$

by the universal property of localization. The statements are proved because of (4.7.4) and (4.7.5). □

Lemma 4.7.6. There exists a fully faithful "cylindrization" A_∞ functor $\widetilde{\mathcal{W}}^{\mathrm{prod}}(\overline{X} \times X) \to \mathcal{W}(\overline{X} \times X)$ which takes the diagonal to the diagonal and takes products $K \times L$ to their cylindrization $K \tilde{\times} L$. In particular, this functor is a quasi-equivalence onto $\mathcal{W}_{\mathrm{cyl}}(\overline{X} \times X)$.

Proof. Let $\mathcal{O}_{\overline{X}\times X}$ be the pre-localized category whose quotient defines $\mathcal{W}(\overline{X} \times X)$. The A_∞ functor

$$\widetilde{\mathcal{O}}^{\mathrm{prod}}_{\overline{X}\times X} \to \mathcal{O}_{\overline{X}\times X} \tag{4.7.11}$$

is defined using the extension of the $(\mathcal{O}^{\mathrm{prod}}_{\overline{X}\times X}, \mathcal{O}_{\overline{X}\times X})$–bimodule $\mathcal{B}$ from [GPS18b, Sec. 6.6] by adding the diagonal Lagrangian. Note that by our choice of wrappings of the

diagonal, all $\Delta^{(i)}$ for $i \geq 0$ are cylindrical and obviously the diagonal is taken to the diagonal. The comparisons between Floer cochain complexes as in [GPS18b, Lem. 6.3] carry over without change so the induced functor from $\mathcal{B}$ is an equivalence onto its image. The desired functor is constructed by localization and it is a quasi-equivalence onto $\mathcal{W}_{\mathrm{cyl}}(\overline{X} \times X)$ by definition. □

Proof of Theorem 4.7.1. The functor $\mathcal{F}$ is constructed as the composition of the inverse of the embedding from Lemma 4.7.6 and the functor $\mathcal{F}^{\mathrm{prod}}$. It satisfies the desired properties by Proposition 4.7.5 and Lemma 4.7.6. □

Remark 4.7.1. Our proof is simply a modification of [GPS18b, Sec. 6] and [Gan12, Sec. 9] and we include the proof in this appendix just for completeness. The situation is more interesting with the presence of a mostly Legendrian stop $\mathfrak{f} \subset \partial_\infty X$. The induced stop in $\partial_\infty(\overline{X} \times X)$ is then given by $(\bar{\mathfrak{f}} \times \mathfrak{c}_X) \cup (\bar{\mathfrak{f}} \times \mathfrak{f} \times \mathbb{R}) \cup (\mathfrak{c}_{\overline{X}} \times \mathfrak{f})$ and the strict diagonal Δ necessarily intersects the stop. To define an admissible Lagrangian from Δ, one can either wrap it across the stop positively to get a Lagrangian Δ^+ or negatively to a Lagrangian Δ^-. The Lagrangians Δ^+ and Δ^- do not define the same object in $\mathcal{W}(\overline{X} \times X, (\bar{\mathfrak{f}} \times \mathfrak{c}_X) \cup (\bar{\mathfrak{f}} \times \mathfrak{f} \times \mathbb{R}) \cup (\mathfrak{c}_{\overline{X}} \times \mathfrak{f}))$ in general. The diagonal bimodule is actually induced by the Lagrangian Δ^+ rather than Δ^-. Our quantitative intersection result 4.6.6 has a counterpart in this setting. However, the relative Lagrangian skeleton $\mathfrak{c}_{(X,\mathfrak{f})}$ of the pair $(X, \mathfrak{f})$ is non-compact so the intersection number between $\mathfrak{c}_{(X,\mathfrak{f})}$ with its generic Hamiltonian push-off should be defined by adding numbers of "short" Reeb chords in $\partial_\infty X$. This could also be used to prove the folklore fact that the partially wrapped Fukaya category of a Weinstein pair is smooth. However, the details are beyond the scope of this chapter.

Remark 4.7.2. Continuing the discussion from the previous remark, the bimodule induced by Δ^- also has interesting geometric properties. When $\mathfrak{f}$ is the core of a Weinstein hypersurface in $\partial_\infty X$ coming from a generic fiber of a Weinstein Lefschetz

fibration $\pi : X \to \mathbb{C}$, this bimodule is related to the graph bimodule of the inverse of the global monodromy of π, which is further identified with the Serre bimodule of $\mathcal{W}(X; \mathfrak{f})$ by the work of Seidel [Sei18]. More generally, such property should hold for *swappable stops* in the sense of [Syl19b, Def. 1.1]. Abusing the notation, let $\Delta^{\pm}$ be the induced bimodules over $\mathcal{W}(X; \mathfrak{f})$. Suppose $\mathfrak{f}$ is the core of some Weinstein hypersurface $A \subset \partial_{\infty} X$. The mapping cone $\mathrm{Cone}(\Delta^- \to \Delta^+)$ induced by the positive isotopy is closely related to the geometry of the Orlov functor $\cup : \mathcal{W}(A) \to \mathcal{W}(X; A)$. One can speculate that this is the key of understanding the conjectural *relative left Calabi–Yau structures* [BD19, Def. 1.8] on the Orlov functor $\cup$.

Chapter 5

An integral Euler cycle in normally complex orbifolds and $\mathbb{Z}$-valued Gromov–Witten type invariants

5.1 Introduction

Let (X, ω) be a closed symplectic manifold and suppose $J : TX \to TX$ is an almost complex structure tamed by ω. Given a homology class $A \in H_2(X; \mathbb{Z})$ and an integer $g \geq 0$, the moduli space of stable J-holomorphic maps of class A and genus g with n marked points $\overline{\mathcal{M}}_{g,n}(X, J, A)$ carries a rational virtual fundamental class $[\overline{\mathcal{M}}_{g,n}(X, J, A)]^{\mathrm{vir}}$ as constructed by Fukaya–Ono [FO99] using the theory of Kuranishi structures, along with other constructions by Li–Tian [LT98], Ruan [Rua99], Hofer–Wysocki–Zehnder [HWZ17], Pardon [Par16], and many others. The Gromov–Witten invariants are defined by pairing $[\overline{\mathcal{M}}_{g,n}(X, J, A)]^{\mathrm{vir}}$ with cohomology classes obtained from pulling back classes in $H^*(\overline{\mathcal{M}}_{g,n}; \mathbb{Q})$ and $H^*(X; \mathbb{Q})$ under the forgetful map and the evaluation maps.

The Gromov–Witten invariants are in general $\mathbb{Q}$-valued because of the presence of

non-trivial automorphisms of stable maps, though genus-0 Gromov–Witten invariants are $\mathbb{Z}$-valued for semi-positive symplectic manifolds as proved by Ruan–Tian [RT95] and studied extensively in [MS04]. Technically speaking, rational numbers show up due to the use of abstract multi-valued perturbations of the $\overline{\partial}_J$-equations as from [FO99, LT98, HWZ17, MW17], or because Poincaré duality type constructions only hold over $\mathbb{Q}$ rather than $\mathbb{Z}$ for orbifolds as used in [Rua99, Par16] (see also [LT98, Remark 4] for a discussion about when the invariants are integers).

Motivated by seeking for $\mathbb{Z}$-valued Gromov–Witten type invariants, this chapter explains how to construct integral "Euler classes" for an orbifold complex vector bundle over an orbifold endowed with a "normal complex structure" following a proposal of Fukaya–Ono [FO97] and the subsequent development by Parker [Par13]. Combined with recent advances in algebraic topology [Par19, Par20] and regularization of moduli spaces of J-holomorphic maps [AMS21], we can study generalized bordism theories over orbispaces and define $\mathbb{Z}$-valued Gromov–Witten type invariants for general symplectic manifolds.

5.1.1 Statement of main results

Recall that an orbifold is called *effective* if the local uniformizer group from any orbifold chart acts faithfully.

Theorem 5.1.1. Suppose X is a compact effective orbifold without boundary endowed with an almost complex structure and let $\mathcal{E} \to X$ be a complex orbibundle. Denote by $X^{\text{free}} \subset X$ the suborbifold consisting of points with trivial isotropy group. Then there exist *normally complex* (Definition 5.3.18) smooth sections $s : X \to \mathcal{E}$ which is *strongly transverse* (Definition 5.6.1) such that $s^{-1}(0) \cap X^{\text{free}}$ defines a pseudocycle (Definition 5.5.6). Moreover, given a pair of such sections s_1 and s_2, the pseudocycles $s_1^{-1}(0) \cap X^{\text{free}}$ and $s_2^{-1}(0) \cap X^{\text{free}}$ are cobordant. Therefore, homology class defined by

the cobordism class of the pseudocycle

$$\chi^{\mathrm{FOP}}_{\mathrm{free}}(X,\mathcal{E}) := [s^{-1}(0)\cap X^{\mathrm{free}}] \in H_*(X;\mathbb{Z}), \tag{5.1.1}$$

which we call by the Fukaya–Ono–Parker–Euler class, is an invariant of $\mathcal{E}\to X$.

Remark 5.1.2. Zinger [Zin08] proved that any (oriented) pseudocycle in a smooth manifold defines an integral homology class and the space of pseudocycles modulo cobordism is naturally isomorphic to the integral homology. In Section 5.5, we show that a pseudocycle in any *Thom–Mather stratified space* (including any orbifold equipped with the *isotropy stratification*, see Section 5.3.2), defines an integral homology class.

Theorem 5.1.1 is a special case of Theorem 5.6.2 in Section 5.6.2, and we comment on how it generalizes Theorem 5.1.1. Firstly, the existence of (almost) complex structures can be weakened into the existence of *normal complex structures*. Given $x\in X$, the tangent space $T_xX = (T_xX)^{\mathrm{triv}}\oplus(T_xX)^{\mathrm{non}}$ is a representation of the isotropy group G_x and a normal complex structure on X is a family of complex structures on the direct sum of nontrivial G_x isotypic pieces $(T_xX)^{\mathrm{non}}$ satisfying certain compatibility conditions, see Section 5.3.5. In particular, the Fukaya–Ono–Parker–Euler class could have an odd degree. Secondly, the compactness assumption on X can be dropped as follows. Instead of considering a pair $(X,\mathcal{E})$, we can start with a triple $(X,\mathcal{E},s)$ where $s: X\to\mathcal{E}$ is a smooth section of an complex orbibundle $\mathcal{E}\to X$ over an arbitrary effective normally complex orbifold X, but the zero locus $s^{-1}(0)$ is required to be compact. Then we can construct strongly transverse normally complex perturbations of s "relative to ∞" to define an element in $\chi^{\mathrm{FOP}}_{\mathrm{free}}(X,\mathcal{E},s)\in H_*(X;\mathbb{Z})$ independent of the perturbation. Thirdly, a strongly transverse normally complex section actually defines an Euler-type class of $\mathcal{E}$ for any given stratum of the isotropy stratification, and the class $\chi^{\mathrm{FOP}}_{\mathrm{free}}(X,\mathcal{E})$ is the one associated to the stratum consisting

of points with trivial isotropy group. These classes could be viewed as refinements of the rational orbifold Euler class $\chi^{\mathrm{orb}}(X, \mathcal{E}) \in H_*(X; \mathbb{Q})$ which could be constructed by multi-valued perturbations.

The first main application of Theorem 5.6.2 concerns about the algebraic topology of bordisms of (derived) orbifolds. To state this result, we introduce the following definition. Note that we do not impose the effectiveness on D.

Definition 5.1.3. A *derived orbifold chart* (with boundary) is a triple (D, E, s) where D is a smooth orbifold (with boundary), $E \to D$ is a smooth orbibundle and $s : D \to E$ is a smooth section. (D, E, s) is said to be compact if $s^{-1}(0)$ is compact.

Following [Joy07, Par20], we can study geometric bordism type invariants constructed from derived orbifold charts. Let Y be an orbispace, i.e. a topological stack locally modeled on the quotient of a topological space by a finite group (see Definition 5.7.4), then one can consider compact derived orbifold charts of the form (D, E, s) together with a map $f : D \to Y$. We introduce an equivalence relation among such quadruples (D, E, s, f) generated by the following relations:

1. (*Restriction*) $(D, E, s, f) \sim (D', E', s', f')$ if $D' \subset D$ is an open subset with $s^{-1}(0) \subset D'$ and $E' = E|_{D'}$, $s' = s|_{D'}$, and $f' = f|_{D'}$.

2. (*Stabilization*) $(D, E, s, f) \sim (D', E', s', f')$ if D' is the total space of a vector bundle $\pi_F : F \to D$, $E' = \pi_F^* E \oplus \pi_F^* F$, $s' = \pi_F^* s \oplus \tau_F$ where $\tau_F : F \to \pi_F^* F$ is the tautological section, and $f' = f \circ \pi_F$.

3. (*Cobordism*) $(D, E, s, f) \sim (D', E', s', f')$ if there is a bordism between them (see Section 5.7.3).

In fact, we restrict our attention to derived orbifold charts (D, E, s) endowed with a *stable complex structure*, which is defined to be a stable complex structure on the

virtual orbibundle $TD - E$. Then stable complex derived orbifold charts mapped to Y modulo the above equivalence relation defines the *stable complex derived orbifold bordism* of Y, written as $\overline{\Omega}_*^{\mathbb{C},\mathrm{der}}(Y)$. This is a generalization of the stable complex bordism group over manifolds. A more detailed and precise exposition is presented in Section 5.7. In particular, any map $f : M \to Y$ from a closed almost complex manifold M represents an element of $\overline{\Omega}_*^{\mathbb{C},\mathrm{der}}(Y)$.

Theorem 5.1.4. (= Theorem 5.7.22) Denote by Γ the set of equivalence classes of triples $[G, V, W]$ where G is a finite group and V, W are finite-dimensional complex G-representations which do not contain any trivial summand, and where the equivalence is induced from stabilization:

$$(G, V, W) \sim (G, V \oplus R, W \oplus R).$$

For any $[\gamma] \in \Gamma$, there is a well-defined map

$$\mathcal{FOP}_{[\gamma]} : \overline{\Omega}_*^{\mathbb{C},\mathrm{der}}(Y) \to H_{*-\dim V+\dim W}(Y; \mathbb{Z}) \tag{5.1.2}$$

which restricts to the pushforward of the usual fundamental class for the trivial $[\gamma]$ and elements represented by maps from compact almost complex manifolds.

We call the set Γ from the above theorem the set of *stabilized isotropy types*. In fact, together with a relative version of the map $\mathcal{FOP}_{[\gamma]}$, we can obtain a natural transformation between two (generalized) homology theories for (orbi)spaces, see Section 5.7.5. Theorem 5.1.4 is already nontrivial even when G is the trivial group, as the map $\mathcal{FOP}_{[\gamma]}$ can be used to endow any compact stably complex orbifold which is not necessarily effective an integral "fundamental class". It is well-known to experts that the set $\overline{\Omega}_*^{\mathbb{C},\mathrm{der}}(Y)$ could be very large and complicated. However, we expect that the integral shadow (5.1.2) is more amenable to manipulate but still contains considerable amount of information.

Our second application of Theorem 5.6.2 is the construction of $\mathbb{Z}$-valued Gromov–Witten type invariants in symplectic topology as proposed by Fukaya–Ono [FO97]. Besides the transversality issue resolved in proving Theorem 5.6.2, an accompanying fundamental difficulty in the construction is the lack of smoothness of the moduli space of pseudo-holomorphic curves. One possible approach for a systematic solution to the smoothness problem is to either apply the polyfold theory of Hofer–Wysocki–Zehnder [HWZ07] or the delicate exponential decay estimates of Fukaya–Oh–Ohta–Ono [FOOO16]. Instead, in this chapter we appeal to a recent work of Abouzaid–McLean–Smith [AMS21] which ingeniously constructed a global Kuranishi chart with a smoothing on the genus zero Gromov–Witten type moduli space using a geometric perturbation. Using the language of derived orbifold charts, one main result from [AMS21] can be formulated as follows.

Proposition 5.1.5 ([AMS21, Proposition 5.35, 6.31]). Let (X, ω) be a closed symplectic manifold and J be a compatible almost complex structure. Let $A \in H_2(X; \mathbb{Z})$ be a homology class and denote by $\overline{\mathcal{M}}_{0,n}(X, J, A)$ the moduli space of genus-0 stable J-holomorphic maps to X with $n \geq 0$ marked points. Then after choosing certain auxiliary data, there exists a smooth stable complex derived orbifold chart (D, E, s) along with a map

$$\widetilde{\mathrm{st}} \times \widetilde{\mathrm{ev}} : D \to \overline{\mathcal{M}}_{0,n} \times X^n$$

such that the zero locus $s^{-1}(0)$ is isomorphic to $\overline{\mathcal{M}}_{0,n}(X, J, A)$ and the restriction of $\widetilde{\mathrm{st}} \times \widetilde{\mathrm{ev}}$ along $s^{-1}(0)$ coincides with the product of the stabilization map and the evaluation map. Moreover, for different choices of J and auxiliary data, the corresponding derived orbifold charts together with $\widetilde{\mathrm{st}} \times \widetilde{\mathrm{ev}}$ define the same element in $\overline{\Omega}_*^{\mathbb{C},\mathrm{der}}(\overline{\mathcal{M}}_{0,n} \times X^n)$.

Remark 5.1.6. The paper [AMS21] uses the notion of *global Kuranishi charts* instead of derived orbifold charts. A topological/smooth global Kuranishi chart is a quadruple

(G, D, E, s) where G is a compact Lie group, D is a topological/smooth manifold endowed with an almost free continuous/smooth G-action, E is a G-equivariant vector bundle over D, and s is a G-equivariant continuous/smooth section of $E \to D$. The equivalence relations introduced in the definition of $\overline{\Omega}_*^{\mathbb{C},\mathrm{der}}$ have counterparts in this context, but there is an additional equivalence relation between global Kuranishi charts, namely identifying (G, D, E, s) with $(G' \times G, P, q^*E, q^*s)$ where $q : P \to D$ is a G-equivariant principal G'-bundle. This change-of-group operation should be interpreted as choosing different global quotient presentation of the orbifold $[D/G]$, so the arguments in [AMS21] indeed prove the statements in the form of Proposition 5.1.5.

Theorem 5.1.7. Suppose (X, ω) is a closed symplectic manifold and $A \in H_2(X; \mathbb{Z})$. Fix a non-negative integer n. Given any $[\gamma] \in \Gamma$ as in Theorem 5.1.4, there is a well-defined integral homology class

$$[\overline{\mathcal{M}}_{0,n}(X, J, A)]^{\mathrm{vir}}_{[\gamma]} \in H_*(\overline{\mathcal{M}}_{0,n} \times X^n; \mathbb{Z}) \tag{5.1.3}$$

defined by virtually "counting" J-holomorphic maps in $\overline{\mathcal{M}}_{0,n}(X, J, A)$ in the stratum indexed by $[\gamma]$.

Proof. This is a combination of Theorem 5.1.4 and Proposition 5.1.5. □

By pairing with classes in $H^*(\overline{\mathcal{M}}_{0,n}; \mathbb{Z}) \otimes_{\mathbb{Z}} H^*(X; \mathbb{Z})^{\otimes n}$, we can define $\mathbb{Z}$-valued Gromov–Witten type invariants indexed by the stable isotropy type Γ. When (X, ω) is semi-positive and $[\gamma]$ is the trivial stabilized isotropy type, as argued in [LT98, Remark 4], our invariants should match the usual genus zero Gromov–Witten invariants.

Remark 5.1.8. As communicated from M. Swaminathan, in the forthcoming work of Hirschi–Swaminathan, following [AMS21], smooth global Kuranishi charts for moduli spaces of J-holomorphic map of higher genera will be constructed and such constructions together with stabilization maps and evaluation maps lead to well-defined ele-

ments in $\overline{\Omega}_*^{\mathbb{C},\mathrm{der}}(\overline{\mathcal{M}}_{g,n}\times X^n)$ independent of intermediate auxiliary data. In particular, one could define higher genus integral Gromov–Witten type invariants as in Theorem 5.1.7. We expect these invariants to satisfy variants of the Kontsevich–Manin axioms and to define versions of cohomological field theories on $H^*(X;\mathbb{Z})$.

Another application of our construction is an alternate proof of the main theorem of [AMS21]. This theorem asserts that if $P \to S^2$ is a smooth Hamiltonian fibration over the 2-sphere with fiber given by a closed symplectic manifold (X,ω), there is an additive splitting isomorphism

$$H^*(P;\mathbb{Z}) \cong H^*(X;\mathbb{Z}) \otimes_{\mathbb{Z}} H^*(S^2;\mathbb{Z}). \tag{5.1.4}$$

Via using the virtual class indexed by the trivial group, we provide a simple proof of the isomorphism (5.1.4) in Section 5.2 without using Morava K-theories or Atiyah duality for orbifolds. We remark that [AMS21] shows that the splitting (5.1.4) actually holds for any complex oriented generalized cohomology theory but our proof only works over $\mathbb{Z}$.

5.1.2 A brief discussion of the proof

Our results have their roots in a brilliant proposal of Fukaya–Ono [FO97]. We explain Fukaya–Ono's idea in technical terms and discuss how to realize their proposal using extra technical inputs.

We consider the following simplified situation. Suppose the moduli space is contained in a single Kuranishi chart (G,U,E,S) where G is a finite group, U is a smooth G-manifold, E (the obstruction space) is a finite-dimensional G-representation, and $S : U \to E$ (the Kuranishi map) is a G-equivariant smooth map. Then the moduli space is $S^{-1}(0)/G$. Let $U^G \subset U$ be the set of G-fixed points. We further assume that the normal bundle to U^G is trivial with fibre being a G-representation V and

$U = U^G \times V$. We also assume that the obstruction space E does not have a trivial G-representation as a summand.

Under these assumptions, the map $S : U \to E$ can be identified with a map with target in the space of smooth G-equivariant maps between V and E

$$f : U^G \to C^\infty(V, E)^G, \tag{5.1.5}$$

where for each $p \in M$, $f(p)$ is the restriction of S to the fiber $\{p\} \times V$. For any such map f, we denote the associated Kuranishi map by S_f, i.e.

$$S_f(p, v) = f(p)(v) \in E, \ \forall (p, v) \in U^G \times V. \tag{5.1.6}$$

We wish to define a (relative) homology class of expected degree $\dim U^G + \dim V - \dim E$ represented by $S_f^{-1}(0)/G$ for a transverse section S_f. In general, due to the failure of equivariant transversality, such a transverse section does not exist. Indeed S_f necessarily vanishes along M whose dimension could be much higher than the expected dimension.

Fukaya–Ono proposed a solution by considering *normally complex polynomial perturbations*. Assume that both V and E are *complex* G-representations. In applications, such complex structures comes from the nature of the Cauchy–Riemann operator. Then under the correspondence in (5.1.5) and (5.1.6) we consider perturbations of S_P corresponding to maps of the form

$$P : U^G \to \mathrm{Poly}_d^G(V, E)$$

where $\mathrm{Poly}_d^G(V, E) \subset C^\infty(V, E)^G$ is the space of G-equivariant complex polynomial maps of degree at most d. When d is large enough, Fukaya–Ono showed in [FO97] that for a generic choice of an element $P \in \mathrm{Maps}(U^G, \mathrm{Poly}_d^G(V, E))$ the zero locus

along the stratum with trivial isotropy group under the G-action

$$\mathcal{M}_P^{\text{free}} := s_P^{-1}(0) \cap (U^G \times V^{\text{free}})/G \tag{5.1.7}$$

is a smooth manifold of the expected dimension. Most crucially, because V and E are both complex, the set-theoretic boundary of $\mathcal{M}_P^{\text{free}}$ has codimension at least 2. Therefore, using the notion of pseudocycles or any other essentially equivalent language, an integral homology class can be defined.

However, the right notion of transversality is not only that the free part (5.1.7) is cut out transversely, but also that its boundary (strata) are cut out transversely. Two technical difficulties arise then. First, although the interior transversality behaves well when we increase the cut-off degree d, it was not clear if one has a proper notion of boundary transversality which also behaves well with respect to the *change of degrees*. More seriously, when the action of G on V has additional orbit types, the pointwise transversality condition depends crucially on the stabilizers. It was not clear if the transversality at one point implies the transversality at nearby points with smaller stabilizers. More precisely, for any nontrivial subgroup $H < G$, let $V_H \subset V$ be the space of H-fixed points and $N_H \subset V$ the complement of V_H; write $v \in V$ as (v', v''). Then a map $P : U^G \to \text{Poly}_d^G(V, E)$ induces another normally polynomial section

$$P_H : U^G \times V_H \to \text{Poly}_d^H(N_H, E), \quad P_H(p, v') := P(p)(v', \cdot).$$

One thus needs a good notion of boundary transversality which behaves well with respect to the *change of groups*, i.e., P is transverse implies P_H is transverse (or vice versa). Otherwise we will not have the openness of the transversality condition, which could cause serious difficulties to construct global perturbations.

The technical core of our construction is the resolution of the above issues by defining the right notion of boundary transversality using the language of *Whitney*

stratifications. Indeed, the transversality notion for normally polynomial sections is related to the complex affine variety

$$Z_d^G(V,E) := \{(v,P) \in V \times \mathrm{Poly}_d^G(V,E) \mid P(v) = 0\}$$

and how one decomposes it into strata of manifolds. Specifically, in this chapter we define a canonical Whitney stratification on $Z_d^G(V,E)$ which is a decomposition into algebraic submanifolds that respects the natural *action stratification* of V:

$$V = \bigsqcup_{H<G} V_H^*, \text{ where } V_H^* := \{v \in V \mid \text{the stabilizer of } v \text{ is } H\}.$$

A map $P : U^G \to \mathrm{Poly}_d^G(V,E)$ (and its associated Kuranishi map S_P) is said to be *strongly transverse* (cf. Definition 5.6.1) if it is transverse to the image of all strata of $Z_d^G(V,E)$ under the map $Z_d^G(V,E) \to \mathrm{Poly}_d^G(V,E)$. This Whitney stratification is rather robust so that it transforms naturally with respect to the natural inclusion map

$$Z_d^G(V,E) \to Z_{d+d'}^G(V,E) \tag{5.1.8}$$

related to the change of degree (see Theorem 5.4.9) and the natural restriction map

$$Z_d^G(V,E) \to V_H^* \times Z_d^H(N_H,E) \tag{5.1.9}$$

related to the change of groups (see Theorem 5.4.13). In this way, the notion of strong transversality of normally polynomial sections is intrinsic. The existence of globally defined strongly transverse normal polynomial perturbations and the relevant cobordism invariance results thus follow by the technical arguments in Section 5.6. Therefore, the integral cycles and invariants expected by Fukaya–Ono [FO97] can be defined. Note that the use of the canonical Whitney stratification allows us to speak of transversality along each stratum of the isotropy stratification of the ambient

orbifold. As a result, we actually obtain a system of integral homology classes indexed by isotropy types.

Remark 5.1.9. As mentioned in the course of the proof, there are several crucial points inspired by the unpublished work of Parker [Par13]. Unlike our approach, Parker tried to achieve transversality using a less canonical notion of *nice Whitney stratifications* on the variety $Z_d^G(V,E)$. Unfortunately, it is unclear to us how to remove the dependence on the choice of these auxiliary Whitney stratifications. The canonical Whitney stratification on $Z_d^G(V,E)$ we define in this chapter is intrinsic to $Z_d^G(V,E)$ as it is the "minimal" Whitney stratification respecting the action stratification, and it has better functorial properties.

Remark 5.1.10. Parker also introduced the notion of *normally complex sections* in [Par13] which generalizes normally polynomial sections as from [FO97]. Although normally polynomial sections suffice our needs as demonstrated in Lemma 5.6.2, working with normally complex sections can provide extra flexibility at various points. Hence we work with normally complex sections throughout the chapter and call them the Fukaya–Ono–Parker (FOP for short) sections (Definition 5.3.16).

Note that showing the existence of strongly transverse sections requires us to work with normally complex effective orbifolds and normally complex orbibundles. To establish Theorem 5.1.4 and its corollary Theorem 5.1.7, we make use of the flexibility provided by the stabilization equivalences combined with the "enough vector bundle" theorem proved by Pardon [Par19, Theorem 1.1] to identify a stable complex derived orbifold chart with an effective derived orbifold chart such that the ambient orbifold is almost complex and the orbibundle is complex.

5.1.3 Relations with other work

$\mathbb{Z}$-valued enumerative invariants in symplectic topology

In symplectic topology there are other Gromov–Witten type invariants which take integer values. Firstly, as mentioned above, for semi-positive symplectic manifolds, Ruan–Tian's construction ([RT95]) shows that the genus-0 Gromov–Witten invariants are indeed $\mathbb{Z}$-valued. Secondly, there are the celebrated $\mathbb{Z}$-valued Gopakumar–Vafa invariants of Calabi–Yau 3-folds [GV98a, GV98b] which govern the $\mathbb{Q}$-valued Gromov–Witten invariants as demonstrated in full generality by Ionel–Parker [IP18]. It was speculated by some experts (e.g. [Joy07, Section 6.3]) that the invariants defined using Fukaya–Ono's normally polynomial sections are related to the Gopakumar–Vafa invariants. However, recent developments [DW19a], [BS21], which should be viewed as continuations of the groundbreaking work of Taubes in dimension 4 [Tau96b], construct $\mathbb{Z}$-valued invariants of Calabi–Yau 3-folds by (virtually) counting *embedded* pseudo-holomorphic curves and these invariants seems to be better connected with the Gopakumar–Vafa invariants. Lastly, the K-theoretic Gromov–Witten invariants, defined by Givental and Lee [Lee04] for algebraic objects (which are expected to exist for general symplectic manifolds in light of [AMS21, Section 6.12]), are also $\mathbb{Z}$-valued.

Topological construction and the Floer-theoretic counterpart

Recall that the typical method of defining the ($\mathbb{Q}$-valued) traditional Gromov–Witten invariants goes through the inductive construction of transverse multisections on a Kuranishi atlas (good coordinate system) over a moduli space. As one only needs transversality in the topological sense, smoothness of coordinate changes are not necessary for the construction (see [LT98] and the case of gauged linear sigma model in [TX18, TX21]). It is expected that one can use normally polynomial sections to define Floer homology with $\mathbb{Z}$ or $\mathbb{Z}_2$-coefficients. A relevant work along this line is [FOOO13]

which uses the notion of "normally conical perturbations" to define Lagrangian Floer homology for spherically positive symplectic manifolds. We believe that our integral Euler classes are topological invariants, but to carry out the constructions we need to use topological transversality and the theory of micro-bundles. Working out all the arguments in the topological category would ease up the technical difficulties in deriving Floer-theoretic applications including the weak Arnol'd conjecture over $\mathbb{Z}$, which improves the main result from [AB21].

$\mathbb{Z}$-valued invariants in algebraic geometry

There are several types of enumerative invariants constructed using algebraic geometry, most notably the Donaldson–Thomas invariants [Tho00] and the Pandharipande–Thomas invariants [PT09]. Although these invariants are closely tied with the Gromov–Witten invariants [MNOP06a, MNOP06b], their constructions depend on sheaf theory. Fukaya–Ono's proposal is differential-topological in nature, as it is still a variant of the general position argument. It would be interesting to see if the virtual classes $[\overline{\mathcal{M}}_{0,n}(X,J,A)]^{\mathrm{vir}}_{[\gamma]}$ admit a purely algebro-geometric interpretation. As remarked before, the classes $[\overline{\mathcal{M}}_{0,n}(X,J,A)]^{\mathrm{vir}}_{[\gamma]}$ are refinements of the traditional rational virtual class $[\overline{\mathcal{M}}_{0,n}(X,J,A)]^{\mathrm{vir}}$, thus an algebro-geometric construction of them might shed light on refinements of other curve-counting invariants.

Wasserman's theorem and stable homotopy theory

A renowned theorem of Wasserman [Was69], recaptured in [Par20, Theorem 5.6], provides a sufficient condition for equivariant transversality to hold. Our result can be interpreted as a variant of Wasserman's theorem given the presence of normal complex structures, whose existence allows us to extract more regularity. Rather than developing a theory of stable *normal* complex bordisms, we confine ourselves with stable complex bordisms of derived orbifolds because the Pontryagin–Thom

construction which identifies geometric bordisms with homotopical cobordisms holds in this case [Par20, Theorem 1.4]. The homotopical cobordism perspective gives far-reaching corollaries of Wasserman's theorem, see [Sch18, Theorem 6.2.33]. As mentioned by [Par20, Remark 5.7], it is an interesting question to understand Fukaya–Ono's proposal on the homotopical cobordism side.

5.1.4 Plan of the chapter

We start with a new proof of Abouzaid–McLean–Smith's cohomological splitting theorem over $\mathbb{Z}$ in Section 5.2 to demonstrate the utility of our construction. Then we introduce basic notions related to orbifolds in Section 5.3, especially the notion of *normal complex structure* (Definition 5.3.13) and *FOP section* (Definition 5.3.18). In Section 5.4, we study the canonical Whitney stratification on the universal zero locus $Z_d^G(V,W)$ and understand its behavior under the change of degrees and groups. Section 5.5 proves that any pseudocycle in a Thom–Mather stratified space defines an integral homology class. The main theorem is established in Section 5.6. In Section 5.7, we discuss how to use FOP perturbations to study stable complex derived bordisms over orbispaces. Finally, the necessary background and results on Whitney stratifications are discussed in Appendix 5.8.

Remark 5.1.11. When defining the $\mathbb{Z}$-valued genus zero curve counting invariants, or equivalently the homology class (5.1.3), as $\overline{\mathcal{M}}_{0,n}$ is a manifold, we do not need the general result about pseudocycles in Thom–Mather stratified spaces. We also do not need to the more general algebraic topology result Theorem 5.1.4. Hence results in Section 5.5 and Section 5.7 can be bypassed.

5.1.5 Acknowledgements

We would like to thank Mohammed Abouzaid, Kenji Fukaya, Helmut Hofer, John Pardon, Paul Seidel, Mohan Swaminathan, and Dingyu Yang for useful correspon-

dences and discussions at various stages of this project.

5.2 A simple proof of Abouzaid–McLean–Smith's splitting theorem

Let (X, ω) be a closed symplectic manifold. Denote by $\mathrm{Ham}(X, \omega)$ the group of Hamiltonian diffeomorphisms of X. Given an element $\phi \in \pi_1(\mathrm{Ham}(X, \omega))$, one can construct a fibration $P_\phi \to S^2$ with fiber symplectomorphic to X and structure group $\mathrm{Ham}(X, \omega)$ from the clutching construction using ϕ. We use Theorem 5.6.2 to give an alternative proof of [AMS21, Theorem 1.1]:

Theorem 5.2.1. There is an additive splitting $H^*(P_\phi; \mathbb{Z}) \cong H^*(X; \mathbb{Z}) \otimes_{\mathbb{Z}} H^*(S^2; \mathbb{Z})$.

Denote by $res : H^*(P_\phi; \mathbb{Z}) \to H^*(X; \mathbb{Z})$ the restriction map induced by the inclusion of a fiber $X \hookrightarrow P_\phi$. The proof is based on constructing a map $s_\phi : H^*(X; \mathbb{Z}) \to H^*(P_\phi; \mathbb{Z})$ such that the composition $res \circ s_\phi$ is the identity map on $H^*(X; \mathbb{Z})$. Following [AMS21, Section 3], we introduce the following auxiliary spaces. Let $\mathbb{S}$ be the one-point blow-up of $\mathbb{CP}^1 \times \mathbb{CP}^1$. The composition of the blow-down map and the projection to the second factor defines a map $\pi_B : \mathbb{S} \to B = \mathbb{CP}^1$ which is a $\mathbb{CP}^1$ fibration with one reducible fiber over $0 \in B$. For each $t \in B$, write $\mathbb{S}_t := \pi_B^{-1}(t)$. Then there exists a symplectic fibration $\pi_{\mathbb{S}} : \tilde{P} \to \mathbb{S}$ such that:

1. over a neighborhood W_∞ of $\infty \in B$, the space $\tilde{P}$ is isomorphic to the trivial symplectic fibration $(S^2 \times W_\infty \times X, \omega_{S^2} \oplus \omega_{S^2}|_{W_\infty} \oplus \omega)$ and the projection $\pi_B : \pi_B^{-1}(W_\infty) \to W_\infty$ coincides with the natural projection $S^2 \times W_\infty \to W_\infty$;

2. the preimage of the reducible fiber $\mathbb{S}_0 \cong \mathbb{CP}^1 \vee \mathbb{CP}^1$ under $\pi_{\mathbb{S}}$ is given by $P_\phi \cup_X P_{\phi^{-1}}$ such that each component carries the canonical deformation class of the symplectic structure. Denote $\mathbb{S}_0 = \mathbb{S}_\phi \vee \mathbb{S}_{\phi^{-1}}$ where $\mathbb{S}_\phi$ resp. $\mathbb{S}_{\phi^{-1}}$ is the base of P_ϕ resp. $P_{\phi^{-1}}$.

Denote by S^2_{hor} the image of a holomorphic section of $\pi_B : \mathbb{S} \to B = \mathbb{CP}^1$ which passes through $\mathbb{S}_\phi \setminus \mathbb{S}_{\phi^{-1}}$. Let $(S^2 \times X)_{hor}$ be the preimage of S^2_{hor} under $\pi_{\mathbb{S}}$. Then $(S^2 \times X)_{hor}$ is symplectically deformation equivalent to the product $(S^2 \times X, \omega_{S^2} \oplus \omega)$. Define X_t to be the the fiber of the projection $\pi_{\mathbb{S}}|_{(S^2 \times X)_{hor}} : (S^2 \times X)_{hor} \to S^2_{hor}$ over $t \in S^2_{hor}$. Denote by P_t the space $\tilde{P}|_{\mathbb{S}_t}$ for $t \neq 0$. See Figure 5.1 for illustration.

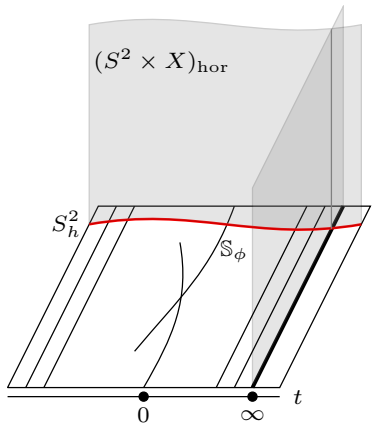


Figure 5.1: The fibration $\tilde{P}$.

