

An algebraic approach to virtual fundamental cycles on moduli spaces of pseudo-holomorphic curves

JOHN PARDON

We develop techniques for defining and working with virtual fundamental cycles on moduli spaces of pseudo-holomorphic curves which are not necessarily cut out transversally. Such techniques have the potential for applications as foundations for invariants in symplectic topology arising from “counting” pseudo-holomorphic curves.

We introduce the notion of an *implicit atlas* on a moduli space, which is (roughly) a convenient system of local finite-dimensional reductions. We present a general intrinsic strategy for constructing a canonical implicit atlas on any moduli space of pseudo-holomorphic curves. The main technical step in applying this strategy in any particular setting is to prove appropriate gluing theorems. We require only topological gluing theorems, that is, smoothness of the transition maps between gluing charts need not be addressed. Our approach to virtual fundamental cycles is algebraic rather than geometric (in particular, we do not use perturbation). Sheaf-theoretic tools play an important role in setting up our functorial algebraic “VFC package”.

We illustrate the methods we introduce by giving definitions of Gromov–Witten invariants and Hamiltonian Floer homology over \mathbb{Q} for general symplectic manifolds. Our framework generalizes to the S^1 -equivariant setting, and we use S^1 -localization to calculate Hamiltonian Floer homology. The Arnold conjecture (as treated by Floer, by Hofer and Salamon, by Ono, by Liu and Tian, by Ruan, and by Fukaya and Ono) is a well-known corollary of this calculation.

37J10, 53D35, 53D40, 53D45, 57R17; 53D37, 53D42, 54B40

1 Introduction

In this paper, we develop a collection of tools and techniques for defining and working with virtual fundamental cycles on compact moduli spaces of pseudo-holomorphic curves (in the sense of Gromov [36]) which are not necessarily cut out transversally. Such techniques have a myriad of potential applications in symplectic geometry by providing foundations for invariants obtained by “counting” pseudo-holomorphic curves:

$$(1.0.1) \quad \text{Symplectic manifold} \implies \text{Moduli space(s) of pseudo-holomorphic curves} \xrightarrow{\text{VFC}} \text{Desired invariant}$$

In this paper, we build a general framework which can potentially be applied to give rigorous foundations for the wide variety of invariants defined using (1.0.1). We hope that this framework may also be applicable to moduli spaces of solutions to other nonlinear elliptic PDEs which give rise to interesting invariants.

When a moduli space is not cut out transversally, its topological structure does not determine its virtual fundamental cycle; rather it must be endowed (canonically) with some additional extra structure. We introduce the notion of an *implicit atlas* on a compact Hausdorff space, which serves as this extra structure on moduli spaces of pseudo-holomorphic curves. We use implicit atlases as a layer of abstraction between the two steps in (1.0.1), making them logically independent.

Our notion of an implicit atlas and our constructions of implicit atlases on moduli spaces of pseudo-holomorphic curves constitute a reworking of existing ideas, with convenient canonicity and functoriality properties. Our construction of virtual fundamental cycles from implicit atlases is more novel (using algebraic rather than geometric methods), and also has good functoriality properties which are useful in applications. It is noteworthy that this algebraic VFC setup requires only topological gluing theorems as input.

The basic idea of using an atlas of charts of the form (1.1.1) on a moduli space to construct its virtual fundamental cycle has existed since the inception of this problem; see for example Li and Tian [53], Liu and Tian [54], Fukaya and Ono [33], Ruan [71], Lu and Tian [56], Fukaya, Oh, Ohta and Ono [29; 30; 31; 32] (the theory of Kuranishi structures), and McDuff and Wehrheim [63; 60]. The polyfolds project of Hofer, Wysocki and Zehnder [41; 42; 43; 44; 45; 46; 47] gives another method for defining virtual fundamental cycles by describing moduli spaces via a generalized infinite-dimensional Fredholm setup.

1.1 Implicit atlases

An implicit atlas organizes together a collection of local charts for a compact moduli space X . A local chart for X is a diagram

$$(1.1.1) \quad X \xleftarrow{\text{open}} s_\alpha^{-1}(0) \xrightarrow{\text{closed}} X_\alpha \xrightarrow{s_\alpha} E_\alpha,$$

where E_α is a finite-dimensional vector space (called the *obstruction space*), X_α is an auxiliary moduli space (called the α -*thickened*¹ *moduli space*), and s_α is called the *Kuranishi map*. For the purpose of constructing the virtual fundamental cycle of X ,

¹Perhaps a better name would be “ α -stabilized moduli space”, though we have decided not to risk confusing this notion of stabilization (ie product with a vector space) and the notion of stabilizing a Riemann surface (ie adding marked points).

such a local chart (1.1.1) is useful over $X \cap X_\alpha^{\text{reg}}$, where $X_\alpha^{\text{reg}} \subseteq X_\alpha$ (called the *regular locus*) is the locus where X_α is cut out transversally (and thus, in particular, X_α^{reg} is a finite-dimensional manifold).

An implicit atlas is an index set A (whose elements are called *thickening datums*) along with obstruction spaces E_α (for all $\alpha \in A$) and I -thickened moduli spaces X_I (for all finite subsets $I \subseteq A$) fitting together globally in a natural generalization of (1.1.1), where the \emptyset -thickened moduli space X_\emptyset is identified with the original moduli space X . An implicit atlas also includes the data of open subsets $X_I^{\text{reg}} \subseteq X_I$ which are manifolds and are required to cover all of X . In particular, an implicit atlas carries a parameter $d \in \mathbb{Z}$, the *virtual dimension*, and we require that $\dim X_I^{\text{reg}} = d + \dim E_I$ for all $I \subseteq A$ (where $E_I := \bigoplus_{\alpha \in I} E_\alpha$). Implicit atlases also allow charts (1.1.1) which incorporate the action of a finite group Γ_α (so that such charts exist on spaces X with nontrivial isotropy), though we will introduce the necessary notation later in the paper.

There is a natural notion of an “implicit atlas with boundary” (or corners) and of the “product implicit atlas” on a product of spaces equipped with implicit atlases (with boundary/corners). These notions enable us to treat Floer-type homology theories via implicit atlases.

We use only “topological” implicit atlases in this paper (ie we only require that the X_I^{reg} are topological manifolds), since the topological structure is sufficient to construct virtual fundamental cycles. There is, of course, a parallel notion of a smooth implicit atlas, which we will not need here.

Remark 1.1.1 From a theoretical standpoint, it would be desirable to endow the moduli space X with the canonical structure of a “derived manifold” (a notion which should be more intrinsic than the notion of an implicit atlas on X). A good notion of a “derived smooth manifold” exists (due to Spivak [75], Borisov and Noel [10] and Joyce [50; 49]), and it is reasonable to expect that a parallel topological theory exists as well. By nature, the theory of derived manifolds uses the language of higher category theory. An implicit atlas on X may be thought of as giving a “presentation” of X as a derived manifold (just as a collection of open sets of \mathbb{R}^n and gluing data can be used to present an ordinary manifold).

The “atlas” approach which we follow here, while being less intrinsic, has the advantage of being more concrete and more elementary. We believe that a more intrinsic approach is unlikely to lead to any simplification in the construction of this extra structure (implicit atlas or derived manifold structure) on moduli spaces of pseudo-holomorphic curves. However, it would likely make it easier to work with (and, in particular, calculate) virtual fundamental cycles on such spaces.

1.2 Construction of implicit atlases

Implicit atlases are designed to encode a system of charts which can be constructed naturally and intrinsically on moduli spaces of pseudo-holomorphic curves in wide generality. Moreover, the basic ingredients which go into the construction of implicit atlases are all familiar in the field. We think of our specific examples of constructions of implicit atlases as special cases of a general strategy which produces on any moduli space of pseudo-holomorphic curves a canonical implicit atlas. For this general strategy to succeed, there are (essentially) two steps which require setting-specific arguments.

The first step requiring setting-specific arguments is *domain stabilization*. For any pseudo-holomorphic curve $u: C \rightarrow M$ in X , we must show that there is a smooth codimension-two submanifold (possibly with boundary) $D \subseteq M$ which intersects it transversally such that adding $u^{-1}(D)$ as added marked points on C makes C stable.² This is an important ingredient in verifying the *covering axiom* of an implicit atlas.

The second step requiring setting-specific arguments is *formal regularity implies topological regularity*. For the I -thickened moduli spaces X_I , we denote by $X_I^{\text{reg}} \subseteq X_I$ the locus where the relevant linearized operator is surjective. We must show that X_I^{reg} is open and is a topological manifold of the correct dimension (we also need a certain topological submersion condition on how X_I^{reg} is cut out inside X_J^{reg} for $I \subseteq J$). These are the *openness and submersion axioms* of an implicit atlas. This is the step where we must appeal to serious analytic results (in particular, this is where gluing of pseudo-holomorphic curves takes place). Note, though, that we only require topological gluing results (ie smoothness of transition maps between gluing charts need not be addressed). We should also point out that these axioms are local statements about the spaces X_I^{reg} , and hence are independent of any auxiliary group action.

Our general strategy constructs an implicit atlas on X which is *canonical* (in the sense that we do not need to make any choices during its construction). This is achieved simply by defining A to be the set of *all* possible thickening datums (of which there are uncountably many), where a thickening datum is a choice of divisor D , an obstruction space E , plus some additional data. Note that it is nontrivial to formulate a notion of “atlas” which allows this type of “universal” construction. Having such a canonical procedure is very useful in a number of places, and this aspect of our approach appears to be new.

Our general strategy works well for the construction of S^1 -equivariant implicit atlases on the moduli spaces of stable Floer trajectories for S^1 -invariant Hamiltonians (a key

²One can sometimes get away with a little less if this specific form of domain stabilization does not hold. We will encounter such a situation when constructing S^1 -equivariant implicit atlases on moduli spaces of stable Floer trajectories for S^1 -invariant Hamiltonians.

step towards the Arnold conjecture). The domain stabilization step is harder in the equivariant setting (we must use a divisor inside M instead of inside $M \times S^1$, and we must be satisfied with not stabilizing Morse-type components of trajectories), but is not difficult. That formal regularity implies topological regularity is the *same* as in the non-equivariant case: the analysis is *independent of the S^1 -equivariance*, and thus is (essentially) identical to that in the non-equivariant case. In particular, we never need to check that S^1 acts smoothly on anything.

1.3 Construction of virtual fundamental cycles

We develop a “VFC package” for any space X equipped with an implicit atlas A . The tools we develop are primarily *algebraic* (chain complexes and sheaves) rather than geometric or topological. Furthermore, these tools have nice functorial properties which allow them to be applied essentially independently of how their internals are constructed. As mentioned earlier, this set of tools does not require a smooth structure on the atlas A .

Let us now discuss the components of the VFC package. The main object we develop is the *virtual cochain complex* $C_{\text{vir}}^\bullet(X; A)$ defined whenever A is *finite*.³ It comes with a canonical isomorphism⁴

$$(1.3.1) \quad \check{H}^\bullet(X) \xrightarrow{\sim} H_{\text{vir}}^\bullet(X; A)$$

(H_{vir}^\bullet is the homology of C_{vir}^\bullet) and with a canonical map

$$(1.3.2) \quad C_{\text{vir}}^{d+\bullet}(X; A) \xrightarrow{s_*} C_{\dim E_A - \bullet}(E_A, E_A \setminus 0)$$

(where $E_A := \bigoplus_{\alpha \in A} E_\alpha$ and d is the virtual dimension of A). Since

$$H_{\dim E_A}(E_A, E_A \setminus 0) = \mathbb{Z},$$

combining (1.3.1) and (1.3.2) yields a map $\check{H}^d(X) \rightarrow \mathbb{Z}$. We define the virtual fundamental class to be this element $[X]_A^{\text{vir}} \in \check{H}^d(X)^\vee$; if X is cut out transversally (that is, $X = X^{\text{reg}}$), then it agrees with the usual fundamental class of X as a closed manifold. This complex C_{vir}^\bullet generalizes naturally to “implicit atlases with boundary” as well.

We use C_{vir}^\bullet for much more than just defining the virtual fundamental class. Since it is a complex (rather than just a sequence of homology groups), it is sufficiently rich to provide a useful notion of virtual fundamental *cycle* (more precisely, the map s_* (1.3.2)

³If $A \subseteq A'$, then $C_{\text{vir}}^\bullet(X; A)$ and $C_{\text{vir}}^\bullet(X; A')$ are equivalent for the purposes of defining virtual fundamental cycles, so given an infinite implicit atlas we can just work with any choice of finite subatlas.

⁴In the present discussion, we ignore issues about orientations.

can be thought of as the chain-level virtual fundamental cycle). This enables us to use the VFC package to treat Floer-type homology theories, which requires something like a “coherent system of virtual fundamental cycles” over a large system of spaces (a “flow category”) equipped with implicit atlases.

Let us now explain and motivate the definition of C_{vir}^\bullet . First, let us imagine we have a single chart (1.1.1) which is global (ie $X = s_\alpha^{-1}(0)$) and $s_\alpha^{-1}(0) \subseteq X_\alpha^{\text{reg}}$. Then we consider the diagram

$$(1.3.3) \quad \check{H}^{\dim X_\alpha^{\text{reg}} - \dim E_\alpha}(X) \xrightarrow{\sim} H_{\dim E_\alpha}(X_\alpha^{\text{reg}}, X_\alpha^{\text{reg}} \setminus X) \xrightarrow{(s_\alpha)_*} H_{\dim E_\alpha}(E_\alpha, E_\alpha \setminus 0) = \mathbb{Z}.$$

The first isomorphism is Poincaré–Lefschetz duality. The virtual fundamental class is the resulting element $[X]^{\text{vir}} \in \check{H}^d(X)^\vee$ (recall $d = \dim X_\alpha^{\text{reg}} - \dim E_\alpha$) which is easily seen to agree with the class defined via perturbation. Thus in this case the complex $C_{\dim X_\alpha^{\text{reg}} - \bullet}(X_\alpha^{\text{reg}}, X_\alpha^{\text{reg}} \setminus X)$ plays the role of $C_{\text{vir}}^\bullet(X; A)$.

In general, $C_{\text{vir}}^\bullet(X; A)$ (which may be thought of as a “global finite-dimensional reduction up to homotopy”) is built out of the singular chain complexes of the X_I^{reg} (or, more accurately, of some auxiliary spaces $X_{I,J,A}$ defined from the X_I^{reg} using a *deformation to the normal cone*). We glue together these singular chain complexes using an appropriate homotopy colimit (a generalized sort of mapping cone); this is technically more convenient than gluing together the spaces themselves. The latter could probably be made to work as well, as long as one is careful to glue using homotopy colimits⁵ (one can run into point-set topological issues with certain natural topological gluings).⁶

We use the language of sheaves to give an especially efficient construction of the key isomorphism (1.3.1). In particular, we reduce (1.3.1) to a statement which can be checked *locally* on X . In a word, we define a “homotopy \mathcal{K} -sheaf” $K \mapsto C_{\text{vir}}^\bullet(K; A)$ on X , and we show that the stalk cohomology $H_{\text{vir}}^\bullet(\{p\}; A)$ is isomorphic to \mathbb{Z} concentrated in degree zero. The isomorphism (1.3.1) then follows from rather general sheaf-theoretic arguments (and in fact there is a corresponding map of complexes). It

⁵Here is a baby example of how one may use mapping cones to bypass point-set topological issues. The long exact sequence of the pair $\cdots \rightarrow H_\bullet(A) \rightarrow H_\bullet(X) \rightarrow H_\bullet(X, A) \rightarrow \cdots$ is valid for an arbitrary pair (X, A) of topological spaces (meaning $A \subseteq X$ has the subspace topology), and relative homology $H_\bullet(X, A)$ is always naturally isomorphic to the reduced homology of the mapping cone $(cA \sqcup X)/\sim$. On the other hand, understanding $H_\bullet(X/A)$ usually requires some niceness assumptions on (X, A) (to ensure that the natural map $H_\bullet(X, A) \rightarrow H_\bullet(X/A, \text{pt})$ is an isomorphism). On the other hand, if one is content working with $H_\bullet(X, A)$ or with the mapping cone, then such niceness assumptions are unnecessary.

⁶McDuff and Wehrheim [63, Examples 3.1.14 and 3.1.15] give examples of natural topological quotients which fail to be Hausdorff, fail to be locally compact, or fail to be locally metrizable.

should not be surprising that sheaves can be used effectively in this setting, since the problem we are facing is precisely to patch together local homological information (from charts (1.1.1)) into global homological information.

Moreover, we find in this paper that this sheaf-theoretic formalism continues to play a key role in the study and application of the complexes C_{vir}^\bullet , beyond simply constructing the fundamental isomorphism (1.3.1). Hence we believe that the sheaf-theoretic formalism is of more importance than the precise manner of definition of C_{vir}^\bullet . In particular, checking commutativity of certain diagrams of homology groups can often be reduced to checking that a certain corresponding diagram of sheaves commutes (which is then just a *local* calculation). This is a key proof technique in many places where we use the VFC package. This is perhaps surprising since the virtual fundamental class itself has no local homological characterization (though see Remark 1.3.3).

Remark 1.3.1 Though our definition of the virtual fundamental cycle does not involve perturbation, we do in some sense show that perturbation is a valid way to compute the virtual fundamental class (in fact, this is an easy corollary of some of its formal properties).

Remark 1.3.2 In this paper, we only construct virtual fundamental cycles in ordinary homology, which is thus inadequate for applications of moduli spaces of pseudo-holomorphic curves involving their fundamental class in smooth (or framed) bordism (eg as in Abouzaid [1] or Ekholm and Smith [20]). However, in discussions with Abouzaid, we realized that by working on the space level, one can probably upgrade our construction to yield a virtual fundamental cycle in the (Steenrod) framed bordism group of X , twisted by a certain “virtual spherical normal bundle”.

Remark 1.3.3 In broad philosophical outline, the strategy we follow to define the virtual fundamental cycle is the following. The functor $U \mapsto C_\bullet(U \text{ rel } \partial U)$ (or Borel–Moore chains $C_\bullet^{\text{BM}}(U)$) is a homotopy sheaf on X , and the virtual fundamental cycle is a global section thereof. Hence we may define the virtual fundamental cycle by *specifying it on a convenient finite open cover and giving patching data on all higher overlaps* (in our case, the particular cover being $X = \bigcup_{I \subseteq A} \psi_{\emptyset I}((s_I | X_I^{\text{reg}})^{-1}(0))$ for finite A). This philosophy is very natural and could apply quite generally.

Remark 1.3.4 The language of ∞ –categories as developed by Lurie [57] seems to be a natural setting for the VFC package and, more generally, for a good theory of derived manifolds and their virtual fundamental cycles (indeed, the existing theory of derived smooth manifolds is by necessity written in the language of higher category theory). We have avoided ∞ –categories in this paper for the sake of concreteness (though at

the cost of needing to use lots of explicit homotopy colimits). However, we expect that in this more abstract framework, one could ultimately develop the most flexible calculational tools (computing virtual fundamental cycles directly from our definition seems rather difficult, due to the inexplicit nature of the isomorphism (1.3.1)).

1.4 Example applications

We use the framework developed in this paper to give new VFC foundations for classical results which rely on virtual moduli cycle techniques. We define Gromov–Witten invariants for general symplectic manifolds (originally due in this generality to Li and Tian [53], Fukaya and Ono [33] and Ruan [71]). We define Hamiltonian Floer homology over \mathbb{Q} for general closed symplectic manifolds, and we use S^1 -localization methods to calculate Hamiltonian Floer homology (originally due in this generality to Liu and Tian [54], Fukaya and Ono [33] and Ruan [71]). The Arnold conjecture on Hamiltonian fixed points is a standard corollary of this calculation.

We hope that the examples we treat here may persuade the reader that it is reasonable to expect to be able to construct implicit atlases on moduli spaces of pseudo-holomorphic curves in considerable generality, and hence that our VFC package is applicable to other curve counting invariants in symplectic topology.

1.5 How to read this paper

The rough logical dependence between the sections of this paper is illustrated in Figure 1. Solid arrows indicate a strong dependency (logical and conceptual); dashed arrows indicate isolated logical dependence (quoting results as black boxes). Obviously, the choice of solid/dashed is rather subjective. In the bottom row, we develop the VFC package and apply it abstractly to spaces equipped with implicit atlases. In the middle row, we construct implicit atlases on various moduli spaces of pseudo-holomorphic curves and define the desired invariants by applying the abstract results developed in the bottom row. The top row contains gluing results necessary to prove certain key axioms for the implicit atlases constructed in the middle row.

We now summarize the contents of each section, from which the reader may decide which sections to read in detail.

Section 2 is of an introductory flavor (it does not contain any definitions or results to be used elsewhere); it aims to develop technical intuition for our approach without getting bogged down in details. We give a simplified definition of an implicit atlas, and we give some simple examples. We give some prototypical constructions of implicit atlases in simplified settings. We define the virtual fundamental class algebraically from some

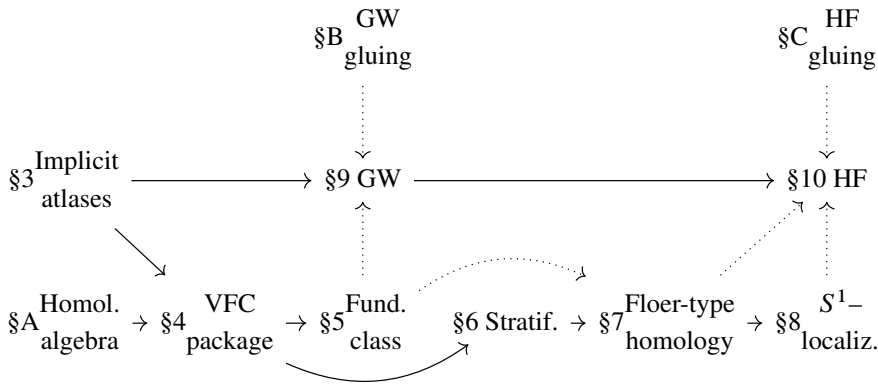


Figure 1

simple implicit atlases. We also give a simplified outline of how we apply the VFC package to construct Floer-type homology theories from moduli spaces equipped with implicit atlases.

In Section 3, we give the definition of an implicit atlas (and its variant with boundary).

In Section 4, we construct the VFC package. For a space X equipped with a finite implicit atlas A with boundary, we define and study the *virtual cochain complexes* $C_{\text{vir}}^\bullet(X; A)$ and $C_{\text{vir}}^\bullet(X \text{ rel } \partial; A)$. The algebraic and sheaf-theoretic foundations from Appendix A play a key role.

In Section 5, we use the VFC package to define the virtual fundamental class. We also derive some of its basic properties and provide some calculation tools.

In Section 6, we introduce and study implicit atlases with stratification. We use the VFC package to obtain an inductive “stratum-by-stratum” understanding of virtual fundamental cycles.

In Section 7, we use the VFC package to define homology groups from a suitably compatible system of implicit atlases on a “flow category”. We also use the results concerning stratifications from Section 6.

In Section 8, we construct an S^1 -equivariant VFC package and use it to prove some S^1 -localization results for the virtual fundamental classes of Section 5 and the homology groups of Section 7.

In Section 9, we construct Gromov–Witten invariants by constructing an implicit atlas on the moduli space of stable pseudo-holomorphic maps and using the results of Section 5.

In Section 10, we construct Hamiltonian Floer homology by constructing implicit atlases on the moduli spaces of stable Floer trajectories and using the results of Section 7. We also construct S^1 -equivariant implicit atlases on the moduli spaces for time-independent Hamiltonians. Applying the results of Section 8, we calculate Hamiltonian Floer homology and thus deduce the Arnold conjecture.

In Appendix A, we review and develop foundational results about sheaves, homotopy sheaves, Čech cohomology, Poincaré duality, and homotopy colimits.

In Appendix B, we prove the gluing theorem needed in the construction of implicit atlases on Gromov–Witten moduli spaces in Section 9.

In Appendix C, we prove the gluing theorem needed in the construction of implicit atlases on Hamiltonian Floer moduli spaces in Section 10.

Acknowledgements I owe much thanks to my advisor Yasha Eliashberg for suggesting this problem and for many useful conversations. I also thank Mohammed Abouzaid, Dusa McDuff, and the anonymous referees in particular for many specific and detailed comments on this work. I am grateful for the useful exchanges I had with Tobias Ekholm, Søren Galatius, Helmut Hofer, Michael Hutchings, Eleny Ionel, Dominic Joyce, Daniel Litt, Patrick Massot, Rafe Mazzeo, Paul Seidel, Dennis Sullivan, Arnav Tripathy, Ravi Vakil, and Katrin Wehrheim. I thank the Simons Center for Geometry and Physics for its hospitality during visits where I presented and discussed this work. This paper is part of the author’s PhD thesis at Stanford University.

The author was partially supported by a National Science Foundation Graduate Research Fellowship under grant number DGE–1147470.

2 Technical introduction

We now give a more technical introduction to the main ideas of this paper. We feel free to make simplifying assumptions for the sake of clarity of exposition, though we do aim to highlight some important technical points. A full treatment is deferred to the body of the paper, where everything is properly defined.

In Section 2.1, we familiarize the reader with the notion of an implicit atlas. In Section 2.2, we give some prototypical constructions of implicit atlases. In Section 2.3, we explain how to construct the virtual fundamental cycle from an implicit atlas in some simple cases. In Section 2.4, we show how our methods can be applied to construct Floer-type homology theories. In Section 2.5, we explain a rudimentary S^1 -localization result for virtual fundamental cycles.

2.1 Implicit atlases

2.1.1 Implicit atlases (proto version) We now introduce (a simplified version of) implicit atlases and give some intuition for what they mean geometrically. The impatient reader may also wish to refer to the true definition of an implicit atlas in Section 3 (Definition 3.1.1). Here we have simplified things by assuming $\Gamma_\alpha = 1$ (“trivial covering groups”), $X_I^{\text{reg}} = X_I$ for $\#I \geq 1$, and $X_\emptyset^{\text{reg}} = \emptyset$.

After the giving the definition (which appears rather complicated at first), we explain it further with examples.

Definition 2.1.1 (Implicit atlas (proto version)) Let X be a compact Hausdorff space. A (proto) implicit atlas of dimension d on X is an index set A along with the following data:

- (i) **Obstruction spaces** A finite-dimensional \mathbb{R} vector space E_α for all $\alpha \in A$ (let $E_I := \bigoplus_{\alpha \in I} E_\alpha$).
- (ii) **Thickenings** A Hausdorff space X_I for all finite $I \subseteq A$, and a homeomorphism $X \rightarrow X_\emptyset$.
- (iii) **Kuranishi maps** A function $s_\alpha: X_I \rightarrow E_\alpha$ for all $\alpha \in I \subseteq A$ (for $I \subseteq J$, let $s_I: X_J \rightarrow E_I$ denote $\bigoplus_{\alpha \in I} s_\alpha$).
- (iv) **Footprints** An open set $U_{IJ} \subseteq X_I$ for all $I \subseteq J \subseteq A$.
- (v) **Footprint maps** A function $\psi_{IJ}: (s_{J \setminus I}|X_J)^{-1}(0) \rightarrow U_{IJ}$ for all $I \subseteq J \subseteq A$.

These data must satisfy the following “compatibility axioms”:

- (i) $\psi_{IJ}\psi_{JK} = \psi_{IK}$ and $\psi_{II} = \text{id}$.
- (ii) $s_I\psi_{IJ} = s_I$.
- (iii) $U_{IJ_1} \cap U_{IJ_2} = U_{I, J_1 \cup J_2}$ and $U_{II} = X_I$.
- (iv) $\psi_{IJ}^{-1}(U_{IK}) = U_{JK} \cap (s_{J \setminus I}|X_J)^{-1}(0)$.
- (v) **Homeomorphism axiom** ψ_{IJ} is a homeomorphism.

They must also satisfy the following “transversality axioms”:

- (vi) **Submersion axiom** $s_{J \setminus I}: X_J \rightarrow E_{J \setminus I}$ is locally modeled on the projection $\mathbb{R}^{d + \dim E_J} \rightarrow \mathbb{R}^{\dim E_{J \setminus I}}$ over $0 \in E_{J \setminus I}$ for $\#I \geq 1$ (in particular, taking $I = J$ implies that X_I is a topological manifold of dimension $d + \dim E_I$ for $\#I \geq 1$).
- (vii) **Covering axiom** $X_\emptyset = \bigcup_{\emptyset \neq I \subseteq A} \psi_{\emptyset I}((s_I|X_I)^{-1}(0))$.




Figure 2: Illustrations of implicit atlases. Left: $A = \{\alpha\}$, $d = 1$, $E_\alpha = \mathbb{R}$. Illustrated are X (black) and X_α (red; light/dark according to the sign of s_α). Right: $A = \{\alpha, \beta\}$, $d = 1$, $E_\alpha = E_\beta = \mathbb{R}$. Illustrated are X (black), X_α (red; light/dark according to the sign of s_α), X_β (blue; light/dark according to the sign of s_β), and $X_{\alpha\beta}$ (green boundary).

Let us unpack this definition a bit. We have manifolds X_α (meaning X_I for $I = \{\alpha\}$) indexed by $\alpha \in A$. Each is equipped with a function $s_\alpha: X_\alpha \rightarrow E_\alpha$ and a homeomorphism $\psi_\alpha: s_\alpha^{-1}(0) \rightarrow U_\alpha$ (meaning $\psi_{\emptyset\{\alpha\}}: (s_\alpha|_{X_\alpha})^{-1}(0) \rightarrow U_{\emptyset\{\alpha\}}$) where $X = \bigcup_{\alpha \in A} U_\alpha$ is an open cover. Thus so far this is nothing more than a set of charts (the “basic charts”) of a particular form (1.1.1) covering X and indexed by $\alpha \in A$. If $A = \{\alpha\}$, then the implicit atlas is simply a single global chart (1.1.1), and this is illustrated in Figure 2, left.

Now for every pair of basic charts $\alpha, \beta \in A$, there is an “overlap chart” $X_{\alpha\beta}$ with footprint $U_\alpha \cap U_\beta$ and obstruction space $E_\alpha \oplus E_\beta$. Furthermore, (open subsets $U_{\alpha,\alpha\beta}$ and $U_{\beta,\alpha\beta}$ of) the original charts X_α and X_β are identified (via inclusion maps $\psi_{\alpha,\alpha\beta}^{-1}$ and $\psi_{\beta,\alpha\beta}^{-1}$) as the zero sets of s_β and s_α , respectively. Note that they are cut out transversally by the submersion axiom, though they may not intersect each other transversally (they do so at precisely those points where X is cut out transversally). Such a system of charts in the case $A = \{\alpha, \beta\}$ is illustrated in Figure 2, right.

More generally, we have charts indexed by the lattice of finite subsets of A . The compatibility axioms relating U_{IJ} , ψ_{IJ} , and s_α are all just various aspects of the charts being suitably compatible with each other. The submersion axiom is the precise property (which in practice would follow from every X_I being cut out transversally for $\#I \geq 1$) we need in order to glue together the “local virtual fundamental cycles”. The

covering axiom simply says the charts cover all of X , so we have enough information to recover its global virtual fundamental cycle.

Remark 2.1.2 The basic charts of a Kuranishi structure are indexed by the points of X . Hence to define a Kuranishi structure on a space X , one must make a choice for each $p \in X$. The charts of an implicit atlas are indexed by an abstract set A , and hence we can define implicit atlases without making any choices (see Section 2.2).

Remark 2.1.3 Most of the axioms of an implicit atlas are stated without reference to whether I is empty or nonempty. This contrasts with other approaches, where there is an axiomatic and notational distinction between the X_I , ψ_{IJ} , or U_{IJ} depending on whether I is empty or nonempty. We believe that our uniform treatment makes implicit atlases simpler notationally and conceptually, and this is a novel aspect of our approach.

Remark 2.1.4 The requirement that X_I be a manifold *whenever I is nonempty* is rather unnatural (see Remark 2.1.3) and is too strong of an assumption for two different reasons. First, we do not know how to construct implicit atlases with this property on moduli spaces of pseudo-holomorphic curves (it is a subtle question of choosing good neighborhoods to ensure transversality over all X_I if $\#I \geq 1$). Second, with this axiom we cannot form the “product implicit atlas” (Definition 6.3.1) which is crucial for understanding coherence of virtual fundamental cycles between moduli spaces when treating Floer-type homology theories. Thus in the real definition of an implicit atlas (Definition 3.1.1), there is no such requirement on X_I . Rather, we specify open subsets $X_I^{\text{reg}} \subseteq X_I$ for all $I \subseteq A$ (morally, this is the locus where X_I is cut out transversally), and we modify the transversality axioms appropriately.

Remark 2.1.5 The most natural notion of “equivalence” between two implicit atlases A and B on a space X seems to be the existence of a chain of inclusions of implicit atlases $A \subseteq C_1 \supseteq \cdots \subseteq C_n \supseteq B$. It is common to speak of “cobordisms” between Kuranishi structures or Kuranishi atlases; the analogous notion of cobordism between implicit atlases (say on the same index set A) on a space X is an implicit atlas with boundary on $X \times [0, 1]$ whose restriction to the boundary coincides with the first implicit atlas on $X \times \{0\}$ and with the second on $X \times \{1\}$.

It is an easy consequence of the VFC machinery that the virtual fundamental class is invariant under both notions of equivalence.

2.1.2 Implicit atlases on spaces with nontrivial isotropy We now give a simple construction of a convenient system of charts on any smooth orbifold. This system

of charts is a special case of (and motivates the definition of) an implicit atlas (with nontrivial covering groups, so as to apply to spaces with nontrivial isotropy).

Fix a smooth orbifold X and let $\{X_\alpha/\Gamma_\alpha = V_\alpha \subseteq X\}_{\alpha \in A}$ be an open cover, where each X_α is a smooth manifold with a smooth action by a finite group Γ_α (let us call this the “covering group”). Then for any finite subset $I = \{\alpha_1, \dots, \alpha_n\} \subseteq A$, there is an “overlap chart”

$$(2.1.1) \quad X_I := X_{\alpha_1} \times_X \cdots \times_X X_{\alpha_n},$$

$$(2.1.2) \quad \Gamma_I := \Gamma_{\alpha_1} \times \cdots \times \Gamma_{\alpha_n}$$

(where (2.1.1) is the “orbifold fiber product”; see Remark 2.1.6 below). It is easy to check that X_I is a smooth manifold with a smooth action by Γ_I and that

$$(2.1.3) \quad X_I/\Gamma_I = V_{\alpha_1} \cap \cdots \cap V_{\alpha_n} \subseteq X.$$

As an exercise, the reader may check that this system of charts described above gives an implicit atlas in the sense of Definition 3.1.1 where every $E_\alpha = 0$.

Remark 2.1.6 (Orbifold fiber product) Let $\mathcal{X}, \mathcal{Y}, \mathcal{Z}$ be orbifolds, and fix maps of orbifolds $\mathcal{X}, \mathcal{Y} \rightarrow \mathcal{Z}$. Then the *orbifold fiber product* $\mathcal{X} \times_{\mathcal{Z}} \mathcal{Y}$ is simply the categorical 2–fiber product in the weak 2–category of orbifolds. It exists whenever $\mathcal{X} \times \mathcal{Y} \rightarrow \mathcal{Z} \times \mathcal{Z}$ is transverse to the diagonal $\mathcal{Z} \rightarrow \mathcal{Z} \times \mathcal{Z}$ (in particular, it exists if at least one of the maps $\mathcal{X}, \mathcal{Y} \rightarrow \mathcal{Z}$ is *étale* as in the above example of $X_\alpha \rightarrow X$).

Thurston [78, Proof of Proposition 13.2.4] gives an explicit hands-on definition of $\mathcal{X} \times_{\mathcal{Z}} \mathcal{Y}$ in the case $\mathcal{X}, \mathcal{Y} \rightarrow \mathcal{Z}$ are both *étale* by working locally on \mathcal{Z} and then gluing. In a more modern perspective (defining orbifolds as certain stacks on the site of smooth manifolds), we may define the orbifold fiber product by the universal property

$$(2.1.4) \quad \text{Hom}(S, \mathcal{X} \times_{\mathcal{Z}} \mathcal{Y}) := \text{Hom}(S, \mathcal{X}) \times_{\text{Hom}(S, \mathcal{Z})} \text{Hom}(S, \mathcal{Y})$$

for any smooth manifold S . The right-hand side denotes the 2–fiber product in the weak 2–category of groupoids, which admits the explicit description

$$(2.1.5) \quad G_1 \times_H G_2 := \{(g_1, g_2, \theta) \mid g_1 \in G_1, g_2 \in G_2, \theta: f_1(g_1) \rightarrow f_2(g_2)\}$$

(with the obvious notion of isomorphism between triples (g_1, g_2, θ)) for groupoids G_1, G_2, H and functors $f_i: G_i \rightarrow H$.

The orbifold fiber product usually does not coincide with the fiber product of the underlying topological spaces $X \times_Y Z$, though there is at least always a map $\mathcal{X} \times_{\mathcal{Z}} \mathcal{Y} \rightarrow X \times_Z Y$. The construction (2.1.1) would not work if we used the fiber product of topological spaces.

2.2 Constructions of implicit atlases

2.2.1 Zero set of a smooth section We now give a simple example of the construction of an implicit atlas. This example is also universal in the sense that all constructions of implicit atlases in this paper are to be thought of as generalizations of this construction.

Suppose we have a smooth manifold B , a smooth vector bundle $p: E \rightarrow B$, and a smooth section $s: B \rightarrow E$ with $s^{-1}(0)$ compact. Let us now construct an implicit atlas A of dimension $d := \dim B - \dim E$ on $X := s^{-1}(0)$. We will revisit this example in Section 5.3, so the reader may also refer there for more details.

We set A to be the set of all *thickening datums*, where a thickening datum α is a triple $(V_\alpha, E_\alpha, \lambda_\alpha)$ consisting of:

- (i) An open set $V_\alpha \subseteq B$.
- (ii) A finite-dimensional vector space E_α .
- (iii) A smooth homomorphism of vector bundles $\lambda_\alpha: V_\alpha \times E_\alpha \rightarrow p^{-1}(V_\alpha)$.

Now our thickenings are

$$(2.2.1) \quad X_I := \left\{ (x, \{e_\alpha\}_{\alpha \in I}) \in \bigcap_{\alpha \in I} V_\alpha \times \bigoplus_{\alpha \in I} E_\alpha \mid s(x) + \sum_{\alpha \in I} \lambda_\alpha(x, e_\alpha) = 0 \right\}.$$

The function $s_\alpha: X_I \rightarrow E_\alpha$ is simply projection to the E_α component. The set $U_{IJ} \subseteq X_I$ is the locus where $x \in \bigcap_{\alpha \in J} V_\alpha$, and the footprint map ψ_{IJ} is simply the natural map forgetting e_α for $\alpha \in J \setminus I$. It may be a good exercise for the reader to verify the compatibility axioms in this particular case.

The transversality axioms as we have stated them in Definition 2.1.1 do not hold because X_I might not be cut out transversally for $\#I \geq 1$. The best way to fix this is to use the real definition of an implicit atlas (Definition 3.1.1) where we keep track of the locus $X_I^{\text{reg}} \subseteq X_I$ where it is cut out transversally and modify the transversality axioms appropriately (see Remark 2.1.4).

Remark 2.2.1 The reader may rightfully object that the index set A defined above is not a set but rather a groupoid (just as there is no “set of all finite sets” or “set of all compact smooth manifolds”). There are two ways of resolving this issue. The simplest solution is to add appropriate “rigidifying data” to turn A into a set (eg we could add the data of an isomorphism $E_\alpha \xrightarrow{\sim} \mathbb{R}^{\dim E_\alpha}$ to the definition of a thickening datum). Another (more cumbersome, and ultimately unnecessary) solution would be to just check that the notion of an implicit atlas and the accompanying theory of virtual fundamental cycles remains valid for A a groupoid instead of a set (see also Remark 3.1.5 for more details).

Remark 2.2.2 The index set A consists of *all* possible thickening datums. It is thus canonical in the sense that it does not depend on any auxiliary choice. For the purposes of extracting the virtual fundamental cycle from an implicit atlas, the only choice we will need to make is that of a finite subatlas; the independence of this choice is part of the VFC machinery.

Remark 2.2.3 To construct the overlap charts X_I , it is crucial that we are able to take the abstract direct sum of the obstruction spaces E_α . Hence, it is better to think of them as abstract vector spaces equipped with maps to E over V_α , rather than as trivialized subbundles of E over V_α (which also imposes the unnecessary restriction that λ_α be injective), since the latter category is not closed under direct sum.

To illustrate this distinction further, let us mention a problematic alternative version of (2.2.1):

$$(2.2.2) \quad X_I := \left\{ x \in \bigcap_{\alpha \in I} V_\alpha \mid s(x) \in \left(\bigoplus_{\alpha \in I} \lambda_\alpha \right) (E_I) \right\}.$$

This definition agrees with (2.2.1) if the map

$$(2.2.3) \quad \bigoplus_{\alpha \in I} \lambda_\alpha: \bigcap_{\alpha \in I} V_\alpha \times \bigoplus_{\alpha \in I} E_\alpha \rightarrow p^{-1} \left(\bigcap_{\alpha \in I} V_\alpha \right)$$

is injective, but in general it can be different, and indeed, if (2.2.3) fails to be injective, the definition (2.2.2) isn't particularly useful. Note that this can occur even if we add the requirement that each λ_α be injective.

The importance of using the direct sum was independently observed by McDuff and Wehrheim [63], and it is implicit in Fukaya, Oh, Ohta and Ono [31].

2.2.2 Moduli space of pseudo-holomorphic curves We now give an example of the construction of an implicit atlas on a moduli space of pseudo-holomorphic curves. This construction can be fruitfully interpreted as a generalization⁷ of the construction from Section 2.2.1.

Specifically, we construct an implicit atlas A on the moduli space of stable maps $\overline{\mathcal{M}}_{0,0}(X, B)$ (we fix a symplectic manifold X , a smooth ω -tame almost complex

⁷Moduli spaces of pseudo-holomorphic curves do not fit literally into the setting of Section 2.2.1 (generalized appropriately to Banach manifolds/bundles) because of three main issues: gluing (nodal domain curves), orbifold structure (nontrivial isotropy groups), and varying complex structures on the domain (nondifferentiability of the reparametrization action). The polyfolds project of Hofer, Wysocki and Zehnder aims to set up an infinite-dimensional Fredholm framework in which moduli spaces of pseudo-holomorphic curves may be described.

structure J , and a homology class $B \in H_2(X; \mathbb{Z})$). The reader impatient for the full details may also wish to refer to Section 9 where we give a full treatment. Here we have simplified things by assuming that $(g, n) = (0, 0)$, $\Gamma_\alpha = S_{r_\alpha}$ and $\bar{\mathcal{M}}_\alpha = \bar{\mathcal{M}}_{0,r_\alpha}$ (which is a smooth manifold!).

We define A to be the set of all *thickening datums*, where a thickening datum is a 4-tuple $(r_\alpha, D_\alpha, E_\alpha, \lambda_\alpha)$ consisting of:

- (i) An integer $r_\alpha > 2$; let $\Gamma_\alpha := S_{r_\alpha}$.
- (ii) A smooth compact codimension-two submanifold $D_\alpha \subseteq X$ with boundary.
- (iii) A finite-dimensional $\mathbb{R}[S_{r_\alpha}]$ -module E_α .
- (iv) A S_{r_α} -equivariant map

$$\lambda_\alpha: E_\alpha \rightarrow C^\infty(\bar{\mathcal{C}}_{0,r_\alpha} \times X, \Omega_{\bar{\mathcal{C}}_{0,r_\alpha}/\bar{\mathcal{M}}_{0,r_\alpha}}^{0,1} \otimes_{\mathbb{C}} TX)$$

(where $\bar{\mathcal{C}}_{0,r_\alpha} \rightarrow \bar{\mathcal{M}}_{0,r_\alpha}$ is the universal family over the Deligne–Mumford moduli space) supported away from the nodes and marked points of the fibers.

Let us remark that the analogue of the open set V_α from Section 2.2.1 is the set of maps $u: C \rightarrow X$ satisfying the conditions appearing in (i) below.

Now our thickening $\bar{\mathcal{M}}_{0,0}(X, B)_I$ (any finite $I \subseteq A$) is defined as the set whose elements consist of:

- (i) A smooth map $u: C \rightarrow X$, where C is a nodal curve of genus 0, so that for all $\alpha \in I$, we have $u \pitchfork D_\alpha$ (meaning $u^{-1}(\partial D_\alpha) = \emptyset$ and for all $p \in u^{-1}(D_\alpha)$, the differential $(du)_p: T_p C \rightarrow T_{u(p)} X / T_{u(p)} D_\alpha$ is surjective and p is not a node of C) with $\#u^{-1}(D_\alpha) = r_\alpha$ such that adding these r_α intersections as marked points makes C stable.
- (ii) Elements $e_\alpha \in E_\alpha$ for all $\alpha \in I$
- (iii) Labelings of $u^{-1}(D_\alpha)$ by $\{1, \dots, r_\alpha\}$ for all $\alpha \in I$.

These data must satisfy the equation

$$(2.2.4) \quad \bar{\partial}u + \sum_{\alpha \in I} \lambda_\alpha(e_\alpha)(\phi_\alpha, u) = 0,$$

where $\phi_\alpha: C \rightarrow \bar{\mathcal{C}}_{0,r_\alpha}$ is the unique isomorphism onto a fiber respecting the labeling of $u^{-1}(D_\alpha)$. There is an action of $\Gamma_\alpha = S_{r_\alpha}$ on $\bar{\mathcal{M}}_{0,0}(X, B)_I$ (for $\alpha \in I$) given by changing the labeling of $u^{-1}(D_\alpha)$ and by its given action on $e_\alpha \in E_\alpha$.

There are obvious projection and forgetful maps

$$s_\alpha: \overline{\mathcal{M}}_{0,0}(X, B)_I \rightarrow E_\alpha \quad \text{and} \quad \psi_{IJ}: (s_{J \setminus I} |_{X_J})^{-1}(0) \rightarrow U_{IJ},$$

respectively, where $U_{IJ} \subseteq \overline{\mathcal{M}}_{0,0}(X, B)_I$ is the locus of curves satisfying the condition in (i) above for all $\alpha \in J$. Thus we have specified the atlas data for A . The compatibility axioms are rather trivial (as in Section 2.2.1), though the homeomorphism axiom requires a bit of thought.

Now let us discuss the transversality axioms, which have much nontrivial content. To verify the covering axiom, we need to show *domain stabilization*, namely that for any J -holomorphic $u: C \rightarrow X$, there is a divisor D_α with $u \pitchfork D_\alpha$ and so that adding $u^{-1}(D_\alpha)$ to C as marked points makes C stable. Given domain stabilization, the rest of the proof of the covering axiom is rather standard (choose $(E_\alpha, \lambda_\alpha)$ big enough to cover the cokernel of the linearized operator $D\bar{\partial}$ at u). The submersion axiom asserts (in particular) that the regular locus $\overline{\mathcal{M}}_{0,0}(X, B)_I^{\text{reg}} \subseteq \overline{\mathcal{M}}_{0,0}(X, B)_I$ is a topological manifold. Proving the submersion axiom is not too difficult over the locus where the domain curve C is smooth (it follows immediately from the implicit function theorem for Banach manifolds), but to show it near a nodal domain curve amounts to proving a gluing theorem (which we do in Appendix B).

Remark 2.2.4 The thickened moduli spaces $\overline{\mathcal{M}}_{0,0}(X, B)_I$ are defined as moduli spaces of solutions to the “ I -thickened $\bar{\partial}$ -equation” (2.2.4). With this intrinsic definition, the convenient overlap properties of the charts (the compatibility axioms) follow rather trivially. The atlas also clearly does not depend on any choice of Sobolev norms or “gluing profile”.

On this point, it is useful to compare with other approaches, which often take the perspective of defining thickened moduli spaces as subsets of some particular Banach manifold of maps. In this context, achieving good overlap properties seems to be more difficult and less conceptual. Another inconvenience of this setting is the lack of differentiability of the reparametrization action (and large reparametrization groups for bubbles).

Remark 2.2.5 In our approach, the most serious analytic questions are encountered in verifying the openness and submersion axioms (that is, in proving the necessary gluing theorems). We remark that these only concern the local properties of the thickened moduli spaces, and, in particular, are separated from the other technical aspects of the construction of the implicit atlas (eg the compatibility axioms, the action of the groups Γ_I , or the action of a larger symmetry group on the entire atlas).

Remark 2.2.6 Standard gluing techniques suffice to verify the openness and submersion axioms for the implicit atlases we construct on moduli spaces of pseudo-holomorphic curves. In fact, the transition maps between gluing charts are clearly smooth when restricted to each stratum (ie for a fixed topological type of the domain), and this yields a canonical “stratified smooth structure” on each X_I^{reg} . If one wanted to obtain a smooth structure on the X_I^{reg} , one would need to construct gluing charts whose transition maps are truly smooth. This would require a choice of “gluing profile” (on which the resulting smooth structure would depend) and is slightly more delicate (see Fukaya, Oh, Ohta and Ono [30; 31] or Hofer, Wysocki and Zehnder [46]). Such methods might yield smooth implicit atlases (see Definition 3.1.3).

2.3 Construction of virtual fundamental cycles

Let us now describe concretely some simple cases of our algebraic definition of the virtual fundamental class of a space X with implicit atlas A . While the cases we treat (one chart and two charts) are admittedly rather basic, they nevertheless illustrate the main ideas necessary to deal with arbitrary implicit atlases. We will see that certain chain complexes play a key role; they will turn out to be the *virtual cochain complexes* C_{vir}^\bullet , which are the central objects of the “VFC package”.

The reader interested in the details of our treatment in full generality should refer to Section 4 (where we construct the VFC package) and Section 5 (where we define the virtual fundamental class).

For the purposes of this section, we use implicit atlases in the sense of Definition 2.1.1. We will ignore issues about orientations (as they can be dealt with rather trivially by introducing the relevant orientation sheaves/groups).

Here we work over \mathbb{Z} ; in the main body of the paper we consider any ground ring in which the orders of all relevant “covering groups” are invertible. It seems plausible that, with some more work, this could be weakened to assuming only invertibility of $\#\Gamma_x$ for all $x \in X$, where Γ_x denotes the isotropy group of $x \in X$ (ie the stabilizer of any inverse image of x under $\psi_{\emptyset, I}$ lying in X_I^{reg}).

2.3.1 What “homology group” does the virtual fundamental class live in? Since X is just a compact Hausdorff space,⁸ we must be careful about what homology theory we use to house the virtual fundamental class (Example 2.3.1 below shows that the singular homology $H_\bullet(X)$ of X is insufficient for this purpose).

⁸The existence of an implicit atlas on X does impose some additional restriction on the topology on X . For example, if X admits an implicit atlas then it is locally metrizable and hence metrizable by Smirnov’s theorem (a paracompact Hausdorff locally metrizable space is metrizable).




Figure 3: The union of $\overline{\{(x, \sin \frac{\pi}{x}) : 0 < x < 1\}}$ and a path from $(0, 0)$ to $(1, 0)$, known as the Warsaw circle $W \subseteq \mathbb{R}^2$.

The dual of Čech cohomology $\check{H}^\bullet(X)^\vee := \text{Hom}(\check{H}^\bullet(X), \mathbb{Z})$ is a good candidate (and it is what we choose to use, though see Remark 5.0.2 for further discussion). For example, a map of spaces $f: X \rightarrow Y$ induces a pushforward $f_*: \check{H}^\bullet(X)^\vee \rightarrow \check{H}^\bullet(Y)^\vee$ (defined as the dual of pullback $f^*: \check{H}^\bullet(Y) \rightarrow \check{H}^\bullet(X)$). Moreover, for a finite CW-complex Z , we have $\check{H}^\bullet(Z) = H^\bullet(Z)$, so $\check{H}^\bullet(Z)^\vee = H_\bullet(Z)/\text{tors}$. It follows that (for many purposes) a virtual fundamental class $[X]_A^{\text{vir}} \in \check{H}^d(X)^\vee$ can be used in the same way one would use the fundamental class $[X] \in H_d(X)$ if X were a closed manifold of dimension d .

Example 2.3.1 (Insufficiency of singular homology) Consider the Warsaw circle $W \subseteq \mathbb{R}^2$ as illustrated in Figure 3; note that singular $H_1(W)$ vanishes. Now $\mathbb{R}^2 \setminus W$ has two connected components; let $s: \mathbb{R}^2 \rightarrow \mathbb{R}$ be positive on one component and negative on the other; this gives the data of an implicit atlas on $W = s^{-1}(0)$. Using any reasonable definition, we certainly want $[W]^{\text{vir}} \in \check{H}^1(W)^\vee \cong \mathbb{Z}$ to be a generator, however this is clearly not in the image of singular homology under the natural map $H_\bullet(W) \rightarrow H^\bullet(W)^\vee \rightarrow \check{H}^\bullet(W)^\vee$. Alternatively, the pushforward of $[W]^{\text{vir}}$ to a small annular neighborhood $A \subseteq \mathbb{R}^2$ of W should be a generator of $H_1(A) \cong \mathbb{Z}$ (as one can see by perturbing s).

2.3.2 Virtual fundamental class from a single chart We have a space X , and the implicit atlas $A = \{\alpha\}$ consists of the following data. We have a topological manifold X_α (not necessarily compact), a vector space E_α , and a continuous function $s_\alpha: X_\alpha \rightarrow E_\alpha$. We also have an identification $X = s_\alpha^{-1}(0)$ (see Figure 2, left).

We define the virtual fundamental class via the diagram

$$(2.3.1) \quad \check{H}^d(X) = H_{\dim E_\alpha}(X_\alpha, X_\alpha \setminus X) \xrightarrow{(s_\alpha)_*} H_{\dim E_\alpha}(E_\alpha, E_\alpha \setminus 0) \xrightarrow{[E_\alpha] \mapsto 1} \mathbb{Z},$$

which we explain below (recall that $\dim X_\alpha = d + \dim E_\alpha$). Poincaré duality gives a canonical⁹ identification $\check{H}^\bullet(X) = H_{\dim X_\alpha - \bullet}(X_\alpha, X_\alpha \setminus X)$, which gives the first equality in (2.3.1) (this is a rather strong version of Poincaré duality, valid for any compact subset X of a manifold X_α ; we will say more about how to prove it in Section 2.3.5). We think of the composite of the maps in (2.3.1) as an element $[X]_A^{\text{vir}} \in \check{H}^d(X)^\vee$, which we call the virtual fundamental class.

In this setting, one can also define the virtual fundamental class using perturbation (under the additional assumption that X_α carries a smooth structure). Specifically, one can perturb s_α to \tilde{s}_α so that it is a submersion over $0 \in E_\alpha$; we consider $\tilde{s}_\alpha^{-1}(0)$ as a “perturbed moduli space” near X . Using the continuity property of Čech cohomology, one can then make sense of $\lim_{\tilde{s}_\alpha \rightarrow s_\alpha} [\tilde{s}_\alpha^{-1}(0)]$ as an element of $\check{H}^d(X)^\vee$.

The algebraic approach is more general (it does not require a smooth structure on X_α), and it is easy to see that it gives the same answer as the perturbation approach when X_α has a smooth structure.

Example 2.3.2 Let $X = \{0\}$, $X_\alpha = E_\alpha = \mathbb{R}$, and $s_\alpha(x) = x^n$ ($n \geq 1$). The reader may easily check that with our definition, the virtual fundamental class is 1 if n is odd and 0 if n is even (as is consistent with the perturbation picture).

Let $X = \{0\}$, $X_\alpha = E_\alpha = \mathbb{C}$ (considered as an \mathbb{R} -vector space), and $s_\alpha(z) = z^n$ ($n \geq 1$). The reader may check that in this case, the virtual fundamental class is n .

2.3.3 Virtual fundamental class from a single chart (with covering group) We now describe how the construction from Section 2.3.2 must be modified in the presence of nontrivial covering groups (as in Section 2.1.2). We have not yet introduced implicit atlases with nontrivial covering groups, so we will simply say explicitly what this means in the present situation of a single chart (the reader may also refer to Definition 3.1.1).

We have a space X , and the implicit atlas $A = \{\alpha\}$ consists of the following data. We have a topological manifold X_α (not necessarily compact), a vector space E_α , and a continuous function $s_\alpha: X_\alpha \rightarrow E_\alpha$. We have a finite group Γ_α acting on X_α and (linearly) on E_α so that s_α is Γ_α -equivariant. Finally, we have an identification $X = s_\alpha^{-1}(0)/\Gamma_\alpha$. We must now work over the coefficient ring $R = \mathbb{Z}[1/\#\Gamma_\alpha]$.

To define the virtual fundamental class we consider

$$(2.3.2) \quad \check{H}^d(s_\alpha^{-1}(0); R)^{\Gamma_\alpha} = H_{\dim E_\alpha}(X_\alpha, X_\alpha \setminus s_\alpha^{-1}(0); R)^{\Gamma_\alpha} \xrightarrow{(s_\alpha)_*} H_{\dim E_\alpha}(E_\alpha, E_\alpha \setminus 0; R)^{\Gamma_\alpha} \xrightarrow{[E_\alpha] \mapsto 1} R.$$

⁹For the moment, we ignore the necessary twist by the orientation sheaf.

Now it is a general property of Čech cohomology that $\check{H}^\bullet(Y; R)^\Gamma = \check{H}^\bullet(Y/\Gamma; R)$ for a compact Hausdorff space Y acted on by a finite group Γ as long as R is a module over $\mathbb{Z}[1/\#\Gamma]$. We precompose (2.3.2) with the canonical isomorphism $\check{H}^d(X; R) \rightarrow \check{H}^d(s_\alpha^{-1}(0); R)^{\Gamma_\alpha}$ given as $1/\#\Gamma_\alpha$ times the pullback. This gives an element $[X]_A^{\text{vir}} \in \check{H}^d(X; R)^\vee$, which we call the virtual fundamental class.

2.3.4 Virtual fundamental class from two charts Generalizing the approach in Section 2.3.2 to multiple charts leads immediately to the heart of the problem of defining virtual fundamental cycles. Namely, we must figure out how to glue together the information contained in each local chart to define the global virtual fundamental class. We explain our solution to this problem in the simple case of $A = \{\alpha, \beta\}$, which nevertheless illustrates most of the ideas necessary to deal with the general case (which in addition just requires good organization of the combinatorics). Since the problem we are facing is precisely to glue together local homological information into global homological information, it should not be surprising that sheaf-theoretic tools and homological algebra are useful.

Our space is X , and the implicit atlas $A = \{\alpha, \beta\}$ amounts to the following data:

$$(2.3.3) \quad s_\alpha: X_\alpha \rightarrow E_\alpha, \quad s_\alpha^{-1}(0) = U_\alpha \subseteq X \text{ (open subset),}$$

$$(2.3.4) \quad s_\beta: X_\beta \rightarrow E_\beta, \quad s_\beta^{-1}(0) = U_\beta \subseteq X \text{ (open subset),}$$

$$(2.3.5) \quad s_\alpha \oplus s_\beta: X_{\alpha\beta} \rightarrow E_\alpha \oplus E_\beta, \quad (s_\alpha \oplus s_\beta)^{-1}(0) = U_{\alpha\beta} = U_\alpha \cap U_\beta \subseteq X,$$

which fit together as outlined in Section 2.1.1 (and as illustrated in Figure 2, right).

We would like to generalize the approach in Section 2.3.2, specifically Equation (2.3.1). For this, we need a replacement for the group $H_\bullet(X_\alpha, X_\alpha \setminus X)$. To construct such a replacement, we would like to “glue together” $C_\bullet(X_\alpha, X_\alpha \setminus U_\alpha)$ and $C_\bullet(X_\beta, X_\beta \setminus U_\beta)$ along $C_\bullet(X_{\alpha\beta}, X_{\alpha\beta} \setminus U_{\alpha\beta})$ (as remarked earlier, it is easier to glue together these *complexes* rather than glue together the corresponding *spaces*). The resulting complex should calculate the Čech cohomology of X (by some version of Poincaré duality) and also have a natural map to $C_\bullet(E_\alpha \oplus E_\beta, E_\alpha \oplus E_\beta \setminus 0)$. If we can construct a complex with these two properties, then we can define the virtual fundamental class just as in (2.3.1). This complex we construct will be called the *virtual cochain complex*.

Remark 2.3.3 The complex $C_{\dim X_\alpha - \bullet}(X_\alpha, X_\alpha \setminus U_\alpha)$ calculates $\check{H}_c^\bullet(U_\alpha)$, and the map $s_*: C_\bullet(X_\alpha, X_\alpha \setminus U_\alpha) \rightarrow C_\bullet(E_\alpha, E_\alpha \setminus 0)$ should be thought of as the chain-level “local virtual fundamental cycle” $[X]_A^{\text{vir}} \in \check{H}_c^{\dim X_\alpha - \dim E_\alpha}(U_\alpha)^\vee$, which we would like to glue together into the global virtual fundamental cycle.

Remark 2.3.4 Technically speaking, it is very important to have a uniform functorial definition of the virtual cochain complexes (one which does not require making any extra choices).

As a first try towards gluing the desired complexes together, let us consider using the mapping cone of the chain map

$$(2.3.6) \quad C_{\dim X_{\alpha\beta}-\bullet}(X_{\alpha\beta}, X_{\alpha\beta} \setminus U_{\alpha\beta}) \xrightarrow{\cap[s_{\beta}^{-1}(0)] \oplus \cap[s_{\alpha}^{-1}(0)]} C_{\dim X_{\alpha}-\bullet}(X_{\alpha}, X_{\alpha} \setminus U_{\alpha}) \oplus C_{\dim X_{\beta}-\bullet}(X_{\beta}, X_{\beta} \setminus U_{\beta})$$

where the maps are intersections of chains with the (transversely cut out!) submanifolds $s_{\beta}^{-1}(0)$ and $s_{\alpha}^{-1}(0)$ of $X_{\alpha\beta}$. There is the question, though, of how these maps are to be defined on the chain level. There are various direct ways to define these maps,¹⁰ however (at least when attempted by the author) the multitude of “choices” one has to make invariably leads to a big mess.

As a second try, let us try to glue together the complexes

$$(2.3.7) \quad C_{\dim X_{\alpha\beta}-\bullet}(X_{\alpha} \times E_{\beta}, [X_{\alpha} \setminus U_{\alpha}] \times E_{\beta}),$$

$$(2.3.8) \quad C_{\dim X_{\alpha\beta}-\bullet}(E_{\alpha} \times X_{\beta}, E_{\alpha} \times [X_{\beta} \setminus U_{\beta}]),$$

$$(2.3.9) \quad C_{\dim X_{\alpha\beta}-\bullet}(X_{\alpha\beta}, X_{\alpha\beta} \setminus U_{\alpha\beta}).$$

This should let us let us avoid the cap-product maps (since there is no “dimension shift” between these complexes). Each of these complexes has a canonical map to $C_{\dim X_{\alpha\beta}-\bullet}(E_{\alpha} \oplus E_{\beta}, E_{\alpha} \oplus E_{\beta} \setminus 0)$, but it is not clear how to glue them together in a manner compatible with this map. In particular, there is no reason for there to exist commuting diagrams:

$$(2.3.10) \quad \begin{array}{ccc} U_{\alpha,\alpha\beta} \times E_{\beta} & \longrightarrow & X_{\alpha\beta} \\ & \searrow s_{\alpha} \oplus \text{id} & \swarrow s_{\alpha\beta} \\ & E_{\alpha} \oplus E_{\beta} & \end{array} \quad \begin{array}{ccc} E_{\alpha} \times U_{\beta,\alpha\beta} & \longrightarrow & X_{\alpha\beta} \\ & \searrow \text{id} \oplus s_{\beta} & \swarrow s_{\alpha\beta} \\ & E_{\alpha} \oplus E_{\beta} & \end{array}$$

Moreover, it is technically very inconvenient (functoriality of the construction is a mess) to have a complex which depends on a choice of maps (2.3.10) (even if this turns out to be a contractible choice, and thus morally irrelevant).

¹⁰One could use very fine chains, do a (chain-level) cap product with a choice of cochain-level Poincaré dual of the relevant submanifold and then project “orthogonally” onto the submanifold. Alternatively, one could just use “generic” chains (or perturb the chains) so they are transverse to the submanifold, and then triangulate the intersection in a suitable way.

As a third try (which ends up working nicely), let us consider the following “deformation to the normal cone”:

$$(2.3.11) \quad Y_{\alpha\beta} := \{(e_\alpha, e_\beta, t, x) \in E_\alpha \times E_\beta \times [0, 1] \times X_{\alpha\beta} \mid s_\alpha(x) = t \cdot e_\alpha, s_\beta(x) = (1 - t) \cdot e_\beta\}.$$

We think of $Y_{\alpha\beta}$ as a family of spaces parametrized by $[0, 1]$ (via the projection $Y_{\alpha\beta} \rightarrow [0, 1]$). Observe that if $t \in (0, 1)$, then e_α, e_β are determined uniquely by x . Therefore, over the open interval $(0, 1)$, we see that $Y_{\alpha\beta}$ is a trivial product space $X_{\alpha\beta} \times (0, 1)$. Over the point $t = 0$, though, the fiber is $E_\alpha \times U_{\beta,\alpha\beta}$, which is the “normal cone” of the submanifold $U_{\beta,\alpha\beta} \subseteq X_{\alpha\beta}$ cut out (transversally!) by the equation $s_\alpha = 0$. Similarly, over the point $t = 1$, the fiber is $U_{\alpha,\alpha\beta} \times E_\beta$. Also observe that $Y_{\alpha\beta}$ is a manifold by the submersion axiom.

Now we consider the mapping cone of

$$(2.3.12) \quad \left[\begin{array}{c} C_{\dim X_{\alpha\beta}-\bullet}(U_{\alpha,\alpha\beta} \times E_\beta, U_{\alpha,\alpha\beta} \times E_\beta \setminus U_\alpha \times 0) \\ \oplus \\ C_{\dim X_{\alpha\beta}-\bullet}(E_\alpha \times U_{\beta,\alpha\beta}, E_\alpha \times U_{\beta,\alpha\beta} \setminus 0 \times U_\beta) \end{array} \right] \longrightarrow \left[\begin{array}{c} C_{\dim X_{\alpha\beta}-\bullet}(X_\alpha \times E_\beta, X_\alpha \times E_\beta \setminus U_\alpha \times 0) \\ \oplus \\ C_{\dim X_{\alpha\beta}-\bullet}(Y_{\alpha\beta}, Y_{\alpha\beta} \setminus 0 \times [0, 1] \times U_{\alpha\beta}) \\ \oplus \\ C_{\dim X_{\alpha\beta}-\bullet}(E_\alpha \times X_\beta, E_\alpha \times X_\beta \setminus 0 \times U_\beta) \end{array} \right].$$

The maps are simply pushforwards along the maps of spaces (with appropriate signs)

$$(2.3.13) \quad U_{\alpha,\alpha\beta} \times E_\beta \hookrightarrow X_\alpha \times E_\beta,$$

$$(2.3.14) \quad U_{\alpha,\alpha\beta} \times E_\beta \hookrightarrow Y_{\alpha\beta} \quad (\text{isomorphism onto } t = 1 \text{ fiber}),$$

$$(2.3.15) \quad E_\alpha \times U_{\beta,\alpha\beta} \hookrightarrow Y_{\alpha\beta} \quad (\text{isomorphism onto } t = 0 \text{ fiber}),$$

$$(2.3.16) \quad E_\alpha \times U_{\beta,\alpha\beta} \hookrightarrow E_\alpha \times X_\beta.$$

There is also an evident map from (the mapping cone of) (2.3.12) to

$$C_\bullet(E_\alpha \oplus E_\beta, E_\alpha \oplus E_\beta \setminus 0)$$

(namely pushforward on the three factors on the right-hand side and zero on the left-hand side).

Now to complete the definition of the virtual fundamental cycle of X using the mapping cone of (2.3.12), we need an argument to show that its homology is canonically isomorphic to the Čech cohomology of X (a sort of Poincaré duality isomorphism). We discuss this next in Section 2.3.5.

2.3.5 Homotopy \mathcal{K} -sheaves in the theory of virtual fundamental cycles As we have already mentioned, the central objects we use to understand virtual fundamental cycles are the *virtual cochain complexes* $C_{\text{vir}}^\bullet(X; A)$ (for example, the mapping cone of (2.3.12) plays the role of the virtual cochain complex in Section 2.3.4). A crucial ingredient in this approach is an isomorphism

$$(2.3.17) \quad H_{\text{vir}}^\bullet(X; A) = \check{H}^\bullet(X).$$

In Section 2.3.2, where we used $C_{\text{vir}}^\bullet(X; A) := C_{\dim X_\alpha - \bullet}(X_\alpha, X_\alpha \setminus X)$, this isomorphism was simply (a strong form of) Poincaré duality.

Let us now discuss our general approach to the isomorphism (2.3.17), which we think of as a generalized form of Poincaré duality. An efficient approach (in fact, the only approach known to the author) to constructing this isomorphism is through the language of homotopy \mathcal{K} -sheaves, and so this is the way we present it. We develop the necessary sheaf-theoretic foundations in Appendix A, so the reader may also wish to refer to that section for more details.

As an introduction to the language of sheaves and homotopy sheaves, let us first use it to give a proof of ordinary Poincaré duality (in fact, a strong version for arbitrary compact subsets of a manifold).¹¹ The following argument appears in full detail in Lemma A.6.4.

Fix a closed manifold M of dimension n . For any compact subset $K \subseteq M$, let $\mathcal{F}^\bullet(K)$ denote the complex $C_{n-\bullet}(M, M \setminus K)$. This object \mathcal{F}^\bullet is a \mathcal{K} -presheaf¹² of complexes (or *complex of \mathcal{K} -presheaves*) on M , which just means that we have natural maps $\mathcal{F}^\bullet(K) \rightarrow \mathcal{F}^\bullet(K')$ for $K \supseteq K'$, which are suitably compatible with each other. Now \mathcal{F}^\bullet satisfies the following key properties:

- (i) **\mathcal{F}^\bullet is a homotopy \mathcal{K} -sheaf** The total complex of the double complex

$$(2.3.18) \quad \mathcal{F}^\bullet(K_1 \cup K_2) \rightarrow \mathcal{F}^\bullet(K_1) \oplus \mathcal{F}^\bullet(K_2) \rightarrow \mathcal{F}^\bullet(K_1 \cap K_2)$$

is acyclic. This is essentially a restatement of the Mayer-Vietoris exact sequence.

- (ii) **\mathcal{F}^\bullet is pure and $H^0 \mathcal{F}^\bullet = \mathfrak{o}_M$** The homology of $\mathcal{F}^\bullet(\{p\})$ (ie $H_{n-\bullet}(M, M \setminus p)$) is concentrated in degree zero, where it can be canonically identified with the fiber of \mathfrak{o}_M (the orientation sheaf of M) at $p \in M$.

¹¹An easier proof (using the fact that Čech cohomology satisfies the “continuity axiom”) is available if M has a smooth structure (along the lines of [68, pages 887–888, Lemmas 3.1 and 3.3]). This approach does not seem to apply to the more general setting we need to treat here.