We choose a compatible almost complex structure J on $\tilde{P}$ making the projection $\pi_{\mathbb{S}} : \tilde{P} \to \mathbb{S}$ pseudo-holomorphic. Moreover, we require that over $S^2 \times W_\infty \times X$, the almost complex structure J is the direct sum of a compatible almost complex structure on X and the standard integrable complex structures on the S^2 factors. Let $A \in H_2(\tilde{P}; \mathbb{Z})$ be the homology class represented by $S^2 \times \{\infty\} \times \{pt\} \subset S^2 \times W_\infty \times X$. We consider the moduli space $\overline{\mathcal{M}}_{0,2}(X, J, A)$ with the evaluation map $ev : \overline{\mathcal{M}}_{0,2}(X, J, A) \to \tilde{P} \times \tilde{P}$. Define

$$\overline{\mathcal{M}}_h := ev^{-1}(\tilde{P} \times (S^2 \times X)_{hor})$$

and its subspaces $\overline{\mathcal{M}}_\infty := ev^{-1}(P_\infty \times (S^2 \times X)_{hor})$ and $\overline{\mathcal{M}}_\phi := ev^{-1}(P_\phi \times (S^2 \times X)_{hor})$.

After applying the evaluation map, we can view $\overline{\mathcal{M}}_\phi$ as a correspondence

$$\begin{array}{ccc} & \overline{\mathcal{M}}_\phi & \\ {\scriptstyle \pi_\phi}\swarrow & & \searrow{\scriptstyle \pi_{hor}} \\ P_\phi & & (S^2\times X)_{hor}. \end{array} \tag{5.2.1}$$

Then given any homology class $\alpha \in H_*(\overline{\mathcal{M}}_\phi;\mathbb{Z})$, we can define a homomorphism

$$\Phi_\alpha : H^*(S^2\times X;\mathbb{Z}) \to H^*(P_\phi;\mathbb{Z}) \tag{5.2.2}$$

by the composition

$$H^*((S^2\times X)_{hor};\mathbb{Z}) \xrightarrow{\pi_{hor}^*} H^*(\overline{\mathcal{M}}_\phi;\mathbb{Z}) \xrightarrow{\cap\alpha} H_*(\overline{\mathcal{M}}_\phi;\mathbb{Z}) \xrightarrow{(\pi_\phi)_*} H_*(P_\phi;\mathbb{Z}) \xrightarrow{\mathrm{PD}} H^*(P_\phi;\mathbb{Z}),$$

where PD is the Poincaré duality isomorphism. The map $s_\phi : H^*(X;\mathbb{Z}) \to H^*(P_\phi;\mathbb{Z})$ leading to a proof of Theorem 5.2.1 is constructed by composing certain Φ_α with the map $pr^* : H^*(X;\mathbb{Z}) \to H^*(S^2\times X;\mathbb{Z})$ induced by the projection $pr : S^2\times X \to X$. To be more precise, such a class α is produced using the FOP perturbation and it actually lies inside a thickening of the moduli space $\overline{\mathcal{M}}_\phi$.

Lemma 5.2.1. There exists a smooth stable complex compact effective derived orbifold chart (D,E,s) such that there exists an isomorphism $s^{-1}(0) \cong \overline{\mathcal{M}}_{0,2}(X,J,A)$. Moreover, the evaluation maps $ev_i : \overline{\mathcal{M}}_{0,2}(X,J,A) \to \tilde{P}$ for $i=1,2$ extend to

$$\begin{array}{ccc} & D & \\ {\scriptstyle \tilde{ev}_1}\swarrow & & \searrow{\scriptstyle \tilde{ev}_2} \\ \tilde{P} & & \tilde{P}. \end{array}$$

which are smooth submersions.

Proof. This is a reformulation of the results from [AMS21, Section 5.8]. The construction of stable complex structure is standard in Gromov–Witten theory. □

Lemma 5.2.2. Define D_∞ by the fiber square

$$\begin{array}{ccc} D_\infty & \longrightarrow & D \\ \downarrow & & \downarrow{\scriptstyle \tilde{ev}_1 \times \tilde{ev}_2} \\ P_\infty \times (S^2 \times X)_{hor} & \xrightarrow[\text{inclusion}]{} & \tilde{P} \times \tilde{P} \end{array}.$$

Then the triple $(D_\infty, E_\infty, s_\infty) := (D_\infty, E|_{D_\infty}, s|_{D_\infty})$ is a derived orbifold chart with stable (normal) complex structure such that $s_\infty^{-1}(0) \cong \overline{\mathcal{M}}_\infty$. Similarly, define D_h by the fiber square

$$\begin{array}{ccc} D_h & \longrightarrow & D \\ \downarrow & & \downarrow{\scriptstyle \tilde{ev}_1 \times \tilde{ev}_2} \\ \tilde{P} \times (S^2 \times X)_{hor} & \xrightarrow[\text{inclusion}]{} & \tilde{P} \times \tilde{P} \end{array}.$$

Then $(D_h, E_h, s_h) := (D_h, E|_{D_h}, s|_{D_h})$ is a derived orbifold chart with stable (normal) complex structure such that $s_h^{-1}(0) \cong \overline{\mathcal{M}}_h$.

Proof. This follows from [AMS21, Corollary 5.36]. Note that the submersive property of the evaluation maps is used here. □

Proof of Theorem 5.2.1. Consider the stable (normal) complex derived orbifold chart (D_h, E_h, s_h) constructed above. Using Theorem 5.6.2 associated to the trivial isotropy type, we can find a homology class $\alpha_h \in H_*(D_h; \mathbb{Z})$ by choosing a generic FOP perturbation of the section s_h. Apply the correspondence construction to the diagram

$$\begin{array}{ccccc} & & D_h & & \\ & \swarrow^{\tilde{ev}_1} & & \searrow^{\tilde{ev}_2} & \\ \tilde{P} & & & & (S^2 \times X)_{hor} \end{array}$$

we obtain a homomorphism $\Psi_h : H^*(S^2 \times X; \mathbb{Z}) \to H^*(\tilde{P}; \mathbb{Z})$.

Using the inclusion maps $\iota_\infty : P_\infty \hookrightarrow \tilde{P}$ and $\iota_\phi : P_\phi \hookrightarrow \tilde{P}$, define $\Psi_\infty := \iota_\infty^* \circ \Psi_h$

and $\Psi_\phi := \iota_\phi^* \circ \Psi_h$. They fit into the following commutative diagram

$$\begin{array}{ccccccc} & & & & H^*(P_\infty;\mathbb{Z}) & & \\ & & \overset{\Psi_\infty}{\nearrow} & & \uparrow{\scriptstyle \iota_\infty^*} & \overset{res}{\searrow} & \\ H^*(X;\mathbb{Z}) & \xrightarrow{pr^*} & H^*((S^2\times X)_{hor};\mathbb{Z}) & \xrightarrow{\Psi_h} & H^*(\tilde{P};\mathbb{Z}) & \xrightarrow{res} & H^*(X;\mathbb{Z}) \\ & & \underset{\Psi_\phi}{\searrow} & & \downarrow{\scriptstyle \iota_\phi^*} & \underset{res}{\nearrow} & \\ & & & & H^*(P_\phi;\mathbb{Z}) & & \end{array} \tag{5.2.3}$$

where the rightmost arrows are induced from the inclusion of a fiber. The split right inverse s_ϕ of the map $res : H^*(P_\phi;\mathbb{Z}) \to H^*(X;\mathbb{Z})$ is defined to be $\Psi_\phi \circ pr^*$. To prove $res \circ s_\phi = id$, it suffices to show that $res \circ \Psi_\infty \circ pr^* = id$.

Let us look at the moduli space $\overline{\mathcal{M}}_h$ geometrically. Recall that over a neighborhood W_∞ of $\infty \in B$, the fibration $\pi_{\mathbb{S}} : \tilde{P} \to \mathbb{S}$ is identified with $W_\infty \times S^2 \times X \to S^2 \times W_\infty$ and the almost complex structure J splits. Moreover, the degree A of the curves is generated by the S^2-factor of the product. As a result, one has a homeomorphism

$$\overline{\mathcal{M}}_{W_\infty} := ev_2^{-1}(\pi_{\mathbb{S}}^{-1}(W_\infty \times S^2)) \cap \overline{\mathcal{M}}_h \cong W_\infty \times S^2 \times X. \tag{5.2.4}$$

In particular, this open subset of $\overline{\mathcal{M}}_h$ is regular and consists of elements with trivial automorphism groups. Therefore, for the derived orbifold chart (D_h, E_h, s_h) constructed for the moduli space $\overline{\mathcal{M}}_h$, the section s_h is transverse over $\overline{\mathcal{M}}_{W_\infty}$. Notice that s_h is also a strongly regular FOP section near $\overline{\mathcal{M}}_{W_\infty}$ because of the triviality of the automorphism groups.

Now we prove that $res \circ \Psi_\infty \circ pr^* = id$. Fix $a \in H^*(X)$. Then by the main theorem of [Zin08], the Poincaré dual $\mathrm{PD}(a) \in H_*(X)$ can be represented by a pseudocycle $\iota_a : N \to X$. The pullback $\tilde{\iota} : S^2 \times N \to (S^2 \times X)_h \cong S^2 \times X$ then represents the Poincaré dual of $pr^* a$. We may choose a strongly transverse FOP section $s_h' : D_h \to E_h$ such that it agrees with s_h near $\overline{\mathcal{M}}_{W_\infty}$ and such that the evaluation map $ev_1 : (s_h')^{-1}(0) \to (S^2 \times X)_h$ is transverse to $\tilde{\iota}(S^2 \times N)$ as well as its boundary. Then

the class $ev_1^*(pr^*a) \cap \alpha_h$ is represented by the pseudocycle $ev_1^*(\tilde{\iota}(S^2 \times N)) \cap (s_h')^{-1}(0)$. By our choice of the FOP perturbation s_h', we see that near W_∞, this pseudocycle is identified with $W_\infty \times S^2 \times N$ with respect to the homeomorphism (5.2.4). Then after pushing forward by the evaluation map ev_2 and intersecting with the fibre $P_\infty \in \tilde{P}$, one can see that the cohomology class $\iota_\infty^* \Psi_h(pr^*a) \in H^*(P_\infty)$ is the Poincaré dual to the pseudocycle $S^2 \times N$ in $P_\infty \cong S^2 \times X$. Then after restricting to a fibre $\{pt\} \times X \subset S^2 \times X$, the resulting cohomology class is the Poincaré dual of N, i.e., a itself. $\square$

Remark 5.2.2. Following Theorem 5.1.7, if we denote by $\Lambda_{\mathbb{Z}}$ the universal Novikov ring over $\mathbb{Z}$

$$\{\sum a_i T^{\lambda_i} | a_i \in \mathbb{Z}, \lambda_i \in \mathbb{R}, \lambda_i \xrightarrow{i \to \infty} +\infty\},$$

one can define a quantum product over $QH_{\mathbb{Z}}^*(X) := H^*(X; \Lambda_{\mathbb{Z}})$ after reducing the $\mathbb{Z}$-grading to $\mathbb{Z}/2$-grading by considering the 3-pointed virtual fundamental class associated to the trivial isotropy type. Given an element $\phi \in \pi_1(\mathrm{Ham}(X, \omega))$, one could construct an invertible element in $QH_{\mathbb{Z}}^*(X)$ by counting pseudo-holomorphic sections of P_ϕ, which is the analogue of Seidel's element in our setting. After proving a suitable gluing theorem, one can prove Theorem 5.2.1 using Seidel's representation as in [LMP99].

5.3 Normally complex orbifolds and bundles

5.3.1 Orbifolds and orbifold vector bundles

We recall the basic definition of effective orbifolds. We follow the definition of [ALR07, Section 1.1]. Working with effective orbifolds allows us to use orbifold charts exclusively without appealing to the language of groupoids. We remark that for our applications in Section 5.7, we can drop effectiveness by choosing an effective presentation of any given derived orbifold.

Let X be a Hausdorff and second countable topological space. An n-dimensional orbifold chart of X is a triple

$$C = (U, \Gamma, \psi)$$

where $U \subset \mathbb{R}^n$ is a nonempty open subset, Γ is a finite group acting effectively and smoothly on U, and $\psi : U \to X$ is a Γ-invariant continuous map such that the induced map

$$\underline{\psi} : U/\Gamma \to X$$

is a homeomorphism onto an open subset of X. If $x \in \psi(U)$ we also say that x is contained in the chart C.

A chart embedding from another chart $C' = (U', \Gamma', \psi')$ to C is a smooth (open) embedding $\iota : U' \hookrightarrow U$ such that

$$\psi \circ \iota = \psi'.$$

It follows that (see [ALR07, Page 3]) given a chart embedding ι as above there exists a canonical group injection $\Gamma' \hookrightarrow \Gamma$ such that ι is equivariant. Therefore we often include the group injection as part of the data of a chart embedding.

As we are in the smooth category, we can always find "linear" charts around any point. An orbifold chart $C = (U, \Gamma, \psi)$ is called linear if Γ acts linearly on $\mathbb{R}^n$ and $U \subset \mathbb{R}^n$ is an invariant open subset. We say that a linear chart is centered at $x \in X$ if $0 \in U$ and $x = \psi(0)$.

We say two charts $C_i = (U_i, \Gamma_i, \psi_i)$, $i = 1, 2$ are compatible if for each $p \in \psi_1(U_1) \cap \psi_2(U_2)$, there exists an orbifold chart $C_p = (U_p, \Gamma_p, \psi_p)$ and chart embeddings into both C_1 and C_2.

An orbifold atlas $\mathcal{A} = \{C_i \mid i \in I\}$ on X is a collection of mutually compatible charts C_i which cover X. We say an atlas $\mathcal{A}' = \{C'_j \mid j \in J\}$ refines $\mathcal{A}$, equivalently, $\mathcal{A}'$ is a refinement of $\mathcal{A}$, if for each C'_j there exists a chart embedding $C'_j \hookrightarrow C_i$ for some

$i \in I$. We say two orbifold atlases are equivalent if they have a common refinement. A topological space X together with an equivalence class of orbifold atlases is called a smooth effective orbifold.

Let X be an effective orbifold. A continuous function $f : X \to \mathbb{R}$ is called smooth if its pullback to each chart is a smooth function.

Remark 5.3.1. One can see that orbifolds are all locally compact. As we also assume they are Hausdorff and second countable, they are paracompact spaces. Hence for any open cover by charts, there exists a subordinate continuous partition of unity. Moreover, as one can approximate continuous functions by smooth functions on each chart, there always exist a subordinate smooth partition of unity.

Remark 5.3.2. We also need the notion of orbifolds with boundary. In that case, the domain of a chart $C = (U, \Gamma, \psi)$ is allowed to be a smooth manifold with boundary. The Γ-action is required to fixed the boundary $\partial U \subset U$ set-wise; moreover, for each $x \in \partial U$, the stabilizer of x is required to fix ∂U pointwise.

Similarly we can define the notion of orbifold vector bundles. Let X be an orbifold, $\mathcal{E}$ be a topological space, and $\pi : \mathcal{E} \to X$ be a continuous map. A bundle chart of $\mathcal{E} \to X$ consists of an orbifold chart $C = (U, \Gamma, \psi)$ of X, a Γ-equivariant smooth vector bundle $\pi_C : E \to U$, and a Γ-invariant continuous map $\hat{\psi} : E \to \mathcal{E}$ such that the following diagram commutes:

$$\begin{array}{ccc} E & \xrightarrow{\hat{\psi}} & \mathcal{E} \\ {\scriptstyle \pi_C}\downarrow & & \downarrow{\scriptstyle \pi} \\ U & \xrightarrow[\psi]{} & X. \end{array}$$

In notation we will use a quadruple $\hat{C} = (U, \Gamma, E, \hat{\psi})$ to denote the bundle chart where the map $\psi : U \to X$ is determined by the map $\hat{\psi} : E \to \mathcal{E}$. We can similarly define the notions of chart embeddings, compatibility, atlases for bundles. Then an orbifold vector bundle structure over $\pi : \mathcal{E} \to X$ is defined to be an equivalence class of bundle

atlases as before. We spell out the definition of sections of an orbifold vector bundle because of their importance in this chapter.

Definition 5.3.3 (Sections)**.** Let $\mathcal{E} \to X$ be an orbifold vector bundle.

1. Let $\hat{C}_i = (U_i, E_i, \Gamma_i, \hat{\psi}_i)$, $i = 1, 2$ be two bundle charts. We say that a Γ_1-equivariant section $s_1 : U_1 \to E_1$ and a Γ_2-equivariant section $s_2 : U_2 \to E_2$ are compatible if for any bundle chart $\hat{C}_0 = (U_0, E_0, \Gamma_0, \hat{\psi}_0)$ of E and chart embeddings $\hat{\iota}_i : \hat{C}_0 \hookrightarrow \hat{C}_i$, $i = 1, 2$ there holds

$$\hat{\iota}_1^{-1} \circ s_1 \circ \iota_1 = \hat{\iota}_2^{-1} \circ s_2 \circ \iota_2$$

   as sections of $E_0 \to U_0$.

2. Suppose $\hat{\mathcal{A}} = \{\hat{C}_i = (U_i, E_i, \Gamma_i, \hat{\psi}_i) \mid i \in I\}$ is a bundle atlas on $\mathcal{E}$. A section of $(\mathcal{E}, \hat{\mathcal{A}})$ is a collection $s = (s_i)_{i \in I}$ where $s_i : U_i \to E_i$ is a Γ_i-equivariant smooth section such that each pair in this collection are compatible.

3. We say two atlases together with sections $(\hat{\mathcal{A}}_1, s_1)$ and $(\hat{\mathcal{A}}_2, s_2)$ are equivalent if $\hat{\mathcal{A}}_1$ and $\hat{\mathcal{A}}_2$ are equivalent as orbifold atlases and their local representations are all pairwise compatible.

4. A section of $\mathcal{E}$ is an equivalence classes of pairs $(\hat{\mathcal{A}}, s)$.

Lemma 5.3.1. For each section of $\mathcal{E} \to X$ and each bundle atlas $\hat{\mathcal{A}}$, there exists a representative $(\hat{\mathcal{A}}, s)$ of the given section.

Proof. Straightforward. □

5.3.2 The isotropy prestratification

One can see from the definition of effective orbifolds that the isomorphism class of the stabilizer of a point x in an orbifold chart $C = (U, \Gamma, \psi)$ only depends on the point

$\psi(x) \in X$. One can use this isomorphism class to decompose the orbifold. Moreover, we can actually obtain a more refined decomposition by using the information in the normal direction.

Definition 5.3.4. Consider triples (G, V, W) where G is a finite group, V and W are finite-dimensional (real) representations of G such that the decompositions of V and W into irreducible representations contain no trivial summands. Two triples (G, V, W) and (G', V', W') are called *isomorphic* if there is a group isomorphism $\varphi_G : G \to G'$ and equivariant linear isomorphisms $\varphi_V : V \to V'$, $\varphi_W : W \to W'$. An isomorphism class of triples is called an *isotropy type*, denoted by γ.

Now consider an effective orbifold X with an orbifold vector bundle $\mathcal{E} \to X$. For each $p \in X$, consider a bundle chart $\hat{C} = (U, G, E, \hat{\psi})$ centered at p. Let $U^G \subset U$ be the G-fixed point locus. Then over U^G one has the G-equivariant decomposition

$$TU|_{U^G} \cong TU^G \oplus NU^G$$

where the first resp. the second summand is the direct sum of trivial resp. nontrivial summands of the splitting of fibers into irreducible G-representations. One also has the decomposition

$$E|_{U^G} \cong \mathring{E}^G \oplus \check{E}^G$$

where $\mathring{E}^G$ is a trivial G-representation and $\check{E}^G$ is the direct sum of nontrivial irreducible G-representations. Define γ_p to be the isotropy type represented by the triple (G, V, W) where V is fiber of $NU^G \to U^G$ and W is the fiber of $\check{E}^G|_{U^G}$. Then γ_p is well-defined, independent of the choice of charts.

Lemma 5.3.2. For each isotropy type γ represented by a triple $(G_\gamma, V_\gamma, W_\gamma)$, define

$$X^*_\gamma = \{p \in X \mid \gamma_p = \gamma\}$$

Then X^*_γ is a topological manifold of dimension $\dim X - \dim V_\gamma$ equipped with a natural smooth structure.

Proof. By definition, for each $p \in X^*_\gamma$, there exists a bundle chart $(U, G, E, \hat{\psi})$ centered at p such that γ is represented by (G, V, W) where V is the fibre of the normal bundle to $U^G \subset U$ and W is the fibre of $\check{E}^G$ at 0. Then U^G has dimension equal to $\dim X - \dim V$. Moreover, a neighborhood of p in X^*_γ is homeomorphic to $\psi(U^G)$. Hence X^*_γ is a topological manifold and $\psi^{-1} : X^*_\gamma \cap \psi(U) \to U^G \subset U$ is a local chart. It is easy to verify that the atlas on X^*_γ obtained in this way is a smooth atlas. □

By Lemma 5.3.2, we obtain a decomposition of the orbifold X as a locally finite disjoint union of locally closed subsets

$$X = \bigsqcup_\gamma X^*_\gamma. \tag{5.3.1}$$

An argument similar to the proof of Lemma 5.3.2 shows that the frontier condition is satisfied. Therefore, (5.3.1) defines a prestratification (see Appendix 5.8) called the *isotropy prestratification*.

For each isotropy type γ represented by $(G_\gamma, V_\gamma, W_\gamma)$, define

$$n_\gamma := \dim X - \mathrm{rank}\mathcal{E} - \dim V_\gamma + \dim W_\gamma \in \mathbb{Z}. \tag{5.3.2}$$

This is the expected dimension of the intersection of X^*_γ and the zero locus of a single-valued section $s : X \to \mathcal{E}$ (see Proposition 5.6.4).

5.3.3 Riemannian metric

We need to specify certain auxiliary data to aid our later construction. These data include Riemannian metrics on orbifolds and connections on orbifold vector bundles.

Definition 5.3.5. A Riemannian metric on an effective orbifold X is a collection of invariant Riemannian metrics g_C on all orbifold charts $C = (U, \Gamma, \psi)$ such that every chart embedding

$$(U', \Gamma', \psi') \hookrightarrow (U, \Gamma, \psi)$$

is isometric.

One can use the standard way to construct Riemannian metrics on orbifolds using smooth partition of unity on orbifolds (see Remark 5.3.1).

Bundle metrics

We need some technical discussion on metrics on total spaces of vector bundles. Suppose (M, g^{TM}) is a Riemannian manifold. Let $\pi : E \to M$ be a vector bundle equipped with a metric h^E and a metric-preserving connection ∇^E. On the total space E we define a Riemannian metric g^{TE} as follows. Via the horizontal-vertical decomposition $TE \cong \pi^*TM \oplus \pi^*E$ induced by the connection ∇^E we define $g^{TE} = \pi^*g^{TM} \oplus \pi^*h^E$. We call g^{TE} the *bundle metric* induced from g^{TM}, h^E, and ∇^E.

Lemma 5.3.3. Suppose E is endowed with the bundle metric as above. Then

1. Each fiber $E_p \subset E$ is totally geodesic.
2. We identify the normal bundle of the 0-section $(TM)^\perp \to M$ naturally with E. Then the connection on $(TM)^\perp$ induced from the Levi–Civita connection of g^{TE} coincides with ∇^E.

Proof. (1) As the restriction of g^{TE} to each fiber is Euclidean, geodesics in each fiber are straight line segments. To show that a fiber is totally geodesic, one only needs to show that any (short) straight line segments are also geodesics in the totally space. Fix $x \in M$. Choose a local orthonormal frame $e_1, \ldots, e_k$ of E defined over a neighborhood $U_x \subset M$ of x. For any $p \in E_x$, choose $r > 0$ sufficiently small such that for any

$q \in B_r(p, E_x)$, the shortest geodesic connecting p and q are contained in $\pi^{-1}(U_x)$. We also assume that U_x supports a local coordinate chart $\varphi = (x^1, \ldots, x^m) : U_x \to \mathbb{R}^m$. One only needs to show that the straight line segment between p and any $q \in B_r(p, E_x)$ is the shortest geodesic. Suppose $\gamma : [0, d] \to E$ is the shortest geodesic parametrized by arc length. Let $y^1, \ldots, y^k$ be the bundle coordinates associated to the frame $e_1, \ldots, e_k$ and write

$$\gamma(t) = (x^1(t), \ldots, x^m(t), y^1(t), \ldots, y^k(t)).$$

We can decompose $\dot{\gamma}(t)$ orthogonally as

$$\dot{\gamma}(t) = \dot{\gamma}^H(t) + \dot{\gamma}^V(t)$$

where $\dot{\gamma}^H(t)$ resp. $\dot{\gamma}^V(t)$ is the horizontal resp. vertical part. Then one has

$$1 = |\dot{\gamma}(t)| = \sqrt{|\dot{\gamma}^H(t)|^2 + |\dot{\gamma}^V(t)|^2} \geq |\dot{\gamma}^V(t)| = \sqrt{\dot{y}^1(t)^2 + \cdots + \dot{y}^k(t)^2}.$$

Hence

$$d \geq \int_0^d \sqrt{\dot{y}^1(t)^2 + \cdots + \dot{y}^k(t)^2} dt \geq d_{E_x}(p, q).$$

However, as γ is the shortest geodesic, the above must be an equality. Hence $\gamma(t)$ is contained in the fiber E_x and $\gamma(t)$ is the straight line segment. Hence the fiber is totally geodesic.

(2) Choose local coordinates $(x^1, \ldots, x^m, y^1, \ldots, y^k)$ on the total space as above where the bundle coordinates y^α correspond to the choice of local orthogonal frame $e_1, \ldots, e_k$ of E. Suppose

$$\nabla^E_{\partial_i} e_\alpha = \omega^\beta_{i,\alpha} e_\beta.$$

To show that the induced connection on the normal bundle $(TM)^\perp \cong E$ coincides

with ∇^E, one only needs to verify that

$$\langle \nabla_i^{TE} \frac{\partial}{\partial y^\alpha}, \frac{\partial}{\partial y^\beta} \rangle = \omega_{i,\alpha}^\beta.$$

Indeed, by the formula for the Levi–Civita connection, one has

$$\begin{aligned}
\langle \nabla_i^{TE} \frac{\partial}{\partial y^\alpha}, \frac{\partial}{\partial y^\beta} \rangle &= \frac{1}{2} \left(\frac{\partial}{\partial x^i} \langle \frac{\partial}{\partial y^\alpha}, \frac{\partial}{\partial y^\beta} \rangle + \frac{\partial}{\partial y^\alpha} \langle \frac{\partial}{\partial x^i}, \frac{\partial}{\partial y^\beta} \rangle - \frac{\partial}{\partial y^\beta} \langle \frac{\partial}{\partial x^i}, \frac{\partial}{\partial y^\alpha} \rangle \right) \\
&= \frac{1}{2} \left(\frac{\partial}{\partial y^\alpha} \langle \sum_{\gamma,\delta} \omega_{i,\gamma}^\delta y^\gamma \frac{\partial}{\partial y^\delta}, \frac{\partial}{\partial y^\beta} \rangle - \frac{\partial}{\partial y^\beta} \langle \sum_{\gamma,\delta} \omega_{i,\gamma}^\delta y^\gamma \frac{\partial}{\partial y^\delta}, \frac{\partial}{\partial y^\alpha} \rangle \right) \\
&= \frac{1}{2} \left(\omega_{i,\alpha}^\beta - \omega_{i,\beta}^\alpha \right) \\
&= \omega_{i,\alpha}^\beta.
\end{aligned}$$

Therefore $\nabla^{(TM)^\perp}$ coincides with ∇^E. □

We would like to identify tubular neighborhoods of fixed point locus with disk bundles of the normal bundle and equip the tubular neighborhoods with bundle metrics. To globalize such bundle metric construction we need to verify certain properties of bundle metrics related to orthogonal decompositions of vector bundles. Let (M, g^{TM}) be a Riemannian manifold and $\pi_{\mathring{E}} : \mathring{E} \to M$, $\pi_{\check{E}} : \check{E} \to M$ be two vector bundles equipped with metrics $h^{\mathring{E}}$, $h^{\check{E}}$ and metric preserving connections $\nabla^{\mathring{E}}$, $\nabla^{\check{E}}$. Denote $E = \mathring{E} \oplus \check{E}$ equipped with the product metric $h^E = h^{\mathring{E}} \oplus h^{\check{E}}$ and the product connection $\nabla^E = \nabla^{\mathring{E}} \oplus \nabla^{\check{E}}$. Then over the total space of E there is the induced bundle metric g^{TE}. On the other hand, the total space of E can also be viewed as the total space of the bundle

$$\pi_{\mathring{E}}^* \check{E} \to \mathring{E}.$$

There is hence another bundle metric $g^{T\pi_{\mathring{E}}^* \check{E}}$ induced from the base metric $g^{T\mathring{E}}$ on $\mathring{E}$, the fiberwise bundle metric $\pi_{\mathring{E}}^* h^{\check{E}}$, and the connection $\pi_{\mathring{E}}^* \nabla^{\check{E}}$. We would like to show that g^{TE} and $g^{T\pi_{\mathring{E}}^* \check{E}}$ coincide.

Corollary 5.3.6. The following items are true.

1. Over $\mathring{\check{E}} \subset E$, the direct sum decomposition

$$TE|_{\mathring{\check{E}}} \cong T\mathring{\check{E}} \oplus \pi_{\mathring{\check{E}}}^* \check{E}$$

   is orthogonal with respect to the bundle metric g^{TE}. Moreover,

$$g^{TE}|_{\pi_{\mathring{\check{E}}}^* \check{E}} = \pi_{\mathring{\check{E}}}^* h^{\check{E}}.$$

2. The connection on $\pi_{\mathring{\check{E}}}^* \check{E}$ (viewed as a subbundle of $TE|_{\mathring{\check{E}}}$) induced from the Levi–Civita connection of g^{TE} coincides with the pullback $\pi_{\mathring{\check{E}}}^* \nabla^{\check{E}}$.

3. For $(x, \mathring{v}) \in \mathring{\check{E}}|_x$ and $\check{v} \in \check{E}|_x$, viewing the latter as a normal vector in $(T_{(x,\mathring{v})}\mathring{\check{E}})^\perp$, one has

$$\exp_{(x,\mathring{v})}(\check{v}) = (x, \mathring{v} + \check{v}).$$

4. The bundle metric g^{TE} on E (viewed as the total space of $E \to M$) coincides with the bundle metric $g^{T\pi_{\check{E}}^* \check{E}}$ on E (viewed as the total space of $\pi_{\check{E}}^* \check{E}$).

Proof. (1) is straightforward, following from the definition of the bundle metric. (2) is a consequence of the second item of Lemma 5.3.3. (3) follows from the first item of Lemma 5.3.3. (4) is a straightforward check using (1)—(3). □

Straightened metrics

Now we consider Riemannian metrics on an orbifold chart. Let U be a smooth manifold acted on by a finite group Γ. Let g be a Γ-invariant Riemannian metric on U. Then for each subgroup $G \subseteq \Gamma$, the metric induces the following objects.

1. The orthogonal decomposition splitting $TU|_{U^G} \cong TU^G \oplus NU^G$ which agrees

with the decomposition into the direct sum of trivial G-representations and nontrivial G-representations.

2. A metric on the normal bundle $NU^G \to U^G$. Let $N_\epsilon(U^G) \subset NU^G$ be the ϵ-disk bundle.

3. The restriction of the Levi–Civita connection on NU^G. This connection then induces a splitting of TNU^G into horizontal and vertical distributions and hence a Riemannian *bundle* metric on the total space of NU^G.

4. The exponential map $N_\epsilon(U^G) \to U$ which pushes forward the metric on the total space of NU^G to a metric on a tubular neighborhood of U^G.

Definition 5.3.7. Let $G \subseteq \Gamma$ be a subgroup. A Γ-invariant Riemannian metric $\tilde{g}$ on U is called *straightened along* NU^G if within a G-invariant neighborhood of $U^G \subset U$ the push-forward metric $\tilde{g}^G$ coincides with $\tilde{g}$. A Γ-invariant metric $\tilde{g}$ is called *straightened* if it is straightened along NU^G for all subgroups $G \subseteq \Gamma$.

Lemma 5.3.4. Suppose $\tilde{g}$ is straightened along NU^G for some subgroup $G \subseteq \Gamma$. Then there is a G-invariant neighborhood $V \subset U$ of U^G such that for every subgroup $H \subseteq G$, the restriction of $\tilde{g}$ to V is straightened along NV^H.

Proof. It is a direct consequence of Corollary 5.3.6. □

Definition 5.3.8. A Riemannian metric on an effective orbifold X is called *straightened* resp. *straightened near a subset* $Y \subset X$ if its pullback to each chart is straightened resp. if its restriction to an open neighborhood of Y is straightened.

It is easy to see that for a Riemannian metric on an orbifold, being straightened is an intrinsic property independent of the choice of an orbifold chart, therefore the above definition is well-defined. Before we show that straightened metrics exist, we consider the following good property of them.

Definition 5.3.9. Suppose W is a finite-dimensional real representation of a finite group G. Then W can be decomposed as the direct sum of irreducible representations. We call the decomposition

$$W = \mathring{W}^G \oplus \check{W}^G$$

where $\mathring{W}^G$ is the direct sum of all trivial summands and $\check{W}^G$ is the direct sum of all nontrivial summands the *basic decomposition* of W with respect to G.

Lemma 5.3.5. Suppose X is equipped with a straightened metric. Then for each orbifold chart $C = (U, \Gamma, \psi)$ with the pullback metric g_C, for each pair of subgroups $H < G < \Gamma$, for each $x \in U^G$ and $v \in N_\epsilon U^G \cap N_x U^G$ for ϵ sufficiently small, the following holds. Using the (orthogonal) basic decomposition with respect to H

$$NU^G \cong (TU^H \cap NU^G) \oplus NU^H|_{U^G}$$

we write

$$v = (\mathring{v}_H, \check{v}_H).$$

Moreover, we use the exponential map $\exp : N_\epsilon U^G \to U$ to identify $\check{v}_H \in T_x U$ with a vector of $NU^H|_{\exp_x \mathring{v}_H}$ and still denote it by $\check{v}_H$. Then one has

$$\exp_x v = \exp_{\exp_x \mathring{v}_H}(\check{v}_H).$$

Proof. As the metric is straightened, locally U is isometric to the total space NU^G. Then the property follows from Corollary 5.3.6. □

Now we prove the existence of straightened metrics by a prototypical induction argument.

Lemma 5.3.6. Let X be an effective orbifold and $Y \subset X$ be a compact subset. Then there exists a Riemannian metric g on X which is straightened near Y.

Proof. We define a filtration

$$X_0 \subset X_1 \subset \cdots \subset X_n = X$$

where $n = \dim X$ as follows: $p \in X_d$ if for any orbifold chart (U, Γ, ψ) and any $x \in U$ with $\psi(x) = p$, the dimension of U^{Γ_x} through x is at most d, where $\Gamma_x \subset \Gamma$ is the isotropy group of x. Then X_d is closed. Let g be an arbitrary Riemannian metric on X. We modify g inductively such that the modification g_d is straightened near $X_d \cap Y$. For $d = 0$, X_0 is discrete. Then for each $p \in X_0 \cap Y$, define a Riemannian metric g_p as follows. Choose a linear orbifold chart $C = (U, \Gamma, \psi)$ centered at p. Let $\tilde{g}$ be the metric on U obtained by pulling back g to the chart C. Define a metric $\tilde{g}_p$ on U whose restriction to a neighborhood of $0 \in U$ is equal to the pushforward of the Euclidean metric on $T_0U \cong \mathbb{R}^n$ (induced from $\tilde{g}$) via the exponential map on U centered at 0. Then $\tilde{g}_p$ induces a Riemannian metric in a neighborhood of $p \in X$. As $\tilde{g}_p$ is Euclidean near 0, it is straightened near p. Using cut-off functions one can obtain a metric g_0 which is straightened near $X_0 \cap Y$.

Suppose we have obtained a metric g_{d-1} straightened over an open neighborhood $V_{d-1} \supset X_{d-1} \cap Y$ for some $d \geq 1$. Then by the compactness of Y, one can find finitely many linear orbifold charts $C_i = (U_i, \Gamma_i, \psi_i)$ for $i = 1, 2, \ldots, N$ with U_i being a radius $2r_i$-ball of $\mathbb{R}^n$ such that

1. $x_i = \psi_i(0) \in (X_d \setminus X_{d-1}) \cap Y$.
2. The union of $F_i' := \psi_i(U_i')$ with U_i' being the radius r_i ball in U_i covers $(Y \cap X_d) \setminus V_{d-1}$.

We claim that one can find a metric $g_{d,k}$ on X which is straightened near $Y \cap X_d \cap (V_{d-1} \cup F_1' \cup \cdots F_k')$. Indeed, suppose we have find $g_{d,k-1}$ for some $k \geq 1$. For the chart C_k, let $\tilde{g}_{d,k-1}$ be the corresponding Γ_k-invariant metric on U_k. Then we can use the exponential map centered at the origin of U_k to pushforward the induced

bundle metric on the ϵ-disk bundle $N_\epsilon U_k^{\Gamma_k}$ (associated to $\tilde{g}_{d,k-1}$) to a neighborhood of $U_k^{\Gamma_k}$, and use a Γ_k-invariant cut-off function to extend it to U_k. This provides a Γ_k-invariant metric $\tilde{g}_{C_k}$ on the manifold U_k. Then $\tilde{g}_{C_k}$ is straightened near $U_k^{\Gamma_k}$. Let g_{C_k} be the corresponding orbifold metric on the open set $F_k = \psi_k(U_k) \subset X$. Now choose a cut-off function $\rho_k : X \to [0,1]$ supported in F_k such that $\rho_k \equiv 1$ near $\overline{F_k'}$ and ρ_k vanishes on $X \setminus F_k$. Then define

$$g_{d,k} = (1 - \rho_k) g_{d,k-1} + \rho_k g_{C_k}$$

which is a new Riemannian metric on X. We only need to show that this metric is straightened near any point $p \in Y \cap X_d \cap (V_{d-1} \cup F_1' \cup \cdots \cup F_k')$. Indeed, if $p \in F_k'$, then locally $g_{d,k} = g_{C_k}$ which is straightened near p. If $p \notin F_k$, then locally $g_{d,k} = g_{d,k-1}$ which is also straightened near p. If $p \in F_k \setminus F_k'$, then $p \in Y \cap X_d \cap (V_{d-1} \cup F_1' \cup \cdots \cup F_{k-1}')$. Then near p the metric $g_{d,k-1}$ is already straightened. Let $x \in U_k$ be a point such that $\psi_k(x) = p$. Then $x \in U_k^{\Gamma_k}$. Then near x the metrics $\tilde{g}_{d,k-1}$ and $\tilde{g}_{C_k}$ are identical. Hence $g_{d,k} = g_{d,k-1}$ near x so $g_{d,k}$ is straightened near x. Then inductively, one can obtain a metric g_d which is straightened near $Y \cap X_d$. The induction on d provides a metric with desired property. □

From the construction used in the proof one can see the following: for a closed subset $Q \subset X$, if we start with a Riemannian metric g on X which is already straightened near $Y \cap Q$, then one can obtain a metric g' which is straightened near Y such that it coincides with g in a neighborhood of $Y \cap Q$. Therefore one can obtain a "relative version" of Lemma 5.3.6. In particular, one can connect any two straightened metric via a one-parameter family.

Lemma 5.3.7. For any pair of Riemannian metrics g_0 and g_1 on X which are straightened near Y, there exists a straightened Riemannian metric on $X \times [0,1]$ which coincides with $g_i + dt^{\otimes 2}$ near $X \times \{i\}$ for $i = 0, 1$.

Proof. Define $\tilde{X} = X \times [0,1]$ and $\tilde{Y} = Y \times [0,1]$. Let t be the coordinate on $[0,1]$. Consider an arbitrary Riemannian metric $\tilde{g}'$ on $\tilde{X}$ which coincides with $g_0 + dt^{\otimes 2}$ on $X \times [0, \frac{1}{3}]$ and coincides with $g_1 + dt^{\otimes 2}$ on $X \times [\frac{2}{3}, 1]$. Then g is straightened near $Y \times (X \times ([0, \frac{1}{4}] \cup [\frac{3}{4}, 1]))$ Then the inductive construction of Lemma 5.3.6 produces a straightened metric $\tilde{g}$ which coincides with $\tilde{g}$ near $Y \times (X \times ([0, \frac{1}{4}] \cup [\frac{3}{4}, 1]))$. □

5.3.4 Connections

We would also like the orbifold vector bundle to be "straightened" in the direction normal to fixed point loci in a way analogous to Riemannian metrics. We first look at connections on a chart. Let U be a smooth manifold acted on effectively by Γ, and $E \to U$ is a Γ-equivariant vector bundle. Let ∇^E be a Γ-invariant connection on E. Suppose also U is equipped with a straightened metric. Then the connection ∇^E together with the Riemannian metric induce the following objects. For each subgroup $H < \Gamma$, there is a neighborhood of U^H identified with a disk bundle $N_r U^H \subset NU^H \to U^H$. Let $\pi_H : NU^H \to U^H$ be the projection. Then the parallel transport of E along normal geodesics induces an H-equivariant bundle isomorphism

$$\rho^H : \pi_H^*(E|_{U^H}) \cong E|_{N_r U^H}.$$

Therefore the H-equivariant splitting $E|_{U^H} \cong \mathring{E}^H \oplus \check{E}^H$ extends to an H-equivariant splitting

$$E|_{N_r U^H} \cong \mathring{E}^H \oplus \check{E}^H. \tag{5.3.3}$$

Definition 5.3.10. (Straightened connections) Let U be a smooth manifold acted on effectively by a finite group Γ. Let $E \to U$ be a Γ-equivariant vector bundle. Suppose U is equipped with a Γ-invariant straightened Riemannian metric. For a subgroup $H < \Gamma$, a Γ-equivariant connection ∇^E on E is said to be *straightened along* NU^H if

in a neighborhood $N_r U^H$ of U^H there holds

$$(\rho^H)^*(\pi^H)^*\nabla^{E|_{U^H}} = \nabla^E.$$

The following statement guarantees that splittings of the form (5.3.3) induced by straightened connections behave well with respect to group injections $H \hookrightarrow G$.

Lemma 5.3.8. Suppose ∇^E is straightened along NU^G. Then there exists a G-invariant neighborhood $V \subset U$ of U^G such that for each subgroup $H \subseteq G$, the restriction of ∇^E to V is straightened along NV^H.

Proof. Basically, this lemma follows from the functoriality of pullback connections. Indeed, as a pullback connection has no curvature in each fiber, one has

$$(\pi_H^*(\rho_G|_{V^H})) \circ \rho_H = \rho_G : E \to \pi_G^* E|_{U^G}$$

near U^G. Therefore,

$$\begin{aligned}\nabla^E = \rho_G^*(\nabla^{E|_{U^G}}) = \rho_H^*(\pi_H^*(\rho_G|_{V^H}))^*(\nabla^{E|_{U^G}}) \\ = \rho_H^*\big(\pi_H^*(\rho_G^*(\nabla^{E|_{U^G}}))|_{V^H}\big) = \rho_H^*(\pi_H^*\nabla^{E|_{V^H}}).\end{aligned}$$

Hence ∇^E is also straightened along NV^H. □

We define the notion of straightenedness for connections on effective orbifolds.

Definition 5.3.11. Let X be an effective orbifold equipped with a straightened Riemannian metric. Let $\mathcal{E} \to X$ be an orbifold vector bundle. Then a connection $\nabla^{\mathcal{E}}$ on $\mathcal{E}$ is called *straightened* if for every bundle chart $\hat{C} = (U, E, \Gamma, \hat{\psi})$ and every subgroup $G < \Gamma$, the Γ-equivariant connection ∇^E on $E \to U$ obtained by pulling back $\nabla^{\mathcal{E}}$ is straightened along NU^G.

Lemma 5.3.9. Let X be an effective orbifold with or without boundary equipped with a straightened Riemannian metric. Let $\mathcal{E} \to X$ be a vector bundle. Let $Y \subset X$ be a compact subset and $Q \subset X$ be a closed set. Suppose $\nabla^{\mathcal{E}}$ is a connection on $\mathcal{E}$ which is straightened near $Y \cap Q$. Then there exists a connection on $\mathcal{E}$ which is straightened near Y and which coincides with $\nabla^{\mathcal{E}}$ near $Y \cap Q$.

Proof. The proof is similar to that of Lemma 5.3.6. □

Definition 5.3.12. Let X be an effective orbifold and $\mathcal{E} \to X$ be an orbifold vector bundle. A *straightened structure* on the pair $(X, \mathcal{E})$ consists of a straightened Riemannian metric on X and a straightened connection on $\mathcal{E}$ with respect to the straightened Riemannian metric. If $(X, \mathcal{E})$ is equipped with a straightened structures, then we say that $(X, \mathcal{E})$ is *straightened*.

5.3.5 Normal complex structures

Now we introduce the most important geometric condition which plays the central role in our construction. In applications, normal complex structures appear naturally as the Cauchy–Riemann operator has a complex linear principal symbol.

Definition 5.3.13. Let X be an effective orbifold. A *normal complex structure* on X is a collection of H-equivariant complex structures

$$I_H : NU^H \to NU^H$$

on the normal bundles $NU^H \to U^H$ for each orbifold chart (U, Γ, ψ) and each subgroup $H \subset \Gamma$. Moreover, the collection must satisfy the following compatibility conditions.

1. For each chart (U, Γ, ψ), each $y \in U$ with stabilizer $\Gamma_y < \Gamma$, and each pair of

subgroups $H < G \subset \Gamma_y$, notice that we have the H-equivariant decomposition

$$N_y U^G \cong (T_y U^H \cap N_y U^G) \oplus N_y U^H$$

into trivial and nontrivial summands of H-representations. This decomposition is also I_G-invariant. Then we require that the restriction of I_G to the second summand $N_y U^H$ of the above decomposition agrees with I_H at y.

2. For each chart embedding from (V, Π, ϕ) to (U, Γ, ψ) given by an injection $\Pi \hookrightarrow \Gamma$ and an equivariant open embedding $\iota : V \hookrightarrow U$, for any $z \in V$ sent to $y = \iota(z)$, with $\Gamma_y \cong \Pi_z$, one has an equivariant isomorphism $T_z V \cong T_y U$ sending $N_z V^{\Pi_z}$ isomorphically to $N_y U^{\Gamma_y}$, we require that this isomorphism is complex linear.

Similarly we can define normal complex structures on bundles. Let X be an effective orbifold and $\mathcal{E} \to X$ be a vector bundle. For each bundle chart $\hat{C} = (U, E, \Gamma, \psi)$ and each subgroup $G < \Gamma$, one has the basic decomposition

$$E|_{U^G} \cong \mathring{E}^G \oplus \check{E}^G$$

with respect to G.

Definition 5.3.14. A normal complex structure on $\mathcal{E}$ is a collection of G-equivariant complex structures

$$J_G : \check{E}^G \to \check{E}^G$$

on the bundle $\check{E}^G \to U^G$ for each bundle chart $\hat{C} = (U, E, \Gamma, \hat{\psi})$ and each subgroup $G < \Gamma$. Moreover, the collection must satisfy the following compatibility conditions.

1. For each bundle chart $\hat{C} = (U, E, \Gamma, \hat{\psi})$, each $y \in U$ with stabilizer $\Gamma_y < \Gamma$, and each pair of subgroups $H \subseteq G < \Gamma_y$, notice that we have the H-equivariant

decomposition

$$\check{E}^G|_y \cong (\check{E}^G_y)^H \oplus \check{E}^H_y.$$

Notice that the decomposition is J_G-invariant. We require that the restriction of J_G to the summand $\check{E}^H_y$ coincides with J_H.

2. For each chart embedding from $(V, F, \Pi, \hat{\phi})$ to $(U, E, \Gamma, \hat{\psi})$ given by an injection $\Pi \hookrightarrow \Gamma$ and an equivariant bundle embedding $\hat{\iota} : F \hookrightarrow E$ covering a base embedding $\iota : V \hookrightarrow U$, for any $z \in V$ (with stabilizer G_z) sent to $y = \iota(z)$ (with stabilizer G_y), one has an equivariant isomorphism $F|_z \cong E|_y$. We require that the induced isomorphism between $\check{F}^{G_z}|_z \cong \check{E}^{G_y}|_y$ is complex linear.

Remark 5.3.15. If (X, J) is an almost complex orbifold, J also endows X with a normal complex structure. The notion of normal complex structure is convenient for the discussion of cobordisms between (derived) almost complex orbifolds, and it also plays an important role in the discussion of certain invariance result.

5.3.6 Normally complex sections

Now we discuss the notion of normally complex sections originally introduced by Parker [Par13] which generalizes Fukaya–Ono's notion of normally polynomial sections [FO97]. We first discuss the notion of lifts within a single orbifold chart. Suppose G is a finite group. Let V, W be finite-dimensional complex representations of G. Let $\mathrm{Poly}^G(V, W)$ be the space of G-equivariant polynomial maps and for each $d \geq 0$, define

$$\mathrm{Poly}^G_d(V, W) \subset \mathrm{Poly}^G(V, W)$$

to be the subspace of polynomial maps of degrees at most d. Then there is an evaluation map

$$\mathrm{ev} : V \times \mathrm{Poly}^G(V, W) \to W. \tag{5.3.4}$$

These notions can also be defined for the parametrized case. Let M be a smooth manifold, $\mathcal{V}, \mathcal{W} \to M$ be smooth complex vector bundles equipped with fiberwise complex linear G-actions. Then one has the bundle of fiberwise polynomial maps

$$\mathrm{Poly}_d^G(\mathcal{V}, \mathcal{W}) \to M$$

and the corresponding evaluation map

$$\mathrm{ev} : \mathcal{V} \oplus \mathrm{Poly}_d^G(\mathcal{V}, \mathcal{W}) \to \mathcal{W}. \tag{5.3.5}$$

Now let X be an effective orbifold and $\mathcal{E} \to X$ be a vector bundle. Suppose $(X, \mathcal{E})$ is normally complex and is straightened. Then for each bundle chart $\hat{C} = (U, G, E, \hat{\psi})$ of $\mathcal{E}$ and any subgroup $H < G$, using the straightened structures, we can identify a neighborhood of U^H with a disk bundle $N_r U^H$ and we can extend the basic decomposition

$$E|_{U^H} \cong \mathring{E}^H \oplus \check{E}^H$$

to a decomposition of E near U^H. From now on such identifications and decompositions will be assumed implicitly.

Definition 5.3.16. Let $\hat{C} = (U, G, E, \hat{\psi})$ be a bundle chart of $\mathcal{E} \to X$ and let $s : U \to E$ be a smooth Γ-equivariant section. A *local normally complex lift* of s of degree at most d near x is a G_x-equivariant smooth bundle map (over U^{G_x})

$$\check{\mathfrak{s}}_x : N_\epsilon U^{G_x} \to \mathrm{Poly}_d^{G_x}(NU^{G_x}, \check{E}^{G_x}) \tag{5.3.6}$$

satisfying the following condition. We define the graph of $\check{\mathfrak{s}}_x$ to be the bundle map

$$\mathrm{graph}(\check{\mathfrak{s}}_x) : N_\epsilon U^{G_x} \to NU^{G_x} \oplus \mathrm{Poly}_d^{G_x}(NU^{G_x}, \check{E}^{G_x}),\ v \mapsto (v, \check{\mathfrak{s}}_x(v)).$$

Using the basic decomposition of E near U^{G_x}, we can write s as

$$s = (\mathring{s}_x, \check{s}_x) : N_\epsilon U^{G_x} \to \mathring{E}^{G_x} \oplus \check{E}^{G_x} \cong E|_{N_\epsilon U^G},$$

then near x one has

$$\check{s}_x = \mathrm{ev} \circ \mathrm{graph}(\check{\mathfrak{s}}_x).$$

By abuse of notions, we also call the bundle map

$$\tilde{s}_x := (\mathring{s}_x, \check{\mathfrak{s}}_x) : N_\epsilon U^{G_x} \to \mathring{E}^{G_x} \oplus \mathrm{Poly}_d^{G_x}(NU^{G_x}, \check{E}^{G_x})$$

or the bundle map

$$\tilde{S}_x := (\mathring{s}_x, \mathrm{graph}(\check{\mathfrak{s}}_x)) : N_\epsilon U^{G_x} \to \mathring{E}^{G_x} \oplus \left(NU^{G_x} \oplus \mathrm{Poly}_d^{G_x}(NU^{G_x}, \check{E}^{G_x})\right)$$

a local normally complex lift of s near x.

If $s : U \to E$ has a local normally complex lift (of degree at most d) near each point of U, then we say that s is a *normally complex section* or an *FOP section*[1] (of degree at most d).

Remark 5.3.17. In equation (5.3.6), we use the exponential map implicitly to identify $N_\epsilon U^{G_x}$ with a neighborhood of the zero section in NU^{G_x}. In other words, a local normally complex lift *depends* on the straightened structure and the normal complex structure. On the other hand, one can see that being an FOP section is a condition invariant under chart embeddings. Hence it is a condition intrinsic for the orbifold bundle, the normal complex structure, and the straightened structures.

Definition 5.3.18. A smooth section $s : X \to \mathcal{E}$ is called a *FOP section* or a *normally complex section* (of degree at most d) if for each $p \in X$ there is a bundle

[1]Stands for Fukaya–Ono–Parker.

chart $\hat{C} = (U, \Gamma, E, \hat{\psi})$ of $\mathcal{E} \to X$ around p such that the pullback of s to U is a normally complex section (of degree at most d) in the sense of Definition 5.3.16.

Remark 5.3.19. There is a more restricted notion called *normally polynomial sections* considered by Fukaya–Ono [FO97]. A section is normally polynomial if in each chart $\hat{C} = (U, \Gamma, E, \hat{\psi})$, for each group $G < \Gamma$, near $U^G \subset U$ the section $\check{s} : N_\epsilon U^G \to \check{E}^G$ is of the form

$$\check{s} : U^G \to \mathrm{Poly}_d^G(NU^G, \check{E}^G).$$

In particular, there is no ambiguity of choosing lifts. In fact it is enough to use normally complex sections to achieve the required transversality (see the proof of Proposition 5.6.3). However, the more flexible notion of normally complex sections, introduced by Parker [Par13], is very convenient to work with. For example one can do cut and paste as shown below. More importantly, this flexible notion allows us to prove that our notion of transversality is well-behaved even for normally polynomial sections (see Section 5.4).

Lemma 5.3.10. The space of normally complex sections of $E \to X$ is a module over $C^\infty(X)$.

Proof. Left to the reader. □

Remark 5.3.20. In general the set of FOP sections is strictly contained in the set of smooth sections. Consider $\mathbb{C}$ acted nontrivially by $\mathbb{Z}_2$. The space of equivariant polynomial maps from $\mathbb{C}$ to itself is generated by monomials $z, z^3, z^5, \cdots$. Hence a FOP map from $\mathbb{C}$ to $\mathbb{C}$ is of the form

$$z \mapsto f_1(z)z + f_3(z)z^3 + \cdots + f_{2k-1}z^{2k-1}$$

where $f_i(z)$ is an even smooth function. We can also write such a map as $z \mapsto f(z)z$ with f an even function. However, not all equivariant smooth maps are of this form,

such as the map $z \mapsto \overline{z}$.

Now we show that smooth sections can be approximated by FOP sections. To measure the distance between sections, choose a metric on $\mathcal{E}$. We remind the reader that X is assumed to be effective throughout this subsection.

Lemma 5.3.11. Let $s : X \to \mathcal{E}$ be a smooth section with $s^{-1}(0)$ compact. Let $D \subset X$ be a precompact open neighborhood of $s^{-1}(0)$. Then there exists $d \geq 0$ such that for any $\epsilon > 0$, there exists a smooth section $s' : X \to \mathcal{E}$ such that

1. s' is an FOP section near $\overline{D}$ of degree at most d.
2. There holds the estimate
$$\|s - s'\|_{C^0(D)} < \epsilon. \tag{5.3.7}$$

Proof. By the compactness of $\overline{D}$, one can find a finite collection of bundle charts $\hat{C}_i^+ = (U_i^+, E_i, \Gamma_i, \hat{\psi}_i)$ $(i = 1, \ldots, N)$ centered at $x_i \in \overline{D}$ satisfying the following conditions.

1. The bundle chart $\hat{C}_i^+$ is linear and $U_i^+ \subset \mathbb{R}^n$ is a radius $2r_i$-ball centered at the origin.
2. The collection $F_i := \psi_i(U_i)$, $i = 1, \ldots, N$ cover $\overline{D}$ where $U_i \subset U_i^+$ is the radius r_i-ball centered at the origin.

Define $F_\infty := X \setminus \overline{D}$. Then choose a smooth partition of unity subordinate to the open cover $\{F_1, \ldots, F_N, F_\infty\}$, denoted by $\rho_1, \ldots, \rho_\infty$.

Let $s_i : U_i^+ \to E_i$ be the pullback of s to the linear chart $\hat{C}_i^+$. We first show that s_i can be approximated by normally complex sections. Indeed we can write

$$s_i = (\mathring{s}_i, \check{s}_i) : U_i^+ \to \mathring{E}_i^{\Gamma_i} \oplus \check{E}_i^{\Gamma_i}.$$

We would like to approximate the smooth map $\check{s}_i$ by a smooth map $\check{\mathfrak{s}}_i : U_i^+ \to \mathrm{Poly}_d^{\Gamma_i}(NU_i^{\Gamma_i}, \check{E}_i^{\Gamma_i})$ near $\overline{U_i}$. Indeed, for each $x \in \overline{U_i}$, suppose its stabilizer is $G_x \subset$

Γ_i. Then $\check{s}_i(x) \in (\check{E}_i^{\Gamma_i})^{G_x}$. Then for a sufficiently large d (which only depends on the group Γ_i) there exists a degree at most d equivariant polynomial map $P_x \in \mathrm{Poly}_d^{\Gamma_i}(NU_i^{\Gamma_i}, \check{E}_i^{\Gamma_i})$ such that $P_x(x) = \check{s}_i(x)$ (see Lemma 5.4.3). Then one can find $r_x > 0$ such that

$$\sup_{x' \in B_{r_x}(x)} |\check{s}_i(x) - \check{s}_i(x')| + \sup_{x' \in B_{r_x}(x)} |P_x(x) - P_x(x')| < \epsilon.$$

As $\overline{U_i}$ is compact, one can find finitely many points $x_j \in \overline{U_i}$ such that

$$\overline{U_i} \subset \bigcup_{j=1}^{l} B_{r_{x_j}}(x_j).$$

Choose a partition of unity subordinate to $B_{r_{x_j}}(x_j)$ and $U_i^+ \setminus \overline{U_i}$. Define

$$\check{\mathfrak{s}}_i' : U_i^+ \to \mathrm{Poly}_d^{\Gamma_i}(NU_i^{\Gamma_i}, \check{E}_i^{\Gamma_i}),\ \check{\mathfrak{s}}_i(x) = \sum_{j=1}^{l} \rho_j(x) P_{x_j}.$$

We first check that this (not necessarily Γ_i-equivariant map) is close to the original section $\check{s}_i$. Indeed, for each $x \in \overline{U_i}$, then one has

$$|\check{s}_i(x) - \mathrm{ev}(x, \check{\mathfrak{s}}_i'(x))| \leq \sum_{j=1}^{l} \rho_j(x) |\check{s}_i(x) - P_{x_j}(x)|.$$

If $x \notin B_{r_{x_j}}(x)$, then $\rho_j(x) = 0$ and the corresponding summand above vanishes; if $x \in B_{r_{x_j}}(x)$, then

$$|\check{s}_i(x) - P_{x_j}(x)| \leq |\check{s}_i(x) - \check{s}_i(x_j)| + |P_{x_j}(x) - P_{x_j}(x_j)| < \epsilon.$$

Hence one has

$$\sup_{x \in \overline{U_i}} |\check{s}_i(x) - \mathrm{ev}(x, \check{\mathfrak{s}}_i'(x))| < \epsilon.$$

Now we make $\check{s}_i$ invariant by setting

$$\check{\mathfrak{s}}_i(x) = \frac{1}{|\Gamma_i|} \sum_{g \in \Gamma_i} \check{\mathfrak{s}}'_i(gx).$$

Then for each $x \in \overline{U_i}$, one has

$$\begin{aligned} |\check{s}_i(x) - \mathrm{ev}(x, \check{\mathfrak{s}}_i(x))| &= \frac{1}{|\Gamma_i|} \left| \sum_{g \in \Gamma_i} g^{-1} \check{s}_i(gx) - \sum_{g \in \Gamma_i} \mathrm{ev}(x, \check{\mathfrak{s}}'_i(gx)) \right| \\ &\leq \frac{1}{|\Gamma_i|} \sum_{g \in \Gamma_i} \left| g^{-1} \check{s}_i(gx) - g^{-1} \mathrm{ev}(gx, \check{\mathfrak{s}}'_i(gx)) \right| < \epsilon. \end{aligned}$$

Then $\check{\mathfrak{s}}_i$ together with $\mathring{s}_i$ defines an FOP section of $E_i \to U_i^+$ and hence an FOP section s'_i of the orbifold bundle $\mathcal{E} \to X$ over $\psi_i(U_i^+)$. Then define

$$s' := \sum_{i=1}^{N} \rho_i s'_i + \rho_\infty s.$$

By Lemma 5.3.10, this is a smooth section of $\mathcal{E}$ which is an FOP section in a neighborhood of $\overline{D}$. The chartwise estimates imply that (5.3.7) holds. □

5.4 Whitney stratifications on the variety Z

The purpose of this section is to specify the notion of transversality for normally complex sections via the language of Whitney stratifications. The fundamental idea is from Fukaya–Ono [FO97]. On the other hand, it is rather technical to prove that this transversality notion is intrinsic, i.e., behaves well with respect to chart embeddings. It is the work of Parker [Par13] from which we learned how to use normally complex sections to compare chartwise transversality notions and how to prove another similar property, i.e. the transversality notion is independent of the choice of the cut-off degree of polynomial maps.

The notion of transversality for FOP sections is based on the intricate study of a particular class of complex affine varieties which we generally refer to as "the variety Z." In this section, let G be a finite group and V, W be two finite-dimensional complex representations of G. We require that the G-action on V is effective. We allow V or W to have trivial G-summands, hence the triple (G, V, W) does not represent an isotropy type in general (see Definition 5.3.4). Define

$$Z^G(V,W) = \{(v,P) \in \mathrm{Poly}^G(V,W) \mid P(v) = 0 \in W\}$$

the zero locus of the evaluation map from (5.3.4). Its cut-off at any degree d is

$$Z^G_d(V,W) := Z^G(V,W) \cap (V \times \mathrm{Poly}^G_d(V,W)).$$

Similarly, one can define the family of the Z-variety for the parametrized case. Given a smooth manifold M and $\mathcal{V}, \mathcal{W} \to M$ smooth complex vector bundles with fiberwise complex linear G-actions, the zero locus of (5.3.5) is denoted by

$$\mathcal{Z}^G(\mathcal{V},\mathcal{W}) := \{(v,P) \in \mathcal{V} \oplus \mathrm{Poly}^G(\mathcal{V},\mathcal{W}) \mid P(v) = 0\}$$

and $\mathcal{Z}^G_d(\mathcal{V},\mathcal{W})$ is defined similarly by considering fiberwise polynomial maps of degree at most d.

5.4.1 The canonical Whitney stratification

We first recall basic notions about Whitney stratifications. More detailed discussions can be found in Appendix 5.8.

Definition 5.4.1. Let M be a smooth manifold and $S \subset M$ be a subset.

1. A *prestratification* of S is a decomposition of S

$$S = \bigsqcup S_\lambda$$

into nonempty locally closed sets such that 1) the decomposition is locally finite and 2) if $S_\lambda \cap \overline{S_\mu} \neq \emptyset$, then $S_\lambda \subset \overline{S_\mu}$. Each member S_λ of this decomposition is called a *strata* of this prestratification.

2. A prestratification is called a *Whitney prestratification* if all strata are smooth submanifolds and each distinct pair (S_λ, S_μ) of strata satisfy Whitney's condition (b) (Definition 5.8.3).

3. A *stratification* of S is a rule $\mathfrak{S}$ which assigns to each $x \in S$ a set-germ S_x containing x such that for each $x \in S$, there is an open neighborhood N and a prestratification of $N \cap S$ such that for all $y \in S \cap N$, the set-germ S_y is the germ of the strata in the prestratification which contains y. In particular, any prestratification of S induces a stratification.

4. A stratification of S is called a *Whitney stratification* if it is induced from a Whitney prestratification.

5. Let N be a smooth manifold and $f : N \to M$ be a smooth map. f is said to be *lemma412* to S with respect to a Whitney stratification $\mathfrak{S}$ if f is transverse to each set-germ of $\mathfrak{S}$.

6. If M is a nonsingular complex algebraic variety and $S \subset M$ is a constructible set, a Whitney prestratification of S is called *complex algebraic* if all of its strata are nonsingular complex algebraic subsets of M. A Whitney stratification is called complex algebraic if it is induced from a complex algebraic Whitney prestratification.

Now consider the variety Z associated to the triple (G, V, W) and a nonnegative integer d. For each $v \in V$, let $G_v \subset G$ be the isotropy subgroup (or stabilizer) of v. The total space V has a prestratification indexed by subgroups of G. More precisely, for each subgroup $H \subset G$, denote

$$V_H := \mathring{V}^H := \{v \in V \mid H \subset G_v\}, \qquad V_H^* := \{v \in V \mid G_v = H\}.$$

The top stratum is also called the isotropy-free part of V, denoted by

$$V^{\text{free}} := \{v \in V \mid G_v = \{e\}\}.$$

Then the decomposition

$$V = \bigsqcup_{H \subseteq G} V_H^*$$

is a Whitney prestratification of V whose strata are all regular complex algebraic sets. We will call it the *action prestratification* on V. The induced prestratification on $V \times \text{Poly}_d^G(V, W)$, where G acts trivially on the second factor, is also called the *action prestratification*.

Definition 5.4.2. A smooth Whitney stratification on $Z := Z_d^G(V, W)$ is said to *respect* the action prestratification if for each $(v, P) \in Z_d^G(V, W)$, if $v \in V_H^*$ for some subgroup $H \subseteq G$, then the set-germ Z_x is contained in $V_H^* \times \text{Poly}_d^G(V, W)$.

This definition is a special case of Definition 5.8.16.

Theorem 5.4.3. For each $d \geq 0$, there exists a unique minimal Whitney stratification $\mathfrak{Z}$ on $Z_d^G(V, W)$ which respects the action prestratification. Moreover, it has the following additional properties.

1. It is complex algebraic, i.e., it is induced from a Whitney prestratification whose strata are all nonsingular complex algebraic sets.

2. It is G-invariant.

3. If a self-diffeomorphism f of $V \times \mathrm{Poly}_d^G(V,W)$ preserves the action prestratification and preserves the set $Z_d^G(V,W)$, then f also preserves this Whitney stratification.

Proof. The construction of such a Whitney stratification, denoted by $\mathfrak{Z}$, is a special case of Theorem 5.8.17. From that theorem one also knows that $\mathfrak{Z}$ is algebraic and is minimal among all smooth ones which respect the action prestratification. The minimality implies the uniqueness (see Definition 5.8.5).

To prove that $\mathfrak{Z}$ is G-invariant, pick any $g \in G$. Then g is a particular diffeomorphism of the vector space $V \times \mathrm{Poly}_d^G(V,W)$ which satisfies assumptions of the map f in Proposition 5.8.19. Hence $g^*\mathfrak{Z} = \mathfrak{Z}$. Similarly, a self-diffeomorphism f of $V \times \mathrm{Poly}_d^G(V,W)$ which preserves the action prestratification and preserves the set $Z_d^G(V,W)$ also satisfies assumptions of Proposition 5.8.19. Hence the last condition is also true. □

Definition 5.4.4. We call the Whitney stratification $\mathfrak{Z}$ of Theorem 5.4.3 the *canonical Whitney stratification* on the variety $Z_d^G(V,W)$.

Remark 5.4.5. To talk about transversality, one does not need to specify a Whitney prestratification. However, when we discuss the pseudocycle property in Section 5.6, it is convenient to have a distinguished set of strata. Indeed, here is a canonically associated prestratification which induces the canonical Whitney stratification (cf. Definition 5.8.14). The *dimension filtration* of $Z_d^G(V,W)$ associated to the canonical Whitney stratification is

$$Z_d^G(V,W)_k := \{x \in Z_d^G(V,W) \mid \dim_{\mathbb{R}} Z_x \le k\},\ k = 0,1,\dots.$$

Here Z_x is the set-germ at x given by the Whitney stratification. Then each $Z_d^G(V,W)_k \setminus Z_d^G(V,W)_{k-1}$ is a nonsingular complex algebraic set of real dimension k. Each of its

connected components (in the Euclidean topology) is also a nonsingular algebraic set (a connected component is also an irreducible component in Zariski topology). Hence the collection of connected components of $Z_d^G(V,W)_k \setminus Z_d^G(V,W)_{k-1}$ for all k is a complex algebraic Whitney prestratification which induces the canonical Whitney stratification. From now on, a *stratum* of $Z_d^G(V,W)$ means a stratum of this canonical Whitney prestratification.

One can also define the notion of the canonical Whitney stratification of the parametrized case. Let M be a base manifold acted on trivially by G and let $\mathcal{V}, \mathcal{W} \to M$ be complex G-vector bundles with fibers isomorphic to V and W respectively. Then the structure group of $\mathcal{V} \oplus \mathrm{Poly}_d^G(\mathcal{V}, \mathcal{W})$ is canonically reduced to $GL(V)^G \times GL(W)^G$, whose elements are diffeomorphisms of $V \times \mathrm{Poly}_d^G(V,W)$ preserving both the action prestratification and the set $Z_d^G(V,W)$. Hence by the last property of Theorem 5.4.3, the canonical Whitney stratification on $Z_d^G(V,W)$ could be patched into a Whitney stratification of $\mathcal{Z}_d^G(\mathcal{V}, \mathcal{W})$ which is "locally trivial". We also call this stratification on $\mathcal{Z}_d^G(\mathcal{V}, \mathcal{W})$ the canonical Whitney stratification.

Remark 5.4.6. The canonical Whitney stratification is an example of Parker's *nice Whitney stratification* ([Par13, Definition 4.6]).

Following Parker, we prove two very useful facts concerning the canonical Whitney stratification.

Lemma 5.4.1. (cf. [Par13, Lemma 4.10]) Let $G, M, \mathcal{V}, \mathcal{W}$ be as above. Fix an integer $d > 0$. Let $\mathcal{V}^+ \subset \mathcal{V}$ be an open subset and let $\underline{f}_1, \underline{f}_2 : \mathcal{V}^+ \to \mathrm{Poly}_d^G(\mathcal{V}, \mathcal{W})$ be two smooth bundle maps such that

$$\mathrm{graph}(\underline{f}_1 - \underline{f}_2) \subset \mathcal{Z}_d^G(\mathcal{V}, \mathcal{W}) := \mathcal{Z}_d^G.$$

Then the graph of $\underline{f}_1$ (viewed as a submanifold of $\mathcal{V}^+ \times_M \mathrm{Poly}_d^G(\mathcal{V}, \mathcal{W})$) is transverse (cf. Definition 5.8.6) to $\mathcal{Z}_d^G$ if and only if the graph of $\underline{f}_2$ is transverse to $\mathcal{Z}_d^G$.

Proof. Suppose the graph of $\underline{f}_1$ is transverse to $\mathcal{Z}_d^G$. For any $(v_0, \underline{f}_2(v_0)) \in \mathcal{Z}_d^G$, we would like to show that the graph of $\underline{f}_2$ is transverse to $\mathcal{Z}_d^G$ at this point $(v_0, \underline{f}_2(v_0))$. Choose a compactly supported cut-off function

$$\rho_0 : \mathcal{V}^+ \to [0, 1]$$

which is identically 1 near v_0. Consider the smooth vertical vector field on $\mathcal{V}^+ \times_M \mathrm{Poly}_d^G(\mathcal{V}, \mathcal{W})$ defined by

$$\xi(v, P) = (0, \rho_0(v)(\underline{f}_1(v) - \underline{f}_2(v))). \tag{5.4.1}$$

Then ξ can be regarded as a smooth vector field on $\mathcal{V} \oplus \mathrm{Poly}_d^G(\mathcal{V}, \mathcal{W})$. Moreover, the flow of ξ is the 1-parameter family of fibre-preserving diffeomorphisms

$$\Phi_t(v, P) = (v, P + t\rho_0(v)(\underline{f}_1(v) - \underline{f}_2(v)))$$

of $\mathcal{V} \oplus \mathrm{Poly}_d^G(\mathcal{V}, \mathcal{W})$ which exists for all time t. It is also easy to see that Φ_t preserves the action prestratification on $\mathcal{V} \oplus \mathrm{Poly}_d^G(\mathcal{V}, \mathcal{W})$ and the set $\mathcal{Z}_d^G$. Hence Φ_t pulls back the canonical Whitney stratification on $\mathcal{Z}_d^G$ to itself. Moreover, Φ_1 maps a neighborhood of $(v_0, \underline{f}_2(v_0))$ in $\mathrm{graph}(\underline{f}_2)$ to a neighborhood of $(v_0, \underline{f}_1(v_0))$ in $\mathrm{graph}(\underline{f}_1)$. Hence the graph of $\underline{f}_2$ is also transverse to $\mathcal{Z}_d^G(\mathcal{E}, \mathcal{F})$ at v_0. □

Lemma 5.4.2. Under the same hypothesis as above. Let $f_1, f_2 : \mathcal{V}^+ \to \mathcal{V} \oplus \mathrm{Poly}_d^G(\mathcal{V}, \mathcal{W})$ be the maps defined by

$$f_i(v) = (v, \underline{f}_i(v)).$$

Suppose both the graphs of $\underline{f}_1$ and $\underline{f}_2$ are transverse to $\mathcal{Z}_d^G$ (with respect to the canonical Whitney stratification). Then the Whitney stratifications on $f_1^{-1}(\mathcal{Z}_d^G) = f_2^{-1}(\mathcal{Z}_d^G)$ pulled back from the canonical Whitney stratifications on $\mathcal{Z}_d^G$ via f_1 and f_2

coincide.

Proof. We just need to check locally around a point $v_0 \in f_1^{-1}(\mathcal{Z}_d^G) = f_2^{-1}(\mathcal{Z}_d^G)$. Let Φ_1 be the time-1 map of the flow of the vector field (5.4.1). Then in a neighborhood of v_0 one has

$$f_1 = \Phi_1 \circ f_2.$$

On the other hand, as Φ_1 preserves the canonical Whitney stratification $\mathfrak{Z}$ on $\mathcal{Z}_d^G$, it follows that near v_0, $f_1^*\mathfrak{Z} = f_2^*\Phi_1^*\mathfrak{Z} = f_2^*\mathfrak{Z}$. □

5.4.2 Regularity of the isotropy-free part

We want the isotropy-free part of the variety Z to be transversely cut-out. This is true if d is sufficiently large as proved by Fukaya–Ono [FO97]. For the convenience of the reader we rewrite their proof here.

Proposition 5.4.7. ([FO97, Lemma 5],[FOOO, Proposition 35.3]) For any finite group G, there exists a positive integer d_0 satisfying the following conditions. Let V, W be finite-dimensional complex representations. Suppose G acts on V faithfully. Then for any $d \geq d_0$ and for each subgroup $H \subset G$, the set

$$Z_{d,H}^* := Z_d^G(V,W) \cap (V_H^* \times \mathrm{Poly}_d^G(V,W))$$

is a nonsingular complex algebraic set of complex dimension

$$\dim_{\mathbb{C}} Z_{d,H}^* = \dim_{\mathbb{C}} \mathrm{Poly}_d^G(V,W) + \dim_{\mathbb{C}} \mathring{V}^H - \dim_{\mathbb{C}} \mathring{W}^H.$$

We first prove a lemma.

Lemma 5.4.3. [FO97, Lemma 5] There exists $d_0 > 0$ such that for all $\mathring{v}_H \in \mathring{V}_H^*$ and $\mathring{w}_H \in \mathring{W}_H$, there exists $P \in \mathrm{Poly}_{d_0}^G(V,W)$ such that $P(\mathring{v}_H) = \mathring{w}_H$.