¹²The prefix “ \mathcal{K} -” indicates sections are given over compact sets instead of open sets. For technical reasons, it is more convenient to work with \mathcal{K} -presheaves rather than presheaves, though at the conceptual level, the reader may safely ignore the difference.

(The precise notions of a homotopy \mathcal{K} -sheaf and of purity are given in Definitions A.2.5 and A.5.1; in the above statements they have been simplified for the sake of exposition.) A formal consequence (Proposition A.5.4) of the fact that \mathcal{F}^\bullet is a pure homotopy \mathcal{K} -sheaf with $H^0\mathcal{F}^\bullet = \mathfrak{o}_M$ is that there is a canonical isomorphism

$$(2.3.19) \quad H_{n-i}(M, M \setminus K) = H^i\mathcal{F}^\bullet(K) = \check{H}^i(K; \mathfrak{o}_M).$$

Specializing to $K = M$, we have derived the Poincaré duality isomorphism $H_{n-i}(M) = \check{H}^i(M; \mathfrak{o}_M)$.

This argument generalizes as follows to prove the isomorphism (2.3.17). For the sake of concreteness, let us take $C_{\text{vir}}^\bullet(X; A)$ to be the mapping cone of (2.3.12), though the general case is not much different. First of all, we observe that there is a natural complex of \mathcal{K} -presheaves \mathcal{F}^\bullet on X whose complex of global sections is $C_{\text{vir}}^\bullet(X; A)$. Namely, to get $\mathcal{F}^\bullet(K)$ we simply replace every occurrence of U_α , U_β , or $U_{\alpha\beta}$ in (2.3.12) by its intersection with K . Now \mathcal{F}^\bullet is a homotopy \mathcal{K} -sheaf (extensions of homotopy \mathcal{K} -sheaves are homotopy \mathcal{K} -sheaves by Lemma A.2.11, and each of the individual complexes appearing in (2.3.12) gives a homotopy \mathcal{K} -sheaf by Mayer-Vietoris). To prove that \mathcal{F}^\bullet is pure and to identify its H^0 , we can calculate $H^i\mathcal{F}^\bullet(\{p\})$ using a spectral sequence which degenerates at the E_2 term (this is the argument in Lemma A.8.2). It is this second step where it matters that we glued the complexes together “correctly”. Since \mathcal{F}^\bullet is a pure homotopy \mathcal{K} -sheaf and we have identified its H^0 , this gives the desired isomorphism.

Let us also mention that the fact that the virtual cochain complex $C_{\text{vir}}^\bullet(X; A)$ is naturally the complex of global sections of a homotopy \mathcal{K} -sheaf on X plays a key role in proving other crucial properties in addition to (2.3.17).

2.4 Floer-type homology theories

In this section, we introduce the basic ideas needed to apply our methods to construct Hamiltonian Floer homology in the setting of non-degenerate periodic orbits and non-transverse moduli spaces of Floer trajectories. A toy example of the same flavor (which we mention only for the sake of analogy) is the problem of defining Morse homology from a Morse function with gradient-like vector field which is not necessarily Morse-Smale.

The methods developed thus far (the VFC package and the framework for constructing implicit atlases) are robust in that they apply to moduli spaces of Floer trajectories without much modification. The main task is to add a layer of (rather intricate) combinatorics and algebra to properly organize together the information they yield. Essentially what we must do is execute the key diagram (1.3.1)–(1.3.2) on the chain level.

We approach the problem in two logically separate steps. In Section 2.4.1, we describe the implicit atlases we put on the moduli spaces of Floer trajectories. In Section 2.4.2, we describe how to use the VFC package to define Floer-type homology groups from an appropriate abstract collection of “flow spaces” equipped with implicit atlases.

2.4.1 The system of implicit atlases We describe the “compatible system of implicit atlases” we put on the moduli spaces of Floer trajectories relevant for defining Hamiltonian Floer homology. For the sake of exposition, we will imagine we are in an artificially simplified setting (the reader may refer to Section 10 for the full details).

We assume there are just three periodic orbits $p < q < r$ (ordered by action). Hence there are just three moduli spaces we have to deal with: $\overline{\mathcal{M}}(p, q)$, $\overline{\mathcal{M}}(q, r)$, and $\overline{\mathcal{M}}(p, r)$ (which are all compact). There is a single concatenation map

$$(2.4.1) \quad \overline{\mathcal{M}}(p, q) \times \overline{\mathcal{M}}(q, r) \rightarrow \overline{\mathcal{M}}(p, r)$$

and it is a homeomorphism onto its image. We describe the construction of implicit atlases ((i)–(iv) below) which are enough to provide a robust notion of “coherent system of virtual fundamental cycles” with which we can define Floer homology.

The moduli spaces $\overline{\mathcal{M}}(p, q)$ and $\overline{\mathcal{M}}(q, r)$ do not contain any “broken trajectories”. There is thus a straightforward generalization of the construction in Section 2.2.2 by which we may define:

- (i) An implicit atlas $A(p, q)$ on $\overline{\mathcal{M}}(p, q)$.
- (ii) An implicit atlas $A(q, r)$ on $\overline{\mathcal{M}}(q, r)$.

Now since $\overline{\mathcal{M}}(p, r)$ contains the codimension-one boundary stratum $\overline{\mathcal{M}}(p, q) \times \overline{\mathcal{M}}(q, r)$, we cannot expect to equip it with an implicit atlas in the sense of Definitions 2.1.1 or 3.1.1. Rather, we would like to equip it with an *implicit atlas with boundary*. Such an atlas on a space X is given by the same data as an implicit atlas, except that in addition we specify closed subsets $\partial X_I \subseteq X_I$ for all $I \subseteq A$ which are compatible with the ψ_{IJ} , and we modify the transversality axioms to assert that X_I^{reg} is a manifold with boundary $\partial X_I \cap X_I^{\text{reg}}$. In particular, there should already be a natural choice of $\partial X \subseteq X$, which in the present case is simply the definition $\partial \overline{\mathcal{M}}(p, r) := \overline{\mathcal{M}}(p, q) \times \overline{\mathcal{M}}(q, r)$. The notion of an implicit atlas with boundary is formulated so that the natural generalization of the construction in Section 2.2.2 defines:

- (iii) An implicit atlas with boundary $A(p, r)$ on $\overline{\mathcal{M}}(p, r)$.

For this atlas, the closed subsets $\partial X_I \subseteq X_I$ are simply the loci where the thickened trajectory has a break at q .

Of course, to have any reasonable notion of “coherent virtual fundamental cycles” for the spaces $\overline{\mathcal{M}}(\cdot, \cdot)$, we need to “relate” the atlases (i)–(iii). It turns out that there is a natural way to do this “over $\overline{\mathcal{M}}(p, q) \times \overline{\mathcal{M}}(q, r)$ ” which we now describe.

First, let us observe that (i) and (ii) naturally give rise to an implicit atlas on $\overline{\mathcal{M}}(p, q) \times \overline{\mathcal{M}}(q, r)$. This is a special case of a general observation: implicit atlases A on X and A' on X' induce an implicit atlas $A \sqcup A'$ (disjoint union of index sets) on $X \times X'$, simply by defining $(X \times X')_{I \sqcup I'} := X_I \times X'_{I'}$. Hence there is a “product implicit atlas” $A(p, q) \sqcup A(q, r)$ on $\overline{\mathcal{M}}(p, q) \times \overline{\mathcal{M}}(q, r)$.

Second, let us observe that (iii) naturally gives rise to an implicit atlas on $\partial\overline{\mathcal{M}}(p, r) = \overline{\mathcal{M}}(p, q) \times \overline{\mathcal{M}}(q, r)$. This is also a special case of a general observation: an implicit atlas with boundary on X induces an implicit atlas (with the same index set) on ∂X , simply by defining $(\partial X)_I := \partial X_I$. Hence there is a “restriction to the boundary” implicit atlas $A(p, r)$ on $\overline{\mathcal{M}}(p, q) \times \overline{\mathcal{M}}(q, r)$.

Now, a good notion of “compatibility” between two implicit atlases A and B on a space X is the existence of an implicit atlas on X with index set $A \sqcup B$, whose subatlases A and B coincide with the given atlases. More importantly, this seems to be the notion of compatibility which arises most naturally in practice (in particular, there is usually a canonical choice for the atlas $A \sqcup B$ which does not depend on any “extra choices”). Hence the final implicit atlas we need is:

- (iv) An implicit atlas $A(p, q) \sqcup A(q, r) \sqcup A(p, r)$ on $\overline{\mathcal{M}}(p, q) \times \overline{\mathcal{M}}(q, r)$ whose subatlas $A(p, q) \sqcup A(q, r)$ coincides with the “product implicit atlas” above and whose subatlas $A(p, r)$ coincides with the “restriction to the boundary” above.

This atlas (iv) is constructed essentially using the same ideas from Section 2.2.2, however a few remarks are in order about its definition.

The thickened moduli spaces for the atlas (iv) are defined in the usual way, as moduli spaces of broken Floer trajectories $u: C \rightarrow M \times S^1$ with some intersection conditions with divisors D_α , satisfying a “thickened $\bar{\partial}$ -equation”. Let us discuss these conditions more precisely. The broken Floer trajectories in question are trajectories from p to r broken at q , or, equivalently, a pair of trajectories, one from p to q and one from q to r . Let us denote the entire broken trajectory as $u: C_{p,r} \rightarrow M \times S^1$, and let us use $C_{p,q}, C_{q,r} \subseteq C_{p,r}$ to denote the closed subcurves representing the portions of the trajectory from p to q and from q to r , respectively (so $C_{p,r} = (C_{p,q} \sqcup C_{q,r}) / \sim$, where \sim identifies a point on $C_{p,q}$ with a point on $C_{q,r}$). Then the important points are:

- (i) For thickening datums $\alpha \in A(p, q)$ we require that $(u|_{C_{p,q}}) \pitchfork D_\alpha$, and we label the intersections with $\{1, \dots, r_\alpha\}$, inducing a unique map $\phi_\alpha: C_{p,q} \rightarrow \overline{\mathbb{C}}_{0,2+r_\alpha}$.

- (ii) For thickening datums $\alpha \in A(q, r)$ we require that $(u|_{C_{q,r}}) \pitchfork D_\alpha$, and we label the intersections with $\{1, \dots, r_\alpha\}$, inducing a unique map $\phi_\alpha: C_{q,r} \rightarrow \bar{\mathcal{C}}_{0,2+r_\alpha}$.
- (iii) For thickening datums $\alpha \in A(p, r)$ we require that $(u|_{C_{p,r}}) \pitchfork D_\alpha$, and we label the intersections with $\{1, \dots, r_\alpha\}$, inducing a unique map $\phi_\alpha: C_{p,r} \rightarrow \bar{\mathcal{C}}_{0,2+r_\alpha}$.
- (iv) The thickened $\bar{\partial}$ -equation we impose is still written in the form (2.2.4), though the term $\lambda_\alpha(e_\alpha)(\phi_\alpha, u)$ is defined to be zero outside the domain of ϕ_α (namely $C_{p,q}$, $C_{q,r}$, or $C_{p,r}$, depending on whether α comes from $A(p, q)$, $A(q, r)$ or $A(p, r)$).

A good exercise in understanding this definition is to check that the subatlases $A(p, q) \sqcup A(q, r)$ and $A(p, r)$ are the “product implicit atlas” and “restriction to the boundary”, respectively (this is just a matter of matching up definitions).

Remark 2.4.1 In theory, there is no way to go from a pair of implicit atlases A and B on a space X to an implicit atlas $A \sqcup B$ on X whose subatlases A and B coincide with the given atlases (and there is no uniqueness for atlases $A \sqcup B$ with this property). However, in practice there is often a natural choice of such an atlas, which moreover is essentially the only reasonable choice. The atlas (iv) is a good example of this.

The main ideas necessary to define the compatible system of implicit atlases in the general case are all present above; the only real difference is that there are more moduli spaces and more atlases to keep track of.

2.4.2 Applying the VFC package We now explain how to use the VFC package to define homology groups from a system of moduli spaces (equipped with implicit atlases) such as appear in a Morse-type setup. The extra complicating factor in this construction (compared with constructing virtual fundamental classes) is that we must make “coherent” choices for each moduli space (ie choices with certain compatibility properties with respect to the maps between moduli spaces). Such choices give rise to “counts” for the 0-dimensional moduli spaces and thus to a differential. Thus we must show that such choices always exist and that the resulting homology groups are independent of this choice.

In our presentation here, we make a number of simplifying assumptions (which we try to point out when relevant) for the sake of exposition. The full details appear in Section 7.

Let us begin by fixing some notation. We fix a finite set \mathcal{P} (the “set of generators”) equipped with a partial order $<$ and a grading $\text{gr}: \mathcal{P} \rightarrow \mathbb{Z}$. For all pairs $p < q$ in \mathcal{P} ,

we fix a compact Hausdorff space $\mathcal{X}(p, q)$ (the “space of broken trajectories from p to q ”). These spaces $\mathcal{X}(p, q)$ are also equipped with “concatenation maps”

$$(2.4.2) \quad \mathcal{X}(p, q) \times \mathcal{X}(q, r) \rightarrow \mathcal{X}(p, r)$$

which satisfy some natural properties, in particular associativity.

We call such a pair $(\mathcal{P}, \mathcal{X})$ (when defined precisely) a *flow category*. This terminology is due to Cohen, Jones and Segal [14], who used it to mean something closer to what we prefer to call a *Morse–Smale flow category*, namely a flow category in which every $\mathcal{X}(p, q)$ is a manifold with corners of dimension $\text{gr}(q) - \text{gr}(p) - 1$ in a manner compatible with the concatenation maps (2.4.2). The basic example of a flow category is the flow category of a Morse function: \mathcal{P} is the set of critical points, and $\mathcal{X}(p, q)$ is the space of (broken) Morse trajectories from p to q ; if the Morse function is Morse–Smale, then this flow category is Morse–Smale.

Given a Morse–Smale flow category $(\mathcal{P}, \mathcal{X})$, one can construct a map $d: \mathbb{Z}[\mathcal{P}] \rightarrow \mathbb{Z}[\mathcal{P}]$ by counting those \mathcal{X} of dimension 0, and one can prove that $d^2 = 0$ by considering those \mathcal{X} of dimension 1. Our goal is to generalize this construction to flow categories equipped with implicit atlases (meaning the spaces $\mathcal{X}(p, q)$ carry suitably compatible implicit atlases with boundary, with dimension $\text{gr}(q) - \text{gr}(p) - 1$). Specifically, let us assume that we have fixed implicit atlases $\mathcal{A}(p, q)$ on $\mathcal{X}(p, q)$ and that

$$\mathcal{A}(p, q) \sqcup \mathcal{A}(q, r) = \mathcal{A}(p, r)$$

for all $p < q < r \in \mathcal{P}$.

Remark 2.4.2 There are many different choices for what one could mean by a “compatible system of implicit atlases” $\mathcal{A}(p, q)$ on the spaces $\mathcal{X}(p, q)$. For the sake of exposition, we have chosen here the structure for which it is *easiest to apply the VFC package*. In the actual construction in Section 7, we use the structure which is the *easiest to construct in practice* (via the natural generalization of Section 2.4.1). As a result, various complexes that here *coincide* are only *canonically quasi-isomorphic* in Section 7. Basically, this means that in Section 7, we will have to take lots of homotopy colimits to make certain maps well-defined on the chain level. A systematic use of ∞ -categories of complexes would tame the resulting explosion of notation, at the cost of relying on that more abstract language/machinery.

We now review what the VFC package gives to us for a flow category $(\mathcal{P}, \mathcal{X})$ equipped with a compatible system of implicit atlases. For a space X equipped with an implicit atlas with boundary, the VFC package provides *virtual cochain complexes* $C_{\text{vir}}^\bullet(X \text{ rel } \partial)$

and $C_{\text{vir}}^\bullet(X)$ (defined using the ideas from Section 2.3). Moreover, giving ∂X the “restriction to the boundary” implicit atlas, there is a natural map

$$(2.4.3) \quad C_{\text{vir}}^{\bullet-1}(\partial X) \rightarrow C_{\text{vir}}^\bullet(X \text{ rel } \partial)$$

whose mapping cone is $C_{\text{vir}}^\bullet(X)$ (by definition). There are canonical isomorphisms

$$(2.4.4) \quad H_{\text{vir}}^\bullet(X) \simeq \check{H}^\bullet(X; \mathfrak{o}_X),$$

$$(2.4.5) \quad H_{\text{vir}}^\bullet(X \text{ rel } \partial) \simeq \check{H}^\bullet(X; \mathfrak{o}_{X \text{ rel } \partial}),$$

where \mathfrak{o}_X is the “virtual orientation sheaf” of X . There is also a map

$$(2.4.6) \quad C_{\text{vir}}^\bullet(X \text{ rel } \partial) \xrightarrow{s^*} C_{\dim E_A - \bullet}(E_A, E_A \setminus 0)$$

(which, when combined with (2.4.5), can be thought of as pairing against the virtual fundamental cycle of X). Now, there are also product maps

$$(2.4.7) \quad C_{\text{vir}}^\bullet(\mathcal{X}(p, q) \text{ rel } \partial) \otimes C_{\text{vir}}^\bullet(\mathcal{X}(q, r) \text{ rel } \partial) \rightarrow C_{\text{vir}}^\bullet([\mathcal{X}(p, q) \times \mathcal{X}(q, r)] \text{ rel } \partial),$$

which, when combined with the concatenation maps (2.4.2), induce maps

$$(2.4.8) \quad C_{\text{vir}}^\bullet(\mathcal{X}(p, q) \text{ rel } \partial) \otimes C_{\text{vir}}^\bullet(\mathcal{X}(q, r) \text{ rel } \partial) \rightarrow C_{\text{vir}}^\bullet(\partial\mathcal{X}(p, r)).$$

From construction, it is clear that the diagram

$$(2.4.9) \quad \begin{array}{ccc} C_{\text{vir}}^\bullet(\mathcal{X}(p, q) \text{ rel } \partial) \otimes C_{\text{vir}}^\bullet(\mathcal{X}(q, r) \text{ rel } \partial) & \xrightarrow{(2.4.8)} & C_{\text{vir}}^\bullet(\partial\mathcal{X}(p, r)) \\ \downarrow s_* \otimes s_* & & \downarrow s_* \\ C_\bullet(E_{\mathcal{A}(p,q)}, E_{\mathcal{A}(p,q)} \setminus 0) & \longrightarrow & C_\bullet(E_{\mathcal{A}(p,r)}, E_{\mathcal{A}(p,r)} \setminus 0) \\ \otimes C_\bullet(E_{\mathcal{A}(q,r)}, E_{\mathcal{A}(q,r)} \setminus 0) & & \end{array}$$

commutes, where the bottom map is simply the cartesian product on chains (recall that $\mathcal{A}(p, r) = \mathcal{A}(p, q) \sqcup \mathcal{A}(q, r)$).

Now, let us describe the construction of a boundary operator $d: \mathbb{Z}[\mathcal{P}] \rightarrow \mathbb{Z}[\mathcal{P}]$ given the flow category $(\mathcal{P}, \mathcal{X})$ and its system of implicit atlases. To define d , we need to choose:¹³

- (i) **Chain-level coherent orientations** Cochains $\lambda(p, q) \in C_{\text{vir}}^0(\mathcal{X}(p, q) \text{ rel } \partial)$ satisfying the following property. Define

$$(2.4.10) \quad \mu(p, r) := \sum_{p < q < r} \lambda(p, q) \cdot \lambda(q, r) \in C_{\text{vir}}^0(\partial\mathcal{X}(p, r))$$

¹³That these should be contractible choices is suggested by the Dold–Thom–Almgren theorem [4; 37, page 430].

(where we implicitly use (2.4.8) on the right-hand side). We require that $d\lambda(p, r)$ equal (the image under (2.4.3) of) $\mu(p, r)$. Thus $(\mu(p, r), \lambda(p, r))$ defines a cycle in the mapping cone of (2.4.3), and thus a homology class in $H_{\text{vir}}^0(\mathcal{X}(p, r)) = \check{H}^0(\mathcal{X}(p, r), \sigma_{\mathcal{X}(p, r)})$ (by (2.4.4)). We require that this homology class coincide with the (given) orientation on $\mathcal{X}(p, r)$.

(ii) **Thom cocycles** Cocycles

$$[[E_{\mathcal{A}(p,q)}]] \in C^{\dim E_{\mathcal{A}(p,q)}}(E_{\mathcal{A}(p,q)}, E_{\mathcal{A}(p,q)} \setminus 0)$$

whose pairing with $[E_{\mathcal{A}(p,q)}]$ is 1 and which are compatible in the sense that the following diagram commutes:

$$(2.4.11) \quad \begin{array}{ccc} C_{\bullet}(E_{\mathcal{A}(p,q)}, E_{\mathcal{A}(p,q)} \setminus 0) & & \\ \otimes C_{\bullet}(E_{\mathcal{A}(q,r)}, E_{\mathcal{A}(q,r)} \setminus 0) & \longrightarrow & C_{\bullet}(E_{\mathcal{A}(p,r)}, E_{\mathcal{A}(p,r)} \setminus 0) \\ \downarrow \llbracket E_{\mathcal{A}(p,q)} \rrbracket \otimes \llbracket E_{\mathcal{A}(q,r)} \rrbracket & & \swarrow \llbracket E_{\mathcal{A}(p,r)} \rrbracket \\ \mathbb{Z} & \longleftarrow & \end{array}$$

Given such choices, we define the boundary operator $d: \mathbb{Z}[\mathcal{P}] \rightarrow \mathbb{Z}[\mathcal{P}]$ to have “matrix coefficients” $\llbracket E_{\mathcal{A}(p,q)} \rrbracket (s_* \lambda(p, q))$. One can show that $d^2 = 0$ by using $d\lambda(p, q) = \mu(p, q)$ and the compatibilities (2.4.9) and (2.4.11).

Let us now sketch the proof of the existence of λ and $\llbracket E \rrbracket$ as above. We will work by induction on (p, q) ; that is, we show that valid $\lambda(p, q)$ and $\llbracket E_{\mathcal{A}(p,q)} \rrbracket$ exist given that we have fixed valid $\lambda(p', q')$ and $\llbracket E_{\mathcal{A}(p',q')} \rrbracket$ for all $p \preceq p' \prec q' \preceq q$ (other than $(p', q') = (p, q)$).

To construct $\lambda(p, q)$, argue as follows. Notice that $\mu(p, q)$ is automatically a cycle (apply (2.4.10) to expand its boundary and everything cancels). Now from a general statement about cycles in mapping cones, the existence of $\lambda(p, q)$ inducing the correct homology class in $H_{\text{vir}}^0(\mathcal{X}(p, q)) = \check{H}^0(\mathcal{X}(p, q); \sigma_{\mathcal{X}(p,q)})$ reduces to showing that the homology class

$$[\mu(p, q)] \in H_{\text{vir}}^0(\partial\mathcal{X}(p, q)) = \check{H}^0(\partial\mathcal{X}(p, q), \sigma_{\partial\mathcal{X}(p,q)})$$

is correct (namely, that it coincides with the image of the desired class under the coboundary map $\check{H}^0(\mathcal{X}(p, q); \sigma_{\mathcal{X}(p,q)}) \rightarrow \check{H}^0(\partial\mathcal{X}(p, q), \sigma_{\partial\mathcal{X}(p,q)})$). The key observation is that *this can be checked locally* since $\sigma_{\partial\mathcal{X}(p,q)}$ is a sheaf. Over the top strata of $\partial\mathcal{X}(p, q)$ (that is, those trajectories that split only once), the agreement is clear by the induction hypothesis on λ and the compatibility of the given coherent orientations. Unfortunately, the top strata may not be dense in $\partial\mathcal{X}(p, q)$, so we need to work harder (see Proposition 6.2.3). In the end, we must use the induction hypothesis for *all* $\lambda(p', q')$.

To construct $[[E_{\mathcal{A}(p,r)}]]$, argue as follows. First, observe that the “homology diagram”

$$(2.4.12) \quad \begin{array}{ccc} H_{\bullet}(E_{\mathcal{A}(p,q)}, E_{\mathcal{A}(p,q)} \setminus 0) & & \\ \otimes H_{\bullet}(E_{\mathcal{A}(q,r)}, E_{\mathcal{A}(q,r)} \setminus 0) & \longrightarrow & H_{\bullet}(E_{\mathcal{A}(p,r)}, E_{\mathcal{A}(p,r)} \setminus 0) \\ [E_{\mathcal{A}(p,q)}] \otimes [E_{\mathcal{A}(q,r)}] \mapsto 1 \downarrow & & \swarrow [E_{\mathcal{A}(p,r)}] \mapsto 1 \\ \mathbb{Z} & & \end{array}$$

trivially commutes. If the horizontal map in (2.4.11) were a cofibration (think “injective”) in a suitable sense, then the commutativity of (2.4.12) would be sufficient to imply the existence of $[[E_{\mathcal{A}(p,r)}]]$. Unfortunately, it is far from clear that this map is a cofibration; moreover, its failure to be a cofibration is a genuine obstruction to defining $[[E_{\mathcal{A}(p,r)}]]$ inductively. For the correct inductive construction, we must use “cofibrant replacements” for the system of complexes $C_{\bullet}(E_{\mathcal{A}(p,q)}, E_{\mathcal{A}(p,q)} \setminus 0)$ (the details of which we leave for Section 7).

Finally, we should argue that the homology groups defined via a choice of $(\lambda, [[E]])$ are in fact independent of that choice. For this, we use the usual strategy of constructing a chain map between $(\mathbb{Z}[\mathcal{P}], d)$ and $(\mathbb{Z}[\mathcal{P}], d')$ for any d and d' arising from $(\lambda, [[E]])$ and $(\lambda', [[E']])$ respectively (plus appropriate chain homotopies). To construct such a chain map, we use a similar inductive procedure, starting from the base choices $(\lambda, [[E]])$ and $(\lambda', [[E']])$. Note that for this argument for independence of choice to work, the inductive nature of our construction is crucial.

2.5 S^1 -localization

We now explain our strategy for proving S^1 -localization results for virtual fundamental cycles. The full details of our treatment appear in Section 8.

The most basic setting in which our results apply (and are interesting) is that of a free S^1 -space X (ie a space with a free action of S^1) equipped with an S^1 -equivariant implicit atlas A . Such an atlas is simply an implicit atlas where all the thickenings X_I are equipped with an action of S^1 so that the functions ψ_{IJ} are S^1 -equivariant and the functions s_I are S^1 -invariant. This last point bears repeating: S^1 does not act on the obstruction spaces E_I (or, alternatively, it acts trivially on them).

In the above setup, our S^1 -localization result states that $\pi_*[X]^{\text{vir}} = 0$, where $\pi: X \rightarrow X/S^1$ is the quotient map and $\pi_*: \check{H}^{\bullet}(X)^{\vee} \rightarrow \check{H}^{\bullet}(X/S^1)^{\vee}$ is the dual of the pullback $\pi^*: \check{H}^{\bullet}(X/S^1) \rightarrow \check{H}^{\bullet}(X)$. Observe that π_* is an isomorphism for $\bullet = 0$ (since S^1 is connected), hence this implies that if the virtual dimension d of X is zero, then $[X]^{\text{vir}} = 0$.

One should expect this result to be true if one believes that one can choose a “chain representative” of $[X]^{\text{vir}}$ which is S^1 -invariant (as the pushforward of such a chain representative is clearly null-homologous). For instance, in the perturbation approach, this result would follow if one could construct S^1 -invariant transverse perturbations. Conversely, one can interpret our vanishing result $\pi_*[X]^{\text{vir}} = 0$ as a sense in which our $[X]^{\text{vir}}$ is S^1 -invariant at the chain level.

Remark 2.5.1 A natural strategy for proving S^1 -localization results is to consider X_I/S^1 as forming an atlas on X/S^1 of virtual dimension one less, and then “pulling back” to X the virtual fundamental cycle on X/S^1 (this is the approach taken by Fukaya, Oh, Ohta and Ono [33; 31]). In the general setting where the S^1 -action is merely continuous, the implicit atlas on X does not induce an implicit atlas on X/S^1 because of the existence of free S^1 -actions on topological manifolds whose quotients are not manifolds,¹⁴ though we still consider this “quotient implicit atlas” in spirit. This extra generality is convenient, since it means we do not need to construct an S^1 -equivariant gluing map (providing the local slice necessary to show that the “quotient implicit atlas” exists).

Let us now prove our assertion $\pi_*[X]^{\text{vir}} = 0$ in the simple case of a single chart (for which we defined $[X]^{\text{vir}}$ in Section 2.3.2). In this case, the implicit atlas $A = \{\alpha\}$ consists of a topological manifold X_α , a function $s_\alpha: X_\alpha \rightarrow E_\alpha$, and an identification $X = s_\alpha^{-1}(0)$. This atlas being S^1 -equivariant means that X_α is equipped with an S^1 -action for which s_α is S^1 -invariant and which induces the given action on X . The desired statement follows from the following diagram, as we explain below:

$$\begin{array}{ccccc} \check{H}^d(X/S^1) & \cong & H_{\dim E_\alpha - 1}^{S^1}(X_\alpha, X_\alpha \setminus X) & \xrightarrow{(s_\alpha)_*} & H_{\dim E_\alpha - 1}^{S^1}(E_\alpha, E_\alpha \setminus 0) \\ \downarrow \pi^* & & \downarrow \pi^! & & \downarrow \pi^! \\ \check{H}^d(X) & \cong & H_{\dim E_\alpha}(X_\alpha, X_\alpha \setminus X) & \xrightarrow{(s_\alpha)_*} & H_{\dim E_\alpha}(E_\alpha, E_\alpha \setminus 0) \xrightarrow{[E_\alpha] \mapsto 1} \mathbb{Z} \end{array}$$

The map π^* is pullback under $\pi: X \rightarrow X/S^1$. The maps $\pi^!$ are the maps from the Gysin long exact sequence for an S^1 -space:

$$(2.5.1) \quad \dots \xrightarrow{\cap e} H_{\bullet-1}^{S^1}(Z) \xrightarrow{\pi^!} H_\bullet(Z) \xrightarrow{\pi_*} H_\bullet^{S^1}(Z) \xrightarrow{\cap e} H_{\bullet-2}^{S^1}(Z) \xrightarrow{\pi^!} \dots$$

¹⁴Such an action may be constructed out of any non-manifold X for which $X \times \mathbb{R}$ is a manifold (namely the obvious action on $X \times S^1$). Many examples of such spaces are known, the first being due to Bing [9]; see also Cannon [11].

The two leftmost horizontal identifications are a form of Poincaré duality (which can be proven, along with the commutativity of the square, with homotopy \mathcal{K} -sheaves as in Section 2.3.5); it is at this step where we use the fact that S^1 acts freely on X .

Now $[X]^{\text{vir}}$ is by definition the composition of the bottom row. Hence $\pi_*[X]^{\text{vir}}$ is the map from the upper left corner to the bottom right. On the other hand, the rightmost vertical map $\pi^!: H_{\dim E_{\alpha-1}}^{S^1}(E_{\alpha}, E_{\alpha} \setminus 0) \rightarrow H_{\dim E_{\alpha}}(E_{\alpha}, E_{\alpha} \setminus 0)$ vanishes since S^1 acts trivially on E_{α} . It follows that $\pi_*[X]^{\text{vir}} = 0$ as desired.

To generalize this approach to arbitrary implicit atlases, we introduce “ S^1 -equivariant virtual cochain complexes” which play the role of $H_{\dim E_{\alpha-1}}^{S^1}(X_{\alpha}, X_{\alpha} \setminus X)$ above. Morally speaking, these S^1 -equivariant virtual cochain complexes play the role of the virtual cochain complexes of the (non-existent; see Remark 2.5.1) induced implicit atlas on X/S^1 . To define the S^1 -equivariant virtual cochain complexes, we use the same definition as for the ordinary virtual cochain complexes, except with (shifted) S^1 -equivariant chains $C_{\bullet-1}^{S^1}$ in place of chains C_{\bullet} .

Remark 2.5.2 We do not construct an S^1 -equivariant virtual fundamental class, nor do we address S^1 -localization for actions which are not free or almost free (having finite stabilizer at every point). However, the machinery we develop could potentially be used for these purposes; see Remark 8.6.3.

3 Implicit atlases

3.1 Implicit atlases

Definition 3.1.1 (Implicit atlas) Let X be a compact Hausdorff space. An *implicit atlas of dimension $d = \text{vdim}_A X$* on X is an index set A along with the following data:

- (i) **Covering groups** A finite group Γ_{α} for all $\alpha \in A$ (let $\Gamma_I := \prod_{\alpha \in I} \Gamma_{\alpha}$).
- (ii) **Obstruction spaces** A finitely generated $\mathbb{R}[\Gamma_{\alpha}]$ -module E_{α} for all $\alpha \in A$ (let $E_I := \bigoplus_{\alpha \in I} E_{\alpha}$).
- (iii) **Thickenings** A Hausdorff Γ_I -space X_I for all finite $I \subseteq A$, and a homeomorphism $X \rightarrow X_{\emptyset}$.
- (iv) **Kuranishi maps** A Γ_{α} -equivariant function $s_{\alpha}: X_I \rightarrow E_{\alpha}$ for all $\alpha \in I \subseteq A$ (for $I \subseteq J$, let $s_I: X_J \rightarrow E_I$ denote $\bigoplus_{\alpha \in I} s_{\alpha}$).
- (v) **Footprints** A Γ_I -invariant open set $U_{IJ} \subseteq X_I$ for all $I \subseteq J \subseteq A$.
- (vi) **Footprint maps** A Γ_J -equivariant function $\psi_{IJ}: (s_{J \setminus I}|_{X_J})^{-1}(0) \rightarrow U_{IJ}$ for all $I \subseteq J \subseteq A$.
- (vii) **Regular locus** A Γ_I -invariant subset $X_I^{\text{reg}} \subseteq X_I$ for all $I \subseteq A$.

These data must satisfy the following “compatibility axioms”:

- (i) $\psi_{IJ}\psi_{JK} = \psi_{IK}$ and $\psi_{II} = \text{id}$.
- (ii) $s_I\psi_{IJ} = s_I$.
- (iii) $U_{IJ_1} \cap U_{IJ_2} = U_{I, J_1 \cup J_2}$ and $U_{II} = X_I$.
- (iv) $\psi_{IJ}^{-1}(U_{IK}) = U_{JK} \cap (s_{J \setminus I}|X_J)^{-1}(0)$.¹⁵
- (v) **Homeomorphism axiom** The map $(s_{J \setminus I}|X_J)^{-1}(0)/\Gamma_{J \setminus I} \rightarrow U_{IJ}$ induced by ψ_{IJ} is a homeomorphism.

They must also satisfy the following “transversality axioms”:

- (vi) $\psi_{IJ}^{-1}(X_I^{\text{reg}}) \subseteq X_J^{\text{reg}}$.
- (vii) $\Gamma_{J \setminus I}$ acts freely on $\psi_{IJ}^{-1}(X_I^{\text{reg}})$.
- (viii) **Openness axiom** $X_I^{\text{reg}} \subseteq X_I$ is open.
- (ix) **Submersion axiom** The map $s_{J \setminus I}: X_J \rightarrow E_{J \setminus I}$ is locally modeled on the projection $\mathbb{R}^{d+\dim E_I} \times \mathbb{R}^{\dim E_{J \setminus I}} \rightarrow \mathbb{R}^{\dim E_{J \setminus I}}$ over $\psi_{IJ}^{-1}(X_I^{\text{reg}}) \subseteq X_J$.
- (x) **Covering axiom** $X_\emptyset = \bigcup_{I \subseteq A} \psi_{\emptyset I}((s_I|X_I^{\text{reg}})^{-1}(0))$.

Remark 3.1.2 The VFC machinery in this paper would go through admitting a slight weakening of the axioms of an implicit atlas. For example, we only ever use the fact that the openness and submersion axioms hold in a neighborhood of $(s_I|X_I)^{-1}(0)$. We will not make this precise here, however, since the constructions of implicit atlases we know of would not be made any easier by such a weakening of the axioms.

Definition 3.1.3 (Smooth implicit atlas) A *smooth structure* on an implicit atlas consists of a smooth structure on each X_I^{reg} such that:

- (i) Γ_I acts smoothly on X_I^{reg} .
- (ii) s_I is smooth over X_I^{reg} .
- (iii) $s_{J \setminus I}: X_J \rightarrow E_{J \setminus I}$ is a smooth submersion over $\psi_{IJ}^{-1}(X_I^{\text{reg}})$.
- (iv) ψ_{IJ} is a local diffeomorphism over $\psi_{IJ}^{-1}(X_I^{\text{reg}})$.

Remark 3.1.4 The VFC machinery in this paper applies to implicit atlases (without a smooth structure), though the notion of a smooth implicit atlas may be useful for other applications of implicit atlases.

¹⁵Added in proof: Abouzaid recently pointed out that this axiom is superfluous; it is implied by the other axioms.

Remark 3.1.5 (Using finite $I \rightarrow A$ instead of finite $I \subseteq A$) An implicit atlas consists of data parametrized by the *category of finite subsets of A* (objects: finite subsets, morphisms: inclusions). A direct modification of the definition allows one to instead use the *category of finite sets over A* (objects: finite sets $I \rightarrow A$, morphisms: injective maps $I \hookrightarrow J$ compatible with the maps to A). In fact, all constructions of implicit atlases we know of yield implicit atlases in this generalized sense. We won't need this generalization in this paper, but let us point out some reasons why it may be useful to keep in mind.

With the definition as it stands now (using finite subsets of A), we can “pull back” an implicit atlas along any *injection* $B \hookrightarrow A$. If we instead use finite sets over A , then we can pull back an implicit atlas along any *map* $B \rightarrow A$ (in fact, we can pull back along any *coproduct-preserving functor* from finite sets over B to finite sets over A , which amounts to the specification of a finite set $I_\beta \rightarrow A$ for all $\beta \in B$). Also, in the category of finite subsets of A we can only take the disjoint union of subsets which are already disjoint; however in the category of finite sets over A there exist arbitrary abstract finite disjoint unions. This allows some extra (though currently unneeded) flexibility in certain constructions, since we do not need to ensure certain sets are disjoint or that certain maps are injective.

Using the category of sets over A is also the natural perspective to take if we wanted to allow A to be a groupoid instead of a set (then an object is a finite set I along with a collection of objects $\{\alpha_i\}_{i \in I}$ of A , and a morphism $(I, \{\alpha_i\}_{i \in I}) \rightarrow (J, \{\alpha_j\}_{j \in J})$ is an injection $j: I \hookrightarrow J$ along with isomorphisms $\alpha_i \xrightarrow{\sim} \alpha_{j(i)}$ in A).

3.2 Implicit atlases with boundary

Definition 3.2.1 (Implicit atlas with boundary) Let X be a compact Hausdorff space together with a closed subset denoted $\partial X \subseteq X$. An *implicit atlas of dimension d with boundary* on X consists of the same data as an implicit atlas, except that in addition we specify a Γ_I -invariant closed subset $\partial X_I \subseteq X_I$ for all $I \subseteq A$, such that $\partial X_\emptyset = \partial X$. We add the following “compatibility axiom”:

$$(i) \quad \psi_{IJ}^{-1}(\partial X_I) = (s_{J \setminus I} | \partial X_J)^{-1}(0).$$

We also modify one “transversality axiom”:

(ii) **Submersion axiom** We allow an additional local model

$$\mathbb{R}_{\geq 0} \times \mathbb{R}^{d + \dim E_I - 1} \times \mathbb{R}^{\dim E_{J \setminus I}} \rightarrow \mathbb{R}^{\dim E_{J \setminus I}},$$

and $\partial X_J^{\text{reg}} \subseteq X_J^{\text{reg}}$ must correspond to the boundary of the local model.

Remark 3.2.2 Just as a manifold is a special case of a manifold with boundary, an implicit atlas is a special case of an implicit atlas with boundary (namely where $\partial X_I = \emptyset$ for all I).

Definition 3.2.3 (Restriction of implicit atlas to boundary) Let X be a space with implicit atlas A of dimension d with boundary. Then this induces an implicit atlas A (the same index set) of dimension $d - 1$ on ∂X , simply by setting $(\partial X)_I := \partial X_I$ and restricting the rest of the data to these subspaces.

4 The VFC package

In this section, we develop the *VFC package*, which is the algebraic machinery we will apply in later sections to work with virtual fundamental cycles. The reader should be comfortable with the material from Appendix A, where we recall and develop the necessary foundational language of sheaves and homological algebra.

Convention 4.0.1 In this section, we work over a fixed ground ring R , and everything takes place in the category of R -modules. We restrict to implicit atlases A for which $\#\Gamma_\alpha$ is invertible in R for all $\alpha \in A$.

Let us now introduce the formalism of our VFC package.

For any space X equipped with a *finite locally orientable* implicit atlas A of dimension d with boundary, we define *virtual cochain complexes* $C_{\text{vir}}^\bullet(X; A)$ and $C_{\text{vir}}^\bullet(X \text{ rel } \partial; A)$. We construct canonical isomorphisms

$$(4.0.1) \quad H_{\text{vir}}^\bullet(X; A) \simeq \check{H}^\bullet(X; \sigma_X),$$

$$(4.0.2) \quad H_{\text{vir}}^\bullet(X \text{ rel } \partial; A) \simeq \check{H}^\bullet(X; \sigma_{X \text{ rel } \partial})$$

(H_{vir}^\bullet denotes the cohomology of C_{vir}^\bullet) for certain *virtual orientation sheaves* σ_X and $\sigma_{X \text{ rel } \partial}$ on X . There is a canonical map

$$(4.0.3) \quad C_{\text{vir}}^{d+\bullet}(X \text{ rel } \partial; A) \xrightarrow{S^*} C_{\dim E_A - \bullet}(E_A, E_A \setminus 0; \sigma_{E_A}^\vee)^{\Gamma_A}$$

which can be thought of as the (chain-level) *virtual fundamental cycle*.

To study the virtual cochain complexes (in particular, to construct the isomorphisms (4.0.1)–(4.0.2)), we define complexes of \mathcal{K} -presheaves on X

$$(4.0.4) \quad K \mapsto C_{\text{vir}}^\bullet(K; A),$$

$$(4.0.5) \quad K \mapsto C_{\text{vir}}^\bullet(K \text{ rel } \partial; A)$$

whose global sections are the virtual cochain complexes.^{16 17} We show that they are pure homotopy \mathcal{K} -sheaves, and that there are canonical isomorphisms of sheaves on X

$$(4.0.6) \quad H_{\text{vir}}^0(-; A) = \mathfrak{o}_X,$$

$$(4.0.7) \quad H_{\text{vir}}^0(- \text{rel } \partial; A) = \mathfrak{o}_{X \text{ rel } \partial}.$$

The isomorphisms (4.0.1)–(4.0.2) then follow from Proposition A.5.4.

The fact that the virtual cochain complexes are the global sections of pure homotopy \mathcal{K} -sheaves with known H^0 will also play a key role in the applications of the VFC package.

Another useful fact we prove here is that the isomorphisms (4.0.1)–(4.0.2) are compatible with the long exact sequence of the pair $(X, \partial X)$ in \check{H}^\bullet and a corresponding long exact sequence of H_{vir}^\bullet .

4.1 Orientations

Recall the notion of the orientation sheaf of a topological manifold (possibly with boundary) given in Definition A.6.2.

Definition 4.1.1 (Orientation module of a vector space) Let E be a finite-dimensional vector space over \mathbb{R} . We let $\mathfrak{o}_E = H_{\dim E}(E, E \setminus 0)$ denote the orientation module of E (a free R -module of rank 1).¹⁸

Definition 4.1.2 (Locally orientable implicit atlas) Let X be a space with implicit atlas A with boundary. We say that A is *locally orientable*¹⁹ if and only if for every $I \subseteq A$ and every $x \in (s_I|X_I^{\text{reg}})^{-1}(0)$, the stabilizer $(\Gamma_I)_x$ acts trivially on $(\mathfrak{o}_{X_I^{\text{reg}}})_x \otimes \mathfrak{o}_{E_I}^\vee$ (this action is always by a sign $(\Gamma_I)_x \rightarrow \{\pm 1\}$). This notion is independent of the ring R (due to our restriction that $\#\Gamma_\alpha$ be invertible in R).

Definition 4.1.3 (Virtual orientation sheaf \mathfrak{o}_X of a space with implicit atlas) Let X be a space with locally orientable implicit atlas A with boundary. Then there exists a sheaf $\mathfrak{o}_{X,A}$ on X equipped with Γ_I -equivariant isomorphisms $\psi_{\partial I}^* \mathfrak{o}_{X,A} \xrightarrow{\sim} \mathfrak{o}_{X_I^{\text{reg}}} \otimes \mathfrak{o}_{E_I}^\vee$ over $(s_I|X_I^{\text{reg}})^{-1}(0)$ for every $I \subseteq A$, which are compatible with the maps ψ_{IJ} . We

¹⁶Note that this is a certain abuse of notation, as A is not an implicit atlas on $K \subsetneq X$.

¹⁷Note that the map s_* (4.0.3) is global; it does not exist on $C_{\text{vir}}^\bullet(K \text{ rel } \partial; A)$ for $K \subsetneq X$ or on any $C_{\text{vir}}^\bullet(K; A)$.

¹⁸According to the previous definition, we should really call this $(\mathfrak{o}_E)_0$ (the stalk at $0 \in E$ of the orientation sheaf of E considered as a manifold), though we do not anticipate this abuse causing any particular confusion.

¹⁹This is analogous to the notion of an orbifold being locally orientable.

call $\mathfrak{o}_{X,A}$ the *virtual orientation sheaf* of X (it is unique up to unique isomorphism); it is locally isomorphic to the constant sheaf \underline{R} . We write \mathfrak{o}_X for $\mathfrak{o}_{X,A}$ when the atlas is clear from context. We let $\mathfrak{o}_{X \text{ rel } \partial} := j_! j^* \mathfrak{o}_X$ where $j: X \setminus \partial X \hookrightarrow X$.

4.2 Virtual cochain complexes $C_{\text{vir}}^\bullet(X; A)$ and $C_{\text{vir}}^\bullet(X \text{ rel } \partial; A)$

Definition 4.2.1 (Deformation to the normal cone) Let X be a space with finite implicit atlas A with boundary. For $I \subseteq J \subseteq A$, we define

$$X_{I,J,A} := \left\{ (e, t, x) \in E_A \times \mathbb{R}_{\geq 0}^A \times X_J^{\text{reg}} \left| \begin{array}{l} t_\alpha = 0 \text{ for } \alpha \in A \setminus I \\ s_\alpha(x) = t_\alpha e_\alpha \text{ for } \alpha \in J \\ \psi_{\{\alpha \in A: t_\alpha > 0\}, J}(x) \in X_{\{\alpha \in A: t_\alpha > 0\}}^{\text{reg}} \end{array} \right. \right\}.$$

The third condition ensures that we only deform to the normal cone of those zero sets $s_{J \setminus I'}^{-1}(0)$ which are cut out transversally ($I' \subseteq I$). Clearly $(\partial X)_{I,J,A} \subseteq X_{I,J,A}$ is the subset where $x \in \partial X_J^{\text{reg}}$ (the former being with respect to the restriction of A to ∂X). There are compatible maps

$$(4.2.1) \quad X_{I,J,A} \times E_{A' \setminus A} \rightarrow X_{I',J',A'} \quad \text{for } I \subseteq I' \subseteq J' \subseteq J \subseteq A \subseteq A'$$

given by $t_{A' \setminus A} = 0$ and $x \mapsto \psi_{J',J}(x)$. For $K \subseteq X$, we let $X_{I,J,A}^K \subseteq X_{I,J,A}$ denote the subspace where $e = 0$ and $x \in \psi_{\emptyset, J}^{-1}(K)$. Note that Γ_A acts on $X_{I,J,A}$ (acting on X_J^{reg} via the projection $\Gamma_A \rightarrow \Gamma_J$ and on E_A).

Remark 4.2.2 (Chains and cochains) We use $C_\bullet(X)$ and $C^\bullet(X)$ to denote singular simplicial chains and cochains on a space X , respectively, and we use $C_\bullet(X, Y)$ to denote the cokernel of $C_\bullet(Y) \hookrightarrow C_\bullet(X)$ for $Y \subseteq X$ (“relative chains”).

Let us also recall the “Eilenberg–Zilber map” $C_\bullet(X) \otimes C_\bullet(Y) \rightarrow C_\bullet(X \times Y)$ for spaces X and Y , corresponding to the standard subdivision of $\Delta^n \times \Delta^m$ into $\binom{n+m}{n}$ copies of Δ^{n+m} . The Eilenberg–Zilber map is associative (in the sense that it gives rise to a unique map $C_\bullet(X) \otimes C_\bullet(Y) \otimes C_\bullet(Z) \rightarrow C_\bullet(X \times Y \times Z)$) and commutative (in the sense that the diagram

$$(4.2.2) \quad \begin{array}{ccc} C_\bullet(X) \otimes C_\bullet(Y) & \rightarrow & C_\bullet(X \times Y) \\ \downarrow & & \downarrow \\ C_\bullet(Y) \otimes C_\bullet(X) & \rightarrow & C_\bullet(Y \times X) \end{array}$$

commutes for all X and Y). We should point out that this commutativity fails for some other common models of singular chains, for example singular cubical chains (modulo degeneracies).

Remark 4.2.3 (Independence of chain model) The particular choice of *singular simplicial* chains is not particularly important. The virtual fundamental classes, etc, resulting from our theory should be unchanged by using any other model of singular chains. This independence would follow immediately (and all issues about finding chain models with good chain-level functoriality properties would go away) if we set up the VFC package using ∞ -categories.

Definition 4.2.4 (Fundamental cycles of vector spaces) Let I be a finite subset of an implicit atlas A . We define

$$(4.2.3) \quad C_{\bullet}(E; I) := C_{\dim E_I + \bullet}(E_I, E_I \setminus 0; \mathfrak{o}_{E_I}^{\vee})^{\Gamma_I}.$$

There is a canonical isomorphism $H_{\bullet}(E; I) = R$ (concentrated in degree zero), and we denote the canonical generator by $[E_I] \in H_{\bullet}(E; I)$.

Definition 4.2.5 (Partial virtual cochain complexes $C_{\text{vir}}^{\bullet}(-; A)_{IJ}$ and $C_{\text{vir}}^{\bullet}(-\text{rel } \partial; A)_{IJ}$) Let X be a space with finite implicit atlas A of dimension d with boundary. For any compact $K \subseteq X$, we define

$$(4.2.4) \quad C_{\text{vir}}^{\bullet}(K \text{ rel } \partial; A)_{IJ} := C_{d + \dim E_A - \bullet}(X_{I,J,A}, X_{I,J,A} \setminus X_{I,J,A}^K; \mathfrak{o}_{E_A}^{\vee})^{\Gamma_A},$$

$$(4.2.5) \quad C_{\text{vir}}^{\bullet}(K; A)_{IJ} := \left[\begin{array}{c} C_{d + \dim E_A - 1 - \bullet}((\partial X)_{I,J,A}, (\partial X)_{I,J,A} \setminus (\partial X)_{I,J,A}^{K \cap \partial X}; \mathfrak{o}_{E_A}^{\vee})^{\Gamma_A} \\ \downarrow \\ C_{d + \dim E_A - \bullet}(X_{I,J,A}, X_{I,J,A} \setminus X_{I,J,A}^K; \mathfrak{o}_{E_A}^{\vee})^{\Gamma_A} \end{array} \right].$$

It is clear that $K \mapsto C_{\text{vir}}^{\bullet}(K; A)_{IJ}$ and $K \mapsto C_{\text{vir}}^{\bullet}(K \text{ rel } \partial; A)_{IJ}$ are both complexes of \mathcal{K} -presheaves on X .

There are canonical maps

$$(4.2.6) \quad C_{\text{vir}}^{\bullet}(K \text{ rel } \partial; A)_{IJ} \rightarrow C_{\text{vir}}^{\bullet}(K; A)_{IJ},$$

$$(4.2.7) \quad C_{\text{vir}}^{d+\bullet}(X \text{ rel } \partial; A)_{IJ} \xrightarrow{s_*} C_{-\bullet}(E; A),$$

$$(4.2.8) \quad C_{\text{vir}}^{\bullet}(-; A)_{IJ} \rightarrow C_{\text{vir}}^{\bullet}(-; A)_{I',J'} \quad \text{for } I \subseteq I' \subseteq J' \subseteq J$$

((4.2.7) is induced by the projection $X_{I,J,A} \rightarrow E_A$, and (4.2.8) is induced by (4.2.1) with $A = A'$). These are compatible with each other in that certain obvious diagrams commute.

Definition 4.2.6 (Virtual cochain complexes $C_{\text{vir}}^{\bullet}(-; A)$ and $C_{\text{vir}}^{\bullet}(-\text{rel } \partial; A)$) Let X be a space with finite implicit atlas A of dimension d with boundary. For any compact

$K \subseteq X$, we define

$$(4.2.9) \quad C_{\text{vir}}^\bullet(K; A) := \text{hocolim}_{I \subseteq J \subseteq A} C_{\text{vir}}^\bullet(K; A)_{IJ},$$

$$(4.2.10) \quad C_{\text{vir}}^\bullet(K \text{ rel } \partial; A) := \text{hocolim}_{I \subseteq J \subseteq A} C_{\text{vir}}^\bullet(K \text{ rel } \partial; A)_{IJ},$$

where $\text{hocolim}_{I \subseteq J \subseteq A}$ is the homotopy colimit (Definition A.7.2) over 2^A with structure maps given by $(\#\Gamma_{J \setminus J'})^{-1}$ times (4.2.8). It is clear that $K \mapsto C_{\text{vir}}^\bullet(K; A)$ and $K \mapsto C_{\text{vir}}^\bullet(K \text{ rel } \partial; A)$ are both complexes of \mathcal{K} -presheaves on X .

There are canonical maps

$$(4.2.11) \quad C_{\text{vir}}^\bullet(K \text{ rel } \partial; A) \rightarrow C_{\text{vir}}^\bullet(K; A),$$

$$(4.2.12) \quad C_{\text{vir}}^{d+\bullet}(X \text{ rel } \partial; A) \xrightarrow{S^*} C_{-\bullet}(E; A)$$

(induced by (4.2.6)–(4.2.7)). More precisely, (4.2.12) is given by (4.2.7) on the $p = 0$ part of the hocolim and is zero on the rest.

Definition 4.2.7 (Maps $C_{\text{vir}}^\bullet(-; A) \rightarrow C_{\text{vir}}^\bullet(-; A')$) Let X be a space with finite implicit atlases $A \subseteq A'$ with boundary. There are canonical maps

$$(4.2.13) \quad C_{\text{vir}}^\bullet(-; A)_{IJ} \otimes C_{-\bullet}(E; A' \setminus A) \rightarrow C_{\text{vir}}^\bullet(-; A')_{I', J'}$$

for $I \subseteq I' \subseteq J' \subseteq J \subseteq A \subseteq A'$ induced by (4.2.1) (of which (4.2.8) is a special case). These are compatible with each other, and thus extend to the homotopy colimit over $I \subseteq J \subseteq A$. Hence we get canonical maps

$$(4.2.14) \quad C_{\text{vir}}^\bullet(-; A) \otimes C_{-\bullet}(E; A' \setminus A) \rightarrow C_{\text{vir}}^\bullet(-; A')$$

which are compatible with (4.2.11)–(4.2.12).

Note that all of the complexes defined here are free R -modules (this follows from our assumption that $\#\Gamma_\alpha$ is invertible in R) and that everything in sight is compatible with base change $\otimes_R S$ for ring homomorphisms $R \rightarrow S$.

4.3 Isomorphisms $H_{\text{vir}}^\bullet(X; A) = \check{H}^\bullet(X; \mathfrak{o}_X)$ (also rel ∂)

Lemma 4.3.1 ($C_{\text{vir}}^\bullet(-; A)_{IJ}$ are pure homotopy \mathcal{K} -sheaves) *Let X be a space with finite locally orientable implicit atlas A with boundary. Then $C_{\text{vir}}^\bullet(-; A)_{IJ}$ and $C_{\text{vir}}^\bullet(- \text{ rel } \partial; A)_{IJ}$ are pure homotopy \mathcal{K} -sheaves on X . Furthermore, there are canonical isomorphisms of sheaves on X*

$$(4.3.1) \quad H_{\text{vir}}^0(-; A)_{IJ} = j!j^* \mathfrak{o}_X,$$

$$(4.3.2) \quad H_{\text{vir}}^0(- \text{ rel } \partial; A)_{IJ} = j!j^* \mathfrak{o}_{X \text{ rel } \partial},$$

where $j: V_I \cap V_J \hookrightarrow X$ for $V_I := \psi_{\emptyset I}((s_I|X_I^{\text{reg}})^{-1}(0)) \subseteq X$ (an open subset).

Proof Applying Lemmas A.6.3 and A.2.11, we see that both $C_{\text{vir}}^\bullet(-; A)_{IJ}$ and $C_{\text{vir}}^\bullet(-\text{rel } \partial; A)_{IJ}$ are homotopy \mathcal{K} -sheaves.

Now let us calculate $H_{\text{vir}}^\bullet(K; A)_{IJ}$ and $H_{\text{vir}}^\bullet(K \text{ rel } \partial; A)_{IJ}$. It follows from the submersion axiom that $X_{I,J,A}$ is a topological manifold of dimension $\text{vdim}_A X + \dim E_A + \#I$ with boundary (not necessarily second countable or paracompact).²⁰ The boundary $\partial(X_{I,J,A})$ is a (not necessarily disjoint) union of two pieces, namely the locus where $x \in \partial X_J^{\text{reg}}$ (which is precisely $(\partial X)_{I,J,A}$) and the locus where $t_\alpha = 0$ for some $\alpha \in I$. It is easy to see that the first piece $(\partial X)_{I,J,A} \subseteq \partial(X_{I,J,A})$ is a closed tamely embedded codimension-zero submanifold with boundary. Hence we may apply Poincaré–Lefschetz duality in the form of Lemma A.6.4 to see that

$$(4.3.3) \quad \begin{aligned} H_{\text{vir}}^\bullet(K \text{ rel } \partial; A)_{IJ} &= \check{H}_c^{\bullet+\#I}(\psi_{\partial J}^{-1}(K \cap V_I \cap V_J) \times \mathbb{R}_{>0}^I, \mathfrak{o}_{\mathbb{R}}^{\otimes I} \otimes \mathfrak{o}_{E_J}^\vee \otimes \mathfrak{o}_{X_J^{\text{reg}}} \text{rel } \partial)^{\Gamma_A}, \end{aligned}$$

$$(4.3.4) \quad \begin{aligned} H_{\text{vir}}^\bullet(K; A)_{IJ} &= \check{H}_c^{\bullet+\#I}(\psi_{\partial J}^{-1}(K \cap V_I \cap V_J) \times \mathbb{R}_{>0}^I, \mathfrak{o}_{\mathbb{R}}^{\otimes I} \otimes \mathfrak{o}_{E_J}^\vee \otimes \mathfrak{o}_{X_J^{\text{reg}}})^{\Gamma_A} \end{aligned}$$

(the orientation sheaf of $X_{I,J,A}$ is given by $\mathfrak{o}_{X_J^{\text{reg}}} \otimes \mathfrak{o}_{E_A} \otimes \mathfrak{o}_{E_J}^\vee \otimes \mathfrak{o}_{\mathbb{R}}^{\otimes I}$). Now we apply the Künneth formula to conclude that

$$(4.3.5) \quad H_{\text{vir}}^\bullet(K \text{ rel } \partial; A)_{IJ} = \check{H}_c^\bullet(\psi_{\partial J}^{-1}(K \cap V_I \cap V_J), \mathfrak{o}_{E_J}^\vee \otimes \mathfrak{o}_{X_J^{\text{reg}}} \text{rel } \partial)^{\Gamma_A},$$

$$(4.3.6) \quad H_{\text{vir}}^\bullet(K; A)_{IJ} = \check{H}_c^\bullet(\psi_{\partial J}^{-1}(K \cap V_I \cap V_J), \mathfrak{o}_{E_J}^\vee \otimes \mathfrak{o}_{X_J^{\text{reg}}})^{\Gamma_A}.$$

Since X is locally orientable, we have $\mathfrak{o}_{X_J^{\text{reg}}} \otimes \mathfrak{o}_{E_J}^\vee = \psi_{\partial J}^* \mathfrak{o}_X$, so we can write

$$(4.3.7) \quad H_{\text{vir}}^\bullet(K \text{ rel } \partial; A)_{IJ} = \check{H}_c^\bullet(\psi_{\partial J}^{-1}(K \cap V_I \cap V_J), \psi_{\partial J}^* \mathfrak{o}_{X \text{ rel } \partial})^{\Gamma_A},$$

$$(4.3.8) \quad H_{\text{vir}}^\bullet(K; A)_{IJ} = \check{H}_c^\bullet(\psi_{\partial J}^{-1}(K \cap V_I \cap V_J), \psi_{\partial J}^* \mathfrak{o}_X)^{\Gamma_A}.$$

Now Γ_A acts through the projection $\Gamma_A \rightarrow \Gamma_J$, and

$$\psi_{\partial J}: \psi_{\partial J}^{-1}(K \cap V_I \cap V_J) \rightarrow K \cap V_I \cap V_J$$

is exactly the quotient by Γ_J . Hence by Lemma A.4.9 we get isomorphisms

$$(4.3.9) \quad H_{\text{vir}}^\bullet(K \text{ rel } \partial; A)_{IJ} = \check{H}_c^\bullet(K \cap V_I \cap V_J, \mathfrak{o}_{X \text{ rel } \partial}),$$

$$(4.3.10) \quad H_{\text{vir}}^\bullet(K; A)_{IJ} = \check{H}_c^\bullet(K \cap V_I \cap V_J, \mathfrak{o}_X).$$

²⁰In fact, it carries a natural structure of a manifold with corners, which for the present purpose is mostly irrelevant.

Now use Lemma A.4.7 to write this as

$$(4.3.11) \quad H_{\text{vir}}^\bullet(K \text{ rel } \partial; A)_{IJ} = \check{H}^\bullet(K, j_! j^* \mathfrak{o}_{X \text{ rel } \partial}),$$

$$(4.3.12) \quad H_{\text{vir}}^\bullet(K; A)_{IJ} = \check{H}^\bullet(K, j_! j^* \mathfrak{o}_X).$$

Thus $C_{\text{vir}}^\bullet(-; A)_{IJ}$ and $C_{\text{vir}}^\bullet(- \text{ rel } \partial; A)_{IJ}$ are both pure, and we manifestly have the desired isomorphisms (4.3.1)–(4.3.2). \square

Lemma 4.3.2 (Local description of isomorphisms $H_{\text{vir}}^0(-; A)_{II} = j_! j^* \mathfrak{o}_X$) *Let X be a space with finite locally orientable implicit atlas A of dimension d with boundary. Consider the maps of complexes of \mathcal{K} -presheaves on X*

$$(4.3.13) \quad C_{\text{vir}}^\bullet(K; A)_{II} \leftarrow \left[\begin{array}{c} C_{d+\dim E_A-1-\bullet}(E_A \setminus I \times \partial X_I^{\text{reg}}, E_A \setminus I \times (\partial X_I^{\text{reg}} \setminus \psi_{\partial I}^{-1}(K)); \mathfrak{o}_{E_A}^\vee)^{\Gamma_A} \\ \downarrow \\ C_{d+\dim E_A-\bullet}(E_A \setminus I \times X_I^{\text{reg}}, E_A \setminus I \times (X_I^{\text{reg}} \setminus \psi_{\partial I}^{-1}(K)); \mathfrak{o}_{E_A}^\vee)^{\Gamma_A} \end{array} \right],$$

$$(4.3.14) \quad C_{\text{vir}}^\bullet(K \text{ rel } \partial; A)_{II} \leftarrow C_{d+\dim E_A-\bullet}(E_A \setminus I \times X_I^{\text{reg}}, E_A \setminus I \times (X_I^{\text{reg}} \setminus \psi_{\partial I}^{-1}(K)); \mathfrak{o}_{E_A}^\vee)^{\Gamma_A}$$

(induced by the corresponding maps on spaces). These maps are quasi-isomorphisms, and the isomorphisms (4.3.1)–(4.3.2) from Lemma 4.3.1 coincide with the Poincaré duality isomorphisms for the complexes on the right above.

Proof Clear from the proof of Lemma 4.3.1. \square

Proposition 4.3.3 ($C_{\text{vir}}^\bullet(-; A)$ are pure homotopy \mathcal{K} -sheaves) *Let X be a space with finite locally orientable implicit atlas A with boundary. Then $C_{\text{vir}}^\bullet(-; A)$ and $C_{\text{vir}}^\bullet(- \text{ rel } \partial; A)$ are pure homotopy \mathcal{K} -sheaves on X . Furthermore, there are canonical isomorphisms of sheaves on X*

$$(4.3.15) \quad H_{\text{vir}}^0(-; A) = \mathfrak{o}_X,$$

$$(4.3.16) \quad H_{\text{vir}}^0(- \text{ rel } \partial; A) = \mathfrak{o}_{X \text{ rel } \partial}.$$

Proof This is a special case of Lemma A.8.2; we just need to make sure all of the hypotheses are satisfied. The open cover in question is $V_I := \psi_{\partial I}((s_I|X_I^{\text{reg}})^{-1}(0))$; it follows from the axioms of an implicit atlas that $V_I \cap V_K \subseteq V_J$ for $I \subseteq J \subseteq K$ and $V_I \cap V_{I'} \subseteq V_{I \cup I'}$ for all I, I' . Now we just need to check that the system of isomorphisms from Lemma 4.3.1 is compatible with the maps of the homotopy diagram. These being maps of sheaves, it suffices to check compatibility locally (on stalks) and thus is a straightforward calculation (this is where the extra normalization factor of $(\#\Gamma_{J \setminus J'})^{-1}$ in Definition 4.2.6 is important). \square

Theorem 4.3.4 (Calculation of H_{vir}^\bullet) *Let X be a space with finite locally orientable implicit atlas A with boundary. Then there are canonical isomorphisms*

$$(4.3.17) \quad H_{\text{vir}}^\bullet(X; A) = \check{H}^\bullet(X; \mathfrak{o}_X),$$

$$(4.3.18) \quad H_{\text{vir}}^\bullet(X \text{ rel } \partial; A) = \check{H}^\bullet(X; \mathfrak{o}_{X \text{ rel } \partial}).$$

Proof By Proposition A.5.4, this is a consequence of Proposition 4.3.3. □

4.4 Long exact sequence for the pair $(X, \partial X)$

In this subsection, we compare two long exact sequences for the pair $(X, \partial X)$, namely the one in Čech cohomology and one coming from the virtual cochain complexes.

It follows from the definition that $C_{\text{vir}}^\bullet(K; A)_{IJ}$ is the mapping cone of the obvious map $C_{\text{vir}}^{\bullet-1}(K \cap \partial X; A)_{IJ} \rightarrow C_{\text{vir}}^\bullet(K \text{ rel } \partial; A)_{IJ}$. Since homotopy colimit commutes with the formation of mapping cones, we see that the same is true dropping the IJ subscript. Hence there are natural maps

$$(4.4.1) \quad \dots \rightarrow C_{\text{vir}}^{\bullet-1}(K \cap \partial X; A) \rightarrow C_{\text{vir}}^\bullet(K \text{ rel } \partial; A) \rightarrow C_{\text{vir}}^\bullet(K; A) \rightarrow \dots$$

and they induce a long exact sequence on cohomology.

Similarly, there is a sequence of sheaves on X which is exact on stalks:

$$(4.4.2) \quad 0 \rightarrow \mathfrak{o}_{X \text{ rel } \partial} \rightarrow \mathfrak{o}_X \rightarrow i_* \mathfrak{o}_{\partial X} \rightarrow 0$$

(where $i: \partial X \hookrightarrow X$). This induces a long exact sequence on Čech cohomology (Lemma A.4.20). Note that $\check{H}^\bullet(X; i_* \mathfrak{o}_{\partial X}) = \check{H}^\bullet(\partial X; \mathfrak{o}_{\partial X})$ (Lemma A.4.8).

Remark 4.4.1 To be slightly pedantic about orientations, it would be more precise to say that $C_{\text{vir}}^\bullet(-; A) = [C_{\text{vir}}^\bullet(- \cap \partial X; A) \otimes \mathfrak{o}_{\mathbb{R}} \rightarrow C_{\text{vir}}^\bullet(- \text{ rel } \partial; A)]$, where we identify $\mathfrak{o}_{\mathbb{R}}$ canonically with the orientation line of the normal bundle of $\partial X_I^{\text{reg}} \subseteq X_I^{\text{reg}}$. Also, note that we should either say that the last map in (4.4.2) is odd or that it is really $\mathfrak{o}_X \rightarrow i_* \mathfrak{o}_{\partial X} \otimes \mathfrak{o}_{\mathbb{R}}$ (and is even).

Proposition 4.4.2 (Compatibility of long exact sequences of the pair $(X, \partial X)$) *The following diagram commutes:*

$$(4.4.3) \quad \begin{array}{ccccccc} H_{\text{vir}}^\bullet(X \text{ rel } \partial; A) & \xrightarrow{(4.4.1)} & H_{\text{vir}}^\bullet(X; A) & \xrightarrow{(4.4.1)} & H_{\text{vir}}^\bullet(\partial X; A) & \xrightarrow{(4.4.1)} & H_{\text{vir}}^{\bullet+1}(X \text{ rel } \partial; A) \\ \parallel \text{Thm 4.3.4} & & \parallel \text{Thm 4.3.4} & & \parallel \text{Thm 4.3.4} & & \parallel \text{Thm 4.3.4} \\ \check{H}^\bullet(X; \mathfrak{o}_{X \text{ rel } \partial}) & \xrightarrow{(4.4.2)} & \check{H}^\bullet(X; \mathfrak{o}_X) & \xrightarrow{(4.4.2)} & \check{H}^\bullet(\partial X; \mathfrak{o}_{\partial X}) & \xrightarrow{(4.4.2)} & \check{H}^{\bullet+1}(X; \mathfrak{o}_{X \text{ rel } \partial}) \end{array}$$

Proof By the definition of the vertical identifications, the commutativity of the first two squares of (4.4.3) reduces to the commutativity of the diagram of sheaves:

$$(4.4.4) \quad \begin{array}{ccccc} H_{\text{vir}}^0(-\text{rel } \partial; A) & \xrightarrow{(4.4.1)} & H_{\text{vir}}^0(-; A) & \xrightarrow{(4.4.1)} & H_{\text{vir}}^0(-; A) \\ \parallel \text{Prop 4.3.3} & & \parallel \text{Prop 4.3.3} & & \parallel \text{Prop 4.3.3} \\ \mathfrak{o}_{X \text{ rel } \partial} & \xrightarrow{(4.4.2)} & \mathfrak{o}_X & \xrightarrow{(4.4.2)} & i_* \mathfrak{o}_{\partial X} \end{array}$$

It suffices to check commutativity on stalks, and to do this, we can work with the complexes on the right in Lemma 4.3.2, for which the commutativity is clear.