Proof. By decomposing W into irreducible components, we may assume that W is an irreducible representation of G. Define the G-vector space

$$\boldsymbol{W} := \bigoplus_{\gamma \in G} \mathbb{C}\{\langle \gamma \rangle\}$$

with G-action defined as

$$g\left(\sum_{\gamma} c_\gamma \cdot \langle \gamma \rangle\right) = \sum_{\gamma} c_\gamma \cdot \langle \gamma g^{-1} \rangle = \sum_{\gamma} c_{\gamma g} \cdot \langle \gamma \rangle.$$

Since $\boldsymbol{W}$ is a regular representation, there is a G-equivariant homomorphism $\Psi : \boldsymbol{W} \to W$ and an element $\boldsymbol{w} \in \boldsymbol{W}$ such that

$$\Psi(\boldsymbol{w}) = \Psi\left(\sum_{\gamma} w_\gamma \cdot \langle \gamma \rangle\right) = \mathring{w}_H.$$

Since $\mathring{w}_H \in \mathring{W}_H$, for all $h \in H$, one has

$$\Psi(h\boldsymbol{w}) = h\Psi(\boldsymbol{w}) = \mathring{w}_H.$$

Hence by taking average over H, one may assume that

$$\gamma' H = \gamma'' H \Longrightarrow w_{\gamma'} = w_{\gamma''}.$$

Now we claim that for some $d_0 > 0$ which only depends on G, one can choose a polynomial $f : V \to \mathbb{C}$ (not necessarily G-invariant) of degree at most d_0 such that

$$\forall \gamma \in G,\ f(\gamma \mathring{v}_H) = w_\gamma.$$

Indeed, there are $n := |G/H|$ distinct elements in the G-orbit of $\mathring{v}_H$. One can choose a linear decomposition $V = V_1 \oplus V_2$ such that V_1 is one-dimensional and that the

projection of these n distinct elements are still distinct in V_1. Then by Lagrange's method of interpolation, one can find a complex polynomial $f : V_1 \to \mathbb{C}$ of degree at most $|G|$ taking the prescribed values w_γ at the corresponding projection image of $\gamma \mathring{v}_H$ in V_1. Extend f trivially to V one obtains a polynomial $f : V \to \mathbb{C}$ satisfying the required conditions. Now define $P : V \to W$ by

$$P(v) = \Psi \left(\sum_{\gamma \in G} f(\gamma v) \cdot \langle \gamma \rangle \right).$$

Then this is a G-equivariant polynomial map sending $\mathring{v}_H$ to $\mathring{w}_H$. □

Proof of Proposition 5.4.7. For each subgroup $H \subset G$, we can write any polynomial map $P \in \mathrm{Poly}_d^G(V, W)$ as

$$P = (\mathring{P}^H, \check{P}^H) : V \to \mathring{W}^H \oplus \check{W}^H.$$

Then the equivariance implies that

$$\check{P}^H|_{\mathring{V}_H} \equiv 0.$$

Therefore

$$Z^*_{d,H} := \{(v, P) \in \mathring{V}^*_H \times \mathrm{Poly}_d^G(V, W) \mid \mathring{P}^H(v) = 0\}.$$

Then when $d \geq d_0$, Lemma 5.4.3 and the *faithfulness* of the G-action on V imply that $Z^*_{d,H}$ is a nonsingular complex algebraic set of dimension

$$\dim_{\mathbb{C}} \mathrm{Poly}_d^G(V, W) + \dim_{\mathbb{C}} \mathring{V}^H - \dim_{\mathbb{C}} \mathring{W}^H.$$

Lastly, by the construction of the canonical Whitney stratification of Z_d^G which respects the action prestratification (see Theorem 5.8.17), as Z_d^{free} is smooth of dimension k, it is entirely contained in $Z^*_{d,k}$. □

Remark 5.4.8. If W is the trivial representation, then d_0 can be taken to be 0. In fact, all constant maps from V to W is G-equivariant and $Z^0_G(V,W) = V \times \{0\} \subset V \times \mathrm{Poly}^G_0(V,W) \cong V \times W$ so the

$$Z^*_{0,H} = V^*_H \times \{0\}.$$

5.4.3 Change of degrees

Here we prove that the canonical Whitney stratifications from different cut-off degrees are compatible with each other. For each $d \geq 0$, consider the inclusion map

$$\phi : V \times \mathrm{Poly}^G_d(V,W) \to V \times \mathrm{Poly}^G_{d+1}(V,W).$$

Then obviously

$$\phi^{-1}(Z^G_{d+1}(V,W)) = Z^G_d(V,W).$$

Our main theorem of this subsection is the following.

Theorem 5.4.9. ϕ is transverse to $Z^G_{d+1}(V,W)$ and pulls back the canonical Whitney stratification on $Z^G_{d+1}(V,W)$ to the canonical Whitney stratification on $Z^G_d(V,W)$.

The idea of the proof is similar to the discussion in [Par13] using Parker's notion of nice Whitney stratifications. First we prove an algebraic result which was also used in [Par13] without providing a proof or reference.

Lemma 5.4.4. $\mathrm{Poly}^G(V,W)$ is a finitely generated $\mathbb{C}[V]^G$-module.

Proof. The proof follows from [(k], we present it here for completeness. By Hilbert's basis theorem, given any finite-dimensional complex G-representation V', the ring of G-invariant polynomials $\mathbb{C}[V']^G$ is finitely generated. Now let $V' = V \oplus W^\vee$ where $W^\vee$ is the dual to W endowed with the corresponding G-action. The ring $\mathbb{C}[V']^G$ has a bi-grading by keeping track of the degree of the V-coordinates and $W^\vee$-coordinates re-

spectively. Choose $h_1, \ldots, h_r, h_{r+1}, \ldots, h_{r+m}, \ldots, h_n$ which generate $\mathbb{C}[V']^G$ such that $h_1, \ldots, h_r$ have $W^\vee$-degree 0, $h_{r+1}, \ldots, h_{r+m}$ have $W^\vee$-degree 1, and $h_{r+m+1}, \ldots, h_n$ have higher $W^\vee$-degrees.

Note that $\mathrm{Poly}^G(V, W)$ can be identified with the subset of $\mathbb{C}[V']^G$ consisting of elements with W-degree 1. Then any element in $\mathrm{Poly}^G(V, W)$ could be written as a linear combination of products of $h_1, \ldots, h_r, h_{r+1}, \ldots, h_{r+m}$. Observe that $\mathrm{Poly}^G(V, W)$ is naturally a module over $\mathbb{C}[V]^G$, therefore it is generated by $h_{r+1}, \ldots, h_{r+m}$ as a $\mathbb{C}[V]^G$-module. □

Next we construct a left inverse of the map ϕ.

Lemma 5.4.5. (cf. [Par13, Lemma 4.11]) Given G, V, W. There exists an integer $d_0 = d_0(G, V, W)$ such that for all $d \geq d_0$, there is a G-invariant map

$$\underline{\phi}' : V \times \mathrm{Poly}^G_{d+1}(V, W) \to \mathrm{Poly}^G_d(V, W)$$

satisfying the following properties. Denote

$$\phi' : V \times \mathrm{Poly}^G_{d+1}(V, W) \to V \times \mathrm{Poly}^G_d(V, W), \ (v, P) \mapsto (v, \underline{\phi}'(v, P)). \tag{5.4.2}$$

Then

1. $\phi' \circ \phi = \mathrm{Id}_{V \times \mathrm{Poly}^G_d(V,W)}$.
2. For each $(v, P) \in V \times \mathrm{Poly}^G_{d+1}(V, W)$, $\mathrm{ev}(\phi'(v, P)) = \mathrm{ev}(v, P)$. In particular, ϕ' maps Z^G_{d+1} surjectively onto $Z^G_d(V, W)$
3. $\phi \circ \phi'$ is transverse to Z^G_{d+1} and pulls back the canonical Whitney stratification on Z^G_{d+1} to itself.

Proof. As $\mathrm{Poly}^G(V, W)$ is finitely generated over $\mathbb{C}[V]^G$ (see Lemma 5.4.4), one can find a sufficiently large d_0 such that $\mathrm{Poly}^G_{d_0}(V, W)$ contains a set of generators $Q_1, \ldots, Q_m$.

Then when $d \geq d_0$, let $\mathrm{Homo}^G_{d+1}(V,W) \subset \mathrm{Poly}^G_{d+1}(V,W)$ be the subset of G-equivariant homogeneous polynomial maps of degree $d+1$. Then one has

$$\mathrm{Poly}^G_{d+1}(V,W) = \mathrm{Poly}^G_d(V,W) \oplus \mathrm{Homo}^G_{d+1}(V,W).$$

Choose a basis $P_1, \ldots, P_k$ of $\mathrm{Homo}^G_{d+1}(V,W)$. Then there exist polynomials $h_{ij} \in \mathbb{C}[V]^G$ for $1 \leq i \leq k$, $1 \leq j \leq m$ such that

$$P_i = \sum_{j=1}^{m} h_{ij} Q_j.$$

Then for each $P \in \mathrm{Poly}^G_{d+1}(V,W)$, let $P' \in \mathrm{Homo}^G_{d+1}(V,W)$ be its degree $d+1$ part which can be uniquely written as

$$P' = \sum_{i=1}^{k} a_i P_i.$$

Then define

$$\underline{\phi}'(v,P) = (P - P' + \sum a_i h_{ij}(v) Q_j). \tag{5.4.3}$$

This map (which is not canonical) is linear in the variable P. Then it is easy to see that for the associated map ϕ' defined by (5.4.2) there holds

$$\phi' \circ \phi = \mathrm{Id}_{V \times \mathrm{Poly}^G_d(V,W)}$$

and for all $(v,P) \in V \times \mathrm{Poly}^G_{d+1}(V,W)$ one has

$$\mathrm{ev}(\phi(\phi'(v,P))) = \mathrm{ev}(\phi'(v,P)) = \mathrm{ev}(v,P). \tag{5.4.4}$$

Therefore ϕ' is surjective and maps Z^G_{d+1} to Z^G_d.

Now we prove the last property. Define $\mathcal{V}$ and $\mathcal{W}$ to be the trivial bundles over

$M_{d+1} := \mathrm{Poly}^G_{d+1}(V,W)$ with fibers V and W respectively. Then we have two bundle maps $\underline{f}_1, \underline{f}_2 : \mathcal{V} \cong V \times M_{d+1} \to \mathrm{Poly}^G_{d+1}(\mathcal{V},\mathcal{W}) \cong \mathrm{Poly}^G_{d+1}(V,W) \times M_{d+1}$ defined by

$$\underline{f}_1(v,P) = (P,P), \qquad\qquad \underline{f}_2(v,P) = (\underline{\phi}(\phi'(v,P)),P)$$

where $\underline{\phi}(v,P) = P$. Then (5.4.4) implies that

$$\mathrm{graph}(\underline{f}_1 - \underline{f}_2) \subset \mathcal{Z}^G_{d+1}(\mathcal{V},\mathcal{W}).$$

As the identity map of $V \times \mathrm{Poly}^G_{d+1}(V,W)$ is transverse to Z_{d+1} (which means the graph of $\underline{f}_1$ is transverse to $\mathcal{Z}_{d+1}$), it follows from Lemma 5.4.1 that the graph of $\underline{f}_2$ is also transverse to $\mathcal{Z}_{d+1}$. It is equivalent to say that the map $\phi \circ \phi'$ is transverse to Z_{d+1}. Moreover, as the identity map preserves the canonical Whitney stratification, by Lemma 5.4.2, $\phi \circ \phi'$ also pulls back the canonical Whitney stratification on Z_{d+1} to itself. □

Proof of Theorem 5.4.9. Still abbreviate $Z^G_d(V,W)$ by Z_d. By Lemma 5.4.5, we know that $\phi \circ \phi'$ resp. $\phi' \circ \phi$ is transverse to Z_{d+1} resp. Z_d, it follows that ϕ resp. ϕ' is transverse to Z_{d+1} resp. Z_d along images of ϕ' resp. ϕ. As ϕ' is surjective, it follows that ϕ is transverse to Z_{d+1}. Furthermore, notice that

$$\phi' = \phi' \circ \phi \circ \phi'.$$

Hence the image of $d\phi'$ at any point coincides with the image of $d\phi'$ at some point on the image of ϕ. As ϕ' is transverse to Z_d along the image of ϕ, it follows that ϕ' is transverse to Z_d everywhere.

Now we show that ϕ resp. ϕ' pulls back the canonical Whitney stratification to the canonical one. We prove inductively that for each subgroup $H \subseteq G$, ϕ resp. ϕ' pulls back the canonical Whitney stratification on $Z^*_{d+1,H}$ resp. $Z^*_{d,H}$ to the canonical

Whitney stratification on $Z^*_{d,H}$ resp. $Z^*_{d+1,H}$. For H being the trivial group, notice that the pair of maps ϕ and ϕ' satisfies the assumptions of the absolute case of Lemma 5.8.2 for $f = \phi$, $g = \phi'$, $M = V^{\text{free}} \times \text{Poly}^G_d(V, W)$, $N = V^{\text{free}} \times \text{Poly}^G_{d+1}(V, W)$, and $Y = Y^\circ = Z^{\text{free}}_d$, $Z = Z^\circ = Z^{\text{free}}_{d+1}$. Hence by Lemma 5.8.2, ϕ resp. ϕ' pulls back the canonical Whitney stratification on Z^{free}_{d+1} resp. Z^{free}_d to the canonical Whitney stratification on Z^{free}_d resp. Z^{free}_{d+1}.[2] Suppose for a subgroup $H \subset G$, one has proved the claim for all proper subgroups $H' \subsetneq H$. Then by the relative case of Lemma 5.8.2, one can show that ϕ pulls back the canonical Whitney stratification on $Z^*_{d+1,H}$ resp. $Z^*_{d,H}$ to the canonical Whitney stratification on $Z^*_{d,H}$ resp. $Z^*_{d+1,H}$. □

Remark 5.4.10. If W is a trivial G-representation, then the statement of Theorem 5.4.9 is true for $d = 0$. In fact for any $d > 0$, the projection map

$$\underline{\phi}' : V \times \text{Poly}^G_d(V, W) \to \text{Poly}^G_0(V, W) \cong W$$

constructed in the above proof coincides with the evaluation map.

By Theorem 5.4.9, for any sufficiently large cut-off degree d, there is a canonical Whitney stratification on $Z^G_d(V, W)$ which is natural with respect to inclusion maps. We need a slight extension of this result. Suppose W_1 and W_2 are two complex G-representations, then for all sufficiently large d_1 and d_2, consider the set

$$\begin{aligned} Z^G_{d_1,d_2}(V, W_1, W_2) &= Z^G_{d_1}(V, W_1) \times_V Z^G_{d_2}(V, W_2) \\ &:= \{(v, P_1, P_2) \in V \times \text{Poly}^G_{d_1}(V, W_1) \times \text{Poly}^G_{d_2}(V, W_2) \mid P_1(v) = 0,\ P_2(v) = 0\}. \end{aligned}$$

Abbreviate it by Z_{d_1,d_2}. Then the construction of Theorem 5.8.17 provides a minimal Whitney stratification on Z_{d_1,d_2} which respects the action prestratification, which is also a nice Whitney stratification. We also call it the canonical Whitney stratification.

[2]Because of Proposition 5.4.7, the base case of the induction does not need Lemma 5.8.2.

Consider $d \geq d_1, d_2$ and the natural inclusion map

$$\phi : V \times \mathrm{Poly}^G_{d_1}(V, W_1) \times \mathrm{Poly}^G_{d_2}(V, W_2) \to V \times \mathrm{Poly}^G_d(V, W_1 \oplus W_2)$$

which maps Z_{d_1,d_2} into $Z^G_d(V, W_1 \oplus W_2)$.

Proposition 5.4.11. When d_1, d_2 are sufficiently large, ϕ pulls back the canonical Whitney stratification on $Z^G_d(V, W_1 \oplus W_2)$ to the canonical one on Z_{d_1,d_2}.

Proof. One can construct a similar projection map

$$\underline{\phi}' : V \times \mathrm{Poly}^G_d(V, W_1 \oplus W_2) \to \mathrm{Poly}^G_{d_1}(V, W_1) \times \mathrm{Poly}^G_{d_2}(V, W_2).$$

The rest of the argument is identical to the proof of Theorem 5.4.9. □

Corollary 5.4.12. Let $\mathring{W}$ be a trivial complex G-representation and $\check{W}$ be another complex G-representation. Let V be an effective complex G-representation. Let M be a base manifold. Let $\mathcal{V}$, $\mathring{\mathcal{W}}$, $\check{\mathcal{W}}$ be the trivial bundles over M with fibers being V, $\mathring{W}$, $\check{W}$ respectively. Let $\mathcal{V}^* \subset \mathcal{V}$ be a G-invariant open subset. Let d be sufficiently large and let

$$\underline{f} = (\mathring{\underline{f}}, \check{\underline{f}}) : \mathcal{V}^* \to \mathrm{Poly}^G_d(\mathcal{V}, \mathring{\mathcal{W}} \oplus \check{\mathcal{W}})$$

be a smooth bundle map. Consider the partial evaluation of $\underline{f}$ which is

$$\underline{f}' : \mathcal{V}^* \to \mathring{W} \times \mathrm{Poly}^G_d(\mathcal{V}, \check{\mathcal{W}})$$

defined by

$$\underline{f}'(v) = (\mathring{\underline{f}}(v)(v), \check{\underline{f}}).$$

Then the graph of $\underline{f}$ is transverse to $\mathcal{Z}^G_d(\mathcal{V}^*, \mathring{\mathcal{W}} \oplus \check{\mathcal{W}})$ if and only if the graph of $\underline{f}'$ is transverse to $\{0\} \times \mathcal{Z}^G_d(\mathcal{V}^*, \check{\mathcal{W}})$.

5.4.4 Change of groups

Let V, W be complex G-spaces. In this subsection, fix a proper subgroup $H \subsetneq G$. Then we have the basic decomposition

$$V = \mathring{V}_H \oplus \check{V}_H$$

where $\check{V}_H$ is the direct sum of nontrivial irreducible H-spaces. There is a *restriction map*

$$\underline{\theta} : V \times \mathrm{Poly}^G(V, W) \to \mathrm{Poly}^H(\check{V}_H, W)$$

defined by (denoting $v = (\mathring{v}_H, \check{v}_H) \in \mathring{V}_H \oplus \check{V}_H$)

$$\underline{\theta}(v, P) = P|_{\{\mathring{v}_H\} \times \check{V}_H}.$$

Define

$$\begin{aligned} \theta : V \times \mathrm{Poly}^G(V, W) &\to V \times \mathrm{Poly}^H(\check{V}_H, W) \\ (v, P) &\mapsto (v, \underline{\theta}(v, P)). \end{aligned} \tag{5.4.5}$$

If we view $\mathring{V}_H$ as a base manifold with the trivial H-action, then over $\mathring{V}_H$ one has the product vector bundles $\check{\mathcal{V}}_H = \mathring{V}_H \times \check{V}_H$ and $\mathcal{W} = \mathring{V}_H \times W$. Then the target of θ can be identified with the fiber product $\check{\mathcal{V}}_H \times_{\mathring{V}_H} \mathrm{Poly}^H(\check{\mathcal{V}}_H, \mathcal{W})$. Then for each degree d one has

$$\theta(Z_d^G(V, W)) \subset \mathcal{Z}_d^H(\check{\mathcal{V}}_H, \mathcal{W}) \subset \check{\mathcal{V}}_H \times_{\mathring{V}_H} \mathrm{Poly}_d^H(\check{\mathcal{V}}_H, \mathcal{W}). \tag{5.4.6}$$

We will compare the canonical Whitney stratification on $Z_d^G(V, W)$ and the canonical Whitney stratification on $\mathcal{Z}_d^H(\check{\mathcal{V}}_H, \mathcal{W})$, where the latter is the parametrized version induced from the canonical Whitney stratification on $Z_d^H(\check{V}_H, W)$.

Denote

$$V_H^+ := \{v \in V \mid G_v \subset H\}$$

which is the complement of the union of finitely many subspaces. Let $\check{\mathcal{V}}_H^+ \subset \check{\mathcal{V}}_H$ be the corresponding open subset of $\check{\mathcal{V}}_H$. Denote

$$Z_d^G(V_H^+, W) := Z_d^G(V, W) \cap (V_H^+ \times \mathrm{Poly}_d^G(V, W))$$

and

$$\mathcal{Z}_d^H(\check{\mathcal{V}}_H^+, \mathcal{W}) := \mathcal{Z}_d^H(\check{\mathcal{V}}, \mathcal{W}) \cap (\check{\mathcal{V}}_H^+ \times_{\check{V}_H} \mathrm{Poly}_d^H(\check{\mathcal{V}}_H, \mathcal{W})),$$

which are Zariski open subsets of $Z_d^G(V, W)$ and $\mathcal{Z}_d^H(\check{\mathcal{V}}_H, \mathcal{W})$ respectively. Hence they have the restricted canonical Whitney stratifications. Our main theorem of this subsection is the following.

Theorem 5.4.13. θ is transverse to $\mathcal{Z}_d^H(\check{\mathcal{V}}_H^+, \mathcal{W})$ and pulls back the canonical Whitney stratification on $\mathcal{Z}_d^H(\check{\mathcal{V}}_H^+, \mathcal{W})$ to the canonical one on $Z_d^G(V_H^+, W)$.

The main idea of the proof also comes from [Par13] and is analogous to the proof of Theorem 5.4.9. We first define a smooth map

$$\underline{\theta}' : V_H^+ \times \mathrm{Poly}^H(\check{V}_H, W) \to \mathrm{Poly}^G(V, W)$$

which is roughly an "inverse" of θ. Using the Lagrange interpolation method, one can find an integer $d_0 = d(G, V) > 0$ such that for each $v \in V_H^+$, there is a polynomial $L_v : V \to \mathbb{C}$ of degree at most d_0 satisfying

$$L_v(gv) = \begin{cases} 1, & \text{if } g \in H, \\ 0, & \text{if } g \in G \setminus H. \end{cases} \tag{5.4.7}$$

L_v can be made smoothly dependent on $v \in V_H^+$. After averaging over H (which is a

linear transformation) we can require that L_v is H-invariant. Also, using the decomposition $V = \mathring{V}_H \oplus \check{V}_H$, regard any $P \in \mathrm{Poly}^H(\check{V}_H, W)$ an H-equivariant polynomial map from V to W which is constant in $\mathring{V}_H$-direction. Then define

$$\underline{\theta}''(v, P) = \frac{1}{|H|} \sum_{g \in G} g^{-1} \circ (L_v \cdot P) \circ g.$$

This map is clearly G-equivariant. Moreover, (5.4.7) implies that

$$\underline{\theta}''(v, P)(v) = P(\check{v}_H) \quad \forall v = (\mathring{v}_H, \check{v}_H) \in V_H^+ \subset V_H \times N_H.$$

Moreover, for each $d > 0$, one has

$$\underline{\theta}''(V_H^+ \times \mathrm{Poly}_d^H(\check{V}_H, W)) \subseteq \mathrm{Poly}_{d+d_0}^G(V, W).$$

Now, we use iterations of the map $\underline{\phi}'$ defined by (5.4.3) to reduce the degree. More precisely, by the construction of the map $\underline{\phi}'$, if $d > 0$ is large enough, there exists a G-equivariant map

$$\underline{\phi}'' : V \times \mathrm{Poly}_{d+d_0}^G(V, W) \to \mathrm{Poly}_d^G(V, W)$$

satisfying the following conditions.

1. $\underline{\phi}''$ is linear in the second variable such that the associated map
2. The associated map

$$\phi'' : V \times \mathrm{Poly}_{d+d_0}^G(V, W) \to V \times \mathrm{Poly}_d^G(V, W), \ (v, P) \mapsto (v, \underline{\phi}''(v, P))$$

   preserves the evaluation map.
3. The restriction of ϕ'' to $V \times \mathrm{Poly}_d^G(V, W)$ is the identity map.

Then define

$$\underline{\theta}' = \underline{\phi}'' \circ \underline{\theta}'' : V_H^+ \times \mathrm{Poly}_d^H(\check{V}_H, W) \to \mathrm{Poly}_d^G(V, W).$$

We denote

$$\theta'(v, P) = (v, \underline{\theta}'(v, P)).$$

Note that θ' is not canonically defined, but we will just use the existence of such maps to deduce propositions independent of the choice of θ'.

Before proving the next result, we do some preparations. Consider the compositions

$$\underline{\theta} \circ \theta' : V_H^+ \times \mathrm{Poly}_d^H(\check{V}_H, W) \to \mathrm{Poly}_d^H(\check{V}_H, W)$$

and

$$\underline{\theta}' \circ \theta : V_H^+ \times \mathrm{Poly}_d^G(V, W) \to \mathrm{Poly}_d^G(V, W).$$

For any $Q \in \mathrm{Poly}_d^H(\check{V}_H, W)$, define a vector field on $V_H^+ \times \mathrm{Poly}_d^H(\check{V}_H, W)$ by

$$v_Q(x, P) = (0, Q - \underline{\theta}(\theta'(x, Q)));$$

for any $Q \in \mathrm{Poly}_d^G(V, W)$, define a vector field on $V_H^+ \times \mathrm{Poly}_d^G(V, W)$ by

$$w_Q(x, P) = (0, Q - \underline{\theta}'(\theta(x, Q))).$$

Their flows are

$$\Phi_{Q,t}(x, P) = (x, P + t(Q - \underline{\theta}(\theta'(x, Q)))),$$
$$\Psi_{Q,t}(x, P) = (x, P + t(Q - \underline{\theta}'(\theta(x, Q))))$$

By the construction of θ', one can see that the flows preserve the corresponding Z-varieties.

Lemma 5.4.6. For any $Q \in \mathrm{Poly}_d^H(\check{V}_H, W)$ resp. $Q \in \mathrm{Poly}_d^G(V, W)$ and $t \in \mathbb{R}$, consider

the map

$$F_{Q,t} := \theta' \circ \Phi_{Q,t} \circ \theta : V_H^+ \times \mathrm{Poly}_d^G(V,W) \to V_H^+ \times \mathrm{Poly}_d^G(V,W)$$

resp.

$$G_{Q,t} := \theta \circ \Psi_{Q,t} \circ \theta' : V_H^+ \times \mathrm{Poly}_d^H(\check{V}_H,W) \to V_H^+ \times \mathrm{Poly}_d^H(\check{V}_H,W).$$

Then the following is true.

1. $F_{Q,t}^{-1}(Z_d^G(V_H^+,W)) = Z_d^G(V_H^+,W)$ and $G_{Q,t}^{-1}(\mathcal{Z}_d^H(\check{\mathcal{V}}_H^+,\mathcal{W})) = \mathcal{Z}_d^H(\check{\mathcal{V}}_H^+,\mathcal{W})$.
2. $F_{Q,t}$ resp. $G_{Q,t}$ is transverse to $Z_d^G(V_H^+,W)$ resp. $\mathcal{Z}_d^H(\check{\mathcal{V}}_H^+,\mathcal{W})$ and pulls back the canonical Whitney stratification to itself.

Proof. Notice that θ, θ', $\Phi_{Q,t}$, and $\Psi_{Q,t}$ all preserve the evaluation map. Hence one has

$$\mathrm{ev}(F_{Q,t}(v,P)) = \mathrm{ev}(v,P) = P(v),\ \mathrm{ev}(G_{Q,t}(v,P)) = \mathrm{ev}(v,P) = P(v).$$

Therefore, the first item is true.

To prove (2), regard $\mathrm{Poly}_d^G(V,W)$ as a base manifold M_d^G with trivial G-action, over which we have the trivial vector bundles

$$\mathcal{E} := M_d^G \times V, \qquad \mathcal{F} := M_d^G \times W.$$

Denote

$$\mathcal{E}_H^+ := V_H^+ \times M_d^G \subset \mathcal{E}.$$

Then $F_{Q,t}$ can be viewed as a bundle map

$$\underline{f}_{Q,t} : \mathcal{E}_H^+ \to \mathrm{Poly}_d^G(\mathcal{E},\mathcal{F}) \cong M_d^G \times \mathrm{Poly}_d^G(V,W).$$

On the other hand, there is the natural inclusion map

$$\underline{\iota} : \mathcal{E}_H^+ \to \mathrm{Poly}_d^G(\mathcal{E}, \mathcal{F}),\ \underline{\iota}(P, v) := (P, P).$$

Then the construction implies that

$$\mathrm{graph}(\underline{f}_{Q,t} - \underline{\iota}) \subset \mathcal{Z}_d^G(\mathcal{E}, \mathcal{F}).$$

As the graph of $\underline{\iota}$ is transverse to the canonical Whitney stratification on $\mathcal{Z}_d^G(\mathcal{E}, \mathcal{F})$, by Lemma 5.4.1 the graph of $\underline{f}_{Q,t}$ is also transverse to $\mathcal{Z}_d^G(\mathcal{E}, \mathcal{F})$. Equivalently, it means that the composition $F_{Q,t}$ is transverse to $Z_d^G(V_H^+, W)$. The case of $G_{Q,t}$ is the same. □

Proof of Theorem 5.4.13. We first prove that θ' resp. θ is transverse to $Z_d^G(V_H^+, W)$ resp. $\mathcal{Z}_d^H(\check{\mathcal{V}}_H^+, \mathcal{W})$. Consider θ'. The above lemma implies that θ' is transverse to $Z_d^G(V_H^+, W)$ along the image of $\Phi_{Q,t} \circ \theta$. The global transversality of θ' follows if we can show that any point in $Z_d^H(V_H^+, W)$ lies in the image of $\Phi_{Q,t} \circ \theta$ for some Q and some t. Indeed, for any $(x, Q) \in Z_d^H(V_H^+, W)$, one has

$$(x, Q) = \Phi_{Q,1}(x, \underline{\theta}(\theta'(x, Q))) = \Phi_{Q,1}(\theta(x, \underline{\theta}'(x, Q))) \in \mathrm{im}(\Phi_{Q,1} \circ \theta).$$

Hence θ' is transverse to $Z_d^G(V, W)$ at (x, Q). As (x, Q) is an arbitrary point of $Z_d^H(V_H^+, W)$, it follows that θ' is transverse to $Z_d^G(V, W)$. The transversality statement for θ is similar.

It is then similar to Theorem 5.4.9 to show that θ resp. θ' pulls back the canonical Whitney stratification to the canonical one, using Lemma 5.8.2 and an induction on strata of the action prestratification. Here are the details. Abbreviate $\mathrm{Poly}_d^G(V, W)$ by P_d^G and $\mathrm{Poly}_d^H(\check{V}_H, W)$ by P_d^H; abbreviate $Z_d^G(V_H^+, W)$ by Z_d^G and $\mathcal{Z}_d^H(\check{\mathcal{V}}_H^+, \mathcal{W})$ by

$\mathcal{Z}_d^H$. First notice that

$$Z_d^G(V_H^+, W) \subset V_H^+ \times P_d^G, \qquad \mathcal{Z}_d^H(\check{\mathcal{V}}_H^+, \mathcal{W}) \subset V_H^+ \times P_d^H.$$

The the strata of the action prestratifications on $V_H^+ \times P_d^G$ and $V_H^+ \times P_d^H$ are both indexed by the set

$$B = \{I < G \mid I \subseteq H\}.$$

Then by Remark 5.8.18, the canonical Whitney stratifications on Z_d^G and $\mathcal{Z}_d^H$ are still the minimal smooth Whitney stratifications which respect the action prestratifications. For each $I \in B$, denote

$$Z_{d,I}^G := Z_d^G \cap (V_I^* \times P_d^G), \qquad \mathcal{Z}_{d,I}^H := \mathcal{Z}_d^H \cap (V_I^* \times P_d^H).$$

Then

$$\theta^{-1}(\mathcal{Z}_{d,I}^H) = Z_{d,I}^G, \qquad (\theta')^{-1}(Z_{d,I}^G) = \mathcal{Z}_{d,I}^H.$$

Then we prove inductively the following statement.

• Suppose $I \in B$. Then θ resp. θ' pulls back the restriction of the canonical Whitney stratifications on $\mathcal{Z}_{d,I}^H$ resp. $Z_{d,I}^G$ to the canonical one on $Z_{d,I}^G$ resp. $\mathcal{Z}_{d,I}^H$.

Notice that we just verified condition (1) of Lemma 5.8.2. On the other hand, by the $Q = 0$, $t = 0$ case of Lemma 5.4.6, conditions (3) and (4) of Lemma 5.8.2 are checked. Hence by the absolute case of Lemma 5.8.2, the base case of the induction is verified. Suppose we have proved that for some $I_0 \in B$, the above statement is true for all $I \in B$ with $I \subsetneq I_0$. Then the above statement follows from the relative case of

Lemma 5.8.2 where

$$Y^\bullet = \bigsqcup_{I \subsetneq I_0} Z^G_{d,I}, \qquad\qquad Z^\bullet = \bigsqcup_{I \subsetneq I_0} \mathcal{Z}^H_{d,I}$$

as the induction hypothesis verifies condition (5) of Lemma 5.8.2.

□

5.5 Pseudocycles in Thom–Mather stratified spaces

In this section we define the notion of pseudocycles in Thom–Mather stratified spaces and extend half of the main results of Zinger [Zin08] about pseudocycles in manifolds. This discussion is necessary as we would like to define the homology classes supported on the perturbed zero loci in effective orbifolds, which are naturally Thom–Mather stratified spaces. However, we remark that we do not need the result of this section to do intersection theory such as to define Gromov–Witten type invariants or Floer homology.

Definition 5.5.1. [Tho64, Tho69][Mat12] A *Thom–Mather stratified space* is a triple $(X, \mathcal{S}, \mathfrak{J})$ satisfying the following axioms (A1)–(A7).

(A1) X is a locally compact, Hausdorff, and second countable topological space.

(A2) $\mathcal{S}$ is a prestratification of X with strata $X_\alpha \subset X$.

(A3) Each member of $\mathcal{S}$ is a smooth manifold.

(A4) $\mathfrak{J}$ (called the *collection of control data*) is a triple $\{(N_\alpha), (\pi_\alpha), (\rho_\alpha)\}$ where for each $X_\alpha \in \mathcal{S}$, N_α is an open neighborhood of X_α, $\pi_\alpha : N_\alpha \to X_\alpha$ is a continuous retraction onto X_α, and $\rho_\alpha : N_\alpha \to [0, +\infty)$ is a continuous function.

(A5) $X_\alpha = \{v \in N_\alpha \mid \rho_\alpha(v) = 0\}$.

(A6) For each pair of strata $X_\alpha, X_\beta \in \mathcal{S}$, define $N_{\beta\alpha} = N_\alpha \cap X_\beta$ (which is an open subset of the manifold X_β), $\pi_{\beta\alpha} = \pi_\alpha|_{N_{\beta\alpha}}$, and $\rho_{\beta\alpha} = \rho_\alpha|_{N_{\beta\alpha}}$. We require that the map

$$(\pi_{\beta\alpha}, \rho_{\beta\alpha}) : N_{\beta\alpha} \to X_\alpha \times (0, +\infty)$$

is a smooth submersion.

(A7) For any three strata $X_\alpha, X_\beta, X_\gamma$ one has

$$\pi_{\beta\alpha} \circ \pi_{\gamma\beta} = \pi_{\gamma\alpha},$$
$$\rho_{\beta\alpha} \circ \pi_{\gamma\beta} = \rho_{\gamma\alpha}$$

whenever both sides of the equations are defined.

For each $\epsilon > 0$ sufficiently small, denote

$$N_\alpha(\epsilon) = \{v \in N_\alpha \mid \rho_\alpha(v) < \epsilon\}, \qquad S_\alpha(\epsilon) = \{v \in N_\alpha \mid \rho_\alpha(v) = \epsilon\},$$

and

$$X_\beta^0(\epsilon) := X_\beta \setminus \bigcup_{\alpha \leq \beta} N_\alpha(\epsilon).$$

Definition 5.5.2. (Family of lines)[Gor76, Gor78a] Given a Thom–Mather stratified space $(X, \mathcal{S}, \mathfrak{J})$, a *family of lines* consists of a positive number $\delta > 0$ and a system of retractions

$$r_\alpha(\epsilon) : N_\alpha \setminus X_\alpha \to S_\alpha(\epsilon),\ \epsilon \in (0, \delta),\ X_\alpha \in \mathcal{S}$$

satisfying the following conditions. Let $X_\alpha < X_\beta$.

1. $r_\alpha(\epsilon) \circ r_\beta(\epsilon') = r_\beta(\epsilon') \circ r_\alpha(\epsilon)$.
2. $\rho_\alpha \circ r_\beta(\epsilon) = \rho_\alpha$.
3. $\rho_\beta \circ r_\alpha(\epsilon) = \rho_\beta$.

4. $\pi_\alpha \circ r_\beta(\epsilon) = \pi_\alpha$.

5. If $0 < \epsilon < \epsilon' < \delta$ then $r_\alpha(\epsilon') \circ r_\alpha(\epsilon) = r_\alpha(\epsilon')$.

6. $\pi_\alpha \circ r_\alpha(\epsilon) = \pi_\alpha$.

7. $r_\alpha(\epsilon)|_{N_\alpha(\epsilon)\cap X_\beta} : N_\alpha(\epsilon) \cap Y \to S_\alpha(\epsilon) \cap Y$ is smooth.

Example 5.5.1.

1. Mather [Mat12] proved that a Whitney stratified set S in a smooth manifold M is a Thom–Mather stratified space.

2. Let G be a compact Lie group acting smoothly on a smooth manifold M. Then Lellmann [Lel75] proved that the orbit space M/G is a Thom–Mather stratified space.[3] As all effective orbifolds can be expressed as a global quotient by a compact Lie group, an effective orbifold admits the structure of Thom–Mather stratified space whose strata are labelled by orbit types. Given an orbifold with a uniform bound the orders of the isotropy groups, it can also be presented as a global quotient of a compact Lie group [Par19, Corollary 1.3] so it is also a Thom–Mather stratified space as above.

Theorem 5.5.3. [Gor76][Gor78a] Any Thom–Mather stratified space admits a family of lines.

The family of lines induces a system of maps

$$\begin{aligned} h_\alpha : N_\alpha \setminus X_\alpha &\to S_\alpha(\epsilon) \times \mathbb{R}_+ \\ v &\mapsto (r_\alpha(\epsilon)(v), \rho_\alpha(v)) \end{aligned}$$

which is a homeomorphism onto its image and which is a diffeomorphism on each stratum.

[3]The authors cannot access the original paper of Lellmann. A complete proof can be found in [DS06, Section II.4].

5.5.1 Pseudocycles

We generalize the notion of pseudocycles in smooth manifolds to the case of Thom–Mather stratified spaces.

Definition 5.5.4. Let X be a Thom–Mather stratified space with strata X_α. Let W be a smooth manifold. A continuous map $F : W \to X$ is called *smooth* if for each α, $W_\alpha := F^{-1}(X_\alpha)$ is a smooth submanifold of W and the restriction $F_\alpha := F|_{W_\alpha} : W_\alpha \to X_\alpha$ is a smooth map.

Definition 5.5.5. A subset $A \subset X$ is said to *have dimension* $\leq k$ if there exists a smooth manifold W and a smooth map $F : W \to X$ such that $A \subseteq F(W)$ such that for each strata X_α, $\dim W_\alpha \leq k$.

Moreover, for any continuous map $f : A \to B$ between topological spaces. The *boundary* of f is defined to be

$$\mathrm{Bd} f := \bigcap_{K \subset A \text{ compact}} \overline{f(A \setminus K)}. \tag{5.5.1}$$

Definition 5.5.6. (Pseudocycles and cobordism)

1. A smooth map $F : W \to X$ is called a *k-dimensional pseudocycle* in X if W is oriented, $F(W)$ is precompact in X, $\dim W = k$, and $\mathrm{Bd} F$ has dimension at most $k - 2$.

2. Two k-dimensional pseudocycles $F_i : W_i \to X$ $(i = 0, 1)$ is called *cobordant* if there exist an oriented $k+1$-dimensional manifold $\tilde{W}$ with boundary such that $\partial \tilde{W} = W_0 \sqcup W_1$ compatible with the orientations, and there exists a smooth map $\tilde{F} : \tilde{W} \to X$ such that $\tilde{F}(\tilde{W})$ is precompact in X and such that

$$\tilde{F}|_{W_0} = F_0, \ \tilde{F}|_{W_1} = F_1, \ \mathrm{Bd} \tilde{F} \leq k - 1.$$

One can define an abelian group $\mathcal{H}_k(X)$ whose elements are cobordism classes of oriented pseudocycles in X and whose addition is induced from disjoint unions of pseudocycles. Our main conclusion is

Theorem 5.5.7. There exists a homomorphism

$$\Psi_X : \mathcal{H}_k(X) \to H_*(X;\mathbb{Z})$$

such that if an element $A \in \mathcal{H}_k(X)$ is represented by a smooth map $F : W \to X$ where X is compact and oriented, then

$$\Psi_X(A) = F_*([W]).$$

When X is a smooth manifold, this theorem is rigorously proved by Zinger [Zin08]. In fact Zinger showed that Ψ is an isomorphism.

5.5.2 A technical lemma

The proof of Theorem 5.5.7 follows the same line as in [Zin08]. The key step is a generalization of [Zin08, Proposition 2.2].

Lemma 5.5.1. Suppose $\dim W = k$ and $F : W \to X$ is a smooth map. Then there exists an open neighborhood $U \subset X$ of $F(W)$ such that

$$H_s(U;\mathbb{Z}) = 0, \ \forall s > k.$$

We first introduce a few new notations. Given a pair of strata $X_\alpha \leq X_\beta$, define $S_{\beta\alpha}(\epsilon) := S_\alpha(\epsilon) \cap X_\beta$. Choose a sufficiently small $d_0 > 0$ and consider the open set

$$W_{\beta\alpha}(d_0) := F^{-1}(N_\alpha(d_0)) \cap W_\beta = F^{-1}(N_\alpha(d_0) \cap X_\beta) \subset W_\beta.$$

Then the composition

$$\pi_{\beta\alpha} \circ F_\beta : W_{\beta\alpha}(d_0) \to X_\alpha$$

is smooth.

Lemma 5.5.2. There exists a collection of real numbers $\{d_\alpha \in (0, d_0) \mid \alpha \in \mathcal{S}\}$ satisfying the following conditions.

1. For each stratum β and $\alpha_1, \ldots, \alpha_s \leq \beta$, the intersection

$$S_{\beta\alpha_1}(d_{\alpha_1}) \cap \cdots \cap S_{\beta\alpha_s}(d_{\alpha_s})$$

   is transverse (in X_β).

2. For each pair of strata $\beta \leq \beta'$ and for $\alpha_1, \ldots, \alpha_s \leq \beta$, the map

$$\pi_{\beta'\beta} \circ F_{\beta'} : W_{\beta'\beta}(d_0) \to X_\beta$$

   is transverse to $S_{\beta\alpha_1}(d_{\alpha_1}) \cap \cdots \cap S_{\beta\alpha_s}(d_{\alpha_s})$.

Proof. Sard's theorem. □

Now by rescaling the control data by positive constants, we may assume that all d_α are equal to some $d \in (0, d_0)$. We can also assume that d is sufficiently small so that

$$N_\alpha(d) \cap N_\beta(d) \neq \emptyset \Longrightarrow \alpha \leq \beta \text{ or } \beta \leq \alpha.$$

Now following Goresky [Gor78a], we would like to construct an "interior triangulation."

Definition 5.5.8. Let $(X, \mathcal{S}, \mathfrak{J})$ be a Thom–Mather stratified space. For $d > 0$ sufficiently small, a (smooth) *interior d-triangulation* consists of a simplicial complex K, a continuous map $\eta : |K| \to X$ satisfying the following conditions.

1. η is a homeomorphism onto the closed set $\bigcup_{\alpha\in\mathcal{S}} X^0_\alpha(d)$.

2. For each α, $K_\alpha := \eta^{-1}(X_\alpha)$ is a subcomplex.

3. $\eta : |K_\alpha| \to X_\alpha$ is smooth over interior of simplices.

4. The map $\eta^{-1} \circ \pi_\alpha \circ \eta : \eta^{-1}(S_\alpha(d)) \to \eta^{-1}(X_\alpha)$ is a simplicial map.

Lemma 5.5.3. There exists a smooth interior d-triangulation satisfying the following conditions. For each pair of strata $\alpha \leq \beta$, the map $\pi_{\beta\alpha} \circ F_\beta : W_{\beta\alpha}(d_0) \to X_\alpha$ is transverse to the interior of each simplex of the subcomplex K_α.

Proof. In [Gor76] and [Gor78a] Goresky showed that there exists an interior d-triangulation. His construction is based on an increasing induction: given any interior triangulation upto the i-th stratum, one can extend it to an interior triangulation of the $i+1$-st stratum. Then in each step of this induction, we may require the corresponding transversality condition. □

We order the strata of X increasingly as $\alpha_1, \ldots, \alpha_m$. Namely

$$\alpha_i \leq \alpha_j \Longrightarrow i \leq j.$$

We can decompose the domain W as follows. First, notice that

$$X = \bigcup_{i=1}^{m} N_{\alpha_i}(d) = \bigsqcup_{i=1}^{m} \left(N_{\alpha_i}(d) \setminus \bigcup_{j<i} N_{\alpha_j}(d) \right).$$

Lemma 5.5.4. There holds

$$\pi_{\alpha_i} \left(N_{\alpha_i}(d) \setminus \bigcup_{j<i} N_{\alpha_j}(d) \right) = X^0_{\alpha_i}(d).$$

Hence

$$N_{\alpha_i}(d) \setminus \bigcup_{j<i} N_{\alpha_j}(d) = \pi^{-1}_{\alpha_i} \left(\eta_{\alpha_i}(|K_{\alpha_i}|) \right).$$

Proof. This follows from the commutation relation of the control data. □

For each simplex $\sigma \in K$, as $\eta(|\sigma|)$ is contained in a stratum (say X_α), one can define its codimension as

$$\mathrm{codim}\sigma = \mathrm{dim}X_\alpha - \mathrm{dim}\sigma.$$

Then by the transversality condition guaranteed by Lemma 5.5.3, one has

$$F(W) \cap \left(N_{\alpha_i}(d) \setminus \bigcup_{j<i} N_{\alpha_j}(d) \right) \subset \pi_{\alpha_i}^{-1} \left(\bigcup_{\substack{\sigma \in K_{\alpha_i} \\ \mathrm{codim}\sigma \leq k}} \eta(\mathrm{Int}\sigma) \right) \cap N_{\alpha_i}(d),$$

where recall that k is the dimension of W. For each m, define

$$K_\alpha(m) = \{\sigma \in K_\alpha \mid \mathrm{codim}\sigma = m,\ \pi_\alpha(F(W) \cap N_\alpha(d_0)) \cap \eta(\mathrm{Int}\sigma) \neq \emptyset\}.$$

Define

$$\underline{U}_i = \bigcup_{m=0}^{k} \bigcup_{\sigma \in K_{\alpha_i}(m)} \eta(\mathrm{Int}\sigma) \subset X^0_{\alpha_i}(d)$$

and

$$U_i := \pi_{\alpha_i}^{-1}(\underline{U}_i) \cap N_{\alpha_i}(d) \subset X.$$

Then we see $F(W) \subset U_1 \cup \cdots \cup U_m$. What remains to be proven is that $U_1 \cup \cdots \cup U_m$ is open and that

$$H_s(U_1 \cup \cdots \cup U_m; \mathbb{Z}) = 0,\ \forall s > k.$$

We prove inductively that $U_1 \cup \cdots \cup U_i$ is open and that $H_s(U_1 \cup \cdots \cup U_i; \mathbb{Z}) = 0$ for all $s > k$ and i. For the base case $i = 1$, as X_{α_1} is a lowest stratum, U_1 is open.

Lemma 5.5.5. For each i, one has

$$H_s(U_i; \mathbb{Z}) = 0\ \forall s > k.$$

Proof. We follow Zinger's argument. Let $\mathrm{sd}K$ be the barycentric subdivision of a simplicial complex K. Then one has

$$\underline{U}_i = \bigcup_{m=0}^{k} \bigcup_{\sigma \in K_{\alpha_i}(m)} \eta(\mathrm{Int}\sigma) = \bigcup_{m=0}^{k} \bigcup_{\sigma \in K_{\alpha_i}(m)} \eta(\mathrm{St}(b_\sigma, \mathrm{sd}K_{\alpha_i})).$$

Here b_σ is the barycenter of the simplex σ and $\mathrm{St}(b_\sigma, \mathrm{sd}K_{\alpha_i})$ is the star of the 0-simplex b_σ in the complex $\mathrm{sd}K_{\alpha_i}$. Then notice that

$$\eta(\mathrm{St}(b_\sigma, \mathrm{sd}K_{\alpha_i})) \cap \eta(\mathrm{St}(b_{\sigma'}, \mathrm{sd}K_{\alpha_i}) \neq \emptyset \Longleftrightarrow \sigma \subset \sigma' \text{ or } \sigma' \subset \sigma.$$

Then we can write

$$\bigcup_{m=0}^{k} \bigcup_{\sigma \in K_{\alpha_i}(m)} \eta(\mathrm{St}(b_\sigma, \mathrm{sd}K_{\alpha_i})) =: \bigcup_{m=0}^{k} \underline{U}_{i,m}.$$

Each $\underline{U}_{i,m}$ is the disjoint union of contractible sets. Moreover, their mutual intersections are also disjoint union of contractible sets. Hence by the Mayer–Vietoris principle, one has that

$$H_s(\underline{U}_i; \mathbb{Z}) = 0, \ \forall s > k.$$

As U_i is homotopy equivalent to $\underline{U}_i$, $H_s(U_i; \mathbb{Z}) = 0$ for $s > k$ as well. □

Before proceeding to the induction step, we need another lemma.

Lemma 5.5.6. For each pair $j < i$, one has

$$U_j \cap N_{\alpha_i}(d) = \pi_{\alpha_i}^{-1}(U_j \cap X_{\alpha_i}) \cap N_{\alpha_i}(d).$$

Proof. It follows from the commutation relations of the control data. More precisely, suppose $x \in U_j \cap N_{\alpha_i}(d)$. Then $\underline{x}(j) := \pi_{\alpha_j}(x) \in \underline{U}_j$. Set $\underline{x}(i) = \pi_{\alpha_i}(x) \in X_{\alpha_i}$. Then

by the commutation relation $\pi_{\alpha_j} \circ \pi_{\alpha_i} = \pi_{\alpha_j}$, one has

$$\pi_{\alpha_j}(\underline{x}(i)) = \pi_{\alpha_j}(x) = \underline{x}(j) \in \underline{U}_j.$$

Moreover, by the relation $\rho_{\alpha_j} \circ \pi_{\alpha_i} = \rho_{\alpha_j}$, one has that

$$\rho_{\alpha_j}(\underline{x}(i)) = \rho_{\alpha_j}(x) < d.$$

Hence $\underline{x}(i) \in U_j$. Therefore one has shown that

$$U_j \cap N_{\alpha_i}(d) \subseteq \pi_{\alpha_i}^{-1}(U_j \cap X_{\alpha_i}) \cap N_{\alpha_i}(d).$$

The inclusion in the other direction is similar. □

Now we assume that we have shown that $U_1 \cup \cdots \cup U_{i-1}$ is open with vanishing integral homology for degrees $s > k$. We would like to construct a "thickening" $U_i^+ \supset U_i$ which is open such that

$$U_1 \cup \cdots \cup U_{i-1} \cup U_i = U_1 \cup \cdots \cup U_{i-1} \cup U_i^+,$$

such that U_i is a deformation retract of U_i^+, and such that

$$H_s((U_1 \cup \cdots \cup U_{i-1}) \cap U_i^+; \mathbb{Z}) = 0 \ \forall s > k-1.$$

Recall that by Lemma 5.5.2, $X_{\alpha_i}^0(d)$ is a manifold with corner. Moreover, each face of $X_{\alpha_i}^0(d)$ is a subcomplex. Then

$$\underline{U}_i \cap \partial X_{\alpha_i}^0(d) = \bigcup \big\{ x \in \eta(\mathrm{Int}\sigma) \mid \sigma \in K_{\alpha_i}(m),\ 0 \leq m \leq k,\ \eta(\mathrm{Int}\sigma) \subset \partial X_{\alpha_i}^0(d) \big\}.$$

Notice that $\partial X_{\alpha_i}^0(d)$ is a topological manifold and the triangulation (K, η) induces

a triangulation of $\partial X^0_{\alpha_i}(d)$. Then each simplex σ appeared in the above union has codimension at most $k-1$ in $\partial X^0_{\alpha_i}(d)$. Hence using the same argument as in the proof of Lemma 5.5.5, one has

$$H_s(\underline{U}_i \cap \partial X^0_{\alpha_i}(d); \mathbb{Z}) = 0 \ \forall s > k-1. \tag{5.5.2}$$

Now we define a thickening. Choose $\epsilon > 0$ sufficiently small. Define $\underline{U}^+_{i,j}$ and deformation retractions

$$\nu_{i,j} : \underline{U}^+_{i,j} \to \underline{U}_i$$

inductively for $j = i, i-1, \ldots, 1$ as follows. First, set

$$\underline{U}^+_{i,i} := \underline{U}_i.$$

Then define

$$\underline{U}^+_{i,j-1} := \underline{U}^+_{i,j} \cup \Big\{ h^{-1}_{\alpha_{j-1}}(x,t) \mid x \in \underline{U}^+_{i,j} \cap \partial N_{\alpha_{j-1}}(d),\ t \in (d-\epsilon, d] \Big\}.$$

Here

$$h_\alpha : N_\alpha \setminus X_\alpha \to S_\alpha(d) \times \mathbb{R}_+$$

is the system of maps induced from the family of lines (for $\epsilon = d$). Intuitively, at the j-th thickening step we extend a little into $X_{\alpha_i} \cap N_{\alpha_j}(d)$. Define

$$\underline{U}^+_i = \underline{U}^+_{i,1}, \qquad\qquad U^+_i := \pi^{-1}_{\alpha_i}(\underline{U}^+_i) \cap N_{\alpha_i}(d)$$

From the construction one can easily derive the following conclusion.

Lemma 5.5.7. 1. $\underline{U}^+_i$ is open in X_{α_i}.

2. $\underline{U}_i$ is a deformation retract of $\underline{U}^+_i$.

3. There holds

$$U_1 \cup \cdots \cup U_i = U_1 \cup \cdots \cup U_{i-1} \cup U_i^+.$$

Corollary 5.5.9. For all $s > k$ there holds

$$H_s(U_1 \cup \cdots \cup U_i; \mathbb{Z}) = 0.$$

Proof. Abbreviate $V_i = U_1 \cup \cdots \cup U_i$. By Mayer–Vietoris, in integral coefficients, one has

$$\cdots \longrightarrow H_s(V_{i-1}) \oplus H_s(U_i^+) \longrightarrow H_s(V_{i-1} \cup U_i^+) \longrightarrow H_{s-1}(V_{i-1} \cap U_i^+) \longrightarrow \cdots.$$

Suppose $s > k$. By the induction hypothesis, one has $H_s(V_{i-1}) = 0$. By the fact that U_i is a deformation retract of U_i^+ and by Lemma 5.5.5, one has $H_s(U_i^+) = H_s(U_i) = 0$.

Claim. $V_{i-1} \cap U_i^+ = \pi_{\alpha_i}^{-1}(V_{i-1} \cap \underline{U}_i^+) = \pi_{\alpha_i}^{-1}(\underline{U}_i^+ \setminus \underline{U}_i)$.

Proof of the claim. The first equality follows from Lemma 5.5.6. To prove the second equality, it suffices to show that $V_{i-1} \cap \underline{U}_i^+ = \underline{U}_i^+ \setminus \underline{U}_i$. Notice that the inclusion $\underline{U}_i^+ \setminus \underline{U}_i \subseteq V_{i-1} \cap \underline{U}_i^+$ is obvious. On the other hand, given $x \in \underline{U}_i^+ \setminus \underline{U}_i$, by the construction there exist a sequence of indices $j_1 < j_2 < \cdots < j_s < i$ $(s \geq 1)$ such that

$$x \in N_{\alpha_{j_a}}(d) \setminus N_{\alpha_{j_a}}(d - \epsilon),\ \forall a = 1, \ldots, s.$$

Define

$$x' = (r_{\alpha_{j_s}} \circ \cdots \circ r_{\alpha_{j_1}})(x) \in \underline{U}_i$$

where r_α is the radial projections provided by the family of lines for $\epsilon = d$. Now by the definition of $\underline{U}_i$, there exist a simplex $\sigma \in K_{\alpha_i}$ and a point $y \in W$ satisfying

$$x' \in \eta(\mathrm{Int}\sigma), \qquad F(y) \in N_{\alpha_i}(d_0) \text{ and } \pi_{\alpha_i}(F(y)) \in \eta(\mathrm{Int}\sigma).$$

Then one has that

$$\pi_{\alpha_{j_s}}(F(y)) = \pi_{\alpha_{j_s}\alpha_i}(\pi_{\alpha_i}(F(y))) = X^0_{\alpha_{j_s}}(d).$$

As the map $\pi_{\alpha_i\alpha_{j_s}} : S_{\alpha_{j_s}}(d) \cap X_{\alpha_i} \to X^0_{\alpha_{j_s}}(d)$ is simplicial, $\eta(\mathrm{Int}\sigma)$ projects to the interior of some simplex $\sigma' \in K_{\alpha_{j_s}}$. Moreover, $\eta(\mathrm{Int}\sigma')$ intersects the image of $\pi_{\alpha_{j_s}} \circ F$. Hence

$$\begin{aligned} x \in \pi^{-1}_{\alpha_i\alpha_{j_s}}(\eta(\mathrm{Int}\sigma')) \cap N_{\alpha_{j_s}}(d) &\subseteq \pi^{-1}_{\alpha_{j_s}}(\eta(\mathrm{Int}\sigma')) \cap N_{\alpha_{j_s}}(d) \\ &\subseteq \pi^{-1}_{\alpha_{j_s}}(\underline{U}_{j_s}) \cap N_{\alpha_{j_s}}(d) \subseteq V_{i-1}. \end{aligned}$$

Therefore $\underline{U}^+_i \setminus \underline{U}_i \subseteq V_{i-1} \cap \underline{U}^+_i$. *End of the proof of the claim.*

Hence

$$H_{s-1}(V_{i-1} \cap U^+_i) = H_{s-1}(\underline{U}^+_i \setminus \underline{U}_i).$$

As $\underline{U}^+_i \setminus \underline{U}_i$ deformation retracts to $\underline{U}_i \cap \partial X^0_{\alpha_i}(d)$, by (5.5.2), one has

$$H_{s-1}(V_{i-1} \cap U^+_i) = 0.$$

Hence $H_s(V_{i-1} \cup U^+_i) = H_s(V_i) = 0$. □

Now Lemma 5.5.1 follows by induction.

Proof of Theorem 5.5.7. Given Lemma 5.5.1, the proof is almost identical as those of [Zin08, Lemma 3.5 & 3.6]. We omit the details. □

5.6 Construction of the integral Euler cycle

In this section we finish our construction of the integral Euler classes. We work under the setting of derived orbifold charts (see Definition 5.1.3). Let $(X, \mathcal{E}, s)$ be a compact

derived orbifold chart (without boundary). Assume that it is normally complex, i.e., X and $\mathcal{E}$ are normally complex (see Definition 5.3.13 and Definition 5.3.14).

5.6.1 Strong transversality

Suppose the pair $(X, \mathcal{E})$ is straightened (see Definition 5.3.12). Then one has the notion of FOP (normally complex) sections of $\mathcal{E}$ (see Definition 5.3.16).

Definition 5.6.1. An FOP section $s : X \to \mathcal{E}$ is called *strongly transverse* at a point $p \in X$ if there is a bundle chart $\hat{C} = (U, E, G, \hat{\psi})$ centered at p satisfying the following conditions. By abuse of notation, let $s : U \to E$ be the pullback of s. With respect to the basic decomposition

$$E \cong \mathring{E}^G \oplus \check{E}^G$$

near the G-fixed locus U^G (which is induced from the straightened structures near U^G), we can write $s = (\mathring{s}, \check{s})$. Notice that both NU^G and $\check{E}^G$ are G-equivariant complex vector bundles. Then in the disk bundle $N_r U^G \subset NU^G$ there exists a smooth normally complex lift (for a sufficiently large d)

$$\check{\mathfrak{s}} : N_r U^G \to \mathrm{Poly}_d^G(NU^G, \check{E}^G)$$

of $\check{s}$ satisfying the following condition. Consider the bundle map

$$\begin{aligned} \tilde{S} : N_r U^G &\to \mathring{E}^G \oplus \big(NU^G \oplus \mathrm{Poly}_d^G(NU^G, \check{E}^G)\big) \\ x &\mapsto \big(\mathring{s}(x), (x, \check{\mathfrak{s}}(x)\big). \end{aligned}$$

Then $\tilde{S}$ is transverse to $\{0\} \oplus \mathcal{Z}_d^G(NU^G, \check{E}^G)$ with respect to the canonical Whitney stratification at the point $0 \in U$.

If s is strongly transverse at every point of a subset $A \subset X$ resp. an open neighborhood of $A \subset X$, then we say that s is strongly transverse along A resp. near

A. If $A = X$, then we say that s is a strongly transverse FOP section.

One can check that the strong transversality condition at a point does not depend on the choice of charts centered at that point nor on the choice of local normal complex lifts.

Next we prove that the set of points where an FOP section is strongly transverse is an open subset of X.

Lemma 5.6.1. (cf. [Par13, Lemma 4.18]) If an FOP section $s : X \to \mathcal{E}$ is strongly transverse at $p \in X$, then there is an open neighborhood $W \subset X$ of p such that s is strongly transverse at each point of W.

Proof. We only need to consider the situation inside a bundle chart $(U, E, G, \hat{\psi})$ centered at p. We also assume that U coincides with the disk bundle $N_r U^G \subset NU^G$. The section s is pulled back to a G-equivariant section of $E \to U$, still denoted by s. We may assume that $s(0) = 0$; otherwise the situation is trivial. Using the straightened structures, we write $s : U \to E$ as $(\mathring{s}, \check{s})$ where $\mathring{s} : U \to \mathring{E}^G$ and $\check{s} : U \to \check{E}^G$ are smooth. Then by the strong transversality condition at p, there is a smooth normally complex lift

$$\check{\mathfrak{s}} : U \to \mathrm{Poly}_d^G(NU^G, \check{E}^G)$$

of $\check{s}$ such that the bundle map

$$\tilde{S} = \mathring{s} \oplus \mathrm{id} \oplus \check{\mathfrak{s}} : U \to \mathring{E}^G \oplus \left(NU^G \oplus \mathrm{Poly}_d^G(NU^G, \check{E}^G)\right)$$

is transverse to $\{0\} \oplus \mathcal{Z}_d^G(NU^G, \check{E}^G)$ at $0 \in U$. Then as the transversality condition to a Whitney stratified set is an open condition (see [Tro78]), $\tilde{S}$ is also transverse to $\{0\} \oplus \mathcal{Z}_d^G(NU^G, \check{E}^G)$ at any point sufficiently closed to $0 \in U$.

Now we consider the strong transversality condition at a nearby point $q \in \psi(U)$. Suppose q is represented by a point $x \in U$ with stabilizer $H \subseteq G$. When $G = H$

and x is sufficiently close to 0, as $\check{\mathfrak{s}}$ is also a local normally complex lift of $\mathfrak{s}$ near q, it follows that the section s is strongly transverse at q. So we assume that $H \subsetneq G$. Now, apply the restriction map θ defined by (5.4.5). Then one obtains a smooth map

$$\underline{\theta} \circ \check{\mathfrak{s}} : N_\epsilon U^H \to \mathrm{Poly}_d^H(NU^H, \check{E}^G).$$

Together with the restriction of $\mathring{s}$ to $N_\epsilon U^H \subset U$ one obtains a section

$$\tilde{s}' : N_\epsilon U^H \to \mathring{E}^G \oplus \mathrm{Poly}_d^H(NU^H, \check{E}^G).$$

Theorem 5.4.13 implies that this map is transverse to $\{0\} \times \mathcal{Z}_d^H(NU^H, \check{E}^G)$ (where we only view $\check{E}^G$ as an H-equivariant bundle) at the point x. As we have the decomposition

$$\check{E}^G \cong (\check{E}^G \cap \mathring{E}^H) \oplus \check{E}^H$$

we may regard $\tilde{s}'$ as a map

$$\tilde{s}' : N_\epsilon U^H \to \mathring{E}^G \oplus \mathrm{Poly}_d^H(NU^H, \check{E}^G \cap \mathring{E}^H) \oplus \mathrm{Poly}_d^H(NU^H, \check{E}^H)$$

whose second component is denoted by $\mathring{\mathfrak{s}} : N_\epsilon U^H \to \mathrm{Poly}_d^H(NU^H, \check{E}^G \cap \mathring{E}^H)$ and whose third component is denoted by $\check{\mathfrak{s}}_x : N_\epsilon U^H \to \mathrm{Poly}_d^H(NU^H, \check{E}^H)$. Define

$$\mathring{s}_x : N_\epsilon U^H \to \mathring{E}^G \oplus (\check{E}^G \cap \mathring{E}^H) \cong \mathring{E}^H$$

by

$$\mathring{s}_x = (\mathring{s}, \mathrm{ev} \circ \mathring{\mathfrak{s}}).$$

Then $\tilde{s}_x = (\mathring{s}_x, \check{\mathfrak{s}}_x)$ is a local normally complex lift of $s : U \to E$ near $x \in U^H \subset U$. Then by Corollary 5.4.12, the graph of $\tilde{s}_x$ is also transverse to $\{0\} \times \mathcal{Z}_d^H(NU^H, \check{E}^H)$. Therefore, by Definition 5.6.1, s is strongly transverse at $q = \psi(x) \in X$. As q runs

over all points in a neighborhood of p, this lemma is proved. □

5.6.2 The integral Euler cycles

Now we state the first main theorem of this chapter. Let $(X, \mathcal{E}, s)$ be a compact normally complex and oriented derived orbifold chart. Fix a reference metric on $\mathcal{E}$. Choose straightened structures on $(\mathcal{E}, X)$ (by Lemma 5.3.6 and Lemma 5.3.9, one can always assume that a precompact neighborhood of $s^{-1}(0)$ is straightened).

Theorem 5.6.2. Given ϵ and a precompact open neighborhood $D \subset X$ of $s^{-1}(0)$, there exists a smooth section $s_\epsilon : X \to \mathcal{E}$ satisfying the following conditions.

1. s_ϵ is an FOP section in a neighborhood of $\overline{D}$ and it is strongly transverse near $\overline{D}$.

2. $\|s_\epsilon - s\|_{C^0(D)} < \epsilon$.

Moreover, when $s_\epsilon^{-1}(0) \cap D$ is compact (in particular, when ϵ is sufficiently small), for each isotropy type γ (represented by (G, V, W)), the intersection $s_\epsilon^{-1}(0) \cap X_\gamma^*$ is a smooth pseudocycle in X of dimension

$$\dim_{\mathbb{R}}(s_\epsilon^{-1}(0) \cap X_\gamma^*) = n_\gamma := \dim_{\mathbb{R}} X - \operatorname{rank}_{\mathbb{R}} \mathcal{E} - \dim_{\mathbb{R}} V + \dim_{\mathbb{R}} W$$

hence (by Theorem 5.5.7) represents a homology class

$$\chi_\gamma^{\mathrm{FOP}}(X, \mathcal{E}, s) := [s_\epsilon^{-1}(0) \cap X_\gamma^*] \in H_{n_\gamma}(X; \mathbb{Z}).$$

We call it the γ-th Fukaya–Ono–Parker–Euler class of the derived orbifold chart $(X, \mathcal{E}, s)$. Moreover, the following items are true.

1. $\chi_\gamma^{\mathrm{FOP}}(X, \mathcal{E}, s)$ is independent of the choice of the precompact open neighborhood D, the (sufficiently small) ϵ, and the smooth section s_ϵ.

2. $\chi_\gamma^{\mathrm{FOP}}(X, \mathcal{E}, s)$ is independent of the choice of straightened structures.

Hence $\chi_\gamma^{\mathrm{FOP}}(X, \mathcal{E}, s)$ only depends on the normal complex structure.

Proof. Proposition 5.6.3, Proposition 5.6.4, and Proposition 5.6.5. □

5.6.3 Existence of strongly transverse perturbations

Now we prove the existence of strongly transverse FOP perturbations. We first specify a sufficiently large cut-off degree of FOP sections. For each isotropy type γ represented by $(G_\gamma, V_\gamma, W_\gamma)$, define

$$d_0(\gamma) \in \mathbb{Z}_{\geq 0}$$

to be the maximum among the d_0 of Lemma 5.4.3, the d_0 of Theorem 5.4.9, and the d_0 of Theorem 5.4.13. For the compact normally complex derived orbifold chart $(X, \mathcal{E}, s)$, define

$$d_0(X, \mathcal{E}, s) = \sup_{p \in s^{-1}(0)} d_0(\gamma_p).$$

By the compactness of $s^{-1}(0)$, this is a finite integer.

We restate the part of Theorem 5.6.2 which concerns the existence of strongly transverse sections.

Proposition 5.6.3. Let $(X, \mathcal{E}, s)$ be an effective, compact, normally complex derived orbifold chart. Suppose $(X, \mathcal{E})$ is normally complex and straightened. Choose an auxiliary metric on $\mathcal{E}$ and an auxiliary precompact open neighborhood D of $s^{-1}(0)$. Then for $d \geq d_0(X, \mathcal{E}, s)$ and for each $\epsilon > 0$, there exists a smooth section $s_\epsilon : X \to \mathcal{E}$ satisfying the following conditions.

1. In an open neighborhood of $\overline{D}$, s_ϵ is an FOP section of degree at most d and is strongly transverse near $\overline{D}$.

2. There holds

$$\|s - s_\epsilon\|_{C^0(D)} < \epsilon.$$

3. When ϵ is small enough, $s_\epsilon^{-1}(0)$ is compact.

The proof is based on an two-layer induction procedure. The basic ingredient for the induction is the following chartwise relative transversality result.

Lemma 5.6.2. For any isotropy type γ the following is true. Let $C = (U, E, G, \hat{\psi})$ be a local chart of E centered at a point $p \in X_\gamma$. Suppose that U can be identified with the G-invariant disk bundle $N_r U^G$ of the normal bundle NU^G of the G-fixed point locus $U^G \subset U$. Choose auxiliary metric on the bundles $\mathring{E}^G$ and for each $d \geq d_0(\gamma)$ an auxiliary metric on $\mathrm{Poly}_d^G(NU^G, \check{E}^G)$.

Let $U' \subset U$ be a G-invariant precompact open subset. Suppose $d \geq d_0(\gamma)$ and

$$\tilde{s} : N_r U^G \to \mathring{E}^G \oplus \mathrm{Poly}_d^G(NU^G, \check{E}^G)$$

is a smooth G-invariant bundle map over U^G whose graph $\tilde{S}$ is transverse to $\{0\} \oplus \mathcal{Z}_d^G(NU^G, \check{E}^G)$ near $\overline{U'} \cap U^G$. Then for any $\epsilon > 0$, there exists another smooth G-invariant bundle map $\tilde{s}' : N_r U^G \to \mathring{E}^G \oplus \mathrm{Poly}_d^G(NU^G, \check{E}^G)$ whose graph is transverse to $\{0\} \oplus \mathcal{Z}_d^G(NU^G, \check{E}^G)$ near U^G which agrees with $\tilde{s}$ near $\overline{U'} \cap U^G$ such that

$$\|\tilde{s} - \tilde{s}'\|_{C^0(U)} < \epsilon.$$

Proof. For any map $\tilde{s} = (\mathring{s}, \check{\mathfrak{s}}) : N_r U^G \to \mathring{E}^G \oplus \mathrm{Poly}_d^G(NU^G, \check{E}^G)$, the graph of $\tilde{s}$ being transverse to $\{0\} \times \mathcal{Z}_d^G(NU^G, \check{E}^G)$ is equivalent to $\mathring{s}$ being transverse to the zero section of $\mathring{E}^G$ and the graph of the restriction of $\check{\mathfrak{s}}$ to $\mathring{s}^{-1}(0)$ is transverse to $\mathcal{Z}_d^G(NU^G, \check{E}^G)$. We first construct a transverse perturbation $\mathring{s}'$ of $\mathring{s}$ which agrees with $\mathring{s}$ near $U^G \cap \overline{U'}$. Indeed, the G-invariance of $\mathring{s}$ implies that the derivatives of $\mathring{s}$ in directions normal to U^G all vanish. Hence $\mathring{s}$ being transverse to the zero section at a point $x \in U^G$ is equivalent to the restriction $\mathring{s}|_{U^G}$ being transverse to the zero section of $\mathring{E}^G|_{U^G}$ at x. Therefore, using the standard transversality argument, one can find a smooth transverse section $\mathring{s}' : U^G \to \mathring{E}^G|_{U^G}$ which agrees with $\mathring{s}$ near $U^G \cap \overline{U'}$. Then

extend $\mathring{s}'$ to a G-invariant section of $\mathring{E}^G$ to U, still denoted by $\mathring{s}'$ such that it agrees with $\mathring{s}$ near $U^G \cap \overline{U'}$. Then $\mathring{s}'$ is transverse near $U^G \cap \overline{U'}$. Moreover, for any $\epsilon > 0$, we may require that

$$\|\mathring{s} - \mathring{s}'\|_{C^0(U)} < \epsilon.$$

Now we only need to consider transverse perturbations of $\check{\mathfrak{s}}$ restricted to $(\mathring{s}')^{-1}(0)$. Therefore, one may just assume that $\mathring{E}^G = 0$. We will consider perturbations of the form

$$\check{s} + \check{f}'$$

where $\check{f}' : N_r U^G \to \mathrm{Poly}_d^G(NU^G, E)$ is the pullback of a smooth map

$$f' : U^G \to \mathrm{Poly}_d^G(NU^G, E).$$

We may further assume without loss of generality that NU^G and E are both trivial with fibers being V and W respectively. Then $\check{s}$ can be regarded as a smooth map

$$\check{s} : U^G \times V \to \mathrm{Poly}_d^G(V, W)$$

and f' is a smooth map $f' : U^G \to \mathrm{Poly}_d^G(V, W)$. Moreover, denote the restriction of $\check{s}$ to $U^G = U^G \times \{0\}$ by

$$f : U^G \to \mathrm{Poly}_d^G(V, W).$$

Claim. Let $Z_\alpha \subset Z_d^G(V, W)$ be a stratum of the canonical Whitney prestratification. The graph of $\check{s} + \check{f}'$ is transverse to $U^G \times Z_\alpha$ at a point $x_0 \in U^G$ if and only if the map $f + f' : U^G \to \mathrm{Poly}_d^G(V, W)$ is transverse to $\mathrm{proj}_2(Z_\alpha)$ at x_0 where

$$\mathrm{proj}_2 : Z_d^G(V, W) \hookrightarrow V \times \mathrm{Poly}_d^G(V, W) \to \mathrm{Poly}_d^G(V, W)$$

is the projection to the second factor.

Proof of the claim. Denote $P_0 = f(x_0)+f'(x_0) \in \mathrm{Poly}_d^G(V,W)$. Suppose $(0,P_0) \in Z_\alpha$. Then the graph of $\check{s}+\check{f}'$ is transverse to $U^G \times Z_\alpha$ at $(x_0,0,P_0)$ means that

$$\left\{\left(\xi, D_x f'(x_0)(\eta) + D_x\check{s}(x_0,0)(\eta) + D_v\check{s}(x_0,0)(\xi)\right) \mid \xi \in V,\ \eta \in T_{x_0}U^G\right\} + T_{(0,P_0)}Z_\alpha$$
$$= V \oplus \mathrm{Poly}_d^G(V,W).$$

Here $D_x\check{s}$ resp. $D_v\check{s}$ is the partial derivative of $\check{s}$ in the U^G resp. V direction. The G-invariance of $\check{s}$ implies that $D_v\check{s}(x_0,0) \equiv 0$. Hence the above implies that

$$(D_x f'(x_0) + D_x f(x_0))(T_{x_0}U^G) + \mathrm{proj}_2(T_{(0,P_0)}Z_\alpha) = \mathrm{Poly}_d^G(V,W),$$

namely, $f+f'$ is transverse to $\mathrm{proj}_2(Z_\alpha)$ at x_0. The implication in the reversed direction is also easy to check. *End of the proof of the claim.*

Now we know that for each stratum $Z_\alpha \subset Z_d^G(V,W)$, f is transverse to $\mathrm{proj}_2(Z_\alpha)$ near $U^G \cap \overline{U'}$. Hence by the standard transversality argument, one can find a map $f' : U^G \to \mathrm{Poly}_d^G(V,W)$ supported away from $U^G \cap \overline{U'}$ such that $f+f'$ is transverse to $\mathrm{proj}_2(Z_\alpha)$ over U^G for all strata Z_α. Therefore, the induced map $\check{s}' = \check{s} + \check{f}'$ has its graph transverse to $U^G \times Z_d^G(V,W)$ with respect to the canonical Whitney stratification near U^G. Moreover, the C^0-norm of f' can be as small as we require. $\square$

Proof of Proposition 5.6.3. By Lemma 5.3.11, we may assume that s is itself an FOP section (of degree at most d). By the compactness of $\overline{D}$, there exists a finite list of isotropy types γ such that $\overline{X_\gamma^*} \cap \overline{D} \neq \emptyset$. Moreover, isotropy types form a partially ordered sets so that

$$X_\delta^* \cap \overline{X_\gamma^*} \neq \emptyset \Longrightarrow \delta \leq \gamma.$$

We order these finitely many isotropy types as

$$\gamma_1, \ldots, \gamma_N$$

such that $\gamma_i \leq \gamma_j \Longrightarrow i \leq j$. For each $n \geq 0$, denote

$$X_{(n)} := \bigcup_{i \leq n} X^*_{\gamma_i}.$$

Fix $d \geq d_0(X, \mathcal{E}, s)$. Fix $\epsilon > 0$ sufficiently small. The induction hypothesis of the first layer of the induction is the following.