In fact, the commutativity of the last square of (4.4.3) is also a consequence of the commutativity of (4.4.4). To see this, consider the diagram:

$$\begin{array}{ccccc} H_{\text{vir}}^\bullet(\partial X; A) & \xleftarrow{\sim} & H^\bullet[C_{\text{vir}}^{\bullet+1}(X \text{ rel } \partial; A) \rightarrow C_{\text{vir}}^\bullet(X; A)] & \longrightarrow & H_{\text{vir}}^{\bullet+1}(X \text{ rel } \partial; A) \\ \parallel \text{Prop A.5.4} & & \parallel \text{Prop A.5.4} & & \parallel \text{Prop A.5.4} \\ \check{H}^\bullet(\partial X; H_{\text{vir}}^0(-; A)) & & & & \check{H}^{\bullet+1}(X, H_{\text{vir}}^0(-\text{rel } \partial; A)) \\ & \swarrow \sim & & \searrow & \\ & \check{H}^\bullet(X; [H_{\text{vir}}^0(-\text{rel } \partial; A)[1] \rightarrow H_{\text{vir}}^0(-; A)]) & & & \\ & \downarrow \sim & & & \downarrow \sim \\ \check{H}^\bullet(\partial X; \mathfrak{o}_{\partial X}) & \xleftarrow{\sim} & \check{H}^\bullet(X; [\mathfrak{o}_{X \text{ rel } \partial}[1] \rightarrow \mathfrak{o}_X]) & \longrightarrow & \check{H}^{\bullet+1}(X, \mathfrak{o}_{X \text{ rel } \partial}) \end{array}$$

The middle vertical arrow is defined because of the commutativity of the left square in (4.4.4). The top left horizontal arrow is a quasi-isomorphism, and thus every left horizontal arrow is a quasi-isomorphism. The bottom left square is commutative because of the commutativity of the right square in (4.4.4); the other squares are trivially commutative. Reversing the left horizontal arrows, we see that the outermost square commutes, and this is exactly the last square in (4.4.3). \square

5 Virtual fundamental classes

In this section, we use the technical results of Section 4 to define the virtual fundamental class of a space with implicit atlas and derive some of its properties. We also show how these properties can be used to calculate the virtual fundamental class in some special situations (calculation directly from the definition seems prohibitively complicated in all but the simplest of cases). The properties we prove here are sufficient for some rudimentary purposes, and we think they at least demonstrate that the virtual fundamental class we have defined is the “right” one.

For more sophisticated applications than those considered in this paper, one would certainly like to have more properties than those proven here. For example, one would very much like to prove the expected formula for $[X \times_M Y]^{\text{vir}}$ in terms of $[X]^{\text{vir}}$ and $[Y]^{\text{vir}}$ (given some natural “fiber product implicit atlas” on $X \times_M Y$, where X, Y are spaces with implicit atlases and M is a manifold).

Convention 5.0.1 In this section, we work over a fixed ground ring R , and everything takes place in the category of R -modules. We restrict to implicit atlases A for which $\#\Gamma_\alpha$ is invertible in R for all $\alpha \in A$.

Remark 5.0.2 (Comparison of homology theories) There are (at least) three natural “homology groups” which one can assign to a compact Hausdorff space X :

- (i) *Dual of Čech cohomology* $\check{H}^\bullet(X)^\vee$ ($^\vee$ denotes dual, ie $\text{Hom}(-, R)$).
- (ii) *Čech homology* $\check{H}_\bullet(X)$ (the inverse limit of the homology of nerves of finite covers).
- (iii) *Steenrod–Sitnikov homology* $\bar{H}_\bullet(X)$ (the homology of the homotopy/derived inverse limit of nerves of finite covers; see Section A.9).

These are successively more refined, in the sense that there are natural maps

$$(5.0.1) \quad \bar{H}_\bullet(X) \rightarrow \check{H}_\bullet(X) \rightarrow \check{H}^\bullet(X)^\vee.$$

If X is homeomorphic to a finite CW-complex, then there are natural isomorphisms

$$(5.0.2) \quad \bar{H}_\bullet(X) = H_\bullet(X), \quad \check{H}_\bullet(X) = H_\bullet(X), \quad \check{H}^\bullet(X)^\vee = H^\bullet(X)^\vee.$$

It follows that a virtual fundamental class in any of these groups can be used for all the applications we are aware of. On the other hand, there are some potential advantages to working with the more refined homology groups:

- (i) \bar{H}_\bullet and \check{H}_\bullet retain torsion information.
- (ii) \bar{H}_\bullet and \check{H}_\bullet have a natural “extension of scalars” map for any map of rings $R \rightarrow S$.
- (iii) \bar{H}_\bullet has the expected long exact sequence for pairs of spaces.

We first define the virtual fundamental class as an element of $\check{H}_\bullet(X)^\vee$ (with appropriate twisted coefficients) and derive some properties. At the end, we indicate how to define a canonical lift to $\bar{H}_\bullet(X)$ by working at the level of the derived category.

5.1 Definition

Definition 5.1.1 (Virtual fundamental class $[X]^{\text{vir}}$) Let X be a space with locally orientable implicit atlas A of dimension d with boundary. Let $B \subseteq A$ be a finite subatlas (which exists by compactness). We consider the composite

$$(5.1.1) \quad \check{H}^{d+\bullet}(X; \mathfrak{o}_{X \text{ rel } \partial}) = H_{\text{vir}}^{d+\bullet}(X \text{ rel } \partial; B) \xrightarrow{(4.2.12)} H_{-\bullet}(E; B) \xrightarrow{[E_B] \mapsto 1} R,$$

where the equality is the canonical isomorphism from Theorem 4.3.4. We thus get a map $\check{H}^d(X; \mathfrak{o}_{X \text{ rel } \partial}) \rightarrow R$. Now suppose $B \subseteq B'$ are two finite subatlases. Then the following diagram commutes (it suffices to check commutativity on stalks, and to do this, we can work with the complexes on the right in Lemma 4.3.2, for which the commutativity is clear):

$$(5.1.2) \quad \begin{array}{ccc} \mathfrak{o}_{X \text{ rel } \partial} & \xrightarrow{\text{Thm 4.3.4}} & H_{\text{vir}}^0(- \text{rel } \partial; B) \\ \parallel & & \downarrow (4.2.14) \times [E_{B' \setminus B}] \\ \mathfrak{o}_{X \text{ rel } \partial} & \xrightarrow{\text{Thm 4.3.4}} & H_{\text{vir}}^0(- \text{rel } \partial; B') \end{array}$$

Hence the following diagram commutes (the first square following from (5.1.2), the rest being clear):

$$(5.1.3) \quad \begin{array}{ccccc} \check{H}^{d+\bullet}(X; \mathfrak{o}_{X \text{ rel } \partial}) & \xrightarrow{\text{Thm 4.3.4}} & H_{\text{vir}}^{d+\bullet}(X \text{ rel } \partial; B) & \xrightarrow{(4.2.12)} & H_{-\bullet}(E; B) & \xrightarrow{[E_B] \mapsto 1} & R \\ \parallel & & \downarrow (4.2.14) \times [E_{B' \setminus B}] & & \downarrow \times [E_{B' \setminus B}] & & \parallel \\ \check{H}^{d+\bullet}(X; \mathfrak{o}_{X \text{ rel } \partial}) & \xrightarrow{\text{Thm 4.3.4}} & H_{\text{vir}}^{d+\bullet}(X \text{ rel } \partial; B') & \xrightarrow{(4.2.12)} & H_{-\bullet}(E; B') & \xrightarrow{[E_{B'}] \mapsto 1} & R \end{array}$$

This shows that the maps $\check{H}^d(X; \mathfrak{o}_{X \text{ rel } \partial}) \rightarrow R$ induced by B and B' coincide. Since any two $B_1, B_2 \subseteq A$ are contained in a third $B_1 \cup B_2 \subseteq A$, we see that the resulting element

$$(5.1.4) \quad [X]_A^{\text{vir}} \in \check{H}^d(X; \mathfrak{o}_{X \text{ rel } \partial})^\vee$$

is independent of B . We write $[X]^{\text{vir}}$ for $[X]_A^{\text{vir}}$ when the atlas is clear from context.

5.2 Properties

Lemma 5.2.1 (Passing to a subatlas preserves $[X]^{\text{vir}}$) Let X be a space with locally orientable implicit atlas A with boundary. If $B \subseteq A$ is any subatlas, then $[X]_A^{\text{vir}} = [X]_B^{\text{vir}}$.

Proof This follows immediately from the definition. □

Lemma 5.2.2 (Shrinking the charts preserves $[X]^{\text{vir}}$) *Let X be a space with locally orientable implicit atlas A with boundary. Let A' be obtained from A by using instead some open subsets $U'_{IJ} \subseteq U_{IJ}$, $X'_I \subseteq X_I$, $X'^{\text{reg}}_I \subseteq X^{\text{reg}}_I$, and restricting ψ_{IJ} , s_I to these subsets, so that A' is also an implicit atlas. Then $[X]_{A'}^{\text{vir}} = [X]_A^{\text{vir}}$.*

Proof We may assume that A is finite. Certainly there is a map $C_{\text{vir}}^\bullet(X \text{ rel } \partial; A') \rightarrow C_{\text{vir}}^\bullet(X \text{ rel } \partial; A)$ which respects the map (4.2.12). Hence it suffices to show that the diagram

$$(5.2.1) \quad \begin{array}{ccc} H_{\text{vir}}^\bullet(X \text{ rel } \partial; A') & \xrightarrow{\text{Thm 4.3.4}} & \check{H}^\bullet(X; \mathfrak{o}_X \text{ rel } \partial) \\ \downarrow & & \parallel \\ H_{\text{vir}}^\bullet(X \text{ rel } \partial; A) & \xrightarrow{\text{Thm 4.3.4}} & \check{H}^\bullet(X; \mathfrak{o}_X \text{ rel } \partial) \end{array}$$

commutes, where the left vertical arrow is induced by the obvious pushforward on chains. By definition of the isomorphisms above, it suffices to check commutativity of the corresponding diagram of sheaves with $\bullet = 0$. This can be checked locally, where it is clear. \square

Lemma 5.2.3 ($[X \sqcup Y]^{\text{vir}} = [X]^{\text{vir}} \oplus [Y]^{\text{vir}}$) *Let X and Y be spaces equipped with locally orientable implicit atlases with boundary, on the same index set A and of the same virtual dimension. Let us also denote by A the resulting implicit atlas on $X \sqcup Y$ (let $(X \sqcup Y)_I := X_I \sqcup Y_I$). Then A on $X \sqcup Y$ is locally orientable with $\mathfrak{o}_{X \sqcup Y} = (i_X)_* \mathfrak{o}_X \oplus (i_Y)_* \mathfrak{o}_Y$ and*

$$[X \sqcup Y]^{\text{vir}} = [X]^{\text{vir}} \oplus [Y]^{\text{vir}}$$

in

$$\check{H}^\bullet(X \sqcup Y; \mathfrak{o}_{X \sqcup Y} \text{ rel } \partial)^\vee = \check{H}^\bullet(X; \mathfrak{o}_X \text{ rel } \partial)^\vee \oplus \check{H}^\bullet(Y; \mathfrak{o}_Y \text{ rel } \partial)^\vee.$$

Proof We may assume that A is finite. There is a natural isomorphism

$$(5.2.2) \quad C_{\text{vir}}^\bullet(X \sqcup Y \text{ rel } \partial; A) = C_{\text{vir}}^\bullet(X \text{ rel } \partial; A) \oplus C_{\text{vir}}^\bullet(Y \text{ rel } \partial; A)$$

compatible with the map to $C_{-\bullet}(E; A)$. Thus it suffices to show that the following square of isomorphisms commutes:

$$(5.2.3) \quad \begin{array}{ccc} H_{\text{vir}}^\bullet(X \sqcup Y \text{ rel } \partial; A) & \xrightarrow{\text{Thm 4.3.4}} & \check{H}^\bullet(X \sqcup Y; \mathfrak{o}_{X \sqcup Y} \text{ rel } \partial) \\ \parallel & & \parallel \\ H_{\text{vir}}^\bullet(X \text{ rel } \partial; A) & \xrightarrow{\text{Thm 4.3.4}} & \check{H}^\bullet(X; \mathfrak{o}_X \text{ rel } \partial) \\ \oplus H_{\text{vir}}^\bullet(Y \text{ rel } \partial; A) & & \oplus \check{H}^\bullet(Y; \mathfrak{o}_Y \text{ rel } \partial) \end{array}$$

By definition of the horizontal maps, it suffices to check the commutativity of the corresponding diagram of sheaves on $X \sqcup Y$, which is clear. \square

Lemma 5.2.4 ($\partial[X]^{\text{vir}} = [\partial X]^{\text{vir}}$) *Let X be a space with locally orientable implicit atlas A with boundary. Then the dual connecting homomorphism $\delta^\vee: \check{H}^\bullet(X; \mathfrak{o}_{X \text{ rel } \partial})^\vee \rightarrow \check{H}^{\bullet-1}(\partial X; \mathfrak{o}_{\partial X})^\vee$ sends $[X]^{\text{vir}}$ to $[\partial X]^{\text{vir}}$.*

Proof We may assume that A is finite. The map $C_{\text{vir}}^\bullet(\partial X; A) \rightarrow C_{\text{vir}}^{\bullet+1}(X \text{ rel } \partial; A)$ in (4.4.1) commutes with the maps from both of these groups to $C_{d-\bullet}(E; A)$. Hence the result follows from the commutativity of the last square in Proposition 4.4.2. \square

Lemma 5.2.5 ($(\partial X \rightarrow X)_*[\partial X]^{\text{vir}} = 0$) *Let X be a space with locally orientable implicit atlas with boundary. Then the pushforward map $\check{H}_{\text{vir}}^\bullet(\partial X; \mathfrak{o}_{\partial X})^\vee \rightarrow \check{H}^\bullet(X; \mathfrak{o}_X)^\vee$ annihilates $[\partial X]^{\text{vir}}$.*

Proof The composition $\check{H}^{\bullet+1}(X; \mathfrak{o}_{X \text{ rel } \partial})^\vee \rightarrow \check{H}^\bullet(\partial X; \mathfrak{o}_{\partial X})^\vee \rightarrow \check{H}^\bullet(X; \mathfrak{o}_X)^\vee$ is zero, and $[\partial X]^{\text{vir}}$ is in the image of the first map by Lemma 5.2.4. \square

Lemma 5.2.6 (If $X = X^{\text{reg}}$ then $[X]^{\text{vir}} = [X]$) *Let X be a space with locally orientable implicit atlas A with boundary. If $X^{\text{reg}} = X$ (so in particular, X is a compact topological manifold with boundary of dimension $d = \text{vdim}_A X$), then $[X]_A^{\text{vir}}$ is the usual fundamental class of X .*

Proof We may replace A with the finite subatlas $\emptyset \subseteq A$. Now in this case, we have

$$\begin{aligned}
 (5.2.4) \quad C_{\text{vir}}^\bullet(K \text{ rel } \partial; \emptyset) &= C_{\text{vir}}^\bullet(K \text{ rel } \partial; \emptyset)_{\emptyset \emptyset} \\
 &= C_{d-\bullet}(X_{\emptyset, \emptyset, \emptyset}, X_{\emptyset, \emptyset, \emptyset} \setminus X_{\emptyset, \emptyset, \emptyset}^K) \\
 &= C_{d-\bullet}(X, X \setminus K).
 \end{aligned}$$

It follows that the identification

$$\check{H}^\bullet(X; \mathfrak{o}_{X \text{ rel } \partial}) = H_{\text{vir}}^\bullet(X \text{ rel } \partial; \emptyset) = H_{d-\bullet}(X)$$

is simply the usual Poincaré duality isomorphism, and the map

$$H_{-\bullet}(X) \rightarrow H_{-\bullet}(E; \emptyset) \rightarrow R$$

is the usual augmentation on H_0 . Hence the composition $\check{H}^\bullet(X; \mathfrak{o}_{X \text{ rel } \partial}) \rightarrow R$ is the pairing against the usual fundamental class of X . \square

Remark 5.2.7 (Calculating $[X]^{\text{vir}}$ using perturbation) The properties from this section can be used to calculate the virtual fundamental class in the following sense. Let us suppose we are given an “explicit” implicit atlas A on a space X . We may always (for convenience) replace A by a subatlas and/or shrink the charts of A and the virtual fundamental class is preserved (Lemmas 5.2.1 and 5.2.2). Now suppose we can extend A to an implicit atlas with boundary on a space Y with $\partial Y = X \cup X'$ where $X' = (X')^{\text{reg}}$ (if A is a smooth implicit atlas, then such an extension is obtained if one can “coherently perturb” the Kuranishi maps s_α so they become transverse to zero). Then X' is a closed smooth manifold, and $[X']^{\text{vir}} = [X']$ is the naive fundamental class (Lemma 5.2.6). Now Y is a cobordism between X and X' , so we have $[X]^{\text{vir}} = [X']$ as homology classes in Y (Lemmas 5.2.3 and 5.2.5). This allows us to understand the pushforward of $[X]^{\text{vir}}$ under any map $f: X \rightarrow Z$ which extends continuously to Y . We do not claim such perturbations always exist; an affirmative answer is provided in very similar, though not identical, contexts by Fukaya and Ono [33], Fukaya, Oh, Ohta and Ono [30; 31] and McDuff and Wehrheim [63].

5.3 Manifold with obstruction bundle

A natural “test case” (beyond Lemma 5.2.6 and Remark 5.2.7) for our definition of the virtual fundamental class is that of the natural implicit atlas on a “manifold with obstruction bundle” (the expected answer being the Poincaré dual of the Euler class). More generally, there is a natural implicit atlas on the zero set of a section of a vector bundle over a manifold (the “manifold with obstruction bundle” case is when the section is identically zero), and again the expected answer is a type of Euler class.

In this section, we show that our definition of the virtual fundamental class indeed agrees with this expected answer. To prove this, we use only the properties of the virtual fundamental class from Section 5.2.

Definition 5.3.1 (Implicit atlas on the zero set of a smooth section) Let B be a smooth manifold with boundary, let $p: E \rightarrow B$ be a smooth vector bundle, and let $s: B \rightarrow E$ be a smooth section with $s^{-1}(0)$ compact. We define an implicit atlas of dimension $\dim B - \dim E$ with boundary on $X := s^{-1}(0)$ as follows. The index set A consists of all triples $(V_\alpha, E_\alpha, \lambda_\alpha)$ (called *thickening datums*) where:

- (i) $V_\alpha \subseteq B$ is an open subset.
- (ii) E_α is a finite-dimensional vector space.
- (iii) $\lambda_\alpha: E_\alpha \times V_\alpha \rightarrow p^{-1}(V_\alpha)$ is a smooth homomorphism of vector bundles.

The thickened spaces are defined as follows:

$$(5.3.1) \quad X_I := \left\{ (x, \{e_\alpha\}_{\alpha \in I}) \in \bigcap_{\alpha \in I} V_\alpha \times \bigoplus_{\alpha \in I} E_\alpha \mid s(x) + \sum_{\alpha \in I} \lambda_\alpha(x, e_\alpha) = 0 \right\}.$$

The Kuranishi map $s_\alpha: X_I \rightarrow E_\alpha$ is the obvious projection map, the footprint U_{IJ} is the locus where $x \in V_\alpha$ for all $\alpha \in J$, and the footprint maps $\psi_{IJ}: (s_{J \setminus I}|_{X_J})^{-1}(0) \rightarrow U_{IJ}$ are the obvious forgetful maps. The compatibility axioms are an easy exercise.

The regular locus $X_I^{\text{reg}} \subseteq X_I$ is the locus where X_I is ‘‘cut out transversally’’; more precisely, $(x, \{e_\alpha\}_{\alpha \in I}) \in X_I$ is in X_I^{reg} if and only if the map

$$(5.3.2) \quad \left(ds + \sum_{\alpha \in I} d\lambda_\alpha(\cdot, e_\alpha) \right) \oplus \bigoplus_{\alpha \in I} \lambda_\alpha(x, \cdot): T_x B \oplus \bigoplus_{\alpha \in I} E_\alpha \rightarrow E_x$$

is surjective. The transversality axioms are an easy exercise. Thus A is indeed an implicit atlas on X .

Since everything here is in the smooth category, A is in fact a smooth implicit atlas.

Definition 5.3.2 (Euler class) Let B be a space equipped with a vector bundle $p: E \rightarrow B$ of rank k . This induces a canonical locally constant sheaf \mathfrak{o}_E on B whose stalk at $b \in B$ is \mathfrak{o}_{E_b} . Let $\tau_E \in H^k(E, E \setminus 0; p^*\mathfrak{o}_E)$ denote the Thom class of E (characterized uniquely by the property that its restriction to any local trivialization $\mathbb{R}^n \times U \rightarrow U$ is the pullback of the tautological class in $H^n(\mathbb{R}^n, \mathbb{R}^n \setminus 0; \mathfrak{o}_{\mathbb{R}^n})$).

The Euler class $e(E) \in H^k(B; \mathfrak{o}_E)$ is $s^*\tau_E$, where $s: B \rightarrow E$ is any section. Any two sections are homotopic, so $e(E)$ is well-defined. For any section $s: B \rightarrow E$ with $s^{-1}(0)$ compact, the Euler class with compact support $e_c(E, s) \in H_c^k(B; \mathfrak{o}_E)$ is $s^*\tau_E$. If s_0 and s_1 are homotopic through a section $\tilde{s}: B \times [0, 1] \rightarrow E$ with compact support, then $e_c(E, s_0) = e_c(E, s_1)$. The natural map $H_c^k(B; \mathfrak{o}_E) \rightarrow H^k(B; \mathfrak{o}_E)$ sends the Euler class with compact support to the Euler class.

Remark 5.3.3 The Euler class with compact support can be nonzero even for the trivial vector bundle.

Proposition 5.3.4 Let B be a smooth manifold, let $p: E \rightarrow B$ be a smooth vector bundle, and let $s: B \rightarrow E$ be a smooth section with $s^{-1}(0)$ compact. Consider the implicit atlas A on $X := s^{-1}(0)$ from Definition 5.3.1. Then there is a canonical isomorphism of sheaves $\mathfrak{o}_X = \mathfrak{o}_B \otimes \mathfrak{o}_E^\vee$ on X (in particular, A is locally orientable), and the image of the virtual fundamental class $[X]^{\text{vir}}$ in $H^{\dim B - \dim E}(B; \mathfrak{o}_B \otimes \mathfrak{o}_E^\vee)^\vee$ equals $e_c(E, s) \cap [B]$.

Note that $\check{H}^\bullet(X; \mathfrak{o}_X) = \varinjlim_{X \subseteq U} H^\bullet(U; \mathfrak{o}_B \otimes \mathfrak{o}_E^\vee)$ (direct limit over open neighborhoods $U \subseteq B$ of X), by the continuity axiom of Čech cohomology. The proposition applies equally well with U in place of B , and thus determines $[X]^{\text{vir}} \in \check{H}^\bullet(X; \mathfrak{o}_X)^\vee$ uniquely (note the use of either Lemma 5.2.1 or Lemma 5.2.2).

Proof The statement about orientation sheaves is a straightforward calculation which we omit.

Choose some smooth family of sections $\tilde{s}: B \times [0, 1] \rightarrow E$ with $\tilde{s}^{-1}(0)$ compact, so that $\tilde{s}(\cdot, 0) = s$ and $s_1 := \tilde{s}(\cdot, 1)$ is transverse to the zero section.

Now $B \times [0, 1]$ is a smooth manifold with boundary, so let \tilde{A} be the implicit atlas on $\tilde{X} := \tilde{s}^{-1}(0)$ from Definition 5.3.1. Similarly, let A' be the implicit atlas on $X' := (s')^{-1}(0)$.

Now we have two implicit atlases on X , namely A and (the restriction to the boundary of) \tilde{A} . On the other hand, we can exhibit A as a subatlas of \tilde{A} by the map on thickening datums sending V_α to $V_\alpha \times [0, 1]$ and λ_α to its obvious extension. Hence they induce the same virtual fundamental class (Lemma 5.2.1), so we may just write $[X]^{\text{vir}}$ (dropping the subscript indicating which atlas we use). The same reasoning applies to X' .

Now $\partial\tilde{X} = X \sqcup X'$, so we have $[X]^{\text{vir}} = [X']^{\text{vir}}$ as homology classes on \tilde{X} (Lemmas 5.2.3 and 5.2.5). Hence we have

$$(5.3.3) \quad [X]^{\text{vir}} = [X']^{\text{vir}}$$

as elements of $H^{\dim B - \dim E}(B; \mathfrak{o}_B \otimes \mathfrak{o}_E^\vee)^\vee$.

We know $X' = (X')^{\text{reg}}$ (it is cut out transversally), so X' is a closed smooth manifold and $[X']^{\text{vir}} = [X']$ (Lemma 5.2.6). But now $[X'] = e_c(E, s) \cap [B]$ by elementary Poincaré duality. \square

Remark 5.3.5 Proposition 5.3.4 in its present form is useless in practice, because the implicit atlases which arise in interesting examples are not literally isomorphic to the one given in Definition 5.3.1. One could, though, hope to show that they are equivalent (see Remark 2.1.5) and then apply Proposition 5.3.4.

5.4 Lift to Steenrod homology

We now define a virtual fundamental class in Steenrod homology $\overline{H}_\bullet(X, \partial X; \mathfrak{o}_X^\vee)$ refining the class in $\check{H}^\bullet(X; \mathfrak{o}_{X \text{ rel } \partial})^\vee$ constructed in Section 5.1. To do this, it suffices to lift our earlier reasoning at the level of homology groups to the level of objects in the derived category. We will be using the notation and results from Section A.9.

Remark 5.4.1 The natural map $\bar{H}_\bullet(X, \partial X; \mathfrak{o}_X^\vee) \rightarrow \check{H}^\bullet(X; \mathfrak{o}_{X \text{ rel } \partial})^\vee$ is an isomorphism if the base ring R is a field (this follows from Lemma A.9.12), but for $R = \mathbb{Z}$ the kernel can be substantial (in particular, it contains all torsion).

Definition 5.4.2 (Virtual fundamental class in Steenrod homology) Let X be a space with locally orientable implicit atlas A of dimension d with boundary. Let $B \subseteq A$ be a finite subatlas. By Propositions 4.3.3 and A.5.4, there is a canonical isomorphism $\check{C}^\bullet(X; \mathfrak{o}_{X \text{ rel } \partial}) = C_{\text{vir}}^\bullet(X \text{ rel } \partial; B)$ in $D(R)$ (observe that the isomorphisms in the proof of Proposition A.5.4 all come from canonical quasi-isomorphisms on the chain level). There is also a canonical isomorphism $C_{-\bullet}(E; B) = R$ in $D(R)$. Since $C_{\text{vir}}^\bullet(X \text{ rel } \partial; B)$ is free and bounded above, we have

$$(5.4.1) \quad \text{Hom}_R(C_{\text{vir}}^\bullet(X \text{ rel } \partial; B), C_{-\bullet}(E; B)) = R \text{Hom}_{D(R)}(\check{C}^\bullet(X; \mathfrak{o}_{X \text{ rel } \partial}), R).$$

Thus the map $s_*: C_{\text{vir}}^{d+\bullet}(X \text{ rel } \partial; B) \rightarrow C_{-\bullet}(E; B)$ gives rise to an element

$$(5.4.2) \quad [X]_A^{\text{vir}} := [s_*] \in \text{Hom}_{D(R)}(\check{C}^\bullet(X; \mathfrak{o}_{X \text{ rel } \partial}), R[-d]) = \bar{H}_d(X, \partial X; \mathfrak{o}_X^\vee),$$

where the equality is the canonical isomorphism from Lemma A.9.12. The commutativity of (5.1.2) implies that this element is independent of the chosen finite subatlas $B \subseteq A$.

By definition, this fundamental class (5.4.2) projects to the fundamental class of Definition 5.1.1 under the map $\bar{H}_\bullet(X, \partial X; \mathfrak{o}_X^\vee) \rightarrow \check{H}^\bullet(X; \mathfrak{o}_{X \text{ rel } \partial})^\vee$. By chasing the various isomorphisms involved in its definition, it can be checked that $[X]^{\text{vir}} \in \bar{H}_d(X, \partial X; \mathfrak{o}_X^\vee)$ is preserved under extension of scalars.

The statements and proofs of Lemmas 5.2.1–5.2.6 and Proposition 5.3.4 generalize readily to the case of $[X]^{\text{vir}} \in \bar{H}_d(X, \partial X; \mathfrak{o}_X^\vee)$. Note that (the generalization of) Proposition 5.3.4 does not determine $[X]^{\text{vir}} \in \bar{H}_d(X, \partial X; \mathfrak{o}_X^\vee)$ uniquely since the map from the Steenrod homology of X to the inverse limit of the homology of its neighborhoods is not necessarily an isomorphism (see [64, page 87, Theorem 4]). However, it does at least determine the image of $[X]^{\text{vir}}$ under any map from X to a finite CW-complex, since any such map extends to some neighborhood of X .

6 Stratifications

In this section, we introduce *implicit atlases with cell-like stratification*. Roughly speaking, an implicit atlas with cell-like stratification is an implicit atlas on a stratified space, along with suitably compatible stratifications on each of the thickenings. We show how to apply the VFC package in this setting to obtain a “stratum-by-stratum”

understanding of virtual fundamental cycles. We also define the *product implicit atlas*, a natural implicit atlas on $X \times Y$ induced from implicit atlases on X and Y .

Convention 6.0.1 In this section, we work over a fixed ground ring R , and everything takes place in the category of R -modules. We restrict to implicit atlases A for which $\#\Gamma_\alpha$ is invertible in R for all $\alpha \in A$.

6.1 Implicit atlas with cell-like stratification

Definition 6.1.1 (Stratification) Let X be a topological space and let \mathcal{S} be a poset. A *stratification* of X by \mathcal{S} is a lower semicontinuous function $X \rightarrow \mathcal{S}$. We let $X^{\leq s}$, X^s and $X^{\geq s}$ denote the inverse images of $\mathcal{S}^{\leq s}$, $\{s\}$ and $\mathcal{S}^{\geq s}$, respectively; lower semicontinuity of $X \rightarrow \mathcal{S}$ means by definition that every $X^{\geq s}$ is open.

For a pair of spaces $(X, \partial X)$, a stratification $(X, \partial X) \rightarrow (\mathcal{S}, \partial \mathcal{S})$ shall mean a stratification $X \rightarrow \mathcal{S}$ along with a downward closed (ie closed under taking smaller elements) subset $\partial \mathcal{S} \subseteq \mathcal{S}$ such that $X^{\partial \mathcal{S}} = \partial X$ (note that this implies automatically that $\partial X \subseteq X$ is closed).

Definition 6.1.2 (Cell-like stratification) Let $(M, \partial M)$ be a topological manifold with boundary and stratification by $(\mathcal{S}, \partial \mathcal{S})$, and fix a map $\dim: \mathcal{S} \rightarrow \mathbb{Z}$. This stratification is called *cell-like* if and only if each pair $(M^{\leq s}, M^{< s})$ is a topological manifold with boundary of dimension $\dim s$.

If $(M, \partial M) \rightarrow (\mathcal{S}, \partial \mathcal{S})$ is cell-like, then so is $\partial M \rightarrow \partial \mathcal{S}$ (with empty boundary) and $(M^{\leq s}, M^{< s}) \rightarrow (\mathcal{S}^{\leq s}, \mathcal{S}^{< s})$.

Example 6.1.3 Let T be a simplicial complex, and suppose that its geometric realization $M := |T|$ is a topological manifold with boundary (then ∂M necessarily corresponds to a subcomplex $\partial T \subseteq T$). Let $(\mathcal{S}, \partial \mathcal{S}) := (\mathcal{F}(T), \mathcal{F}(\partial T))$ consist of the face posets of T and ∂T . Then the stratification $(M, \partial M) \rightarrow (\mathcal{S}, \partial \mathcal{S})$ is cell-like (note that T need not be a PL manifold).

Example 6.1.4 The natural stratification $\mathbb{R}_{\geq 0}^k \times \mathbb{R}^{n-k} \rightarrow \{(0, \infty) > \{0\}\}^k$ is cell-like. The natural stratification $\mathbb{R}_{\geq 0}^k \times \mathbb{R}^{n-k} \rightarrow \mathbb{Z}$ given by $(n$ minus) the number of zeros in the first k coordinates is *not* cell-like for $k \geq 2$.

Lemma 6.1.5 (Some local properties of cell-like stratifications) *Let $(M, \partial M) \rightarrow (\mathcal{S}, \partial \mathcal{S})$ be cell-like. Then:*

- (i) *If $s < t$ then $M^s \subseteq \partial M^{\leq t}$.*
- (ii) *If $M^s \neq \emptyset$, then $\dim s \leq \dim M$.*

- (iii) If $\dim \mathfrak{s} = \dim M$, then $M^{\mathfrak{s}} \subseteq M$ is open.
- (iv) If $\dim \mathfrak{s} = \dim M - 1$, $\mathfrak{s} \in \partial \mathcal{S}$, and $M^{\mathfrak{s}} \neq \emptyset$, then $\#\mathcal{S}^{>\mathfrak{s}} = 1$.
- (v) If $\dim \mathfrak{s} = \dim M - 1$, $\mathfrak{s} \notin \partial \mathcal{S}$, and $M^{\mathfrak{s}} \neq \emptyset$, then $\#\mathcal{S}^{>\mathfrak{s}} = 2$, and these $M^{\leq \mathfrak{t}}$ give collars on either side of $M^{\mathfrak{s}} \subseteq M$ (so in particular $M^{\mathfrak{s}} \subseteq M$ is locally flat).

Proof Since $\mathfrak{s} < \mathfrak{t}$, we have $M^{\mathfrak{s}} \subseteq M^{<\mathfrak{t}} = \partial M^{\leq \mathfrak{t}}$, giving (i).

Recall Brouwer’s “invariance of domain”, which implies that if a subset $X \subseteq \mathbb{R}^n$ is (in the subspace topology) locally homeomorphic to \mathbb{R}^m , then $m \leq n$, with equality if and only if X is open. This immediately gives (ii), (iii).

We prove (iv). Fix $p \in M^{\mathfrak{s}}$. A neighborhood of p is covered by strata $\succ \mathfrak{s}$, so since $M^{\mathfrak{s}} \subseteq M$ is not open, there exists a stratum $\mathfrak{t} \succ \mathfrak{s}$. Now $M^{\mathfrak{s}} \subseteq \partial M^{\leq \mathfrak{t}}$ is open by (iii) (applied to $M^{<\mathfrak{t}} \rightarrow \mathcal{S}^{<\mathfrak{t}}$). A doubling argument (and invariance of domain) near $p \in \partial M$ shows that $M^{\mathfrak{s}} \cup M^{\mathfrak{t}}$ contains a neighborhood of p . Since we have exhausted a neighborhood of $p \in M^{\mathfrak{s}}$, it follows using (i) that there is no other stratum $\succ \mathfrak{s}$.

We prove (v). Fix $p \in M^{\mathfrak{s}}$. A neighborhood of p is covered by strata $\succ \mathfrak{s}$, so since $M^{\mathfrak{s}} \subseteq M$ is not open, there exists a stratum $\mathfrak{t} \succ \mathfrak{s}$. Now $M^{\mathfrak{s}} \subseteq \partial M^{\leq \mathfrak{t}}$ is open by (iii), as is $M^{\mathfrak{t}} \subseteq M$. This gives a collar on one side of $M^{\mathfrak{s}} \subseteq M$. This still does not fill out a neighborhood of p , so there exists $\mathfrak{t}' \neq \mathfrak{t}$ with $\mathfrak{t}' \succ \mathfrak{s}$, giving a collar on the other side of $M^{\mathfrak{s}} \subseteq M$. Now by invariance of domain, we have thus exhausted a neighborhood of $p \in M^{\mathfrak{s}}$ inside M , and hence it follows using (i) that there are no more strata $\succ \mathfrak{s}$ other than $\mathfrak{t}, \mathfrak{t}'$. □

Definition 6.1.6 (Implicit atlas with cell-like stratification) Let $(X, \partial X)$ be a pair of compact Hausdorff spaces equipped with a stratification $(X, \partial X) \rightarrow (\mathcal{S}, \partial \mathcal{S})$ (\mathcal{S} finite) along with a map $\dim: \mathcal{S} \rightarrow \mathbb{Z}$. Also fix “orientation data”:

- (i) For every $\mathfrak{s} \in \mathcal{S}$, an orientation line $\mathfrak{o}_{\mathfrak{s}}$ (ie a free $\mathbb{Z}/2$ -graded \mathbb{Z} -module of rank one).
- (ii) For $\text{codim}(\mathfrak{s} \leq \mathfrak{t}) = 1$, an odd “coboundary” map $\mathfrak{o}_{\mathfrak{s}} \rightarrow \mathfrak{o}_{\mathfrak{t}}$.

An implicit atlas of dimension d with boundary and cell-like stratification on $(X, \partial X) \rightarrow (\mathcal{S}, \partial \mathcal{S})$ consists of the same data as an implicit atlas with boundary, except that in addition we specify a Γ_I -invariant stratification $(X_I, \partial X_I) \rightarrow (\mathcal{S}, \partial \mathcal{S})$ for all $I \subseteq A$, whose restriction to X_{\emptyset} is the given stratification. We add the following “compatibility axiom”:

- (i) The restriction of the stratification on X_J to $(s_{J \setminus I} | X_J)^{-1}(0)$ coincides with the pullback of the stratification on X_I via ψ_{IJ} .

We also modify one “transversality axiom”:

- (ii) **Submersion axiom** In the local model $\mathbb{R}^{d+\dim E_I} \times \mathbb{R}^{\dim E_{J \setminus I}} \rightarrow \mathbb{R}^{\dim E_{J \setminus I}}$ (or $(\mathbb{R}_{\geq 0} \times \mathbb{R}^{d+\dim E_I - 1}) \times \mathbb{R}^{\dim E_{J \setminus I}} \rightarrow \mathbb{R}^{\dim E_{J \setminus I}}$), we require the stratification on the domain to be pulled back from a cell-like stratification on the first factor (for $\mathfrak{s} \mapsto \dim \mathfrak{s} + \dim E_I$).

We also specify isomorphisms of sheaves $\mathfrak{o}_{(X_I^{\text{reg}})^{\leq \mathfrak{s}}} \otimes \mathfrak{o}_{\mathfrak{s}} \xrightarrow{\sim} \mathfrak{o}_{X_I^{\text{reg}}}$ over $(X_I^{\text{reg}})^{\leq \mathfrak{s}}$ which are compatible with ψ_{IJ} and such that the diagram

$$(6.1.1) \quad \begin{array}{ccc} \mathfrak{o}_{(X_I^{\text{reg}})^{\leq \mathfrak{s}}} \otimes \mathfrak{o}_{\mathfrak{s}} & \longrightarrow & \mathfrak{o}_{X_I^{\text{reg}}} \\ \downarrow & & \parallel \\ \mathfrak{o}_{(X_I^{\text{reg}})^{\leq \mathfrak{t}}} \otimes \mathfrak{o}_{\mathfrak{t}} & \longrightarrow & \mathfrak{o}_{X_I^{\text{reg}}} \end{array}$$

commutes for $\text{codim}(\mathfrak{s} \preceq \mathfrak{t}) = 1$ (the left vertical map is the tensor product of the coboundary map $\mathfrak{o}_{\mathfrak{s}} \rightarrow \mathfrak{o}_{\mathfrak{t}}$ and the inverse of the boundary map $\mathfrak{o}_{(X_I^{\text{reg}})^{\leq \mathfrak{t}}} \rightarrow \mathfrak{o}_{(X_I^{\text{reg}})^{\leq \mathfrak{s}}}$). In other words, we identify $\mathfrak{o}_{\mathfrak{s}}$ with the orientation line of the normal bundle of $(X_I^{\text{reg}})^{\leq \mathfrak{s}} \subseteq X_I^{\text{reg}}$, so that the coboundary maps $\mathfrak{o}_{\mathfrak{s}} \rightarrow \mathfrak{o}_{\mathfrak{t}}$ coincide with the geometric coboundary maps on normal bundles.

Given an implicit atlas with boundary and cell-like stratification A on $(X, \partial X) \rightarrow (\mathcal{S}, \partial \mathcal{S})$, we may obtain by restriction to the corresponding strata an implicit atlas with cell-like stratification on $\partial X \rightarrow \partial \mathcal{S}$ (empty boundary, tensor every orientation line with $\mathfrak{o}_{\mathbb{R}}$) and an implicit atlas with boundary and cell-like stratification on $(X^{\leq \mathfrak{s}}, X^{< \mathfrak{s}}) \rightarrow (\mathcal{S}^{\leq \mathfrak{s}}, \mathcal{S}^{< \mathfrak{s}})$ (tensor every orientation line with $\mathfrak{o}_{\mathfrak{s}}^{\vee}$).

Remark 6.1.7 There should be a slightly more general setting for the results of this section (and their proofs), which takes as input a weakened version of Definition 6.1.2. For example, it is probably enough to require that in a neighborhood of any $p \in M$, the closure of any (local) component of $M^{\mathfrak{s}}$ is a manifold with boundary whose interior is this local component (this is satisfied by both stratifications in Example 6.1.4); basically this allows “non-embedded faces”.

6.2 Stratified virtual cochain complexes

In this section, we apply the VFC package to obtain a “stratum-by-stratum” understanding of virtual fundamental cycles on a space with implicit atlas with cell-like stratification. To do this, we build a complex out of the virtual cochain complexes associated to each stratum and then study the properties of this larger complex. This construction can be viewed as a generalization of the definition of $C_{\text{vir}}^{\bullet}(X; A)$ as the mapping cone $[C_{\text{vir}}^{\bullet}(\partial X; A) \rightarrow C_{\text{vir}}^{\bullet}(X \text{ rel } \partial; A)]$, and the main result here can be viewed as a generalization of Proposition 4.4.2.

Example 6.2.1 As a first step towards understanding the main construction of this section, let us first describe a similar construction in a more familiar setting. Let M be a topological manifold with cell-like stratification by \mathcal{S} . Define

$$(6.2.1) \quad C_\bullet(M; \mathcal{S}) := \bigoplus_{\mathfrak{s} \in \mathcal{S}} C_{\bullet - (\dim M - \dim \mathfrak{s})}(M^{\leq \mathfrak{s}}; \mathfrak{o}_M \otimes \mathfrak{o}_{M^{\leq \mathfrak{s}}}^\vee),$$

equipped with the differential given by the sum over $\text{codim}(\mathfrak{s} \preceq \mathfrak{t}) = 1$ of the pushforwards $C_\bullet(M^{\leq \mathfrak{s}}) \rightarrow C_\bullet(M^{\leq \mathfrak{t}})$ (covered by the dual of the boundary map $\mathfrak{o}_{M^{\leq \mathfrak{t}}} \rightarrow \mathfrak{o}_{M^{\leq \mathfrak{s}}}$) plus the internal differential. This differential squares to zero by Lemma 6.1.5 (the square of the differential is a sum over $\text{codim}(\mathfrak{s} \preceq \mathfrak{t}) = 2$ of maps $C_\bullet(M^{\leq \mathfrak{s}}) \rightarrow C_\bullet(M^{\leq \mathfrak{t}})$, each of which can be seen to vanish by applying Lemma 6.1.5(v) to $M^{< \mathfrak{t}} \rightarrow \mathcal{S}^{< \mathfrak{t}}$). Now there is a natural map

$$(6.2.2) \quad C_\bullet(M; \mathcal{S}) \rightarrow C_\bullet(M),$$

$$(6.2.3) \quad \bigoplus_{\mathfrak{s} \in \mathcal{S}} \gamma_{\mathfrak{s}} \mapsto \sum_{\substack{\mathfrak{s} \in \mathcal{S} \\ \dim \mathfrak{s} = \dim M}} \gamma_{\mathfrak{s}}$$

(it follows from Lemma 6.1.5 that this is a chain map) which we claim is a quasi-isomorphism.

To see that (6.2.2) is a quasi-isomorphism, observe that it is the map on global sections of a corresponding map of complexes of \mathcal{K} -presheaves

$$C_\bullet(M, M \setminus K; \mathcal{S}) \rightarrow C_\bullet(M, M \setminus K)$$

on the one-point compactification M_+ of M . Both are homotopy \mathcal{K} -sheaves by Lemma A.6.3. Thus by Corollary A.4.19, it suffices to show that the map

$$C_\bullet(M, M \setminus p; \mathcal{S}) \rightarrow C_\bullet(M, M \setminus p)$$

is a quasi-isomorphism for every $p \in M$. This holds by the following local argument. Note that $H_\bullet(M, M \setminus p; \mathcal{S})$ is isomorphic to \mathbb{Z} (to see this, consider the filtration by $\dim \mathfrak{s}$), so it suffices to construct a cycle in $C_\bullet(M, M \setminus p; \mathcal{S})$ representing a generator of \mathbb{Z} and show that its image in $H_\bullet(M, M \setminus p) = \mathbb{Z}$ is a generator. Such a cycle may be constructed by induction on \mathcal{S} (starting from the stratum containing p and going up), and its image in $H_\bullet(M, M \setminus p)$ generates since its degree at p coincides with its degree at any nearby point in an open stratum, which is one by construction. A similar argument appears in Barraud and Cornea [6, page 670, Lemma 2.2].

Thus we can think of $C_\bullet(M; \mathcal{S})$ as a model for chains on M (the reader should make sure they understand what this means geometrically).

Definition 6.2.2 (Stratified virtual cochain complexes $C_{\text{vir}}^\bullet(-, \mathcal{S}; A)$) Let $(X, \partial X) \rightarrow (\mathcal{S}, \partial\mathcal{S})$ be equipped with a finite implicit atlas with boundary and cell-like stratification A . For any compact $K \subseteq X$, we define

$$(6.2.4) \quad C_{\text{vir}}^\bullet(K, \mathcal{S}; A) := \bigoplus_{\mathfrak{s} \in \mathcal{S}} C_{\text{vir}}^\bullet(K \cap X^{\leq \mathfrak{s}} \text{ rel } \partial; A) \otimes \mathfrak{o}_{\mathfrak{s}},$$

(on the right, A refers to the restriction of the atlas to $X^{\leq \mathfrak{s}}$) equipped with the differential given by the sum over $\text{codim}(\mathfrak{s} \leq \mathfrak{t}) = 1$ of the pushforwards

$$(6.2.5) \quad C_{\text{vir}}^\bullet(K \cap X^{\leq \mathfrak{s}} \text{ rel } \partial; A) \rightarrow C_{\text{vir}}^{\bullet+1}(K \cap X^{\leq \mathfrak{t}} \text{ rel } \partial; A)$$

tensored with the specified coboundary map $\mathfrak{o}_{\mathfrak{s}} \rightarrow \mathfrak{o}_{\mathfrak{t}}$ (plus the internal differential). The fact that this differential squares to zero follows from Lemma 6.1.5 and the compatibility of the coboundary maps $\mathfrak{o}_{\mathfrak{s}} \rightarrow \mathfrak{o}_{\mathfrak{t}}$ with the geometric boundary maps $\mathfrak{o}_{(X_I^{\text{reg}})^{\leq \mathfrak{t}}} \rightarrow \mathfrak{o}_{(X_I^{\text{reg}})^{\leq \mathfrak{s}}}$.

There is a natural map

$$(6.2.6) \quad C_{\text{vir}}^\bullet(K, \mathcal{S}; A) \xrightarrow{\sim} C^\bullet(K; A),$$

$$(6.2.7) \quad \{\gamma_{\mathfrak{s}}\}_{\mathfrak{s} \in \mathcal{S}} \mapsto \left(\sum_{\substack{\mathfrak{s} \in \partial\mathcal{S} \\ \dim \mathfrak{s} = d-1}} \gamma_{\mathfrak{s}}, \sum_{\substack{\mathfrak{s} \in \mathcal{S} \\ \dim \mathfrak{s} = d}} \gamma_{\mathfrak{s}} \right).$$

Note that for $\dim \mathfrak{s} = d$, the identification $\mathfrak{o}_{(X_I^{\text{reg}})^{\leq \mathfrak{s}}} \otimes \mathfrak{o}_{\mathfrak{s}} = \mathfrak{o}_{X_I^{\text{reg}}}$ gives a (locally constant) isomorphism $\mathfrak{o}_{\mathfrak{s}} = \mathbb{Z}$ over $(X_I^{\text{reg}})^{\leq \mathfrak{s}}$ (and similarly for $\mathfrak{s} \in \partial\mathcal{S}$ with $\dim \mathfrak{s} = d - 1$; see also Remark 4.4.1) which is used implicitly in (6.2.7). This is a chain map by Lemma 6.1.5.

Proposition 6.2.3 ($C_{\text{vir}}^\bullet(-, \mathcal{S}; A) \rightarrow C_{\text{vir}}^\bullet(-; A)$ is a quasi-isomorphism) Suppose $(X, \partial X) \rightarrow (\mathcal{S}, \partial\mathcal{S})$ is equipped with a finite locally orientable implicit atlas with boundary and cell-like stratification A . Then (6.2.6) is a quasi-isomorphism, and the following diagram of sheaves on X commutes:

$$\begin{array}{ccccc}
 H_{\text{vir}}^0(-, \mathcal{S}; A) & \xrightarrow{(6.2.6)} & H_{\text{vir}}^0(-; A) & \xrightarrow{\text{Prop 4.3.3}} & \mathfrak{o}_X \\
 \downarrow & & & & \downarrow \\
 H_{\text{vir}}^0(- \cap X^{\leq \mathfrak{s}}, \mathcal{S}^{\leq \mathfrak{s}}; A) \otimes \mathfrak{o}_{\mathfrak{s}} & \xrightarrow{(6.2.6)} & H_{\text{vir}}^0(- \cap X^{\leq \mathfrak{s}}; A) \otimes \mathfrak{o}_{\mathfrak{s}} & \xrightarrow{\text{Prop 4.3.3}} & \mathfrak{o}_{X^{\leq \mathfrak{s}}} \otimes \mathfrak{o}_{\mathfrak{s}} \\
 \uparrow & & & & \uparrow \\
 H_{\text{vir}}^0(- \cap X^{\leq \mathfrak{s}} \text{ rel } \partial; A) \otimes \mathfrak{o}_{\mathfrak{s}} & \xrightarrow{\text{Prop 4.3.3}} & & & \mathfrak{o}_{X^{\leq \mathfrak{s}} \text{ rel } \partial} \otimes \mathfrak{o}_{\mathfrak{s}}
 \end{array}$$

Note that $C_{\text{vir}}^\bullet(K \cap X^{\leq s} \text{ rel } \partial; A) \rightarrow C_{\text{vir}}^\bullet(K \cap X^{\leq s}, \mathcal{S}^{\leq s}; A)$ is an isomorphism for $K \subseteq X^s$.

Proof Filter $C_{\text{vir}}^\bullet(-, \mathcal{S}; A)$ by $\dim \mathfrak{s}$; the associated graded of this filtration is the direct sum of $C_{\text{vir}}^\bullet(- \cap X^{\leq s} \text{ rel } \partial; A)$. Each these is a homotopy \mathcal{K} -sheaf on X (pushforward from $X^{\leq s}$ to X preserves homotopy \mathcal{K} -sheaves by Definition A.3.4), and hence $C_{\text{vir}}^\bullet(-, \mathcal{S}; A)$ is also a homotopy \mathcal{K} -sheaf (Lemma A.2.11).

Now the map (6.2.6) is a map of homotopy \mathcal{K} -sheaves, so to check that it is a quasi-isomorphism, it suffices to check it is a quasi-isomorphism on stalks (Corollary A.4.19). For this, we may use the argument from Example 6.2.1 (adapted to the case with boundary) along with Lemma 4.3.2. Moreover, this local construction also gives us the desired commutativity of the diagram of sheaves on X . □

6.3 Product implicit atlas

Definition 6.3.1 (Product implicit atlas) Let X_1 and X_2 be spaces with implicit atlases A_1 and A_2 respectively. The *product implicit atlas* $A_1 \sqcup A_2$ on $X_1 \times X_2$ is defined by setting

$$(X_1 \times X_2)_{I_1 \sqcup I_2} := (X_1)_{I_1} \times (X_2)_{I_2} \quad \text{and} \quad (X_1 \times X_2)_{I_1 \sqcup I_2}^{\text{reg}} := (X_1)_{I_1}^{\text{reg}} \times (X_2)_{I_2}^{\text{reg}},$$

with the rest of the data extended in the obvious manner. Of course, this extends naturally to the setting of implicit atlases with boundary and cell-like stratification (given X_1 and X_2 stratified by \mathcal{S}_1 and \mathcal{S}_2 respectively, their product $X_1 \times X_2$ is stratified by $\mathcal{S}_1 \times \mathcal{S}_2$, with $\partial(\mathcal{S}_1 \times \mathcal{S}_2) := (\partial\mathcal{S}_1 \times \mathcal{S}_2) \cup (\mathcal{S}_1 \times \partial\mathcal{S}_2)$).

Definition 6.3.2 Let X_1 and X_2 be spaces with finite implicit atlases A_1 and A_2 ; equip $X_1 \times X_2$ with the product implicit atlas $A_1 \sqcup A_2$. Let us define a canonical map

$$(6.3.1) \quad C_{\text{vir}}^\bullet(X_1 \text{ rel } \partial; A_1) \otimes C_{\text{vir}}^\bullet(X_2 \text{ rel } \partial; A_2) \rightarrow C_{\text{vir}}^\bullet(X_1 \times X_2 \text{ rel } \partial; A_1 \sqcup A_2)$$

which is compatible with the maps (4.2.12) and (4.2.14) and is associative.

There are isomorphisms $(X_1)_{I_1, J_1, A_1} \times (X_2)_{I_2, J_2, A_2} \rightarrow (X_1 \times X_2)_{I_1 \sqcup I_2, J_1 \sqcup J_2, A_1 \sqcup A_2}$ which are compatible with the maps (4.2.1). This induces maps

$$(6.3.2) \quad C_{\text{vir}}^\bullet(X_1 \text{ rel } \partial; A_1)_{I_1, J_1} \otimes C_{\text{vir}}^\bullet(X_2 \text{ rel } \partial; A_2)_{I_2, J_2} \\ \rightarrow C_{\text{vir}}^\bullet(X_1 \times X_2 \text{ rel } \partial; A_1 \sqcup A_2)_{I_1 \sqcup I_2, J_1 \sqcup J_2}.$$

These maps are compatible with the maps (4.2.13) (note that this compatibility uses the commutativity of (4.2.2)). It follows using Definition A.7.5 that they induce the desired map (6.3.1).

Note also that (6.3.1) extends to a collection of maps

$$(6.3.3) \quad C_{\text{vir}}^\bullet(K_1 \text{ rel } \partial; A_1) \otimes C_{\text{vir}}^\bullet(K_2 \text{ rel } \partial; A_2) \rightarrow C_{\text{vir}}^\bullet(K_1 \times K_2 \text{ rel } \partial; A_1 \sqcup A_2)$$

for compact $K_1 \subseteq X_1$ and $K_2 \subseteq X_2$, which are compatible with restriction, thus inducing a map of sheaves $p_1^* \mathfrak{o}_{X_1} \otimes p_2^* \mathfrak{o}_{X_2} \rightarrow \mathfrak{o}_{X_1 \times X_2}$. It can be checked that this is the tautological such map by checking locally (ie for K_1 and K_2 single points) using Lemma 4.3.2.

7 Floer-type homology theories

In this section, we define Floer-type homology groups from a collection of “flow spaces” (equipped with appropriately compatible implicit atlases) such as arise in a Morse-type setup. The necessary VFC machinery has already been set up in Section 4 and Section 6. The main task in this section is to correctly organize everything together algebraically.

Convention 7.0.1 In this section, we work over a fixed ground ring R , and everything takes place in the category of R -modules unless stated otherwise. We restrict to implicit atlases A for which $\#\Gamma_\alpha$ is invertible in R for all $\alpha \in A$.

The main object of study is a *flow category diagram* \mathcal{X}/Z_\bullet , where Z_\bullet is a semisimplicial set. Roughly speaking, this consists of a set of generators \mathcal{P}_z for every vertex $z \in Z_0$, along with a collection of spaces $\mathcal{X}(\sigma, p, q)$, which are to be thought of as the spaces of flows from $p \in \mathcal{P}_{z_0}$ to $q \in \mathcal{P}_{z_n}$ over $\sigma \in Z_n$ spanning vertices z_0, \dots, z_n . Given a flow category diagram \mathcal{X}/Z_\bullet , our goal is to construct:

- (i) For every $\sigma \in Z_0$, a boundary map $R[\mathcal{P}_{z_0}] \rightarrow R[\mathcal{P}_{z_0}]$.
- (ii) For every $\sigma \in Z_1$, a chain map $R[\mathcal{P}_{z_0}] \rightarrow R[\mathcal{P}_{z_1}]$.
- (iii) For every $\sigma \in Z_2$, a chain homotopy between the two maps $R[\mathcal{P}_{z_0}] \rightarrow R[\mathcal{P}_{z_2}]$.
- (iv) For every $\sigma \in Z_3, \dots$

We will refer to such data as a *diagram* $\mathbb{H}: Z_\bullet \rightarrow N_{\text{dg}}(\text{Ch}_R)$ (see Definition 7.6.5).

Indeed, when the flow spaces $\mathcal{X}(\sigma, p, q)$ are compact oriented manifolds with corners of dimension $\text{gr}(q) - \text{gr}(p) + \dim \sigma - 1$ in a compatible manner (that is, \mathcal{X}/Z_\bullet is a *Morse–Smale flow category diagram*), one may obtain such a diagram \mathbb{H} by counting the 0-dimensional flow spaces (one may see that the maps satisfy the required identities by considering the boundary of the 1-dimensional flow spaces).

Our goal is to generalize this construction to the setting where the spaces \mathcal{X} are equipped with compatible implicit atlases. In this generalization, the diagram $\mathbb{H}: Z_\bullet \rightarrow N_{\text{dg}}(\text{Ch}_R)$

is not determined uniquely. Rather, its construction depends on making a certain set of “coherent choices” (of virtual fundamental cycles) for the spaces \mathcal{X} . Hence, the main steps we must take are: formulating precisely what “coherent choices” mean, proving such choices always exist, and proving that the diagram is (in a suitable sense) independent of the choices. Let us now comment briefly on the latter two steps.

Given a flow category diagram \mathcal{X}/Z_\bullet with an implicit atlas, we encode the “space” of coherent choices via a map $\pi: \tilde{Z}_\bullet \rightarrow Z_\bullet$. Namely, $\pi: \tilde{Z}_\bullet \rightarrow Z_\bullet$ is defined by the property that *giving a section $s: Z_\bullet \rightarrow \tilde{Z}_\bullet$ of π is the same as making coherent choices over all of Z_\bullet* . Now, the statement that coherent choices give rise to a diagram $\mathbb{H}: Z_\bullet \rightarrow \text{N}_{\text{dg}}(\text{Ch}_R)$ translates into a canonical diagram $\tilde{\mathbb{H}}: \tilde{Z}_\bullet \rightarrow \text{N}_{\text{dg}}(\text{Ch}_R)$ (defined essentially by the property that the set of coherent choices over Z_\bullet corresponding to a section $s: Z_\bullet \rightarrow \tilde{Z}_\bullet$ gives rise to the diagram $\mathbb{H} := \tilde{\mathbb{H}} \circ s$). Thus, we have constructed

$$(7.0.1) \quad \begin{array}{ccc} \tilde{Z}_\bullet & \xrightarrow{\tilde{\mathbb{H}}} & \text{N}_{\text{dg}}(\text{Ch}_R) \\ \pi \downarrow & & \\ Z_\bullet & & \end{array}$$

Now, the key result we prove is that $\pi: \tilde{Z}_\bullet \rightarrow Z_\bullet$ is a *trivial Kan fibration* (think: “is a bundle with contractible fibers”). From this, we obtain (mostly formally) that coherent choices exist and that the resulting diagram is (up to quasi-isomorphism) independent of the choice (both are incarnations of the fact that “the space of sections of a trivial Kan fibration is contractible”).

Remark 7.0.2 (Restricting to the 2–skeleton of Z_\bullet .) If one is satisfied with working in the homotopy category (ie constructing a diagram $Z_\bullet \rightarrow H^0(\text{Ch}_R)$), then one needs only the 2–skeleton of Z_\bullet . On the other hand, there is little simplification to be gained by using 2–truncated semisimplicial sets instead of semisimplicial sets. Moreover, the “higher homotopies” which are kept track of in $\text{N}_{\text{dg}}(\text{Ch}_R)$ are known to contain interesting information in certain settings (for example, they can be used to obstruct isotopies between symplectic embeddings, as in Floer, Hofer and Wysocki [26]).

7.1 Sets of generators, triples (σ, p, q) , and F-modules

Definition 7.1.1 (Simplicial set and semisimplicial set) Let Δ be the category of finite nonempty totally ordered sets with morphisms weakly order-preserving maps. Let Δ_{inj} be the subcategory of injective morphisms. A *simplicial set* Z_\bullet is a functor $Z: \Delta^{\text{op}} \rightarrow \text{Set}$, and a *semisimplicial set* Z_\bullet is a functor $Z: \Delta_{\text{inj}}^{\text{op}} \rightarrow \text{Set}$. In both cases, we write Z_n for $Z(\{0, \dots, n\})$. For $\sigma \in Z_n$ and $0 \leq j_0 < \dots < j_m \leq n$, we denote by $\sigma[[j_0 \dots j_m]]$ the image of σ under the map $Z_n \rightarrow Z_m$ induced by the map

$\{0, \dots, m\} \rightarrow \{0, \dots, n\}$ given by $i \mapsto j_i$ ($\sigma|[j_0 \dots j_m]$ is called a facet of σ); if $m = 0$ we also write $\sigma(j)$ for $\sigma|[j]$.

Definition 7.1.2 (Set of generators) Let Z be a set. A set of generators \mathcal{P}/Z is a collection of sets $\{\mathcal{P}_z\}_{z \in Z}$, each equipped with a grading $\text{gr}: \mathcal{P}_z \rightarrow \mathbb{Z}$ and an “action” $a: \mathcal{P}_z \rightarrow \mathbb{R}$. Given a set of generators \mathcal{P}/Z and a map $f: Y \rightarrow Z$, we can form the pullback set of generators $f^*\mathcal{P}/Y$ defined by $(f^*\mathcal{P})_y := \mathcal{P}_{f(y)}$.

Definition 7.1.3 (Triples (σ, p, q)) Let Z_\bullet be a semisimplicial set, and let \mathcal{P}/Z_0 be a set of generators. The notation (σ, p, q) always means a triple where $\sigma \in Z_n$ is an n -simplex and $(p, q) \in \mathcal{P}_{\sigma(0)} \times \mathcal{P}_{\sigma(n)}$, where either $\dim \sigma > 0$ or $a(p) < a(q)$. We say that $(\sigma', p', q') < (\sigma, p, q)$ (“strictly precedes”) if and only if one of the following two conditions holds:

- (i) $\sigma' \subsetneq \sigma$ (ie σ' is a facet of positive codimension of σ).
- (ii) $\sigma' = \sigma$ and $a(p) \leq a(p')$ and $a(q') \leq a(q)$ with at least one inequality being strict.

It is easy to see that $<$ is a partial order.

Definition 7.1.4 (F–module) Let Z_\bullet be a semisimplicial set, and let \mathcal{P}/Z_0 be a set of generators. Let \mathcal{C}^\otimes be a monoidal category with an initial object $0 \in \mathcal{C}$ such that $X \otimes 0 = 0 = 0 \otimes X$ for all $X \in \mathcal{C}$.

An $F(\mathcal{P}/Z_\bullet)$ –module \mathcal{W} (often abbreviated “F–module”) in \mathcal{C} is a collection of objects $\mathcal{W} = \{\mathcal{W}(\sigma, p, q) \in \mathcal{C}\}_{(\sigma, p, q)}$ equipped with product maps

$$(7.1.1) \quad \mathcal{W}(\sigma|[0 \dots k], p, q) \otimes \mathcal{W}(\sigma|[k \dots n], q, r) \rightarrow \mathcal{W}(\sigma, p, r) \quad \text{for } 0 \leq k \leq n$$

and face maps

$$(7.1.2) \quad \mathcal{W}(\sigma|[0 \dots \hat{k} \dots n], p, q) \rightarrow \mathcal{W}(\sigma, p, q) \quad \text{for } 0 < k < n$$

which are compatible in a sense we will now describe. Note that for both (7.1.1) and (7.1.2), the triples indexing the domain strictly precede the triple indexing the target (because we always restrict to triples (σ, p, q) with $\dim \sigma > 0$ or $a(p) < a(q)$).

Now given any $\sigma \in Z_\bullet$ spanning vertices $0, \dots, n$ and a choice of

$$(7.1.3) \quad 0 = j_0 < \dots < j_\ell = n,$$

$$(7.1.4) \quad 0 = a_0 \leq \dots \leq a_m = \ell,$$

$$(7.1.5) \quad p_i \in \mathcal{P}_{\sigma(j_{a_i})} \quad \text{for } 0 \leq i \leq m,$$

we can apply the product map $m - 1$ times and the face map $n - \ell$ times in some order to obtain a map

$$(7.1.6) \quad \mathcal{W}(\sigma[[j_{a_0} \dots j_{a_1}], p_0, p_1]) \otimes \dots \otimes \mathcal{W}(\sigma[[j_{a_{m-1}} \dots j_{a_m}], p_{m-1}, p_m]) \rightarrow \mathcal{W}(\sigma, p_0, p_m).$$

We say that the product/face maps are *compatible* if and only if this map is independent of the order in which they are applied (this reduces to three basic commutation identities).

Given an $F(\mathcal{P}/Z_\bullet)$ -module \mathcal{W} and a map $f: Y_\bullet \rightarrow Z_\bullet$, we can form the pullback $f^*\mathcal{W}$, which is an $F(f^*\mathcal{P}/Y_\bullet)$ -module defined by $(f^*\mathcal{W})(\sigma, p, q) := \mathcal{W}(f(\sigma), p, q)$.

The categories \mathcal{C} relevant for this paper are:

- (i) The category of spaces with the product monoidal structure (“F-module space”).
- (ii) The category of posets with the product monoidal structure (“F-module poset”).
- (iii) The category of chain complexes with the tensor product monoidal structure (“F-module complex”).

All are in fact symmetric monoidal (noting that the relevant symmetric monoidal structure on complexes is the super tensor product).

Example 7.1.5 Let Z_\bullet be any semisimplicial set, and let us take as set of generators $\mathcal{P} := Z_0$ (ie a single generator over every vertex of Z_\bullet). We define an $F(\mathcal{P}/Z_\bullet)$ -module space \mathcal{W} by $\mathcal{W}(\sigma, p, q) := \mathcal{F}(\sigma)$, where \mathcal{F} is the space of broken Morse trajectories from Definition 10.1.4. The reader may easily verify that this forms an F-module space, with product/face maps given by (10.1.3)–(10.1.4).

Definition 7.1.6 (Support of an F-module) Let \mathcal{W} be an F-module. We define the *support* of \mathcal{W} , denoted $\text{supp } \mathcal{W}$, as the smallest collection of triples containing those for which $\mathcal{W}(\sigma, p, q) \neq 0$ that is closed under product/face operations (meaning that if the triples on the left side of (7.1.1) or (7.1.2) are in the set, then so is the triple on the right). Equivalently, $(\sigma, p, q) \in \text{supp } \mathcal{W}$ if and only if there is some choice of (7.1.3)–(7.1.5) for which $p_0 = p$ and $p_m = q$ and for which every factor on the left-hand side of (7.1.6) is $\neq 0$.

Definition 7.1.7 (Strata of an F-module) Let \mathcal{W} be an F-module. We let $\mathcal{S}_\mathcal{W}(\sigma, p, q)$ denote the set of choices of (7.1.3)–(7.1.5) for which $p_0 = p$ and $p_m = q$ and for which every factor on the left-hand side of (7.1.6) is in $\text{supp } \mathcal{W}$. We equip $\mathcal{S}_\mathcal{W}(\sigma, p, q)$ with the partial order induced by formally applying product/face maps. The reader may easily convince themselves that $\mathcal{S}_\mathcal{W}$ is itself an F-module poset.

There is an order-reversing map $\text{codim} : \mathcal{S}_{\mathcal{W}}(\sigma, p, q) \rightarrow \mathbb{Z}_{\geq 0}$ defined by $\text{codim } \mathfrak{s} := (m - 1) + (n - \ell)$ for $\mathfrak{s} \in \mathcal{S}_{\mathcal{W}}(\sigma, p, q)$. We let $\mathfrak{s}^{\text{top}} \in \mathcal{S}_{\mathcal{W}}(\sigma, p, q)$ denote the unique maximal element (the only element \mathfrak{s} with $\text{codim } \mathfrak{s} = 0$; it is given by $\ell = n$ and $m = 1$).

For $\mathfrak{s} \in \mathcal{S}_{\mathcal{W}}(\sigma, p, q)$, we let $\mathcal{W}(\sigma, p, q, \mathfrak{s})$ denote the left-hand side of (7.1.6). The *boundary inclusion map*

$$(7.1.7) \quad \text{colim}_{\mathfrak{s} \in \partial \mathcal{S}_{\mathcal{W}}(\sigma, p, q)} \mathcal{W}(\sigma, p, q, \mathfrak{s}) \rightarrow \mathcal{W}(\sigma, p, q)$$

(where $\partial \mathcal{S}_{\mathcal{W}}(\sigma, p, q)$ denotes $\mathcal{S}_{\mathcal{W}}(\sigma, p, q) \setminus \mathfrak{s}^{\text{top}}$) will play an important role.

7.2 Flow category diagrams and their implicit atlases

Definition 7.2.1 (Flow category diagram) Let Z_{\bullet} be a semisimplicial set. A *flow category diagram* \mathcal{X}/Z_{\bullet} (read “ \mathcal{X} over Z_{\bullet} .”) is:

- (i) A set of generators \mathcal{P}/Z_0 .
- (ii) An F -module space \mathcal{X} where each $\mathcal{X}(\sigma, p, q)$ is compact Hausdorff and each $\mathcal{S}_{\mathcal{X}}(\sigma, p, q)$ is finite.
- (iii) A stratification of each $\mathcal{X}(\sigma, p, q)$ by $\mathcal{S}_{\mathcal{X}}(\sigma, p, q)$ which is compatible with the product/face maps and so that $\mathcal{X}(\sigma, p, q, \mathfrak{s}) \rightarrow \mathcal{X}(\sigma, p, q)$ is a homeomorphism onto $\mathcal{X}(\sigma, p, q)^{\leq \mathfrak{s}}$.

The following finiteness properties must also be satisfied:

- (iv) For all σ, p , and $M < \infty$, we have

$$\#\{q : \mathcal{X}(\sigma, p, q) \neq \emptyset \text{ and } a(q) < a(p) + M\} < \infty.$$

- (v) For all σ , we have $\inf\{a(q) - a(p) : \mathcal{X}(\sigma, p, q) \neq \emptyset\} > -\infty$.

Let H be a group. An H -equivariant flow category diagram is a flow category diagram along with:

- (vi) A free action of H on \mathcal{P} .
- (vii) An action of H on \mathcal{X} (meaning compatible maps $h : \mathcal{X}(\sigma, p, q) \rightarrow \mathcal{X}(\sigma, hp, hq)$).
- (viii) Homomorphisms $\text{gr} : H \rightarrow \mathbb{Z}$ and $a : H \rightarrow \mathbb{R}$ such that $\text{gr}(hp) = \text{gr}(h) + \text{gr}(p)$ and $a(hp) = a(h) + a(p)$ for all $h \in H$ and $p \in \mathcal{P}$.

Given an H -equivariant flow category diagram \mathcal{X}/Z_{\bullet} and a map $f : Y_{\bullet} \rightarrow Z_{\bullet}$, we can form the pullback H -equivariant flow category diagram $f^*\mathcal{X}/Y_{\bullet}$.

Remark 7.2.2 (Morse–Smale flow category diagram) A Morse–Smale flow category diagram is one in which each $\mathcal{X}(\sigma, p, q)$ is a (compact) topological manifold with corners of dimension $\text{gr}(q) - \text{gr}(p) + \dim \sigma - 1$ (the corner structure being induced by the stratification by $\mathcal{S}_{\mathcal{X}}(\sigma, p, q)$).

Remark 7.2.3 (∞ –category FlowCat) Let FlowCat be the semisimplicial set which represents the functor $Z_{\bullet} \mapsto \{\text{Flow category diagrams over } Z_{\bullet}\}$. The reader familiar with ∞ –categories may wish to think of FlowCat as an ∞ –category of flow categories (though only in a vague sense, since we have not given it the structure of a simplicial set, nor have we verified the weak Kan condition). All of the constructions in this section involving flow category diagrams over a semisimplicial set Z_{\bullet} are compatible with pullback, and thus can be equivalently thought of as “universal” constructions over $\text{FlowCat}^{\text{IA}}$ (which represents the functor of flow category diagrams equipped with implicit atlases).

Definition 7.2.4 (Implicit atlas on flow category diagram) Let \mathcal{X}/Z_{\bullet} be a flow category diagram. An implicit atlas \mathcal{A} on \mathcal{X}/Z_{\bullet} consists of the following data. We give index sets $\bar{\mathcal{A}}(\sigma, p, q)$, and we define²¹

$$(7.2.1) \quad \mathcal{A}(\sigma, p, q)^{\geq \mathfrak{s}} := \coprod_{\substack{0 \leq i_0 < \dots < i_m \leq n \\ (p', q') \in \mathcal{P}_{\sigma(i_0)} \times \mathcal{P}_{\sigma(i_m)} \\ \exists t \in \mathcal{S}_{\mathcal{X}}(\sigma, p, q)^{\geq \mathfrak{s}} \text{ containing } ([i_0 \dots i_m], p', q')}} \bar{\mathcal{A}}(\sigma \llbracket [i_0 \dots i_m], p', q' \rrbracket).$$

We explain the notation: recall that $\mathcal{S}_{\mathcal{X}}(\sigma, p, q)$ parametrizes the “possible left-hand sides” of (7.1.6); the coproduct is over all $([i_0 \dots i_m], p', q')$ which appear as a factor in some $t \in \mathcal{S}_{\mathcal{X}}(\sigma, p, q)$ with $\mathfrak{s} \preceq t$.

For all (σ, p, q) and $\mathfrak{s} \in \mathcal{S}_{\mathcal{X}}(\sigma, p, q)$, we give an implicit atlas with boundary with cell-like stratification $\mathcal{A}(\sigma, p, q)^{\geq \mathfrak{s}}$ on $\mathcal{X}(\sigma, p, q)^{\leq \mathfrak{s}}$ (stratified by $\mathcal{S}_{\mathcal{X}}(\sigma, p, q)^{\leq \mathfrak{s}}$, of virtual dimension $\text{gr}(q) - \text{gr}(p) + \dim \sigma - 1 - \text{codim } \mathfrak{s}$), for which the stratification conforms to the following local model. Given $\mathfrak{s}' \preceq \mathfrak{s}$, let $G = G(\mathfrak{s}', \mathfrak{s})$ denote the set of possible product/face operations which may be applied to \mathfrak{s}' for which the result is still $\preceq \mathfrak{s}$. There is a tautological isomorphism of posets $2^G \rightarrow \mathcal{S}_{\mathcal{X}}(\sigma, p, q)^{\mathfrak{s}' \preceq \cdot \preceq \mathfrak{s}}$ sending a given set of product/face operations to the result of applying them to \mathfrak{s}' . Now the local model for the stratification on (regular thickened moduli spaces of) $\mathcal{X}(\sigma, p, q)^{\leq \mathfrak{s}}$ near a point of type \mathfrak{s}' is given by $\mathbb{R}_{\geq 0}^G \times \mathbb{R}^N$, stratified in the obvious way by $\mathcal{S}_{\mathcal{X}}(\sigma, p, q)^{\mathfrak{s}' \preceq \cdot \preceq \mathfrak{s}}$. Clearly this stratification is cell-like. Moreover, the normal bundle to the \mathfrak{s}' stratum is canonically identified with $\mathfrak{o}_{\mathbb{R}}^{\otimes G(\mathfrak{s}', \mathfrak{s})}$, and the implicit atlas should use this as the orientation data.

²¹Warning: a particular set $\bar{\mathcal{A}}(\sigma', p', q')$ may appear many times on the right-hand side.

In addition, we give compatible identifications between these atlases as follows:

- (i) Let $\mathfrak{s} \preceq \mathfrak{t} \in \mathcal{S}_{\mathcal{X}}(\sigma, p, q)$. Then by definition

$$\mathcal{A}(\sigma, p, q)^{\geq \mathfrak{t}} \subseteq \mathcal{A}(\sigma, p, q)^{\geq \mathfrak{s}}.$$

Both are implicit atlases on $\mathcal{X}(\sigma, p, q)^{\leq \mathfrak{s}}$ (the former by restriction to this sub-stratum of $\mathcal{X}(\sigma, p, q)^{\leq \mathfrak{t}}$), and we identify the former with the subatlas of the latter corresponding to this tautological inclusion of index sets.

- (ii) Let $\mathfrak{s} \in \mathcal{S}_{\mathcal{X}}(\sigma|[0 \dots \hat{k} \dots n], p, q) \subseteq \mathcal{S}_{\mathcal{X}}(\sigma, p, q)$. Then by definition

$$\mathcal{A}(\sigma|[0 \dots \hat{k} \dots n], p, q)^{\geq \mathfrak{s}} \subseteq \mathcal{A}(\sigma, p, q)^{\geq \mathfrak{s}}.$$

Both are implicit atlases on

$$\mathcal{X}(\sigma|[0 \dots \hat{k} \dots n], p, q)^{\leq \mathfrak{s}} = \mathcal{X}(\sigma, p, q)^{\leq \mathfrak{s}},$$

and we identify the former with the subatlas of the latter corresponding to this tautological inclusion of index sets.

- (iii) Let $\mathfrak{s}_1 \times \mathfrak{s}_2 \in \mathcal{S}_{\mathcal{X}}(\sigma|[0 \dots k], p, q) \times \mathcal{S}_{\mathcal{X}}(\sigma|[k \dots n], q, r) \subseteq \mathcal{S}_{\mathcal{X}}(\sigma, p, r)$. Then by definition

$$\mathcal{A}(\sigma|[0 \dots k], p, q)^{\geq \mathfrak{s}_1} \sqcup \mathcal{A}(\sigma|[k \dots n], q, r)^{\geq \mathfrak{s}_2} \subseteq \mathcal{A}(\sigma, p, r)^{\geq (\mathfrak{s}_1 \times \mathfrak{s}_2)}$$

Both are implicit atlases on

$$\mathcal{X}(\sigma|[0 \dots k], p, q)^{\leq \mathfrak{s}_1} \times \mathcal{X}(\sigma|[k \dots n], q, r)^{\leq \mathfrak{s}_2} = \mathcal{X}(\sigma, p, r)^{\leq (\mathfrak{s}_1 \times \mathfrak{s}_2)},$$

and we identify the former with the subatlas of the latter corresponding to this tautological inclusion of index sets.

An implicit atlas on an H -equivariant flow category \mathcal{X}/Z_{\bullet} is an implicit atlas along with a lift of the action of H to the implicit atlas structure.

Given an implicit atlas \mathcal{A} on \mathcal{X}/Z_{\bullet} and a map $f: Y_{\bullet} \rightarrow Z_{\bullet}$, we can form the pullback implicit atlas $f^*\mathcal{A}$ on $f^*\mathcal{X}/Y_{\bullet}$.

Remark 7.2.5 The above definition has been formulated to reflect the collection of atlases which is the simplest to construct, yet still sufficient to define Floer-type homology groups.

Definition 7.2.6 (Coherent orientations) Let \mathcal{X}/Z_{\bullet} be a flow category diagram with locally orientable implicit atlas \mathcal{A} . A set of *coherent orientations* ω is a choice of

global sections $\omega(\sigma, p, q) \in \check{H}^0(\mathcal{X}(\sigma, p, q); \mathfrak{o}_{\mathcal{X}(\sigma, p, q)})$ with the following property. Note that covering each of the product/face maps

$$(7.2.2) \quad \mathcal{X}(\sigma|[0 \dots k], p, q) \times \mathcal{X}(\sigma|[k \dots n], q, r) \rightarrow \partial\mathcal{X}(\sigma, p, r),$$

$$(7.2.3) \quad \mathcal{X}(\sigma|[0 \dots \hat{k} \dots n], p, q) \rightarrow \partial\mathcal{X}(\sigma, p, q)$$

is an isomorphism of orientation sheaves. We require that ω transform in the following way under this isomorphism:

$$(7.2.4) \quad \omega(\sigma|[0 \dots k], p, q) \times \omega(\sigma|[k \dots n], q, r) = (-1)^{k+\text{gr}(p)} d\omega(\sigma, p, r),$$

$$(7.2.5) \quad -\omega(\sigma|[0 \dots \hat{k} \dots n], p, q) = (-1)^{k+\text{gr}(p)} d\omega(\sigma, p, q),$$

where $d\omega \in \check{H}^0(\partial\mathcal{X}; \mathfrak{o}_{\partial\mathcal{X}})$ is the boundary orientation induced by ω . Each of

$$(7.2.6) \quad \omega(\sigma|[0 \dots k], p, p') \times \omega(\sigma|[k \dots \ell], p', p'') \times \omega(\sigma|[\ell \dots n], p'', q),$$

$$(7.2.7) \quad \omega(\sigma|[0 \dots \hat{k} \dots \ell], p, p') \times \omega(\sigma|[\ell \dots n], p', q),$$

$$(7.2.8) \quad \omega(\sigma|[0 \dots k], p, p') \times \omega(\sigma|[k \dots \hat{\ell} \dots n], p', q),$$

$$(7.2.9) \quad \omega(\sigma|[0 \dots \hat{k} \dots \hat{\ell} \dots n], p, q)$$

can be expressed in terms of $\omega(\sigma, p, q)$ in two different ways using (7.2.4)–(7.2.5). One can easily check that with the choice of signs in (7.2.4)–(7.2.5), these two expressions coincide for each of (7.2.6)–(7.2.9) (the Koszul rule of signs applies to calculating the boundary of a product of orientations, which means that this coherence condition involves $\text{vdim } \mathcal{X}(\sigma|[0 \dots k], p, p')$ for both (7.2.6) and (7.2.8)).

Coherent orientations on an H -equivariant flow category diagram with implicit atlas are coherent orientations which are invariant under the action of H .

Given coherent orientations ω on \mathcal{X}/Z_\bullet and a map $f: Y_\bullet \rightarrow Z_\bullet$, we can form the pullback coherent orientations $f^*\omega$ on $f^*\mathcal{X}/Y_\bullet$.

Remark 7.2.7 One can obtain alternative sign conventions in (7.2.4)–(7.2.5) by “twisting” $\omega(\sigma, p, q)$. For example, multiplying $\omega(\sigma, p, q)$ by (-1) flips the sign of (7.2.4), multiplying by $(-1)^{\dim \sigma}$ flips the sign of (7.2.5), and multiplying by $(-1)^{\text{gr}(p)}$ or $(-1)^{\text{gr}(q)}$ multiplies (7.2.4) by $(-1)^{\text{gr}(q)}$.

7.3 Augmented virtual cochain complexes

We would like to endow $C_{\text{vir}}^\bullet(\mathcal{X}(\sigma, p, q) \text{ rel } \partial)$ with the structure of an F -module complex. More precisely, we would like to construct product maps

$$(7.3.1) \quad C_{\text{vir}}^\bullet(\mathcal{X}(\sigma|[0 \dots k], p, q) \text{ rel } \partial) \otimes C_{\text{vir}}^\bullet(\mathcal{X}(\sigma|[k \dots n], q, r) \text{ rel } \partial) \rightarrow C_{\text{vir}}^\bullet(\partial\mathcal{X}(\sigma, p, r))$$

and face maps

$$(7.3.2) \quad C_{\text{vir}}^\bullet(\mathcal{X}(\sigma|[0 \dots \hat{k} \dots n], p, q) \text{ rel } \partial) \rightarrow C_{\text{vir}}^\bullet(\partial\mathcal{X}(\sigma, p, q))$$

(induced by the corresponding product and face maps of the spaces $\mathcal{X}(\sigma, p, q)$). However, to obtain maps (7.3.1)–(7.3.2) defined on the chain level, we must replace the virtual cochain complexes with certain *augmented virtual cochain complexes* (which are canonically quasi-isomorphic to their “non-augmented” counterparts).

In this subsection, we build the augmented virtual cochain complexes (using a homotopy colimit construction) and then we define the F –module structure on them. We also define the analogue of the map s_* for the augmented virtual cochain complexes. We remark that the proliferation of homotopy colimits could probably be abated (and, indeed, this entire subsection eliminated) at the expense of using more abstract language (specifically, working in a symmetric monoidal ∞ –category of complexes).

Definition 7.3.1 (Augmented virtual cochain complexes $C_{\text{vir}}^\bullet(-; \mathcal{A})^+$) Let \mathcal{X}/Z_\bullet be a flow category diagram with implicit atlas \mathcal{A} , where every $\mathcal{A}(\sigma, p, q)$ is finite and we have fixed fundamental cycles $[E_\alpha] \in C_\bullet(E; \alpha)$ for all $\alpha \in \bar{\mathcal{A}}(\sigma, p, q)$. We define the complexes

$$(7.3.3) \quad C_{\text{vir}}^\bullet(\mathcal{X} \text{ rel } \partial; \mathcal{A})^+(\sigma, p, q) := \text{hocolim}_{\mathfrak{s} \leq \mathfrak{t} \in \mathcal{S}_\mathcal{X}(\sigma, p, q)} C_{\text{vir}}^{\bullet - \text{codim } \mathfrak{s}}(\mathcal{X}(\sigma, p, q)^{\leq \mathfrak{s}} \text{ rel } \partial; \mathcal{A}(\sigma, p, q)^{\geq \mathfrak{t}}),$$

$$(7.3.4) \quad C_{\text{vir}}^{\bullet-1}(\partial\mathcal{X}; \mathcal{A})^+(\sigma, p, q) := \text{hocolim}_{\mathfrak{s} \leq \mathfrak{t} \in \mathcal{S}_\mathcal{X}(\sigma, p, q)} \begin{cases} C_{\text{vir}}^{\bullet-1}(\partial\mathcal{X}(\sigma, p, q); \mathcal{A}(\sigma, p, q)^{\geq \mathfrak{t}}), & \mathfrak{s} = \mathfrak{t} = \mathfrak{s}^{\text{top}}, \\ C_{\text{vir}}^{\bullet - \text{codim } \mathfrak{s}}(\mathcal{X}(\sigma, p, q)^{\leq \mathfrak{s}} \text{ rel } \partial; \mathcal{A}(\sigma, p, q)^{\geq \mathfrak{t}}), & \text{otherwise.} \end{cases}$$

The structure maps of the homotopy diagrams come from the obvious pushforward maps (for increasing \mathfrak{s}) and the maps (4.2.14) using the fixed fundamental cycles $[E_\alpha]$ (for decreasing \mathfrak{t}). These are compatible because of the commutativity of (4.2.2).

Next, let us observe that we have a natural commutative diagram:

$$(7.3.5) \quad \begin{array}{ccc} C_{\text{vir}}^{\bullet-1}(\partial\mathcal{X}(\sigma, p, q); \mathcal{A}(\sigma, p, q)^{\geq \mathfrak{s}^{\text{top}}}) & \xleftarrow{\sim} & C_{\text{vir}}^{\bullet-1}(\partial\mathcal{X}; \mathcal{A})^+(\sigma, p, q) \\ \downarrow & & \downarrow \\ C_{\text{vir}}^\bullet(\mathcal{X}(\sigma, p, q) \text{ rel } \partial; \mathcal{A}(\sigma, p, q)^{\geq \mathfrak{s}^{\text{top}}}) & \xleftarrow{\sim} & C_{\text{vir}}^\bullet(\mathcal{X} \text{ rel } \partial; \mathcal{A})^+(\sigma, p, q) \end{array}$$

The horizontal maps (inclusions of the $\mathfrak{s} = \mathfrak{t} = \mathfrak{s}^{\text{top}}$ subcomplexes) are quasi-isomorphisms by Lemma A.7.3 (which morally says that $\mathfrak{s}^{\text{top}} \in \mathcal{S}_\mathcal{X}(\sigma, p, q)$ acts as a final

object in the homotopy colimits (7.3.3)–(7.3.4), which applies because each of the structure maps from $(\mathfrak{s}, \mathfrak{t})$ to $(\mathfrak{s}, \mathfrak{t}')$ is a quasi-isomorphism.

Note that by definition, the support of $C_{\text{vir}}^\bullet(\mathcal{X} \text{ rel } \partial; \mathcal{A})^+$ and of $C_{\text{vir}}^\bullet(\partial\mathcal{X}; \mathcal{A})^+$ are contained in the support of \mathcal{X} .

Definition 7.3.2 (Product/face maps for $C_{\text{vir}}^\bullet(-; \mathcal{A})^+$) Let \mathcal{X}/Z_\bullet be a flow category diagram with implicit atlas \mathcal{A} , where every $\mathcal{A}(\sigma, p, q)$ is finite and we have fixed fundamental cycles $[E_\alpha]$. Let us now define face and product maps for $C_{\text{vir}}^\bullet(\mathcal{X} \text{ rel } \partial; \mathcal{A})^+$. These will be of degree 1 and will be equipped with a canonical factorization through $C_{\text{vir}}^{\bullet-1}(\partial\mathcal{X}; \mathcal{A})^+ \rightarrow C_{\text{vir}}^\bullet(\mathcal{X} \text{ rel } \partial; \mathcal{A})^+$. In other words, we really are going to construct maps

$$(7.3.6) \quad C_{\text{vir}}^\bullet(\mathcal{X} \text{ rel } \partial; \mathcal{A})^+(\sigma|[0 \dots k], p, q) \otimes C_{\text{vir}}^\bullet(\mathcal{X} \text{ rel } \partial; \mathcal{A})^+(\sigma|[k \dots n], q, r) \rightarrow C_{\text{vir}}^\bullet(\partial\mathcal{X}; \mathcal{A})^+(\sigma, p, r),$$

$$(7.3.7) \quad C_{\text{vir}}^\bullet(\mathcal{X} \text{ rel } \partial; \mathcal{A})^+(\sigma|[0 \dots \hat{k} \dots n], p, q) \rightarrow C_{\text{vir}}^\bullet(\partial\mathcal{X}; \mathcal{A})^+(\sigma, p, q).$$

We construct (7.3.7). The corresponding face map for $\mathcal{S}_\mathcal{X}$ is covered by a corresponding morphism of the homotopy diagrams (7.3.3)–(7.3.4) (namely (4.2.14) using the fixed fundamental cycle $[E_{\mathcal{A}(\sigma, p, q) \geq \mathfrak{t} \setminus \mathcal{A}(\sigma|[0, \dots, \hat{k}, \dots, n], p, q) \geq \mathfrak{t}}}]$). This gives rise to a corresponding map (7.3.7) on homotopy colimits.

We construct (7.3.6). We construct a morphism of homotopy diagrams (7.3.3)–(7.3.4) covering the corresponding product map for $\mathcal{S}_\mathcal{X}$. Using the product operation on homotopy diagrams (Definition A.7.5), it suffices to construct compatible maps

$$(7.3.8) \quad C_{\text{vir}}^\bullet(\mathcal{X}(\sigma|[0, \dots, k], p, q)^{\leq \mathfrak{s}_1} \text{ rel } \partial; \mathcal{A}(\sigma|[0, \dots, k], p, q)^{\geq \mathfrak{t}_1}) \otimes C_{\text{vir}}^\bullet(\mathcal{X}(\sigma|[k, \dots, n], q, r)^{\leq \mathfrak{s}_2} \text{ rel } \partial; \mathcal{A}(\sigma|[k, \dots, n], q, r)^{\geq \mathfrak{t}_2}) \rightarrow C_{\text{vir}}^\bullet(\mathcal{X}(\sigma, p, q)^{\leq \mathfrak{s}_1 \times \mathfrak{s}_2} \text{ rel } \partial; \mathcal{A}(\sigma, p, q)^{\geq \mathfrak{t}_1 \times \mathfrak{t}_2})$$

for $\mathfrak{s}_1 \preceq \mathfrak{t}_1 \in \mathcal{S}(\sigma|[0, \dots, k], p, q)$ and $\mathfrak{s}_2 \preceq \mathfrak{t}_2 \in \mathcal{S}(\sigma|[k, \dots, n], q, r)$. By definition, the subatlas

$$\mathcal{A}(\sigma|[0, \dots, k], p, q)^{\geq \mathfrak{t}_1} \sqcup \mathcal{A}(\sigma|[k, \dots, n], q, r)^{\geq \mathfrak{t}_2} \subseteq \mathcal{A}(\sigma, p, q)^{\geq \mathfrak{t}_1 \times \mathfrak{t}_2}$$

is identified with the product implicit atlas on

$$\mathcal{X}(\sigma, p, q)^{\leq \mathfrak{s}_1 \times \mathfrak{s}_2} = \mathcal{X}(\sigma|[0, \dots, k], p, q)^{\leq \mathfrak{s}_1} \times \mathcal{X}(\sigma|[k, \dots, n], q, r)^{\leq \mathfrak{s}_2}.$$

Thus the desired map is constructed in Definition 6.3.2.

Definition 7.3.3 (Complexes $C_\bullet(E; \mathcal{A})^+$) Let \mathcal{X}/Z_\bullet be a flow category diagram with implicit atlas \mathcal{A} , where every $\bar{\mathcal{A}}(\sigma, p, q)$ is finite and we have fixed fundamental cycles $[E_\alpha]$. We define

$$(7.3.9) \quad C_\bullet(E; \mathcal{A})^+(\sigma, p, q) := \operatorname{hocolim}_{\mathfrak{s} \leq t \in \mathcal{S}_X(\sigma, p, q)} C_\bullet(E; \mathcal{A}^{\geq t}(\sigma, p, q)),$$

where the maps in the homotopy diagram are $-\times [E_{A(\sigma, p, q, t) \setminus A(\sigma, p, q, t)}]$. We equip $C_\bullet(E; \mathcal{A})^+$ with product/face maps just as in Definition 7.3.2. Now there are natural maps

$$(7.3.10) \quad C_{\text{vir}}^{\text{vdim } X(\sigma, p, q) + \bullet} (X \operatorname{rel} \partial; \mathcal{A})^+(\sigma, p, q) \xrightarrow{s_*} C_{-\bullet}(E; \mathcal{A})^+(\sigma, p, q)$$

(induced by (4.2.12)), which are maps of F -modules (that is, they respect the product/face maps).

Note that by definition, the support of $C_\bullet(E; \mathcal{A})^+$ is contained in the support of \mathcal{X} .

7.4 Cofibrant F -module complexes

We introduce the notion of an F -module complex being *cofibrant*, and we introduce a cofibrant replacement functor Q for F -module complexes. The machinery we set up is used only for technical reasons in Definition 7.5.3 so that the proof of Proposition 7.5.5 works correctly.

Recall that an injection of modules with projective cokernel automatically splits.

Definition 7.4.1 (Cofibrations of complexes) We say a complex is *cofibrant* if and only if it is projective (as a module). We say a map of complexes is a *cofibration* if and only if it is injective and its cokernel is cofibrant; we use the arrow \twoheadrightarrow to indicate that a map is a cofibration. It is easy to check that a composition of cofibrations is again a cofibration.

Definition 7.4.2 (Cofibrant F -module complex) Let \mathcal{W}_\bullet be an F -module complex. We say that \mathcal{W}_\bullet is *cofibrant* if and only if for all (σ, p, q) :

- (i) $\mathcal{W}_\bullet(\sigma, p, q)$ is cofibrant.
- (ii) The map

$$(7.4.1) \quad \operatorname{colim}_{\mathfrak{s} \in \partial \mathcal{S}_W(\sigma, p, q)} \mathcal{W}_\bullet(\sigma, p, q, \mathfrak{s}) \twoheadrightarrow \mathcal{W}_\bullet(\sigma, p, q)$$

is a cofibration.

Lemma 7.4.3 *Let \mathcal{W}_\bullet be a cofibrant F -module complex. Then the map*

$$(7.4.2) \quad \operatorname{colim}_{\substack{\mathfrak{s} \in \mathcal{S}_{\mathcal{W}}(\sigma, p, q) \\ \mathfrak{s} \prec \mathfrak{s}_0}} \mathcal{W}(\sigma, p, q, \mathfrak{s}) \twoheadrightarrow \mathcal{W}(\sigma, p, q, \mathfrak{s}_0)$$

is a cofibration for all $\mathfrak{s}_0 \in \mathcal{S}(\sigma, p, q)$.

Proof Write \mathfrak{s}_0 in the form (7.1.3)–(7.1.5). Then we have

$$(7.4.3) \quad \mathcal{S}_{\mathcal{W}}(\sigma, p, q)^{\leq \mathfrak{s}_0} = \mathcal{S}_{\mathcal{W}}(\sigma[[j_{a_0} \dots j_{a_1}], p_0, p_1]) \times \dots \times \mathcal{S}_{\mathcal{W}}(\sigma[[j_{a_{m-1}} \dots j_{a_m}], p_{m-1}, p_m]).$$

Now we consider the m -cubical diagram given by $\bigotimes_{i=1}^m$ of

$$(7.4.4) \quad \operatorname{colim}_{\mathfrak{s} \in \partial \mathcal{S}_{\mathcal{W}}(\sigma[[j_{a_{i-1}} \dots j_{a_i}], p_{i-1}, p_i])} \mathcal{W}_\bullet(\sigma[[j_{a_{i-1}} \dots j_{a_i}], p_{i-1}, p_i, \mathfrak{s}]) \twoheadrightarrow \mathcal{W}_\bullet(\sigma[[j_{a_{i-1}} \dots j_{a_i}], p_{i-1}, p_i])$$

Now (7.4.2) is simply the map to the maximal vertex of the m -cube from the colimit over the m -cube minus the maximal vertex. This is a cofibration for any cubical diagram of the form $\bigotimes_{i=1}^m [A_i \twoheadrightarrow B_i]$ where each $A_i \twoheadrightarrow B_i$ is a cofibration. To see this, write $B_i = A_i \oplus P_i$, where P_i is projective, and then the map is obviously injective with cokernel $\bigotimes_{i=1}^m P_i$. A tensor product of projective modules is projective, as can be seen using either the tensor-hom adjunction or the fact that a module is projective if and only if it is a direct summand of a free module. \square

Lemma 7.4.4 *Let \mathcal{W}_\bullet be a cofibrant F -module complex. Let $\mathcal{S} \subseteq \mathcal{T} \subseteq \mathcal{S}_{\mathcal{W}}(\sigma, p, q)$ be finite downward closed subsets. Then the map*

$$(7.4.5) \quad \operatorname{colim}_{\mathfrak{s} \in \mathcal{S}} \mathcal{W}_\bullet(\sigma, p, q, \mathfrak{s}) \twoheadrightarrow \operatorname{colim}_{\mathfrak{s} \in \mathcal{T}} \mathcal{W}_\bullet(\sigma, p, q, \mathfrak{s})$$

is a cofibration. In particular, $\operatorname{colim}_{\mathfrak{s} \in \mathcal{T}} \mathcal{W}_\bullet(\sigma, p, q, \mathfrak{s})$ is projective (take $\mathcal{S} = \emptyset$ above).

Proof Let us abbreviate $A_{\mathfrak{s}} := \mathcal{W}_\bullet(\sigma, p, q, \mathfrak{s})$.

We proceed by induction on the cardinality of \mathcal{T} , the case $\mathcal{T} = \emptyset$ being clear. Using the fact that a composition of cofibrations is again a cofibration, it suffices to consider the case $(\mathcal{S}, \mathcal{T}) = (\mathcal{S} \setminus \mathfrak{s}_0, \mathcal{S})$, where $\mathfrak{s}_0 \in \mathcal{S}$ is a maximal element. Now $\operatorname{colim}_{\mathfrak{s} \in \mathcal{S}} A_{\mathfrak{s}}$ is the colimit of the diagram

$$(7.4.6) \quad \begin{array}{ccc} & & A_{\mathfrak{s}_0} \\ & \nearrow & \\ \operatorname{colim}_{\substack{\mathfrak{s} \in \mathcal{S} \\ \mathfrak{s} \prec \mathfrak{s}_0}} A_{\mathfrak{s}} & & \\ & \searrow & \\ & & \operatorname{colim}_{\mathfrak{s} \in \mathcal{S} \setminus \mathfrak{s}_0} A_{\mathfrak{s}} \end{array}$$

where the top arrow is a cofibration by Lemma 7.4.3 and the bottom arrow is a cofibration by the induction hypothesis. It follows that $\text{colim}_{s \in \mathcal{S} \setminus s_0} A_s \rightarrow \text{colim}_{s \in \mathcal{S}} A_s$ is injective, with cokernel isomorphic to the cokernel of the top map above. Hence it is a cofibration as needed. \square

Definition 7.4.5 (Cofibrant replacement functor Q) Let \mathcal{W}_\bullet be an F -module complex. Suppose that \mathcal{W}_\bullet satisfies the following properties:

- (i) Each $\mathcal{W}_\bullet(\sigma, p, q)$ is cofibrant.
- (ii) Each $\mathcal{S}_{\mathcal{W}}(\sigma, p, q)$ is finite.

In this case, we define (functorially) an F -module complex $Q\mathcal{W}_\bullet$ (called the *cofibrant replacement*) with the following properties:

- (i) $Q\mathcal{W}_\bullet$ is cofibrant.
- (ii) $\text{supp } Q\mathcal{W}_\bullet = \text{supp } \mathcal{W}_\bullet$.
- (iii) There is a (functorial) surjective quasi-isomorphism $Q\mathcal{W}_\bullet \xrightarrow{\sim} \mathcal{W}_\bullet$ (compatible with product/face maps).

We define $Q\mathcal{W}_\bullet(\sigma, p, q)$ and the product/face maps with target (σ, p, q) by induction on the set of triples (σ, p, q) , equipped with the partial order $\preceq_{\mathcal{W}}$ in which $(\sigma', p', q') \preceq_{\mathcal{W}} (\sigma, p, q)$ if and only if (σ', p', q') appears as a “factor” of some element of $\mathcal{S}_{\mathcal{W}}(\sigma, p, q)$. This partial order is *well-founded* (ie there is no strictly decreasing infinite sequence) since each $\mathcal{S}_{\mathcal{W}}(\sigma, p, q)$ is finite, and thus it is valid for induction. The inductive step for (σ, p, q) works as follows. By the induction hypothesis, we have defined $Q\mathcal{W}_\bullet(\sigma, p, q, s)$ for all $s \in \partial\mathcal{S}_{\mathcal{W}}(\sigma, p, q)$. Hence it suffices to construct (functorially) $Q\mathcal{W}_\bullet(\sigma, p, q)$ fitting into the following commutative diagram:

$$(7.4.7) \quad \begin{array}{ccc} \text{colim}_{s \in \partial\mathcal{S}_{\mathcal{W}}(\sigma, p, q)} Q\mathcal{W}_\bullet(\sigma, p, q, s) & \xrightarrow{\text{product/face}} & ?\exists Q\mathcal{W}_\bullet(\sigma, p, q) \\ \downarrow & & \downarrow \sim \\ \text{colim}_{s \in \partial\mathcal{S}_{\mathcal{W}}(\sigma, p, q)} \mathcal{W}_\bullet(\sigma, p, q, s) & \xrightarrow{\text{product/face}} & \mathcal{W}_\bullet(\sigma, p, q) \end{array}$$

Now we define $Q\mathcal{W}_\bullet(\sigma, p, q)$ to be the mapping cylinder of the diagonal composition. Since the domain and codomain of this map are both cofibrant (for the domain, use Lemma 7.4.4, which applies since $\mathcal{S}_{\mathcal{W}}(\sigma, p, q)$ has been assumed to be finite), it follows that the top map is a cofibration and that $Q\mathcal{W}_\bullet(\sigma, p, q)$ is cofibrant. The product/face maps with target $Q\mathcal{W}_\bullet(\sigma, p, q)$ are defined using the top horizontal map; they are compatible since by construction they factor through the colimit in the upper left corner.

Certainly $\text{supp } Q\mathcal{W}_\bullet \supseteq \text{supp } \mathcal{W}_\bullet$. Conversely, suppose $Q\mathcal{W}_\bullet(\sigma, p, q)$ is nonzero. Then either $\mathcal{W}_\bullet(\sigma, p, q) \neq 0$ (so $(\sigma, p, q) \in \text{supp } \mathcal{W}_\bullet$), or $Q\mathcal{W}_\bullet(\sigma, p, q, \mathfrak{s}) \neq 0$ for some $\mathfrak{s} \in \partial\mathcal{S}(\sigma, p, q)$ (so by induction, the triples comprising \mathfrak{s} are in $\text{supp } \mathcal{W}_\bullet$, and hence so is (σ, p, q)).

Lemma 7.4.6 *Let $\{C_\bullet^k\}_{k=1}^n$ be cofibrant complexes over \mathbb{Z} such that $H_i C_\bullet^k = 0$ for $i < 0$. Then for $i < 0$ we have $H_i(\bigotimes_{k=1}^n C_\bullet^k) = 0$.*

Proof Since the tensor product of projective modules is projective, we may use induction to reduce to the case $k = 2$.

Thus, we have two cofibrant complexes A_\bullet and B_\bullet with $H_i A_\bullet = H_i B_\bullet = 0$ for $i < 0$ and we would like to conclude that $H_i(A_\bullet \otimes B_\bullet) = 0$ for $i < 0$. This follows from the Künneth theorem, specifically in the form of [67, page 301, Theorem 9.16] or [70, page 679, Theorem 10.81] (both of which apply because $\text{Tor}_{\mathbb{Z}}^i(\cdot, \cdot) = 0$ for $i > 1$).²² \square

Lemma 7.4.7 *Let \mathcal{W}_\bullet be a cofibrant F -module complex over \mathbb{Z} such that $H_i \mathcal{W}_\bullet = 0$ for $i < 0$. Then for any finite downward closed subset $\mathcal{S} \subseteq \mathcal{S}_{\mathcal{W}}(\sigma, p, q)$, the map*

$$(7.4.8) \quad H_i \bigoplus_{\mathfrak{s} \in \mathcal{S}} \mathcal{W}_\bullet(\sigma, p, q, \mathfrak{s}) \twoheadrightarrow H_i \text{colim}_{\mathfrak{s} \in \mathcal{S}} \mathcal{W}_\bullet(\sigma, p, q, \mathfrak{s})$$

is surjective for $i \leq 0$. Moreover, both sides vanish for $i < 0$.

Proof Let us abbreviate $A_{\mathfrak{s}} := \mathcal{W}_\bullet(\sigma, p, q, \mathfrak{s})$. Since $A_{\mathfrak{s}}$ is a tensor product of various $\mathcal{W}_\bullet(\sigma', p', q')$, it follows (using Lemma 7.4.6) that $H_i A_{\mathfrak{s}}$ for $i < 0$, so the left-hand side of (7.4.8) vanishes for $i < 0$. Hence it remains just to show that (7.4.8) is surjective for $i \leq 0$.

We proceed by induction on the cardinality of \mathcal{S} , the case $\mathcal{S} = \emptyset$ being clear. Let $\mathfrak{s}_0 \in \mathcal{S}$ be any maximal element. Now $\text{colim}_{\mathfrak{s} \in \mathcal{S}} A_{\mathfrak{s}}$ is the colimit of (7.4.6); it follows that we have the exact sequence

$$(7.4.9) \quad 0 \rightarrow \text{colim}_{\substack{\mathfrak{s} \in \mathcal{S} \\ \mathfrak{s} < \mathfrak{s}_0}} A_{\mathfrak{s}} \rightarrow A_{\mathfrak{s}_0} \oplus \text{colim}_{\mathfrak{s} \in \mathcal{S} \setminus \mathfrak{s}_0} A_{\mathfrak{s}} \rightarrow \text{colim}_{\mathfrak{s} \in \mathcal{S}} A_{\mathfrak{s}} \rightarrow 0.$$

By the induction hypothesis, it suffices to show that the second map above is surjective on H_i for $i \leq 0$. By the long exact sequence induced by (7.4.9), it suffices to show that $H_i \text{colim}_{\substack{\mathfrak{s} \in \mathcal{S} \\ \mathfrak{s} < \mathfrak{s}_0}} A_{\mathfrak{s}} = 0$ for $i < 0$. Now this follows from the induction hypothesis. \square

²²If we were assuming A_\bullet and B_\bullet to be bounded below, then we could apply the Künneth spectral sequence $\bigoplus_{i+j=q} \text{Tor}_p(H_i A_\bullet, H_j B_\bullet) \Rightarrow H_{p+q}(A_\bullet \otimes B_\bullet)$ (see [70, page 686, Theorem 10.90]) to reach the desired conclusion without any assumptions on the ground ring.

Lemma 7.4.8 (Lifting cycles along a fibration) *Let $\tilde{A}^\bullet \rightarrow A^\bullet$ be surjective. Fix a cycle $a \in A^0$ and a homology class $\tilde{a} \in H^0 \tilde{A}^\bullet$ whose images in $H^0 A^\bullet$ coincide. Then there exists a cycle $\tilde{a} \in \tilde{A}^0$ which maps to a and which represents \tilde{a} .*

Proof Pick any cycle $\tilde{a}' \in \tilde{A}^0$ representing \tilde{a} . Then $\tilde{a}' - a \in A^0$ is a boundary db , and lifting $b \in A^{-1}$ to $\tilde{b} \in \tilde{A}^{-1}$, we let $\tilde{a} := \tilde{a}' - d\tilde{b}$. □

Lemma 7.4.9 (Representing homology classes in mapping cones) *Fix $f: A^\bullet \rightarrow B^\bullet$ and $\omega \in H^0[A^\bullet \rightarrow B^{\bullet-1}]$. Let $a \in A^0$ be a cycle representing $\delta\omega \in H^0(A^\bullet)$. Then there exists $b \in B^{-1}$ with $db = f(a)$ such that $a \oplus b$ represents ω .*

Proof This is a special case of Lemma 7.4.8 for the surjection $[A^\bullet \rightarrow B^{\bullet-1}] \twoheadrightarrow A^\bullet$. □

Lemma 7.4.10 (Universal coefficient theorem) *Let A_\bullet be a cofibrant complex over \mathbb{Z} . Then there is a natural short exact sequence*

$$(7.4.10) \quad 0 \rightarrow \text{Ext}^1(H_{i-1}A_\bullet, \mathbb{Z}) \rightarrow H^i \text{Hom}(A_\bullet, \mathbb{Z}) \rightarrow \text{Hom}(H_i A_\bullet, \mathbb{Z}) \rightarrow 0.$$

Proof Well-known. □

Lemma 7.4.11 (Extending cocycles along a cofibration) *Suppose we have the following commuting diagrams (solid arrows) of complexes over \mathbb{Z} and their homology (where \mathbb{Z} is concentrated in degree zero):*

$$(7.4.11) \quad \begin{array}{ccc} A_\bullet & \xrightarrow{\quad} & \tilde{A}_\bullet \\ \downarrow & \swarrow \text{dashed} & \\ \mathbb{Z} & & \end{array} \quad \xRightarrow{H_\bullet} \quad \begin{array}{ccc} H_\bullet A_\bullet & \longrightarrow & H_\bullet \tilde{A}_\bullet \\ \downarrow & \swarrow & \\ \mathbb{Z} & & \end{array}$$

where A_\bullet and \tilde{A}_\bullet are cofibrant and $H_{-1}A_\bullet = 0$. Then there exists a dashed arrow compatible with the rest of the diagram.

Proof By Lemma 7.4.10, we have

$$(7.4.12) \quad H^0 \text{Hom}(A_\bullet, \mathbb{Z}) \xrightarrow{\sim} \text{Hom}(H_0 A_\bullet, \mathbb{Z}),$$

$$(7.4.13) \quad H^0 \text{Hom}(\tilde{A}_\bullet, \mathbb{Z}) \twoheadrightarrow \text{Hom}(H_0 \tilde{A}_\bullet, \mathbb{Z}).$$

Let $f: A_0 \rightarrow \mathbb{Z}$ denote the vertical map in the first diagram. Using the surjectivity of (7.4.13), there exists a cocycle $\tilde{f}: \tilde{A}_0 \rightarrow \mathbb{Z}$ giving the desired map on homology. Denote by $i: A_\bullet \rightarrow \tilde{A}_\bullet$ the inclusion. Now by the commutativity of the second diagram, the difference $f - \tilde{f}i$ acts as zero on $H_0 A_\bullet$. Since (7.4.12) is an isomorphism, this difference is thus a coboundary δg for some $g: A_{-1} \rightarrow \mathbb{Z}$. Now extend g to $\tilde{g}: \tilde{A}_{-1} \rightarrow \mathbb{Z}$ (using the fact that $A_{-1} \rightarrow \tilde{A}_{-1}$ splits) and let the dotted arrow be $\tilde{f} + \delta\tilde{g}$. □

7.5 Resolution $\tilde{Z}_\bullet \rightarrow Z_\bullet$

We now introduce the space $\pi: \tilde{Z}_\bullet \rightarrow Z_\bullet$ (depending on a flow category diagram \mathcal{X}/Z_\bullet equipped with an implicit atlas and coherent orientations) which one may think of as parametrizing coherent choices of virtual fundamental cycles over all of the flow spaces $\mathcal{X}(\sigma, p, q)$.

The main result is that $\pi: \tilde{Z}_\bullet \rightarrow Z_\bullet$ is a trivial Kan fibration.

Definition 7.5.1 (System of chains) Let \mathcal{W}^\bullet be an F -module complex with each $\mathcal{S}_W(\sigma, p, q)$ finite. A *system of chains* $\lambda \in \mathcal{W}^\bullet$ (of degree $d(\sigma, p, q)$) is a collection of elements $\lambda_{\sigma,p,q} \in \mathcal{W}^{d(\sigma,p,q)}(\sigma, p, q)$ satisfying $d\lambda_{\sigma,p,q} = \mu_{\sigma,p,q}$, where

$$(7.5.1) \quad \mu_{\sigma,p,r} := \sum_{k=0}^n \sum_{q \in \mathcal{P}_{\sigma(k)}} (-1)^{k+\text{gr}(p)} \lambda_{\sigma|[0\dots k],p,q} \cdot \lambda_{\sigma|[k\dots n],q,r} - \sum_{k=1}^{n-1} (-1)^{k+\text{gr}(p)} \lambda_{\sigma|[0\dots \hat{k}\dots n],p,r} \in \mathcal{W}^\bullet(\sigma, p, r)$$

(using the product/face maps on the right-hand side). Note that this sum is finite since $\mathcal{S}_W(\sigma, p, q)$ is finite. Also, note that the triples on the right-hand side all strictly precede the triple on the left-hand side. We also require that the parity (in the sense of Convention A.0.3) of $\lambda_{\sigma,p,q}$ equals $\text{gr}(q) - \text{gr}(p) + \dim \sigma - 1 \in \mathbb{Z}/2$. This ensures (via the Koszul rule of signs) that expanding $d\mu_{\sigma,p,q}$ using the identity $d\lambda_{\sigma,p,q} = \mu_{\sigma,p,q}$ yields zero (the signs work out correctly as a consequence of the discussion surrounding (7.2.6)–(7.2.9)). In practice the degrees $d(\sigma, p, q)$ are chosen so that $\mu_{\sigma,p,q}$ (as defined by (7.5.1)) is formally homogeneous of the same degree as $d\lambda_{\sigma,p,q}$.

Remark 7.5.2 There is a natural bijection between systems of chains $\lambda \in \mathcal{W}^\bullet$ and maps of F -modules $R[\mathcal{S}_W] \rightarrow \mathcal{W}^\bullet$, where $R[\mathcal{S}_W](\sigma, p, q)$ is the free R -module on $\mathcal{S}_W(\sigma, p, q)$ with differential $d\mathfrak{s} := \sum_{\text{codim}(\mathfrak{s}' \leq \mathfrak{s})=1} \mathfrak{s}'$ (with appropriate signs) and equipped with the obvious product/face maps coming from the F -module structure on \mathcal{S}_W (more intrinsically, $R[\mathcal{S}_W](\sigma, p, q)$ is the direct sum over $\mathfrak{s} \in \mathcal{S}_W(\sigma, p, q)$ of the orientation lines from Definition 7.2.4). This perspective is relevant for the key step of the proof of Proposition 7.5.5 below.

Definition 7.5.3 (Resolution $\pi: \tilde{Z}_\bullet \rightarrow Z_\bullet$) Let \mathcal{X}/Z_\bullet be an H -equivariant flow category diagram with implicit atlas \mathcal{A} and coherent orientations ω . We construct a *resolution* $\pi: \tilde{Z}_\bullet \rightarrow Z_\bullet$ (which depends on $\mathcal{X}, \mathcal{A}, \omega$) as follows. A simplex $\Delta^n \rightarrow \tilde{Z}_\bullet$ consists of the following data:

- (i) A map $f: \Delta^n \rightarrow Z_\bullet$ (where Δ^n is the semisimplicial n -simplex).

- (ii) An H -invariant finite subatlas $\mathcal{B} \subseteq f^*\mathcal{A}$ (meaning choices of finite subsets $\bar{\mathcal{B}}(\sigma, p, q) \subseteq \bar{\mathcal{A}}(f(\sigma), p, q)$).
- (iii) An H -invariant collection of fundamental cycles $[E_\alpha] \in C_\bullet(E; \alpha)$ for all $\alpha \in \bar{\mathcal{B}}$.
- (iv) An H -invariant system of chains $\lambda \in C_{\text{vir}}^\bullet(f^*\mathcal{X} \text{ rel } \partial; \mathcal{B})^+$ (degree zero and supported inside $\text{supp } \mathcal{X}$) with the following property. Note that (μ, λ) is a cycle in the mapping cone

$$[C_{\text{vir}}^\bullet(\partial f^*\mathcal{X}; \mathcal{B})^+ \rightarrow C_{\text{vir}}^\bullet(f^*\mathcal{X} \text{ rel } \partial; \mathcal{B})^+],$$

whose homology is identified with $\check{H}^\bullet(f^*\mathcal{X}; \mathfrak{o}_{f^*\mathcal{X}})$. We require that the homology class of (μ, λ) equal $f^*\omega \in \check{H}^0(f^*\mathcal{X}; \mathfrak{o}_{f^*\mathcal{X}})$.

- (v) An H -invariant system of chains $\tilde{\lambda}$ for $QC_\bullet(E; \mathcal{B})^+$ (degree $\text{gr}(q) - \text{gr}(p) + \dim \sigma - 1$ and supported inside $\text{supp } \mathcal{X}$) whose image in $C_\bullet(E; \mathcal{B})^+$ coincides with $s_*\lambda$ (recall (7.3.10)).
- (vi) An H -invariant map of F -modules $[[E]]: QC_\bullet(E; \mathcal{B})_{\mathbb{Z}}^+ \rightarrow \mathbb{Z}$ which sends the fundamental class to 1. The left-hand side $QC_\bullet(E; \mathcal{B})_{\mathbb{Z}}^+$ is defined over \mathbb{Z} , and the right-hand side is the F -module which for every (σ, p, q) is \mathbb{Z} concentrated in degree zero (with the product/face maps being multiplication/identity, respectively).

Remark 7.5.4 (Resolution commutes with pullback) Let \mathcal{X}/Z_\bullet be an H -equivariant flow category diagram with implicit atlas \mathcal{A} and coherent orientations ω . Let $f: Y_\bullet \rightarrow Z_\bullet$ be a map, and consider $f^*\mathcal{X}/Y_\bullet$ with implicit atlas $f^*\mathcal{A}$ and coherent orientations $f^*\omega$. Then there is a canonical fiber diagram relating the resolutions:

$$(7.5.2) \quad \begin{array}{ccc} \tilde{Y}_\bullet & \rightarrow & \tilde{Z}_\bullet \\ \downarrow & & \downarrow \\ Y_\bullet & \rightarrow & Z_\bullet \end{array}$$

Proposition 7.5.5 ($\pi: \tilde{Z}_\bullet \rightarrow Z_\bullet$ is a trivial Kan fibration) The map $\pi: \tilde{Z}_\bullet \rightarrow Z_\bullet$ has the right lifting property with respect to the boundary inclusions $\partial\Delta^n \hookrightarrow \Delta^n$ for all $n \geq 0$ (where Δ^n is the semisimplicial n -simplex). In other words, given any commuting diagram

$$(7.5.3) \quad \begin{array}{ccc} \partial\Delta^n & \longrightarrow & \tilde{Z}_\bullet \\ \downarrow & \nearrow \text{dashed} & \downarrow \\ \Delta^n & \longrightarrow & Z_\bullet \end{array}$$

of solid arrows, there exists a dashed arrow making the diagram commute.

Proof It is equivalent to show that, given any $f: \Delta^n \rightarrow Z_\bullet$ along with data (ii)–(vi) over $\partial\Delta^n$, the data can be extended over Δ^n . For ease of notation, let us rename $(f^*\mathcal{X}, f^*\mathcal{A}, f^*\omega)$ as $(\mathcal{X}, \mathcal{A}, \omega)$.

We construct (ii)–(vi) via H –equivariant induction on the set of triples (σ, p, q) , equipped with the partial order $\preceq_{\mathcal{X}}$ in which $(\sigma', p', q') \preceq_{\mathcal{X}} (\sigma, p, q)$ if and only if (σ', p', q') appears as a “factor” of some element of $\mathcal{S}_{\mathcal{X}}(\sigma, p, q)$. This partial order is *well-founded* (ie there is no strictly decreasing infinite sequence) since each $\mathcal{S}_{\mathcal{X}}(\sigma, p, q)$ is finite, and thus it is valid for induction. The induction works H –equivariantly since the action of H on \mathcal{P} is free, and (σ, p, q) and (σ, hp, hq) are always incomparable for $h \in H$ (so we may henceforth ignore the action of H). The inductive step is as follows.

Choosing $\mathcal{B}(\sigma, p, q)$ We must choose a finite $\bar{\mathcal{B}}(\sigma, p, q) \subseteq \bar{\mathcal{A}}(\sigma, p, q)$ so that the corresponding subatlas of $\mathcal{A}(\sigma, p, q) \stackrel{\geq s^{\text{top}}}{=} \bar{\mathcal{A}}(\sigma, p, q)$ on $\mathcal{X}(\sigma, p, q)$ satisfies the covering axiom (and thus is an implicit atlas). This is possible by compactness.

Choosing $[E_\alpha]$ Trivial.

Choosing $[[E]]_{\sigma,p,q}: QC_\bullet(E; \mathcal{B})_{\mathbb{Z}}^+(\sigma, p, q) \rightarrow \mathbb{Z}$ We need to construct $[[E]]_{\sigma,p,q}$ fitting into the following diagram:

$$\begin{array}{ccc}
 \text{colim}_{s \in \partial\mathcal{S}_{\mathcal{X}}(\sigma,p,q)} QC_\bullet(E; \mathcal{B})_{\mathbb{Z}}^+(\sigma, p, q, s) & \xrightarrow{\text{product/face}} & QC_\bullet(E; \mathcal{B})_{\mathbb{Z}}^+(\sigma, p, q) \\
 \downarrow [[E]] & \swarrow \text{?}\exists [[E]]_{\sigma,p,q} & \\
 \mathbb{Z} & &
 \end{array}$$

(7.5.4)

$$\begin{array}{ccc}
 H_0 \text{ colim}_{s \in \partial\mathcal{S}_{\mathcal{X}}(\sigma,p,q)} QC_\bullet(E; \mathcal{B})_{\mathbb{Z}}^+(\sigma, p, q, s) & \xrightarrow{\text{product/face}} & H_0 QC_\bullet(E; \mathcal{B})_{\mathbb{Z}}^+(\sigma, p, q) \\
 \downarrow [[E]] & \swarrow [E] \mapsto 1 & \\
 \mathbb{Z} & &
 \end{array}$$

Since $QC_\bullet(E; \mathcal{B})_{\mathbb{Z}}^+$ is cofibrant, it follows that the top horizontal map is a cofibration between cofibrant complexes (this uses Lemma 7.4.4). Now Lemma 7.4.7 tells us that the direct sum of $H_0 QC_\bullet(E; \mathcal{B})_{\mathbb{Z}}^+(\sigma, p, q, s)$ surjects onto H_0 on the left, and this makes it easy to check that the second diagram commutes. Lemma 7.4.7 also tells us that H_{-1} on the left vanishes. Hence we may apply Lemma 7.4.11 to conclude the existence of a suitable $[[E]]_{\sigma,p,q}$.

Choosing $\lambda_{\sigma,p,q}$ (the key step) According to Lemma 7.4.9, we may choose a $\lambda_{\sigma,p,q}$ with the required property if and only if the homology class of

$$\mu_{\sigma,p,q} \in H_{\text{vir}}^\bullet(\partial\mathcal{X}; \mathcal{B})^+(\sigma, p, q)$$

coincides with

$$d\omega(\sigma, p, q) \in \check{H}^0(\partial\mathcal{X}(\sigma, p, q); \mathfrak{o}_{\partial\mathcal{X}(\sigma, p, q)}).$$

Let us now prove this desired statement $[\mu_{\sigma, p, q}] = d\omega(\sigma, p, q)$. Let us fix $\tau_0 \in \partial\mathcal{S}_X(\sigma, p, q)$ and verify equality over $\mathcal{X}(\sigma, p, q)^{\tau_0}$ (this is clearly enough).

We consider the following diagram, where $\mathcal{X}^{\leq s}$, $\mathcal{B}^{\geq t}$, $\mathcal{S}_X^{\leq \tau}$, $\partial\mathcal{S}_X$, $\partial\mathcal{S}_X^{\leq \tau}$ stand for $\mathcal{X}(\sigma, p, q)^{\leq s}$, $\mathcal{B}(\sigma, p, q)^{\geq t}$, $\mathcal{S}_X(\sigma, p, q)^{\leq \tau}$, $\partial\mathcal{S}_X(\sigma, p, q)$, $\partial\mathcal{S}_X(\sigma, p, q)^{\leq \tau}$, respectively:

$$(7.5.5) \quad \mu_{\sigma, p, q} \in C_{\text{vir}}^\bullet(\partial\mathcal{X}; \mathcal{B})^+(\sigma, p, q)$$

↑

$$(7.5.6) \quad \bigoplus_{\tau \in \partial\mathcal{S}_X} \mu_{\sigma, p, q, \tau} \in \bigoplus_{\tau \in \partial\mathcal{S}_X} \text{hocolim}_{s < t \in \mathcal{S}_X^{\leq \tau}} C_{\text{vir}}^\bullet(\mathcal{X}^{\leq s} \text{ rel } \partial; \mathcal{B}^{\geq t})$$

↓

$$(7.5.7) \quad \bigoplus_{\tau \in \partial\mathcal{S}_X^{\leq \tau_0}} \mu_{\sigma, p, q, \tau} \in \bigoplus_{\tau \in \partial\mathcal{S}_X^{\leq \tau_0}} \text{hocolim}_{s < t \in \mathcal{S}_X^{\leq \tau}} C_{\text{vir}}^\bullet(\mathcal{X}^{\leq s} \text{ rel } \partial; \mathcal{B}^{\geq t})$$

↑

$$(7.5.8) \quad \mu_{\sigma, p, q, \tau_0} \in \text{hocolim}_{s < t \in \mathcal{S}_X^{\leq \tau_0}} C_{\text{vir}}^\bullet(\mathcal{X}^{\leq s} \text{ rel } \partial; \mathcal{B}^{\geq t})$$

The differentials in (7.5.6)–(7.5.7) are given by the sum over $\text{codim}(\tau \preceq \tau') = 1$ of the obvious pushforward maps (plus the internal differentials of each hocolim). The first vertical map is the sum over $\text{codim } \tau = 0$ (ie over maximal elements of $\partial\mathcal{S}_X(\sigma, p, q)$); clearly this is a chain map. The remaining vertical maps are clear. The easiest way to keep track of signs in the present discussion is to use the orientation lines associated to elements of $\mathcal{S}_X(\sigma, p, q)$, as in Definition 6.2.2, though we will suppress them from the notation. Note that there is a natural diagram of complexes of \mathcal{K} –presheaves on $\partial\mathcal{X}(\sigma, p, q)$ whose diagram of global sections is (7.5.5)–(7.5.8). We will see below that these complexes of \mathcal{K} –presheaves are in fact pure homotopy \mathcal{K} –sheaves and that the induced diagram of sheaves (obtained by taking H^0) is given by the tautological maps

$$(7.5.9) \quad \mathfrak{o}_{\partial\mathcal{X}(\sigma, p, q)} \leftarrow \mathfrak{o}_{\partial\mathcal{X}(\sigma, p, q)} \rightarrow \mathfrak{o}_{\mathcal{X}(\sigma, p, q)^{\leq \tau_0}} \leftarrow \mathfrak{o}_{\mathcal{X}(\sigma, p, q)^{\leq \tau_0} \text{ rel } \partial}.$$

First, though, let us argue that this claim implies the desired result. The cycle $\mu_{\sigma, p, q} \in (7.5.5)$ lifts to a cycle $\bigoplus_{\tau \in \partial\mathcal{S}_X(\sigma, p, q)} \mu_{\sigma, p, q, \tau} \in (7.5.6)$, where $\mu_{\sigma, p, q, \tau}$ is obtained by applying the product/face maps to the tensor product of the various λ corresponding to the factors on the left-hand side of (7.1.6) corresponding to τ (this is a cycle by (7.5.1)). Obviously

$$\bigoplus_{\tau \in \partial\mathcal{S}_X(\sigma, p, q)} \mu_{\sigma, p, q, \tau} \in (7.5.6) \quad \text{maps to} \quad \bigoplus_{\tau \in \partial\mathcal{S}_X(\sigma, p, q)^{\leq \tau_0}} \mu_{\sigma, p, q, \tau} \in (7.5.7).$$

On sections over $K \subseteq \mathcal{X}(\sigma, p, q)^{\tau_0}$, the last vertical map is an isomorphism, and (the restriction to K of) $\bigoplus_{\tau \in \partial \mathcal{S}_{\mathcal{X}}(\sigma, p, q)^{\leq \tau_0}} \mu_{\sigma, p, q, \tau} \in (7.5.7)$ lifts to (the restriction to K of) $\mu_{\sigma, p, q, \tau_0} \in (7.5.8)$. Finally, observe that over any $K \subseteq \mathcal{X}(\sigma, p, q)^{\tau_0}$, the homology class of $\mu_{\sigma, p, q, \tau_0} \in (7.5.8)$ coincides with the restriction of $d\omega(\sigma, p, q)$ to the τ_0 stratum (this follows from the induction hypothesis, using the fact that the orientations are coherent and the observation at the end of Definition 6.3.2). Clearly this implies that $\mu_{\sigma, p, q}$ coincides with $d\omega(\sigma, p, q)$ in homology over $\mathcal{X}(\sigma, p, q)^{\tau_0}$.

Thus it remains only to show that (7.5.5)–(7.5.8) are pure homotopy \mathcal{K} -sheaves and that the induced maps on H^0 coincide with the tautological maps on orientation sheaves (7.5.9). Note that by Proposition 6.2.3, the following closely related diagram (using similar simplified notation as in (7.5.5)–(7.5.8)) satisfies all of the desired properties (ie is a diagram of pure homotopy \mathcal{K} -sheaves inducing (7.5.9) on H^0):

$$\begin{array}{c}
 (7.5.10) \quad C_{\text{vir}}^{\bullet}(\partial \mathcal{X}; \mathcal{B}^{\geq s^{\text{top}}}) \\
 \uparrow \\
 (7.5.11) \quad \bigoplus_{\tau \in \partial \mathcal{S}_{\mathcal{X}}} C_{\text{vir}}^{\bullet}(\mathcal{X}^{\leq \tau} \text{ rel } \partial; \mathcal{B}^{\geq s^{\text{top}}}) \\
 \downarrow \\
 (7.5.12) \quad \bigoplus_{\tau \in \partial \mathcal{S}_{\mathcal{X}}^{\leq \tau_0}} C_{\text{vir}}^{\bullet}(\mathcal{X}^{\leq \tau} \text{ rel } \partial; \mathcal{B}^{\geq s^{\text{top}}}) \\
 \uparrow \\
 (7.5.13) \quad C_{\text{vir}}^{\bullet}(\mathcal{X}^{\leq \tau_0} \text{ rel } \partial; \mathcal{B}^{\geq s^{\text{top}}})
 \end{array}$$

Hence it is enough to relate this diagram to (7.5.5)–(7.5.8) via quasi-isomorphisms. The diagram obtained from (7.5.5)–(7.5.8) by replacing every occurrence of $\mathcal{B}(\sigma, p, q)^{\geq t}$ with $\mathcal{B}(\sigma, p, q)^{\geq s^{\text{top}}}$ maps quasi-isomorphically to (7.5.5)–(7.5.8) (by pairing with the fixed fundamental cycles), and it also maps quasi-isomorphically to the diagram above (use the obvious pushforwards on the $p = 0$ level of the homotopy colimits, and zero for $p > 0$), thus giving the desired result.

Choosing $\tilde{\lambda}_{\sigma, p, q}$ We must find $\tilde{\lambda}_{\sigma, p, q}$ lifting $s_*\lambda_{\sigma, p, q}$ and satisfying $d\tilde{\lambda}_{\sigma, p, q} = \tilde{\mu}_{\sigma, p, q}$. Now we know that $\tilde{\mu}_{\sigma, p, q}$ is a cycle (see Definition 7.5.1); furthermore, its image in $C_{\bullet}(E; \mathcal{B})^+$ is $s_*\mu_{\sigma, p, q} = d(s_*\lambda_{\sigma, p, q})$, which is null-homologous. Since $Q\mathcal{W}_{\bullet} \xrightarrow{\sim} \mathcal{W}_{\bullet}$ is always a quasi-isomorphism, it follows that $\tilde{\mu}_{\sigma, p, q}$ is null-homologous. Hence there exists $\tilde{\lambda}'_{\sigma, p, q}$ with $d\tilde{\lambda}'_{\sigma, p, q} = \tilde{\mu}_{\sigma, p, q}$; now we must modify $\tilde{\lambda}'_{\sigma, p, q}$ so that it lifts $s_*\lambda_{\sigma, p, q}$. The difference of $s_*\lambda_{\sigma, p, q}$ and the image of $\tilde{\lambda}'_{\sigma, p, q}$ is a cycle in $C_{\bullet}(E; \mathcal{B})^+$. It can be lifted to a cycle in $QC_{\bullet}(E; \mathcal{B})^+$ using Lemma 7.4.8 (again using that $Q\mathcal{W}_{\bullet} \xrightarrow{\sim} \mathcal{W}_{\bullet}$ is a quasi-isomorphism), which is enough. \square

7.6 Categories of complexes

We review various categories of complexes over graded rings. We use these categories as targets for the Floer-type homology groups we construct.

Definition 7.6.1 (Complexes over graded rings) Let S be a graded ring. A *differential graded S -module* is a graded module A^\bullet over S along with a map $d: A^\bullet \rightarrow A^{\bullet+1}$ satisfying $d^2 = 0$.

Definition 7.6.2 (Category Ch_S) Let S be a graded ring. We let Ch_S denote the category whose objects are differential graded S -modules and whose morphisms are chain maps.

Definition 7.6.3 (Category $H^0(\text{Ch}_S)$) Let S be a graded ring. We let $H^0(\text{Ch}_S)$ denote the category whose objects are objects of Ch_S and whose morphisms are chain maps modulo chain homotopy.

Definition 7.6.4 (∞ -category $\text{N}_{\text{dg}}(\text{Ch}_S)$) Let S be a graded ring. We let $\text{N}_{\text{dg}}(\text{Ch}_S)$ denote the differential graded nerve of Ch_S (see Lurie [58, Construction 1.3.1.6] or Definition 7.6.5 below).

Let us now discuss the relationship between these three categories and in particular explain the definition of $\text{N}_{\text{dg}}(\text{Ch}_S)$ in concrete terms. Despite appearances, the reader need not be familiar with ∞ -categories to understand $\text{N}_{\text{dg}}(\text{Ch}_S)$.

Recall that if \mathcal{C} is a category and X_\bullet is a simplicial set, a *diagram* $X_\bullet \rightarrow \mathcal{C}$ is a map of simplicial sets $X_\bullet \rightarrow N_\bullet \mathcal{C}$ where $N_\bullet \mathcal{C}$ denotes the *nerve*²³ of \mathcal{C} . Concretely, a diagram $F: X_\bullet \rightarrow \mathcal{C}$ is the data of:

- (i) For every vertex $v \in X_0$, an object $A_v \in \mathcal{C}$.
- (ii) For every edge $e \in X_1$ from v_0 to v_1 , a morphism $f_e: A_{v_0} \rightarrow A_{v_1}$ in \mathcal{C} .

Moreover, the following conditions must be satisfied:

- (iii) For every degenerate edge $e \in X_1$ over a vertex v , we have $f_e = \text{id}_{A_v}$.
- (iv) For every face in X_2 spanning edges e_{01}, e_{12}, e_{02} , we have $f_{e_{12}} \circ f_{e_{01}} = f_{e_{02}}$.

²³An n -simplex $\Delta^n \rightarrow N_\bullet \mathcal{C}$ is a diagram $A_0 \xrightarrow{f_{01}} A_1 \xrightarrow{f_{12}} \dots \xrightarrow{f_{n-1,n}} A_n$ in \mathcal{C} . Strictly speaking, $N_\bullet \mathcal{C}$ is not a simplicial *set* unless \mathcal{C} is small.

We often speak of a diagram $X_\bullet \rightarrow \mathcal{C}$ where X_\bullet is only a *semisimplicial* set, in which case we ignore condition (iii).

It should now be clear what we mean by a diagram $X_\bullet \rightarrow \text{Ch}_S$ or $X_\bullet \rightarrow H^0(\text{Ch}_S)$ if X_\bullet is a (semi)simplicial set. Let us now say what we mean by a diagram $X_\bullet \rightarrow \text{N}_{\text{dg}}(\text{Ch}_S)$. Such a diagram is similar to a diagram $X_\bullet \rightarrow \text{Ch}_S$, except that we only require condition (iv) to hold “up to coherent higher homotopy”.

Definition 7.6.5 (Diagram $X_\bullet \rightarrow \text{N}_{\text{dg}}(\text{Ch}_S)$ [58, Construction 1.3.1.6]) Let S be a graded ring and let X_\bullet be a simplicial set. A diagram $X_\bullet \rightarrow \text{N}_{\text{dg}}(\text{Ch}_S)$ consists of:²⁴

- (i) For every $v \in X_0$, a graded S -module A_v^\bullet .
- (ii) For every $\sigma \in X_n$ spanning $v_0, \dots, v_n \in X_0$, a map $f_\sigma: A_{v_0} \rightarrow A_{v_n}$ of degree $1 - n$, such that

$$(7.6.1) \quad \sum_{k=0}^n (-1)^k f_{\sigma|[k\dots n]} \circ f_{\sigma|[0\dots k]} = \sum_{k=1}^{n-1} (-1)^k f_{\sigma|[0\dots\hat{k}\dots n]},$$

and if σ is degenerate, then

$$(7.6.2) \quad f_\sigma = \begin{cases} \text{id}, & \dim \sigma = 1, \\ 0, & \text{otherwise.} \end{cases}$$

If X_\bullet is only a semisimplicial set, then we ignore (7.6.2).

Let us now explain this definition by examining what (7.6.1) says for low-dimensional simplices $\sigma: \Delta^n \rightarrow X_\bullet$ (following [58, Example 1.3.1.8]).

- (i) Let $n = 0$. Then $f_0: A_0^\bullet \rightarrow A_0^\bullet$ has degree 1, and (7.6.1) asserts that $f_0 \circ f_0 = 0$. Thus (A_0, f_0) is a chain complex.
- (ii) Let $n = 1$. Then $f_{01}: A_0^\bullet \rightarrow A_1^\bullet$ has degree zero, and (7.6.1) asserts that $f_{01} \circ f_0 - f_1 \circ f_{01} = 0$. Thus $f_{01}: (A_0^\bullet, f_0) \rightarrow (A_1^\bullet, f_1)$ is a chain map. For a degenerate 1-simplex, (7.6.2) asserts that $f_{01}: (A_0^\bullet, f_0) \rightarrow (A_1^\bullet, f_1)$ is the identity map.
- (iii) Let $n = 2$. Then $f_{012}: A_0^\bullet \rightarrow A_2^\bullet$ has degree -1 , and (7.6.1) asserts that $f_{012}f_0 - f_{12}f_{01} + f_2f_{012} = -f_{02}$. Thus f_{012} is a chain homotopy between $f_{02}, f_{12}f_{01}: A_0^\bullet \rightarrow A_2^\bullet$.

²⁴For the reader familiar with ∞ -categories: the notion of a diagram $X_\bullet \rightarrow \text{N}_{\text{dg}}(\text{Ch}_S)$ suffices to define $\text{N}_{\text{dg}}(\text{Ch}_S)$ by Yoneda’s lemma.

We can now relate $N_{\text{dg}}(\text{Ch}_S)$ to the more familiar categories Ch_S and $H^0(\text{Ch}_S)$ by introducing natural forgetful functors

$$(7.6.3) \quad \text{Ch}_S \rightarrow N_{\text{dg}}(\text{Ch}_S) \rightarrow H^0(\text{Ch}_S),$$

which should be clear given the above interpretation of diagrams to $N_{\text{dg}}(\text{Ch}_S)$. More precisely, a diagram $X_\bullet \rightarrow \text{Ch}_S$ gives rise to a diagram $X_\bullet \rightarrow N_{\text{dg}}(\text{Ch}_S)$ where the higher homotopies $\{f_\sigma\}_{\dim \sigma \geq 2}$ are all zero. A diagram $X_\bullet \rightarrow N_{\text{dg}}(\text{Ch}_S)$ gives rise to a diagram $X_\bullet \rightarrow H^0(\text{Ch}_S)$ which forgets about the choice of higher homotopies $\{f_\sigma\}_{\dim \sigma \geq 2}$, remembering only the existence of $\{f_\sigma\}_{\dim \sigma = 2}$ satisfying (7.6.1).