Hypothesis A_n. There exists an FOP section $s_n : X \to E$ (of degree at most d) such that s_n is strongly transverse near $X_{(n)} \cap \overline{D}$. Moreover,

$$\|s - s_n\|_{C^0(D)} < n\epsilon.$$

Assume Hypothesis A_{n-1} is true and let $V_{(n-1)} \subset X$ be a precompact open neighborhood of $X_{(n-1)} \cap \overline{D}$ over which s_{n-1} is strongly transverse. Choose another precompact open neighborhood $V'_{(n-1)}$ of $X_{(n-1)} \cap \overline{D}$ whose closure is contained in $V_{(n-1)}$. Then we can find a finite collection of bundle charts

$$\hat{C}_j = (U_j, G_j, E_j, \hat{\psi}_j),\ j = 1, \ldots, k_n$$

and precompact open subsets $U'_j \subset U_j$ satisfying the following conditions.

1. Each $\hat{C}_j$ is centered at a point $x_j \in X^*_{\gamma_n}$.
2. Each U_j is identified with the disk bundle $N_{2r_j} U_j^{G_j}$ and U'_j is the disk bundle $N_{r_j}(U_j^{G_j} \cap U'_j)$.
3. The union of $F'_j := \psi_j(U'_j)$ covers the compact set $(\overline{X^*_{\gamma_n}} \cap \overline{D}) \setminus V'_{(n-1)}$.

The induction hypothesis of the second layer of the induction is the following.

Hypothesis $B_{n,k}$. There is an FOP section $s_{n,k} : X \to E$ (of degree at most d) which

is strongly transverse near

$$X_{(n)} \cap \Big(\overline{V'_{(n-1)}} \cup \bigcup_{1\leq j\leq k} \overline{F'_j}\Big).$$

Moreover,

$$\|s - s_{n,k}\|_{C^0(D)} < (n - 1 + \frac{k}{k_n})\epsilon.$$

Notice that A_N implies (1) and (2) of this proposition and B_{n,k_n} implies A_n. We also declare $B_{n,0} = A_{n-1}$. Therefore, one only needs to prove that $B_{n,k-1}$ implies $B_{n,k}$. Consider the pullback of $s_{n,k-1}$ to the chart $\hat{C}_k$, which can be lifted to a smooth bundle map

$$\tilde{s}_{n,k-1} = (\mathring{s}_{n,k-1}, \check{s}_{n,k-1}) : N_{2r_k}U_k^{G_k} \to \mathring{E}_k^{G_k} \oplus \mathrm{Poly}_d^{G_k}(NU_k^{G_k}, \check{E}_k^{G_k}).$$

By the induction hypothesis $B_{n,k-1}$, the graph of the map $\tilde{s}_{n,k-1}$ is transverse to $\{0\} \times \mathcal{Z}_d^G(NU_k^{G_k}, \check{E}_k^{G_k})$ near the G_k-invariant compact set

$$Q_k := \psi_k^{-1}\left(\overline{V'_{(n-1)}} \cup \overline{F'_1} \cup \cdots \cup \overline{F'_{k-1}}\right) \cap U_k^{G_k} \subset U_k.$$

By Lemma 5.6.2, there exists a smooth G-invariant bundle map

$$\tilde{s}'_{n,k} : N_r U_k^{G_k} \to \mathring{E}_k^{G_k} \oplus \mathrm{Poly}_d^{G_k}(NU_k^{G_k}, \check{E}_k^{G_k})$$

whose graph is transverse to $\{0\} \times \mathcal{Z}_d^{G_k}(NU_k^{G_k}, \check{E}_k^{G_k})$ near $U_k^{G_k}$ and which agrees with $\tilde{s}_{n,k-1}$ near Q_k. Moreover, we may require that

$$\|\tilde{s}'_{n,k} - \tilde{s}_{n,k-1}\|_{C^0(U_k)} \leq \frac{\epsilon}{k_n}.$$

Then using a G_k-invariant cut-off function on U_k supported near Q_k, one obtains a

smooth FOP section $s_{n,k}$ which agrees with $s_{n,k-1}$ near

$$X_{(n)} \cap \Big(\overline{V'_{(n-1)}} \cup \bigcup_{1\leq j\leq k-1} \overline{F'_j}\Big)$$

and hence strongly transverse near

$$X_{(n)} \cap \Big(\overline{V'_{(n-1)}} \cup \bigcup_{1\leq j\leq k} \overline{F'_j}\Big).$$

Moreover,

$$\|s_{n,k} - s_{n,k-1}\|_{C^0(D)} \leq \frac{\epsilon}{k_n}.$$

So $B_{n,k}$ is established. Therefore, the section s_ϵ can be constructed inductively.

Lastly, choose another precompact open neighborhood D' of $s^{-1}(0)$ whose closure is contained in the interior of D. Set

$$\epsilon_0 := \inf_{x\in D\setminus D'} |s(x)| > 0.$$

Then when $\epsilon < \epsilon_0$, one can see that $s_\epsilon^{-1}(0)$ is contained in $\overline{D'}$, hence is compact. □

5.6.4 The pseudocycle property

Recall that (see Example 5.5.1) an orbifold is a Thom–Mather stratified space whose strata can be indexed by isotropy types. Hence one has the notion of pseudocycles in an orbifold (see Definition 5.5.6).

Proposition 5.6.4. Under the assumptions of Theorem 5.5.7. Suppose $s : X \to \mathcal{E}$ is a strongly transverse FOP section with $s^{-1}(0)$ being compact. Then for each isotropy type γ (see Definition 5.3.4), the set

$$s^{-1}(0) \cap X^*_\gamma$$

is a smooth submanifold of X^*_γ of real dimension n_γ (see (5.3.2)). Moreover, the inclusion map $s^{-1}(0) \cap X^*_\gamma \hookrightarrow X$ is a pseudocycle in the sense of Definition 5.5.6.

Proof. We first show that $s^{-1}(0) \cap X^*_\gamma$ is a smooth submanifold of X^*_γ. Indeed, for each $p \in s^{-1}(0) \cap X^*_\gamma$, choose a bundle chart $\hat{C} = (U, G, E, \hat{\psi})$ centered at p. Let $s : U \to E$ be the pullback of s. Then using the straightened structures, we can write $s = (\mathring{s}, \check{s})$. As s is an FOP section, there exists a smooth lift

$$\check{\mathfrak{s}} : U \to \mathrm{Poly}^G_d(NU^G, \check{E}^G)$$

of $\check{s}$. As s is strongly transverse, the graph of the bundle map $\tilde{s} := (\mathring{s}, \check{\mathfrak{s}})$ is transverse to $\{0\} \oplus \mathcal{Z}^G_d(NU^G, \check{E}^G)$. It implies that $\mathring{s}$ is transverse to the zero section of $\mathring{E}^G$ and hence the restriction of $\mathring{s}|_{U^G}$ is transverse to the zero section of $\mathring{E}^G|_{U^G}$. Notice that $\check{s}$ necessarily vanishes along U^G. Hence

$$\psi^{-1}(s^{-1}(0) \cap X^*_\gamma) = \mathring{s}^{-1}(0) \cap U^G$$

which is a smooth submanifold of U^G. This provides a local chart of $s^{-1}(0) \cap X^*_\gamma$. Moreover, as $Z_{d,G} = \{0\} \times \mathrm{Poly}^G_d(V, W)$, one has that

$$\begin{aligned}\dim_{\mathbb{R}}(s^{-1}(0) \cap X^*_\gamma) &= \dim_{\mathbb{R}} U^G - \mathrm{rank}_{\mathbb{R}} \mathring{E}^G + \dim_{\mathbb{R}} \mathrm{Poly}^G_d(V, W) - \dim_{\mathbb{R}} Z_{d,G} \\ &= \dim_{\mathbb{R}} X - \mathrm{rank}_{\mathbb{R}} \mathcal{E} - \dim_{\mathbb{R}} V + \dim_{\mathbb{R}} W = n_\gamma.\end{aligned}$$

Now we prove the pseudocycle property. Suppose γ can be represented by a triple (G, V, W). Consider the boundary of $s^{-1}(0) \cap X^*_\gamma$ (see (5.5.1)). By using the isotropy prestratification on the orbifold X, one can see that

$$\mathrm{Bd}(s^{-1}(0) \cap X^*_\gamma) = \overline{s^{-1}(0) \cap X^*_\gamma} \setminus (s^{-1}(0) \cap X^*_\gamma) \subset \bigcup_{\gamma' \leq \gamma} s^{-1}(0) \cap X^*_{\gamma'}.$$

By the compactness of the above boundary set, there are finitely many bundle charts

$$\hat{C}_i = (U_i, E_i, G_i, \hat{\psi}_i),\ i = 1, \ldots, N$$

centered at $p_i \in \mathrm{Bd}(s^{-1}(0) \cap X^*_\gamma) \cap X^*_{\gamma_i}$ with $\gamma_i \lneq \gamma$ such that

$$\mathrm{Bd}(s^{-1}(0) \cap X^*_\gamma) \subset \bigcup_{i=1}^{N} \psi_i(U_i).$$

Moreover there is an integer d such that over each U_i, the pullback of s, denoted by $s_i : U_i = N_r U_i^{G_i} \to E_i$, has a normally complex lift

$$\tilde{s}_i = (\mathring{s}_i, \check{\mathfrak{s}}_i) : N_r U_i^{G_i} \to \mathring{E}_i^{G_i} \oplus \mathrm{Poly}_d^{G_i}(NU_i^{G_i}, \check{E}_i^{G_i}).$$

Let the fiber of $NU_i^{G_i}$ at 0 be V_i and the fiber of $\check{E}_i^{G_i}$ at 0 be W_i. Then (G_i, V_i, W_i) is a representative of the isotropy type γ_i. Define

$$T_i := \left\{(G_i', V_i', W_i') \mid G_i' \subset G_i,\ V_i' = \check{V}_i^{G_i'},\ W_i' = \check{W}_i^{G_i'},\ \gamma = [G_i', V_i', W_i']\right\}.$$

As x_i lies in the boundary of $s^{-1}(0) \cap X^*_\gamma$, T_i is nonempty. Consider

$$\bigcup_{(G_i', V_i', W_i') \in T_i} \mathrm{Bd} Z^*_{d, G_i'} \subset Z_d^{G_i}(V_i, W_i).$$

Now as $Z^*_{d,G_i'}$ is the union of strata associated to the canonical Whitney stratification on $Z_d^{G_i}(V_i, W_i)$ and s is strongly transverse, one can see that the associated map

$$\tilde{S}_i = (\mathring{s}_i, \mathrm{graph}\check{\mathfrak{s}}_i) : N_r U_i^{G_i} \to \mathring{E}_i^{G_i} \oplus (NU_i^{G_i} \oplus \mathrm{Poly}_d^{G_i}(NU_i^{G_i}, \check{E}_i^{G_i}))$$

is transverse to $\{0\} \oplus \mathcal{Z}^*_{d,G_i'}(NU_i^{G_i}, \check{E}_i^{G_i})$. Moreover,

$$\psi_i^{-1}(s^{-1}(0) \cap X_\gamma^*) = \bigcup_{(G_i',V_i',W_i') \in T_i} \tilde{S}_i^{-1}\left(\{0\} \oplus \mathcal{Z}^*_{d,G_i'}(NU_i^{G_i}, \check{E}_i^{G_i})\right)$$

and

$$\dim_{\mathbb{R}}(s^{-1}(0) \cap X_\gamma^*) = \dim_{\mathbb{R}} U_i^{G_i} - \mathrm{rank}\mathring{E}_i^{G_i} + \dim_{\mathbb{R}} Z^*_{d,G_i'} - \dim_{\mathbb{R}} \mathrm{Poly}_d^{G_i}(V_i, W_i). \quad (5.6.1)$$

On the other hand, as $Z^*_{d,G_i'}$ is the union of strata, there exists algebraic submanifolds

$$Z_{i,\alpha} \subset Z_d^{G_i}(V_i, W_i) \text{ s.t. } \bigcup_{(G_i',V_i',W_i') \in T_i} \mathrm{Bd} Z^*_{d,G_i'} \subset \bigcup_\alpha Z_{i,\alpha}.$$

Each $Z_{i,\alpha}$ is also a stratum of $Z_d^{G_i}(V_i, W_i)$. Moreover, as the canonical Whitney stratification is *complex* algebraic, one has

$$\dim_{\mathbb{R}} Z_{i,\alpha} \leq \dim_{\mathbb{R}} Z^*_{d_i,G_i'} - 2.$$

Consider

$$W_{i,\alpha} := \tilde{S}_i^{-1}\left(\{0\} \oplus Z_{i,\alpha}\right) \subset U_i.$$

As s is strongly transverse, the above set is a smooth manifold. Moreover, by (5.6.1)

$$\begin{aligned} \dim_{\mathbb{R}} W_{i,\alpha} &= \dim_{\mathbb{R}} U_i^{G_i} - \mathrm{rank}_{\mathbb{R}} \mathring{E}_i^{G_i} + \dim_{\mathbb{R}} Z_{i,\alpha} - \dim_{\mathbb{R}} \mathrm{Poly}_d^{G_i}(V_i, W_i) \\ &\leq \dim_{\mathbb{R}} U_i^{G_i} - \mathrm{rank}_{\mathbb{R}} \mathring{E}_i^{G_i} + \dim_{\mathbb{R}} Z^*_{d_i,G_i'} - \dim_{\mathbb{R}} \mathrm{Poly}_d^{G_i}(V_i, W_i) - 2 \quad (5.6.2) \\ &= \dim_{\mathbb{R}}(s^{-1}(0) \cap X_\gamma^*) - 2. \end{aligned}$$

Define

$$f_{i,\alpha} : W_{i,\alpha} \to M, \ f_{i,\alpha} := \psi_i|_{W_{i,\alpha}}.$$

Then $f_{i,\alpha}$ is smooth (in the sense of Definition 5.5.4) and

$$\mathrm{Bd}(s^{-1}(0) \cap X^*_\gamma) \subset \bigcup_{1 \le i \le N} \bigcup_{\alpha} f_{i,\alpha}(W_{i,\alpha}).$$

Then by (5.6.2) and the definition of pseudocycle (Definition 5.5.6), we have shown that $s^{-1}(0) \cap X^*_\gamma$ is indeed a pseudocycle. □

5.6.5 Invariance of the Euler cycles

We prove that the pseudocycle up to cobordism is independent of the choice of auxiliary structures and the choice of strongly transverse perturbations.

Proposition 5.6.5. Let $(X, \mathcal{E}, s)$ be the same as in Proposition 5.6.3. Choose an auxiliary metric on $\mathcal{E}$ and an auxiliary precompact open neighborhood $D \subset X$ of $s^{-1}(0)$. Then there exists $\epsilon > 0$ satisfying the following conditions.

Suppose X and $\mathcal{E}$ are endowed with two possibly different straightened structures. Let s_0 and s_1 be two different strongly transverse FOP sections with respect to these two straightened structures of level d_0 and d_1 respectively with $d_0, d_1 \ge d_0(X, \mathcal{E}, s)$, such that

$$\|s_i - s\|_{C^0(D)} < \epsilon, \text{ for } i = 0, 1.$$

Then for each isotropy type γ of X, there is a pseudocycle cobordism between $s_0^{-1}(0) \cap X^*_\gamma$ and $s_1^{-1}(0) \cap X^*_\gamma$.

Proof. First one can construct a concordance between the two sets of straightened structures. Let $(g_i^{TX}, \nabla_i^{\mathcal{E}})$, $i = 0, 1$ be the two straightened structures. Then by Lemma 5.3.7, one can construct a Riemannian metric $g^{T\tilde{X}}$ on $\tilde{X} := X \times [0, 1]$ which is straightened near $\overline{D} \times [0, 1]$ and which is of product type near $X \times \{0, 1\}$, whose restriction to $X \times \{i\}$ is g_i^{TX} for $i = 0, 1$. Then by Lemma 5.3.9, one can find a connection on the product $\tilde{\mathcal{E}} := \mathcal{E} \times [0, 1] \to X \times [0, 1]$ which is straightened near

$\overline{D} \times [0,1]$ with respect to $g^{T\tilde{X}}$ and which is of product type near $X \times [0,1]$, whose restriction to $X \times \{i\}$ is $\nabla_i^{\mathcal{E}}$. One can then interpolate between s_0 and s_1 to obtain an FOP section $\tilde{s}$ of $\tilde{\mathcal{E}}$ (with respect to the product normal complex structure and the concordance straightened structure we just constructed). Moreover, as both s_0 and s_1 are C^0-close to s over D, by the proof of Lemma 5.3.11, one can require that

$$\|\tilde{s} - \tilde{s}^*\|_{C^0(D\times[0,1])} < 3\epsilon.$$

Here $\tilde{s}^*$ is the pullback of the original smooth section s to $\tilde{X}$. Notice that $\tilde{s}$ is already strongly transverse near $\overline{D} \times \{0,1\}$. Then by the proof of Proposition 5.6.3, one can find a strongly transverse FOP perturbation $\tilde{s}'$ of $\tilde{s}$ which agrees with $\tilde{s}$ near $\overline{D}\times\{0,1\}$ and $\|\tilde{s} - \tilde{s}'\|_{C^0(D\times[0,1])} < \epsilon$. As ϵ is small, the zero locus of $\tilde{s}'$ is still compact. Then similar to the proof of the pseudocycle property (Proposition 5.6.4), for each isotropy type γ, one can see that

$$(\tilde{s}')^{-1}(0) \cap \tilde{X}^*_\gamma$$

is a cobordism of pseudocycles between $s_0^{-1}(0) \cap X^*_\gamma$ and $s_1^{-1}(0) \cap X^*_\gamma$. □

5.7 Stable complex derived orbifold bordism

We explain how the previous constructions can be upgraded to study the *stable complex derived orbifold bordisms*. This section is algebro-topological in nature. Moreover, we will not restrict ourselves to effective orbifolds but also allow *orbispaces*. Namely, we study certain (generalized) homology theory of topological stacks locally modeled on the action groupoid of a not necessarily effective finite group acting on a topological space. Our expositions closely follow [Par19, Par20].

5.7.1 Orbispaces

Denote by **Top** the category of topological spaces with morphisms given by continuous maps. Let **Grpd** be the 2-category of essentially small groupoids.

Definition 5.7.1. A functor $F : \mathbf{Top}^{\mathrm{op}} \to \mathbf{Grpd}$ is called a *stack* if it satisfies descent. Namely, for any topological space U and any open cover $U = \cup_i U_i$, the natural functor

$$F(U) \to \mathrm{Eq}\ (\prod_i F(U_i) \rightrightarrows \prod_{i,j} F(U_i \cap U_j) \mathrel{\substack{\rightarrow\\[-1ex]\rightarrow\\[-1ex]\rightarrow}} \prod_{i,j,k} F(U_i \cap U_j \cap U_k)) \tag{5.7.1}$$

is an equivalence, where Eq is the equalizer. Denote by $\mathrm{Shv}(\mathbf{Top}, \mathbf{Grpd})$ the 2-category of stacks, with morphisms given by natural transformations. This 2-category has a natural monoidal structure by taking products in the obvious way.

The Yoneda functor $\mathbf{Top} \to \mathrm{Shv}(\mathbf{Top}, \mathbf{Grpd})$ is continuous (commutes with limits), and it admits a left adjoint $|\cdot| : \mathrm{Shv}(\mathbf{Top}, \mathbf{Grpd}) \to \mathbf{Top}$ called passing to the *coarse space*. For a topological space X, one can identify it with the associated stack $\mathrm{Hom}(-, X)$. Similarly, for $F \in \mathrm{Shv}(\mathbf{Top}, \mathbf{Grpd})$, we do not distinguish $F(X)$ and natural transformations $\mathrm{Hom}(-, X) \to F(-)$.

Let $f : X \to Y$ be a continuous map between topological spaces. It is said to *admit local sections* if there exists an open cover $Y = \cup_i U_i$ such that for each restriction $f|_{f^{-1}(U_i)} : f^{-1}(U_i) \to U_i$ there is $s_i : U_i \to f^{-1}(U_i)$ continuous with $f|_{f^{-1}(U_i)} \circ s_i = id|_{U_i}$. A morphism between stacks $F_1 \to F_2$ admits local sections if for every topological space X and morphism $X \to F_2$, the induced map $F_1 \times_{F_2} X \to X$ admits local sections. Using this process of "base change", one can similarly define when a morphism $F_1 \to F_2$ is injective, surjective, closed, open, étale, separated, or proper.

Definition 5.7.2. A stack is *representable* if it is in the essential image of the Yoneda embedding, i.e. it is isomorphic to a topological space. A morphism between two

stacks $F_1 \to F_2$ is called representable if for every map $X \to F_2$ where X is a topological space, the fiber product $F_1 \times_{F_2} X$ is representable.

Definition 5.7.3. A stack F is called *topological* if there exists a representable map $U \to F$ admitting local sections where U is a topological space, and such a map is called an *atlas*.

A *point* of a stack F is defined to be a map $* \to F$, or equivalently an object of the groupoid $F(*)$ where $*$ is the topological space of one single point. The *isotropy group/stabilizer* of a point is defined to be the automorphism group of the corresponding object in $F(*)$, i.e. the points of the fiber product $* \times_F *$. In particular, the isotropy groups have a natural topological group structure if F is topological.

A standard example of topological stacks is the action groupoid. Suppose V is a topological space and G is a topological group which acts continuously on V. Then the action groupoid $G \times V \rightrightarrows V$ associates a topological space X the groupoid whose objects are given by open covers $X = \cup_i U_i$ together with maps $f_i : U_i \to V$ and $f_{ij} : U_i \cap U_j \to V \times G$ such that $f_i|_{U_i \cap U_j} = \pi_V \circ f_{ij}$ and $f_j|_{U_i \cap U_j} = \pi_G \circ f_{ij} \cdot \pi_V \circ f_{ij}$. An (iso)morphism in $[G \times V \rightrightarrows V](X)$ consists of maps $U_i \cap U_i'$ satisfying certain compatibility conditions for two covers $X = \cup_i U_i$ and $X = \cup_i U_i'$. The coarse space of $G \times V \rightrightarrows V$ is V/G and the stabilizers are the stabilizers of the G-action on V in the usual sense. The action groupoid $G \times V \rightrightarrows V$ is also denoted by V/G.

Definition 5.7.4. A stack F is a called *separated orbispace* if F admits an étale atlas $U \to F$ and the diagonal map $F \to F \times F$ is separated and proper.

In the following discussions, we work with separated orbispaces exclusively without explicitly mentioning separability. We will use the letter X instead of F to denote orbispaces. One equivalent characterization of a topological stack X being an orbispace is to require the coarse space $|X|$ to be Hausdorff, and the existence of a cover of X by open substacks of the form V/G where G is a finite discrete group acting

continuously on a Hausdorff space V. Note that the stabilizers of orbispaces are all finite. If X_1 and X_2 are orbispaces, a map $X_1 \to X_2$ is representable if and only if it induces an injection on stabilizers [Par19, Corollary 3.6]. We will use this assertion as the working definition of the representability of maps between orbispaces.

Definition 5.7.5. 1. An orbispace X is called a *topological orbifold* if there is an étale atlas $U \to F$ such that U is locally homeomorphic to $\mathbb{R}^n$ and n is called the dimension of X.

2. A *smooth structure* on a topological orbifold X is a choice of étale atlas $U \to X$ and a smooth structure on U such that the two smooth structures on $U \times_X U$ obtained by the two pullbacks $U \times_X U \to U$ coincide.

3. A *smooth orbifold* is a topological orbifold X with a smooth structure.

It follows from Definition 5.7.5 that smooth orbifolds are locally modeled on U/G where $U \subset \mathbb{R}^n$ is an open subset (equipped with the induced smooth structure) and G acts on U smoothly. Note that the smoothness of the group action is equivalent to its linearity. One can also talk about topological/smooth orbifolds with boundary by allowing étale charts locally modeled on open subsets in $\mathbb{R}^n_{x_n \geq 0}$.

It is a standard fact that the notion of smooth orbifolds from Definition 5.7.5 and the one from Section 5.3 if we assume that G acts faithfully on each étale local chart, see [MP97, Theorem 4.1]. Similarly, we can define vector bundles over orbispaces (actually, the same definition works for stacks but we stick with orbispaces for concreteness) without using local charts explicitly.

Definition 5.7.6. A real or complex vector bundle over an orbispace X is a representable map $E \to X$ together with maps $\mathbb{R}$ or $\mathbb{C} \times E \to E$ and $E \times_X E \to X$, such that the pullback to any topological space defines a vector bundle and the two maps coincide with fiberwise scalar multiplication and addition.

A section of a vector bundle $\pi : E \to X$ is then defined to be a map (not necessarily representable!) $s : X \to E$ such that $\pi \circ s$ agrees with the identity morphism on X. When X is a smooth orbifold, a vector bundle $E \to X$ is smooth if its pullback to a smooth étale atlas $U \to X$ defines a smooth vector bundle on U. The smoothness of a section $s : X \to E$ is also tested on any smooth étale atlas. When we present an effective orbifold using orbifold charts, these notions coincide with the ones in Section 5.3. When X is paracompact, any vector bundle $E \to X$ admits an Euclidean/Hermitian inner product. When X is a smooth orbifold, its tangent bundle TX is characterized by requiring the pullback of TX to every étale local chart to give the tangent bundle of the chart.

A principal H-bundle over X is a representable map $P \to X$ along with a map $H \times P \to P$ such that its pullback to any topological space defines a principle H-bundle with the structural H-action.

Theorem 5.7.7 ([Par19, Theorem 1.1]). Suppose X is a compact orbispace, i.e. the coarse space $|X|$ is a compact topological space. Then there exists a complex vector bundle $V \to X$, whose fiber over $x \in X$ (an object in $X(*)$ using the previous notation) is isomorphic to a direct sum of copies of the regular representation of the isotropy group G_x.

An important consequence of the above theorem is, any compact orbispace is equivalent to the action groupoid of a compact Lie group, by taking the the unitary frame bundle of a vector bundle as above endowed with the natural action of a unitary group. Note that the unit disc bundle of V is an *effective* orbifold with boundary. This observation will be used later to connect our study of integral Euler classes on effective orbifolds to derived orbifold bordisms.

5.7.2 Stable complex structures

Definition 5.7.8. Let X be an orbispace. A *stable vector bundle* on X is a formal difference $E_1 - E_2$ where $E_i \to X$ is a vector bundle over X for $i = 1, 2$. A pair of stable vector bundles $E_1 - E_2$ and $E_1' - E_2'$ are isomorphic if there exists a vector bundle E and an isomorphism

$$E_1 \oplus E_2' \oplus E \xrightarrow{\sim} E_1' \oplus E_2 \oplus E. \tag{5.7.2}$$

Stable vector bundles over X and isomorphisms of them form a groupoid. When X is a topological space, Definition 5.7.8 coincides with the usual notion of stable vector bundles.

Definition 5.7.9. Let $E \to X$ be a real vector bundle. A *complex structure* on E consists of a complex vector bundle $E' \to X$ along with an isomorphism $E \xrightarrow{\sim} E'$.

Remark 5.7.10. Given a finite group G, denote by $\mathbf{B}G$ the action groupoid $*/G$. Then *orbi-CW-complexes* are generalizations of CW-complexes with local building blocks $D^k \times \mathbf{B}G$, and the attaching maps are required to be representable, see [Par20, Section 2.4]. Given $n \in \mathbb{Z}_{\geq 0}$, there is an *orbispace* $\mathbf{B}O(n)$ [Par20, Section 4.1] such that isomorphism classes of real vector bundles over an *orbi-CW-complex* X are in bijection with homotopy classes of maps from X to $\mathbf{B}O(n)$, by pulling back the tautological bundle over $\mathbf{B}O(n)$. A similar construction holds for the unitary group with tautological complex bundles $\xi_n \to \mathbf{B}U(n)$. There is a natural map between orbispaces $\mathbf{B}U(n) \to \mathbf{B}O(2n)$. Using these notions, a complex structure on E^{2n} is equivalently defined by a map $f : X \to \mathbf{B}U(n)$ together with an isomorphism $E \xrightarrow{\sim} f^*\xi_n$. Note that both the construction of $\mathbf{B}O(n)$ and the classification result are based on obstruction theory, so this homotopical definition works well only if X is an orbi-CW-complex. On the contrary, Definition 5.7.9 makes sense for arbitrary orbispaces, or more generally, topological stacks.

Definition 5.7.11. 1. A *stable complex structure* on a stable vector bundle $E_1 - E_2$ over X consists of a vector bundle $E \to X$ together with complex structures on $E_1 \oplus E$ and $E_2 \oplus E$ if the difference of the ranks of E_1 and E_2 is even; if such difference is odd, a stable complex structure on $E_1 - E_2$ consists of a vector bundle $E \to X$ and complex structures on $(E_1 \oplus \underline{\mathbb{R}}) \oplus E$ and $E_2 \oplus E$.

2. Suppose $E' \to X$ is a complex vector bundle. Then the *E'-stabilization* of a (stable) complex stable vector bundle $E_1 - E_2$ is defined to be the stable vector bundle $(E_1 \oplus E') - (E_2 \oplus E')$ equipped with the obvious stable complex structure.

3. Consider two stable complex structures on $E_1 - E_2$ given by vector bundles $E^{(i)}$ and complex structures on $E_1 \oplus E^{(i)}$ and $E_2 \oplus E^{(i)}$, $i = 1, 2$. We say that the two stable complex structures are isomorphic if there exist complex vector bundles $F^{(1)}, F^{(2)} \to X$ such that there are isomorphisms $E_1 \oplus E^{(1)} \oplus F^{(1)} \cong E_1 \oplus E^{(2)} \oplus F^{(2)}$ and $E_2 \oplus E^{(1)} \oplus F^{(1)} \cong E_2 \oplus E^{(2)} \oplus F^{(2)}$ which intertwine with the complex structures.

Remark 5.7.12. When the stable vector bundle is actually a vector bundle, i.e. E_2 is trivial, a stable complex structure on E_1 consists of a vector bundle $E \to X$, a complex structure on $E_1 \oplus E$, and a complex structure on E.

Remark 5.7.13. It is important to note that the stabilization in Definition 5.7.11 could be *arbitrary* vector bundles over X. This is in contrast with the notion of (stable) complex structures of vector bundles over topological spaces, where the stabilizations are of the form $\oplus\, \underline{\mathbb{R}}^k$. The latter definition (called *coarsely stable complex structure* by [Par20, section3.3]) is stronger than the former in general. For the construction of stable complex structures on (virtual) tangent bundles of Gromov–Witten type moduli spaces, it is crucial to use Definition 5.7.11.

5.7.3 Derived orbifolds and bordism groups

We discuss bordism-type constructions over orbispaces following [Par20]. Similar ideas probably date back to the work of Joyce [Joy07, Section 5]. Recall that the notion of derived orbifold chart is defined in Definition 5.1.3. Given a derived orbifold chart (D, E, s), define its *virtual dimension* to be $\dim D - \mathrm{rank} E$.

Definition 5.7.14. A stable complex structure on a derived orbifold chart (D, E, s) is an isomorphism class of stable complex structures on the stable vector bundle $TD - E$ over D.

Remark 5.7.15. If the virtual dimension of a derived orbifold chart (D, E, s) is odd and it is endowed with a stable complex structure, we can find an orbibundle E' over D such that the tangent bundle of the total space $E' \oplus \underline{\mathbb{R}}$ admits a complex structure. Due to the triviality of the factor $\underline{\mathbb{R}}$, the total space E', accordingly its open unit disc bundle with respect to any metric, is a *normally complex* orbifold (see Definition 5.3.13) if E' is an effective orbifold.

Given a derived orbifold chart (D, E, s) and an open subset $U \subset D$ with $s^{-1}(0) \subset U$, the *restriction* of (D, E, s) to U is given by the derived orbifold chart $(U, E|_U, s|_U)$. If $E' \to D$ is a vector bundle, the *stabilization* of (D, E, s) by E' is defined to be $(E', E \oplus E', s \oplus id)$. After equipping E' with a metric, we can also take the open unit disc bundle of E' to get a derived orbifold chart such that the ambient space is precompact, provided D is compact. If two derived orbifold charts are isomorphic after restriction and stabilization, we say they are *equivalent* to each other. Furthermore, a stable complex structure on (D, E, s) also naturally induces a stable complex structure on any restriction or any stabilization.

Proposition 5.7.16. Suppose (D, E, s) is a compact and stable complex derived orbifold chart. Then up to a restriction of the derived orbifold chart, there is a stabilization (D', E', s') together with complex structures on TD' (if the virtual dimension is

even, otherwise a complex structure on $TD' \oplus \underline{\mathbb{R}}$) and E' such that D' is an effective orbifold and that the stable complex structure on (D', E', s') induced from the complex structures coincides with the stable complex structure induced from that on (D, E, s).

Proof. By restricting to a precompact open subset of D containing the compact orbispace $s^{-1}(0)$ if necessary, we can stabilize (D, E, s) by a complex vector bundle from Theorem 5.7.7 such that the ambient orbifold is effective. Using Definition 5.7.8, we can take a further stabilization to obtain a derived orbifold chart (D', E', s') so that TD' and E' are endowed with complex structures. □

Remark 5.7.17. Applying restrictions if necessary, any two different effective and (almost) complex derived orbifold charts constructed from Proposition 5.7.16 admit a common stabilization compatible with the complex structures, see the proof of [Par20, Proposition 5.1], which uses the enough vector bundles Theorem 5.7.7.

Let X be a compact orbispace. Let us consider a derived orbifold chart (D, E, s) together with a map $f : D \to X$ (in reality, f is just required to be defined over a neighborhood of $s^{-1}(0) \subset D$ and we can apply restriction). A *bordism* between two such pairs (D_1, E_1, s_1, f_1) and (D_2, E_2, s_2, f_2) is given by a derived orbifold chart with boundary $(W, \tilde{E}, \tilde{s})$, a codimension-0 embedding $D_1 \coprod D_2 \hookrightarrow \partial W$, and a map $\tilde{f} : W \to X$ such that the restriction of the corresponding data along D_i agrees with (D_i, E_i, s_i, f_i) for $i = 1, 2$. When the boundaries (D_i, E_i, s_i) have stable complex structures, the bordism $(W, \tilde{E}, \tilde{s}, \tilde{f})$ is called a *stable complex bordism* if there exists a stable complex structure on $(TW \oplus \underline{\mathbb{R}}) - \tilde{E}$ whose restriction to D_i is stably isomorphic to the given stable complex structures on $TD_i \oplus \underline{\mathbb{C}} - E_i$ for $i = 1, 2$. The factor $\underline{\mathbb{C}}$ comes from the direct sum of the normal bundle (with the inward-pointing orientation) of $D_i \hookrightarrow W$ and the trivial real vector bundle $\underline{\mathbb{R}}$.

Definition 5.7.18. 1. Suppose X is an orbispace. The *stable complex derived orbifold bordism* of X, written as $\overline{\Omega}_*^{\mathbb{C},\mathrm{der}}(X)$, is the set consists of compact stable complex derived orbifold charts (D, E, s) together with a map $f : D \to X$ modulo equivalence of derived orbifold charts and compact stable complex bordism. Here if $\pi : E' \to D$ is a vector bundle, the induced map from E' to X of the derived orbifold chart $(E', E \oplus E', s \oplus id)$ is $f \circ \pi$.

2. The *representable stable complex derived orbifold bordism* $\Omega_*^{\mathbb{C},\mathrm{der}}(X)$ is defined similarly as $\overline{\Omega}_*^{\mathbb{C},\mathrm{der}}(X)$ with the additional requirements that the map $f : D \to X$ and its extension over the bordism $\tilde{f} : W \to X$ are both representable. Both $\Omega_*^{\mathbb{C},\mathrm{der}}(X)$ and $\overline{\Omega}_*^{\mathbb{C},\mathrm{der}}(X)$ have an abelian group structure under disjoint union.

In both cases, (D, E, s, f) defines an element in $\overline{\Omega}_k^{\mathbb{C},\mathrm{der}}(X)$ or $\Omega_k^{\mathbb{C},\mathrm{der}}(X)$ if the virtual dimension of (D, E, s) is k.

Remark 5.7.19. Given a derived orbifold chart (D, E, s), if $E = 0$ (which forces $s = 0$) and D is a manifold, a stable complex structure on (D, E, s) is equivalent to a stable complex structure on TD and a complex structure on the stabilizing vector bundle (modulo an $\underline{\mathbb{R}}$-factor if D is odd-dimensional). After forgetting the latter structure, the tautological map $D \to *$ defines an element in the classical stable complex bordism group.

Remark 5.7.20. Denote by OrbSpc the category of orbispaces and (homotopy classes of) maps between them, and define RepOrbSpc to be the category of orbispaces with (homotopy classes of) representable maps. Then there is a natural faithful functor RepOrbSpc $\to$ OrbSpc which admits a right adjoint

$$R : \mathrm{OrbSpc} \to \mathrm{RepOrbSpc}, \tag{5.7.3}$$

as proved in [Par20, Proposition 3.13]. In other words, for any orbispace X, we have

an isomorphism

$$\overline{\Omega}_*^{\mathbb{C},\mathrm{der}}(X) \cong \Omega_*^{\mathbb{C},\mathrm{der}}(R(X)). \tag{5.7.4}$$

From the point view of algebraic topology, it is more natural to study $\Omega_*^{\mathbb{C},\mathrm{der}}(X)$ and the not necessarily representable version $\overline{\Omega}_*^{\mathbb{C},\mathrm{der}}(X)$ can be recovered from the former via the isomorphism (5.7.4). However, to construct integer-valued Gromov–Witten type invariants, especially without the insertion from $H^*(\overline{\mathcal{M}}_{g,n})$, it is better to go through $\overline{\Omega}_*^{\mathbb{C},\mathrm{der}}(X)$, as what we did in Theorem 5.1.7.

5.7.4 Integral homology classes from stable complex derived orbifold bordism

We show how to use normally complex perturbations to extract integral homology classes from stable complex derived orbifold bordism. To formulate our result, we introduce the notion of *stabilized isotropy types*.

Definition 5.7.21. Define an equivalence relation on the set of isotropy types induced from the following relation

$$(G, V, W) \sim (G, V \oplus R, W \oplus R)$$

where R is a finite-dimensional complex G-representation which does not contain trivial G-representations. The set of equivalence classes Γ, whose elements are denoted by $[\gamma] = [G, V, W]$, is called the set of *stabilized isotropy types*.