7.7 Definition

We now define the Floer-type homology groups of an H -equivariant flow category diagram \mathcal{X}/Z_\bullet with implicit atlas and coherent orientations. More precisely, given such data we construct a diagram $\mathbb{H}(\mathcal{X}): Z_\bullet \rightarrow H^0(\text{Ch}_{R[[H]])}$.

We first define a diagram $\tilde{\mathbb{H}}(\mathcal{X}): \tilde{Z}_\bullet \rightarrow N_{\text{dg}}(\text{Ch}_{R[[H]])}$ (this is straightforward from the definition of \tilde{Z}_\bullet). Schematically:

$$(7.7.1) \quad \begin{array}{ccc} \tilde{Z}_\bullet & \xrightarrow{\tilde{\mathbb{H}}(\mathcal{X})} & N_{\text{dg}}(\text{Ch}_{R[[H]])} \\ \pi \downarrow & & \\ Z_\bullet & & \end{array}$$

We then use Proposition 7.5.5 (that $\pi: \tilde{Z}_\bullet \rightarrow Z_\bullet$ is a trivial Kan fibration) to show that $\tilde{\mathbb{H}}(\mathcal{X})$ descends uniquely to a diagram $\mathbb{H}(\mathcal{X}): Z_\bullet \rightarrow H^0(\text{Ch}_{R[[H]])}$.

For the descent argument, we need to assume that Z_\bullet is a *simplicial set*. The main nontrivial step is to show that certain lifts of degenerate edges in Z_\bullet to \tilde{Z}_\bullet are sent by $\tilde{\mathbb{H}}(\mathcal{X})$ to the identity map in $H^0(\text{Ch}_{R[[H]])}$.

Remark 7.7.1 There should also be a (more refined) descent with target $N_{\text{dg}}(\text{Ch}_{R[[H]])}$ (see also Remark 7.0.2), though we have decided to omit this for the sake of brevity (the correct uniqueness statement is more complicated to state).

Definition 7.7.2 (Novikov rings) Let T be a set equipped with a “grading” $\text{gr}: T \rightarrow \mathbb{Z}$ and an “action” $a: T \rightarrow \mathbb{R}$. We let $R[[T]]$ denote the graded R -module consisting of formal sums $\sum_{t \in T} c_t \cdot t$ with $c_t \in R$, satisfying the following two finiteness conditions:

- (i) $\#\{n \in \mathbb{Z} : \exists t \text{ such that } c_t \neq 0 \text{ and } \text{gr}(t) = n\} < \infty$.
- (ii) $\#\{t \in T : c_t \neq 0 \text{ and } a(t) < M\} < \infty$ for all $M < \infty$.

In other words, $R[[T]]$ is the graded completion of $R[T]$ with respect to the non-archimedean a -adic norm $|t|_a := e^{-a(t)}$. If T is a group and gr, a are group homomorphisms, then $R[[T]]$ is a graded ring.

Definition 7.7.3 ((Co)homology groups $\tilde{\mathbb{H}}(\mathcal{X}): \tilde{Z}_\bullet \rightarrow \text{N}_{\text{dg}}(\text{Ch}_{R[[H]])}$) Let \mathcal{X}/Z_\bullet be an H -equivariant flow category diagram with implicit atlas \mathcal{A} and coherent orientations ω . We define a diagram

$$(7.7.2) \quad \tilde{\mathbb{H}}(\mathcal{X})_{\mathcal{A},\omega}: \tilde{Z}_\bullet \rightarrow \text{N}_{\text{dg}}(\text{Ch}_{R[[H]])}.$$

We will write $\tilde{\mathbb{H}}(\mathcal{X})$ for $\tilde{\mathbb{H}}(\mathcal{X})_{\mathcal{A},\omega}$ when the atlas and orientations are clear from context.

To a vertex of \tilde{Z}_\bullet we associate $R[[\mathcal{P}_z]]$, where z is the corresponding vertex in Z_\bullet . This is clearly a graded $R[[H]]$ -module.

Now for a simplex $\sigma \in \tilde{Z}_n$, we aim to define the map $f_\sigma: R[[\mathcal{P}_{\sigma(0)}]] \rightarrow R[[\mathcal{P}_{\sigma(n)}]]$ by the formula

$$(7.7.3) \quad f_\sigma(p) := \sum_{q \in \mathcal{P}_{\sigma(n)}} ([E]_{\sigma,p,q} \otimes \text{id}_R)(\tilde{\lambda}_{\sigma,p,q}) \cdot q.$$

Let us now argue that (7.7.3) makes sense and that the resulting maps f_σ satisfy (7.6.1).

First, observe that since $\tilde{\lambda}_{\sigma,p,q}$ is of degree $\text{gr}(q) - \text{gr}(p) + \dim \sigma - 1$ and $[E]_{\sigma,p,q}$ has degree zero, all terms on the right-hand side of (7.7.3) are of degree $\text{gr}(p) + 1 - \dim \sigma$. Now, $\tilde{\lambda}_{\sigma,p,q} = 0$ if $\mathcal{X}(\sigma, p, q) = \emptyset$. Hence finiteness condition Definition 7.2.1(iv) implies that (7.7.3) converges in the a -adic topology, and hence defines an R -linear map $f_\sigma: R[[\mathcal{P}_{\sigma(0)}]] \rightarrow R[[\mathcal{P}_{\sigma(n)}]]$. Finiteness condition (v) implies that there exists $M < \infty$ such that $|f_\sigma(p)|_a \leq M|p|_a$. Thus f_σ extends uniquely to a continuous (in fact, bounded) R -linear map $f_\sigma: R[[\mathcal{P}_{\sigma(0)}]] \rightarrow R[[\mathcal{P}_{\sigma(n)}]]$. Clearly f_σ is $R[[H]]$ -linear since $[E]$ and $\tilde{\lambda}$ are both H -invariant.

Now it remains to verify that the f_σ defined by (7.7.3) satisfy (7.6.1). To see this, write

$$\begin{aligned} 0 &= (-1)^{\text{gr}(p)} \cdot [E]_{\sigma,p,r}(d\tilde{\lambda}_{\sigma,p,r}) \\ &= (-1)^{\text{gr}(p)} \cdot [E]_{\sigma,p,r}(\tilde{\mu}_{\sigma,p,r}) \\ &= [E]_{\sigma,p,r} \left(\sum_{k=0}^n \sum_{q \in \mathcal{P}_{\sigma(k)}} (-1)^k \tilde{\lambda}_{\sigma|[0\dots k],p,q} \cdot \tilde{\lambda}_{\sigma|[k\dots n],q,r} \right. \\ &\quad \left. - \sum_{k=1}^{n-1} (-1)^k \tilde{\lambda}_{\sigma|[0\dots \hat{k}\dots n],p,r} \right) \end{aligned}$$

$$\begin{aligned}
 &= \sum_{k=0}^n \sum_{q \in \mathcal{P}_{\sigma(k)}} (-1)^k \llbracket E \rrbracket_{\sigma|[0\dots k],p,q}(\tilde{\lambda}_{\sigma|[0\dots k],p,q}) \cdot \llbracket E \rrbracket_{\sigma|[k\dots n],q,r}(\tilde{\lambda}_{\sigma|[k\dots n],q,r}) \\
 &\qquad\qquad\qquad - \sum_{k=1}^{n-1} (-1)^k \llbracket E \rrbracket_{\sigma|[0\dots \hat{k}\dots n],p,r}(\tilde{\lambda}_{\sigma|[0\dots \hat{k}\dots n],p,r}) \\
 &= \text{coefficient of } r \text{ in } \sum_{k=0}^n (-1)^k f_{\sigma|[k\dots n]}(f_{\sigma|[0\dots k]}(p)) - \sum_{k=1}^{n-1} (-1)^k f_{\sigma|[0\dots \hat{k}\dots n]}(p).
 \end{aligned}$$

The first equality follows because $\llbracket E \rrbracket_{\sigma,p,r}$ is a chain map, the second equality follows from the definition of a system of chains, the third equality is (7.5.1), the fourth equality follows from the fact that $\llbracket E \rrbracket$ is compatible with the product/face maps, and the fifth equality follows from the definition of f_{σ} (7.7.3).

Proposition 7.7.4 (Degenerate edges in \tilde{Z}_{\bullet} give the identity map up to homotopy) *Let $Z_{\bullet} = *$ be the semisimplicial set with a single simplex σ^i in dimension i for all $i \geq 0$ (ie the simplicial 0–simplex). Let \mathcal{X}/Z_{\bullet} be an H –equivariant flow category diagram with implicit atlas and coherent orientations. Suppose that:*

- (i) $\mathcal{X}(\sigma^1, p, p) = \mathcal{X}(\sigma^1, p, p)^{\text{reg}}$ is a single point and $\omega(\sigma^1, p, p) = 1$, for all p .
- (ii) $\mathcal{X}(\sigma^0, p, q) = \emptyset \implies \mathcal{X}(\sigma^i, p, q) = \emptyset$ for all i and all $p \neq q$.

Then for any $\sigma \in \tilde{Z}_1$ whose two vertices coincide, the associated map in $\tilde{\mathbb{H}}(\mathcal{X})$ is homotopic to the identity map.

Proof Let $d: R[\mathcal{P}] \rightarrow R[\mathcal{P}]$ denote the boundary operator associated to the vertex of σ , and let $1 - \epsilon: R[\mathcal{P}] \rightarrow R[\mathcal{P}]$ denote the chain map associated to σ . We must show that ϵ is chain homotopic to the zero map. Since $\tilde{Z}_{\bullet} \rightarrow Z_{\bullet}$ is a trivial Kan fibration (Proposition 7.5.5), there exists a 2–simplex in \tilde{Z}_{\bullet} all of whose edges are σ . Associated to this 2–simplex is a chain homotopy $h: R[\mathcal{P}] \rightarrow R[\mathcal{P}]$ between $(1 - \epsilon)$ and $(1 - \epsilon) \circ (1 - \epsilon)$; in other words,

$$(7.7.4) \qquad \epsilon = dh + hd + \epsilon^2.$$

Now we claim that the only nonzero “matrix coefficients” $c_{p,q}$ of d, ϵ, h are those for which $\mathcal{X}(\sigma^0, p, q) \neq \emptyset$. By definition, f_{σ} can have nonzero matrix coefficients $c_{p,q}$ only for $\mathcal{X}(\sigma, p, q) \neq \emptyset$. Hence hypothesis (ii) gives the desired claim as long as $p \neq q$. For the diagonal matrix coefficients $c_{p,p}$, we argue separately as follows. By degree considerations, only ϵ can have nonzero $c_{p,p}$. It thus suffices to show that the matrix coefficient $c_{p,p} = \llbracket E \rrbracket_{\sigma^1,p,p}(\tilde{\lambda}_{\sigma^1,p,p})$ of $f_{\sigma^1} = 1 - \epsilon$ equals 1. Note that no product/face maps have target (σ^1, p, p) , so $\tilde{\lambda}_{\sigma^1,p,p}$ is a cycle representing $\omega(\sigma^1, p, p) = 1$. Hence $\tilde{\lambda}_{\sigma^1,p,p}$ is a cycle representing $[E]$, and so $\llbracket E \rrbracket_{\sigma^1,p,p}(\tilde{\lambda}_{\sigma,p,p}) = 1$ as needed. Hence the claim is valid. From the claim, we make the following two observations:

- (i) We have $|d(p)|_a, |\epsilon(p)|_a, |h(p)|_a \leq |p|_a$.
- (ii) For all $M < \infty$ and $p \in \mathcal{P}$, there exists $N < \infty$ such that any length- N composition of d, ϵ, h applied to p has a -adic norm $\leq e^{-M}$.

(The second observation also uses Definition 7.2.1(iv).)

Now iterating the identity $\epsilon = dh + hd + \epsilon^2$, we are led to the infinite series²⁵

$$(7.7.5) \quad \epsilon = \sum_{n=1}^{\infty} \frac{1}{n} \binom{2n-2}{n-1} (hd + dh)^n,$$

which by observations (i)–(ii) converges in the a -adic topology when applied to any element of $R[[\mathcal{P}]]$. Now we have

$$(hd + dh)^n = \sum_{k=0}^n (dh)^k (hd)^{n-k} = H_n d + dH_n,$$

where $H_n = \sum_{k=1}^n h(dh)^{k-1} (hd)^{n-k}$. This gives the desired chain homotopy between ϵ and zero (again using observations (i)–(ii) to justify convergence of infinite sums). \square

Lemma 7.7.5 (Criterion for descent along a trivial Kan fibration) *Let $\pi: \tilde{Z}_\bullet \rightarrow Z_\bullet$ be a trivial Kan fibration of semisimplicial sets, and let $\tilde{\mathbb{H}}: \tilde{Z}_\bullet \rightarrow \mathcal{C}$ be a diagram in some category \mathcal{C} . A descent of $\tilde{\mathbb{H}}$ to Z_\bullet is a diagram $\mathbb{H}: Z_\bullet \rightarrow \mathcal{C}$ along with an isomorphism $\tilde{\mathbb{H}} \rightarrow \mathbb{H} \circ \pi$:*

$$(7.7.6) \quad \begin{array}{ccc} \tilde{Z}_\bullet & \xrightarrow{\tilde{\mathbb{H}}} & \mathcal{C} \\ \pi \downarrow & \nearrow \mathbb{H} & \\ Z_\bullet & & \end{array}$$

Suppose that:

- (i) Z_\bullet is a simplicial set.
- (ii) For any edge σ^1 in \tilde{Z}_\bullet whose endpoints coincide and which projects to a degenerate edge in Z_\bullet , the associated map in \mathcal{C} is the identity map.

Then a descent exists and is unique up to unique isomorphism.

Proof The proof below shows that we may define $\mathbb{H} := \tilde{\mathbb{H}} \circ s$ for any section $s: Z_\bullet \rightarrow \tilde{Z}_\bullet$.

²⁵The coefficients $\frac{1}{n} \binom{2n-2}{n-1}$ are integers (Catalan numbers C_{n-1}), so writing this expression does not make any implicit assumptions on the ring R .

Defining \mathbb{H} over 0–simplices Fix some 0–simplex v of Z_\bullet . For every lift \tilde{v} , we have an object $\tilde{\mathbb{H}}(\tilde{v})$. Furthermore, for every lift \tilde{e} of the degenerate 1–simplex e over v , we get a map

$$\tilde{\mathbb{H}}(\tilde{v}_1) \xrightarrow{\tilde{\mathbb{H}}(\tilde{e})} \tilde{\mathbb{H}}(\tilde{v}_2),$$

which is the identity map if $\tilde{v}_1 = \tilde{v}_2$. Finally, for every lift \tilde{f} of the degenerate 2–simplex f over v , the diagram

$$(7.7.7) \quad \begin{array}{ccc} & \tilde{\mathbb{H}}(\tilde{v}_2) & \\ \tilde{\mathbb{H}}(\tilde{e}_{12}) \nearrow & & \searrow \tilde{\mathbb{H}}(\tilde{e}_{23}) \\ \tilde{\mathbb{H}}(\tilde{v}_1) & \xrightarrow{\tilde{\mathbb{H}}(\tilde{e}_{13})} & \tilde{\mathbb{H}}(\tilde{v}_3) \end{array}$$

commutes, where $\tilde{e}_{12}, \tilde{e}_{23}, \tilde{e}_{13}$ are the edges of \tilde{f} . Using that $\tilde{Z}_\bullet \rightarrow Z_\bullet$ is a trivial Kan fibration to conclude that we can always find lifts with specified boundary conditions, it follows that all $\tilde{\mathbb{H}}(\tilde{v})$ are canonically isomorphic. Thus there exists a unique choice for $\mathbb{H}(v)$.

Defining \mathbb{H} over 1–simplices Fix some 1–simplex e of Z_\bullet with vertices v_1, v_2 . If \tilde{e} is any lift of e , then we get a map

$$\mathbb{H}(v_1) = \tilde{\mathbb{H}}(\tilde{v}_1) \xrightarrow{\tilde{\mathbb{H}}(\tilde{e})} \tilde{\mathbb{H}}(\tilde{v}_2) = \mathbb{H}(v_2).$$

Furthermore, this map $\mathbb{H}(v_1) \rightarrow \mathbb{H}(v_2)$ is seen to be independent of the choice of lift \tilde{e} by lifting degenerate 2–simplices over e . Thus there exists a unique choice for $\mathbb{H}(e)$.

Defining \mathbb{H} over n –simplices for $n \geq 2$ We just need to check that the diagram

$$(7.7.8) \quad \begin{array}{ccc} & \mathbb{H}(v_2) & \\ \mathbb{H}(e_{12}) \nearrow & & \searrow \mathbb{H}(e_{23}) \\ \mathbb{H}(v_1) & \xrightarrow{\mathbb{H}(e_{13})} & \mathbb{H}(v_3) \end{array}$$

commutes for all 2–simplices f in Z_\bullet , where e_{12}, e_{23}, e_{13} are the edges of f . This follows from lifting f and the commutativity of (7.7.7). □

Definition 7.7.6 ((Co)homology groups $\mathbb{H}(\mathcal{X}): Z_\bullet \rightarrow H^0(\text{Ch}_R[[H]])$) Let Z_\bullet be a simplicial set. Let \mathcal{X}/Z_\bullet be an H –equivariant flow category diagram with implicit atlas \mathcal{A} and coherent orientations ω . Suppose that for any vertex σ^0 of Z_\bullet , we have:

- (i) $\mathcal{X}(\sigma^1, p, p) = \mathcal{X}(\sigma^1, p, p)^{\text{reg}}$ is a single point and $\omega(\sigma^1, p, p) = 1$, for all p .
- (ii) $\mathcal{X}(\sigma^0, p, q) = \emptyset \implies \mathcal{X}(\sigma^i, p, q) = \emptyset$ for all i and $p \neq q$.

(Here σ^i denotes the completely degenerate i -simplex over σ^0 .) We have

$$(7.7.9) \quad \begin{array}{ccc} \tilde{Z}_\bullet & \xrightarrow{\tilde{\mathbb{H}}(\mathcal{X})} & H^0(\text{Ch}_{R[[H]]}) \\ \pi \downarrow & & \\ Z_\bullet & & \end{array}$$

(abusing notation and using $\tilde{\mathbb{H}}(\mathcal{X}): \tilde{Z}_\bullet \rightarrow H^0(\text{Ch}_{R[[H]])$) to denote the composition of $\tilde{\mathbb{H}}(\mathcal{X})$ with the forgetful functor $\text{N}_{\text{dg}} \rightarrow H^0$). The hypotheses of Lemma 7.7.5 are satisfied by Propositions 7.5.5 and 7.7.4, and hence we get a canonical descent

$$(7.7.10) \quad \mathbb{H}(\mathcal{X})_{A,\omega}: Z_\bullet \rightarrow H^0(\text{Ch}_{R[[H]])}$$

We will write $\mathbb{H}(\mathcal{X})$ for $\mathbb{H}(\mathcal{X})_{A,\omega}$ when the atlas and orientations are clear from context.

7.8 Properties

Lemma 7.8.1 (Passing to a subatlas preserves $\mathbb{H}(\mathcal{X})$) *Let $Z_\bullet, \mathcal{X}/Z_\bullet, \mathcal{A}$ and ω be as in Definition 7.7.6. If $\mathcal{B} \subseteq \mathcal{A}$ is any subatlas, then there is a canonical isomorphism $\mathbb{H}(\mathcal{X})_{\mathcal{A}} = \mathbb{H}(\mathcal{X})_{\mathcal{B}}$.*

Proof Indeed, we have $\tilde{Z}_\bullet^{\mathcal{B}} \subseteq \tilde{Z}_\bullet^{\mathcal{A}}$ compatibly with $\tilde{\mathbb{H}}(\mathcal{X})$. □

Lemma 7.8.2 (Shrinking the charts preserves $\mathbb{H}(\mathcal{X})$) *Let $Z_\bullet, \mathcal{X}/Z_\bullet, \mathcal{A}$ and ω be as in Definition 7.7.6. Let \mathcal{A}' be obtained from \mathcal{A} by using instead some open subsets $U'_{IJ} \subseteq U_{IJ}, X'_I \subseteq X_I, X_I^{\text{reg}'} \subseteq X_I^{\text{reg}}$, and restricting ψ_{IJ}, s_I to these subsets, so that \mathcal{A}' is also an implicit atlas. Then there is a canonical isomorphism $\mathbb{H}(\mathcal{X})_{\mathcal{A}} = \mathbb{H}(\mathcal{X})_{\mathcal{A}'}$.*

Proof There is a natural map $\tilde{Z}_\bullet^{\mathcal{A}'} \rightarrow \tilde{Z}_\bullet^{\mathcal{A}}$ for which the pullback of $\tilde{\mathbb{H}}(\mathcal{X})_{\mathcal{A}}$ is canonically isomorphic to $\tilde{\mathbb{H}}(\mathcal{X})_{\mathcal{A}'}$. This is enough. □

Lemma 7.8.3 (Universal coefficients) *Let $Z_\bullet, \mathcal{X}/Z_\bullet, \mathcal{A}, \omega$ be as in Definition 7.7.6. Let $R \rightarrow S$ a homomorphism of base rings. Then there is a canonical isomorphism*

$$\mathbb{H}(\mathcal{X})^S = \mathbb{H}(\mathcal{X})^R \otimes_{R[[H]]} S[[H]]$$

(denoting $- \otimes_{R[[H]]} S[[H]]: H^0(\text{Ch}_{R[[H]])} \rightarrow H^0(\text{Ch}_{S[[H]])}$).

Proof There is a natural base change map $b: \tilde{Z}_\bullet^R \rightarrow \tilde{Z}_\bullet^S$ and a canonical isomorphism $\tilde{\mathbb{H}}(\mathcal{X})^S \circ b = \tilde{\mathbb{H}}(\mathcal{X})^R \otimes_{R[[H]]} S[[H]]$. Now the result follows since $\tilde{\mathbb{H}}(\mathcal{X})^S \circ b$ and $\tilde{\mathbb{H}}(\mathcal{X})^S$ have the same descent to Z_\bullet . □

Proposition 7.8.4 (If $\mathcal{X} = \mathcal{X}^{\text{reg}}$, then $\mathbb{H}(\mathcal{X})$ is given by counting 0-dimensional flow spaces) *Let $Z_\bullet, \mathcal{X}/Z_\bullet, A$, and ω be as in Definition 7.7.6. Suppose that $\mathcal{X}(\sigma, p, q) = \mathcal{X}(\sigma, p, q)^{\text{reg}}$ for all (σ, p, q) (so, in particular, \mathcal{X}/Z_\bullet is Morse–Smale in the sense of Remark 7.2.2).*

Define a diagram $\mathbb{H}'(\mathcal{X}): Z_\bullet \rightarrow \text{N}_{\text{dg}}(\text{Ch}_{R[[H]]})$ as follows. The graded $R[[H]]$ -module associated to a vertex z of \tilde{Z}_\bullet is $R[[\mathcal{P}_z]]$. For a simplex $\sigma \in \tilde{Z}_n$, we define the map $f_\sigma: R[[\mathcal{P}_{\sigma(0)}]] \rightarrow R[[\mathcal{P}_{\sigma(n)}]]$ by the formula

$$(7.8.1) \quad f_\sigma(p) := \sum_{\substack{q \in \mathcal{P}_{\sigma(n)} \\ \text{gr}(q) - \text{gr}(p) = 1 - \dim \sigma}} \langle \omega(\sigma, p, q), [\mathcal{X}(\sigma, p, q)] \rangle \cdot q.$$

It is easy to verify that the maps f_σ are well-defined and satisfy (7.6.1).

Now there is a canonical isomorphism $\mathbb{H}(\mathcal{X}) = \mathbb{H}'(\mathcal{X})$ (where on the right-hand side we implicitly compose with the forgetful functor $\text{N}_{\text{dg}} \rightarrow H^0$).

Proof We will show an equality of “matrix coefficients”

$$(7.8.2) \quad ([E]_{\sigma,p,q} \otimes \text{id}_R)(\tilde{\lambda}_{\sigma,p,q}) = \langle \omega(\sigma, p, q), [\mathcal{X}(\sigma, p, q)] \rangle$$

for $(\sigma, p, q) \in \tilde{Z}_\bullet$ with $\text{vdim } \mathcal{X}(\sigma, p, q) = 0$. Thus comparing (7.7.3) and (7.8.1), it follows that the two diagrams $\tilde{Z}_\bullet \rightarrow \text{N}_{\text{dg}}(\text{Ch}_{R[[H]]})$ in question, namely $\tilde{\mathbb{H}}(\mathcal{X})$ and $\mathbb{H}'(\mathcal{X}) \circ \pi$, coincide. The desired isomorphism thus follows from the definition of $\mathbb{H}(\mathcal{X})$ as the descent of $\tilde{\mathbb{H}}(\mathcal{X})$.

To prove (7.8.2), argue as follows. Since $\mathcal{X}(\sigma, p, q) = \mathcal{X}(\sigma, p, q)^{\text{reg}}$, we know that if $\text{vdim } \mathcal{X}(\sigma, p, q) < 0$ then $\mathcal{X}(\sigma, p, q) = \emptyset$ and hence $\lambda_{\sigma,p,q} = 0$ and $\tilde{\lambda}_{\sigma,p,q} = 0$. Now when $\text{vdim } \mathcal{X}(\sigma, p, q) = 0$, we see that $\mu_{\sigma,p,q} = 0$ and $\tilde{\mu}_{\sigma,p,q} = 0$ (because all of the terms defining them involve a triple with negative dimension), so $\lambda_{\sigma,p,q}$ and $\tilde{\lambda}_{\sigma,p,q}$ are cycles. Since the homology class of $\lambda_{\sigma,p,q}$ coincides with $\omega(\sigma, p, q)$, the left-hand side of (7.8.2) coincides with the evaluation of $\omega(\sigma, p, q)$ on $[\mathcal{X}(\sigma, p, q)]^{\text{vir}} \in \check{H}^0(\mathcal{X}(\sigma, p, q); \mathfrak{o}_{\mathcal{X}(\sigma,p,q)})^\vee$. Now we are done since $[\mathcal{X}(\sigma, p, q)]^{\text{vir}} = [\mathcal{X}(\sigma, p, q)]$ by Lemma 5.2.6. \square

8 S^1 -localization

In this section, we prove vanishing results for virtual fundamental cycles on almost free²⁶ S^1 -spaces equipped with an S^1 -equivariant implicit atlas.

Convention 8.0.1 In this section, we work over a fixed ground ring R , and everything takes place in the category of R -modules. We restrict to implicit atlases A for which

²⁶An “almost free” action is one for which the stabilizer group of every point is finite.

$\#\Gamma_\alpha$ is invertible in R for all $\alpha \in A$. We restrict to S^1 -equivariant implicit atlases for which $\#(S^1)_p$ is invertible in R for all $p \in X_\emptyset$.

For this purpose, we introduce the S^1 -equivariant virtual cochain complexes $C_{S^1, \text{vir}}^\bullet$, which enjoys properties similar to those of C_{vir}^\bullet . There are canonical ‘‘comparison maps’’

$$(8.0.1) \quad C_{S^1, \text{vir}}^\bullet(X; A) \rightarrow C_{\text{vir}}^\bullet(X; A),$$

$$(8.0.2) \quad C_{S^1, \text{vir}}^\bullet(X \text{ rel } \partial; A) \rightarrow C_{\text{vir}}^\bullet(X \text{ rel } \partial; A)$$

and a canonical commutative diagram

$$(8.0.3) \quad \begin{array}{ccc} C_{S^1, \text{vir}}^{d+\bullet}(X \text{ rel } \partial; A) & \xrightarrow{s_*} & C_{\dim E_A - \bullet - 1}^{S^1}(E_A, E_A \setminus 0; \mathfrak{o}_{E_A}^\vee)^{\Gamma_A} \\ \downarrow & & \downarrow \pi^! \\ C_{\text{vir}}^{d+\bullet}(X \text{ rel } \partial; A) & \xrightarrow{s_*} & C_{\dim E_A - \bullet}(E_A, E_A \setminus 0; \mathfrak{o}_{E_A}^\vee)^{\Gamma_A} \end{array}$$

(S^1 acting trivially on E_A). Furthermore, if X is an almost free S^1 -space and A is locally S^1 -orientable, then there are canonical isomorphisms

$$(8.0.4) \quad H_{S^1, \text{vir}}^\bullet(X; A) = \check{H}^\bullet(X/S^1; \pi_* \mathfrak{o}_X),$$

$$(8.0.5) \quad H_{S^1, \text{vir}}^\bullet(X \text{ rel } \partial; A) = \check{H}^\bullet(X/S^1; \pi_* \mathfrak{o}_{X \text{ rel } \partial}).$$

The construction of $C_{S^1, \text{vir}}^\bullet$ and the proof of these properties are the main technical ingredients for the S^1 -localization statements we prove.

To prove the desired vanishing results, we consider using $C_{S^1, \text{vir}}^\bullet$ in place of C_{vir}^\bullet to define virtual fundamental cycles. The properties and compatibilities above then show that the desired statements follow essentially from the vanishing (on homology) of the right vertical map in (8.0.3). This basic strategy works easily to give the desired vanishing results for the virtual fundamental classes from Section 5. We also apply this strategy to prove results for the Floer-type homology groups from Section 7 in the presence of an S^1 -action on the flow spaces (to the effect that flow spaces on which the action is almost free may be ignored).

We do not construct an S^1 -equivariant virtual fundamental cycle, though the machinery we set up is a step in this direction (see Remark 8.6.3).

8.1 Background on S^1 -equivariant homology

Definition 8.1.1 (Gysin sequence) Let $\pi: E \rightarrow B$ be a principal S^1 -bundle. Analysis of the associated Serre spectral sequence gives the Gysin long exact sequence

$$(8.1.1) \quad \cdots \xrightarrow{\cap e} H_{\bullet-1}(B) \xrightarrow{\pi^!} H_\bullet(E) \xrightarrow{\pi_*} H_\bullet(B) \xrightarrow{\cap e} H_{\bullet-2}(B) \xrightarrow{\pi^!} \cdots$$

The $\cap e$ map is cap product with the Euler class $e(E) \in H^2(B)$. To see this, observe that the sequence (8.1.1) coincides with the long exact sequence of the pair for (the mapping cone of) $\pi: E \rightarrow B$, and that there is a natural isomorphism

$$H_{\bullet+2}(B, E) \xrightarrow{\cap \tau} H_{\bullet}(B),$$

where $\tau = \tau(E) \in H^2(B, E)$ is the Thom class (eg as argued in [39, page 444] for the corresponding sequence in cohomology).

For any S^1 -space X , there is a principal S^1 -bundle $\pi: X \times ES^1 \rightarrow (X \times ES^1)/S^1$, and thus a long exact sequence

$$(8.1.2) \quad \cdots \xrightarrow{\cap e} H_{\bullet-1}^{S^1}(X) \xrightarrow{\pi^!} H_{\bullet}(X) \xrightarrow{\pi_*} H_{\bullet}^{S^1}(X) \xrightarrow{\cap e} H_{\bullet-2}^{S^1}(X) \xrightarrow{\pi^!} \cdots$$

for $e(X) \in H_{S^1}^2(X)$. The same reasoning applies for pairs of spaces, so the same exact sequence exists for relative homology as well.

Lemma 8.1.2 *Let X be a trivial S^1 -space. Then the $\pi^!$ map in the Gysin sequence vanishes, turning it into a short exact sequence*

$$(8.1.3) \quad 0 \rightarrow H_{\bullet}(X) \xrightarrow{\pi_*} H_{\bullet}^{S^1}(X) \xrightarrow{\cap e} H_{\bullet-2}^{S^1}(X) \rightarrow 0.$$

The same statement applies to relative homology of trivial S^1 -spaces.

Proof The composition $H_{\bullet}(X) \rightarrow H_{\bullet}(X \times BS^1) \rightarrow H_{\bullet}(X)$ is the identity map, and so the first map is injective. On the other hand, this map is precisely $\pi_*: H_{\bullet}(X) \rightarrow H_{\bullet}^{S^1}(X)$ since the S^1 -action is trivial, so π_* is injective, which is sufficient. The same argument applies in the relative setting as well. \square

Lemma 8.1.3 *Let X be a locally compact Hausdorff S^1 -space which is almost free at $p \in X$, and suppose that the order of the stabilizer $\#(S^1)_p$ is invertible in the ground ring R . Then there exists an S^1 -invariant neighborhood $S^1 p \subseteq K \subseteq X$ so that the Euler class $e(K) \in H_{S^1}^2(K)$ vanishes.*

Proof Apply the Tietze extension theorem to the identity map $S^1 p \rightarrow S^1 p$ to obtain an S^1 -invariant neighborhood K of $S^1 p$ and a retraction $r: K \rightarrow S^1 p$. By averaging and passing to a smaller neighborhood, we may assume without loss of generality that r is S^1 -equivariant. Now by the naturality of the Euler class, we have $e(K) = r^* e(S^1 p)$. Thus it suffices to show that $e(S^1 p) \in H_{S^1}^2(S^1 p)$ vanishes. Now we have

$$H_{S^1}^2(S^1 p) = H_{(S^1)_p}^2(p) = H^2((S^1)_p),$$

where the latter is the group cohomology of the finite stabilizer group $(S^1)_p$. It is a standard fact that the group cohomology of a finite group is annihilated by the order of the group. Hence our assumption that $\#(S^1)_p$ is invertible in R guarantees that this cohomology group vanishes. \square

Lemma 8.1.4 *Let M be a topological S^1 -manifold of dimension d which is almost free near $p \in M$. Suppose that $\#(S^1)_p$ is invertible in the ground ring R . Then there is a canonical isomorphism*

$$(8.1.4) \quad H_{\bullet}^{S^1}(M, M \setminus S^1 p) = \begin{cases} H^1(S^1 p; \mathfrak{o}_M), & \bullet = d - 1, \\ 0, & \bullet \neq d - 1 \end{cases}$$

(note also that $H^1(S^1 p; \mathfrak{o}_M) = H^0(S^1 p; \mathfrak{o}_M \otimes \mathfrak{o}_{S^1 p}^{\vee})$).

Proof By Poincaré duality, we have canonical isomorphisms

$$(8.1.5) \quad H_{\bullet}(M, M \setminus S^1 p) = \begin{cases} H^0(S^1 p; \mathfrak{o}_M), & \bullet = d, \\ H^1(S^1 p; \mathfrak{o}_M), & \bullet = d - 1, \\ 0, & \text{otherwise.} \end{cases}$$

By excision and Lemma 8.1.3, the Gysin sequence reduces to a short exact sequence

$$(8.1.6) \quad 0 \rightarrow H_{\bullet-1}^{S^1}(M, M \setminus S^1 p) \xrightarrow{\pi^!} H_{\bullet}(M, M \setminus S^1 p) \xrightarrow{\pi_*} H_{\bullet}^{S^1}(M, M \setminus S^1 p) \rightarrow 0.$$

Combining this with (8.1.5), we see that $H_{\bullet}^{S^1}(M, M \setminus S^1 p) = 0$ for $\bullet \neq d - 1$, and that both maps in the sequence

$$(8.1.7) \quad H_{d-1}(M, M \setminus S^1 p) \xrightarrow{\pi_*} H_{d-1}^{S^1}(M, M \setminus S^1 p) \xrightarrow{\pi^!} H_d(M, M \setminus S^1 p)$$

are isomorphisms, which yields the desired result. \square

8.2 S^1 -equivariant implicit atlases

Definition 8.2.1 (S^1 -equivariant implicit atlas) Let X be an S^1 -space. An S^1 -equivariant implicit atlas A on X is an implicit atlas A along with an action of S^1 on each thickening X_I (commuting with the Γ_I -action) such that each map ψ_{IJ} is S^1 -equivariant, each function s_{α} is S^1 -invariant, and each subset X_I^{reg} is S^1 -invariant.

Similarly, we define an S^1 -equivariant implicit atlas with boundary and/or stratification by in addition requiring that the boundary loci ∂X_I and/or stratifications $X_I \rightarrow \mathcal{S}$ be S^1 -invariant.

Note that in the above definition S^1 does not act on any of the obstruction spaces E_{α} (or, alternatively, it acts trivially on them).

8.3 S^1 -equivariant orientations

We begin with the trivial observation that if X is equipped with a locally orientable S^1 -equivariant implicit atlas A , then the S^1 -action on X lifts canonically to the orientation sheaf \mathfrak{o}_X .

Definition 8.3.1 (Locally S^1 -orientable implicit atlas) Let X be an S^1 -space with S^1 -equivariant implicit atlas with boundary. We say that A is *locally S^1 -orientable* if and only if it is locally orientable and for all $p \in X$, the stabilizer $(S^1)_p$ acts trivially on $(\mathfrak{o}_X)_p$ (this action is always by a sign $(S^1)_p \rightarrow \{\pm 1\}$). This notion is independent of the ring R (due to our restriction that $\#(S^1)_p$ be invertible in R).

It is easy to see that if A is locally S^1 -orientable, then $\pi_*\mathfrak{o}_X$ is locally isomorphic to the constant sheaf \underline{R} (where $\pi: X \rightarrow X/S^1$).

Remark 8.3.2 If A is locally orientable and \mathfrak{o}_X has a global section, then A is automatically locally S^1 -orientable (S^1 is connected, so the section must be S^1 -invariant, hence the claim).

8.4 S^1 -equivariant virtual cochain complexes $C_{S^1, \text{vir}}^\bullet(X; A)$ (also $\text{rel } \partial$)

To define the S^1 -equivariant virtual cochain complexes, we must first fix a model $C_{\bullet}^{S^1}$ of S^1 -equivariant chains to work with.

Remark 8.4.1 If we used the language of ∞ -categories, there would be no need to construct models of chains with good (chain-level) functoriality properties (see Remark 4.2.3).

We begin by stating the (chain-level) properties we would like our model $C_{\bullet}^{S^1}$ to satisfy. We want a functor $C_{\bullet}^{S^1}$ from spaces to chain complexes of free \mathbb{Z} -modules (and then we can tensor up to any base ring R); we also demand that $C_{\bullet}^{S^1}(A) \rightarrow C_{\bullet}^{S^1}(X)$ be injective for $A \subseteq X$ (and then we define relative chains $C_{\bullet}^{S^1}(X, A)$ as the cokernel). Now we need there to be functorial maps

$$(8.4.1) \quad C_{\bullet}^{S^1}(X) \xrightarrow{\pi^!} C_{\bullet+1}(X),$$

$$(8.4.2) \quad C_{\bullet}^{S^1}(X) \otimes C_{\bullet}(Y) \rightarrow C_{\bullet}^{S^1}(X \times Y),$$

realizing (respectively) the Gysin map and the obvious product map. The map (8.4.2) must be compatible with the Eilenberg–Zilber map on C_{\bullet} in that the two ways of building up the map

$$(8.4.3) \quad C_{\bullet}^{S^1}(X) \otimes C_{\bullet}(Y) \otimes C_{\bullet}(Z) \rightarrow C_{\bullet}^{S^1}(X \times Y \times Z)$$

coincide. Moreover, (8.4.1) and (8.4.2) must be compatible in that the diagram

$$(8.4.4) \quad \begin{array}{ccc} C_{\bullet}^{S^1}(X) \otimes C_{\bullet}(Y) & \longrightarrow & C_{\bullet}^{S^1}(X \times Y) \\ \downarrow & & \downarrow \\ C_{\bullet+1}(X) \otimes C_{\bullet}(Y) & \longrightarrow & C_{\bullet+1}(X \times Y) \end{array}$$

commutes. To define $C_{\bullet}^{S^1}$ with the aforementioned properties, let us recall the construction of the Serre spectral sequence due to Dress [17]. For a Serre fibration $\pi: E \rightarrow B$, we consider diagrams of the form:

$$(8.4.5) \quad \begin{array}{ccc} \Delta^p \times \Delta^q & \longrightarrow & E \\ \downarrow & & \downarrow \pi \\ \Delta^p & \longrightarrow & B \end{array}$$

Let $C_{p,q}(\pi: E \rightarrow B)$ denote the free abelian group generated by such diagrams. Then the direct sum of all of these $C_{\bullet,\bullet}(\pi: E \rightarrow B)$ is a double complex (differentials corresponding to the two pieces of boundary $\partial\Delta^p \times \Delta^q$ and $\Delta^p \times \partial\Delta^q$). We let $C_{\bullet}(\pi: E \rightarrow B)$ denote the corresponding total complex. There is a natural map $C_{\bullet}(\pi: E \rightarrow B) \rightarrow C_{\bullet}(E)$, where we subdivide $\Delta^p \times \Delta^q$ in the usual way. By considering the spectral sequence associated to the filtration by q , Dress showed that this map is a quasi-isomorphism. Dress also showed that the spectral sequence associated to the filtration by p is the Serre spectral sequence; in particular, the $E_{p,q}^2$ term is $H_p(B, H_q(F))$.

To define $C_{\bullet}^{S^1}$, fix once and for all an ES^1 . Then for any space X , the map

$$\pi: X \times ES^1 \rightarrow (X \times ES^1)/S^1$$

is a principal S^1 -bundle, so *a fortiori* it is a Serre fibration. We define

$$(8.4.6) \quad C_{\bullet-1}^{S^1}(X) := \{ \gamma \in C_{\bullet}(\pi: X \times ES^1 \rightarrow (X \times ES^1)/S^1) \mid \gamma, d\gamma \text{ have no component with } q\text{-degree} < 1 \}.$$

The inclusion $C_{\bullet-1}^{S^1}(X) \hookrightarrow C_{\bullet}(\pi: X \times ES^1 \rightarrow (X \times ES^1)/S^1)$ is compatible with the p -grading on each. Let us consider the associated morphism of spectral sequences (induced by the p -filtration). The latter has $E_{p,q}^2$ term $H_p^{S^1}(X, H_q(S^1))$ by Dress [17]. It follows from the definition (8.4.6) that the $E_{p,q}^2$ term of the former is the same in degrees $q \geq 1$ and zero otherwise. Since $H_q(S^1)$ is nonzero only for $q \leq 1$, it follows that the former spectral sequence has no further differentials, and we conclude that the homology of $C_{\bullet}^{S^1}(X)$ is indeed $H_{\bullet}^{S^1}(X)$ as needed.

Now let us define the maps (8.4.1) and (8.4.2) and verify the required properties. The map (8.4.1) is obtained by the standard subdivision of $\Delta^p \times \Delta^q$ into simplices along with the projection map $X \times ES^1 \rightarrow X$. The map (8.4.2) is defined as follows. Given maps $\Delta^p \times \Delta^q \rightarrow X \times ES^1$ and $\Delta^{p'} \rightarrow Y$, we obtain a map $\Delta^{p'} \times \Delta^p \times \Delta^q \rightarrow Y \times X \times ES^1$ and we subdivide $\Delta^{p'} \times \Delta^p$. The required properties are then straightforward to verify.

Definition 8.4.2 (S^1 -equivariant virtual cochain complexes $C_{S^1, \text{vir}}^\bullet(-; A)$ (and IJ))
 Let X be an S^1 -space with finite S^1 -equivariant implicit atlas with boundary A . For any S^1 -invariant compact $K \subseteq X$, we define

$$(8.4.7) \quad C_{S^1, \text{vir}}^\bullet(K; A)_{IJ}, \quad C_{S^1, \text{vir}}^\bullet(K \text{ rel } \partial; A)_{IJ},$$

$$(8.4.8) \quad C_{S^1, \text{vir}}^\bullet(K; A), \quad C_{S^1, \text{vir}}^\bullet(K \text{ rel } \partial; A)$$

as in Definitions 4.2.5–4.2.6, except using $C_{-\bullet-1}^{S^1}$ in place of C_\bullet in (4.2.4)–(4.2.5). It is clear that (8.4.7)–(8.4.8) are complexes of \mathcal{K} -presheaves on X/S^1 (replace K with $\pi^{-1}(K)$, where $\pi: X \rightarrow X/S^1$). The Gysin map (8.4.1) induces “comparison maps”

$$(8.4.9) \quad C_{S^1, \text{vir}}^\bullet(-; A)_{(IJ)} \rightarrow C_{\text{vir}}^\bullet(-; A)_{(IJ)}$$

for all flavors (8.4.7)–(8.4.8)

Analogously with (4.2.6)–(4.2.8) and (4.2.11)–(4.2.12), there are natural maps (compatible with (8.4.9))

$$(8.4.10) \quad C_{S^1, \text{vir}}^\bullet(K \text{ rel } \partial; A)_{IJ} \rightarrow C_{S^1, \text{vir}}^\bullet(K; A)_{IJ},$$

$$(8.4.11) \quad C_{S^1, \text{vir}}^{d+\bullet}(X \text{ rel } \partial; A)_{IJ} \xrightarrow{S^*} C_{-\bullet-1}^{S^1}(E; A),$$

$$(8.4.12) \quad C_{S^1, \text{vir}}^\bullet(-; A)_{IJ} \rightarrow C_{S^1, \text{vir}}^\bullet(-; A)_{I', J'},$$

$$(8.4.13) \quad C_{S^1, \text{vir}}^\bullet(K \text{ rel } \partial; A) \rightarrow C_{S^1, \text{vir}}^\bullet(K; A),$$

$$(8.4.14) \quad C_{S^1, \text{vir}}^{d+\bullet}(X \text{ rel } \partial; A) \xrightarrow{S^*} C_{-\bullet-1}^{S^1}(E; A).$$

Analogously with (4.2.13)–(4.2.14), there are natural maps (compatible with (8.4.9))

$$(8.4.15) \quad C_{S^1, \text{vir}}^\bullet(-; A)_{IJ} \otimes C_{-\bullet}(E; A' \setminus A) \rightarrow C_{S^1, \text{vir}}^\bullet(-; A')_{I', J'},$$

$$(8.4.16) \quad C_{S^1, \text{vir}}^\bullet(-; A) \otimes C_{-\bullet}(E; A' \setminus A) \rightarrow C_{S^1, \text{vir}}^\bullet(-; A').$$

8.5 Isomorphisms $H_{S^1, \text{vir}}^\bullet(X; A) = \check{H}^\bullet(X/S^1; \pi_* \circ_X)$ (also rel ∂)

Lemma 8.5.1 ($C_{S^1, \text{vir}}^\bullet(-; A)_{IJ}$ are pure homotopy \mathcal{K} -sheaves) *Let X be an almost free S^1 -space with finite locally S^1 -orientable S^1 -equivariant implicit atlas with*

boundary A . Then $C_{S^1, \text{vir}}^\bullet(-; A)_{IJ}$ and $C_{S^1, \text{vir}}^\bullet(-\text{rel } \partial; A)_{IJ}$ are pure homotopy \mathcal{K} -sheaves on X/S^1 . Furthermore, there are canonical isomorphisms

$$(8.5.1) \quad H_{S^1, \text{vir}}^0(-; A)_{IJ} = \pi_* j_! j^* \circ_X,$$

$$(8.5.2) \quad H_{S^1, \text{vir}}^0(-\text{rel } \partial; A)_{IJ} = \pi_* j_! j^* \circ_{X \text{ rel } \partial}$$

of sheaves on X/S^1 , where $j: V_I \cap V_J \hookrightarrow X$ for $V_I := \psi_{\emptyset I}((s_I|X_I^{\text{reg}})^{-1}(0)) \subseteq X$ and $\pi: X \rightarrow X/S^1$.

Proof As in the proof of Lemma 4.3.1, we use Lemmas A.6.3 and A.2.11 to see that $C_{S^1, \text{vir}}^\bullet(-; A)_{IJ}$ and $C_{S^1, \text{vir}}^\bullet(-\text{rel } \partial; A)_{IJ}$ are homotopy \mathcal{K} -sheaves. More precisely, in the present context we need the statement of Lemma A.6.3 for $C_{\bullet}^{S^1}$ in place of C_{\bullet} . The proof of this lemma applies equally well to $C_{\bullet}^{S^1}$, noting that since $C_{\bullet}^{S^1}$ is a model of S^1 -equivariant chains, it in particular satisfies Mayer–Vietoris in that the total complex

$$(8.5.3) \quad C_{\bullet}^{S^1}(U \cap U') \rightarrow C_{\bullet}^{S^1}(U) \oplus C_{\bullet}^{S^1}(U') \rightarrow C_{\bullet}^{S^1}(U \cup U')$$

is acyclic. Now let us show purity. By Lemma A.5.5, it suffices to show that the restrictions of $C_{S^1, \text{vir}}^\bullet(-; A)_{IJ}$ and $C_{S^1, \text{vir}}^\bullet(-\text{rel } \partial; A)_{IJ}$ to $(V_I \cap V_J)/S^1$ and to $X/S^1 \setminus (V_I \cap V_J)/S^1$ are pure. The latter restriction is trivially pure, since both complexes are simply zero for $K \cap V_I \cap V_J = \emptyset$. Hence it suffices to show that the restrictions of both complexes to $(V_I \cap V_J)/S^1$ are pure.

We consider the comparison maps

$$(8.5.4) \quad H_{S^1, \text{vir}}^\bullet(K; A)_{IJ} \rightarrow H_{\text{vir}}^\bullet(K; A)_{IJ},$$

$$(8.5.5) \quad H_{S^1, \text{vir}}^\bullet(K \text{ rel } \partial; A)_{IJ} \rightarrow H_{\text{vir}}^\bullet(K \text{ rel } \partial; A)_{IJ}.$$

These are Gysin maps $\pi^!$, and hence by the Gysin sequence their kernels coincide (respectively) with the images of

$$(8.5.6) \quad H_{S^1, \text{vir}}^{\bullet-2}(K; A)_{IJ} \xrightarrow{\cap e} H_{S^1, \text{vir}}^\bullet(K; A)_{IJ},$$

$$(8.5.7) \quad H_{S^1, \text{vir}}^{\bullet-2}(K \text{ rel } \partial; A)_{IJ} \xrightarrow{\cap e} H_{S^1, \text{vir}}^\bullet(K \text{ rel } \partial; A)_{IJ},$$

where e is (the pullback of) $e \in H_{S^1}^2(X_J^{\text{reg}}/\Gamma_J)$. Pick any $p \in V_J \subseteq X$. By Lemma 8.1.3, this e vanishes when restricted to small S^1 -invariant compact neighborhoods of

$$S^1 p \subseteq V_J = (s_J|X_J^{\text{reg}})^{-1}(0)/\Gamma_J \subseteq X_J^{\text{reg}}/\Gamma_J.$$

It follows that (8.5.6)–(8.5.7) vanish for small S^1 -invariant compact neighborhoods K of $S^1 p \subseteq V_J$. Hence (8.5.4)–(8.5.5) are injective for such K . Now, using this injectivity and the fact that $C_{\text{vir}}^\bullet(-; A)_{IJ}$ and $C_{\text{vir}}^\bullet(-\text{rel } \partial; A)_{IJ}$ are pure homotopy \mathcal{K} -sheaves (Lemma 4.3.1), it follows that $C_{S^1, \text{vir}}^\bullet(-; A)_{IJ}$ and $C_{S^1, \text{vir}}^\bullet(-\text{rel } \partial; A)_{IJ}$ are pure on V_J/S^1 , and hence on all of X/S^1 .

It remains to identify $H_{S^1, \text{vir}}^0(-; A)_{IJ}$ and $H_{S^1, \text{vir}}^0(-\text{rel } \partial; A)_{IJ}$. Consider the comparison maps

$$(8.5.8) \quad H_{S^1, \text{vir}}^0(-; A)_{IJ} \rightarrow \pi_* H_{\text{vir}}^0(-; A)_{IJ},$$

$$(8.5.9) \quad H_{S^1, \text{vir}}^0(-\text{rel } \partial; A)_{IJ} \rightarrow \pi_* H_{\text{vir}}^0(-\text{rel } \partial; A)_{IJ},$$

which are maps of sheaves on X/S^1 . It suffices (by Lemma 4.3.1) to show that these are isomorphisms (which we will check on stalks, ie for $K = S^1 p$). Now this is just a calculation, similar to that in the proof of Lemma 4.3.1, except using an S^1 -equivariant version of Poincaré–Lefschetz duality based on Lemma 8.1.4 in place of Lemma A.6.4 (and it is in this calculation where we use the local S^1 -orientability of A). \square

Proposition 8.5.2 ($C_{S^1, \text{vir}}^\bullet(-; A)$ are pure homotopy \mathcal{K} -sheaves) *Let X be an almost free S^1 -space with finite locally S^1 -orientable S^1 -equivariant implicit atlas with boundary A . Then $C_{S^1, \text{vir}}^\bullet(-; A)$ and $C_{S^1, \text{vir}}^\bullet(-\text{rel } \partial; A)$ are pure homotopy \mathcal{K} -sheaves on X/S^1 . Furthermore, there are canonical isomorphisms*

$$(8.5.10) \quad H_{S^1, \text{vir}}^0(-; A) = \pi_* \circ_X,$$

$$(8.5.11) \quad H_{S^1, \text{vir}}^0(-\text{rel } \partial; A) = \pi_* \circ_{X \text{ rel } \partial}$$

of sheaves on X/S^1 , where $\pi: X \rightarrow X/S^1$.

Proof This is exactly analogous to Proposition 4.3.3 and has the same proof (using Lemma 8.5.1 in place of Lemma 4.3.1). \square

Theorem 8.5.3 (Calculation of $H_{S^1, \text{vir}}^\bullet$) *Let X be an almost free S^1 -space with finite locally S^1 -orientable S^1 -equivariant implicit atlas with boundary A . Then there are canonical isomorphisms fitting (as the top horizontal maps) into commutative diagrams:*

$$\begin{array}{ccc} H_{S^1, \text{vir}}^\bullet(X; A) & \xrightarrow{\sim} & \check{H}^\bullet(X/S^1, \pi_* \circ_X) \\ (8.4.9) \downarrow & & \pi_* \downarrow \\ H_{\text{vir}}^\bullet(X; A) & \xrightarrow{\text{Thm 4.3.4}} & \check{H}^\bullet(X, \circ_X) \end{array}$$

and

$$\begin{array}{ccc}
 H_{S^1, \text{vir}}^\bullet(X \text{ rel } \partial; A) & \xrightarrow{\sim} & \check{H}^\bullet(X/S^1, \pi_* \circ_X \text{ rel } \partial) \\
 \downarrow (8.4.9) & & \downarrow \pi^* \\
 H_{\text{vir}}^\bullet(X \text{ rel } \partial; A) & \xrightarrow{\text{Thm 4.3.4}} & \check{H}^\bullet(X, \circ_X \text{ rel } \partial)
 \end{array}$$

Proof Consider the following diagram:

$$\begin{array}{ccccc}
 H_{S^1, \text{vir}}^\bullet(X; A) & \xrightarrow{(8.4.9)} & & & H_{\text{vir}}^\bullet(X; A) \\
 \downarrow \sim & & & & \downarrow \sim \\
 \check{H}^\bullet(X/S^1; C_{S^1, \text{vir}}^\bullet(-; A)) & \xrightarrow{(8.4.9)} & \check{H}^\bullet(X/S^1, \pi_* C_{\text{vir}}^\bullet(-; A)) & \xrightarrow{\pi^*} & \check{H}^\bullet(X; C_{\text{vir}}^\bullet(-; A)) \\
 \downarrow \sim & & \downarrow \sim & & \downarrow \sim \\
 \check{H}^\bullet(X/S^1; \tau_{\geq 0} C_{S^1, \text{vir}}^\bullet(-; A)) & \xrightarrow{(8.4.9)} & \check{H}^\bullet(X/S^1, \pi_* \tau_{\geq 0} C_{\text{vir}}^\bullet(-; A)) & \xrightarrow{\pi^*} & \check{H}^\bullet(X; \tau_{\geq 0} C_{\text{vir}}^\bullet(-; A)) \\
 \uparrow \sim & & \uparrow & & \uparrow \sim \\
 \check{H}^\bullet(X/S^1; H_{S^1, \text{vir}}^0(-; A)) & \xrightarrow{(8.4.9)} & \check{H}^\bullet(X/S^1, \pi_* H_{\text{vir}}^0(-; A)) & \xrightarrow{\pi^*} & \check{H}^\bullet(X; H_{\text{vir}}^0(-; A)) \\
 \downarrow \sim & & \downarrow \sim & & \downarrow \sim \\
 \check{H}^\bullet(X/S^1; \pi_* \circ_X) & \xlongequal{\quad} & \check{H}^\bullet(X/S^1, \pi_* \circ_X) & \xrightarrow{\pi^*} & \check{H}^\bullet(X; \circ_X)
 \end{array}$$

where the vertical isomorphisms are by Propositions 8.5.2 and 4.3.3 (see also (A.5.3)).

Now each small square of this diagram commutes (only the bottom left requires an argument; one shows that the corresponding diagram of sheaves commutes by checking it locally, where it is just a calculation). Now the first square in Theorem 8.5.3 is the same as the big square above, which commutes by a diagram chase (for which one should be careful of the fact that one of the vertical arrows is not an isomorphism).

An identical argument applies to the second square in Theorem 8.5.3. □

8.6 Localization for virtual fundamental classes

Theorem 8.6.1 (S^1 -localization for $[X]^{\text{vir}}$) *Let X be an almost free S^1 -space with locally S^1 -orientable S^1 -equivariant implicit atlas with boundary A . Then $\pi_* [X]^{\text{vir}} = 0$, where $\pi_* = (\pi^*)^\vee: \check{H}^\bullet(X; \circ_X \text{ rel } \partial)^\vee \rightarrow \check{H}^\bullet(X/S^1; \pi_* \circ_X \text{ rel } \partial)^\vee$.*

Proof We may assume that A is finite. Now by Theorem 8.5.3, we have a commutative diagram:

$$\begin{array}{ccccc}
 \check{H}^{d+\bullet}(X/S^1; \pi_* \mathfrak{o}_{X \text{ rel } \partial}) & \xrightarrow{\text{Thm 8.5.3}} & H_{S^1, \text{vir}}^{d+\bullet}(X \text{ rel } \partial; A) & \xrightarrow{(8.4.14)} & H_{-\bullet-1}^{S^1}(E; A) \\
 \downarrow \pi^* & & \downarrow (8.4.9) & & \downarrow \pi^! \\
 \check{H}^{d+\bullet}(X; \mathfrak{o}_{X \text{ rel } \partial}) & \xrightarrow{\text{Thm 4.3.4}} & H_{\text{vir}}^{d+\bullet}(X \text{ rel } \partial; A) & \xrightarrow{(4.2.12)} & H_{-\bullet}(E; A) \xrightarrow{[E_A] \mapsto 1} R
 \end{array}$$

By definition, $[X]^{\text{vir}}$ is the total composition of the bottom row for $\bullet = 0$. Now the desired statement follows since the rightmost vertical map is zero (by Lemma 8.1.2, because S^1 acts trivially on E_A). \square

Corollary 8.6.2 *Let X be an almost free S^1 -space with locally S^1 -orientable S^1 -equivariant implicit atlas with boundary A of dimension 0. Then $[X]^{\text{vir}} = 0$.*

Proof This follows from Theorem 8.6.1 since the map $\pi^*: \check{H}^0(X/S^1; \pi_* \mathfrak{o}_{X \text{ rel } \partial}) \rightarrow \check{H}^0(X; \mathfrak{o}_{X \text{ rel } \partial})$ is an isomorphism (since S^1 is connected and the atlas A is locally S^1 -orientable). \square

Remark 8.6.3 We expect that the machinery of this section can be used to define an S^1 -equivariant virtual fundamental cycle

$$[X]^{S^1, \text{vir}} \in \check{H}^{d-1}(X/S^1; \pi_* \mathfrak{o}_{X \text{ rel } \partial})^\vee \xrightarrow{(\pi_!)^\vee} \check{H}^d(X; \mathfrak{o}_{X \text{ rel } \partial})^\vee$$

lifting $[X]^{\text{vir}}$, via the diagram

$$\begin{array}{ccccc}
 \check{H}^{d+\bullet}(X; \mathfrak{o}_{X \text{ rel } \partial}) & \xrightarrow{\text{Thm 4.3.4}} & H_{\text{vir}}^{d+\bullet}(X \text{ rel } \partial; A) & \xrightarrow{(4.2.12)} & H_{-\bullet}(E; A) \xrightarrow{[E_A] \mapsto 1} R \\
 \downarrow \pi_! & & \downarrow & & \downarrow \pi_* \parallel \\
 \check{H}^{d+\bullet-1}(X/S^1; \pi_* \mathfrak{o}_{X \text{ rel } \partial}) & \xrightarrow{\text{Thm 8.5.3}} & H_{S^1, \text{vir}}^{d+\bullet-1}(X \text{ rel } \partial; A) & \xrightarrow{(8.4.14)} & H_{-\bullet}^{S^1}(E; A) \xrightarrow{[E_A] \mapsto 1} R
 \end{array}$$

where the second vertical map is induced by pushforward π_* on chains. Note that we would need to show that the first square commutes. This would provide another proof of Theorem 8.6.1. This should also be applicable without any restriction on the S^1 -action on X (provided we use S^1 -equivariant (co)homology in the appropriate places).

8.7 Localization for homology

To prove the desired localization result for Floer-type homology groups, the chain models used in Section 8.4 are inadequate. Specifically, the fact that we need to consider

product maps $\mathcal{X}(p, q) \times \mathcal{X}(q, r) \rightarrow \mathcal{X}(p, r)$ between flow spaces with an S^1 -action means that we need a corresponding map $C_{\bullet}^{S^1}(X) \times C_{\bullet}^{S^1}(Y) \rightarrow C_{\bullet+1}^{S^1}(X \times Y)$ on S^1 -equivariant chains. To obtain such a map, we will end up using different models of chains for every (σ, p, q) . We first introduce the models of chains we will use, we then describe how to use these chain models for the constructions of Section 7, and finally we prove the localization result for homology groups.

Remark 8.7.1 If we used the language of ∞ -categories, there would be no need to construct models of chains with good (chain-level) functoriality properties (see Remark 4.2.3).

8.7.1 MF-sets and S^1 -MF-sets

Definition 8.7.2 A PL manifold with cells (M, \mathcal{S}) is a compact connected nonempty orientable PL-manifold with boundary M with stratification by a finite poset \mathcal{S} (see Definition 6.1.1) such that:

- (i) $M^{\leq \mathfrak{s}}$ is a connected nonempty orientable PL-submanifold with boundary.
- (ii) $(M^{\leq \mathfrak{s}})^{\circ} = M^{\mathfrak{s}}$.
- (iii) \mathcal{S} has a unique maximal element.

For every $\mathfrak{s} \in \mathcal{S}$, there corresponds a face $(M^{\leq \mathfrak{s}}, \mathcal{S}^{\leq \mathfrak{s}})$, which is also a PL manifold with cells.

Definition 8.7.3 An MF-set is a set \mathcal{Y} along with:

- (i) For every $i \in \mathcal{Y}$, a PL manifold with cells (M_i, \mathcal{S}_i) .
- (ii) **Face identifications** For every $i \in \mathcal{Y}$ and every $\mathfrak{s} \in \mathcal{S}_i$, an index $j \in \mathcal{Y}$ and an isomorphism $(M_i^{\leq \mathfrak{s}}, \mathcal{S}_i^{\leq \mathfrak{s}}) \xrightarrow{\sim} (M_j, \mathcal{S}_j)$. These indices and isomorphisms must be (strictly) transitive in the following obvious sense. If $\mathfrak{s} \in \mathcal{S}$ is maximal, then $j = i$ and

$$(M_i, \mathcal{S}_i) = (M_i^{\leq \mathfrak{s}}, \mathcal{S}_i^{\leq \mathfrak{s}}) \xrightarrow{\sim} (M_j, \mathcal{S}_j) = (M_i, \mathcal{S}_i)$$

is the identity map. If $\mathfrak{s}' \preceq \mathfrak{s}$ is a nested pair of faces with identifications

$$(M_i^{\leq \mathfrak{s}}, \mathcal{S}_i^{\leq \mathfrak{s}}) \xrightarrow{\sim} (M_j, \mathcal{S}_j) \quad \text{and} \quad (M_i^{\leq \mathfrak{s}'}, \mathcal{S}_i^{\leq \mathfrak{s}'}) \xrightarrow{\sim} (M_{j'}, \mathcal{S}_{j'}),$$

and the \mathfrak{s}' face of (M_j, \mathcal{S}_j) is identified

$$(M_j^{\leq \mathfrak{s}'}, \mathcal{S}_j^{\leq \mathfrak{s}'}) \xrightarrow{\sim} (M_k, \mathcal{S}_k),$$

then $k = j'$ and the two identifications of $(M_i^{\leq \mathfrak{s}'}, \mathcal{S}_i^{\leq \mathfrak{s}'})$ with $(M_k, \mathcal{S}_k) = (M_{j'}, \mathcal{S}_{j'})$ are the same.

A morphism of MF-sets $f: \mathcal{Y} \rightarrow \mathcal{Y}'$ is a map of sets covered by isomorphisms $(M_i, \mathcal{S}_i) \rightarrow (M_{f(i)}, \mathcal{S}_{f(i)})$, compatible with the face identifications for \mathcal{Y} and \mathcal{Y}' in the obvious way.

The category of MF-sets has a natural symmetric monoidal structure: given MF-sets $\{M_i\}_{i \in \mathcal{Y}}$ and $\{M'_j\}_{j \in \mathcal{Y}'}$, their product is defined to be $\{M_i \times M'_j\}_{(i,j) \in \mathcal{Y} \times \mathcal{Y}'}$, which is again an MF-set.

For an MF-set \mathcal{Y} , let $C_{\bullet}^{\mathcal{Y}}(X)$ denote the complex freely generated²⁷ by maps $M_i \rightarrow X$ (for $i \in \mathcal{Y}$), with differential given by the obvious sum over all codimension-one faces. Note that there is a natural map

$$(8.7.1) \quad C_{\bullet}^{\mathcal{Y}}(X) \otimes C_{\bullet}^{\mathcal{Y}'}(X') \rightarrow C_{\bullet}^{\mathcal{Y} \times \mathcal{Y}'}(X \times X').$$

Remark 8.7.4 It would perhaps be more natural to work with MF-sets in the DIFF category, although in that case it is not clear precisely what sort of stratifications and corner structure one should allow so that the proof of Lemma 8.7.7 goes through.

Example 8.7.5 The collection of standard simplices $\{\Delta^n\}_{n \geq 0}$ equipped with their standard simplicial stratifications and the standard identifications of the faces of Δ^n with the various Δ^i , forms an MF-set.

Definition 8.7.6 An MF-set \mathcal{Y} is called *saturated* if and only if for every PL manifold with cells (M, \mathcal{S}) along with, for every $s \in \mathcal{S}$ of positive codimension, an index $j \in \mathcal{Y}$ and an isomorphism $(M_j, \mathcal{S}_j) \xrightarrow{\sim} (M^{\leq s}, \mathcal{S}^{\leq s})$, such that these indices and isomorphisms are (strictly) transitive in the obvious sense, there exists $i \in \mathcal{Y}$ and an isomorphism $(M_i, \mathcal{S}_i) \xrightarrow{\sim} (M, \mathcal{S})$ respecting these given face identifications.

Every MF-set \mathcal{Y} embeds into a saturated MF-set \mathcal{Y}_{∞} , which may be constructed (non-canonically) as follows. We define a sequence of inclusions $\mathcal{Y} = \mathcal{Y}_{-1} \hookrightarrow \mathcal{Y}_0 \hookrightarrow \mathcal{Y}_1 \hookrightarrow \dots$, and we let $\mathcal{Y}_{\infty} := \varinjlim \mathcal{Y}_n$. To define \mathcal{Y}_n , we consider all PL manifolds with cells of dimension n with face identifications to elements of \mathcal{Y}_{n-1} as in Definition 8.7.6. The collection of isomorphism classes of such data (manifold along with face identifications) forms a set, so we may choose (non-canonically) a set \mathcal{Z}_n parametrizing all of them. We then set $\mathcal{Y}_n := \mathcal{Y}_{n-1} \sqcup \mathcal{Z}_n$. Now it is easy to check that \mathcal{Y}_{∞} is saturated.

Lemma 8.7.7 *Let \mathcal{Y} be a saturated MF-set. Then there is a canonical isomorphism between singular homology $H_{\bullet}(X)$ and $H_{\bullet}^{\mathcal{Y}}(X)$ (the homology of $C_{\bullet}^{\mathcal{Y}}(X)$).*

²⁷A given map $M_i \rightarrow X$ contributes a copy of the orientation module of M_i , which is isomorphic to \mathbb{Z} but not canonically so.

Proof Since \mathcal{Y} is saturated, there exists a morphism of MF-sets $\{\Delta^n\}_{n \geq 0} \rightarrow \{M_i\}_{i \in \mathcal{Y}}$ (where $\{\Delta^n\}_{n \geq 0}$ is as in Example 8.7.5). Since

$$C_\bullet(X) = C_\bullet^{\{\Delta^n\}_{n \geq 0}}(X)$$

by definition, we obtain a chain map $C_\bullet(X) \rightarrow C_\bullet^{\mathcal{Y}}(X)$. The resulting map

$$(8.7.2) \quad H_\bullet(X) \rightarrow H_\bullet^{\mathcal{Y}}(X)$$

is independent of the choice of morphism $\{\Delta^n\}_{n \geq 0} \rightarrow \{M_i\}_{i \in \mathcal{Y}}$, as can be seen as follows. There is an MF-set

$$\{\Delta^n\}_{n \geq 0} \sqcup \{\Delta^n \times [0, 1]\}_{n \geq 0} \sqcup \{\Delta^n\}_{n \geq 0}$$

where each $\{\Delta^n\}_{n \geq 0}$ is as in Example 8.7.5, where $\Delta^n \times [0, 1]$ is given the product stratification ($[0, 1]$ stratified by $\{\{0\}, \{1\}, (0, 1)\}$), and $\Delta^n \times \{0\}$ (resp. $\Delta^n \times \{1\}$) is identified with the first (resp. second) copy of Δ^n . Since \mathcal{Y} is saturated, it follows that for any pair of morphisms $\{\Delta^n\}_{n \geq 0} \rightarrow \{M_i\}_{i \in \mathcal{Y}}$ there is a morphism

$$\{\Delta^n\}_{n \geq 0} \sqcup \{\Delta^n \times [0, 1]\}_{n \geq 0} \sqcup \{\Delta^n\}_{n \geq 0} \rightarrow \{M_i\}_{i \in \mathcal{Y}}$$

whose restriction to the two copies of $\{\Delta^n\}_{n \geq 0}$ are the two given morphisms. From this data one easily constructs a chain homotopy between the two maps $C_\bullet(X) \rightarrow C_\bullet^{\mathcal{Y}}(X)$. Hence the map (8.7.2) is canonical.