Suppose we have a compact derived orbifold chart (D, E, s) such that D is effective and normally complex, and E is endowed with a complex structure. Note that a complex structure on E naturally defines a normal complex structure, see Definition 5.3.14. Then by Theorem 5.6.2, for each isotropy type γ, one has the Fukaya–Ono–

Parker–Euler class

$$\chi_\gamma^{\mathrm{FOP}}(D,E,s) \in H_{n_\gamma}(D;\mathbb{Z}).$$

Then for a stabilized isotropy type $[\gamma] \in \Gamma$, define

$$\chi_{[\gamma]}^{\mathrm{FOP}}(D,E,s) := \sum_{\gamma\in[\gamma]} \chi_\gamma^{\mathrm{FOP}}(D,E,s). \tag{5.7.5}$$

Now we state and prove the main application of the construction of the integral Euler class in algebraic topology.

Theorem 5.7.22. Suppose Y is an orbispace. Then the map $\chi_{[\gamma]}^{\mathrm{FOP}}$ from (5.7.5) can be extended to a well-defined map

$$\mathcal{FOP}_{[\gamma]} : \overline{\Omega}_*^{\mathbb{C},\mathrm{der}}(Y) \to H_*(|Y|;\mathbb{Z}). \tag{5.7.6}$$

Proof. We present the proof for derived orbifold charts of even virtual dimensions. When the virtual dimension is odd, we work with the normal complex structure as from Remark 5.7.15 and the proof proceeds in exactly the same way.

To define the map (5.7.6), suppose (D,E,s) is a compact stable complex derived orbifold chart with a map $f : D \to Y$. By Proposition 5.7.16, there is a compact derived orbifold chart (D_1,E_1,s_1) obtained from stabilizing (D,E,s) via a vector bundle $\pi_1 : F_1 \to D$ such that $D_1 = F_1$ is effective and almost complex, and $E_1 = \pi_1^*TD \oplus \pi_1^*F_1$ is equipped with a complex structure. In this proof we call such a stabilization a *good* stabilization. Define $f_1 : D_1 \to Y$ to be the lift of f. Then (D,E,s,f) and (D_1,E_1,s_1,f_1) represent the same element of $\overline{\Omega}_*^{\mathbb{C},\mathrm{der}}(Y)$. Then we define

$$\widetilde{\mathcal{FOP}}_{[\gamma]}(D,E,s,f) := (f_1)_*\big(\chi_{[\gamma]}^{\mathrm{FOP}}(D_1,E_1,s_1)\big). \tag{5.7.7}$$

First, we show that the right hand side of (5.7.7) is invariant under restrictions of (D_1,E_1,s_1,f_1). Indeed, if $U_1 \subset D_1$ is an open neighborhood of $s_1^{-1}(0)$ and s_1' is

a strongly transverse FOP section of $E_1 \to D_1$ which is sufficiently close to s_1, then $(s_1')^{-1}(0)$ is also contained in U_1. Hence obviously

$$(f_1) * \big(\chi^{\mathrm{FOP}}_{[\gamma]}(D_1, E_1, s_1)\big) = (f_1|_{U_1})_*\big(\chi^{\mathrm{FOP}}_{[\gamma]}(U_1, E_1|_{U_1}, s_1|_{U_1})\big).$$

Second, we show that the association (5.7.7) is independent of the choice of the good stabilization (D_1, E_1, s_1, f_1). We first consider the case when (D_2, E_2, s_2, f_2) is a further stabilization of (D_1, E_1, s_1, f_2) by a complex vector bundle $\pi_2 : F_2 \to D_1$. We may assume that D_1 is indeed a disk bundle of $F_1 \to D$ and $s_1 : D_1 \to E_1$ is already a strongly transverse FOP section with compact zero locus. Then it is obvious that $\pi_2^* s_1 \oplus \tau_{F_2} : D_2 = F_2 \to \pi_2^* E_1 \oplus \pi_2^* F_2$ (where τ_{F_2} is the tautological section of $\pi_2^* F_2 \to F_2$) is also a strongly transverse FOP section. Moreover, it is easy to see that

$$(f_1)_*\big(\chi^{\mathrm{FOP}}_{[\gamma]}(D_1, E_1, s_1)\big) = (f_1 \circ \pi_2)_*\big(\chi^{\mathrm{FOP}}_{[\gamma]}(D_2, E_2, s_2)\big).$$

Now we consider the case when (D_2, E_2, s_2, f_2) is another good stabilization of (D, E, s, f). Then it represents the same element as (D_1, E_1, s_1, f_1) in $\overline{\Omega}_*^{\mathbb{C},\mathrm{der}}(Y)$. Then by [Par20, Proposition 5.1][4], by applying suitable restrictions and stabilizations (by complex vector bundles) to both of them, we can assume that (D_1, E_1, s_1, f_1) and (D_2, E_2, s_2, f_2) are cobordant to each other via a compact bordism $(W, \tilde{E}, \tilde{s}, \tilde{f})$ such that all the relevant stable complex structures are actually complex and the orbifolds are effective. Because the invariance of the pushforward of the Fukaya–Ono–Parker–Euler class under germ equivalences and stabilizations when the section s is a strongly transverse FOP section, the cobordism invariance follows from the same arguments as in the proof of Proposition 5.6.5. Therefore, (5.7.7) is well-defined.

Lastly it is straightforward to verify that (5.7.7) descends to the desired map

[4]The cited proposition was stated and proved for unorientable bordism in the representable case but the proof applies to bordisms with structures and the non-representable case.

(5.7.6) because: 1) a restriction of (D, E, s, f) induces a restriction of a good stabilization; 2) a good stabilization of a stabilization (D', E', s', f') of (D, E, s, f) is also a good stabilization of (D, E, s, f); 3) one also has a notion of good stabilizations of cobordisms. □

5.7.5 Naturality of the $\mathcal{FOP}$ map

Suppose (X, A) is a pair of topological spaces. [5] Then elements in $\overline{\Omega}_*^{\mathbb{C},\mathrm{der}}(X, A)$ are represented by derived orbifold charts with boundary (D, E, s) together with a map $f : (D, \partial D) \to (X, A)$ and we need to choose a stable complex structure on (D, E, s). Note that using the collar of the boundary $\partial D \times [0, 1) \hookrightarrow D$, the derived orbifold chart (without boundary) $(\partial D, E|_{\partial D}, s|_\partial)$ acquires a stable complex structure. Then the stable complex derived orbifold bordism group $\overline{\Omega}_*^{\mathbb{C},\mathrm{der}}(X, A)$ is defined exactly the same as in Definition 5.7.18 which is for the absolute case. The additive structure is given by the disjoint union and the inverse element is induced by the conjugating complex structure on $\mathbb{C}$ as usual.

By restricting to the boundary, we see that there is a well-defined map

$$\partial : \overline{\Omega}_*^{\mathbb{C},\mathrm{der}}(X, A) \to \overline{\Omega}_{*-1}^{\mathbb{C},\mathrm{der}}(A). \tag{5.7.8}$$

Of course, by considering derived orbifold charts with $\partial D = \emptyset$, we obtain the projection

$$p : \overline{\Omega}_*^{\mathbb{C},\mathrm{der}}(X) \to \overline{\Omega}_*^{\mathbb{C},\mathrm{der}}(X, A) \tag{5.7.9}$$

and the inclusion

$$i : \overline{\Omega}_*^{\mathbb{C},\mathrm{der}}(A) \to \overline{\Omega}_*^{\mathbb{C},\mathrm{der}}(X). \tag{5.7.10}$$

[5] We restrict ourselves to *topological spaces* rather than *orbispaces* because the target of the $\mathcal{FOP}$ maps for the latter case are the corresponding coarse spaces, which bring us back to topological spaces in any case.

Proposition 5.7.23. The stable complex derived orbifold bordism group $\overline{\Omega}_*^{\mathbb{C},\mathrm{der}}(-,-)$ defines a generalized homology theory. Namely, it satisfies the Eilenberg–Steenrod axioms:

1. (*Homotopy*) If $g', g'' : (X_1, A_1) \to (X_2, A_2)$ are homotopic, then the induced maps
$$g'_*, g''_* : \overline{\Omega}_*^{\mathbb{C},\mathrm{der}}(X_1, A_1) \to \overline{\Omega}_*^{\mathbb{C},\mathrm{der}}(X_2, A_2).$$
are equal.

2. (Exactness) For $A \hookrightarrow X$, the following sequence between abelian groups are exact
$$\cdots \to \overline{\Omega}_*^{\mathbb{C},\mathrm{der}}(A) \xrightarrow{i} \overline{\Omega}_*^{\mathbb{C},\mathrm{der}}(X) \xrightarrow{p} \overline{\Omega}_*^{\mathbb{C},\mathrm{der}}(X, A) \xrightarrow{\partial} \overline{\Omega}_{*-1}^{\mathbb{C},\mathrm{der}}(A) \to \cdots .$$

3. (Excision) For $U \hookrightarrow A \hookrightarrow X$ such that $\overline{U} \subset Int(A)$, the natural inclusion of the pair $(X \setminus U, A \setminus U) \hookrightarrow (X, A)$ defines an isomorphism
$$\overline{\Omega}_*^{\mathbb{C},\mathrm{der}}(X, A) \xrightarrow{\sim} \overline{\Omega}_*^{\mathbb{C},\mathrm{der}}(X \setminus U, A \setminus U).$$

4. (Additivity) If $(X, A) = \coprod_i (X_i, A_i)$, then the canonical morphism
$$\overline{\Omega}_*^{\mathbb{C},\mathrm{der}}(X, A) \xrightarrow{\sim} \oplus_i \overline{\Omega}_*^{\mathbb{C},\mathrm{der}}(X_i, A_i)$$
is an isomorphism.

Proof. The proof is exactly the same as the one for classical bordism groups after keeping track of the vector bundles and sections, see for example [CF64, I.5]. □

The dimension axiom fails to be true. The group $\overline{\Omega}_*^{\mathbb{C},\mathrm{der}}(-)$ looks quite complicated. Understanding its structure would be very useful for symplectic topology.

Proposition 5.7.24. Suppose $[\gamma]$ is a stabilized isotropy type. Then the map $\mathcal{FOP}_{[\gamma]}$ defines a natural transformation between (generalized) homology theories

$$\mathcal{FOP}_{[\gamma]} : \overline{\Omega}_*^{\mathbb{C},\mathrm{der}}(-) \to H_{*+n_\gamma}(-;\mathbb{Z}).$$

Proof. The definition of the FOP map could be easily extended to the relative case using the collar structure. The functoriality follows from the fact that $\mathcal{FOP}_{[\gamma]}$ is defined by pushing forward the pseudocycle class produced by a generic FOP perturbation under the bordism map. Therefore, it suffices to show the following diagram commutes:

$$\begin{array}{ccc} \overline{\Omega}_{*+1}^{\mathbb{C},\mathrm{der}}(X,A) & \xrightarrow{\mathcal{FOP}_{[\gamma]}} & H_{*+1}(X,A;\mathbb{Z}) \\ \downarrow{\scriptstyle \partial} & & \downarrow{\scriptstyle \partial} \\ \overline{\Omega}_{*}^{\mathbb{C},\mathrm{der}}(A) & \xrightarrow{\mathcal{FOP}_{[\gamma]}} & H_{*}(A;\mathbb{Z}). \end{array} \quad . \tag{5.7.11}$$

Given a stable complex derived orbifold chart with boundary (D,E,s), assume that there is an embedding $\partial D \times [0,1) \to D$. After choosing a straightened structure on $(\partial D, E|_{\partial D}, s|_{\partial D})$, we can extend it to a straightened structure on (D,E,s) which coincides with the product structure on the collar $\partial D \times [0,1)$. After choosing an FOP perturbation $\tilde{s}|_{\partial D}$ of $s|_{\partial D}$, we can extend it to an FOP perturbation $\tilde{s}$ of s just as in the proof of Proposition 5.6.5. After taking out the strata indexed by $[\gamma]$, the restriction of $\tilde{s}^{-1}(0)$, which represents the image of $[(D,E,s)]$ under $\mathcal{FOP}_\gamma$, has the property that the further intersection with ∂D coincides with $\tilde{s}|_{\partial D}^{-1}(0)$ over the strata indexed by $[\gamma]$. This is exactly the content of (5.7.11). □

5.8 Whitney stratifications

In this section we review the classical result of the existence of canonical Whitney stratifications in the complex algebraic setting. We also prove certain relative version of this existence result which will be used in our construction. We follow the narrative

of [Mat73]. This is an appendix of this chapter.

5.8.1 Prestratifications and stratifications

Let M be a smooth manifold and $S \subset M$ be a subset. A decomposition of S into nonempty subsets

$$S = \bigsqcup_{\lambda \in \Lambda} S_\lambda$$

is called a *prestratification* if the following conditions are satisfied.

1. Each S_λ is locally closed.

2. The decomposition is locally finite.

3. (axiom of frontier) If $\overline{S_\lambda} \cap S_\mu \neq \emptyset$, then $S_\mu \subset \overline{S_\lambda}$.

Abusing the notation, we use the letter Λ to denote a decomposition of a given subset S. Each S_λ is called a *stratum* of this prestratification. Given a prestratification Λ, there is a partial order among strata:

$$S_\lambda \leq S_\mu \quad \text{if} \quad S_\lambda \subset \overline{S_\mu}.$$

Definition 5.8.1. Two subsets $A, B \subset M$ are called equivalent at $x \in M$ if there is an open neighborhood $U_x \subset M$ of x such that $A \cap U_x = B \cap U_x$. An equivalence class is called a *set-germ* at x.

Given a prestratification $\{S_\lambda \mid \lambda \in \Lambda\}$ of S, it assigns to each $x \in S$ a set-germ represented by the stratum S_λ which contains x. Two prestratifications of S are called *equivalent* if they give the same set-germ at each point of S.

Definition 5.8.2. A *stratification* of S is a rule $\mathfrak{S} = \{S_x \mid x \in S\}$ which assigns to each point $x \in S$ a set-germ S_x of closed subsets of S satisfying the following property: for each $x \in S$, there is a neighborhood $N \subset M$ of x and a prestratification

Λ of $S \cap N$ such that for each $y \in S \cap N$, the germ S_y is the germ at y of the member of Λ which contains y.

We see from the definition that for a stratification $\mathfrak{S}$ of S, the set-germ S_x contains x. Moreover, any equivalence class of prestratifications defines a stratification.

5.8.2 Whitney stratifications

Now we recall Whitney's conditions (a) and (b). It is well-known that condition (b) implies condition (a) (see [Mat12]).

Definition 5.8.3 (Whitney conditions (a) and (b))**.** Let U and V be smooth submanifolds of M. We say that the pair (U, V) satisfies Whitney's condition (a) at $x \in V$ if the following is true. Given a sequence $x_i \in U$ converging to x and $T_{x_i}U$ converges to a $\dim U$-dimensional subspace $\tau \subset T_xM$, one has $T_xV \subset \tau$. We say that the pair (U, V) satisfies Whitney's condition (b) at $x \in V$ if the following is true. Given a sequence $x_i \in V$ converging to x and $y_i \in U$ converging to x such that $x_i \neq y_i$, the secant line $\overline{x_i y_i}$ [6] converges to a line $l \subset T_xM$, and $T_{x_i}U$ converges to a $\dim U$-dimensional subspace $\tau \subset T_xM$, one has $l \subset \tau$.

Definition 5.8.4. Let M be a smooth manifold and $S \subset M$ is a subset. A prestratification

$$S = \bigsqcup_{\lambda \in \Lambda} S_\lambda$$

is called a (smooth) *Whitney prestratification* if each S_λ is a smooth submanifold of M of a constant dimension and for any $\lambda, \mu \in \Lambda$, the pair (S_λ, S_μ) satisfies Whitney's condition (b) at all points of S_μ.

A stratification of S is called a *Whitney stratification* if it is the induced from a Whitney prestratification.

[6] The secant line depends on a choice of local coordinate system near x, but the condition does not depend on such choices.

A useful fact is that the partial order relation among strata of a Whitney prestratification is compatible with dimensions.

Lemma 5.8.1. Given a Whitney prestratification Λ of S as in Definition 5.8.4, if $S_\lambda < S_\mu$, then $\dim S_\lambda < \dim S_\mu$.

Proof. Suppose $S_\lambda < S_\mu$. Then by definition, $S_\lambda \subset \overline{S_\mu}$. Choose $x \in S_\lambda$. Then one can find a sequence of points $y_i \in S_\mu$ converging to x such that $T_{y_i}S_\mu$ converges to a subspace $U \subset T_xM$. Then by Whitney's condition (a), $T_xS_\lambda \subset U$. Therefore, $\dim S_\lambda \leq \dim U = \dim S_\mu$.

Now choose a Riemannian metric in a neighborhood N of x in M and identify a neighborhood of $S_\lambda \cap N$ with a neighborhood of the zero section of the normal bundle of $S_\lambda \cap N$. Then one can find another sequence $x_i \in S_\lambda$ such that the secant line $\overline{x_i y_i}$ is perpendicular to $T_{x_i}S_\lambda$ and that x_i converges to x. Moreover, by taking a subsequence, we may assume that the secant lines converges to a line l which is perpendicular to T_xS_λ. By Whitney's condition (b), U must contain both T_xS_λ and l. Therefore one must have $\dim S_\mu > \dim S_\lambda$. □

A Whitney stratification defines a *dimension filtration* as follows. For each $k \geq 0$, define

$$S_k := \{x \in S \mid \dim S_x \leq k\}. \tag{5.8.1}$$

Here S_x is the set-term at x specified by the Whitney stratification, which has a well-defined dimension. Lemma 5.8.1 above implies that each S_k is a closed subset of S.

Definition 5.8.5. Let $\mathfrak{S}$ and $\mathfrak{S}'$ be two Whitney stratifications on $S \subset M$. We define $\mathfrak{S} < \mathfrak{S}'$ if there exists k such that $S_k \subsetneq S'_k$ and $S_l = S'_l$ for all $l > k$. We say that a Whitney stratification $\mathfrak{S}$ is *minimal* if for any Whitney stratification $\mathfrak{S}'$ one has $\mathfrak{S} \leq \mathfrak{S}'$.

It is easy to see that two Whitney stratifications are identical if they have identical dimension filtrations. Then it follows that minimal Whitney stratifications are unique.

Definition 5.8.6. Let M be a smooth manifold and $S \subset M$ has a Whitney stratification $\mathfrak{S}$. A smooth map $f : N \to M$ is called *lemma412* to S with respect to $\mathfrak{S}$ (or transverse to $\mathfrak{S}$) if for each $x \in f^{-1}(S)$, f is transverse to the germ $S_{f(x)}$ at x.

If f is transverse to S with respect to $\mathfrak{S}$, then f pulls back the set-germ $S_{f(x)}$ to a set-germ $f^{-1}(S_{f(s)})$ of submanifolds at x.

Theorem 5.8.7. [Mat73, Corollary 8.8] Let X, Y be smooth manifold. Let $S \subset X$ be a closed subset equipped with a Whitney stratification $\mathfrak{S}$. Let $f : Y \to X$ be a smooth map transverse to $\mathfrak{S}$. Then $f^*\mathfrak{S}$ is a Whitney stratification of $f^{-1}(S)$.

Remark 5.8.8. Minimal Whitney stratification is not a functorial object with respect to pullback via transverse maps. More precisely, suppose $\mathfrak{S}$ is the minimal Whitney stratification of $S \subset M$ and $f : N \to M$ is a smooth map transverse to $\mathfrak{S}$. The pullback Whitney stratification of $f^{-1}(S)$ may not be the minimal one.

5.8.3 Invariance of minimal Whitney stratifications

It is straightforward to prove that minimal Whitney stratifications are invariant under diffeomorphisms. For the purpose of our application, we would like to generalize this invariance property to certain relative setting for a broader class of smooth maps.

Definition 5.8.9. Let M be a smooth manifold and $S \subset M$ be a closed subset. Suppose there is a decomposition

$$S = S^\bullet \sqcup S^\circ$$

where S° is closed in M (we allow $S^\bullet = \emptyset$). Suppose also that $S^\bullet$ is equipped with a smooth Whitney stratification. A Whitney stratification on S is said to *extend* the

existing Whitney stratification on $S^\bullet$ if its restriction (viewed as a set-germ valued function) to $S^\bullet$ agrees with the existing one and if its restriction to S° is a Whitney stratification of S° (i.e. for $x \in S^\circ$, the set-germ S_x is contained in S°).

Lemma 5.8.2. Let M, N be smooth manifolds, $Y \subset M$, $Z \subset N$ be closed subsets with decompositions

$$Y = Y^\bullet \sqcup Y^\circ, \qquad\qquad Z = Z^\bullet \sqcup Z^\circ$$

where $Y^\circ \subset M$ resp. $Z^\circ \subset N$ is closed. Suppose Y resp. Z is equipped with a Whitney stratification $\mathfrak{Y}$ resp. $\mathfrak{Z}$ which is the minimal extension of the restriction $\mathfrak{Y}^\bullet := \mathfrak{Y}|_{Y^\bullet}$ resp. $\mathfrak{Z}^\bullet := \mathfrak{Z}|_{Z^\bullet}$.

Let $f : M \to N$ and $g : N \to M$ be smooth maps satisfying the following conditions.

1. f resp. g is transverse to $\mathfrak{Z}$ resp. $\mathfrak{Y}$.
2. $f^{-1}(Z^\bullet) = Y^\bullet$, $f^{-1}(Z^\circ) = Y^\circ$, $g^{-1}(Y^\bullet) = Z^\bullet$, $g^{-1}(Y^\circ) = Z^\circ$.
3. $f \circ g$ is transverse to $\mathfrak{Z}$ and pulls back $\mathfrak{Z}$ to $\mathfrak{Z}$.
4. $g \circ f$ is transverse to $\mathfrak{Y}$ and pulls back $\mathfrak{Y}$ to $\mathfrak{Y}$.
5. f pulls back $\mathfrak{Z}^\bullet$ to $\mathfrak{Y}^\bullet$ and g pulls back $\mathfrak{Y}^\bullet$ to $\mathfrak{Z}^\bullet$.

Then f pulls back $\mathfrak{Z}$ to $\mathfrak{Y}$ and g pulls back $\mathfrak{Y}$ to $\mathfrak{Z}$.

Proof. Denote $m = \dim M$, $n = \dim N$. As f resp. g is transverse to $\mathfrak{Z}$ resp. $\mathfrak{Y}$, by Theorem 5.8.7, the pullback $\check{\mathfrak{Y}} := f^*\mathfrak{Z}$ resp. $\check{\mathfrak{Z}} := g^*\mathfrak{Y}$ is a Whitney stratification on Y resp. Z. By condition (5) and the minimality of $\mathfrak{Y}$ resp. $\mathfrak{Z}$, one has

$$\mathfrak{Y} \leq \check{\mathfrak{Y}}, \qquad\qquad \mathfrak{Z} \leq \check{\mathfrak{Z}}.$$

We prove inductively that for all $l \geq 0$

$$Y_{m-l} = \check{Y}_{m-l}, \qquad\qquad Z_{n-l} = \check{Z}_{n-l}.$$

The $l = 0$ case is automatically true. Suppose the above is true for all $l < k$. We first show that $Y_{m-k} = \check{Y}_{m-k}$. Suppose on the contrary that $Y_{m-k} \neq \check{Y}_{m-k}$. Then as $\mathfrak{Y} \leq \check{\mathfrak{Y}}$, one must have $Y_{m-k} \subsetneq \check{Y}_{m-k}$. Hence there exists a point $y \in Y_{m-k+1} = \check{Y}_{m-k+1}$ and $y \in \check{Y}_{m-k} \setminus Y_{m-k}$. Notice that y must be in Y°. Let Y_y and $\check{Y}_y$ be the corresponding set-germs associated to $\mathfrak{Y}$ and $\check{\mathfrak{Y}}$ respectively. Then

$$\dim Y_y = m - k + 1, \qquad\qquad \dim \check{Y}_y \leq m - k. \tag{5.8.2}$$

The assumption that $g \circ f$ is transverse to $\mathfrak{Y}$ implies that f is transverse to $\check{\mathfrak{Z}}$, and the assumption that $g \circ f$ pulls back $\mathfrak{Y}$ to $\mathfrak{Y}$ implies that f pulls back $\check{\mathfrak{Z}}$ to $\mathfrak{Y}$. Therefore, by (5.8.2) one has

$$\dim \check{Z}_{f(y)} = n - m + \dim Y_y = n - k + 1.$$

where $\check{Z}_{f(y)}$ is the set-germ at $f(y)$ associated to the Whitney stratification $\check{\mathfrak{Z}}$. On the other hand, by (5.8.2) and the transversality of f to $\mathfrak{Z}$, one has

$$\dim Z_{f(y)} = n - m + \dim \check{Y}_y \leq n - k.$$

Therefore,

$$f(y) \in Z_{n-k} \setminus \check{Z}_{n-k} \Longrightarrow Z_{n-k} \nsubseteq \check{Z}_{n-k}.$$

As $Z_{n-l} = \check{Z}_{n-l}$ for all $l < k$, this contradicts the fact that $\mathfrak{Z} \leq \check{\mathfrak{Z}}$. Therefore, $Y_{m-k} = \check{Y}_{m-k}$. Similarly, one also obtains that $Z_{n-k} = \check{Z}_{n-k}$. Therefore $\mathfrak{Y} = \check{\mathfrak{Y}}$ and $\mathfrak{Z} = \check{\mathfrak{Z}}$. □

It is useful to consider a special case.

Corollary 5.8.10. Let M be a smooth manifold and $Y \subset M$ be a closed set with a decomposition

$$Y = Y^\bullet \sqcup Y^\circ$$

with Y° being closed. Suppose $\mathfrak{Y}$ is a smooth Whitney stratification on Y which is the minimal extension of $\mathfrak{Y}^\bullet := \mathfrak{Y}|_{Y^\bullet}$.

Suppose $f : M \to M$ is a diffeomorphism satisfying 1) $f(Y) = Y$, 2) $f(Y^\bullet) = Y^\bullet$, and 3) f pulls back $\mathfrak{Y}^\bullet$ to itself. Then f pulls back $\mathfrak{Y}$ to itself.

Proof. This follows from Lemma 5.8.2 where $M = N$, $Y = Z$, and $g = f^{-1}$. □

5.8.4 Existence results in the complex algebraic case

Now assume that M is nonsingular complex algebraic variety. A subset of M is called *constructible* if it is a finite union of locally closed sets (in Zariski topology). A constructible subset $U \subset M$ is a *nonsingular algebraic subset* if for each $x \in U$ there are rational functions $f_1, \ldots, f_a$, regular on a Zariski neighborhood V of x in M, such that

$$U \cap V = \{y \in V \mid f_1(y) = \cdots = f_a(y) = 0\}$$

and such that the differentials $df_1, \ldots, df_a$ are linearly independent at x.

Let $S \subset M$ be a constructible subset. A point $p \in S$ is called a *smooth* regular point if there is an open neighborhood $N \subset M$ of p (in the Euclidean topology) such that $N \cap S$ is a smooth submanifold.

Lemma 5.8.3. $p \in S$ is a smooth regular point if and only if there exists a Zariski open neighborhood $U \subset M$ of p such that $S \cap U$ is a nonsingular algebraic set.

Proof. See [Mil68, p.13, Remark]. □

Therefore we do not need to distinguish smooth regular points and algebraic regular points. A regular point p of a constructible set S, has a well-defined local dimension $\dim_p S$. In this appendix "dimension" always means "real dimension." A point $p \in S$ is called *singular* if it is not regular. In that case we define the local dimension to be $-\infty$. For each dimension k, let

$$S_{k,\mathrm{reg}} \subset S = \{p \in S \text{ regular} \mid \dim_p S = k\}.$$

Then $S_{k,\mathrm{reg}}$ for all k, together with the set of singular points S_{sing}, are all constructible sets. We also denote

$$\dim S \leq k \iff \dim_p S \leq k, \ \forall p \in S.$$

Now we give a proof of the existence results of minimal Whitney stratifications following Mather [Mat73]. We restrict our consideration to the case of constructible sets inside a smooth complex algebraic variety M. The construction relies on the following fundamental result, whose original version was proved by Whitney [Whi65] in the complex analytic setting.

Theorem 5.8.11. (see [Mat73, Theorem 4.1, Addendum 4.4]) Let U and V be disjoint nonsingular algebraic subsets of M. Suppose $\dim V < \dim U$. Let $S_b(U, V)$ be the set of points $x \in V \cap \overline{U}$ where Whitney's condition (b) fails for the triple (U, V, x). Then $S_b(U, V)$ is constructible. Moreover, if $V \subset \overline{U}$, then $S_b(U, V)$ is nowhere dense in V.

Now we can construct a canonical Whitney stratification on a constructible set. Indeed the following theorem is the "absolute" version of Theorem 5.8.15 below. However we would like to give a separate proof to show the reader the basic idea of the construction.

Theorem 5.8.12. [Mat73, Theorem 4.9] Let S be a constructible set of a smooth complex algebraic variety M. Then there exists a minimal smooth Whitney stratifi-

cation $\mathfrak{S}$ of S, and this stratification is algebraic, i.e., it is induced from a Whitney prestratification whose strata are all nonsingular algebraic subsets.

Proof. Let $n = \dim_{\mathbb{R}} M$. We construct inductively a decreasing sequence of subsets

$$S = \Sigma_n \supseteq \Sigma_{n-1} \supseteq \cdots \supseteq \Sigma_k$$

such that for each $l \geq k$ the following conditions are satisfied.

1. Σ_l is a closed algebraic subset of M of real dimension at most l.

2. $\Sigma_{l+1} \setminus \Sigma_l$ is a nonsingular algebraic subset of real dimension $l+1$ (empty if $l+1$ is odd).

3. For each $x \in \Sigma_{l+1} \setminus \Sigma_l$ and $m > l$, the pair $(\Sigma_{m+1} \setminus \Sigma_m, \Sigma_{l+1} \setminus \Sigma_l)$ satisfies Whitney's condition (b) at x.

We start with $\Sigma_n = S$. Suppose we have constructed Σ_l for all $l \geq k$ satisfying the induction hypothesis. Then define Σ_{k-1} to be the closure of points $x \in \Sigma_k$ satisfying one of the following conditions.

1. x is a singular point of Σ_k or a regular point with $\dim_x \Sigma_k < k$;

2. x is a regular point of Σ_k with local dimension equal to k, and there exists $l > k$ such that the pair $(\Sigma_l \setminus \Sigma_{l-1}, \Sigma_{k,\mathrm{reg}})$ does not satisfy Whitney's condition (b) at x (in particular $x \in \overline{\Sigma_l \setminus \Sigma_{l-1}}$).

The set of points satisfying (1) above is constructible; as Euclidean closures of constructible sets are also Zariski closed, the closure of the set of points satisfying (1) is a closed algebraic set of dimension at most $k-1$. On the other hand, by Theorem 5.8.11 above, the closure of $S_b(\Sigma_l \setminus \Sigma_{l-1}, \Sigma_{k,\mathrm{reg}})$ is a closed algebraic set and has dimension at most $k-1$. Hence Σ_{k-1} is a closed algebraic subset of M of dimension at most $k-1$ and $\Sigma_k \setminus \Sigma_{k-1}$ is a nonsingular algebraic subset of dimension k. Other

properties required for the induction hypothesis for Σ_{k-1} are automatically satisfied. This completes the induction step.

Now let Λ be the collection of connected components (with respect to the Euclidean topology) of $\Sigma_k \setminus \Sigma_{k-1}$ for all k. Then the collection Λ is locally finite. Then Λ satisfies all conditions for a Whitney prestratification except for the axiom of frontier. The fact that Λ satisfies the axiom of frontier follows from Proposition 5.8.13 below.

Now take the Whitney stratification induced from Λ. One needs to show that it is minimal among all smooth Whitney stratifications. Indeed, let $\mathfrak{S}'$ be another smooth Whitney stratification with dimension filtration $\{\Sigma'_k \mid k = 0, 1, \cdots\}$. We need to show that either $\Sigma_l = \Sigma'_l$ for all l or it is true for $l > k$ and $\Sigma_k \subsetneq \Sigma'_k$. We prove this claim inductively. For $n = \dim_{\mathbb{R}} M$, one has $\Sigma_n = \Sigma'_n = S$. Suppose we have proved that $\Sigma_l = \Sigma'_l$ for all $l > k$ and suppose by contradiction that $\Sigma_k \nsubseteq \Sigma'_k$. Then there exists $x \in \Sigma_k$ but $x \notin \Sigma'_k$. Hence $x \in \Sigma'_{k+1} \setminus \Sigma'_k$. As Σ'_k is closed, it follows $\Sigma'_{k+1} = \Sigma_{k+1}$ is locally a smooth submanifold near x of dimension $k+1$. Therefore, by Lemma 5.8.3, x is also an algebraic regular point of Σ_{k+1}, hence $x \in \Sigma_{k+1,\mathrm{reg}}$. Then by the construction of Σ_k, there must be some $m > k$ such that $(\Sigma_{m+1} \setminus \Sigma_m, \Sigma_{k+1,\mathrm{reg}})$ does not satisfy Whitney's condition (b) at x. As $\Sigma_{m+1} \setminus \Sigma_m = \Sigma'_{m+1} \setminus \Sigma'_m$, this contradicts the assumption that $\mathfrak{S}'$ is a smooth Whitney stratification of S.

Finally, we show that this Whitney stratification is complex algebraic. Indeed, for each k, $\Sigma_k \setminus \Sigma_{k-1}$ is the union of its irreducible components, each of which are nonsingular algebraic subsets. As $\Sigma_k \setminus \Sigma_{k-1}$ is nonsingular, different irreducible components do not intersect. On the other hand, each irreducible component is connected as it is regular. Therefore, the Whitney prestratification Λ considered above has its strata being nonsingular algebraic subsets. Therefore, the minimal Whitney stratification is algebraic. □

Proposition 5.8.13. [Mat73, Proposition 8.7] Let X be a C^2 manifold and $A \subset X$

be a closed subset. Let Λ be a partition of A into subsets, each of which is a C^2 submanifold of X. Suppose that Λ is locally finite and if $A_\lambda, A_\mu \in \Lambda$ then the pair (A_λ, A_μ) satisfies Whitney's condition (b). If all members of Λ are connected, then Λ satisfies the axiom of the frontier and hence is a Whitney prestratification.

Definition 5.8.14. We call the Whitney stratification of a constructible subset $S \subset M$ constructed in the proof of Theorem 5.8.12 the *canonical* Whitney stratification and the prestratification Λ whose strata are connected components of $\Sigma_k \setminus \Sigma_{k-1}$ the *canonical Whitney prestratification*.

5.8.5 Relative minimal Whitney stratifications

For application, we need certain Whitney stratifications which are minimal in some relative sense. Suppose M is a smooth complex algebraic variety. Let $S \subset M$ be a closed algebraic set with a decomposition

$$S = S^\bullet \sqcup S^\circ$$

where S° itself is a closed algebraic set. Suppose $S^\bullet$ is equipped with an algebraic Whitney stratification, denoted by $\mathfrak{S}^\bullet$.

Theorem 5.8.15. Among all smooth Whitney stratifications of S which extend the existing Whitney stratification on $S^\bullet$, there is a unique minimal one. In addition, the minimal one is algebraic.

Proof. We follow the same strategy of the proof of Theorem 5.8.12. Let $\{\Omega_k\}$ be the dimension filtration of $S^\bullet$ associated to the existing Whitney stratification. As the existing Whitney stratification is algebraic, for each k, $\Omega_k \setminus \Omega_{k-1}$ is a nonsingular algebraic subset. Suppose $\dim M = n$. For each $k \leq n$, we would like to define a

decreasing sequence of closed analytic subsets

$$S^\circ = \Sigma_n \supseteq \Sigma_{n-1} \cdots \supseteq \Sigma_k$$

such that for all $l \geq k$ the following conditions are satisfied.

1. Σ_l is a closed algebraic set of dimension at most l.
2. $\Sigma_{l+1} \setminus \Sigma_l$ is a nonsingular algebraic subset of dimension $l+1$.
3. $\Sigma_{l+1} \setminus \Sigma_l$ is disjoint from the closure of Ω_{l+1}.
4. For each $x \in \Sigma_{l+1} \setminus \Sigma_l$ and $m > l$, the pair $(\Sigma_{m+1} \setminus \Sigma_m, \Sigma_{l+1} \setminus \Sigma_l)$ and the pair $(\Omega_{m+1} \setminus \Omega_m, \Sigma_{l+1} \setminus \Sigma_l)$ satisfy Whitney's condition (b) at x.

Now we start with $\Sigma_n := S^\circ$. Suppose we have constructed $\Sigma_n, \ldots, \Sigma_k$ satisfying the above conditions. Now consider the closure of points $x \in \Sigma_k$ satisfying one of the following conditions, denoted by $\Sigma_{k-1} \subset \Sigma_k$.

1. x is contained in the closure of Ω_k.
2. x is either a singular point of Σ_k or a regular point with $\dim_x \Sigma_k < k$.
3. There exists $m > k$ such that the pair $(\Sigma_m \setminus \Sigma_{m-1}, \Sigma_{k,\mathrm{reg}})$ does not satisfy Whitney's condition (b) at x.
4. There exists $m > k$ such that the pair $(\Omega_m \setminus \Omega_{m-1}, \Sigma_{k,\mathrm{reg}})$ does not satisfy Whitney's condition (b) at x.

Then following the same argument as the proof of Theorem 5.8.12, $\Sigma_k \setminus \Sigma_{k-1}$ is a nonsingular algebraic subset and Σ_{k-1} is a closed algebraic subset of real dimension at most $k-1$. This finishes the inductive step. Then similar to Theorem 5.8.12, one can find a Whitney prestratification (of S°) by taking all connected components of $\Sigma_k \setminus \Sigma_{k-1}$. Moreover, the induced Whitney stratification is complex algebraic.