Now let us show that (8.7.2) is an isomorphism. Fix a map $\{\Delta^n\}_{n \geq 0} \rightarrow \{M_i\}_{i \in \mathcal{Y}}$. Fix triangulations T_i of M_i for which each $M_i^{\leq s}$ is a union of simplices, which are compatible with the face identifications, and which restrict to the tautological triangulation of $\{\Delta^n\}_{n \geq 0}$. Such triangulations may be constructed by induction. By triangulation, we mean a triangulation in which each simplex is equipped with a total order on its set of vertices, compatible with its faces (ie a semisimplicial set rather than a simplicial complex). Such triangulations induce a map of complexes $C_\bullet^{\mathcal{Y}}(X) \rightarrow C_\bullet(X)$, and the composition $C_\bullet(X) \rightarrow C_\bullet^{\mathcal{Y}}(X) \rightarrow C_\bullet(X)$ is clearly the identity map. It suffices to show that the other composition is chain homotopic to the identity map.

Now let us define a new MF-set in terms of the map $\{\Delta^n\}_{n \geq 0} \rightarrow \{M_i\}_{i \in \mathcal{Y}}$ and the triangulations T_i . Note that the stratification of M_i by the face poset $\mathcal{F}(T_i)$ of T_i refines the stratification by \mathcal{S}_i ; in other words, we have maps $M_i \rightarrow \mathcal{F}(T_i) \rightarrow \mathcal{S}_i$. The objects of the new MF-set are $\{M_i\}_{i \in \mathcal{Y}} \sqcup \{M_i \times [0, 1]\}_{i \in \mathcal{Y}}$, where $M_i \times [0, 1]$ is

stratified by the following strata:

$$(8.7.3) \quad M_i^s \times \{0\}, \quad s \in \mathcal{S}_i,$$

$$(8.7.4) \quad M_i^s \times (0, 1), \quad s \in \mathcal{S}_i,$$

$$(8.7.5) \quad M_i^t \times \{1\}, \quad t \in \mathcal{F}(T_i).$$

Thus the poset of strata of $M_i \times [0, 1]$ is $\mathcal{S}_i \sqcup \mathcal{S}_i \sqcup \mathcal{F}(T_i)$. For the face identifications, $M_i \times \{0\}$ is identified tautologically with M_i , and each of the closed strata $M_i^{\leq t} \times \{1\}$ ($t \in \mathcal{F}(T_i)$) is identified with the corresponding Δ^n (considered as an object of $\{M_i\}_{i \in \mathcal{Y}}$ via the inclusion $\{\Delta^n\}_{n \geq 0} \rightarrow \{M_i\}_{i \in \mathcal{Y}}$).

Now we have a chain of inclusions

$$\{\Delta^n\}_{n \geq 0} \hookrightarrow \{M_i\}_{i \in \mathcal{Y}} \hookrightarrow \{M_i\}_{i \in \mathcal{Y}} \sqcup \{M_i \times [0, 1]\}_{i \in \mathcal{Y}}.$$

Since \mathcal{Y} is saturated, there in fact exists a map backwards

$$\{M_i\}_{i \in \mathcal{Y}} \sqcup \{M_i \times [0, 1]\}_{i \in \mathcal{Y}} \rightarrow \{M_i\}_{i \in \mathcal{Y}}$$

(acting identically on $\{M_i\}_{i \in \mathcal{Y}}$). Using this map, we obtain a chain map $C_{\bullet}^{\mathcal{Y}}(X) \rightarrow C_{\bullet+1}^{\mathcal{Y}}(X)$ (precompose chains with the projection $M_i \times [0, 1] \rightarrow M_i$). This map is a chain homotopy between the identity map and the composition $C_{\bullet}^{\mathcal{Y}}(X) \rightarrow C_{\bullet}(X) \rightarrow C_{\bullet}^{\mathcal{Y}}(X)$, so we are done. \square

Equip S^1 with its standard PL structure, for which the group operations are PL.

Definition 8.7.8 An S^1 -MF-set is an MF-set \mathcal{Y} along with a principal S^1 -bundle $(ES^1)_{\mathcal{Y}} \rightarrow (BS^1)_{\mathcal{Y}}$ where $(BS^1)_{\mathcal{Y}}$ (and hence also $(ES^1)_{\mathcal{Y}}$) is an increasing union of compact polyhedra (we do not require $(ES^1)_{\mathcal{Y}}$ to be contractible), and for every $i \in \mathcal{Y}$, a pullback diagram

$$(8.7.6) \quad \begin{array}{ccc} M_i & \longrightarrow & (ES^1)_{\mathcal{Y}} \\ \downarrow & & \downarrow \\ N_i & \longrightarrow & (BS^1)_{\mathcal{Y}} \end{array}$$

with PL maps where N_i (and hence also M_i) is a PL manifold, and where the stratification on M_i is pulled back from N_i , such that these diagrams are compatible with the face identifications. A map of S^1 -MF-sets $\mathcal{Y} \rightarrow \mathcal{Y}'$ is said to be injective if and only if it is injective as a map of sets and the map $(ES^1)_{\mathcal{Y}} \rightarrow (ES^1)_{\mathcal{Y}'}$ is injective.

There is a forgetful functor from S^1 -MF-sets to MF-sets, where we remember M_i (of course, there is another natural forgetful functor remembering N_i , though we will

never use it). When speaking of a morphism $\mathcal{Y} \rightarrow \mathcal{Y}'$ where \mathcal{Y} is an S^1 -MF-set and \mathcal{Y}' is an MF-set, we implicitly apply the forgetful functor to \mathcal{Y} .

The category of S^1 -MF-sets has a natural symmetric monoidal structure: given S^1 -MF-sets $\{M_i\}_{i \in \mathcal{Y}}$ and $\{M'_j\}_{j \in \mathcal{Y}'}$, we may define their product $\{M_i \times M'_j\}_{(i,j) \in \mathcal{Y} \times \mathcal{Y}'}$, which is again an S^1 -MF-set, via the diagonal S^1 -action on $M_i \times M'_j$ and

$$(ES^1)_{\mathcal{Y} \times \mathcal{Y}'} := (ES^1)_{\mathcal{Y}} \times (ES^1)_{\mathcal{Y}'}$$

with the diagonal action. The forgetful functor from S^1 -MF-sets to MF-sets is clearly a symmetric monoidal functor. It also makes sense to take the product of an S^1 -MF-set and an MF-set.

For an S^1 -MF-set \mathcal{Y} , let $C_{\bullet}^{S^1, \mathcal{Y}}(X)$ (for X an S^1 -space) denote the complex generated by commuting diagrams of S^1 -equivariant maps

$$(8.7.7) \quad \begin{array}{ccccc} M_i & \longrightarrow & X \times (ES^1)_{\mathcal{Y}} & \longrightarrow & (ES^1)_{\mathcal{Y}} \\ \downarrow & & \downarrow & & \downarrow \\ N_i & \longrightarrow & (X \times (ES^1)_{\mathcal{Y}})/S^1 & \longrightarrow & (BS^1)_{\mathcal{Y}} \end{array}$$

($i \in \mathcal{Y}$), where the outer square coincides with (8.7.6), with differential given by the obvious sum over all codimension-one faces. A generator (8.7.7) resides in degree $\dim N_i$. Note that there are natural compatible maps

$$(8.7.8) \quad C_{\bullet}^{S^1, \mathcal{Y}}(X') \rightarrow C_{\bullet+1}^{\mathcal{Y}}(X),$$

$$(8.7.9) \quad C_{\bullet}^{\mathcal{Y}}(X) \otimes C_{\bullet}^{S^1, \mathcal{Y}'}(X') \rightarrow C_{\bullet}^{S^1, \mathcal{Y} \times \mathcal{Y}'}(X \times X'),$$

$$(8.7.10) \quad C_{\bullet}^{S^1, \mathcal{Y}}(X) \otimes C_{\bullet}^{S^1, \mathcal{Y}'}(X') \rightarrow C_{\bullet+1}^{S^1, \mathcal{Y} \times \mathcal{Y}'}(X \times X').$$

Definition 8.7.9 An S^1 -MF-set \mathcal{Y} is called *saturated* if and only if $(ES^1)_{\mathcal{Y}}$ is contractible and every pullback diagram

$$(8.7.11) \quad \begin{array}{ccc} M & \longrightarrow & (ES^1)_{\mathcal{Y}} \\ \downarrow & & \downarrow \\ N & \longrightarrow & (BS^1)_{\mathcal{Y}} \end{array}$$

where N (and thus M) is a PL manifold with cells, along with strictly transitive identifications of the faces of positive codimension with elements of \mathcal{Y} (as in Definition 8.7.8), is isomorphic to some $i \in \mathcal{Y}$.

Note that the notions of saturation for MF-sets and S^1 -MF-sets are different: an S^1 -MF-set is *never* saturated as an MF-set.

Every S^1 -MF-set \mathcal{Y} embeds into a saturated S^1 -MF-set \mathcal{Y}_∞ , which may be constructed (again, non-canonically) by first embedding $(ES^1)_\mathcal{Y}$ into something contractible, and then proceeding as in the case of MF-sets.

Lemma 8.7.10 *Let \mathcal{Y} be a saturated S^1 -MF-set. Then there is a canonical isomorphism between S^1 -equivariant singular homology $H_\bullet^{S^1}(X)$ and $H_\bullet^{S^1, \mathcal{Y}}(X)$ (the homology of $C_\bullet^{S^1, \mathcal{Y}}(X)$). Furthermore, the natural map $C_\bullet^{S^1, \mathcal{Y}}(X) \rightarrow C_{\bullet+1}^{\mathcal{Y}'}(X)$ induces the Gysin map on homology, for \mathcal{Y}' a saturated MF-set.*

Proof All S^1 -MF-sets in this proof will share the same $(ES^1)_\mathcal{Y} \rightarrow (BS^1)_\mathcal{Y}$, so we will omit the subscript \mathcal{Y} from the notation.

Let $\{\Delta^n\}_{n \geq 0, f: \Delta^n \rightarrow BS^1}$ denote the S^1 -MF-set indexed by pairs (n, f) consisting of an integer $n \geq 0$ and a PL map $f: \Delta^n \rightarrow BS^1$, where the stratifications and face identifications are as in Example 8.7.5. Now the complex

$$C_\bullet^{S^1, \{\Delta^n\}_{n \geq 0, f: \Delta^n \rightarrow BS^1}}(X)$$

is freely generated by maps $\Delta^n \rightarrow (X \times ES^1)/S^1$ whose composition with the projection $(X \times ES^1)/S^1 \rightarrow BS^1$ is PL. A straightforward approximation argument shows that the inclusion of this complex into the complex generated by all maps $\Delta^n \rightarrow (X \times ES^1)/S^1$ is a quasi-isomorphism. Hence there is a canonical isomorphism

$$(8.7.12) \quad H_\bullet^{S^1, \{\Delta^n\}_{n \geq 0, f: \Delta^n \rightarrow BS^1}}(X) \simeq H_\bullet((X \times ES^1)/S^1) = H_\bullet^{S^1}(X).$$

As in the proof of Lemma 8.7.7, since \mathcal{Y} is saturated, there exists a morphism of S^1 -MF-sets $\{\Delta^n\}_{n \geq 0, f: \Delta^n \rightarrow BS^1} \rightarrow \mathcal{Y}$. This induces a chain map

$$C_\bullet^{S^1, \{\Delta^n\}_{n \geq 0, f: \Delta^n \rightarrow BS^1}}(X) \rightarrow C_\bullet^{S^1, \mathcal{Y}}(X),$$

which as before induces a canonical map on homology

$$(8.7.13) \quad H_\bullet^{S^1, \{\Delta^n\}_{n \geq 0, f: \Delta^n \rightarrow BS^1}}(X) \rightarrow H_\bullet^{S^1, \mathcal{Y}}(X),$$

which is independent of the choice of morphism $\{\Delta^n\}_{n \geq 0, f: \Delta^n \rightarrow BS^1} \rightarrow \mathcal{Y}$. It suffices to show that this map is an isomorphism.

Fix a map $\{\Delta^n\}_{n \geq 0, f: \Delta^n \rightarrow BS^1} \rightarrow \mathcal{Y}$, which we observe is necessarily injective. Fix triangulations T_i of M_i as in the proof of Lemma 8.7.7. Such triangulations induce a map of complexes

$$C_\bullet^{S^1, \mathcal{Y}}(X) \rightarrow C_\bullet^{S^1, \{\Delta^n\}_{n \geq 0, f: \Delta^n \rightarrow BS^1}}(X).$$

The composition

$$C_{\bullet}^{S^1, \{\Delta^n\}_{n \geq 0, f: \Delta^n \rightarrow BS^1}}(X) \rightarrow C_{\bullet}^{S^1, \mathcal{Y}}(X) \rightarrow C_{\bullet}^{S^1, \{\Delta^n\}_{n \geq 0, f: \Delta^n \rightarrow BS^1}}(X)$$

is clearly the identity map. It thus suffices to show that the reverse composition is chain homotopic to the identity.

As in the proof of Lemma 8.7.7, we construct a new S^1 -MF-set, namely

$$\{M_i\}_{i \in \mathcal{Y}} \sqcup \{M_i \times [0, 1]\}_{i \in \mathcal{Y}},$$

from the inclusion $\{\Delta^n\}_{n \geq 0, f: \Delta^n \rightarrow BS^1} \hookrightarrow \{M_i\}_{i \in \mathcal{Y}}$ and the triangulations T_i . There are inclusions

$$\{\Delta^n\}_{n \geq 0, f: \Delta^n \rightarrow BS^1} \hookrightarrow \{M_i\}_{i \in \mathcal{Y}} \hookrightarrow \{M_i\}_{i \in \mathcal{Y}} \sqcup \{M_i \times [0, 1]\}_{i \in \mathcal{Y}}.$$

The desired chain homotopy may then be constructed as in the proof of Lemma 8.7.7.

That the induced map $H_{\bullet}^{S^1}(X) \rightarrow H_{\bullet+1}(X)$ is the Gysin map is left to the reader. \square

8.7.2 F-modules valued in MF-sets and S^1 -MF-sets Plugging the category of MF-sets into Definition 7.1.4, we may talk about F-modules valued in MF-sets. In other words, an F-module MF-set is a collection of MF-sets $\mathcal{Y}(\sigma, p, q)$ along with product/face maps

$$(8.7.14) \quad \mathcal{Y}(\sigma|[0 \dots \hat{k} \dots n], p, q) \rightarrow \mathcal{Y}(\sigma, p, q),$$

$$(8.7.15) \quad \mathcal{Y}(\sigma|[0 \dots k], p, q) \times \mathcal{Y}(\sigma|[k \dots n], q, r) \rightarrow \mathcal{Y}(\sigma, p, r),$$

which are compatible in the sense of Definition 7.1.4. Similarly, we may talk about F-modules of S^1 -MF-sets. We may also talk about morphisms from F-module S^1 -MF-sets to F-module MF-sets using the forgetful functor described earlier.

Proposition 8.7.11 *Let \mathcal{X}/Z_{\bullet} be a flow category diagram. There exists an F-module of MF-sets \mathcal{Y} and an F-module of S^1 -MF-sets \mathcal{Y}_{S^1} (both supported inside $\text{supp } \mathcal{X}$) along with a morphism $\mathcal{Y}_{S^1} \rightarrow \mathcal{Y}$, satisfying the following property. For all $(\sigma, p, q) \in \text{supp } \mathcal{X}$, both $\mathcal{Y}(\sigma, p, q)$ and $\mathcal{Y}_{S^1}(\sigma, p, q)$ are saturated, and the tautologous maps*

$$(8.7.16) \quad \text{colim}_{\mathfrak{s} \in \partial \mathcal{S}_{\mathcal{X}}(\sigma, p, q)} \mathcal{Y}(\sigma, p, q, \mathfrak{s}) \hookrightarrow \mathcal{Y}(\sigma, p, q),$$

$$(8.7.17) \quad \text{colim}_{\mathfrak{s} \in \partial \mathcal{S}_{\mathcal{X}}(\sigma, p, q)} \mathcal{Y}_{S^1}(\sigma, p, q, \mathfrak{s}) \hookrightarrow \mathcal{Y}_{S^1}(\sigma, p, q)$$

are injective (and the colimits on the left exist). Moreover, such $\mathcal{Y}_{S^1} \rightarrow \mathcal{Y}$ may be constructed by induction on (σ, p, q) , partially ordered by $\preceq_{\mathcal{X}}$ (which is well-founded; see Definition 7.4.5 or the proof of Proposition 7.5.5).

Note the similarity of (8.7.16)–(8.7.17) with Definition 7.4.2. One should think of these conditions as being: “ \mathcal{Y} and \mathcal{Y}_{S^1} are both cofibrant”.

Proof It suffices to perform the inductive step for a given (σ, p, q) . First, we show that the colimits (8.7.16)–(8.7.17) exist, essentially by rewriting the proofs of Lemmas 7.4.3 and 7.4.4 in the present context. We actually show the more general statement that the colimits

$$(8.7.18) \quad \operatorname{colim}_{t \in \mathcal{T}} \mathcal{Y}(\sigma, p, q, t) \quad \text{and} \quad \operatorname{colim}_{t \in \mathcal{T}} \mathcal{Y}_{S^1}(\sigma, p, q, t)$$

exist for any $\mathcal{T} \subseteq \partial \mathcal{S}_X(\sigma, p, q)$ which is downward closed. We show this by induction on the cardinality of \mathcal{T} . If $\mathcal{T} = \emptyset$, existence is trivial. For \mathcal{T} nonempty, let $t_0 \in \mathcal{T}$ be any maximal element, and observe that (8.7.18) can be written as the colimit of the diagram:

$$(8.7.19) \quad \begin{array}{ccc} & & \mathcal{Y}(\sigma, p, q, t_0) \\ & \nearrow & \\ \operatorname{colim}_{\substack{t \in \mathcal{T} \\ t < t_0}} \mathcal{Y}(\sigma, p, q, t) & & \\ & \searrow & \\ & & \operatorname{colim}_{t \in \mathcal{T} \setminus t_0} \mathcal{Y}(\sigma, p, q, t) \end{array}$$

Let us show that the top arrow is injective. Write t_0 in the form (7.1.3)–(7.1.5). Then we have

$$(8.7.20) \quad \begin{aligned} \mathcal{S}_X(\sigma, p, q)^{\leq t_0} \\ = \mathcal{S}_X(\sigma | [j_{a_0} \dots j_{a_1}], p_0, p_1) \times \dots \times \mathcal{S}_X(\sigma | [j_{a_{m-1}} \dots j_{a_m}], p_{m-1}, p_m). \end{aligned}$$

Now we consider the m -cubical diagram given by

$$(8.7.21) \quad \prod_{i=1}^m \operatorname{colim}_{s \in \partial \mathcal{S}_X(\sigma | [j_{a_{i-1}} \dots j_{a_i}], p_{i-1}, p_i)} \mathcal{Y}(\sigma | [j_{a_{i-1}} \dots j_{a_i}], p_{i-1}, p_i, s) \hookrightarrow \mathcal{Y}(\sigma | [j_{a_{i-1}} \dots j_{a_i}], p_{i-1}, p_i).$$

Now the top arrow of (8.7.19) is simply the map to the maximal vertex of this m -cube from the colimit over the m -cube minus the maximal vertex (this is where we use the fact that \mathcal{T} is downward closed). This is injective for any cubical diagram of the form $\prod_{i=1}^m [A_i \hookrightarrow B_i]$ where each $A_i \hookrightarrow B_i$ is injective. Now that the top arrow in (8.7.19) is injective, the existence of the colimit is clear (first take the colimit of underlying sets and then the rest of the structure extends in an obvious way), and hence the colimit in (8.7.18) exists. This reasoning applies equally well to \mathcal{Y}_{S^1} (including $(ES^1)_{\mathcal{Y}} \rightarrow (BS^1)_{\mathcal{Y}}$).

Now that we have shown that the colimits on the left-hand side of (8.7.16)–(8.7.17) exist, it remains to define $\mathcal{Y}(\sigma, p, q)$ and $\mathcal{Y}_{S^1}(\sigma, p, q)$.

We observed earlier that any S^1 –MF-set injects into a saturated S^1 –MF-set. Let $\mathcal{Y}_{S^1}(\sigma, p, q)$ be any saturated S^1 –MF-set with an injection (8.7.17). Note that with this definition, the product/face maps with target $\mathcal{Y}_{S^1}(\sigma, p, q)$ are obvious, as is their compatibility.

To define $\mathcal{Y}(\sigma, p, q)$, consider the following diagram:

$$(8.7.22) \quad \begin{array}{ccc} \operatorname{colim}_{\mathfrak{s} \in \partial \mathcal{S}_X(\sigma, p, q)} \mathcal{Y}(\sigma, p, q, \mathfrak{s}) & & \\ \uparrow & & \\ \operatorname{colim}_{\mathfrak{s} \in \partial \mathcal{S}_X(\sigma, p, q)} \mathcal{Y}_{S^1}(\sigma, p, q, \mathfrak{s}) & \longrightarrow & \mathcal{Y}_{S^1}(\sigma, p, q) \end{array}$$

The lower colimit is in the category of S^1 –MF-sets. It is clear from the inductive proof of its existence that this particular colimit commutes with the forgetful functor to MF-sets. Hence we may equally well think of this colimit as taking place in the category of MF-sets, and with this perspective the definition of the vertical map is clear. Now since the horizontal map is injective, the colimit of this diagram (in the category of MF-sets) clearly exists (first take the colimit in sets and then the rest of the structure is obvious). We pick any saturated MF-set into which the colimit of (8.7.22) embeds, and we call this saturated MF-set $\mathcal{Y}(\sigma, p, q)$. Via the embedding of the upper colimit into $\mathcal{Y}(\sigma, p, q)$, it is tautological that the face/product maps with target $\mathcal{Y}(\sigma, p, q)$ exist and are compatible. The map $\mathcal{Y}_{S^1}(\sigma, p, q) \rightarrow \mathcal{Y}(\sigma, p, q)$ is similarly tautological, as is the fact that it is compatible with the product/face maps. \square

8.7.3 Augmented virtual cochain complexes from F–module MF-sets and S^1 –MF-sets We now describe a modified version of the complexes from Definitions 7.3.1 and 7.3.3 and their S^1 –equivariant versions using a choice of $\mathcal{Y}_{S^1} \rightarrow \mathcal{Y}$ as in Proposition 8.7.11.

Let us give alternative definitions of the complexes

$$(8.7.23) \quad C_{\text{vir}}^\bullet(X \text{ rel } \partial; \mathcal{A})^+(\sigma, p, q), \quad C_{\text{vir}}^\bullet(\partial X; \mathcal{A})^+(\sigma, p, q), \quad C_\bullet(E; \mathcal{A})^+(\sigma, p, q).$$

The “fixed fundamental cycles” $[E_\alpha] \in C_\bullet(E; \alpha)$ still live in ordinary singular simplicial chains. Now (8.7.23) are defined in terms of the (relative) singular chains on certain spaces; the modification we make is just to use a different model of singular chains (depending on (σ, p, q)) which we now describe. Our model of chains on a space X

is generated by the set of maps

$$(8.7.24) \quad M_i \times \prod_{\alpha \in \coprod \bar{\mathcal{A}}(\sigma', p', q')} \Delta^{i_\alpha} \rightarrow X,$$

where $i \in \mathcal{Y}(\sigma, p, q)$, $i_\alpha \in \mathbb{Z}_{\geq 0}$ (all but finitely many must be zero), and $\coprod \bar{\mathcal{A}}(\sigma', p', q')$ stands for

$$(8.7.25) \quad \coprod_{\substack{0 \leq i_0 < \dots < i_m \leq n \\ (p', q') \in \mathcal{P}_{\sigma(i_0)} \times \mathcal{P}_{\sigma(i_m)} \\ \exists t \in \mathcal{S}_X(\sigma, p, q) \text{ containing } ([i_0 \dots i_m], p', q')}} \bar{\mathcal{A}}(\sigma|[i_0 \dots i_m], p', q')$$

(this is very similar to (7.2.1)). We define relative chains as usual: $C_\bullet(X; Y) := C_\bullet(X)/C_\bullet(Y)$ for $Y \subseteq X$. It is a straightforward (though tedious) exercise to verify that with this definition of singular chains, the structure maps for the homotopy colimits used to construct (8.7.23) are all well-defined and appropriately compatible. We define the product $- \times [E_\alpha]$ via the Eilenberg–Zilber subdivision of $\Delta^p \times \Delta^q$ (as in Remark 4.2.2) at the index α . It may also be verified in a similar manner that the product/face maps for (8.7.23) as defined in Definitions 7.3.2 and 7.3.3 make sense and are compatible given this model of chains. For these maps, we use the product/face maps of \mathcal{Y} along with Eilenberg–Zilber (separately at each index α).

We also define S^1 –equivariant versions

$$(8.7.26) \quad C_{S^1, \text{vir}}^\bullet(\mathcal{X} \text{ rel } \partial; \mathcal{A})^+(\sigma, p, q), \quad C_{S^1, \text{vir}}^\bullet(\partial \mathcal{X}; \mathcal{A})^+(\sigma, p, q), \\ C_\bullet^{S^1}(E; \mathcal{A})^+(\sigma, p, q)$$

of (8.7.23) using $C_{S^1, \text{vir}}^\bullet$ in place of C_{vir}^\bullet (and $C_\bullet^{S^1}$ in place of C_\bullet). The fixed fundamental cycles $[E_\alpha]$ again live in ordinary singular simplicial chains. We use the following as our model for (relative) S^1 –equivariant chains (for $C_{S^1, \text{vir}}^\bullet$). Our model of S^1 –equivariant chains on an S^1 –space X is generated by the set of commuting diagrams of S^1 –equivariant maps

$$(8.7.27) \quad \begin{array}{ccc} M_i \times \prod_{\alpha \in \coprod \bar{\mathcal{A}}(\sigma', p', q')} \Delta^{i_\alpha} & \longrightarrow & X \times (ES^1)_{\mathcal{Y}_{S^1}(\sigma, p, q)} \\ \downarrow & & \downarrow \\ M_i & \longrightarrow & (ES^1)_{\mathcal{Y}_{S^1}(\sigma, p, q)} \end{array}$$

where $i \in \mathcal{Y}_{S^1}(\sigma, p, q)$, the vertical maps are the projections, the bottom horizontal map is the given structure map, and $\coprod \bar{\mathcal{A}}(\sigma', p', q')$ is as before. It is then a straightforward (though tedious) exercise to verify that with this definition of S^1 –equivariant chains, the complexes (8.7.26) are well-defined, have well-defined product/face maps, and that

the obvious forgetful maps to their non-equivariant versions (8.7.23) are well-defined and compatible with the product/face maps.

Now it remains only to show that our models of chains and S^1 -equivariant chains described above do actually calculate singular homology and S^1 -equivariant singular homology.²⁸ Note that once we do this, we can use our alternative definitions of (8.7.23) and (8.7.26) in place of the originals in all the arguments of Section 7.

We show that our model of chains (generated by diagrams (8.7.24)) calculates singular homology as follows. Let us abbreviate $\mathcal{Y} = \mathcal{Y}(\sigma, p, q)$. Our complex generated by maps (8.7.24) is in fact a double complex with bigrading $p = \sum_{\alpha} i_{\alpha}$ and $q = \dim M_i$. Now let us consider the associated spectral sequence. We calculate the $E_{p,q}^1$ term as follows. By definition $E_{p,q}^1$ is simply the homology of the complex with only the differentials decreasing q . Clearly this breaks up as a direct sum over tuples $\{i_{\alpha} \in \mathbb{Z}_{\geq 0}\}$. For a given tuple $\{i_{\alpha} \in \mathbb{Z}_{\geq 0}\}$, we must calculate the homology of the complex

$$(8.7.28) \quad \bigoplus_{\substack{i \in \mathcal{Y} \\ M_i \times \prod_{\alpha} \Delta^{i_{\alpha}} \rightarrow X}} \mathbb{Z}.$$

We claim that this is the same as the homology of the complex

$$(8.7.29) \quad \bigoplus_{\substack{i \in \mathcal{Y} \\ M_i \rightarrow X}} \mathbb{Z}.$$

There is clearly a map (8.7.29) \rightarrow (8.7.28), given by precomposing with the projection

$$M_i \times \prod_{\alpha} \Delta^{i_{\alpha}} \rightarrow M_i,$$

and picking a point $x \in \prod_{\alpha} \Delta^{i_{\alpha}}$ gives a map in the opposite direction. The composition (8.7.29) \rightarrow (8.7.28) \rightarrow (8.7.29) is clearly the identity. The other composition is chain homotopic to the identity by the following argument. Since \mathcal{Y} is saturated, we can choose a map $I: \mathcal{Y} \rightarrow \mathcal{Y}$ and isomorphisms $(M_{I(i)}, \mathcal{S}_{I(i)}) \xrightarrow{\simeq} (M_i \times [0, 1], \mathcal{S}_i \times \mathcal{S})$ compatible with the face identifications, where $([0, 1], \mathcal{S})$ is the PL manifold $[0, 1]$ with strata $\{\{0\}, \{1\}, (0, 1)\}$, and $M_i \times \{0\}$, $M_i \times \{1\}$ have the obvious face identifications with M_i (construct I and the isomorphisms by induction on dimension). Using this coherent choice of “cylinder objects”, it is easy to define a chain homotopy between the identity map and the composition (8.7.28) \rightarrow (8.7.29) \rightarrow (8.7.28) (using the fact that $\prod_{\alpha} \Delta^{i_{\alpha}}$ is contractible). It follows that the canonical map (8.7.29) \rightarrow (8.7.28)

²⁸Technically speaking, we must also show that the pushforward maps, product maps, and Gysin maps have the expected action on homology, though this verification is safely left to the reader.

is a quasi-isomorphism. By Lemma 8.7.7, the complex (8.7.29) calculates singular homology. Hence we conclude that the E^1 term is

$$(8.7.30) \quad E_{p,q}^1 = \bigoplus_{\sum_{\alpha} i_{\alpha} = p} H_q(X).$$

The differentials on the E^1 page are easy to understand (essentially we have $H_q(X)$ tensored with the restricted tensor product $\otimes_{\alpha} C_{\bullet}(\text{pt})$). We conclude that

$$(8.7.31) \quad E_{p,q}^2 = \begin{cases} H_q(X), & p = 0, \\ 0, & p \neq 0. \end{cases}$$

Hence there are no further differentials and we are done.

We now use a similar argument to show that our model of S^1 -equivariant chains calculates S^1 -equivariant singular homology. Let us abbreviate $\mathcal{Y} = \mathcal{Y}_{S^1}(\sigma, p, q)$. Our complex generated by maps (8.7.27) is a double complex as before, and we consider the associated spectral sequence. As before, $E_{p,q}^1$ splits up as a direct sum over tuples $\{i_{\alpha} \in \mathbb{Z}_{\geq 0}\}$. For a given tuple $\{i_{\alpha} \in \mathbb{Z}_{\geq 0}\}$, the corresponding direct summand is the homology of the complex freely generated by diagrams

$$(8.7.32) \quad \begin{array}{ccc} M_i \times \prod_{\alpha} \Delta^{i_{\alpha}} & \longrightarrow & X \times (ES^1)_{\mathcal{Y}_{S^1}} \\ \downarrow & & \downarrow \\ M_i & \longrightarrow & (ES^1)_{\mathcal{Y}_{S^1}} \end{array} \quad \text{with } i \in \mathcal{Y}.$$

As in the non-equivariant case, this is canonically quasi-isomorphic to the complex freely generated by diagrams

$$(8.7.33) \quad \begin{array}{ccc} M_i & \longrightarrow & X \times (ES^1)_{\mathcal{Y}_{S^1}} \\ \downarrow & & \downarrow \\ M_i & \longrightarrow & (ES^1)_{\mathcal{Y}_{S^1}} \end{array} \quad \text{with } i \in \mathcal{Y}.$$

Thus by Lemma 8.7.10 the E^1 term is

$$(8.7.34) \quad E_{p,q}^1 = \bigoplus_{\sum_{\alpha} i_{\alpha} = p} H_q^{S^1}(X).$$

The differentials on the E^1 page are as before, and thus

$$(8.7.35) \quad E_{p,q}^2 = \begin{cases} H_q^{S^1}(X), & p = 0, \\ 0, & p \neq 0, \end{cases}$$

so we are done.

8.7.4 Floer-type homology groups from F-module MF-sets and S^1 -MF-sets

We now give a new definition of Floer-type homology groups identical to Definition 7.7.6 except using the chain models from Section 8.7.3. Precisely, we define a new resolution $\tilde{Z}_\bullet \rightarrow Z_\bullet$ by including the data²⁹ of F-module $(S^1\text{-})\text{MF-sets } \mathcal{Y}_{S^1} \rightarrow \mathcal{Y}$ over Δ^n for $f^*\mathcal{X}$ satisfying the conclusion of Proposition 8.7.11. Now Proposition 8.7.11 implies that this new resolution is still a trivial Kan fibration, and thus the rest of the construction of Floer-type homology groups in Section 7 applies as written.

Remark 8.7.12 These Floer-type homology groups are *a priori* different from the ones constructed in Section 7. However, they share all the same properties, and we certainly expect them to be canonically isomorphic (see Remark 4.2.3). We won't pursue this here, though, since it is not necessary for our intended application.

8.7.5 Localization The following localization result concerns the Floer-type homology groups from Section 8.7.4.

Theorem 8.7.13 (S^1 -localization for $\mathbb{H}(\mathcal{X})$) *Let \mathcal{X}/Z_\bullet be an H -equivariant flow category diagram with implicit atlas \mathcal{A} and coherent orientations ω , satisfying the hypotheses of Definition 7.7.6. Suppose that S^1 acts compatibly on this entire structure (S^1 -actions on $\mathcal{X}(\sigma, p, q)$ and S^1 -equivariant implicit atlases, so that all the relevant structure maps are S^1 -equivariant).*

Fix partitions (into disjoint closed subsets)

$$(8.7.36) \quad \mathcal{X}(\sigma, p, q) = \mathcal{X}(\sigma, p, q)^0 \sqcup \mathcal{X}(\sigma, p, q)^1$$

where S^1 acts almost freely on $\mathcal{X}(\sigma, p, q)^1$, and suppose that the product/face maps specialize to maps

$$(8.7.37) \quad \mathcal{X}(\sigma|[0 \dots k], p, q)^0 \times \mathcal{X}(\sigma|[k \dots n], q, r)^0 \rightarrow \mathcal{X}(\sigma, p, r)^0,$$

$$(8.7.38) \quad \mathcal{X}(\sigma|[0 \dots k], p, q)^0 \times \mathcal{X}(\sigma|[k \dots n], q, r)^1 \rightarrow \mathcal{X}(\sigma, p, r)^1,$$

$$(8.7.39) \quad \mathcal{X}(\sigma|[0 \dots k], p, q)^1 \times \mathcal{X}(\sigma|[k \dots n], q, r)^0 \rightarrow \mathcal{X}(\sigma, p, r)^1,$$

$$(8.7.40) \quad \mathcal{X}(\sigma|[0 \dots k], p, q)^1 \times \mathcal{X}(\sigma|[k \dots n], q, r)^1 \rightarrow \mathcal{X}(\sigma, p, r)^1,$$

$$(8.7.41) \quad \mathcal{X}(\sigma|[0 \dots \hat{k} \dots n], p, q)^0 \rightarrow \mathcal{X}(\sigma, p, q)^0,$$

$$(8.7.42) \quad \mathcal{X}(\sigma|[0 \dots \hat{k} \dots n], p, q)^1 \rightarrow \mathcal{X}(\sigma, p, q)^1.$$

²⁹For the reader concerned with set-theoretic issues: so that the collection of all possible choices of this data forms a set, one may add “rigidifying data” to the definition of an $(S^1\text{-})\text{MF-set}$, eg require \mathcal{Y} to be countable, equip \mathcal{Y} with an injection into ω_1 (the first uncountable ordinal), and equip each M_i and $(ES^1)_{\mathcal{Y}}$ with an embedding into \mathbb{R}^∞ (note that this rigidifying data is not required to be compatible in any way with the rest of the structure, and thus it does not affect any of our previous arguments).

Assume in addition that for any vertex $\sigma^0 \in Z_\bullet$, we have

$$\mathcal{X}(\sigma^0, p, q)^0 = \emptyset \implies \mathcal{X}(\sigma^i, p, q)^0 = \emptyset$$

for all i and $p \neq q$.

Then there is a canonical isomorphism $\mathbb{H}(\mathcal{X})_{\mathcal{A}} = \mathbb{H}(\mathcal{X}^0)_{\mathcal{A}}$, where \mathcal{X}^0 is the H -equivariant flow category diagram with product/face maps given by (8.7.37) and (8.7.41), equipped with the implicit atlas obtained by removing \mathcal{X}^1 from every thickening.

Proof Part I We begin by reducing to a situation in which the partition (8.7.36) is extended to every thickened moduli space, so that the product/face maps on thickened moduli spaces also take the form (8.7.37)–(8.7.42). Precisely, consider a new implicit atlas $\tilde{\mathcal{A}}$ (on the same index set \mathcal{A}) in which the thickened moduli spaces are

$$(8.7.43) \quad \mathcal{X}(\sigma, p, q)_{\tilde{I}}^{\leq s} := [\mathcal{X}(\sigma, p, q)_{\tilde{I}}^{\leq s} \setminus \mathcal{X}(\sigma, p, q)^1] \sqcup [\mathcal{X}(\sigma, p, q)_{\tilde{I}}^{\leq s} \setminus \mathcal{X}(\sigma, p, q)^0].$$

Now these thickened moduli spaces are by definition equipped with a partition

$$\mathcal{X}(\sigma, p, q)_{\tilde{I}}^{\leq s} = (\mathcal{X}(\sigma, p, q)_{\tilde{I}}^{\leq s})^0 \sqcup (\mathcal{X}(\sigma, p, q)_{\tilde{I}}^{\leq s})^1$$

(namely the partition (8.7.43)), and the natural product/face maps for the implicit atlas $\tilde{\mathcal{A}}$ take the form (8.7.37)–(8.7.42).

The thickened moduli spaces of $\tilde{\mathcal{A}}$ are equipped with natural maps to those of \mathcal{A} (compatible with the product/face maps); these maps are not open embeddings, but nevertheless the proof of Lemma 7.8.2 applies to give a canonical isomorphism $\mathbb{H}(\mathcal{X})_{\mathcal{A}} = \mathbb{H}(\mathcal{X})_{\tilde{\mathcal{A}}}$. Thus it suffices to produce an isomorphism $\mathbb{H}(\mathcal{X})_{\tilde{\mathcal{A}}} = \mathbb{H}(\mathcal{X}^0)_{\tilde{\mathcal{A}}^0}$, where $\tilde{\mathcal{A}}^0$ is the implicit atlas on \mathcal{X}^0 with thickened moduli spaces

$$(\mathcal{X}(\sigma, p, q)_{\tilde{I}}^{\leq s})^0.$$

From now on, we will work exclusively with the atlas $\tilde{\mathcal{A}}$, which we now rather abusively rename as \mathcal{A} .

Part II We will define an isomorphism $\mathbb{H}(\mathcal{X})_{\mathcal{A}} = \mathbb{H}(\mathcal{X}^0)_{\mathcal{A}^0}$, where \mathcal{A}^0 is the implicit atlas on \mathcal{X}^0 with thickened moduli spaces

$$(\mathcal{X}(\sigma, p, q)_{\tilde{I}}^{\leq s})^0.$$

Note that since we have coherent orientations ω and S^1 is connected, it follows that all the flow spaces are locally S^1 -orientable (see Remark 8.3.2), so we may use the S^1 -localization machinery freely.

Let $\tilde{Z}_\bullet \rightarrow Z_\bullet$ denote the resolution (as in Section 8.7.4) associated to \mathcal{X}/Z_\bullet with the implicit atlas \mathcal{A} , and let $\tilde{Z}_\bullet^0 \rightarrow Z_\bullet$ denote the resolution associated to \mathcal{X}^0/Z_\bullet with the implicit atlas \mathcal{A}^0 .

We will construct a diagram of the following shape:

$$(8.7.44) \quad \begin{array}{ccccccc} \tilde{Z}_\bullet & \leftarrow & \tilde{Z}_\bullet^a & \rightarrow & \tilde{Z}_\bullet^b & \leftarrow & \tilde{Z}_\bullet^c & \rightarrow & \tilde{Z}_\bullet^0 \\ & & \searrow & & \downarrow & & \swarrow & & \nearrow \\ & & & & Z_\bullet & & & & \end{array}$$

and natural diagrams $\tilde{\mathbb{H}}, \tilde{\mathbb{H}}^a, \tilde{\mathbb{H}}^b, \tilde{\mathbb{H}}^c, \tilde{\mathbb{H}}^0$ from $\tilde{Z}_\bullet, \tilde{Z}_\bullet^a, \tilde{Z}_\bullet^b, \tilde{Z}_\bullet^c, \tilde{Z}_\bullet^0$ to $N_{\text{dg}}(\text{Ch}_R[[H]])$, respectively. We will also construct natural isomorphisms relating these and their pullbacks under the horizontal maps in (8.7.44). Finally, we will show that each of the vertical maps (8.7.44) satisfies the conclusions of Propositions 7.5.5 and 7.7.4. The desired result follows easily from these statements; let us now start on the proof.

We may consider $\mathcal{A}(\sigma, p, q)$ as a system of implicit atlases on $\mathcal{X}(\sigma, p, q)^0$ and on $\mathcal{X}(\sigma, p, q)^1$ separately. Now consider the following maps (where the complexes on the left are defined analogously to those on the right as in Section 8.7.3):

$$(8.7.45) \quad C_{\text{vir}}^\bullet(\partial\mathcal{X}^0, \mathcal{A})^+ \oplus C_{S^1, \text{vir}}^\bullet(\partial\mathcal{X}^1, \mathcal{A})^+ \rightarrow C_{\text{vir}}^\bullet(\partial\mathcal{X}; \mathcal{A})^+,$$

$$(8.7.46) \quad C_{\text{vir}}^\bullet(\mathcal{X}^0 \text{ rel } \partial, \mathcal{A})^+ \oplus C_{S^1, \text{vir}}^\bullet(\mathcal{X}^1 \text{ rel } \partial, \mathcal{A})^+ \rightarrow C_{\text{vir}}^\bullet(\mathcal{X} \text{ rel } \partial; \mathcal{A})^+.$$

We give $C_{\text{vir}}^\bullet(\mathcal{X}^0 \text{ rel } \partial, \mathcal{A})^+ \oplus C_{S^1, \text{vir}}^\bullet(\mathcal{X}^1 \text{ rel } \partial, \mathcal{A})^+$ the structure of an F–module with product/face maps

$$(8.7.47) \quad \begin{aligned} C_{\text{vir}}^\bullet(\mathcal{X}^0 \text{ rel } \partial; \mathcal{A})^+(\dots) \otimes C_{\text{vir}}^\bullet(\mathcal{X}^0 \text{ rel } \partial; \mathcal{A})^+(\dots) \\ \rightarrow C_{\text{vir}}^\bullet(\partial\mathcal{X}^0; \mathcal{A})^+(\dots), \end{aligned}$$

$$(8.7.48) \quad \begin{aligned} C_{\text{vir}}^\bullet(\mathcal{X}^0 \text{ rel } \partial; \mathcal{A})^+(\dots) \otimes C_{S^1, \text{vir}}^\bullet(\mathcal{X}^1 \text{ rel } \partial; \mathcal{A})^+(\dots) \\ \rightarrow C_{S^1, \text{vir}}^\bullet(\partial\mathcal{X}^1; \mathcal{A})^+(\dots), \end{aligned}$$

$$(8.7.49) \quad \begin{aligned} C_{S^1, \text{vir}}^\bullet(\mathcal{X}^1 \text{ rel } \partial; \mathcal{A})^+(\dots) \otimes C_{\text{vir}}^\bullet(\mathcal{X}^0 \text{ rel } \partial; \mathcal{A})^+(\dots) \\ \rightarrow C_{S^1, \text{vir}}^\bullet(\partial\mathcal{X}^1; \mathcal{A})^+(\dots), \end{aligned}$$

$$(8.7.50) \quad \begin{aligned} C_{S^1, \text{vir}}^\bullet(\mathcal{X}^1 \text{ rel } \partial; \mathcal{A})^+(\dots) \otimes C_{S^1, \text{vir}}^\bullet(\mathcal{X}^1 \text{ rel } \partial; \mathcal{A})^+(\dots) \\ \rightarrow C_{S^1, \text{vir}}^\bullet(\partial\mathcal{X}^1; \mathcal{A})^+(\dots), \end{aligned}$$

$$(8.7.51) \quad C_{\text{vir}}^\bullet(\mathcal{X}^0 \text{ rel } \partial; \mathcal{A})^+(\dots) \rightarrow C_{\text{vir}}^\bullet(\partial\mathcal{X}^0; \mathcal{A})^+(\dots),$$

$$(8.7.52) \quad C_{S^1, \text{vir}}^\bullet(\mathcal{X}^1 \text{ rel } \partial; \mathcal{A})^+(\dots) \rightarrow C_{S^1, \text{vir}}^\bullet(\partial\mathcal{X}^1; \mathcal{A})^+(\dots),$$

just as in Definition 7.3.2 (and using the chain models from Section 8.7.3). These are compatible with (8.7.45)–(8.7.46) and (7.3.6)–(7.3.7).

Next, we consider $C_{\bullet}(E; \mathcal{A})^+ \oplus C_{\bullet-1}^{S^1}(E; \mathcal{A})^+$ (as in Section 8.7.3). Now, there are two natural maps

$$(8.7.53) \quad C_{\bullet}(E; \mathcal{A})^+ \oplus C_{\bullet-1}^{S^1}(E; \mathcal{A})^+ \xrightarrow[\text{id} \oplus 0]{\text{id} \oplus \pi^!} C_{\bullet}(E; \mathcal{A})^+.$$

We give $C_{\bullet}(E; \mathcal{A})^+ \oplus C_{\bullet-1}^{S^1}(E; \mathcal{A})^+$ the structure of an F -module with product/face maps of the shape (8.7.47)–(8.7.52) just as in Definition 7.3.3. Then both maps (8.7.53) are maps of F -modules. There is also a pushforward map

$$C_{\text{vir}}^{\text{vdim } \mathcal{X} + \bullet}(\mathcal{X}^0 \text{ rel } \partial, \mathcal{A})^+ \oplus C_{S^1, \text{vir}}^{\text{vdim } \mathcal{X} + \bullet}(\mathcal{X}^1 \text{ rel } \partial, \mathcal{A})^+ \rightarrow C_{-\bullet}(E; \mathcal{A})^+ \oplus C_{-\bullet-1}^{S^1}(E; \mathcal{A})^+,$$

which is a map of F -modules.

Now let us define \tilde{Z}_{\bullet}^a , \tilde{Z}_{\bullet}^b , and \tilde{Z}_{\bullet}^c . We modify the definition of $\tilde{Z}_{\bullet} \rightarrow Z_{\bullet}$ (ie Definition 7.5.3 as amended in Section 8.7.4) as follows. For $\tilde{Z}_{\bullet}^a = \tilde{Z}_{\bullet}^c$ we replace Definition 7.5.3(iv)–(v) with (i)–(ii) below, and for \tilde{Z}_{\bullet}^b we replace Definition 7.5.3(iv)–(vi) with (i)–(iii) below.

- (i) An H -invariant system of chains

$$\lambda = \lambda^0 \oplus \lambda^1 \in C_{\text{vir}}^{\bullet}(\mathcal{X}^0 \text{ rel } \partial, \mathcal{B})^+ \oplus C_{S^1, \text{vir}}^{\bullet}(\mathcal{X}^1 \text{ rel } \partial, \mathcal{B})^+$$

(degree zero and supported inside $\text{supp } \mathcal{X}$) with the following property. Note that (μ, λ) is a cycle in the mapping cone

$$\begin{array}{ccc} C_{\text{vir}}^{\bullet}(\partial f^* \mathcal{X}^0; \mathcal{B})^+ & & C_{\text{vir}}^{\bullet}(f^* \mathcal{X}^0 \text{ rel } \partial; \mathcal{B})^+ \\ \oplus & \longrightarrow & \oplus \\ C_{S^1, \text{vir}}^{\bullet}(\partial f^* \mathcal{X}^1; \mathcal{B})^+ & & C_{S^1, \text{vir}}^{\bullet}(f^* \mathcal{X}^1 \text{ rel } \partial; \mathcal{B})^+ \end{array}$$

whose homology is identified with

$$\check{H}^{\bullet}(f^* \mathcal{X}^0; \mathfrak{o}_{f^* \mathcal{X}^0}) \oplus \check{H}^{\bullet}((f^* \mathcal{X}^1)/S^1; \pi_* \mathfrak{o}_{f^* \mathcal{X}^1}),$$

which in degree zero simply equals $\check{H}^0(f^* \mathcal{X}; \mathfrak{o}_{f^* \mathcal{X}})$. We require that the homology class of (μ, λ) equal $f^* \omega \in \check{H}^0(f^* \mathcal{X}; \mathfrak{o}_{f^* \mathcal{X}})$.

- (ii) An H -invariant system of chains $\tilde{\lambda}$ for $Q[C_{\bullet}(E; \mathcal{A})^+ \oplus C_{\bullet-1}^{S^1}(E; \mathcal{A})^+]$ (degree $\text{gr}(q) - \text{gr}(p) + \dim \sigma - 1$ and supported inside $\text{supp } \mathcal{X}$) whose image in $C_{\bullet}(E; \mathcal{A})^+ \oplus C_{\bullet-1}^{S^1}(E; \mathcal{A})^+$ coincides with the image of λ .
- (iii) An H -invariant map

$$[[E]]: Q[C_{\bullet}(E; \mathcal{A})_{\mathbb{Z}}^+ \oplus C_{\bullet-1}^{S^1}(E; \mathcal{A})_{\mathbb{Z}}^+] \rightarrow \mathbb{Z}$$

of F -modules which sends the fundamental class in $H_{\bullet}(E; \mathcal{A})_{\mathbb{Z}}^+$ to 1 and is zero on $H_{\bullet-1}^{S^1}(E; \mathcal{A})_{\mathbb{Z}}^+$.

The maps $\tilde{Z}_\bullet \leftarrow \tilde{Z}_\bullet^a \rightarrow \tilde{Z}_\bullet^b \leftarrow \tilde{Z}_\bullet^c \rightarrow \tilde{Z}_\bullet^0$ are defined as follows:

$$(8.7.54) \quad \begin{array}{ccccccc} \tilde{Z}_\bullet & \longleftarrow & \tilde{Z}_\bullet^a & \longrightarrow & \tilde{Z}_\bullet^b & \longleftarrow & \tilde{Z}_\bullet^c \longrightarrow \tilde{Z}_\bullet^0 \\ \lambda & \xleftarrow{(8.7.46)} & \lambda & & \lambda & \xrightarrow{\lambda = \lambda^0 \oplus \lambda^1} & \lambda^0 \\ \tilde{\lambda} & \xleftarrow{Q[\text{id} \oplus \pi^1]} & \tilde{\lambda} & & \tilde{\lambda} & \xrightarrow{Q[\text{id} \oplus 0]} & \tilde{\lambda}^0 \\ & & & & \llbracket E \rrbracket & \xrightarrow{Q[\text{id} \oplus \pi^1]} & \llbracket E \rrbracket \xleftarrow{Q[\text{id} \oplus 0]} \llbracket E \rrbracket \end{array}$$

Lemma 8.1.2 shows that the map $\tilde{Z}^a \rightarrow \tilde{Z}^b$ above produces an $\llbracket E \rrbracket$ satisfying (iii).

Now $\tilde{\mathbb{H}}, \tilde{\mathbb{H}}^a, \tilde{\mathbb{H}}^b, \tilde{\mathbb{H}}^c, \tilde{\mathbb{H}}^0$ are defined via the matrix coefficients $c_{\sigma,p,q}$, defined as

$$c_{\sigma,p,q} := \begin{cases} (\llbracket E \rrbracket_{\sigma,p,q} \otimes \text{id}_R)(\tilde{\lambda}_{\sigma,p,q}) & \text{for } \tilde{\mathbb{H}}, \\ (\llbracket E \rrbracket_{\sigma,p,q} \otimes \text{id}_R)(\text{id} \oplus \pi^1)(\tilde{\lambda}_{\sigma,p,q}) & \text{for } \tilde{\mathbb{H}}^a, \\ (\llbracket E \rrbracket_{\sigma,p,q} \otimes \text{id}_R)(\tilde{\lambda}_{\sigma,p,q}) & \text{for } \tilde{\mathbb{H}}^b, \\ (\llbracket E \rrbracket_{\sigma,p,q} \otimes \text{id}_R)(\text{id} \oplus 0)(\tilde{\lambda}_{\sigma,p,q}) & \text{for } \tilde{\mathbb{H}}^c, \\ (\llbracket E \rrbracket_{\sigma,p,q} \otimes \text{id}_R)(\tilde{\lambda}_{\sigma,p,q}) & \text{for } \tilde{\mathbb{H}}^0. \end{cases}$$

These matrix coefficients give rise to diagrams in $\text{N}_{\text{dg}}(\text{Ch}_R \llbracket H \rrbracket)$ since $\llbracket E \rrbracket$ is a map of F -modules. The isomorphisms between these diagrams and their pullbacks under the maps (8.7.54) are evident.

It now remains to show that each of the maps $\tilde{Z}_\bullet^a, \tilde{Z}_\bullet^b, \tilde{Z}_\bullet^c \rightarrow \tilde{Z}_\bullet$ satisfies the conclusions of Propositions 7.5.5 and 7.7.4. The proof of Proposition 7.5.5 applies to these maps as written (for \tilde{Z}_\bullet^b , the extension of $\llbracket E \rrbracket$ step uses the fact that $H_\bullet(E; \mathcal{A}) \oplus H_{\bullet-1}^{S_1}(E; \mathcal{A})$ is concentrated in degrees ≥ 0 so that Lemma 7.4.7 still applies). The proof of Proposition 7.7.4 also applies to these maps as written. Now the uniqueness of descent from Lemma 7.7.5 shows that there are canonical isomorphisms of the descents $\mathbb{H} = \mathbb{H}^a = \mathbb{H}^b = \mathbb{H}^c = \mathbb{H}^0$. □

9 Gromov–Witten invariants

In this section, we define Gromov–Witten invariants for a general smooth closed symplectic manifold (X, ω) (which we now fix). This has been treated in the literature by Li and Tian [53], Fukaya and Ono [33] and Ruan [71].

More specifically, the main subject of this section is the construction of an implicit atlas on the moduli space of stable J -holomorphic maps $\overline{\mathcal{M}}_{g,n}^\beta(X)$; the same method also yields an implicit atlas on $\overline{\mathcal{M}}_{g,n}^\beta(X, J_{[0,1]})$ (the moduli space associated to a family of J parametrized by $[0, 1]$). Hence the virtual fundamental class $[\overline{\mathcal{M}}_{g,n}^\beta(X)]^{\text{vir}}$ (Definition 5.1.1) is defined, and the properties derived in Section 5 show that its image

in $H_\bullet(\overline{\mathcal{M}}_{g,n} \times X^n)$ is independent of J . The necessary gluing results (which we isolate in Proposition 9.2.6) are proved in Appendix B.

It would be interesting to use implicit atlases to prove that the invariants defined here satisfy the Kontsevich–Manin axioms [52], as proved in Fukaya and Ono [33] for their definition (to do this, one would need to show additional properties of the virtual fundamental class, eg as suggested in Section 5).

9.1 Moduli space $\overline{\mathcal{M}}_{g,n}^\beta(X)$

Let us now fix a smooth almost complex structure J on X which is tamed by ω .

Definition 9.1.1 (Nodal curve of type (g, n)) A nodal curve of type (g, n) is a compact nodal Riemann surface C of arithmetic genus g along with an injective function $l: \{1, \dots, n\} \rightarrow C$ (the “ n marked points”) whose image is disjoint from the nodes of C . An isomorphism $(C, l) \rightarrow (C', l')$ of curves of type (g, n) is an isomorphism of Riemann surfaces $\iota: C \rightarrow C'$ such that $l' = \iota \circ l$. We usually omit l from the notation. Such a curve is called *stable* if and only if its automorphism group is finite.

We denote by $\overline{\mathcal{M}}_{g,n}$ the Deligne–Mumford moduli space of stable nodal curves of type (g, n) , and we denote by $\overline{\mathcal{C}}_{g,n} \rightarrow \overline{\mathcal{M}}_{g,n}$ the universal family (which coincides with the “forget the last marked point and stabilize” map $\overline{\mathcal{M}}_{g,n+1} \rightarrow \overline{\mathcal{M}}_{g,n}$). The moduli space $\overline{\mathcal{M}}_{g,n}$ is a compact complex analytic orbifold.

Definition 9.1.2 (J –holomorphic map) A J –holomorphic map of type (g, n) is a pair (C, u) where C is a nodal curve of type (g, n) and $u: C \rightarrow X$ is smooth and satisfies $\bar{\partial}u = 0$. An isomorphism $(C, u) \rightarrow (C', u')$ of J –holomorphic maps of type (g, n) is an isomorphism $\iota: C \rightarrow C'$ of curves of type (g, n) such that $u = u' \circ \iota$. We say a J –holomorphic map is *stable* if and only if its automorphism group (ie group of self-isomorphisms) is finite.

Definition 9.1.3 (Moduli space of stable maps; introduced by Kontsevich [51]) Let $\beta \in H_2(X, \mathbb{Z})$. We define $\overline{\mathcal{M}}_{g,n}^\beta(X)$ as the set of stable J –holomorphic maps of type (g, n) for which $u_*[C] = \beta$. We equip $\overline{\mathcal{M}}_{g,n}^\beta(X)$ with the Gromov topology, which is well-known to be compact Hausdorff.

For completeness, let us recall the definition of the Gromov topology on $\overline{\mathcal{M}}_{g,n}^\beta(X)$. A neighborhood base³⁰ at a pair (C, u) may be obtained as follows. We choose some

³⁰A neighborhood of a point x in a topological space X is a subset $N \subseteq X$ such that $x \in N^\circ$ (the interior). A neighborhood base at a point $x \in X$ is a collection of neighborhoods $\{N_\alpha\}$ of x such that

additional ℓ marked points on C so that it has no automorphisms fixing these points, and we consider the graph $\Gamma_u \subseteq X \times \bar{C}_{g,n+\ell}$, where $\bar{C}_{g,n+\ell} \rightarrow \bar{M}_{g,n+\ell}$ denotes the universal curve. A neighborhood base at (C, u) is obtained by taking all J -holomorphic maps from curves in $\bar{M}_{g,n+\ell}$ whose graph is close to Γ_u in the Hausdorff topology (and forgetting the ℓ extra marked points). Choosing different ℓ' marked points yields an equivalent neighborhood base.

9.2 Implicit atlas on $\bar{M}_{g,n}^\beta(X)$

We define an implicit atlas A^{GW} on $\bar{M}_{g,n}^\beta(X)$, proceeding in several steps. Note that the space on which we will define an implicit atlas is no longer denoted X , and this leads to a few (evident) notational differences from the earlier sections where we considered implicit atlases abstractly.

Definition 9.2.1 (Index set A^{GW}) A (Gromov–Witten) thickening datum α is a 6-tuple $(D_\alpha, r_\alpha, \Gamma_\alpha, \bar{M}_\alpha, E_\alpha, \lambda_\alpha)$, where:

- (i) $D_\alpha \rightarrow X$ is a compact smooth embedded codimension-two submanifold with boundary.
- (ii) $r_\alpha \geq 0$ is an integer such that $2g + n + r_\alpha > 2$.
- (iii) Γ_α is a finite group.
- (iv) \bar{M}_α is a smooth Γ_α -manifold³¹ equipped with an isomorphism of orbifolds between $\bar{M}_\alpha/\Gamma_\alpha$ and an open subset of $\bar{M}_{g,n+r_\alpha}/S_{r_\alpha}$ (let $\bar{C}_\alpha \rightarrow \bar{M}_\alpha$ be the pullback³² of the universal family; the action of Γ_α on \bar{M}_α lifts canonically to \bar{C}_α by the universal property of the pullback).
- (v) E_α is a finitely generated $\mathbb{R}[\Gamma_\alpha]$ -module.
- (vi) $\lambda_\alpha: E_\alpha \rightarrow C^\infty(\bar{C}_\alpha \times X, \Omega_{\bar{C}_\alpha/\bar{M}_\alpha}^{0,1} \otimes_{\mathbb{C}} TX)$ is a Γ_α -equivariant linear map supported away from the nodes and marked points of the fibers of $\bar{C}_\alpha \rightarrow \bar{M}_\alpha$.

Let A^{GW} denote the set³³ of all thickening datums.

for every open $U \subseteq X$ containing x , there exists some $N_\alpha \subseteq U$. A neighborhood base is necessarily nonempty and filtered, ie for all α, β , there exists γ with $N_\gamma \subseteq N_\alpha \cap N_\beta$. Conversely, given a set X and for every $x \in X$ a nonempty filtered collection of subsets $\{N_\alpha^x\}$ each containing x , there is a unique topology on X such that $\{N_\alpha^x\}$ is a neighborhood base at x for all $x \in X$.

³¹Note that we do not assume Γ_α acts effectively on \bar{M}_α . Indeed, $\bar{M}_{1,1}$ and $\bar{M}_{2,0}$ are ineffective orbifolds, so we must allow ineffective actions if we want $\bar{M}_\alpha/\Gamma_\alpha$ to be isomorphic to an open subset of one of these spaces.

³²Pullback here means the orbifold fiber product (see Remark 2.1.6). Note that the fibers of $\bar{C}_\alpha \rightarrow \bar{M}_\alpha$ are nodal curves of type $(g, n + r_\alpha)$ (not quotiented by their automorphism group).

³³The reader concerned with set-theoretic issues may wish to add “rigidifying data” to the definition of a thickening datum as in Remark 2.2.1.

Definition 9.2.2 (Transversality of smooth maps) Let $u: C \rightarrow X$ be a smooth map from a nodal curve of type (g, n) , and let $D \subseteq X$ be a smooth codimension-two submanifold with boundary. We say u is *transverse to D* (written $u \pitchfork D$) if and only if $u^{-1}(\partial D) = \emptyset$ and for all $p \in u^{-1}(D)$, the derivative $du: T_p C \rightarrow T_{u(p)} X / T_{u(p)} D$ is surjective and p is neither a node nor a marked point of C .

Definition 9.2.3 (I -thickened J -holomorphic map) Let $I \subseteq A^{\text{GW}}$ be a finite subset. An I -thickened J -holomorphic map of type (g, n) is a tuple $(C, u, \{\phi_\alpha\}_{\alpha \in I}, \{e_\alpha\}_{\alpha \in I})$, where:

- (i) C is a nodal curve of type (g, n) .
- (ii) $u: C \rightarrow X$ is a smooth map such that $u \pitchfork D_\alpha$ with exactly r_α intersections for all $\alpha \in I$.
- (iii) $\phi_\alpha: C \rightarrow \bar{C}_\alpha$ an isomorphism between C (with r_α extra marked points $u^{-1}(D_\alpha)$) and some fiber of $\bar{C}_\alpha \rightarrow \bar{M}_\alpha$.
- (iv) $e_\alpha \in E_\alpha$.
- (v) The I -thickened $\bar{\partial}$ -equation

$$(9.2.1) \quad \bar{\partial}u + \sum_{\alpha \in I} \lambda_\alpha(e_\alpha)(\phi_\alpha, u) = 0$$

is satisfied. This equation takes place in

$$C^\infty(\tilde{C}, \Omega_{\tilde{C}}^{0,1} \otimes_{\mathbb{C}} u^*TX).$$

Certainly $\bar{\partial}u$ is a smooth section of $\Omega_{\tilde{C}}^{0,1} \otimes_{\mathbb{C}} u^*TX$ over \tilde{C} . To interpret $\lambda_\alpha(e_\alpha)(\phi_\alpha, u)$ as such a section, consider the composition

$$(9.2.2) \quad C \xrightarrow{(\phi_\alpha, u)} \bar{C}_\alpha \times X \xrightarrow{\lambda(e_\alpha)} \Omega_{\bar{C}_\alpha/\bar{M}_\alpha}^{0,1} \otimes_{\mathbb{C}} TX.$$

This sends a point $p \in C$ to an element of

$$(\Omega_{\bar{C}_\alpha/\bar{M}_\alpha}^{0,1})_{\phi_\alpha(p)} \otimes T_{u(p)}X,$$

which we identify via ϕ_α with $(\Omega_C^{0,1} \otimes_{\mathbb{C}} u^*TX)_p$. This is what we mean by $\lambda_\alpha(e_\alpha)(\phi_\alpha, u)$.

An *isomorphism* $\iota: (C, u, \{\phi_\alpha\}_{\alpha \in I}, \{e_\alpha\}_{\alpha \in I}) \rightarrow (C', u', \{\phi'_\alpha\}_{\alpha \in I}, \{e'_\alpha\}_{\alpha \in I})$ between two I -thickened J -holomorphic maps of type (g, n) is an isomorphism $\iota: C \rightarrow C'$ of curves of (g, n) -type such that $u = u' \circ \iota$, $\phi_\alpha = \phi'_\alpha \circ \iota$ and $e_\alpha = e'_\alpha$. We say an I -thickened J -holomorphic map is *stable* if and only if its automorphism group (ie group of self-isomorphisms) is finite.

Definition 9.2.4 (Atlas data for A^{GW} on $\overline{\mathcal{M}}_{g,n}^\beta(X)$) We define $\overline{\mathcal{M}}_{g,n}^\beta(X)_I$ as the set of isomorphism classes of stable I -thickened J -holomorphic maps of type (g, n) such that $u_*[C] = \beta$. Equip $\overline{\mathcal{M}}_{g,n}^\beta(X)_I$ with the Gromov topology³⁴ for (u, ϕ_α) and with the obvious topology for e_α . It is clear by definition that $\overline{\mathcal{M}}_{g,n}^\beta(X) = \overline{\mathcal{M}}_{g,n}^\beta(X)_\emptyset$.

There is an evident action Γ_I on $\overline{\mathcal{M}}_{g,n}^\beta(X)_I$, namely

$$\{g_\alpha\}_{\alpha \in I} \cdot (u, \{\phi_\alpha\}_{\alpha \in I}, \{e_\alpha\}_{\alpha \in I}) = (u, \{g_\alpha \cdot \phi_\alpha\}_{\alpha \in I}, \{g_\alpha \cdot e_\alpha\}_{\alpha \in I}),$$

which works since $\lambda_\alpha(e_\alpha)(\phi_\alpha, u) = \lambda_\alpha(g_\alpha \cdot e_\alpha)(g_\alpha \cdot \phi_\alpha, u)$ by Γ_α -equivariance.

There are evident maps $s_\alpha: \overline{\mathcal{M}}_{g,n}^\beta(X)_I \rightarrow E_\alpha$ simply picking out e_α .

For $I \subseteq J \subseteq A^{GW}$, there is an obvious forgetful map

$$\psi_{IJ}: (s_{J \setminus I} | \overline{\mathcal{M}}_{g,n}^\beta(X)_J)^{-1}(0) \rightarrow \overline{\mathcal{M}}_{g,n}^\beta(X)_I.$$

Let $U_{IJ} \subseteq \overline{\mathcal{M}}_{g,n}^\beta(X)_I$ consist of those elements such that $u \pitchfork D_\alpha$ with exactly r_α intersections for all $\alpha \in J \setminus I$ and such that adding these extra marked points makes C isomorphic to a fiber of $\overline{\mathcal{C}}_\alpha \rightarrow \overline{\mathcal{M}}_\alpha$. It's not too hard to see that this is an open set (using elliptic regularity).

The compatibility axioms are all immediate, though let us justify the homeomorphism axiom. First, we claim that the map

$$(s_{J \setminus I} | \overline{\mathcal{M}}_{g,n}^\beta(X)_J)^{-1}(0) / \Gamma_{J \setminus I} \rightarrow U_{IJ} \subseteq \overline{\mathcal{M}}_{g,n}^\beta(X)_I$$

induced by ψ_{IJ} is a bijection. This is more or less clear: fixing an element of U_{IJ} , an inverse image must have $e_\alpha = 0$ for $\alpha \in J \setminus I$, and by definition of U_{IJ} , there exists a suitable ϕ_α which is clearly unique up to the action of Γ_α for $\alpha \in J \setminus I$. That the topologies coincide can be checked directly from their definitions.

Definition 9.2.5 (Regular locus for A^{GW} on $\overline{\mathcal{M}}_{g,n}^\beta(X)$) We define $\overline{\mathcal{M}}_{g,n}^\beta(X)_I^{\text{reg}} \subseteq \overline{\mathcal{M}}_{g,n}^\beta(X)_I$.

Let $(C, u_0, \{\phi_\alpha\}_{\alpha \in I}, \{e_\alpha\}_{\alpha \in I}) \in \overline{\mathcal{M}}_{g,n}^\beta(X)_I$. Roughly speaking, this point is contained in $\overline{\mathcal{M}}_{g,n}^\beta(X)_I^{\text{reg}}$ if and only if it has trivial automorphism group and the “vertical” (ie from a fixed domain curve) linearized version of the I -thickened $\bar{\partial}$ -equation (9.2.1) is surjective. Let us now make this precise.

Consider the smooth Banach manifold $W^{k,p}(C, X)$ for some large integer k and some $p \in (1, \infty)$. Note that the condition $u \pitchfork D_\alpha$ with exactly r_α intersections making C

³⁴When $I \neq \emptyset$, this topology is easy to define: it is simply given by using the Hausdorff distance on the image $C \subseteq X \times \prod_{\alpha \in I} \overline{\mathcal{C}}_\alpha$.

isomorphic to a fiber of $\bar{C}_\alpha/\bar{M}_\alpha$ (for all $\alpha \in I$) forms an open subset of $W^{k,p}(C, X)$ containing u_0 . Furthermore, there is a unique continuous choice of

$$(9.2.3) \quad C \times W^{k,p}(C, X) \rightarrow \bar{C}_\alpha$$

over C times a neighborhood of u_0 which extends ϕ_α on $C \times \{u_0\}$ (this holds because $\bar{M}_\alpha \rightarrow \bar{M}_{g,n+r_\alpha}/S_{r_\alpha}$ is an étale map of orbifolds). One can also check that this map (9.2.3) is highly differentiable (this depends on k being large).

Now we consider the smooth Banach bundle whose fiber over a pair $(u, \{e_\alpha\}_{\alpha \in I}) \in W^{k,p}(C, X) \times E_I$ is

$$W^{k-1,p}(\tilde{C}, \Omega_{\tilde{C}}^{0,1} \otimes_C u^*TX)$$

(where \tilde{C} is the normalization of C). Now the left-hand side of (9.2.1) (using (9.2.3) in place of ϕ_α) is a highly differentiable section of this bundle. We say that $(C, u_0, \{\phi_\alpha\}_{\alpha \in I}, \{e_\alpha\}_{\alpha \in I}) \in \bar{M}_{g,n}^\beta(X)_I^{\text{reg}}$ if and only if the following two conditions hold:

- (i) This section is transverse to the zero section at (u_0, e_α) .
- (ii) The automorphism group of $(C, u_0, \{\phi_\alpha\}_{\alpha \in I}, \{e_\alpha\}_{\alpha \in I})$ is trivial.

It is an easy exercise in elliptic regularity to show that the first condition is independent of the choice of (k, p) (as long as k is sufficiently large so that the condition makes sense).

Let

$$(9.2.4) \quad \text{vdim } \bar{M}_{g,n}^\beta(X) := \dim \bar{M}_{g,n} + (1 - g) \dim X + 2\langle c_1(X), \beta \rangle.$$

Let us now verify the transversality axioms. Freeness of the action of $\Gamma_{J \setminus I}$ on $\psi_{IJ}^{-1}(X_I^{\text{reg}})$ follows from the fact that points in X_I^{reg} have trivial automorphism group. The openness and submersion axioms follow from the following result, whose proof is given in Appendix B.

Proposition 9.2.6 (Formal regularity implies topological regularity) *For all $I \subseteq J \subseteq A^{\text{GW}}$, we have:*

- (i) $\bar{M}_{g,n}^\beta(X)_I^{\text{reg}} \subseteq \bar{M}_{g,n}^\beta(X)_J$ is an open subset.
- (ii) The map $s_{J \setminus I}: \bar{M}_{g,n}^\beta(X)_J \rightarrow E_{J \setminus I}$ is locally modeled on the projection

$$(9.2.5) \quad \mathbb{R}^{\text{vdim } \bar{M}_{g,n}^\beta(X) + \dim E_I} \times \mathbb{R}^{\dim E_{J \setminus I}} \rightarrow \mathbb{R}^{\dim E_{J \setminus I}}$$

over $\psi_{IJ}^{-1}(\bar{M}_{g,n}^\beta(X)_I^{\text{reg}}) \subseteq \bar{M}_{g,n}^\beta(X)_J$.

- (iii) There is a canonical identification of the orientation local system $\mathfrak{o}_{\bar{M}_{g,n}^\beta(X)_I^{\text{reg}}}$ with \mathfrak{o}_{E_I} (by the standard reduction to the canonical orientation of a complex linear Fredholm operator as in McDuff and Salamon [61; 62]).

Near points of $\overline{\mathcal{M}}_{g,n}^\beta(X)_I^{\text{reg}}$ with smooth domain curve, the proposition follows from standard techniques (implicit function theorem, elliptic regularity, and an index theorem). The real content of the proposition is that it holds near points with nodal domain curve.

Finally, let us verify the covering axiom.

Lemma 9.2.7 *Let $f: N^n \rightarrow M^m$ be a smooth map of smooth manifolds with $n \leq m$. Then for every $p \in N$ with df_p injective and every neighborhood U of $f(p)$, there exists a codimension- n smooth submanifold with boundary $D \subseteq M$ contained in U such that $f \pitchfork D$ and $f^{-1}(D)$ contains a point arbitrarily close to p .*

Proof After possibly shrinking U , choose a local projection $\pi: U \rightarrow T_p N$. Then let D be π^{-1} of a regular value of $\pi \circ f$ (which exists by Sard’s theorem). □

Lemma 9.2.8 (Domain stabilization for stable J -holomorphic maps) *Let $u: C \rightarrow X$ be a stable J -holomorphic map. Then there exists a smooth codimension-two submanifold with boundary $D \subseteq X$ such that $u \pitchfork D$ and adding $u^{-1}(D)$ to C as extra marked points makes C stable.*

Proof Let us observe that if $C_0 \subseteq C$ is any unstable irreducible component, we must have $u: C_0 \rightarrow X$ is nonconstant. Hence since u is J -holomorphic, it follows that there is a point on C_0 (which is not a node or marked point) where du is injective. It follows using Lemma 9.2.7 that there exists $D \subseteq X$ such that $u \pitchfork D$ and adding $u^{-1}(D)$ to C as marked points makes C stable. □

Lemma 9.2.9 (Covering axiom for A^{GW} on $\overline{\mathcal{M}}_{g,n}^\beta(X)$) *We have*

$$(9.2.6) \quad \overline{\mathcal{M}}_{g,n}^\beta(X) = \bigcup_{I \subseteq A^{\text{GW}}} \psi_{\emptyset I}((s_I | \overline{\mathcal{M}}_{g,n}^\beta(X)_I^{\text{reg}})^{-1}(0)).$$

Proof Fix a point in $\overline{\mathcal{M}}_{g,n}^\beta(X)$ (that is, a curve $u: C \rightarrow X$). We will construct $\alpha \in A^{\text{GW}}$ so that this point is contained in $\psi_{\emptyset \{\alpha\}}((s_\alpha | \overline{\mathcal{M}}_{g,n}^\beta(X)_{\{\alpha\}}^{\text{reg}})^{-1}(0))$.

First, pick $D_\alpha \subseteq X$ satisfying the conclusion of Lemma 9.2.8.

Now let $r_\alpha = \#u^{-1}(D_\alpha)$. Adding $u^{-1}(D_\alpha)$ as extra marked points to C gives a point in $\overline{\mathcal{M}}_{g,n+r_\alpha}/S_{r_\alpha}$, and we pick some local orbifold chart $\overline{\mathcal{M}}_\alpha/\Gamma_\alpha \rightarrow \overline{\mathcal{M}}_{g,n+r_\alpha}/S_{r_\alpha}$ covering this point.

Now let us consider the linearized $\bar{\partial}$ operator

$$(9.2.7) \quad D\bar{\partial}(u, \cdot): C^\infty(C, u^*TX) \rightarrow C^\infty(\tilde{C}, \Omega_{\tilde{C}}^{0,1} \otimes_{\mathbb{C}} u^*TX)$$

(\tilde{C} being the normalization of C) where $C^\infty(C, u^*TX) \subseteq C^\infty(\tilde{C}, u^*TX)$ is the subspace of functions descending continuously to C . If $(k, p) \in \mathbb{Z}_{\geq 1} \times (1, \infty)$ with $kp > 2$, then we get a corresponding $D\bar{\partial}(u, \cdot)$ map $W^{k,p} \rightarrow W^{k-1,p}$ (the restriction on (k, p) comes from the need to define $W^{k,p}(C, u^*TX) \subseteq W^{k,p}(\tilde{C}, u^*TX)$). This operator is Fredholm; in particular its cokernel is finite-dimensional.

Suppose we have a finite-dimensional vector space E_0 equipped with a linear map

$$(9.2.8) \quad \tilde{\lambda}_0: E_0 \rightarrow C^\infty(\tilde{C}, \Omega_{\tilde{C}}^{0,1} \otimes_{\mathbb{C}} u^*TX)$$

supported away from (the inverse image in \tilde{C} of) the nodes and marked points of C . Then the following conditions are equivalent (exercise using elliptic regularity):

- (i) (9.2.8) is surjective onto $\text{coker}(D\bar{\partial}(u, \cdot)): C^\infty \rightarrow C^\infty$.
- (ii) (9.2.8) is surjective onto $\text{coker}(D\bar{\partial}(u, \cdot)): W^{k,p} \rightarrow W^{k-1,p}$ for some (k, p) .
- (iii) (9.2.8) is surjective onto $\text{coker}(D\bar{\partial}(u, \cdot)): W^{k,p} \rightarrow W^{k-1,p}$ for all (k, p) .

There exists such a pair E_0 and $\tilde{\lambda}_0$ satisfying these conditions since we may choose $k = 1$ and $p > 2$ and then remember that C^∞ functions supported away from the nodes and marked points are dense in $W^{0,p} = L^p$.

Now pick any isomorphism to a fiber $\phi_\alpha: C \rightarrow \bar{C}_\alpha$ and extend $\tilde{\lambda}_0$ to a map

$$\lambda_0: E_0 \rightarrow C^\infty(\bar{C}_\alpha \times X, \Omega_{\bar{C}_\alpha/\bar{M}_\alpha}^{0,1} \otimes_{\mathbb{C}} TX)$$

supported away from the nodes and marked points of the fibers. Define $E_\alpha = E_0[\Gamma]$ and define

$$\lambda_\alpha: E_\alpha \rightarrow C^\infty(\bar{C}_\alpha \times X, \Omega_{\bar{C}_\alpha/\bar{M}_\alpha}^{0,1} \otimes_{\mathbb{C}} TX)$$

by Γ_α -linear extension from λ_0 . It is now clear by definition that $u: C \rightarrow X$ is covered by $\bar{M}_{g,n}^\beta(X)_{\{\alpha\}}^{\text{reg}}$, where $\alpha \in A^{\text{GW}}$ is the element just constructed. □

We have now shown the following.

Theorem 9.2.10 A^{GW} is an implicit atlas on $\bar{M}_{g,n}^\beta(X)$.

9.3 Definition of Gromov–Witten invariants

Definition 9.3.1 (Gromov–Witten invariants) Fix nonnegative integers g and n with $2g + n > 2$, and fix $\beta \in H_2(X; \mathbb{Z})$. By Proposition 9.2.6, the virtual orientation sheaf induced by A^{GW} is canonically trivialized. Thus the implicit atlas A^{GW} induces a virtual fundamental class

$$[\bar{M}_{g,n}^\beta(X)]^{\text{vir}} \in \check{H}^\bullet(\bar{M}_{g,n}^\beta(X); \mathbb{Q})^\vee.$$

We define the *Gromov–Witten invariant*

$$(9.3.1) \quad \text{GW}_{g,n}^\beta(X) \in H_\bullet(\overline{\mathcal{M}}_{g,n} \times X^n; \mathbb{Q})$$

as the pushforward of $[\overline{\mathcal{M}}_{g,n}^\beta(X)]^{\text{vir}}$ under the tautological map $\overline{\mathcal{M}}_{g,n}^\beta(X) \rightarrow \overline{\mathcal{M}}_{g,n} \times X^n$. This is well-defined by the next lemma.

Lemma 9.3.2 $\text{GW}_{g,n}^\beta(X) \in H_\bullet(\overline{\mathcal{M}}_{g,n} \times X^n)$ is independent of the choice of J .

Proof Let J_0 and J_1 be any two smooth ω -tame almost complex structures on X . We denote by $\overline{\mathcal{M}}_{g,n}^\beta(X, J_0)$ and $\overline{\mathcal{M}}_{g,n}^\beta(X, J_1)$ the corresponding moduli spaces of stable maps (denoted earlier by simply $\overline{\mathcal{M}}_{g,n}^\beta(X)$ when we considered just a fixed J).

There exists a smooth path of ω -tame almost complex structures $J_{[0,1]} = \{J_t\}_{t \in [0,1]}$ connecting J_0 and J_1 . Let us consider the corresponding “parametrized” moduli space of stable maps $\overline{\mathcal{M}}_{g,n}^\beta(X, J_{[0,1]})$. The construction from Section 9.2 gives an implicit atlas with boundary A^{GW} on $\overline{\mathcal{M}}_{g,n}^\beta(X, J_{[0,1]})$ whose restriction to

$$\partial \overline{\mathcal{M}}_{g,n}^\beta(X, J_{[0,1]}) := \overline{\mathcal{M}}_{g,n}^\beta(X, J_0) \sqcup \overline{\mathcal{M}}_{g,n}^\beta(X, J_1)$$

agrees with A^{GW} on these spaces. Hence it follows using Lemmas 5.2.3 and 5.2.5 that

$$[\overline{\mathcal{M}}_{g,n}^\beta(X, J_0)]^{\text{vir}} = [\overline{\mathcal{M}}_{g,n}^\beta(X, J_1)]^{\text{vir}} \quad \text{in } \check{H}^\bullet(\overline{\mathcal{M}}_{g,n}^\beta(X, J_{[0,1]}))^\vee,$$

which is enough. □

10 Hamiltonian Floer homology

In this section, we define Hamiltonian Floer homology for a general closed symplectic manifold M (which we now fix). We also calculate Hamiltonian Floer homology using the S^1 -localization idea of Floer, and we derive the Arnold conjecture from this calculation. These results (in this generality) are originally due to Liu and Tian [54], Fukaya and Ono [33] and Ruan [71]. For a general introduction to Hamiltonian Floer homology, the reader may consult Salamon [72] (we assume some familiarity with the basic theory).

The main content of this section is the construction of implicit atlases on the relevant spaces of stable pseudo-holomorphic cylinders. Once we do this, the definition from Section 7 gives the desired homology groups. We also construct S^1 -equivariant implicit atlases on the moduli spaces for time-independent Hamiltonians. This allows us to use the S^1 -localization results of Section 8 to show that Hamiltonian Floer homology coincides with Morse homology (a standard corollary of which is the Arnold conjecture).

The necessary gluing results are stated in Propositions 10.3.3 and 10.6.2 and are proved in Appendix C.

It would be interesting to define this isomorphism as in Piunikhin, Salamon and Schwarz [69] using their moduli spaces of “spiked disks” (this route avoids the use of S^1 -localization).

10.1 Preliminaries

Definition 10.1.1 (Abelian cover of free loop space) Let L_0M denote the space of null-homotopic smooth maps $S^1 \rightarrow M$, and let $\widetilde{L_0M}$ denote the space of such loops together with a homology class of bounding 2-disk. Then $\widetilde{L_0M} \rightarrow L_0M$ is a π -cover, where $\pi = \text{im}(\pi_2(M) \rightarrow H_2(M; \mathbb{Z}))$.³⁵

Definition 10.1.2 (Hamiltonian flows) For a smooth function $H: M \times S^1 \rightarrow \mathbb{R}$, let $X_H: M \times S^1 \rightarrow TM$ denote the Hamiltonian vector field induced by H , and let $\phi_H: M \rightarrow M$ denote the time-1 flow map of X_H . A *periodic orbit* of H is a smooth function $\gamma: S^1 \rightarrow M$ satisfying $\gamma'(t) = X_{H(t)}(\gamma(t))$. Let $C^\infty(M \times S^1)^{\text{reg}} \subseteq C^\infty(M \times S^1)$ denote those functions H for which ϕ_H has non-degenerate fixed points.

Definition 10.1.3 (Simplicial sets of H and J) Define the simplicial set $H_\bullet(M)$, where $H_n(M)$ is the set of smooth functions $H: \Delta^n \rightarrow C^\infty(M \times S^1)$ which are constant near the vertices and send the vertices to $C^\infty(M \times S^1)^{\text{reg}}$. Define the simplicial set $J_\bullet(M)$ where $J_n(M)$ is the set of smooth functions $J: \Delta^n \rightarrow J(M)$ which are constant near the vertices ($J(M)$ is the space of smooth almost complex structures tamed by ω) and which send the vertices to almost complex structures which are ω -compatible.

It is easy to see that $H_\bullet(M)$ and $J_\bullet(M)$ are both *contractible Kan complexes*. A semisimplicial set Z_\bullet is a contractible Kan complex if and only if every map $\partial\Delta^n \rightarrow Z_\bullet$ can be extended to a map $\Delta^n \rightarrow Z_\bullet$ for all $n \geq 0$ (where Δ^n is the semisimplicial n -simplex).

Let $JH_\bullet(M) = J_\bullet(M) \times H_\bullet(M)$, which of course is also a contractible Kan complex.

Definition 10.1.4 (Standard Morse function on Δ^n) For this definition, let us view the n -simplex Δ^n as

$$(10.1.1) \quad \Delta^n = \{x \in [0, 1]^{n+1} : 0 = x_0 \leq \dots \leq x_n \leq 1\}.$$

³⁵We could as easily work on the smaller cover corresponding to the image of $\omega \oplus c_1(M): \pi_2(M) \rightarrow \mathbb{R} \oplus \mathbb{Z}$. The corresponding equivalence relation is that $f_1, f_2: D^2 \rightarrow M$ with $f_1|_{S^1} = f_2|_{S^1}$ are equivalent if and only if ω and $c_1(M)$ both vanish on “ $f_1 - f_2 \in \pi_2(M)$ ”.

The i^{th} vertex of Δ^n is given by $x_{n-i} = 0$ and $x_{n-i+1} = 1$. We now consider the Morse function on Δ^n given by $f(x) := \sum_{i=1}^n \cos \pi x_i$. Its gradient

$$(10.1.2) \quad \nabla f(x) = \sum_{i=1}^n \pi \sin(\pi x_i) \frac{\partial}{\partial x_i}$$

is tangent to the boundary of Δ^n , and its critical points are precisely the vertices of Δ^n , the index at vertex i being $n - i$. Note also that for any facet inclusion $\Delta^k \hookrightarrow \Delta^n$, the pushforward of ∇f is again ∇f .

Let us consider the space $\mathcal{F}(\Delta^n)$ of broken Morse flow lines from vertex 0 (index n) to vertex n (index 0). This space is homeomorphic to a cube $[0, 1]^{n-1}$, the factors $[0, 1]$ being naturally indexed by the “middle vertices” $1, \dots, n - 1$ of Δ^n . A flow line is broken at a vertex $i \in \{1, \dots, n - 1\}$ if and only if the corresponding coordinate in $[0, 1]^{n-1}$ equals 1. For any simplex σ there are canonical compatible product/face maps

$$(10.1.3) \quad \mathcal{F}(\sigma|[0 \dots k]) \times \mathcal{F}(\sigma|[k \dots n]) \rightarrow \mathcal{F}(\sigma),$$

$$(10.1.4) \quad \mathcal{F}(\sigma|[0 \dots \hat{k} \dots n]) \rightarrow \mathcal{F}(\sigma).$$

See also Section C.13.1 for more details. Adams [3] considered spaces of paths on Δ^n with the same key properties (though it is important that our flow lines are smooth, whereas Adams’ are not).

10.2 Moduli space of Floer trajectories

Let us now define the flow category diagram (Definition 7.2.1) which gives rise to Hamiltonian Floer homology.

Definition 10.2.1 For $H \in C^\infty(M \times S^1)^{\text{reg}}$, let $\mathcal{P}_H \subseteq \widetilde{L_0 M}$ denote the set of null-homotopic periodic orbits equipped with a homology class of bounding 2-disk.

In the following definition, the reader may prefer to focus on the cases $n = 0$ (Floer trajectories relevant for the differential), $n = 1$ (Floer trajectories relevant for the continuation maps), and $n = 2$ (Floer trajectories relevant for the homotopies between continuation maps).

Definition 10.2.2 (Floer trajectory) Let $\sigma \in J_n(M) \times H_n(M)$ be an n -simplex; we denote by $H^\sigma: \Delta^n \times M \times S^1 \rightarrow \mathbb{R}$ and $J^\sigma: \Delta^n \rightarrow J(M)$ the corresponding smooth families. Let $p \in \mathcal{P}_{H_0}$ and $q \in \mathcal{P}_{H_n}$ be periodic orbits, where $H_i = H^\sigma(i \in \Delta^n, \cdot, \cdot)$ is the Hamiltonian associated to the i^{th} vertex of Δ^n . A Floer trajectory of type (σ, p, q) is a triple (C, ℓ, u) , where:

- (i) C is a nodal curve of type $(0, 2)$. Let us call the two marked points $x^-, x^+ \in C$, and let $k = k(C)$ be the number of vertices (irreducible components of C) on the unique path from x^- to x^+ in the dual graph of C .
- (ii) $\ell: \coprod_{i=1}^k \mathbb{R} \rightarrow \Delta^n$ is a broken Morse flow line from vertex 0 to vertex n (for the Morse function in Definition 10.1.4). Let $0 = v_0 \leq \dots \leq v_k = n$ be the corresponding sequence of vertices. We allow ℓ to contain constant flow lines, ie we allow $v_i = v_{i+1}$.
- (iii) $u: C \rightarrow M \times S^1 \times \coprod_{i=1}^k \mathbb{R}$ is a smooth *building* of type (σ, p, q) , by which we mean the following. Let C° be C punctured at $\{x^-, x^+\}$ and at the nodes corresponding to the edges in the unique path in the dual graph of C from x^- to x^+ . The connected components $\{C_1^\circ, \dots, C_k^\circ\}$ are naturally ordered (x^- on the first component and x^+ on the k^{th} component). There must be periodic orbits $\{\gamma_0, \dots, \gamma_k\}$ where $\gamma_i \in \mathcal{P}_{H, v_i}$ with $\gamma_0 = p$ and $\gamma_k = q$. Then the negative and positive ends of C_i° must be asymptotic to $(\gamma_{i-1}(t), t)$ and $(\gamma_i(t), t)$, respectively (with multiplicity one). We also require that u have “finite energy”. In addition, $u|_{C_i^\circ}$ must be in the correct homology class: the element of $\pi_2(M)$ obtained by gluing together $u|_{C_i^\circ}$ (resolving any nodes of C_i°) with the given disks bounding γ_i and γ_{i-1} must vanish in homology.
- (iv) u is pseudo-holomorphic with respect to the almost complex structure on $M \times S^1 \times \coprod_{i=1}^k \mathbb{R}$ defined as follows. Use coordinates $(t, s) \in S^1 \times \coprod_{i=1}^k \mathbb{R}$. Fix the standard almost complex structure on $S^1 \times \coprod_{i=1}^k \mathbb{R}$, namely

$$J_{S^1 \times \coprod_{i=1}^k \mathbb{R}} \left(\frac{\partial}{\partial s} \right) = \frac{\partial}{\partial t}.$$

Also fix the (s -dependent) almost complex structure $J^\sigma(\ell(s))$ on M . We let $A: T[S^1 \times \coprod_{i=1}^k \mathbb{R}] \rightarrow TM$ be defined by $A(\partial/\partial s) = X_{H^\sigma(\ell(s), t, \cdot)}$ and extended anti-holomorphically. Now we use the following almost complex structure on $M \times S^1 \times \coprod_{i=1}^k \mathbb{R}$:

$$(10.2.1) \quad J = \begin{pmatrix} J^\sigma(\ell(s)) & A \\ 0 & J_{S^1 \times \coprod_{i=1}^k \mathbb{R}} \end{pmatrix}.$$

Note that the projection $M \times S^1 \times \coprod_{i=1}^k \mathbb{R} \rightarrow S^1 \times \coprod_{i=1}^k \mathbb{R}$ is holomorphic.

An *isomorphism* $v: (C, \ell, u) \rightarrow (C', \ell', u')$ of Floer trajectories is an isomorphism $\iota_1: C \rightarrow C'$ of curves of type $(0, 2)$ and an isomorphism $\iota_2: \coprod_{i=1}^k \mathbb{R} \rightarrow \coprod_{i=1}^k \mathbb{R}$ (acting by translation on each factor and respecting the ordering of the terms; note that the existence of ι_1 implies that $k(C) = k(C')$) such that $u = (\text{id}_{M \times S^1} \times \iota_2^{-1}) \circ u' \circ \iota_1$ and $\ell = \ell' \circ \iota_2$. We say a Floer trajectory is *stable* if and only if its automorphism group (ie group of self-isomorphisms) is finite.

Definition 10.2.3 We define a π -equivariant flow category diagram $\overline{\mathcal{M}}/JH_\bullet(M)$ as follows.

- (i) For a vertex $(J, H) \in J_0(M) \times H_0(M)$, we let $\mathcal{P}_{(J,H)} = \mathcal{P}_H$.
- (ii) The grading $\text{gr}: \mathcal{P} \rightarrow \mathbb{Z}$ is the usual Conley–Zehnder index, and $\text{gr}: \pi \rightarrow \mathbb{Z}$ is given by $\text{gr}(h) = 2\langle c_1(TM), h \rangle$.
- (iii) The action $a: \mathcal{P} \rightarrow \mathbb{R}$ is the usual symplectic action, and $a: \pi \rightarrow \mathbb{R}$ is given by $a(h) = \langle \omega, h \rangle$.
- (iv) We let $\overline{\mathcal{M}}(\sigma, p, q)$ be the set of stable Floer trajectories of type (σ, p, q) , equipped with the Gromov topology. It is well-known that $\overline{\mathcal{M}}(\sigma, p, q)$ is compact Hausdorff. The finiteness conditions required on $\overline{\mathcal{M}}(\sigma, p, q)$ also follow from Gromov compactness. The product/face maps on $\overline{\mathcal{M}}(\sigma, p, q)$ are evident.
- (v) The action of π on everything is clear.

10.3 Implicit atlas

Let us now define an implicit atlas on the flow category diagram $\overline{\mathcal{M}}/JH_\bullet(M)$ (recall Definition 7.2.4). This construction follows the same outline as the construction of an implicit atlas on the moduli space of stable maps in Section 9.2 (the main difference being that here there is more notation to keep track of). Note that the flow category diagram on which we will define an implicit atlas is no longer denoted \mathcal{X}/Z_\bullet , and this leads to a few (evident) notational differences from Section 7, where we considered implicit atlases on flow category diagrams abstractly.

Definition 10.3.1 (Index set $A^{\text{HF}}(\Delta^n)$) A (Hamiltonian Floer) thickening datum α on the simplex Δ^n is a quadruple $(D_\alpha, r_\alpha, E_\alpha, \lambda_\alpha)$, where:

- (i) $D_\alpha \subseteq M \times S^1 \times \Delta^n$ is a compact smooth submanifold with corners locally modeled on

$$\mathbb{R}_{\geq 0}^N \times \mathbb{R}^{N'} \subseteq \mathbb{R}_{\geq 0}^N \times \mathbb{R}^{N'+2} \quad \text{or} \quad \mathbb{R}_{\geq 0}^{N+1} \times \mathbb{R}^{N'} \subseteq \mathbb{R}_{\geq 0}^N \times \mathbb{R}^{N'+3}.$$

Let us denote by $\partial^{\text{ess}} D_\alpha \subseteq \partial D_\alpha$ the closure of $\partial D_\alpha \setminus [M \times S^1 \times \partial \Delta^n]$ (which is precisely the set of points with local model of the second type).

- (ii) $r_\alpha \geq 1$ is an integer; let $\Gamma_\alpha = S_{r_\alpha}$.
- (iii) E_α is a finitely generated $\mathbb{R}[S_{r_\alpha}]$ -module.
- (iv) The map

$$\lambda_\alpha: E_\alpha \rightarrow C^\infty(\overline{\mathcal{C}}_{0,2+r_\alpha} \times M, \Omega_{\overline{\mathcal{C}}_{0,2+r_\alpha}/\overline{\mathcal{M}}_{0,2+r_\alpha}}^{0,1} \otimes_{\mathbb{R}} TM)$$

is an S_{r_α} -equivariant linear map supported away from the nodes and marked points of the fibers of $\overline{\mathcal{C}}_{0,2+r_\alpha} \rightarrow \overline{\mathcal{M}}_{0,2+r_\alpha}$ (the universal family).

Let $A^{\text{HF}}(\Delta^n)$ denote the set of all thickening datums on Δ^n .

We define a π -equivariant implicit atlas \mathcal{A} on the flow category diagram $\overline{\mathcal{M}}/JH_\bullet(M)$ as follows. We set $\overline{\mathcal{A}}(\sigma, p, q) := A^{\text{HF}}(\sigma)$. Let us now define the implicit atlas on $\overline{\mathcal{M}}(\sigma, p, q)^{\leq s}$ with index set $\mathcal{A}(\sigma, p, q)^{\geq s}$ (built from $\overline{\mathcal{A}}(\sigma, p, q)$ as in (7.2.1)).

Definition 10.3.2 (Atlas data for $\mathcal{A}(\sigma, p, q)^{\geq s}$ on $\overline{\mathcal{M}}(\sigma, p, q)^{\leq s}$) For a subset $I \subseteq \mathcal{A}(\sigma, p, q)^{\geq s}$, an I -thickened Floer trajectory of type $(\sigma, p, q)^{\leq s}$ is defined to be a tuple $(C, \ell, u, \{\phi_\alpha\}_{\alpha \in I}, \{e_\alpha\}_{\alpha \in I})$, where:

- (i) C is a nodal curve of type $(0, 2)$.
- (ii) $\ell: \coprod_{i=1}^k \mathbb{R} \rightarrow \Delta^n$ is a broken flow line from vertex 0 to vertex n .
- (iii) $u: C \rightarrow M \times S^1 \times \coprod_{i=1}^k \mathbb{R}$ is a smooth building of type (σ, p, q) (in the sense of Definition 10.2.2(iii)), with combinatorial type of u belonging to $\mathcal{S}_{\overline{\mathcal{M}}}(\sigma, p, q)^{\leq s}$. Recall that by definition (see (7.2.1)), $\mathcal{A}(\sigma, p, q)^{\geq s}$ is a disjoint union of various $\overline{\mathcal{A}}(\sigma|[i_0 \dots i_n], p', q') := A^{\text{HF}}(\sigma|[i_0 \dots i_n])$. Hence any given $\alpha \in I$ comes from one of these, say $\overline{\mathcal{A}}(\sigma|[i_0^\alpha \dots i_n^\alpha], p'_\alpha, q'_\alpha)$. Let $C_\alpha \subseteq C$ denote the union of irreducible components corresponding to this triple, which exists because the combinatorial type of u belongs to $\mathcal{S}_{\overline{\mathcal{M}}}(\sigma, p, q)^{\leq s}$ (this C_α is a key notion for the present construction of an implicit atlas).
- (iv) For all $\alpha \in I$, we must have $u|_{C_\alpha} \pitchfork D_\alpha$ with exactly r_α intersections. By $u|_{C_\alpha} \pitchfork D_\alpha$, we mean that under the map

$$(\text{id}_{M \times S^1} \times \ell) \circ u: C_\alpha \rightarrow M \times S^1 \times \Delta^{\sigma|[i_0 \dots i_m]}$$

we have³⁶ $\overline{C_\alpha} \cap \partial^{\text{ess}} D_\alpha = \emptyset$ and for every point $p \in C_\alpha$ mapping to D_α , the derivative

$$d((\text{id}_{M \times S^1} \times \ell) \circ u): T_p C \rightarrow T_{u(p)}[M \times S^1 \times \Delta^{\sigma|[i_0 \dots i_m]}/T_{u(p)} D_\alpha$$

is surjective and p is not a node or marked point of C .

- (v) $\phi_\alpha: C_\alpha \rightarrow \overline{\mathbb{C}}_{0, 2+r_\alpha}$ is an isomorphism with a fiber (where C_α is considered to have two marked points x^-, x^+ corresponding to p'_α, q'_α plus the r_α marked points $(u|_{C_\alpha})^{-1}(D_\alpha)$).
- (vi) $e_\alpha \in E_\alpha$.
- (vii) The I -thickened $\bar{\partial}$ -equation

$$(10.3.1) \quad \bar{\partial}u + \sum_{\alpha \in I} \lambda_\alpha(e_\alpha)(\phi_\alpha, u) = 0$$

³⁶The closure of the image $\overline{C_\alpha}$ is precisely the image of C_α union the asymptotic periodic orbits.

is satisfied, where we use the almost complex structure on $M \times S^1 \times \coprod_{i=1}^k \mathbb{R}$ defined in (10.2.1). The term $\lambda_\alpha(e_\alpha)(\phi_\alpha, u)$ only makes sense over C_α ; we interpret it as zero over the rest of C . Note that for (10.3.1), we project λ_α onto $\Omega_{\mathbb{C}^{0,2+r_\beta}/\overline{\mathcal{M}}_{0,2+r_\beta}}^{0,1} \otimes_{\mathbb{C}} TM_{J^\sigma}(\ell(s))$.

An *isomorphism* between two I -thickened Floer trajectories $(C, \ell, u, \{e_\alpha\}, \{\phi_\alpha\})$ and $(C', \ell', u', \{e'_\alpha\}, \{\phi'_\alpha\})$ is an isomorphism $\iota_1: C \rightarrow C'$ of curves of type $(0, 2)$ and an isomorphism $\iota_2: \coprod_{i=1}^k \mathbb{R} \rightarrow \coprod_{i=1}^k \mathbb{R}$ (acting by translation on each factor and respecting the ordering of the terms; note that the existence of ι_1 implies that $k(C) = k(C')$) such that

$$u = (\text{id}_{M \times S^1} \times \iota_2^{-1}) \circ u' \circ \iota_1, \quad \ell = \ell' \circ \iota_2, \quad \phi_\alpha = \phi'_\alpha \circ \iota_1, \quad e_\alpha = e'_\alpha \quad \text{for all } \alpha \in I.$$

We say an I -thickened Floer trajectory is *stable* if and only if its automorphism group (ie group of self-isomorphisms) is finite.

Let $\overline{\mathcal{M}}(\sigma, p, q)_{\overline{I}}^{\leq s}$ denote the set of stable I -thickened Floer trajectories of type $(\sigma, p, q)^{\leq s}$, and equip it with the Gromov topology. The actions of Γ_I on the thickened moduli spaces, the functions s_I , the projections ψ_{IJ} , and the sets U_{IJ} are all defined as in Definition 9.2.4 from the Gromov–Witten setting.

The stratification of $\overline{\mathcal{M}}(\sigma, p, q)_{\overline{I}}^{\leq s}$ by $\mathcal{S}_{\overline{\mathcal{M}}}(\sigma, p, q)^{\leq s}$ is evident.

The compatibility axioms for $\mathcal{A}(\sigma, p, q)^{\geq s}$ on $\overline{\mathcal{M}}(\sigma, p, q)^{\leq s}$ are all immediate; the homeomorphism axiom can again be verified directly as in the Gromov–Witten setting.

The regular loci for $\mathcal{A}(\sigma, p, q)^{\geq s}$ on $\overline{\mathcal{M}}(\sigma, p, q)^{\leq s}$ are defined as in Definition 9.2.5, meaning that a 5-tuple $(C, \ell, u, \{\phi_\alpha\}_{\alpha \in I}, \{e_\alpha\}_{\alpha \in I})$ is regular if and only if it has trivial automorphism group and the linearized operator (fixing C and varying u, ℓ , and $\{e_\alpha\}_{\alpha \in I}$) is surjective (see Section C.1.6 for more details). Let

$$\text{vdim } \overline{\mathcal{M}}(\sigma, p, q)^{\leq s} := \text{gr}(q) - \text{gr}(p) + \dim \sigma - 1 - \text{codim } s.$$

Let us now discuss the transversality axioms for $\mathcal{A}(\sigma, p, q)^{\geq s}$ on $\overline{\mathcal{M}}(\sigma, p, q)^{\leq s}$. The openness and submersion axioms follow from the following result, whose proof is given in Appendix C.

Proposition 10.3.3 (Formal regularity implies topological regularity) *For all $I \subseteq J \subseteq \mathcal{A}(\sigma, p, q)^{\geq s}$, we have:*

- (i) $(\overline{\mathcal{M}}(\sigma, p, q)_{\overline{I}}^{\leq s})^{\text{reg}} \subseteq \overline{\mathcal{M}}(\sigma, p, q)_{\overline{I}}^{\leq s}$ is an open subset.
 - (ii) The map $s_{J \setminus I}: \overline{\mathcal{M}}(\sigma, p, q)_{\overline{J}}^{\leq s} \rightarrow E_{J \setminus I}$ over the locus $\psi_{IJ}^{-1}((\overline{\mathcal{M}}(\sigma, p, q)_{\overline{I}}^{\leq s})^{\text{reg}}) \subseteq \overline{\mathcal{M}}(\sigma, p, q)_{\overline{J}}^{\leq s}$ is locally modeled on the projection
- $$(10.3.2) \quad \mathbb{R}^{\text{vdim } \overline{\mathcal{M}}(\sigma, p, q)^{\leq s} + \dim E_I} \times \mathbb{R}^{\dim E_{J \setminus I}} \rightarrow \mathbb{R}^{\dim E_{J \setminus I}}$$

over the top stratum $\mathfrak{s} \in \mathcal{S}_{\overline{\mathcal{M}}}(\sigma, p, q)^{\leq s}$. More generally, the local model (compatible with stratifications) is given by replacing the first factor on the left by $\mathbb{R}_{\geq 0}^n \times \mathbb{R}^{n'}$ stratified appropriately by $\mathcal{S}_{\overline{\mathcal{M}}}(\sigma, p, q)$.

- (iii) There exist π -invariant coherent trivializations of the local systems $\mathfrak{o}_{E_I}^{\vee} \otimes \mathfrak{o}_{(\overline{\mathcal{M}}(\sigma, p, q)_{\overline{J}}^{\leq s})^{\text{reg}}}$ (in the sense of Definition 7.2.6).

Finally, to verify the covering axiom, we use the general strategy from Lemma 9.2.9 in the Gromov–Witten case. To apply this in the present setting, we just need the following stabilization lemma to take the place of Lemma 9.2.8.

Lemma 10.3.4 (Domain stabilization for stable Floer trajectories) *Suppose that $\ell: \coprod_{i=1}^k \mathbb{R} \rightarrow \Delta^n$ and $u: C \rightarrow M \times S^1 \times \coprod_{i=1}^k \mathbb{R}$ is a point in $\overline{\mathcal{M}}(\sigma, p, q)$. Then there exists $D \subseteq M \times S^1 \times \Delta^n$ as in Definition 10.3.1(i) such that $C \pitchfork D$ in the sense of Definition 10.3.2(iv) and so that adding the intersections to C as extra marked points makes C stable.*

Proof As in the proof of Lemma 9.2.8, we use (an appropriate variant for manifolds with corners of) Lemma 9.2.7 (which we may also use to avoid the periodic orbits in question). It thus suffices to show that for any unstable component C_0 of C , there exists a point $p \in C_0$ where $d((\text{id}_{M \times S^1} \times \ell) \circ u): TC_0 \rightarrow TM \times TS^1 \times T\Delta^n$ is injective. To find such a point, we split into two cases.

First, suppose $u: C_0 \rightarrow S^1 \times \mathbb{R}$ is constant. Then $u: C_0 \rightarrow M$ is a (nonconstant!) J -holomorphic sphere, and thus has a point of injectivity of du .

Second, suppose $u: C_0 \rightarrow S^1 \times \mathbb{R}$ is not constant. Then $u: C_0 \rightarrow S^1 \times \mathbb{R}$ is an isomorphism; let us use $(t, s) \in S^1 \times \mathbb{R}$ as coordinates on C_0 . Now $\partial u / \partial t \in TM \times TS^1$ has nonzero coordinate in the TS^1 component (everywhere), and $\partial u / \partial s \in TM \times TS^1$ has zero coordinate in the TS^1 coordinate (everywhere). Hence if du is everywhere noninjective, we find that $\partial u / \partial s = 0 \in TM \times TS^1$ everywhere. It follows that $u: C_0 \rightarrow M \times S^1$ is independent of the \mathbb{R} coordinate, and thus is simply a trivial cylinder mapping onto a periodic orbit. Now the stability condition (the automorphism group being finite) implies that the corresponding (piece of a) flow line $\ell: \mathbb{R} \rightarrow \Delta^n$ is nontrivial. It follows that we have the desired injectivity. \square

Theorem 10.3.5 \mathcal{A} is a π -equivariant implicit atlas on $\overline{\mathcal{M}}/JH_\bullet(M)$.

Proof We have shown above that each individual $\mathcal{A}(\sigma, p, q)^{\geq s}$ is an implicit atlas on $\overline{\mathcal{M}}(\sigma, p, q)^{\leq s}$. The required compatibility isomorphisms between these implicit atlases follow directly from the definition. \square

10.4 Definition of Hamiltonian Floer homology

Definition 10.4.1 (Hamiltonian Floer homology) We have a π -equivariant flow category diagram $\overline{\mathcal{M}}/JH_\bullet(M)$ equipped with an implicit atlas \mathcal{A} . Moreover, we obtain coherent orientations ω from Proposition 10.3.3. Hence according to Definition 7.7.6, we get a diagram $F\mathbb{H}: JH_\bullet(M) \rightarrow H^0(\text{Ch}_{\mathbb{Q}[[\pi]]})$ (the hypotheses of Definition 7.7.6 can be easily verified). Since $JH_\bullet(M)$ is a contractible Kan complex, this is really just a single object $F\mathbb{H}^\bullet(M) \in H^0(\text{Ch}_{\mathbb{Q}[[\pi]]})$ which we call the *Hamiltonian Floer homology of M* .

Remark 10.4.2 The ring $\mathbb{Q}[[\pi]]$ is the *graded completion* of $\mathbb{Q}[\pi]$; see Definition 7.7.2.

10.5 S^1 -invariant Hamiltonians

To calculate the Hamiltonian Floer homology $F\mathbb{H}^\bullet(M)$ as defined above, we consider the case when H is a (time-independent) Morse function on M .

Fix a smooth almost complex structure J on M compatible with ω . This induces a metric on M , so there is a notion of gradient flow line for smooth functions on M .

Let $H: M \rightarrow \mathbb{R}$ be a Morse function for which the time-1 Hamiltonian flow map of H has non-degenerate fixed points, all of which are critical points of H (for example, $H = \epsilon \cdot H_0$ is such a function for any Morse function $H_0: M \rightarrow \mathbb{R}$ and sufficiently small $\epsilon > 0$). Consider the inclusion $* \hookrightarrow JH_\bullet(M)$ (where $*$ is the simplicial 0-simplex, ie the simplicial set with a single n -simplex for all n) defined by mapping everything to the constant families of almost complex structures and Hamiltonians given by J and H . We will restrict attention to the (pullback) flow category diagram $\overline{\mathcal{M}}/*$. With our assumptions on H , the set of generators is $\mathcal{P} = \text{crit}(H) \times \pi$ canonically, and the grading on \mathcal{P} is given by the Morse index on $\text{crit}(H)$ plus $\text{gr}: \pi \rightarrow \mathbb{Z}$.

Now there is a canonical S^1 -action on the spaces of stable Floer trajectories in $\overline{\mathcal{M}}/*$ (postcompose u with a rotation of S^1) which is compatible with the product/face maps (this action exists since H is independent of the S^1 -coordinate). It follows that $\overline{\mathcal{M}}^{S^1}/*$ is also a flow category diagram (defined using the fixed locus $\overline{\mathcal{M}}(\sigma, p, q)^{S^1} \subseteq \overline{\mathcal{M}}(\sigma, p, q)$).

Definition 10.5.1 (Morse trajectory) Let $\sigma \in *$ be the n -simplex, and let $p, q \in \mathcal{P} = \text{crit}(H) \times \pi$. A *Morse trajectory of type (σ, p, q)* is a triple (k, ℓ, u) , where:

- (i) $k \geq 1$ is a positive integer.
- (ii) $\ell: \coprod_{i=1}^k \mathbb{R} \rightarrow \Delta^n$ is a broken Morse flow line from vertex 0 to vertex n (for the Morse function in Definition 10.1.4). Let $0 = v_0 \leq \dots \leq v_k = n$ be the corresponding sequence of vertices. We allow ℓ to contain constant flow lines, ie we allow $v_i = v_{i+1}$.
- (iii) $u: \coprod_{i=1}^k \mathbb{R} \rightarrow M$ is a broken Morse flow line from p to q for the function H . We allow u to contain constant flow lines.
- (iv) The π -components of p and q agree.

An isomorphism $\iota: (k, \ell, u) \rightarrow (k', \ell', u')$ of Morse trajectories is an isomorphism $\iota: \coprod_{i=1}^k \mathbb{R} \rightarrow \coprod_{i=1}^{k'} \mathbb{R}$ (acting by translation on each factor and respecting the ordering of the terms; we require $k = k'$) such that $u = u' \circ \iota$ and $\ell = \ell' \circ \iota$. We say a Morse trajectory is *stable* if and only if its automorphism group (ie group of self-isomorphisms) is finite.

Let $\overline{\mathcal{M}}_{\text{morse}}(\sigma, p, q)$ denote the space of stable Morse trajectories.

Proposition 10.5.2 (Formal regularity implies topological regularity) *Suppose H is Morse–Smale. Then with the stratification by k , the spaces $\overline{\mathcal{M}}_{\text{morse}}(\sigma, p, q)$ are compact topological manifolds with corners.*

For the case of the zero simplex $\sigma = \sigma^0$, this is proved by Wehrheim [80]. In fact, the general case also follows from [80] since $\overline{\mathcal{M}}_{\text{morse}}(\sigma, p, q)$ is the space of broken Morse flow lines on $M \times \Delta^n$ from $p \times 0$ to $q \times n$ for the Morse function $H + f$ (and f is defined and smooth on a neighborhood of $\Delta^n \subseteq \mathbb{R}^n$). Alternatively, one may restrict to H which have a particular normal form near each critical point in which case this result holds by more elementary arguments (see also [80]).

Lemma 10.5.3 *There is canonical homeomorphism $\overline{\mathcal{M}}(\sigma, p, q)^{S^1} = \overline{\mathcal{M}}_{\text{morse}}(\sigma, p, q)$.*

Proof Suppose a stable Floer trajectory (C, ℓ, u) is S^1 -invariant. Then each of the components $\{C_1^\circ, \dots, C_k^\circ\}$ must be smooth (no nodes) and hence isomorphic to $S^1 \times \mathbb{R}$. Using the holomorphic projection $M \times S^1 \times \coprod_{i=1}^k \mathbb{R} \rightarrow S^1 \times \coprod_{i=1}^k \mathbb{R}$, we get holomorphic identifications $C_i^\circ = S^1 \times \mathbb{R}$. Since (C, ℓ, u) is S^1 -invariant, the function $u: C_i^\circ \rightarrow M$ must be independent of the S^1 -coordinate, and hence (examining the $\bar{\partial}$ -equation) is simply a Morse flow line of H . Running this argument in reverse, we also see that every stable Morse trajectory gives rise to a stable Floer trajectory which is S^1 -invariant. □

We will also need the following deeper fact (and henceforth we assume that H is defined as in Lemma 10.5.4):

Lemma 10.5.4 *Fix a Morse function $H_0: M \rightarrow \mathbb{R}$ whose gradient flow is Morse–Smale and suppose $H = \epsilon \cdot H_0$ with $\epsilon > 0$ sufficiently small. Then $\overline{\mathcal{M}}(\sigma, p, q)^{S^1} \subseteq \overline{\mathcal{M}}(\sigma, p, q)$ is open and cut out transversally (meaning $\overline{\mathcal{M}}(\sigma, p, q)^{S^1} \subseteq \overline{\mathcal{M}}(\sigma, p, q)^{\text{reg}}$). Hence we have a partition into disjoint closed subsets*

$$(10.5.1) \quad \overline{\mathcal{M}}(\sigma, p, q) = \overline{\mathcal{M}}(\sigma, p, q)^{S^1} \sqcup [\overline{\mathcal{M}}(\sigma, p, q) \setminus \overline{\mathcal{M}}(\sigma, p, q)^{S^1}]$$

Proof Salamon and Zehnder [73, Theorem 7.3(1)] show that $\overline{\mathcal{M}}(\sigma^0, p, q)^{S^1} \subseteq \overline{\mathcal{M}}(\sigma^0, p, q)^{\text{reg}}$ for the 0–simplex σ^0 and sufficiently small $\epsilon > 0$. Since we only consider constant families of J and H , the linearized operators for the higher σ^i can be written in terms of the linearized operators for σ^0 , and it follows that $\overline{\mathcal{M}}(\sigma, p, q)^{S^1} \subseteq \overline{\mathcal{M}}(\sigma, p, q)^{\text{reg}}$ for all σ .

Now we have $\overline{\mathcal{M}}_{\text{morse}}(\sigma, p, q) = \overline{\mathcal{M}}(\sigma, p, q)^{S^1} \subseteq \overline{\mathcal{M}}(\sigma, p, q)^{\text{reg}}$. It remains to show that this is an open inclusion. This can likely be seen by following closely the gluing argument used to prove Propositions 10.3.3 and 10.5.2, however we can argue directly as follows. Both $\overline{\mathcal{M}}_{\text{morse}}(\sigma, p, q)$ and $\overline{\mathcal{M}}(\sigma, p, q)^{\text{reg}}$ are topological manifolds with corners, and the stratification for the former is the pullback of the stratification for the latter. Their dimensions are dictated by the Morse index and the Conley–Zehnder index respectively, which in this case coincide. Hence we are done by Lemma 10.5.5 below. \square

Lemma 10.5.5 *Let M be a topological manifold with corners and $K \subseteq M$ a closed subset. Suppose that the restriction of the corner stratification on M induces a topological manifold with corners structure on K of the same dimension. Then $K \subseteq M$ is open.*

Proof The question is local on M , so we may assume that $M = \mathbb{R}^n \times \mathbb{R}_{\geq 0}^m$. Let $\tilde{K} \subseteq \mathbb{R}^n \times \mathbb{R}^m$ be obtained by reflecting K across the last m coordinate hyperplanes. The hypotheses then imply that \tilde{K} is a manifold of dimension $n + m$. Now we have that $\tilde{K} \subseteq \mathbb{R}^{n+m}$ is open by Brouwer’s “invariance of domain”. This is enough. \square

10.6 S^1 –equivariant implicit atlas

We constructed in Section 10.3 an implicit atlas \mathcal{A} on $\overline{\mathcal{M}}/JH_{\bullet}(M)$, and thus in particular on $\overline{\mathcal{M}}/\ast$. In this subsection, we modify this construction to define another implicit atlas \mathcal{B}^{S^1} on $\overline{\mathcal{M}}/\ast$, one which is S^1 –equivariant in the sense that the S^1 –action on $\overline{\mathcal{M}}(\sigma, p, q)$ extends canonically to all the thickenings $\overline{\mathcal{M}}(\sigma, p, q)_I^{\leq 5}$ in \mathcal{B}^{S^1} . The key technical step is to perform domain stabilization with S^1 –invariant divisors (Lemma 10.6.4).

Let us first motivate the definition of \mathcal{B}^{S^1} by describing a “first attempt” at defining an S^1 -equivariant implicit atlas on $\overline{\mathcal{M}}/*$. We consider the subatlas $\mathcal{A}^{S^1} \subseteq \mathcal{A}$ consisting of those thickening datums α for which D_α is S^1 -invariant. Now there is clearly a canonical S^1 -action on the thickenings $\overline{\mathcal{M}}(\sigma, p, q)_{\overline{I}}^{\leq s}$ (postcomposition of u with a rotation of S^1) for $I \subseteq \mathcal{A}^{S^1}(\sigma, p, q)^{\geq s}$.³⁷ Now $\mathcal{A}^{S^1} \subseteq \mathcal{A}$ forms an implicit atlas if and only if it satisfies the covering axiom. However, the covering axiom for \mathcal{A}^{S^1} fails: we cannot stabilize the domains of Morse flow lines (points of $\overline{\mathcal{M}}^{S^1}$) using S^1 -invariant divisors D_α (more generally, we cannot stabilize the domain of any broken trajectory containing a Morse flow line).

To fix this issue, we first modify the definition of \mathcal{A} to allow Morse components of Floer trajectories which do not get stabilized.

Definition 10.6.1 (Implicit atlas \mathcal{B} on $\overline{\mathcal{M}}/*$) We define an implicit atlas \mathcal{B} on $\overline{\mathcal{M}}/*$ as follows. On the level of index sets, we define $\mathcal{B} := \mathcal{A}$. However, we modify the definition of an I -thickened Floer trajectory as follows. We require that when C_α is considered with the r_α extra marked points $(u|C_\alpha)^{-1}(D_\alpha)$, the only unstable components are mapped by u to Morse flow lines; let $C_\alpha \rightarrow C_\alpha^{\text{st}}$ be the map contracting all such unstable components. Now instead of $\phi_\alpha: C_\alpha \rightarrow \overline{\mathcal{C}}_{0,2+r_\alpha}$, we use $\phi_\alpha: C_\alpha^{\text{st}} \rightarrow \overline{\mathcal{C}}_{0,2+r_\alpha}$.

The rest of the atlas data is defined analogously with that of \mathcal{A} without any serious difference. Note, however, that to verify that the locus $U_{IJ} \subseteq X_I$ is open, we must appeal to Lemma 10.5.4.

Note that the thickened moduli spaces for \mathcal{A} are open subsets of those for \mathcal{B} , so the covering axiom for \mathcal{A} implies the covering axiom for \mathcal{B} . Now to verify that \mathcal{B} is an implicit atlas, everything is the same as for \mathcal{A} except for the openness and submersion axioms, which follow from the following result (identical to Proposition 10.3.3), whose proof is given in Appendix C.

Proposition 10.6.2 (Formal regularity implies topological regularity) For all $I \subseteq J \subseteq \mathcal{B}(\sigma, p, q)^{\geq s}$, we have:

- (i) $(\overline{\mathcal{M}}(\sigma, p, q)_{\overline{I}}^{\leq s})^{\text{reg}} \subseteq \overline{\mathcal{M}}(\sigma, p, q)_{\overline{I}}^{\leq s}$ is an open subset.
- (ii) The map $s_{J \setminus I}: \overline{\mathcal{M}}(\sigma, p, q)_{\overline{J}}^{\leq s} \rightarrow E_{J \setminus I}$ over the locus

$$\psi_{IJ}^{-1}((\overline{\mathcal{M}}(\sigma, p, q)_{\overline{I}}^{\leq s})^{\text{reg}}) \subseteq \overline{\mathcal{M}}(\sigma, p, q)_{\overline{J}}^{\leq s}$$

³⁷Note that we do not need to put any restrictions on λ_α .

is locally modeled on the projection

$$(10.6.1) \quad \mathbb{R}^{\text{vdim } \overline{\mathcal{M}}(\sigma, p, q) \leq s + \dim E_I} \times \mathbb{R}^{\dim E_{J \setminus I}} \rightarrow \mathbb{R}^{\dim E_{J \setminus I}}$$

over the top stratum $s \in \mathcal{S}_{\overline{\mathcal{M}}}(\sigma, p, q) \leq s$. More generally, the local model (compatible with stratifications) is given by replacing the first factor on the left by $\mathbb{R}_{\geq 0}^n \times \mathbb{R}^{n'}$ stratified appropriately by $\mathcal{S}_{\overline{\mathcal{M}}}(\sigma, p, q)$.

- (iii) There exist π -invariant coherent trivializations of the local systems $\mathfrak{o}_{E_I}^\vee \otimes \mathfrak{o}_{(\overline{\mathcal{M}}(\sigma, p, q)_I^{\leq s})^{\text{reg}}}$ (in the sense of Definition 7.2.6), agreeing by restriction with those for \mathcal{A} , and coinciding with the usual orientations from Morse theory on $\overline{\mathcal{M}}_{\text{morse}}(\sigma, p, q) = \overline{\mathcal{M}}(\sigma, p, q)^{S^1} \subseteq \overline{\mathcal{M}}(\sigma, p, q)^{\text{reg}}$.

Thus \mathcal{B} is an implicit atlas on $\overline{\mathcal{M}}/\ast$.

Definition 10.6.3 (S^1 -equivariant implicit atlas \mathcal{B}^{S^1} on $\overline{\mathcal{M}}/\ast$) Let $\mathcal{B}^{S^1} \subseteq \mathcal{B}$ consist of those thickening datums α for which D_α is S^1 -invariant. There is a canonical S^1 -action on the thickenings $\overline{\mathcal{M}}(\sigma, p, q)_I^{\leq s}$ (postcomposition of u with a rotation of S^1) for $I \subseteq \mathcal{B}^{S^1}(\sigma, p, q)^{\geq s}$, and this S^1 -action is compatible with the rest of the structure.

To verify that $\mathcal{B}^{S^1} \subseteq \mathcal{B}$ is an implicit atlas, we just need to verify the covering axiom. We follow the usual proof of the covering axiom as in Lemma 9.2.9 and use the fact that Morse components are already cut out transversally (Lemma 10.5.4). To complete the proof, we just need Lemma 10.6.4 below, which says that for any stable Floer trajectory, we can stabilize the domain using S^1 -invariant divisors (except, of course, for any irreducible components mapping to Morse flow lines).

Lemma 10.6.4 (S^1 -equivariant domain stabilization for stable Floer trajectories of $\overline{\mathcal{M}}/\ast$) Let $\ell: \coprod_{i=1}^k \mathbb{R} \rightarrow \Delta^n$ and $u: C \rightarrow M \times S^1 \times \coprod_{i=1}^k \mathbb{R}$ be a point in $\overline{\mathcal{M}}(\sigma, p, q)$. Then there exists $D \subseteq M \times S^1 \times \Delta^n$ as in Definition 10.3.1(i) which is S^1 -invariant with $C \pitchfork D$ in the sense of Definition 10.3.2(iv) and so that adding these intersections to C as extra marked points makes C stable, except for irreducible components $S^1 \times \mathbb{R} \subseteq C$ on which u is independent of the S^1 -coordinate (“Morse flow lines”).

Proof Instead of finding an S^1 -invariant $D \subseteq M \times S^1 \times \Delta^n$, we find a $D \subseteq M \times \Delta^n$ (which is clearly equivalent) and we ignore the S^1 factor in the codomain of u . As in Lemma 10.3.4, it suffices to show that for any unstable component $C_0 \subseteq C$, either C_0 is a Morse flow line or there exists a point (other than a node) where $d((\text{id}_M \times \ell) \circ u): TC_0 \rightarrow TM \times T\Delta^n$ is injective.

If the projection $C_0 \rightarrow S^1 \times \mathbb{R}$ is constant, then $u: C_0 \rightarrow M$ is a (nonconstant!) J -holomorphic sphere and we are done as in Lemma 10.3.4. Hence it suffices to treat the case when $u: C_0 \rightarrow S^1 \times \mathbb{R}$ is not constant. Thus $u: C_0 \rightarrow S^1 \times \mathbb{R}$ is an isomorphism, so we can use $(t, s) \in S^1 \times \mathbb{R}$ as coordinates on C_0 . Now we split into two cases.

First, suppose $\ell: \mathbb{R} \rightarrow \Delta^n$ is not constant. If $u: C_0 \rightarrow M$ is independent of the S^1 coordinate, then it is a Morse flow line, and we do not need to stabilize. Otherwise, there is a point where $\partial u / \partial t \neq 0$, and since $\ell: \mathbb{R} \rightarrow \Delta^n$ has nonvanishing derivative everywhere, it follows that du is injective at this point.

Second, suppose $\ell: \mathbb{R} \rightarrow \Delta^n$ is constant. Then our map $u: C_0 \rightarrow M$ satisfies

$$(10.6.2) \quad \frac{\partial u}{\partial t} + J \circ \frac{\partial u}{\partial s} = \nabla H.$$

Certainly $u: C_0 \rightarrow M$ is not constant; otherwise it would be unstable (infinite automorphism group). Thus du is nonzero somewhere. If du has rank two somewhere, then we are done. Thus let us suppose that this is not the case and show that $u|_{C_0}$ is independent of the t coordinate (and thus is a Morse flow line). Thus there exists some open set $U \subseteq C_0 = S^1 \times \mathbb{R}$ where du has rank 1. Inside U , we have that $\ker du \subseteq TC_0 = T(S^1 \times \mathbb{R})$ is an (integrable!) 1-dimensional distribution, so U is equipped with a 1-dimensional foliation and u is constant on the leaves. Thus we have (locally) a factorization $u: S^1 \times \mathbb{R} \xrightarrow{r} (-\epsilon, \epsilon) \xrightarrow{w} M$, the leaves of the foliation being given by $r^{-1}(\delta)$ for $\delta \in (-\epsilon, \epsilon)$. Now (10.6.2) becomes

$$(10.6.3) \quad \left(\frac{\partial r}{\partial t} + \frac{\partial r}{\partial s} \cdot J \right) \cdot w'(r(s, t)) = (\nabla H)(w(r(s, t))).$$

Since du has rank 1, we know that $w'(r(s, t)) \neq 0$. Hence the value of $r(s, t)$ determines the value of its derivative uniquely, ie dr is constant along the leaves $r^{-1}(\delta)$ of the foliation. It follows that the foliation is (locally) linear(!) and that we can follow any leaf infinitely in both directions and it never exits U (since $dr = 0$ outside U). Now if any leaf had nonzero slope, it would force $u: C_0 \rightarrow M$ to be constant, a contradiction. Thus all leaves have slope zero; in other words u is independent of the S^1 -coordinate over $U = S^1 \times U'$. But now we see that $U' = \mathbb{R}$, since if U' had boundary, it would imply that we have a Morse flow line reaching a critical point in finite time. Thus u is (globally) a Morse flow line. □

10.7 Calculation of Hamiltonian Floer homology and the Arnold conjecture

Arnold conjectured that the minimal number of fixed points of a non-degenerate Hamiltonian symplectomorphism $M \rightarrow M$ enjoys a lower bound similar to the minimal

number of critical points of a Morse function on M (known as the Morse number of M). It remains an open problem to obtain a sharp bound on the minimal number of symplectic fixed points, though much progress has been made.

Arnold’s conjecture was proved for surfaces by Eliashberg [21] and for tori by Conley and Zehnder [15]. The existence of at least one fixed point was shown by Gromov [36, page 331, 2.3.B’₄] under the assumption $\omega|_{\pi_2(M)} = 0$.

A major breakthrough was made by Floer [24], who introduced Hamiltonian Floer homology and showed (under some assumptions) that it is isomorphic to singular homology. Floer’s work provides a lower bound on the number of symplectic fixed points of the type predicted by Arnold. Indeed, if Hamiltonian Floer homology can be defined and shown to be isomorphic to singular homology, then we get a lower bound towards the Arnold conjecture of the form

$$(10.7.1) \quad \min\{\text{rk } D_\bullet \mid (D_\bullet, d) \text{ a free differential graded } \Lambda\text{-module} \\ \text{homotopy equivalent to } C_\bullet(M; \Lambda)\},$$

where Λ is the Novikov ring

$$(10.7.2) \quad \Lambda := \mathbb{Z}[\text{im}(\pi_2(M) \xrightarrow{\omega \oplus 2c_1(M)} \mathbb{R} \oplus \mathbb{Z})]^\wedge$$

completed with respect to ω and graded by $2c_1$ (note that if the grading on Λ is nontrivial, a “differential graded Λ -module” is not the same as a “complex of Λ -modules”). One can also adjoin $\pi_1(M)$ to the coefficient ring (as in Fukaya [28] or Abouzaid [2]) to obtain a sharper lower bound in (10.7.1), and furthermore the methods of Sullivan [77] allow one (at least in many cases) to replace “homotopy equivalent” in (10.7.1) with “simple homotopy equivalent”. Note that for $\Lambda = \mathbb{Z}[\pi_1(M)]$ and “simple homotopy equivalent” in place of “homotopy equivalent”, the lower bound (10.7.1) is precisely the stable Morse number of M (see Damian [16, page 424, Corollary 2.6]).

Floer’s original work [24] covered the case of monotone symplectic manifolds (ie $\omega = \lambda c_1$ on $\pi_2(M)$ for some $\lambda > 0$), and the work of Hofer and Salamon [40] and Ono [66] extended this to semipositive symplectic manifolds (ie there do not exist classes $A \in \pi_2(M)$ with $\omega(A) > 0$ and $3 - n \leq c_1(A) < 0$). The case of general symplectic manifolds is due to Liu and Tian [54], Fukaya and Ono [33] and Ruan [71], using virtual techniques (which require rational coefficients) to resolve lack of transversality. We reprove their results below using the VFC machinery developed in this paper.

In the following result, we use the definition of Floer-type homology groups from Section 8.7.4.

Theorem 10.7.1 $F\mathbb{H}^\bullet(M)$ is isomorphic to $H^\bullet(M; \mathbb{Z}) \otimes \mathbb{Q}[[\pi]]$ as modules over $\mathbb{Q}[[\pi]]$.

Proof We use the setup of Section 10.5–10.6.

The homology groups associated to the flow category diagram $\overline{\mathcal{M}}/*$ and the implicit atlas \mathcal{A} are by definition $F\mathbb{H}^\bullet(M)$. Now, as we observed previously, the thickened moduli spaces of \mathcal{A} are open subsets of those of \mathcal{B} , so by Lemma 7.8.2, $F\mathbb{H}^\bullet(M)$ may also be defined using the implicit atlas \mathcal{B} on $\overline{\mathcal{M}}/*$. Now $\mathcal{B}^{S^1} \subseteq \mathcal{B}$ is a subatlas, so by Lemma 7.8.1 it may also be used to define $F\mathbb{H}^\bullet(M)$.

Thus let us restrict attention to the atlas \mathcal{B}^{S^1} on $\overline{\mathcal{M}}/*$. Recall that by Lemma 10.5.4, there is a partition into closed subsets

$$(10.7.3) \quad \overline{\mathcal{M}}(\sigma, p, q) = \overline{\mathcal{M}}(\sigma, p, q)^{S^1} \sqcup [\overline{\mathcal{M}}(\sigma, p, q) \setminus \overline{\mathcal{M}}(\sigma, p, q)^{S^1}].$$

Now we apply S^1 -localization to \mathcal{B}^{S^1} on $\overline{\mathcal{M}}/*$ in the form of Theorem 8.7.13, which applies since S^1 acts with finite stabilizers on $\overline{\mathcal{M}}(\sigma, p, q) \setminus \overline{\mathcal{M}}(\sigma, p, q)^{S^1}$ and our coefficient group is \mathbb{Q} . It follows that $F\mathbb{H}^\bullet(M)$ may be defined using the flow category diagram $\overline{\mathcal{M}}^{S^1}/*$ with the implicit atlas obtained from \mathcal{B}^{S^1} by removing $\overline{\mathcal{M}} \setminus \overline{\mathcal{M}}^{S^1}$ from every thickening.

All the flow spaces of $\overline{\mathcal{M}}^{S^1}/*$ are cut out transversally by Lemma 10.5.4, so by Proposition 7.8.4, the homology groups of $\overline{\mathcal{M}}^{S^1}/*$ can be defined by simply counting the 0-dimensional flow spaces according to the orientations ω . Now $\overline{\mathcal{M}}^{S^1}$ coincides with the Morse flow category diagram $\overline{\mathcal{M}}_{\text{morse}}/*$ of H by Lemma 10.5.3 (with the same orientations by Proposition 10.6.2), and this gives the desired isomorphism. \square

Remark 10.7.2 One expects to be able to show that the isomorphism in Theorem 10.7.1 is canonical by considering continuation maps and chain homotopies associated to families of Morse functions.

Corollary 10.7.3 (Arnold conjecture) *Let $H: M \times S^1 \rightarrow \mathbb{R}$ be a smooth function whose time-1 Hamiltonian flow $\phi_H: M \rightarrow M$ has non-degenerate fixed points. Then $\#\text{Fix } \phi_H \geq \dim H_\bullet(M; \mathbb{Q})$ (in fact, we may replace $\text{Fix } \phi_H$ with those fixed points whose associated periodic orbit is null-homotopic in M).*

Proof Pick any ω -compatible almost complex structure J and consider the vertex $(J, H) \in JH_\bullet(M)$. Over this vertex, pick any complex $FC^\bullet(M)$ which calculates $F\mathbb{H}^\bullet(M)$. Then $FC^\bullet(M)$ is a free $\mathbb{Q}[[\pi]]$ -module of rank

$$(10.7.4) \quad \text{rk}_{\mathbb{Q}[[\pi]]} FC^\bullet(M) \leq \#\text{Fix } \phi_H$$

(its rank equals the number of null-homotopic periodic orbits). By Theorem 10.7.1, the homology $F\mathbb{H}^\bullet(M)$ is free of rank

$$(10.7.5) \quad \text{rk}_{\mathbb{Q}[[\pi]]} F\mathbb{H}^\bullet(M) = \dim H_\bullet(M; \mathbb{Q}).$$

Now by definition, there is a $\mathbb{Q}[[\pi]]$ -linear boundary map $d: FC^\bullet(M) \rightarrow FC^{\bullet+1}(M)$ and by definition $\text{rk}_{\mathbb{Q}[[\pi]]} F\mathbb{H}^\bullet(M) = \ker d / \text{im } d$. Now apply Lemma 10.8.1 to conclude that

$$(10.7.6) \quad \text{rk}_{\mathbb{Q}[[\pi]]} F\mathbb{H}^\bullet(M) \leq \text{rk}_{\mathbb{Q}[[\pi]]} FC^\bullet(M).$$

Thus we are done. □

10.8 A little commutative algebra

Lemma 10.8.1 *Let M be a free module over a commutative ring R , and let $d: M \rightarrow M$ satisfy $d^2 = 0$. If $H = \ker d / \text{im } d$ is free, then $\text{rk } H \leq \text{rk } M$.*

Proof Since H is free, it is projective, so the surjection $\ker d \twoheadrightarrow H$ has a section $H \hookrightarrow \ker d \subseteq M$. Hence there is an injection $H \hookrightarrow M$. Now use Lemma 10.8.2. □

Lemma 10.8.2 *Let $\phi: R^{\oplus A} \hookrightarrow R^{\oplus B}$ be an inclusion of free modules over a commutative ring R . Then $A \leq B$.*

Proof This is a standard yet tricky exercise. We recall one of the (many) standard proofs.

The map ϕ is described by some matrix of $A \times B$ elements of R . Certainly the kernel of this matrix remains zero over the subring of R generated by its entries. Thus we may assume without loss of generality that R is a finitely generated \mathbb{Z} -algebra. Now localize at a prime ideal $\mathfrak{p} \subset R$ of height zero. Localization is exact, so again ϕ remains injective. Thus we may assume without loss of generality that R is a local Noetherian ring of dimension zero, and thus R is an Artin local ring [5, Theorem 8.5]. Since R is Artinian, all finitely generated modules have finite length [5, Propositions 6.5 and 6.8], and hence there is a length homomorphism $K_0(R) \rightarrow \mathbb{Z}$ [5, Proposition 6.9] (which is clearly an isomorphism since R is local). It thus follows from the short exact sequence

$$0 \rightarrow R^{\oplus A} \xrightarrow{\phi} R^{\oplus B} \rightarrow \text{coker } \phi \rightarrow 0$$

that $A \leq B$. □

Appendix A: Homological algebra

In this appendix, we collect some useful facts concerning sheaves, homotopy sheaves and their Čech cohomology. We assume the reader is familiar with most elementary aspects of sheaves. Many of the results in this appendix are also elementary, though for completeness we give most proofs as we do not know of a good reference.

In Section A.1, we recall presheaves and sheaves. In Section A.2, we introduce homotopy sheaves. In Section A.3, we list standard pushforward and pullback operations on (homotopy) sheaves. In Section A.4, we introduce and prove basic properties of Čech cohomology. In Section A.5, we introduce the central notion of a pure homotopy sheaf. In Section A.6, we prove a version of Poincaré–Lefschetz duality using pure homotopy sheaves. In Section A.7, we introduce a certain relevant type of homotopy colimit. In Section A.8, we prove an easy lemma about homotopy colimits of pure homotopy sheaves. In Section A.9, we review the definition of Steenrod homology.

Convention A.0.1 In this appendix, by *space* we mean locally compact Hausdorff space.

Convention A.0.2 By a *complex* C^\bullet we mean a \mathbb{Z} -graded object $\bigoplus_{i \in \mathbb{Z}} C^i$ (in some abelian category) along with a degree-1 endomorphism d with $d^2 = 0$. The *homology* of a complex C^\bullet is denoted $H^\bullet C^\bullet$, defined by

$$H^i C^\bullet := \ker(C^i \xrightarrow{d} C^{i+1}) / \operatorname{im}(C^{i-1} \xrightarrow{d} C^i).$$

A complex is called *acyclic* if and only if its homology vanishes. A map of complexes $f: A^\bullet \rightarrow B^\bullet$ is called a *quasi-isomorphism* if and only if it induces an isomorphism on homology. We will often use the fact that a map of complexes is a quasi-isomorphism if and only if its mapping cone is acyclic.