We need to prove that this canonically constructed Whitney stratification is minimal among all smooth extensions of the existing one. Let S_k be the associated dimension filtration and let S'_k be the associated dimension filtration of another smooth extension $\mathfrak{S}'$. Then one can see

$$S_k = \Omega_k \sqcup \Sigma_k, \qquad\qquad S'_k = \Omega_k \sqcup \Sigma'_k.$$

We need to show that either $\Sigma_l = \Sigma'_l$ for all l or there exists k such that it is true for all $l > k$ and $\Sigma_k \subsetneq \Sigma'_k$. We argue inductively. For $n = \dim M$, one has $\Sigma_n = \Sigma'_n = S^\circ$ by definition. Suppose we have proved that $\Sigma_l = \Sigma'_l$ for all $l > k$. Assume in contradiction that $\Sigma_k \nsubseteq \Sigma'_k$. Then there exists a point $x \in \Sigma_k$ which is not contained in Σ'_k, namely, $x \in \Sigma'_{k+1} \setminus \Sigma'_k$. As Σ'_k is closed, x is a smooth regular point of Σ'_{k+1} of local (real) dimension $k+1$. Then by Lemma 5.8.3, x is also an algebraic regular point. Now because $\mathfrak{S}'$ is a Whitney stratification of the whole set S and $x \in \Sigma'_{k+1} \setminus \Sigma'_k$, it follows from Lemma 5.8.1, $x \notin \overline{\Omega_{k+1}}$. Then as $x \in \Sigma_k$, by the construction of the canonical Whitney stratification, it follows that there is $m > k$ such that either $(\Sigma_{m+1} \setminus \Sigma_m, \Sigma_{k+1,\mathrm{reg}})$ or $(\Omega_{m+1} \setminus \Omega_m, \Sigma_{k+1,\mathrm{reg}})$ does not satisfy Whitney's condition (b) at x. This contradicts the hypothesis that $\mathfrak{S}'$ is a Whitney stratification. □

5.8.6 Relative to a prestratification

Let V be a smooth complex algebraic variety equipped with a finite algebraic Whitney prestratification

$$V = \bigsqcup_{\alpha \in A} V_\alpha. \tag{5.8.3}$$

We denote this prestratification by $\mathcal{A}$.

Example 5.8.1. Suppose G is a finite group and V is a finite dimensional complex representation of G. Given any subgroup $H \subseteq G$, denote by V_H the invariant subspace under the H-action and define $V_H^* \subset V$ to be the subset of V which consists of points

whose isotropy group is given by H. Then $\overline{V_H^*} = V_H$ and the decomposition

$$V = \bigsqcup_{H \subseteq G} V_H^*$$

is an algebraic Whitney prestratification.

Definition 5.8.16. Let $Z \subset V$ be a closed algebraic subset. A smooth Whitney stratification of Z is said to *respect* the prestratification (5.8.3) if for each $x \in Z \cap V_\alpha$ the set-germ Z_x is contained in V_α.

The main result here is that there exists a minimal smooth Whitney stratification that respects the given prestratification, which is also algebraic.

Theorem 5.8.17. There exists an algebraic Whitney stratification on Z which respects $\mathcal{A}$ and which is the minimal among all smooth Whitney stratifications which respect $\mathcal{A}$.

Proof. The desired Whitney stratification is the union of Whitney stratifications of the strata $Z_\alpha := Z \cap V_\alpha$. We inductively construct those Whitney stratifications. Choose an increasing order of $\{V_\alpha \mid \alpha \in A\}$ as

$$V_{\alpha_1}, \ldots, V_{\alpha_m}$$

such that $V_{\alpha_i} \leq V_{\alpha_j}$ implies $i \leq j$. For all k, denote

$$V_{(k)} := V \setminus \bigcup_{i<k} V_{\alpha_i}$$

which is a Zariski open subset of V. Denote

$$Z_{(k)} := Z \cap V_{(k)}.$$

Suppose, inductively, that we have constructed an algebraic Whitney stratification

on $Z_{(k)} \subset V_{(k)}$ which respects A and which is minimal among such smooth Whitney stratifications. Now we need to extend to $Z_{(k-1)}$. Let

$$\{\Sigma_{(k),l} \subset Z_{(k)} \mid l = 0, 1, \ldots\}$$

be the dimension filtration of $Z_{(k)}$. Then for each l, the set

$$\Sigma_{(k),l} \setminus \Sigma_{(k),l-1}$$

is a nonsingular algebraic subset of $V_{(k)}$. As $V_{(k)}$ is Zariski open inside $V_{(k-1)}$, the set $\Sigma_{(k),l} \setminus \Sigma_{(k),l-1}$ is also a nonsingular algebraic subset of $V_{(k-1)}$.[7] Hence the existing Whitney stratification of $Z_{(k)}$ (which is viewed as a subset of $V_{(k-1)}$) is also algebraic. Then by Theorem 5.8.15, there exists a canonical extension of this Whitney stratification to the closed algebraic subset

$$Z_{(k-1)} = Z_{(k)} \sqcup (Z \cap V_{\alpha_{k-1}})$$

of $V_{(k-1)}$ and the extension is still algebraic. Inductively, one has constructed an algebraic Whitney stratification which respects the given prestratification.

Lastly, one can combine induction with the proof of the minimality part of Theorem 5.8.15 to show that the constructed Whitney stratification is minimal among all smooth Whitney stratifications of Z which respect the given prestratification. □

Remark 5.8.18. Suppose $B \subset A$ satisfies

$$\alpha \in A, \beta \in B, \beta \leq \alpha \Longrightarrow \alpha \in B.$$

[7] This is the point where our proof does not extend to the complex analytic setting.

Define

$$V_B := \bigsqcup_{\beta \in B} V_\beta, \qquad\qquad Z_B := Z \cap V_B.$$

Then V_B is Zariski open in V and has the induced prestratification $\mathcal{B}$ whose strata are V_β. In the context of Example 5.8.1, a typical example of V_B is the set

$$V_H^+ := \{v \in V \mid G_v \subset H\}$$

for a given subgroup $H \subseteq G$. Then it is easy to see that the restriction of the minimal Whitney stratification on $Z \subset V$ which respects the prestratification $\mathcal{A}$ to V_B is also the minimal Whitney stratification on $Z_B \subset V_B$ which respects the prestratification $\mathcal{B}$.

Lastly, we have the following invariance property of the canonical Whitney stratification which respects the given prestratification.

Proposition 5.8.19. Let f be a smooth diffeomorphism of V such that $f(Z) = Z$ and such that for some bijection $\rho : A \to A$ there holds

$$f(V_\alpha) = V_{\rho(\alpha)}, \ \forall \alpha \in A.$$

Then f preserves the minimal Whitney stratification which respects $\mathcal{A}$.

Proof. We decompose

$$A = \bigsqcup_{i=1}^{m} A_i$$

where the A_i's are the minimal nonempty ρ-invariant subsets of A. Then $\{A_1, \ldots, A_m\}$ inherits a partial order from A which is still denoted by $\leq$. We also assume the order

$A_1, \ldots, A_m$ is increasing, i.e., $A_i < A_j$ implies $i < j$. Set

$$V_{A_i} := \bigcup_{\alpha \in A_i} V_\alpha, \ i = 1, \ldots, m$$

each of which is f-invariant and set

$$V_{(k)} := \bigcup_{l \geq k} V_{A_l}, \ k = 1, \ldots, m$$

each of which is Zariski open. Then by the absolute case of Corollary 5.8.10 (when $M = V_{(m)}, Y = Y^\circ = Z_{(m)} := Z \cap V_{(m)}$, f preserves the minimal Whitney stratification on $Z_{(m)} \subset V_{(m)}$. Now suppose we have shown that for some k that f preserves the Whitney stratification on $Z_{(k)} := Z \cap V_{(k)}$. Then by Corollary 5.8.10, f preserves the Whitney stratification on $Z_{(k-1)}$. By induction, f preserves the minimal Whitney stratification on Z which respects $\mathcal{A}$. $\square$

Bibliography

[AB21] Mohammed Abouzaid and Andrew J Blumberg, *Arnold conjecture and Morava K-theory*, arXiv preprint arXiv:2103.01507 (2021).

[Abo10] Mohammed Abouzaid, *A geometric criterion for generating the Fukaya category*, Publications Mathématiques de l'IHÉS **112** (2010), 191–240.

[Aea17] Shishir Agrawal et. al., *The CRing project*, 2017.

[AGEN19] Daniel Alvarez-Gavela, Yakov Eliashberg, and David Nadler, *Arborealization II: Geomorphology of Lagrangian ridges*, arXiv preprint arXiv:1912.03439 (2019).

[AGEN20] ———, *Arborealization III: Positive arborealization of polarized Weinstein manifolds*, arXiv preprint arXiv:2011.08962 (2020).

[AGEN21] ———, *Arborealization I: Stability of arboreal models*, arXiv preprint arXiv:2101.04272 (2021).

[ALR07] Alejandro Adem, Johann Leida, and Yongbin Ruan, *Orbifolds and stringy topology*, Cambridge Tracts in Mathematics, vol. 171, Cambridge University Press, 2007.

[AM17] Mohammed Abouzaid and Ciprian Manolescu, *A sheaf-theoretic model for* $\mathrm{SL}(2,\mathbb{C})$ *Floer homology*, arXiv preprint arXiv:1708.00289 (2017).

[AM90] Selman Akbulut and John D McCarthy, *Casson's invariant for oriented homology 3-spheres: an exposition* (1990).

[AMS21] Mohammed Abouzaid, Mark McLean, and Ivan Smith, *Complex cobordism, Hamiltonian loops and global Kuranishi charts*, 2021. arXiv: 2110.14320.

[AS10a] Alberto Abbondandolo and Matthias Schwarz, *Floer homology of cotangent bundles and the loop product*, Geom. Topol. **14** (2010), no. 3, 1569–1722.

[AS10b] Mohammed Abouzaid and Paul Seidel, *Altering symplectic manifolds by homologous recombination*, arXiv preprint arXiv:1007.3281 (2010).

[AS10c] ———, *An open string analogue of Viterbo functoriality*, Geom. Topol. **14** (2010), no. 2, 627–718.

[AT15] Takuma Aihara and Ryo Takahashi, *Generators and dimensions of derived categories of modules*, Comm. Algebra **43** (2015).

[BBD82] A. A. Beilinson, J. Bernstein, and P. Deligne, *Faisceaux pervers*, Analysis and topology on singular spaces, I (Luminy, 1981), 1982, pp. 5–171.

[BD19] Christopher Brav and Tobias Dyckerhoff, *Relative Calabi–Yau structures*, Compositio Mathematica **155** (2019), no. 2, 372–412.

[BDM19] Matthew R. Ballard, Alexander Duncan, and Patrick K. McFaddin, *The toric Frobenius morphism and a conjecture of Orlov*, Eur. J. Math. **5** (2019), no. 3, 640–645.

[BF12] Matthew Ballard and David Favero, *Hochschild dimensions of tilting objects*, International Mathematics Research Notices **2012** (2012), no. 11, 2607–2645.

[BFK14] Matthew Ballard, David Favero, and Ludmil Katzarkov, *A category of kernels for equivariant factorizations, II: further implications*, Journal de Mathématiques Pures et Appliquées **102** (2014), no. 4, 702–757.

[BG17] Benjamin Briggs and Vincent Gélinas, *The A-infinity centre of the Yoneda algebra and the characteristic action of Hochschild cohomology on the derived category*, arXiv preprint arXiv:1702.00721 (2017).

[BH98] Hans U Boden and Christopher M Herald, *The* SU(3) *Casson invariant for integral homology 3-spheres*, Journal of Differential Geometry **50** (1998), no. 1, 147–206.

[BHK01] Hans U Boden, Christopher M Herald, and Paul Kirk, *An integer valued* SU(3) *Casson invariant*, Mathematical Research Letters **8** (2001), no. 5, 589–603.

[BHK05] Hans U Boden, Christopher M Herald, and Paul A Kirk, *The integer valued* SU(3) *Casson invariant for Brieskorn spheres*, Journal of Differential Geometry **71** (2005), no. 1, 23–83.

[BIK08] Dave Benson, Srikanth B. Iyengar, and Henning Krause, *Local cohomology and support for triangulated categories*, Ann. Sci. Éc. Norm. Supér. (4) **41** (2008), no. 4, 573–619.

[BIKO10] Petter Andreas Bergh, Srikanth B. Iyengar, Henning Krause, and Steffen Oppermann, *Dimensions of triangulated categories via Koszul objects*, Math. Z. **265** (2010), no. 4, 849–864.

[BN90] Steven Boyer and Andrew Nicas, *Varieties of group representations and Casson's invariant for rational homology 3-spheres*, Transactions of the American Mathematical Society **322** (1990), no. 2, 507–522.

[Bon89] A. I. Bondal, *Representations of associative algebras and coherent sheaves*, Izv. Akad. Nauk SSSR Ser. Mat. **53** (1989), no. 1, 25–44 (Russian).

[BP01] Jim Bryan and Rahul Pandharipande, *BPS states of curves in Calabi-Yau 3-folds*, Geom. Topol. **5** (2001).

[BS21] Shaoyun Bai and Mohan Swaminathan, *Bifurcations of embedded curves and an extension of Taubes' Gromov invariant to Calabi-Yau 3-folds*, arXiv preprint arXiv:2106.01206 (2021).

[BZ20] Shaoyun Bai and Boyu Zhang, *Equivariant Cerf theory and perturbative* SU(n) *Casson invariants*, arXiv preprint arXiv:2009.01118 (2020).

[CF09] Kai Cieliebak and Urs Adrian Frauenfelder, *A Floer homology for exact contact embeddings*, Pacific J. Math. **239** (2009), no. 2, 251–316.

[CF64] P. E. Conner and E. E. Floyd, *Differentiable periodic maps*, Ergebnisse der Mathematik und ihrer Grenzgebiete, (N.F.), Band 33, Academic Press, Inc., Publishers, New York; Springer-Verlag, Berlin-Göttingen-Heidelberg, 1964.

[CK99] David A Cox and Sheldon Katz, *Mirror symmetry and algebraic geometry*, American Mathematical Soc., 1999.

[CLM02] Sylvain E Cappell, Ronnie Lee, and Edward Y Miller, *A perturbative* SU(3) *Casson invariant*, Commentarii Mathematici Helvetici **77** (2002), no. 3, 491–523.

[CLM90] ———, *A symplectic geometry approach to generalized Casson's invariants of 3-manifolds*, Bulletin of the American Mathematical Society **22** (1990), no. 2, 269–275.

[CLM94] ———, *On the Maslov index*, Communications on Pure and Applied Mathematics **47** (1994), no. 2, 121–186.

[CRGG17] Baptiste Chantraine, Georgios Dimitroglou Rizell, Paolo Ghiggini, and Roman Golovko, *Geometric generation of the wrapped fukaya category of weinstein manifolds and sectors*, arXiv preprint arXiv:1712.09126 (2017).

[CS99] Moira Chas and Dennis Sullivan, *String topology*, arXiv preprint math (1999).

[Cur94] Cynthia L Curtis, *Generalized Casson invariants for* $\mathrm{SO}(3), \mathrm{U}(2), \mathrm{Spin}(4)$ *and* $\mathrm{SO}(4)$, Transactions of the American Mathematical Society (1994), 49–86.

[DIW21] Aleksander Doan, Eleny-Nicoleta Ionel, and Thomas Walpuski, *The Gopakumar-Vafa finiteness conjecture*, arXiv preprint arXiv:2103.08221 (2021).

[Don19] Simon Donaldson, *Deformations of multivalued harmonic functions*, arXiv preprint arXiv:1912.08274 (2019).

[Don87] Simon K Donaldson, *The orientation of Yang-Mills moduli spaces and 4-manifold topology*, Journal of Differential Geometry **26** (1987), no. 3, 397–428.

[Dri04] Vladimir Drinfeld, *Dg quotients of dg categories*, J. Algebra **272** (2004), no. 2, 643–691.

[DS06] Karl Dovermann and Reinhard Schultz, *Equivariant surgery theories and their periodicity properties*, Lecture Notes in Mathematics, vol. 1443, Springer, 2006.

[DS94] Stamatis Dostoglou and Dietmar A. Salamon, *Cauchy-Riemann operators, self-duality, and the spectral flow*, First European Congress of Mathematics, Vol. I (Paris, 1992), 1994, pp. 511–545.

[DU95] Georgios D Daskalopoulos and Karen K Uhlenbeck, *An application of transversality to the topology of the moduli space of stable bundles*, Topology **34** (1995), no. 1, 203–215.

[DW17] Harm Derksen and Jerzy Weyman, *An introduction to quiver representations*, Graduate Studies in Mathematics, vol. 184, American Mathematical Society, Providence, RI, 2017.

[DW19a] Aleksander Doan and Thomas Walpuski, *Counting embedded curves in symplectic 6-manifolds*, arXiv preprint arXiv:1910.12338 (2019).

[DW19b] ———, *On counting associative submanifolds and Seiberg-Witten monopoles*, Pure Appl. Math. Q. **15** (2019), no. 4, 1047–1133.

[DW20] ———, *Equivariant Brill-Noether theory for elliptic operators and super-rigidity of J-holomorphic maps*, 2020.

[DW21] ———, *Castelnuovo's bound and rigidity in almost complex geometry*, Advances in Mathematics **379** (2021), 107550.

[Eft16] Eaman Eftekhary, *On finiteness and rigidity of J-holomorphic curves in symplectic three-folds*, Advances in Mathematics **289** (2016), 1082–1105.

[Eft20] ———, *Counting periodic orbits of vector fields over smooth closed manifolds*, arXiv preprint arXiv:2012.01808 (2020).

[EGL20] Yakov Eliashberg, Sheel Ganatra, and Oleg Lazarev, *Flexible Lagrangians*, Int. Math. Res. Not. IMRN **8** (2020), 2408–2435.

[EL18] A. D. Elagin and V. A. Lunts, *Regular subcategories in bounded derived categories of affine schemes*, Mat. Sb. **209** (2018), no. 12, 87–116.

[EL21] Alexey Elagin and Valery A. Lunts, *Three notions of dimension for triangulated categories*, J. Algebra **569** (2021), 334–376.

[Ela20] Alexey Elagin, *Calculating dimension of triangulated categories: path algebras, their tensor powers and orbifold projective lines*, arXiv preprint arXiv:2004.04694 (2020).

[EXZ21] Alexey Elagin, Junyan Xu, and Shizhuo Zhang, *On cyclic strong exceptional collections of line bundles on surfaces*, Eur. J. Math. **7** (2021), no. 1, 69–115.

[FH94] A. Floer and H. Hofer, *Symplectic homology. I. Open sets in* $\boldsymbol{C}^n$, Math. Z. **215** (1994), no. 1, 37–88.

[FHT01] Yves Félix, Stephen Halperin, and Jean-Claude Thomas, *Rational homotopy theory*, Graduate Texts in Mathematics, vol. 205, Springer-Verlag, New York, 2001.

[Flo88] Andreas Floer, *An instanton-invariant for 3-manifolds*, Communications in mathematical physics **118** (1988), no. 2, 215–240.

[FLTZ12] Bohan Fang, Chiu-Chu Melissa Liu, David Treumann, and Eric Zaslow, *T-duality and homological mirror symmetry for toric varieties*, Adv. Math. **229** (2012), no. 3, 1875–1911.

[FO97] Kenji Fukaya and Kaoru Ono, *Floer homology and Gromov–Witten invariant over integer of general symplectic manifolds - summary* -, Proceeding of the Last Taniguchi conference, 1997.

[FO99] ———, *Arnold conjecture and Gromov–Witten invariants for general symplectic manifolds*, Topology **38** (1999), 933–1048.

[FOOO13] Kenji Fukaya, Yong-Geun Oh, Hiroshi Ohta, and Kaoru Ono, *Lagrangian Floer theory over integers: sperhically positive symplectic manifolds*, Pure and Applied Mathematics Quarterly **9** (2013), no. 2, 189–289.

[FOOO16] ———, *Exponential decay estimates and smoothness of the moduli space of pseudo-holomorphic curves*, arXiv preprint arXiv:1603.07026 (2016).

[FOOO] ———, *Normally polynomial perturbations*. `https://www.math.kyoto-u.ac.jp/~fukaya/normallypolynomial.pdf`.

[FOT08] Yves Félix, John Oprea, and Daniel Tanré, *Algebraic models in geometry*, Oxford Graduate Texts in Mathematics, vol. 17, Oxford University Press, Oxford, 2008.

[FTVP04] Yves Felix, Jean-Claude Thomas, and Micheline Vigué-Poirrier, *The Hochschild cohomology of a closed manifold*, Publ. Math. Inst. Hautes Études Sci. **99** (2004), 235–252.

[FTVP07] Yves Félix, Jean-Claude Thomas, and Micheline Vigué-Poirrier, *Rational string topology*, J. Eur. Math. Soc. (JEMS) **9** (2007), no. 1, 123–156.

[Gam20] Benjamin Gammage, *Mirror symmetry for Berglund-Hübsch Milnor fibers*, arXiv preprint arXiv:2010.15570 (2020).

[Gan12] Sheel Ganatra, *Symplectic Cohomology and Duality for the Wrapped Fukaya Category*, ProQuest LLC, Ann Arbor, MI, 2012. Thesis (Ph.D.)–Massachusetts Institute of Technology.

[Gan21] ———, *Categorical non-properness in wrapped Floer theory*, arXiv preprint arXiv:2104.06516 (2021).

[Gor76] Mark Goresky, *Geometric cohomology and homology of stratified spaces*, Ph.D. Thesis, 1976.

[Gor78a] ———, *Triangulation of stratified objects*, Proceedings of the American Mathematical Society (1978), 193–200.

[Gor78b] R. Mark Goresky, *Triangulation of stratified objects*, Proc. Amer. Math. Soc. **72** (1978), no. 1, 193–200.

[GP17] Emmanuel Giroux and John Pardon, *Existence of Lefschetz fibrations on Stein and Weinstein domains*, Geom. Topol. **21** (2017), no. 2, 963–997.

[GPS18a] Sheel Ganatra, John Pardon, and Vivek Shende, *Microlocal morse theory of wrapped Fukaya categories*, arXiv preprint arXiv:1809.08807 (2018).

[GPS18b] ———, *Sectorial descent for wrapped Fukaya categories*, arXiv preprint arXiv:1809.03427 (2018).

[GPS20] ———, *Covariantly functorial wrapped Floer theory on Liouville sectors*, Publ. Math. Inst. Hautes Études Sci. **131** (2020), 73–200.

[Gro85] M. Gromov, *Pseudo holomorphic curves in symplectic manifolds*, Invent. Math. **82** (1985), no. 2, 307–347.

[GS17] Benjamin Gammage and Vivek Shende, *Mirror symmetry for very affine hypersurfaces*, arXiv preprint arXiv:1707.02959 (2017).

[GS20] Sheel Ganatra and Kyler Siegel, *On the embedding complexity of Liouville manifolds*, arXiv preprint arXiv:2012.04627 (2020).

[GV98a] Rajesh Gopakumar and Cumrun Vafa, *M-theory and topological strings–I*, 1998. `https://arxiv.org/abs/hep-th/9809187`.

[GV98b] ———, *M-theory and topological strings–II*, 1998. `https://arxiv.org/abs/hep-th/9812127`.

[Hat02] Allen Hatcher, *Algebraic topology*, Cambridge University Press, Cambridge, 2002.

[HK20] Paul Hacking and Ailsa Keating, *Homological mirror symmetry for log Calabi–Yau surfaces*, arXiv preprint arXiv:2005.05010 (2020).

[(k] Gregor Kemper (https://mathoverflow.net/users/82616/gregor kemper), *Equivariant polynomial maps*. URL:https://mathoverflow.net/q/223838 (version: 2015-11-18).

[HST01] Shinobu Hosono, Masa-Hiko Saito, and Atsushi Takahashi, *Relative Lefschetz action and BPS state counting*, International Mathematics Research Notices **2001** (2001), no. 15, 783–816.

[Hue95] Johannes Huebschmann, *Symplectic and Poisson structures of certain moduli spaces. I*, Duke Math. J. **80** (1995), no. 3, 737–756.

[HWZ07] H. Hofer, K. Wysocki, and E. Zehnder, *A general Fredholm theory. I. A splicing-based differential geometry*, J. Eur. Math. Soc. (JEMS) **9** (2007), no. 4, 841–876.

[HWZ17] ———, *Applications of polyfold theory I: The polyfolds of Gromov-Witten theory*, Mem. Amer. Math. Soc. **248** (2017), no. 1179, v+218.

[IP18] Eleny-Nicoleta Ionel and Thomas H. Parker, *The Gopakumar-Vafa formula for symplectic manifolds*, Ann. of Math. (2) **187** (2018), no. 1, 1–64.

[Jef94] Lisa C. Jeffrey, *Extended moduli spaces of flat connections on Riemann surfaces*, Math. Ann. **298** (1994), no. 4, 667–692.

[Joy07] Dominic Joyce, *Kuranishi homology and Kuranishi cohomology*, arXiv preprint arXiv:0707.3572 (2007).

[Kap70] Irving Kaplansky, *Commutative rings*, Allyn and Bacon, Inc., Boston, Mass., 1970.

[Kat08] Sheldon Katz, *Genus zero Gopakumar-Vafa invariants of contractible curves*, Journal of Differential Geometry **79** (2008), no. 2, 185–195.

[Kea18] Ailsa Keating, *Homological mirror symmetry for hypersurface cusp singularities*, Selecta Math. (N.S.) **24** (2018), no. 2, 1411–1452.

[KL12] Young-Hoon Kiem and Jun Li, *Categorification of Donaldson-Thomas invariants via perverse sheaves*, arXiv preprint arXiv:1212.6444 (2012).

[KM11] Peter B Kronheimer and Tomasz S Mrowka, *Knot homology groups from instantons*, Journal of Topology **4** (2011), no. 4, 835–918.

[Laz20] Oleg Lazarev, *Simplifying Weinstein Morse functions*, Geom. Topol. **24** (2020).

[Lee04] Yuan-Pin Lee, *Quantum K-theory, I: foundations*, Duke Mathematical Journal **121** (2004), no. 3, 389–424.

[Lel75] N. W. Lellmann, *Orbitraüme von G-Mannigfaltigkeiten und stratifizierte Mengen*, 1975. Diplomarbeit, Universität Bonn.

[Les96] Christine Lescop, *Global surgery formula for the Casson-Walker invariant*, Princeton University Press, 1996.

[LMP99] François Lalonde, Dusa McDuff, and Leonid Polterovich, *Topological rigidity of Hamiltonian loops and quantum homology*, Invent. Math. **135** (1999), no. 2, 369–385.

[LO06] Volodymyr Lyubashenko and Sergiy Ovsienko, *A construction of quotient A_∞-categories*, Homology Homotopy Appl. **8** (2006), no. 2, 157–203.

[Low08] Wendy Lowen, *Hochschild cohomology, the characteristic morphism and derived deformations*, Compos. Math. **144** (2008).

[LP20] Yanki Lekili and Alexander Polishchuk, *Homological mirror symmetry for higher-dimensional pairs of pants*, Compos. Math. **156** (2020), no. 7, 1310–1347.

[LT98] Jun Li and Gang Tian, *Virtual moduli cycles and Gromov–Witten invariants of general symplectic manifolds*, Topics in symplectic 4-manifolds (irvine, ca, 1996), 1998, pp. 47–83.

[LU18] Yanki Lekili and Kazushi Ueda, *Homological mirror symmetry for Milnor fibers via moduli of A_∞-structures*, arXiv preprint arXiv:1806.04345 (2018).

[LU20] ———, *Homological mirror symmetry for Milnor fibers of simple singularities*, arXiv preprint arXiv:2004.07374 (2020).

[Mar88] Alexis Marin, *Un nouvel invariant pour les spheres d'homologie de dimension trois*, Seminare Bourbaki, 1987–88, Asterisque N **693** (1988), 151.

[Mat12] John Mather, *Notes on topological stability*, Bulletin (New Series) of the American Mathematical Society **49** (2012), no. 4, 475–506.

[Mat73] ———, *Stratifications and mappings*, Dynamical systems, 1973, pp. 195–232.

[McD91] Dusa McDuff, *Symplectic manifolds with contact type boundaries*, Invent. Math. **103** (1991), no. 3, 651–671.

[Mil68] John Milnor, *Singular points of complex hypersurfaces*, Annals of Mathematics Studies, vol. 61, Princeton University Press, 1968.

[MNOP06a] D. Maulik, N. Nekrasov, A. Okounkov, and R. Pandharipande, *Gromov-Witten theory and Donaldson-Thomas theory. I*, Compos. Math. **142** (2006), no. 5, 1263–1285.

[MNOP06b] ———, *Gromov-Witten theory and Donaldson-Thomas theory. II*, Compos. Math. **142** (2006), no. 5, 1286–1304.

[MP97] I. Moerdijk and D. A. Pronk, *Orbifolds, sheaves and groupoids*, K-Theory **12** (1997), no. 1, 3–21.

[MRS11] Tomasz Mrowka, Daniel Ruberman, and Nikolai Saveliev, *Seiberg-Witten equations, end-periodic dirac operators, and a lift of Rohlin's invariant*, Journal of Differential Geometry **88** (2011), no. 2, 333–377.

[MS04] Dusa McDuff and Dietmar Salamon, *J-holomorphic curves and symplectic topology*, Colloquium Publications, vol. 52, American Mathematical Society, 2004.

[MS10] Maksim Maydanskiy and Paul Seidel, *Lefschetz fibrations and exotic symplectic structures on cotangent bundles of spheres*, J. Topol. **3** (2010), no. 1, 157–180.

[MS12] Dusa McDuff and Dietmar Salamon, *J-holomorphic curves and symplectic topology*, Second, American Mathematical Society Colloquium Publications, vol. 52, American Mathematical Society, Providence, RI, 2012.

[MS17] ———, *Introduction to symplectic topology*, Oxford University Press, 2017.

[MT18] Davesh Maulik and Yukinobu Toda, *Gopakumar–Vafa invariants via vanishing cycles*, Inventiones mathematicae **213** (2018), no. 3, 1017–1097.

[Mun16] James R Munkres, *Elementary differential topology.(am-54), volume 54*, Princeton University Press, 2016.

[MW12] Ciprian Manolescu and Christopher Woodward, *Floer homology on the extended moduli space*, Perspectives in analysis, geometry, and topology, 2012, pp. 283–329.

[MW17] Dusa McDuff and Katrin Wehrheim, *Smooth Kuranishi atlases with isotropy*, Geom. Topol. **21** (2017), no. 5, 2725–2809.

[Nad17] David Nadler, *Arboreal singularities*, Geometry & Topology **21** (2017), no. 2, 1231–1274.

[Nee01] Amnon Neeman, *Triangulated categories*, Annals of Mathematics Studies, vol. 148, Princeton University Press, Princeton, NJ, 2001.

[Nic95] Liviu I. Nicolaescu, *The Maslov index, the spectral flow, and decompositions of manifolds*, Duke Math. J. **80** (1995), no. 2, 485–533.

[NS17] Frank Neumann and Markus Szymik, *Spectral sequences for Hochschild cohomology and graded centers of derived categories*, Selecta Math. (N.S.) **23** (2017), no. 3, 1997–2018.

[NS20] David Nadler and Vivek Shende, *Sheaf quantization in weinstein symplectic manifolds*, arXiv preprint arXiv:2007.10154 (2020).

[Orl09a] Dmitri Orlov, *Derived categories of coherent sheaves and triangulated categories of singularities*, Algebra, arithmetic, and geometry: in honor of Yu. I. Manin. Vol. II, 2009, pp. 503–531.

[Orl09b] ———, *Remarks on generators and dimensions of triangulated categories*, Mosc. Math. J. **9** (2009), no. 1, 153–159, back matter.

[Par13] Brett Parker, *Integral counts of pseudo-holomorphic curves*, 2013. arXiv: 1309.0585.

[Par16] John Pardon, *An algebraic approach to virtual fundamental cycles on moduli spaces of pseudo-holomorphic curves*, Geom. Topol. **20** (2016), no. 2, 779–1034.

[Par19] ———, *Enough vector bundles on orbispaces*, arXiv preprint arXiv:1906.05816 (2019).

[Par20] ———, *Orbifold bordism and duality for finite orbispectra*, arXiv preprint arXiv:2006.12702 (2020).

[Pir19] Dmitrii Pirozhkov, *Rouquier dimension of some blow-ups*, arXiv preprint arXiv:1908.08283 (2019).

[Pom21] Daniel Pomerleano, *Intrinsic mirror symmetry and categorical crepant resolutions*, arXiv preprint arXiv:2103.01200 (2021).

[Poz99] Marcin Pozniak, *Floer homology, Novikov rings and clean intersections*, Northern california symplectic geometry seminar, 1999, pp. 119–181.

[PT09] R. Pandharipande and R. P. Thomas, *Curve counting via stable pairs in the derived category*, Invent. Math. **178** (2009), no. 2, 407–447.

[Rit13] Alexander F Ritter, *Topological quantum field theory structure on symplectic cohomology*, Journal of Topology **6** (2013), no. 2, 391–489.

[Rou08] Raphaël Rouquier, *Dimensions of triangulated categories*, J. K-Theory **1** (2008), no. 2, 193–256.

[RS17] Alexander F. Ritter and Ivan Smith, *The monotone wrapped Fukaya category and the open-closed string map*, Selecta Math. (N.S.) **23** (2017), no. 1, 533–642.

[RT95] Yongbin Ruan and Gang Tian, *A mathematical theory of quantum cohomology*, Journal of Differential Geometry **42** (1995), 259–367.

[Rua99] Yongbin Ruan, *Virtual neighborhoods and pseudo-holomorphic curves*, Proceedings of 6th Gökova Geometry-Topology Conference, 1999, pp. 161–231.

[Sav11] Nikolai Saveliev, *Lectures on the topology of 3-manifolds: an introduction to the Casson invariant*, Walter de Gruyter, 2011.

[Sch18] Stefan Schwede, *Global homotopy theory*, New Mathematical Monographs, vol. 34, Cambridge University Press, Cambridge, 2018.

[See14] Nora Seeliger, *Loop homology of spheres and complex projective spaces*, Forum Math. **26** (2014), no. 4, 967–981.

[Sei08a] Paul Seidel, *A biased view of symplectic cohomology*, Current developments in mathematics, 2006, 2008, pp. 211–253.

[Sei08b] ———, *Fukaya categories and Picard-Lefschetz theory*, Zurich Lectures in Advanced Mathematics, European Mathematical Society (EMS), Zürich, 2008.

[Sei18] ———, *Fukaya A_∞-structures associated to Lefschetz fibrations. VI*, arXiv preprint arXiv:1810.07119 (2018).

[Sei] ———, *Lectures on categorical dynamics and symplectic topology*, Version 2.0.

[She17] Vivek Shende, *Microlocal category for Weinstein manifolds via h-principle*, arXiv preprint arXiv:1707.07663 (2017).

[SL91] Reyer Sjamaar and Eugene Lerman, *Stratified symplectic spaces and reduction*, Annals of Mathematics (1991), 375–422.

[ST16] Vivek Shende and Alex Takeda, *Calabi–Yau structures on topological Fukaya categories*, arXiv preprint arXiv:1605.02721 (2016).

[Sta18] Laura Starkston, *Arboreal singularities in weinstein skeleta*, Selecta Mathematica **24** (2018), no. 5, 4105–4140.

[Sta21] The Stacks project authors, *The stacks project* (2021).

[SW08] Dietmar Salamon and Katrin Wehrheim, *Instanton Floer homology with Lagrangian boundary conditions*, Geometry & Topology **12** (2008), no. 2, 747–918.

[Swa21] Mohan Swaminathan, *Rel–C^∞ structures on Gromov–Witten moduli spaces*, Journal of Symplectic Geometry **19** (2021), no. 2, 413–473.

[Syl19a] Zachary Sylvan, *On partially wrapped Fukaya categories*, J. Topol. **12** (2019), no. 2, 372–441.

[Syl19b] ———, *Orlov and Viterbo functors in partially wrapped Fukaya categories*, arXiv preprint arXiv:1908.02317 (2019).

[Tab15] Gonçalo Tabuada, *Noncommutative motives*, University Lecture Series, vol. 63, American Mathematical Society, Providence, RI, 2015. With a preface by Yuri I. Manin.

[Tab16] ———, *A note on secondary K-theory*, Algebra Number Theory **10** (2016), no. 4, 887–906.

[Tab18] Goncalo Tabuada, *Noncommutative Weil conjecture*, arXiv preprint arXiv:1808.00950 (2018).

[Tau90] Clifford H Taubes, *Casson's invariant and gauge theory*, Journal of Differential Geometry **31** (1990), no. 2, 547–599.

[Tau96a] ———, *Counting pseudo-holomorphic submanifolds in dimension 4*, Journal of Differential Geometry **44** (1996), no. 4, 818–893.

[Tau96b] Clifford Henry Taubes, *Counting pseudo-holomorphic submanifolds in dimension* 4, J. Differential Geom. **44** (1996), no. 4, 818–893.

[Tho00] R. P. Thomas, *A holomorphic Casson invariant for Calabi-Yau 3-folds, and bundles on $K3$ fibrations*, J. Differential Geom. **54** (2000), no. 2, 367–438.

[Tho64] René Thom, *Local topological properties of differentiable mappings*, Colloquium on differential analysis, 1964, pp. 191–202.

[Tho69] ———, *Ensembles et morphismes stratifiés*, Bulletin of the American Mathematical Society **75** (1969), 240–284.

[Tro78] David Trotman, *Stability of transversality to a stratification implies Whitney (a)-regularity*, Inventiones Mathematicae **50** (1978), no. 3, 273–277.

[TX18] Gang Tian and Guangbo Xu, *Gauged linear sigma model in geometric phases*, 2018. `https://arxiv.org/abs/1809.00424`.

[TX21] ———, *Counting pointlike instantons without gluing*, 2021. `https://arxiv.org/abs/2110.15379`.

[Vak17] Ravi Vakil, *The rising sea: Foundations of algebraic geometry*, preprint (2017).

[Vit99] C. Viterbo, *Functors and computations in Floer homology with applications. i*, Geom. Funct. Anal. **9** (1999), no. 5, 985–1033.

[Voi96] Claire Voisin, *A mathematical proof of a formula of Aspinwall and Morrison*, Compositio Math. **104** (1996), no. 2, 135–151.

[Wal92a] Kevin Walker, *An extension of Casson's invariant*, Princeton University Press, 1992.

[Wal92b] ———, *An extension of Casson's invariant*, Annals of Mathematics Studies, vol. 126, Princeton University Press, Princeton, NJ, 1992.

[Was69] Arthur G. Wasserman, *Equivariant differential topology*, Topology **8** (1969), 127–150.

[Wat73] Keiichi Watanabe, *Some examples of one dimensional gorenstein domains*, Nagoya Math. J. **49** (1973), 101–109.

[Wen21] Chris Wendl, *Transversality and super-rigidity for multiply covered holomorphic curves*, 2021.

[Whi65] Hassler Whitney, *Tangents to an analytic variety*, Annals of Mathematics **81** (1965), 496–549.

[Yan16] Song Yang, *A note on the rouquier dimensions of product varieties*, Journal of Algebra and Its Applications **15** (2016), no. 04, 1650065.

[Zho20] Peng Zhou, *Lagrangian skeleta of hypersurfaces in* $(\mathbb{C}^*)^n$, Selecta Math. (N.S.) (2020).

[Zin08] Aleksey Zinger, *Pseudocycles and integral homology*, Transactions of the American Mathematical Society **360** (2008), 2741–2765.

[Zin11] ———, *A comparison theorem for Gromov–Witten invariants in the symplectic category*, Advances in Mathematics **228** (2011), no. 1, 535–574.