The *shift* of a complex $C^\bullet[n]$ is defined by $(C^\bullet[n])^i := C^{i+n}$. We use the *truncation functors* defined by

$$(\tau_{\geq i} C^\bullet)^j := \begin{cases} C^j, & j > i, \\ \operatorname{coker}(C^{i-1} \xrightarrow{d} C^i), & j = i, \\ 0, & j < i, \end{cases}$$

$$(\tau_{\leq i} C^\bullet)^j := \begin{cases} 0, & j > i, \\ \ker(C^i \xrightarrow{d} C^{i+1}), & j = i, \\ C^j, & j < i. \end{cases}$$

Given a sequence of maps of complexes $A_0^\bullet \rightarrow \dots \rightarrow A_n^\bullet$ such that adjacent maps compose to zero, we denote by $[A_0^\bullet \rightarrow \dots \rightarrow A_n^{\bullet-n}]$ the associated total complex of this double complex. For example, $f: A^\bullet \rightarrow B^\bullet$ denotes a map, and $[A^\bullet \xrightarrow{f} B^{\bullet-1}]$ denotes its mapping cone.

Convention A.0.3 We fix sign conventions by making everything $\mathbb{Z}/2$ -graded and always using the super tensor product \otimes (so the isomorphism $A \otimes B \xrightarrow{\sim} B \otimes A$ is given by $a \otimes b \mapsto (-1)^{|a||b|} b \otimes a$, and $(f \otimes g)(a \otimes b) := (-1)^{|g||a|} f(a) \otimes g(b)$).

We fix $\text{Hom}(A, B) \otimes A \rightarrow B$ as given by $f \otimes a \mapsto f(a)$. Complexes are $(\mathbb{Z}, \mathbb{Z}/2)$ -bigraded; differentials are always odd and chain maps are always even. Note that the $\mathbb{Z}/2$ -grading of a complex is often, but not always, the reduction modulo 2 of the \mathbb{Z} -grading.

Convention A.0.4 Direct and inverse limits always take place over directed sets.

A.1 Presheaves and sheaves

Definition A.1.1 (Presheaf and \mathcal{K} -presheaf³⁸) Let X be a space. A *presheaf* (resp. \mathcal{K} -*presheaf*) on X is a contravariant functor from the category of open (resp. compact) sets of X to the category of abelian groups. A morphism of presheaves is simply a natural transformation of functors. The categories of presheaves and \mathcal{K} -presheaves are denoted $\text{Prshv } X$ and $\text{Prshv}_{\mathcal{K}} X$, respectively.

Definition A.1.2 (Stalk) For a presheaf \mathcal{F} , let $\mathcal{F}_p := \varinjlim_{p \in U} \mathcal{F}(U)$, and for a \mathcal{K} -presheaf \mathcal{F} let $\mathcal{F}_p := \mathcal{F}(\{p\})$. In both cases we say \mathcal{F}_p is the *stalk* of \mathcal{F} at p .

Definition A.1.3 (Sheaf) A *sheaf* is a presheaf \mathcal{F} satisfying the following condition:

$$\text{(Sh)} \quad 0 \rightarrow \mathcal{F}\left(\bigcup_{\alpha \in A} U_{\alpha}\right) \rightarrow \prod_{\alpha \in A} \mathcal{F}(U_{\alpha}) \rightarrow \prod_{\alpha, \beta \in A} \mathcal{F}(U_{\alpha} \cap U_{\beta}) \quad \text{is exact } \forall \{U_{\alpha} \subseteq X\}_{\alpha \in A}.$$

The category of sheaves on a space X is denoted $\text{Shv } X$ (a full subcategory of $\text{Prshv } X$, meaning a morphism between sheaves is the same as a morphism of the corresponding presheaves).

Definition A.1.4 (\mathcal{K} -sheaf) A \mathcal{K} -*sheaf* is a \mathcal{K} -presheaf \mathcal{F} satisfying the following three conditions:³⁹

(Sh $_{\mathcal{K}}$ 1) $\mathcal{F}(\emptyset) = 0$.

(Sh $_{\mathcal{K}}$ 2) The sequence

$$0 \rightarrow \mathcal{F}(K_1 \cup K_2) \rightarrow \mathcal{F}(K_1) \oplus \mathcal{F}(K_2) \rightarrow \mathcal{F}(K_1 \cap K_2)$$

is exact for all $K_1, K_2 \subseteq X$.

(Sh $_{\mathcal{K}}$ 3) The map

$$\varinjlim_{\substack{K \subseteq U \\ U \text{ open}}} \mathcal{F}(\bar{U}) \rightarrow \mathcal{F}(K)$$

is an isomorphism for all $K \subseteq X$.

³⁸Terminology “ \mathcal{K} -” following Lurie [57, Definition 7.3.4.1].

³⁹A similar set of axioms appears in Lurie [57, Definition 7.3.4.1].

The category of \mathcal{K} -sheaves on a space X is denoted $\text{Shv}_{\mathcal{K}} X$ (a full subcategory of $\text{Prshv}_{\mathcal{K}} X$).

Remark A.1.5 It is always the case that $K = \bigcap_{\substack{K \subseteq U \\ U \text{ open}}} \bar{U}$ for compact $K \subseteq X$.

Definition A.1.6 We define functors

$$(A.1.1) \quad \text{Prshv}_{\mathcal{K}} X \begin{matrix} \xrightarrow{\alpha_*} \\ \xleftarrow{\alpha^*} \end{matrix} \text{Prshv} X$$

by the formulas

$$(A.1.2) \quad (\alpha^* \mathcal{F})(K) := \varinjlim_{\substack{K \subseteq U \\ U \text{ open}}} \mathcal{F}(U),$$

$$(A.1.3) \quad (\alpha_* \mathcal{F})(U) := \varprojlim_{\substack{K \subseteq U \\ K \text{ compact}}} \mathcal{F}(K).$$

It is easy to see that there is an adjunction $\text{Hom}_{\text{Prshv}_{\mathcal{K}} X}(\alpha^* \mathcal{F}, \mathcal{G}) = \text{Hom}_{\text{Prshv} X}(\mathcal{F}, \alpha_* \mathcal{G})$ (giving an element of either Hom-set is the same as giving a compatible system of maps $\mathcal{F}(U) \rightarrow \mathcal{G}(K)$ for pairs $K \subseteq U$).

Lemma A.1.7 Let $K_1, \dots, K_n \subseteq X$ be compact. Then $\{U_1 \cap \dots \cap U_n\}_{\text{open } U_i \supseteq K_i}$ forms a cofinal system of neighborhoods of $K_1 \cap \dots \cap K_n$.

Proof We may assume without loss of generality that X is compact. By induction, it suffices to treat the case $n = 2$. The rest is an exercise (use the fact that a compact Hausdorff space is normal). □

Lemma A.1.8 We have

$$(A.1.4) \quad \text{Shv}_{\mathcal{K}} X \begin{matrix} \xrightarrow{\alpha_*} \\ \xleftarrow{\alpha^*} \end{matrix} \text{Shv} X,$$

and this is an equivalence of categories.⁴⁰

Proof Suppose \mathcal{F} is a sheaf, and let us verify $\alpha^* \mathcal{F}$ is a \mathcal{K} -sheaf. Axiom (Sh _{\mathcal{K}} 1) is clear (take $A = \emptyset$ in (Sh)). Axiom (Sh _{\mathcal{K}} 2) follows from (Sh) and Lemma A.1.7 since direct limits are exact. Axiom (Sh _{\mathcal{K}} 3) is clear from the definition also.

⁴⁰A similar result appears in Lurie [57, Corollary 7.3.4.10].

Suppose \mathcal{F} is a \mathcal{K} -sheaf, and let us verify that $\alpha_*\mathcal{F}$ is a sheaf. Let us first observe that (by induction using (Sh $_{\mathcal{K}}$ 1) and (Sh $_{\mathcal{K}}$ 2)) if $\{K_\alpha \subseteq X\}_{\alpha \in A}$ is any collection of compact sets, all but finitely many of which are empty, then the following is exact:

$$(A.1.5) \quad 0 \rightarrow \mathcal{F}\left(\bigcup_{\alpha \in A} K_\alpha\right) \rightarrow \prod_{\alpha \in A} \mathcal{F}(K_\alpha) \rightarrow \prod_{\alpha, \beta \in A} \mathcal{F}(K_\alpha \cap K_\beta).$$

Let us now verify axiom (Sh) for $\alpha_*\mathcal{F}$ for open sets $\{U_\alpha \subseteq X\}_{\alpha \in A}$. Certainly (Sh) is the inverse limit of (A.1.5) over all collections of compact subsets $\{K_\alpha \subseteq U_\alpha\}_{\alpha \in A}$ for which $K_\alpha = \emptyset$ except for finitely many $\alpha \in A$. This is sufficient since inverse limit is left exact.

Now to see that the adjoint pair $\alpha^* \dashv \alpha_*$ is actually an equivalence of categories, it suffices to show that the natural morphisms $\mathcal{F} \rightarrow \alpha_*\alpha^*\mathcal{F}$ and $\alpha^*\alpha_*\mathcal{G} \rightarrow \mathcal{G}$ are isomorphisms. In other words, we must show that the natural maps

$$(A.1.6) \quad \mathcal{F}(U) \rightarrow \varprojlim_{\substack{K \subseteq U \\ K \text{ compact}}} \varinjlim_{\substack{K \subseteq U' \\ U' \text{ open}}} \mathcal{F}(U'),$$

$$(A.1.7) \quad \varinjlim_{\substack{K \subseteq U \\ U \text{ open}}} \varprojlim_{\substack{K' \subseteq U \\ K' \text{ compact}}} \mathcal{G}(K') \rightarrow \mathcal{G}(K)$$

are isomorphisms. The right-hand side of (A.1.6) can be thought of as the inverse limit of $\mathcal{F}(U')$ over all open $U' \subseteq U$ with $\overline{U'} \subseteq U$ and $\overline{U'}$ compact. It is easy to see that the map from $\mathcal{F}(U)$ to this inverse limit is an isomorphism if \mathcal{F} is a sheaf. The left-hand side of (A.1.7) can be identified with the left-hand side of (Sh $_{\mathcal{K}}$ 3), and thus the map to $\mathcal{G}(K)$ is an isomorphism if \mathcal{F} is a \mathcal{K} -sheaf. □

Convention A.1.9 In view of the canonical equivalence of categories $\text{Shv } X = \text{Shv}_{\mathcal{K}} X$ from Lemma A.1.8, we use the single word “sheaf” for an object of either category.

Lemma A.1.10 Let $\{K_\alpha\}_{\alpha \in A}$ be a filtered directed system of compact subsets of X (ie for all $\alpha, \beta \in A$ there exists $\gamma \in A$ with $K_\gamma \subseteq K_\alpha \cap K_\beta$). Then for any \mathcal{F} satisfying (Sh $_{\mathcal{K}}$ 3), the following is an isomorphism:

$$(A.1.8) \quad \varinjlim_{\alpha \in A} \mathcal{F}(K_\alpha) \rightarrow \mathcal{F}\left(\bigcap_{\alpha \in A} K_\alpha\right).$$

Proof Write $K := \bigcap_{\alpha \in A} K_\alpha$. Now consider the diagram

$$(A.1.9) \quad \begin{array}{ccccc} \varinjlim_{\alpha \in A} \varinjlim_{\substack{K_\alpha \subseteq U \\ U \text{ open}}} \mathcal{F}(\bar{U}) & \longrightarrow & \varinjlim_{\substack{\exists \alpha \in A: K_\alpha \subseteq U \\ U \text{ open}}} \mathcal{F}(\bar{U}) & \longrightarrow & \varinjlim_{\substack{K \subseteq U \\ U \text{ open}}} \mathcal{F}(\bar{U}) \\ \downarrow & & & & \downarrow \\ \varinjlim_{\alpha \in A} \mathcal{F}(K_\alpha) & \longrightarrow & & \longrightarrow & \mathcal{F}(K) \end{array}$$

The vertical maps are both isomorphisms by (Sh_{3C}). The first horizontal map is an isomorphism since $\{K_\alpha\}_{\alpha \in A}$ is filtered. It thus remains to show that if $K \subseteq U$ and U is open, then there exists $\alpha \in A$ such that $K_\alpha \subseteq U$. Since $K \subseteq U$, we have $(X \setminus U) \cap K = \emptyset$, so $\bigcap_{\alpha \in A} (X \setminus U) \cap K_\alpha = \emptyset$. This is a filtered directed system of compact sets, and so the intersection being empty implies that one of the terms $(X \setminus U) \cap K_\alpha$ is empty, so $K_\alpha \subseteq U$ as desired. \square

A.2 Homotopy sheaves

Convention A.2.1 Let \mathcal{F}^\bullet be a complex of $(\mathcal{K}\text{-})$ presheaves. The notions from Convention A.0.2 are applied “objectwise”, that is, to each complex $\mathcal{F}^\bullet(U)$ (resp. $\mathcal{F}^\bullet(K)$) individually. For example, the homology $H^i \mathcal{F}^\bullet$ is again a $(\mathcal{K}\text{-})$ presheaf, and a map of $(\mathcal{K}\text{-})$ presheaves $f: \mathcal{F}^\bullet \rightarrow \mathcal{G}^\bullet$ is a quasi-isomorphism if and only if it induces an isomorphism of $(\mathcal{K}\text{-})$ presheaves $H^i \mathcal{F}^\bullet \rightarrow H^i \mathcal{G}^\bullet$ for all i . The homology $(\mathcal{K}\text{-})$ presheaves $H^i \mathcal{F}^\bullet$ should not be confused with the various flavors of Čech (hyper)cohomology $\check{H}^i(X; \mathcal{F}^\bullet)$ we introduce in Section A.4.

Definition A.2.2 (Homotopy sheaf) A *homotopy sheaf* is a complex of presheaves \mathcal{F}^\bullet satisfying the condition

$$(hSh) \quad \left[\mathcal{F}^\bullet \left(\bigcup_{\alpha \in A} U_\alpha \right) \rightarrow \prod_{\alpha \in A} \mathcal{F}^{\bullet-1}(U_\alpha) \rightarrow \prod_{\alpha, \beta \in A} \mathcal{F}^{\bullet-2}(U_\alpha \cap U_\beta) \rightarrow \dots \right]$$

is acyclic for all $\{U_\alpha \subseteq X\}_{\alpha \in A}$.

The category of homotopy sheaves on a space X is denoted $hShv X$ (morphisms are morphisms of complexes of presheaves).

Remark A.2.3 The definition above is given for illustrative purposes only. It is probably only a good definition for \mathcal{F}^\bullet which is bounded below, and it would perhaps be better to impose (hSh) for all *hypercovers* of open subsets $U \subseteq X$.

Example A.2.4 A sheaf \mathcal{F} is *flasque* if and only if all restriction maps $\mathcal{F}(U) \rightarrow \mathcal{F}(U')$ are surjective. It is easy to see that any bounded-below complex of flasque sheaves \mathcal{F}^\bullet is a homotopy sheaf. For example, let $\mathcal{F}^\bullet(U) := C^\bullet(U)$ be the presheaf of singular cochain complexes. Then the sheafification of \mathcal{F}^\bullet is a complex of flasque sheaves, and thus is a homotopy sheaf. One can show (using barycentric subdivision) that the map from \mathcal{F}^\bullet to its sheafification is a quasi-isomorphism, and hence \mathcal{F}^\bullet is a homotopy sheaf as well.

Definition A.2.5 (Homotopy \mathcal{K} -sheaf) A *homotopy \mathcal{K} -sheaf* is a complex of \mathcal{K} -presheaves \mathcal{F}^\bullet satisfying the following three conditions:

(hSh $_{\mathcal{K}}$ 1) $\mathcal{F}^\bullet(\emptyset)$ is acyclic.

(hSh $_{\mathcal{K}}$ 2) The complex

$$[\mathcal{F}^\bullet(K_1 \cup K_2) \rightarrow \mathcal{F}^{\bullet-1}(K_1) \oplus \mathcal{F}^{\bullet-1}(K_2) \rightarrow \mathcal{F}^{\bullet-2}(K_1 \cap K_2)]$$

is acyclic for all $K_1, K_2 \subseteq X$.

(hSh $_{\mathcal{K}}$ 3) The map

$$\varinjlim_{\substack{K \subseteq U \\ U \text{ open}}} \mathcal{F}^\bullet(\bar{U}) \rightarrow \mathcal{F}^\bullet(K)$$

is a quasi-isomorphism for all $K \subseteq X$.

The category of homotopy \mathcal{K} -sheaves on a space X is denoted $\text{hShv}_{\mathcal{K}} X$ (morphisms are morphisms of complexes of \mathcal{K} -presheaves).

Remark A.2.6 Note that (hSh $_{\mathcal{K}}$ 2) gives rise to a ‘‘Mayer–Vietoris’’ long exact sequence in cohomology.

Example A.2.7 A \mathcal{K} -sheaf \mathcal{F} is *soft* if and only if all restriction maps $\mathcal{F}(K) \rightarrow \mathcal{F}(K')$ are surjective. It is easy to see that any complex of soft \mathcal{K} -sheaves \mathcal{F}^\bullet is a homotopy \mathcal{K} -sheaf.

Remark A.2.8 In analogy with Definition A.1.6, we expect there are functors

$$(A.2.1) \quad \text{hShv}_{\mathcal{K}} X \begin{matrix} \xrightarrow{R\alpha_*} \\ \xleftarrow{\alpha_*} \end{matrix} \text{hShv} X,$$

where α^* is the direct limit (A.1.2) and $R\alpha_*$ is the *homotopy* (or *derived*) version $R\varprojlim$ of the inverse limit (A.1.3) (the naive inverse limit functor α_* is the ‘‘wrong’’ functor since inverse limit is not exact).

In analogy with Lemma A.1.8, it seems likely⁴¹ that there is an adjunction $\alpha^* \dashv R\alpha_*$ which is an equivalence (in the sense of model categories or ∞ -categories), though perhaps only after restricting to the (possibly better-behaved) subcategory of homotopy (\mathcal{K} -)sheaves which are bounded below.

Remark A.2.9 Guided by the needs of the rest of the paper, we proceed to focus on homotopy \mathcal{K} -sheaves rather than on homotopy sheaves.

Lemma A.2.10 *Properties (hSh $_{\mathcal{K}1}$)–(hSh $_{\mathcal{K}3}$) are preserved by quasi-isomorphisms.*

Proof For (hSh $_{\mathcal{K}1}$) and (hSh $_{\mathcal{K}3}$) this is trivial. For (hSh $_{\mathcal{K}2}$), suppose $\mathcal{F}^\bullet \xrightarrow{\simeq} \mathcal{G}^\bullet$ is a quasi-isomorphism. Now the sequence (hSh $_{\mathcal{K}2}$) applied to the mapping cone $[\mathcal{F}^\bullet \rightarrow \mathcal{G}^{\bullet-1}]$ is certainly acyclic (since it has a finite filtration whose associated graded is acyclic). Thus (hSh $_{\mathcal{K}2}$) holds for \mathcal{F}^\bullet if and only if it holds for \mathcal{G}^\bullet . \square

Lemma A.2.11 (Extensions of homotopy \mathcal{K} -sheaves are homotopy \mathcal{K} -sheaves) *Let \mathcal{F}^\bullet be a complex of \mathcal{K} -presheaves. If \mathcal{F}^\bullet has a finite filtration whose associated graded is a homotopy \mathcal{K} -sheaf, then \mathcal{F}^\bullet is a homotopy \mathcal{K} -sheaf.*

Proof This follows from the fact that if a complex C^\bullet has a finite filtration whose associated graded is acyclic then C^\bullet is itself acyclic. \square

Lemma A.2.12 (Lowest nonzero homology \mathcal{K} -presheaf of a homotopy \mathcal{K} -sheaf is a \mathcal{K} -sheaf) *If \mathcal{F}^\bullet is a homotopy \mathcal{K} -sheaf and $H^{-1}\mathcal{F}^\bullet = 0$, then $H^0\mathcal{F}^\bullet$ is a \mathcal{K} -sheaf.*

Proof Properties (Sh $_{\mathcal{K}1}$) and (Sh $_{\mathcal{K}3}$) follow directly from (hSh $_{\mathcal{K}1}$) and (hSh $_{\mathcal{K}3}$). To show (Sh $_{\mathcal{K}2}$) for $H^0\mathcal{F}^\bullet$, use the long exact sequence induced by (hSh $_{\mathcal{K}2}$) and the vanishing of $H^{-1}\mathcal{F}^\bullet(K_1 \cap K_2)$. \square

Lemma A.2.13 *Let $\{K_\alpha\}_{\alpha \in A}$ be a filtered directed system of compact subsets of X (ie for all $\alpha, \beta \in A$ there exists $\gamma \in A$ with $K_\gamma \subseteq K_\alpha \cap K_\beta$). Then for any \mathcal{F}^\bullet satisfying (hSh $_{\mathcal{K}3}$), the following is a quasi-isomorphism:*

$$(A.2.2) \quad \varinjlim_{\alpha \in A} \mathcal{F}^\bullet(K_\alpha) \rightarrow \mathcal{F}^\bullet\left(\bigcap_{\alpha \in A} K_\alpha\right).$$

Proof Same as for Lemma A.1.10. \square

⁴¹A similar result appears in Lurie [57, Corollary 7.3.4.10].

A.3 Pushforward, exceptional pushforward, and pullback

Definition A.3.1 (Pushforward of $(\mathcal{K}-)$ presheaves) Let $f: X \rightarrow Y$ be a map of spaces. We define functors:

- (i) $f_*: \text{Prshv } X \rightarrow \text{Prshv } Y$ by $(f_*\mathcal{F})(U) := \mathcal{F}(f^{-1}(U))$.
- (ii) $f_*: \text{Prshv}_{\mathcal{K}} X \rightarrow \text{Prshv}_{\mathcal{K}} Y$ by $(f_*\mathcal{F})(K) := \mathcal{F}(f^{-1}(K))$ (if f is proper).

(The action of f_* on morphism spaces is obvious.)

Lemma A.3.2 Let $f: X \rightarrow Y$ be proper. Then $\{f^{-1}(U)\}_{K \subseteq U}$ is a cofinal system of neighborhoods of $f^{-1}(K)$ for any compact $K \subseteq X$.

Proof Exercise (use the fact that a compact Hausdorff space is normal). □

Lemma A.3.3 f_* preserves $(\text{Sh}), (\text{hSh}), (\text{Sh}_{\mathcal{K}1})-(\text{Sh}_{\mathcal{K}3}),$ and $(\text{hSh}_{\mathcal{K}1})-(\text{hSh}_{\mathcal{K}3})$.

Proof These are trivial except for $(\text{Sh}_{\mathcal{K}3})$ and $(\text{hSh}_{\mathcal{K}3})$, which use Lemma A.3.2. □

Definition A.3.4 (Pushforward of $(\mathcal{K}-)$ sheaves and homotopy \mathcal{K} -sheaves) Let $f: X \rightarrow Y$ be a map of spaces. By Lemma A.3.3, Definition A.3.1 gives rise to functors:

- (i) $f_*: \text{Shv } X \rightarrow \text{Shv } Y$.
- (ii) $f_*: \text{hShv } X \rightarrow \text{hShv } Y$.
- (iii) $f_*: \text{Shv}_{\mathcal{K}} X \rightarrow \text{Shv}_{\mathcal{K}} Y$ (if f is proper).
- (iv) $f_*: \text{hShv}_{\mathcal{K}} X \rightarrow \text{hShv}_{\mathcal{K}} Y$ (if f is proper).

(The action on morphism spaces is induced from the f_* at the level of (complexes of) $(\mathcal{K}-)$ presheaves.)

It is easy to see that f_* commutes with the equivalence $\text{Shv } X = \text{Shv}_{\mathcal{K}} X$ (if f is proper).

Definition A.3.5 (Pullback and exceptional pushforward of sheaves) Let $f: X \rightarrow Y$ be a map of spaces. We define:

- (i) $f^*: \text{Shv } Y \rightarrow \text{Shv } X$, the standard sheaf pullback (namely, $f^*\mathcal{F}$ is the sheafification of the presheaf $U \mapsto \varinjlim_{f(U) \subseteq V} \mathcal{F}(V)$).
- (ii) $f_! : \text{Shv } X \rightarrow \text{Shv } Y$ by $(f_!\mathcal{F})(U) \subseteq (f_*\mathcal{F})(U)$ being the subspace of sections which vanish in a neighborhood of $Y \setminus X$ (if f is the inclusion of an open set).

Definition A.3.6 (Pullback of homotopy \mathcal{K} -sheaves) Let $f: X \rightarrow Y$ be an injective map of spaces. We define:

(i) $f^*: \text{hShv}_{\mathcal{K}} Y \rightarrow \text{hShv}_{\mathcal{K}} X$ by $(f^*\mathcal{F}^\bullet)(K) := \mathcal{F}^\bullet(f(K))$.

We check the properties: (hSh $_{\mathcal{K}}$ 1) is clear, and (hSh $_{\mathcal{K}}$ 2) follows since f is injective. For (hSh $_{\mathcal{K}}$ 3), use Remark A.1.5, the injectivity of f , and Lemma A.2.13.

A.4 Čech cohomology

We introduce various flavors of Čech (hyper)cohomology relevant to our situation.

Remark A.4.1 A refinement of a cover $\{U_\alpha\}_{\alpha \in A}$ is a cover $\{U_\beta\}_{\beta \in B}$ along with a map $f: B \rightarrow A$ such that $U_\beta \subseteq U_{f(\alpha)}$. Refinements are the morphisms used in the directed systems used to define (all flavors of) Čech cohomology. Different refinements $\{U_\alpha\}_{\alpha \in A} \rightarrow \{U_\beta\}_{\beta \in B}$ induce different maps on Čech complexes, but they all agree after passing to cohomology (more precisely, the “space” of such refinements is contractible or empty). In particular, it follows that the directed systems used in defining Čech cohomology are filtered.

Remark A.4.2 The empty covering is a final object in the category of coverings of \emptyset , so we always have $\check{H}^\bullet(\emptyset; -) = 0$.

A.4.1 ... of sheaves

Definition A.4.3 (\check{H}^\bullet and \check{H}_c^\bullet of sheaves) Let \mathcal{F} be a sheaf on a space X . We define the Čech cohomology

$$(A.4.1) \quad \check{H}^\bullet(X; \mathcal{F}) := \varinjlim_{\substack{X = \bigcup_{\alpha \in A} U_\alpha \\ \text{open cover}}} H^\bullet \left[\bigoplus_{p \geq 0} \prod_{\substack{S \subseteq A \\ |S|=p+1}} \mathcal{F} \left(\bigcap_{\alpha \in S} U_\alpha \right) [-p] \right]$$

with the standard Čech differential.⁴² For any compact $K \subseteq X$, define $\check{H}_K^\bullet(X; \mathcal{F})$ (Čech cohomology with supports in K) via (A.4.1) except replacing every instance of $\mathcal{F}(U)$ with $\ker[\mathcal{F}(U) \rightarrow \mathcal{F}(U \setminus K)]$. We let $\check{H}_c^\bullet(X; \mathcal{F}) := \varinjlim_{K \subseteq X} \check{H}_K^\bullet(X; \mathcal{F})$ (Čech cohomology with compact supports).

Lemma A.4.4 (\check{H}^\bullet on a compact space needs only finite open covers) *If X is compact, then the natural map $\check{H}_{\text{fin}}^\bullet(X; \mathcal{F}) \rightarrow \check{H}^\bullet(X; \mathcal{F})$ is an isomorphism, where the left-hand side is defined as in (A.4.1) except using only finite open covers.*

⁴²Technically speaking, so that the signs in the differential can be defined canonically, we should really tensor each term of the direct product with $(\mathbb{Z}o_1 \oplus \mathbb{Z}o_2)/(o_1 + o_2) \cong \mathbb{Z}$ where o_1, o_2 denote the two orientations of the p -simplex on vertex set S .

Proof Since X is compact, it follows that finite open covers are cofinal (every open cover has a finite refinement). □

Lemma A.4.5 ($\check{H}_c^\bullet = \check{H}^\bullet$ on compact space) *If X is compact, then the natural map $\check{H}_c^\bullet(X; \mathcal{F}) \rightarrow \check{H}^\bullet(X; \mathcal{F})$ is an isomorphism.*

Proof Trivial. □

Definition A.4.6 (Pullback on \check{H}^\bullet and \check{H}_c^\bullet) Let $f: X \rightarrow Y$ be a map of spaces. An open cover of Y pulls back to give an open cover of X , and this gives an identification of the Čech complex for the cover of Y with coefficients in $f_*\mathcal{F}$ with the Čech complex for the cover of X with coefficients in \mathcal{F} . Hence we get natural maps:

- (i) $f^*: \check{H}^\bullet(Y; f_*\mathcal{F}) \rightarrow \check{H}^\bullet(X; \mathcal{F})$ for $\mathcal{F} \in \text{Shv } X$.
- (ii) $f^*: \check{H}_c^\bullet(Y; f_*\mathcal{F}) \rightarrow \check{H}_c^\bullet(X; \mathcal{F})$ for $\mathcal{F} \in \text{Shv } X$ (if f is proper).

Lemma A.4.7 (\check{H}_c^\bullet commutes with $f_!$) *Let $f: X \hookrightarrow Y$ be the inclusion of an open subset. Then there is a natural isomorphism $f_!: \check{H}_c^\bullet(X; \mathcal{F}) \rightarrow \check{H}_c^\bullet(Y; f_!\mathcal{F})$.*

Proof For $K \subseteq X$, there are natural maps

$$(A.4.2) \quad \check{H}_K^\bullet(Y; f_!\mathcal{F}) \begin{matrix} \xleftarrow{f^*} \\ \xrightarrow{f_!} \end{matrix} \check{H}_K^\bullet(X; \mathcal{F})$$

(for f^* : pull back the open cover; for $f_!$: add $Y \setminus K$ to the open cover and extend by zero). It is easy to see that $f_!$ and f^* are inverses. The desired map $f_!: \check{H}_c^\bullet(X; \mathcal{F}) \rightarrow \check{H}_c^\bullet(Y; f_!\mathcal{F})$ is defined as the composition

$$(A.4.3) \quad \varinjlim_{K \subseteq X} \check{H}_K^\bullet(X; \mathcal{F}) \stackrel{(A.4.2)}{=} \varinjlim_{K \subseteq X} \check{H}_K^\bullet(Y; f_!\mathcal{F}) \rightarrow \varinjlim_{K \subseteq Y} \check{H}_K^\bullet(Y; f_!\mathcal{F}).$$

We must show that the second map is an isomorphism; to see this, it suffices to show that the following is an isomorphism for all $K \subseteq Y$:

$$(A.4.4) \quad \varinjlim_{K' \subseteq X \cap K} \check{H}_{K'}^\bullet(Y; f_!\mathcal{F}) \rightarrow \check{H}_K^\bullet(Y; f_!\mathcal{F}).$$

We claim that for any Čech cochain β for $f_!\mathcal{F}$ with supports in K subordinate to an open cover of Y , there is a refinement on which the restriction of β has supports in some $K' \subseteq X \cap K$. It follows from the claim (using a cofinality argument) that (A.4.4) is an isomorphism. To prove the claim, argue as follows.

First, choose a refinement for which only finitely many open sets U_1, \dots, U_n intersect K and for which the remaining open sets cover $Y \setminus K$. Pick open sets $V_i \subseteq U_i$ which cover K and for which \bar{V}_i is compact and $\bar{V}_i \subseteq U_i$ (this is always possible).

Now β is a finite collection $\{\beta_I \in (f_! \mathcal{F})(\bigcap_{i \in I} U_i)\}_{\emptyset \neq I \subseteq \{1, \dots, n\}}$. We have

$$(A.4.5) \quad \text{supp} \left(\beta_I \Big|_{\bigcap_{i \in I} V_i} \right) \subseteq \text{supp } \beta_I \cap \bigcap_{i \in I} \bar{V}_i.$$

Since $\text{supp } \beta_I \subseteq \bigcap_{i \in I} U_i$ is (relatively) closed and $\bigcap_{i \in I} \bar{V}_i \subseteq \bigcap_{i \in I} U_i$ is compact, it follows that the right-hand side is compact. Also, we have $\text{supp } \beta_I \subseteq X$ (by definition of $f_! \mathcal{F}$). Hence the right-hand side of (A.4.5) is a compact subset of X . It follows that the restriction of β to the refinement obtained by replacing U_i with V_i for $i = 1, \dots, n$ has support in a compact subset of X . \square

Lemma A.4.8 (f^* is an isomorphism if f has finite fibers) *Let $f: X \rightarrow Y$ be proper with finite fibers. Then $f^*: \check{H}^\bullet(Y; f_* \mathcal{F}) \rightarrow \check{H}^\bullet(X; \mathcal{F})$ is an isomorphism, as is $f^*: \check{H}_c^\bullet(Y; f_* \mathcal{F}) \rightarrow \check{H}_c^\bullet(X; \mathcal{F})$.*

Proof Given an open cover containing $U = U_1 \sqcup \dots \sqcup U_n$ (finite disjoint union), we get a refinement by replacing U with $\{U_1, \dots, U_n\}$. A *partition* of an open cover is a refinement obtained by doing such a replacement on some (possibly infinitely many) open sets of the cover. Note that a partition induces an isomorphism on Čech cochains since \mathcal{F} is a sheaf.

We claim that partitions of pullbacks of open covers of Y form a cofinal system of open covers of X . This is clear using Lemma A.3.2 and the fact that f has finite fibers. It follows from this cofinality that f^* is an isomorphism. \square

Lemma A.4.9 (\check{H}^\bullet and \check{H}_c^\bullet commute with finite quotients) *Let X be a space and let $\pi: X \rightarrow X/\Gamma$ be the quotient map under a finite group action. Let \mathcal{F} be any sheaf of $\mathbb{Z}[1/\#\Gamma]$ -modules on X/Γ . Then the following maps are all isomorphisms:*

$$\begin{aligned} \check{H}^\bullet(X/\Gamma; \mathcal{F}) &\rightarrow \check{H}^\bullet(X/\Gamma; (\pi_* \pi^* \mathcal{F})^\Gamma) \rightarrow \check{H}^\bullet(X/\Gamma; \pi_* \pi^* \mathcal{F})^\Gamma \rightarrow \check{H}^\bullet(X; \pi^* \mathcal{F})^\Gamma, \\ \check{H}_c^\bullet(X/\Gamma; \mathcal{F}) &\rightarrow \check{H}_c^\bullet(X/\Gamma; (\pi_* \pi^* \mathcal{F})^\Gamma) \rightarrow \check{H}_c^\bullet(X/\Gamma; \pi_* \pi^* \mathcal{F})^\Gamma \rightarrow \check{H}_c^\bullet(X; \pi^* \mathcal{F})^\Gamma \end{aligned}$$

(here we note that by functoriality, $\pi^* \mathcal{F}$ is Γ -equivariant, and thus Γ acts on $\pi_* \pi^* \mathcal{F}$).

Proof Isomorphism one: the natural map $\mathcal{F} \rightarrow (\pi_* \pi^*)^\Gamma$ is in fact an isomorphism of sheaves (check on stalks). Isomorphism two: obvious since taking Γ -invariants is exact on $\mathbb{Z}[1/\#\Gamma]$ -modules. Isomorphism three: use Lemma A.4.8, which applies since π is automatically proper. \square

A.4.2 ... of complexes of \mathcal{K} -presheaves

Definition A.4.10 (\check{H}^\bullet of complexes of \mathcal{K} -presheaves) Let \mathcal{F}^\bullet be a complex of \mathcal{K} -presheaves on a compact space X . We define

$$(A.4.6) \quad \check{H}^\bullet(X; \mathcal{F}^\bullet) := \varinjlim_{\substack{X = \bigcup_{i=1}^n K_i \\ \text{finite compact cover}}} H^\bullet \left[\bigoplus_{p \geq 0} \bigoplus_{1 \leq i_0 < \dots < i_p \leq n} \mathcal{F}^{\bullet-p} \left(\bigcap_{j=0}^p K_{i_j} \right) \right]$$

with the standard Čech differential (plus the internal differential of \mathcal{F}^\bullet). We also define $\check{H}^\bullet(X; \mathcal{F})$ for any \mathcal{K} -presheaf \mathcal{F} by viewing it as a complex concentrated in degree zero.

Lemma A.4.11 (Two definitions of \check{H}^\bullet on $\text{Shv } X = \text{Shv}_{\mathcal{K}} X$ agree) *Let \mathcal{F} be a sheaf on a compact space X . Then there is a natural isomorphism $\check{H}^\bullet(X; \mathcal{F}) \rightarrow \check{H}^\bullet(X; \alpha^* \mathcal{F})$ ($\alpha^* \mathcal{F}$ is a \mathcal{K} -sheaf; see Definition A.1.6).*

Proof Let us denote by $\check{H}^\bullet(X; \mathcal{F}; \{U_\alpha\}_{\alpha \in A})$ and $\check{H}^\bullet(X; \alpha^* \mathcal{F}; \{K_i\}_{i=1}^n)$ the arguments of the direct limits (A.4.1) and (A.4.6), respectively.

Since X is compact, every open cover has a finite compact refinement. This gives a map $\check{H}^\bullet(X; \mathcal{F}) \rightarrow \check{H}^\bullet(X; \alpha^* \mathcal{F})$. To show that this map is an isomorphism, it suffices to show that for any fixed finite compact cover $\{K_i\}_{i=1}^n$, the map

$$(A.4.7) \quad \varinjlim_{\substack{X = \bigcup_{\alpha \in A} U_\alpha \\ \text{open cover refined by } \{K_i\}_{i=1}^n}} \check{H}^\bullet(X; \mathcal{F}; \{U_\alpha\}_{\alpha \in A}) \rightarrow \check{H}^\bullet(X; \alpha^* \mathcal{F}; \{K_i\}_{i=1}^n)$$

is an isomorphism. By a cofinality argument, we can change the directed system on the left side to be open covers $\{U_i\}_{i=1}^n$ with $U_i \supseteq K_i$. Hence it suffices to show that

$$(A.4.8) \quad \varinjlim_{\{U_i \supseteq K_i\}_{i=1}^n} \check{H}^\bullet(X; \mathcal{F}; \{U_i\}_{i=1}^n) \rightarrow \check{H}^\bullet(X; \alpha^* \mathcal{F}; \{K_i\}_{i=1}^n)$$

is an isomorphism. This is clear from the definition of α^* and from Lemma A.1.7. \square

Lemma A.4.12 (\check{H}^\bullet preserves quasi-isomorphisms) *Suppose $\mathcal{F}^\bullet \rightarrow \mathcal{G}^\bullet$ is a quasi-isomorphism of complexes of \mathcal{K} -presheaves. Then the induced map $\check{H}^\bullet(X; \mathcal{F}^\bullet) \rightarrow \check{H}^\bullet(X; \mathcal{G}^\bullet)$ is an isomorphism.*

Proof There is clearly a long exact sequence

$$(A.4.9) \quad \dots \rightarrow \check{H}^{\bullet-1}(X; \mathcal{G}^\bullet) \rightarrow \check{H}^\bullet(X; [\mathcal{F}^\bullet \rightarrow \mathcal{G}^{\bullet-1}]) \rightarrow \check{H}^\bullet(X; \mathcal{F}^\bullet) \rightarrow \check{H}^\bullet(X; \mathcal{G}^\bullet) \rightarrow \dots$$

Hence it suffices to show that if \mathcal{F}^\bullet is acyclic then $\check{H}^\bullet(X; \mathcal{F}^\bullet) = 0$. This is true because then each Čech complex has a finite filtration whose associated graded is acyclic. \square

Lemma A.4.13 (Hypercohomology spectral sequence) *Let \mathcal{F}^\bullet be a bounded-below complex of \mathcal{K} -presheaves. Then there is a convergent spectral sequence $E_1^{p,q} = \check{H}^q(X; \mathcal{F}^p) \Rightarrow \check{H}^{p+q}(X; \mathcal{F}^\bullet)$.*

Proof This is just the spectral sequence of the Čech double complex. \square

Proposition A.4.14 (A homotopy \mathcal{K} -sheaf calculates its own \check{H}^\bullet) *If \mathcal{F}^\bullet satisfies (hSh $_{\mathcal{K}}$ 1) and (hSh $_{\mathcal{K}}$ 2), then the canonical map $H^\bullet \mathcal{F}^\bullet(X) \rightarrow \check{H}^\bullet(X; \mathcal{F}^\bullet)$ is an isomorphism.*

Proof We prove that $\mathcal{F}^\bullet(X) \rightarrow \check{C}^\bullet(X; \mathcal{F}^\bullet; K_1, \dots, K_n)$ (the right-hand side denotes the Čech complex for the finite compact cover $X = K_1 \cup \dots \cup K_n$) is a quasi-isomorphism by induction on n . The base case $n = 1$ is obvious since $\check{C}^\bullet(X; \mathcal{F}^\bullet; X) = \mathcal{F}^\bullet(X)$ by definition. For the inductive step, it suffices to show that the natural map

$$\check{C}^\bullet(X; \mathcal{F}^\bullet; K_1 \cup K_2, K_3, \dots, K_n) \rightarrow \check{C}^\bullet(X; \mathcal{F}^\bullet; K_1, K_2, K_3, \dots, K_n)$$

is a quasi-isomorphism. We will show that the mapping cone is acyclic; to see this, let us filter it according to how many of the K_3, \dots, K_n are chosen among i_0, \dots, i_p . This is a finite filtration, so it suffices to show that the associated graded is acyclic. The associated graded is a direct sum of complexes of the form $[\mathcal{F}^\bullet(K) \rightarrow \mathcal{F}^{\bullet-1}(K)]$ (which is obviously acyclic) and

$$[\mathcal{F}^\bullet(K \cap (K_1 \cup K_2)) \rightarrow \mathcal{F}^{\bullet-1}(K \cap K_1) \oplus \mathcal{F}^{\bullet-1}(K \cap K_2) \rightarrow \mathcal{F}^{\bullet-2}(K \cap K_1 \cap K_2)]$$

(which is acyclic by (hSh $_{\mathcal{K}}$ 2)).

The above argument works when $X \neq \emptyset$; if $X = \emptyset$, use (hSh $_{\mathcal{K}}$ 1). \square

Lemma A.4.15 (A \mathcal{K} -sheaf calculates its own \check{H}^0) *If \mathcal{F} satisfies (Sh $_{\mathcal{K}}$ 1) and (Sh $_{\mathcal{K}}$ 2), then the canonical map $\mathcal{F}(X) \rightarrow \check{H}^0(X; \mathcal{F})$ is an isomorphism.*

Proof Use induction as in the proof of Proposition A.4.14. \square

Lemma A.4.16 (\check{H}^\bullet is determined by stalks I) *Let \mathcal{F} satisfy (Sh $_{\mathcal{K}}$ 3). If the stalks of \mathcal{F} vanish, then $\check{H}^\bullet(X; \mathcal{F}) = 0$.*

Proof It suffices to show that for all $\alpha \in \mathcal{F}(K)$, there exists a finite compact cover $X = \bigcup_{i=1}^n K_i$ such that $\mathcal{F}(K) \rightarrow \bigoplus_{i=1}^n \mathcal{F}(K \cap K_i)$ annihilates α . We consider the commutative diagram

$$(A.4.10) \quad \begin{array}{ccc} \varinjlim_{K \subseteq U} \mathcal{F}(\bar{U}) & \longrightarrow & \prod_{p \in K} \varinjlim_{p \in U} \mathcal{F}(\bar{U}) \\ \downarrow & & \downarrow \\ \mathcal{F}(K) & \longrightarrow & \prod_{p \in K} \mathcal{F}_p \end{array}$$

where the vertical maps are both isomorphisms. Now the vanishing of \mathcal{F}_p and a compactness argument shows that there are finitely many open sets $U_i \subseteq X$ covering K such that α vanishes in $\mathcal{F}(K \cap \bar{U}_i)$ for all i . Thus the compact cover

$$X = \left(X \setminus \bigcup_{i=1}^n U_i \right) \cup \bar{U}_1 \cup \dots \cup \bar{U}_n$$

has the desired properties. □

Lemma A.4.17 (\check{H}^\bullet is determined by stalks II) *Let \mathcal{F}^\bullet satisfy (hSh_K3) and $H^i \mathcal{F}^\bullet = 0$ for $i \ll 0$. If \mathcal{F}^\bullet has acyclic stalks, then $\check{H}^\bullet(X; \mathcal{F}^\bullet) = 0$.*

Proof We show that for all i , the map $\check{H}^\bullet(X; \mathcal{F}^\bullet) \rightarrow \check{H}^\bullet(X; \tau_{\geq i} \mathcal{F}^\bullet)$ is an isomorphism (this is sufficient since $\check{H}^j(X; \tau_{\geq i} \mathcal{F}^\bullet) = 0$ for $j < i$). We proceed by induction on i .

Since $H^i \mathcal{F}^\bullet = 0$ for $i \ll 0$, we have that $\mathcal{F}^\bullet \rightarrow \tau_{\geq i} \mathcal{F}^\bullet$ is a quasi-isomorphism, so $\check{H}^\bullet(X; \mathcal{F}^\bullet) \rightarrow \check{H}^\bullet(X; \tau_{\geq i} \mathcal{F}^\bullet)$ is an isomorphism for $i \ll 0$ (Lemma A.4.12). Thus we have the base case of the induction. For the inductive step, it suffices to show that $\check{H}^\bullet(X; H^{i-1} \mathcal{F}^\bullet) = 0$. This follows from Lemma A.4.16. □

Proposition A.4.18 (\check{H}^\bullet is determined by stalks III) *Let $\mathcal{F}^\bullet, \mathcal{G}^\bullet$ satisfy (hSh_K3) and $H^i \mathcal{F}^\bullet = H^i \mathcal{G}^\bullet = 0$ for $i \ll 0$. If $\mathcal{F}^\bullet \rightarrow \mathcal{G}^\bullet$ induces a quasi-isomorphism on stalks, then it induces an isomorphism $\check{H}^\bullet(X; \mathcal{F}^\bullet) \rightarrow \check{H}^\bullet(X; \mathcal{G}^\bullet)$.*

Proof Recall the long exact sequence (A.4.9) and apply Lemma A.4.17 to the mapping cone $[\mathcal{F}^\bullet \rightarrow \mathcal{G}^{\bullet-1}]$. □

Corollary A.4.19 (A map of homotopy \mathcal{K} -sheaves being a quasi-isomorphism can be checked on stalks) *Let $\mathcal{F}^\bullet \rightarrow \mathcal{G}^\bullet$ be a map of homotopy \mathcal{K} -sheaves which satisfy $H^i \mathcal{F}^\bullet = H^i \mathcal{G}^\bullet = 0$ for $i \ll 0$. Then $\mathcal{F}^\bullet \rightarrow \mathcal{G}^\bullet$ is a quasi-isomorphism if and only if $\mathcal{F}_p^\bullet \rightarrow \mathcal{G}_p^\bullet$ is a quasi-isomorphism for all $p \in X$.*

Proof For any $K \subseteq X$, we have a commutative diagram:

$$(A.4.11) \quad \begin{array}{ccc} H^\bullet \mathcal{F}^\bullet(K) & \xrightarrow{\sim} & \check{H}^\bullet(K; \mathcal{F}^\bullet) \\ \downarrow & & \downarrow \\ H^\bullet \mathcal{G}^\bullet(K) & \xrightarrow{\sim} & \check{H}^\bullet(K; \mathcal{G}^\bullet) \end{array}$$

The rows are isomorphisms by Proposition A.4.14. If $\mathcal{F}_p^\bullet \rightarrow \mathcal{G}_p^\bullet$ is a quasi-isomorphism, then the right vertical map is a quasi-isomorphism by Proposition A.4.18. \square

Lemma A.4.20 (Long exact sequence for \check{H}^\bullet) *Let $\mathcal{F}, \mathcal{G}, \mathcal{H}$ satisfy (Sh $_{\mathcal{K}3}$). If $0 \rightarrow \mathcal{F} \rightarrow \mathcal{G} \rightarrow \mathcal{H} \rightarrow 0$ is exact on stalks, then it induces a long exact sequence on \check{H}^\bullet .*

Proof By Lemma A.4.17, we have $\check{H}^\bullet(X; [\mathcal{F} \rightarrow \mathcal{G}[-1] \rightarrow \mathcal{H}[-2]]) = 0$. Now inspection of the hypercohomology spectral sequence (Lemma A.4.13) gives the desired long exact sequence. \square

Definition A.4.21 (Pullback on \check{H}^\bullet) Let $f: X \rightarrow Y$ be a map of compact spaces. A finite compact cover of Y pulls back to give a finite compact cover of X , and this gives an identification of the Čech complex for the cover of Y with coefficients in $f_* \mathcal{F}^\bullet$ with the Čech complex for the cover of X with coefficients in \mathcal{F}^\bullet . Hence we get a natural map:

$$(i) \quad f^*: \check{H}^\bullet(Y; f_* \mathcal{F}^\bullet) \rightarrow \check{H}^\bullet(X; \mathcal{F}^\bullet) \text{ for } \mathcal{F}^i \in \text{Prshv}_{\mathcal{K}} X.$$

A.5 Pure homotopy \mathcal{K} -sheaves

By Proposition A.4.14, a homotopy \mathcal{K} -sheaf can be thought of as a *resolution* (that is, its global sections compute the cohomology of some complex of sheaves). In this section, we introduce the notion of a *pure* homotopy \mathcal{K} -sheaf, which may be thought of as a resolution of a sheaf (as opposed to a complex of sheaves). More specifically, a pure homotopy \mathcal{K} -sheaf \mathcal{F}^\bullet “is” a resolution of $H^0 \mathcal{F}^\bullet$ (which by Lemma A.5.3 is always a sheaf).

Definition A.5.1 We say that a homotopy \mathcal{K} -sheaf \mathcal{F}^\bullet on X is *pure* if and only if:

- (i) **Stalk cohomology** $H^i \mathcal{F}_p^\bullet = 0$ for $i \neq 0$ and all $p \in X$.
- (ii) **Weak vanishing** $H^i \mathcal{F}^\bullet = 0$ for $i \ll 0$ locally on X (meaning that for all $p \in X$, there exists an open set $U \subseteq X$ containing p and an integer $N > -\infty$ such that $H^i \mathcal{F}^\bullet(K) = 0$ for all $K \subseteq U$ and $i \leq N$).

Remark A.5.2 It would be nice to know whether the stalk cohomology condition implies the weak vanishing condition in general (it would be much easier to check purity).

Lemma A.5.3 *Let \mathcal{F}^\bullet be a pure homotopy \mathcal{K} -sheaf. Then:*

- (i) **Strong vanishing** $H^i \mathcal{F}^\bullet = 0$ for $i < 0$.
- (ii) $H^0 \mathcal{F}^\bullet$ is a \mathcal{K} -sheaf.

Proof By Lemma A.2.12, strong vanishing implies that $H^0 \mathcal{F}^\bullet$ is a \mathcal{K} -sheaf. Now let us prove strong vanishing. By restricting to a compact subset, it suffices to treat the case when the underlying space is compact. Now from (hSh \mathcal{K} 2), compactness, and weak vanishing, it follows that $H^i \mathcal{F}^\bullet = 0$ for $i \ll 0$. Now let us prove strong vanishing by induction on $i < 0$ (we have just proven the base case). For the inductive step, observe that $H^i \mathcal{F}^\bullet$ is a sheaf by the induction hypothesis ($H^{i-1} \mathcal{F}^\bullet = 0$) and Lemma A.2.12, and thus $H^i \mathcal{F}_p^\bullet = 0 \implies H^i \mathcal{F}^\bullet = 0$. □

Proposition A.5.4 (\check{H}^\bullet of a pure homotopy \mathcal{K} -sheaf) *Let \mathcal{F}^\bullet be a pure homotopy \mathcal{K} -sheaf. Then there is a canonical isomorphism*

$$(A.5.1) \quad H^\bullet \mathcal{F}^\bullet(X) = \check{H}^\bullet(X; H^0 \mathcal{F}^\bullet).$$

More generally, let $[\mathcal{F}_0^\bullet \rightarrow \mathcal{F}_1^{\bullet-1} \rightarrow \dots \rightarrow \mathcal{F}_n^{\bullet-n}]$ be a complex of \mathcal{K} -presheaves where each \mathcal{F}_i^\bullet is a pure homotopy \mathcal{K} -sheaf. Then there is a canonical isomorphism

$$(A.5.2) \quad H^\bullet[\mathcal{F}_0^\bullet(X) \rightarrow \dots \rightarrow \mathcal{F}_n^{\bullet-n}(X)] = \check{H}^\bullet(X; [H^0 \mathcal{F}_0^\bullet \rightarrow \dots \rightarrow (H^0 \mathcal{F}_n^\bullet)[-n]]).$$

Proof The isomorphism (A.5.2) is defined as the composition of the following isomorphisms:

$$(A.5.3) \quad \begin{array}{c} H^\bullet[\mathcal{F}_0^\bullet(X) \rightarrow \dots \rightarrow \mathcal{F}_n^{\bullet-n}(X)] \\ \downarrow \\ \check{H}^\bullet(X; [\mathcal{F}_0^\bullet \rightarrow \dots \rightarrow \mathcal{F}_n^{\bullet-n}]) \\ \downarrow \\ \check{H}^\bullet(X; [\tau_{\geq 0} \mathcal{F}_0^\bullet \rightarrow \dots \rightarrow \tau_{\geq n} \mathcal{F}_n^{\bullet-n}]) \\ \uparrow \\ \check{H}^\bullet(X; [H^0 \mathcal{F}_0^\bullet \rightarrow \dots \rightarrow H^n \mathcal{F}_n^{\bullet-n}]) \end{array}$$

The maps are isomorphisms for the following reasons. Map one: $[\mathcal{F}_0^\bullet \rightarrow \dots \rightarrow \mathcal{F}_n^{\bullet-n}]$ is a homotopy \mathcal{K} -sheaf (Lemma A.2.11), and a homotopy \mathcal{K} -sheaf calculates its own \check{H}^\bullet (Proposition A.4.14). Map two: the map of coefficient \mathcal{K} -presheaves is a quasi-isomorphism (Lemma A.5.3) and thus induces an isomorphism on \check{H}^\bullet (Lemma A.4.12). Map three: the map of coefficient \mathcal{K} -presheaves is a quasi-isomorphism on stalks (by purity), and thus induces an isomorphism on \check{H}^\bullet (Proposition A.4.18). □

Lemma A.5.5 (Checking purity on a cover) *Let \mathcal{F}^\bullet be a homotopy \mathcal{K} -sheaf. Write $X = U \cup Z$ with U open and Z closed, and suppose that $i^*\mathcal{F}^\bullet$ and $j^*\mathcal{F}^\bullet$ are both pure ($i: Z \hookrightarrow X$ and $j: U \hookrightarrow X$). Then \mathcal{F}^\bullet is pure.*

Proof It suffices to show the weak vanishing property for \mathcal{F}^\bullet . Let $\alpha \in H^i \mathcal{F}^\bullet(K)$ be arbitrary with $i < 0$. By Remark A.1.5 and Lemma A.2.13, the map

$$(A.5.4) \quad \varinjlim_{\substack{Z \subseteq V \\ V \text{ open}}} \mathcal{F}^\bullet(K \cap \bar{V}) \rightarrow \mathcal{F}^\bullet(K \cap Z)$$

is a quasi-isomorphism. Since $H^i \mathcal{F}^\bullet(K \cap Z) = 0$ by strong vanishing for $i^*\mathcal{F}^\bullet$, we see that the image of α in $H^i \mathcal{F}^\bullet(K \cap \bar{V})$ vanishes for some open $V \supseteq Z$. Now applying (hSh $_{\mathcal{K}2}$) to $K = (K \cap \bar{V}) \cup (K \setminus V)$, we see that the vanishing of the image of α in $H^i \mathcal{F}^\bullet(K \cap \bar{V})$ implies that α “comes from” the cohomology of

$$[\mathcal{F}^\bullet(K \setminus V) \rightarrow \mathcal{F}^{\bullet-1}(K \cap \bar{V} \cap (K \setminus V))].$$

On the other hand, this latter group vanishes in degrees $i < 0$ by strong vanishing for $j^*\mathcal{F}^\bullet$. Thus $\alpha = 0$, so we have even shown strong vanishing for \mathcal{F}^\bullet . \square

A.6 Poincaré–Lefschetz duality

We prove a version of Poincaré duality for arbitrary closed subsets of a topological manifold. This proof is a good illustration of the tools we have developed concerning pure homotopy \mathcal{K} -sheaves (which arise naturally in the proof).

We observed in Example A.2.4 that $U \mapsto C^\bullet(U)$ is a homotopy sheaf on any space (and it should be thought of as a resolution of the constant sheaf). To prove Poincaré duality for a topological manifold M of dimension n , we will show that $K \mapsto C_{\dim M - \bullet}(M, M \setminus K)$ is a pure homotopy \mathcal{K} -sheaf and calculate its H^0 as \mathfrak{o}_M (in other words, it should be thought of as a resolution of the orientation sheaf of M).

Convention A.6.1 Throughout this paper, we make no second countability or paracompactness assumptions on manifolds (topological or smooth).

Definition A.6.2 (Orientation sheaf of manifold) Let M be a topological manifold. We let \mathfrak{o}_M denote the *orientation sheaf*⁴³ of M (the corresponding \mathcal{K} -sheaf is defined by $\mathfrak{o}_M(K) := H_{\dim M}(M, M \setminus K)$); it is locally isomorphic to the constant sheaf $\underline{\mathbb{Z}}$ (and

⁴³Note that the fundamental class lies in homology twisted by \mathfrak{o}_M^\vee .

following Convention A.0.3, it has parity $\dim M \in \mathbb{Z}/2$). Now let M be a topological manifold with boundary, and define orientation sheaves on M

$$(A.6.1) \quad \mathfrak{o}_M := j_*\mathfrak{o}_{M \setminus \partial M},$$

$$(A.6.2) \quad \mathfrak{o}_{M \text{ rel } \partial} := j_!\mathfrak{o}_{M \setminus \partial M},$$

where $j: M \setminus \partial M \hookrightarrow M$. Then \mathfrak{o}_M is locally isomorphic to the constant sheaf $\underline{\mathbb{Z}}$, and there is a sequence of sheaves

$$(A.6.3) \quad 0 \rightarrow \mathfrak{o}_{M \text{ rel } \partial} \rightarrow \mathfrak{o}_M \rightarrow i_*\mathfrak{o}_{\partial M} \rightarrow 0$$

which is exact on stalks, where $i: \partial M \hookrightarrow M$, and the second map comes from the boundary map in the long exact sequence of the pair $(M, \partial M)$.

Lemma A.6.3 (Homotopy \mathcal{K} -sheaf axioms for singular chains) *Let X be a topological space. Then we have:*

- (i) $C_\bullet(X, X)$ is acyclic.
- (ii) Let $A, B \subseteq X$ be closed. Then the complex

$$[C_\bullet(X, X \setminus (A \cup B)) \rightarrow C_{\bullet+1}(X, X \setminus A) \oplus C_{\bullet+1}(X, X \setminus B) \rightarrow C_{\bullet+2}(X, X \setminus (A \cap B))]$$

is acyclic.

- (iii) Let $K = \bigcap_{\alpha \in A} K_\alpha$ where $\{K_\alpha\}_{\alpha \in A}$ is a family of closed subsets of X which is filtered in the sense that for all $\alpha_1, \alpha_2 \in A$, there exists $\beta \in A$ with $K_\beta \subseteq K_{\alpha_1} \cap K_{\alpha_2}$. Then $\varinjlim_{\alpha \in A} C_\bullet(X, X \setminus K_\alpha) \rightarrow C_\bullet(X, X \setminus K)$ is a quasi-isomorphism.

Proof Statement (i) is obvious.

Statement (ii) can be deduced from Mayer–Vietoris using a form of the nine lemma, as we now explain. Let us write $U := X \setminus A$ and $V := X \setminus B$. Now consider the following total complex:

$$(A.6.4) \quad \begin{array}{ccccc} C_\bullet(U \cap V) & \longrightarrow & C_{\bullet+1}(U) \oplus C_{\bullet+1}(V) & \longrightarrow & C_{\bullet+2}(U \cup V) \\ \downarrow & & \downarrow & & \downarrow \\ C_{\bullet+1}(X) & \longrightarrow & C_{\bullet+2}(X) \oplus C_{\bullet+2}(X) & \longrightarrow & C_{\bullet+3}(X) \\ \downarrow & & \downarrow & & \downarrow \\ C_{\bullet+2}(X, U \cap V) & \longrightarrow & C_{\bullet+3}(X, U) \oplus C_{\bullet+3}(X, V) & \longrightarrow & C_{\bullet+4}(X, U \cup V) \end{array}$$

The columns are acyclic (by definition of relative chains), and hence the total complex is acyclic as well. The first row is acyclic by Mayer–Vietoris, and the second row is obviously acyclic. Thus the third row is acyclic, as needed.

Statement (iii) is true because the map is in fact an isomorphism on the chain level. It is clearly surjective; to show injectivity we must show that a singular chain on X which is disjoint from K is in fact disjoint from K_α for some $\alpha \in A$. This follows because the standard n -simplex is compact. \square

Lemma A.6.4 (Poincaré–Lefschetz duality) *Let M be a topological manifold of dimension n with boundary. Let $i: X \hookrightarrow M$ be a closed subset. Let $N \subseteq \partial M$ (closed) be a tamely embedded codimension-zero submanifold with boundary.⁴⁴ Then there is a canonical isomorphism*

$$(A.6.5) \quad H^\bullet[C_{n-1-\bullet}(N, N \setminus X) \rightarrow C_{n-\bullet}(M, M \setminus X)] = \check{H}_c^\bullet(X; i^* j_! j^* \mathfrak{o}_M),$$

where $j: M^\circ \cup N^\circ \hookrightarrow M$. In particular, specializing to $N = \emptyset$, there is a canonical isomorphism

$$(A.6.6) \quad H_{n-\bullet}(M, M \setminus X) = \check{H}_c^\bullet(X; i^* \mathfrak{o}_{M \text{ rel } \partial M}).$$

Proof Let X^+ be the one-point compactification of X . Define a complex of \mathcal{K} -presheaves \mathcal{F}^\bullet on X^+ by

$$(A.6.7) \quad \mathcal{F}^\bullet(K) := [C_{n-1-\bullet}(N, N \setminus K) \rightarrow C_{n-\bullet}(M, M \setminus K)]$$

(where on the right-hand side by K we really mean $K \cap X$). Applying Lemmas A.6.3 and A.2.11, we see that \mathcal{F}^\bullet is a homotopy \mathcal{K} -sheaf.

We claim that \mathcal{F}^\bullet is a pure homotopy \mathcal{K} -sheaf. Certainly the homology of \mathcal{F}^\bullet is bounded below, because the singular chain complex of a topological manifold has homology bounded above and \mathcal{F}^\bullet is built out of these. Now, it is easy to calculate

$$(A.6.8) \quad H^\bullet \mathcal{F}_p^\bullet = \begin{cases} \mathbb{Z}, & p \in M^\circ \cup N^\circ, \\ 0, & p \notin M^\circ \cup N^\circ, \end{cases}$$

concentrated in degree zero (see [39, page 231, Section 3.3] for the special case $N = \partial M = \emptyset$). Hence \mathcal{F}^\bullet is a pure homotopy \mathcal{K} -sheaf. In fact, it is not hard to see (using the adjunction $j_! \dashv j^*$) that (A.6.8) lifts to an isomorphism of sheaves $f_! i^* j_! j^* \mathfrak{o}_M \rightarrow H^0 \mathcal{F}^\bullet$, where $f: X \hookrightarrow X^+$.

Now we conclude

$$H^\bullet \mathcal{F}^\bullet(X^+) \stackrel{\text{Prop A.5.4}}{=} \check{H}^\bullet(X^+; f_! i^* j_! j^* \mathfrak{o}_M) \stackrel{\text{Lem A.4.7}}{=} \check{H}_c^\bullet(X; i^* j_! j^* \mathfrak{o}_M).$$

Now observe that $H^\bullet \mathcal{F}^\bullet(X^+)$ is the left-hand side of (A.6.5). \square

⁴⁴In other words, $N \subseteq \partial M$ is a closed subset which locally looks like either $\emptyset \subseteq \mathbb{R}^{n-1}$, $\mathbb{R}_{\geq 0} \times \mathbb{R}^{n-2} \subseteq \mathbb{R}^{n-1}$, or $\mathbb{R}^{n-1} \subseteq \mathbb{R}^{n-1}$.

Remark A.6.5 (Why homotopy \mathcal{K} -sheaves instead of homotopy sheaves?) The naive modification of (A.6.7) substituting open $U \subseteq X$ in place of compact $K \subseteq X$ does not yield a homotopy sheaf. To get a homotopy sheaf, one could apply the proposed functor $R\alpha_*: \text{hShv}_{\mathcal{K}} X \rightarrow \text{hShv } X$ from Remark A.2.8 to the homotopy \mathcal{K} -sheaf (A.6.7). It is somewhat easier, though, to just work directly in the setting of homotopy \mathcal{K} -sheaves (which has some advantages, for example stalks of \mathcal{K} -presheaves are easier to define/understand). It is for this reason that throughout this paper we work with homotopy \mathcal{K} -sheaves instead of homotopy sheaves.

A.7 Homotopy colimits

We make common use of the following type of “homotopy diagram” and its corresponding “homotopy colimit”.

Definition A.7.1 (Homotopy diagram) Let \mathcal{S} be a finite poset. A *homotopy diagram* over \mathcal{S} is a collection of complexes $\{A_{\mathfrak{s},\mathfrak{t}}^\bullet\}_{\mathfrak{s} \leq \mathfrak{t} \in \mathcal{S}}$ equipped with compatible maps $A_{\mathfrak{s},\mathfrak{t}}^\bullet \rightarrow A_{\mathfrak{s}',\mathfrak{t}'}^\bullet$ for $\mathfrak{s} \leq \mathfrak{s}' \leq \mathfrak{t}' \leq \mathfrak{t}$ (meaning that $A_{\mathfrak{s},\mathfrak{t}}^\bullet \rightarrow A_{\mathfrak{s}',\mathfrak{t}'}^\bullet \rightarrow A_{\mathfrak{s}'',\mathfrak{t}''}^\bullet$ and $A_{\mathfrak{s},\mathfrak{t}}^\bullet \rightarrow A_{\mathfrak{s}'',\mathfrak{t}''}^\bullet$ agree).

Definition A.7.2 (Homotopy colimit) Let \mathcal{S} be a finite poset and let $\{A_{\mathfrak{s},\mathfrak{t}}^\bullet\}_{\mathfrak{s} \leq \mathfrak{t} \in \mathcal{S}}$ be a homotopy diagram over \mathcal{S} . We define the *homotopy colimit*

$$(A.7.1) \quad \text{hocolim}_{\mathfrak{s} \leq \mathfrak{t} \in \mathcal{S}} A_{\mathfrak{s},\mathfrak{t}}^\bullet := \bigoplus_{p \geq 0} \bigoplus_{\mathfrak{s}_0 < \dots < \mathfrak{s}_p \in \mathcal{S}} A_{\mathfrak{s}_0, \mathfrak{s}_p}^{\bullet+p}$$

with differential (decreasing p) given by the alternating sum over forgetting one of the \mathfrak{s}_i (plus the internal differential).⁴⁵ Loosely speaking, we are “gluing together” the $\{A_{\mathfrak{s},\mathfrak{s}}^\bullet\}_{\mathfrak{s} \in \mathcal{S}}$ along the “morphisms” $A_{\mathfrak{s},\mathfrak{s}}^\bullet \leftarrow A_{\mathfrak{s},\mathfrak{t}}^\bullet \rightarrow A_{\mathfrak{t},\mathfrak{t}}^\bullet$ for $\mathfrak{s} \leq \mathfrak{t}$.

Lemma A.7.3 (Terminal object for hocolim) *Let \mathcal{S} be a finite poset with unique maximal element $\mathfrak{s}^{\text{top}}$, and let $\{A_{\mathfrak{s},\mathfrak{t}}^\bullet\}_{\mathfrak{s} \leq \mathfrak{t} \in \mathcal{S}}$ be a homotopy diagram with the property that every map $A_{\mathfrak{s},\mathfrak{t}}^\bullet \rightarrow A_{\mathfrak{s}',\mathfrak{t}'}^\bullet$ ($\mathfrak{s} \leq \mathfrak{t}' \leq \mathfrak{t}$) is a quasi-isomorphism. Then the natural inclusion $A_{\mathfrak{s}^{\text{top}},\mathfrak{s}^{\text{top}}}^\bullet \rightarrow \text{hocolim}_{\mathfrak{s} \leq \mathfrak{t} \in \mathcal{S}} A_{\mathfrak{s},\mathfrak{t}}^\bullet$ is a quasi-isomorphism.*

Proof We filter $\text{hocolim}_{\mathfrak{s} \leq \mathfrak{t} \in \mathcal{S}} A_{\mathfrak{s},\mathfrak{t}}^\bullet$ by the number of $\mathfrak{s}_0, \dots, \mathfrak{s}_p$ which are not equal to $\mathfrak{s}^{\text{top}}$. The zeroth associated graded piece is the subcomplex $A_{\mathfrak{s}^{\text{top}},\mathfrak{s}^{\text{top}}}^\bullet$, so it suffices to show that all the other associated graded pieces are acyclic. Each of these is a direct sum of mapping cones $[A_{\mathfrak{s},\mathfrak{s}^{\text{top}}}^{\bullet+1} \rightarrow A_{\mathfrak{s},\mathfrak{t}}^\bullet]$, which are acyclic by assumption. \square

⁴⁵This can be interpreted as the complex of simplicial chains on the nerve of \mathcal{S} using a coefficient system determined by $\{A_{\mathfrak{s},\mathfrak{t}}^\bullet\}$.

Lemma A.7.4 (hocolim preserves quasi-isomorphisms) *Fix a finite poset \mathcal{S} and let $\{A_{s,t}^\bullet\}_{s \leq t \in \mathcal{S}}$ and $\{B_{s,t}^\bullet\}_{s \leq t \in \mathcal{S}}$ be homotopy diagrams over \mathcal{S} . Suppose that there are compatible quasi-isomorphisms $A_{s,t}^\bullet \rightarrow B_{s,t}^\bullet$. Then the induced map $\text{hocolim}_{s \leq t \in \mathcal{S}} A_{s,t}^\bullet \rightarrow \text{hocolim}_{s \leq t \in \mathcal{S}} B_{s,t}^\bullet$ is a quasi-isomorphism.*

Proof Since the functor $\text{hocolim}_{s \leq t \in \mathcal{S}}$ commutes with the formation of mapping cones, it suffices to show that if each $A_{s,t}^\bullet$ is acyclic, then so is $\text{hocolim}_{s \leq t \in \mathcal{S}} A_{s,t}^\bullet$. This holds since in this case it has a finite filtration whose associated graded is acyclic. \square

Definition A.7.5 (Tensor product of homotopy diagrams) Let \mathcal{S} and \mathcal{T} be two finite posets, and let $\{A_{s,s'}^\bullet\}$ and $\{B_{t,t'}^\bullet\}$ be homotopy diagrams over \mathcal{S} and \mathcal{T} , respectively. Their tensor product $\{A_{s,s'}^\bullet \otimes B_{t,t'}^\bullet\}$ is naturally a homotopy diagram over $\mathcal{S} \times \mathcal{T}$. Now there is a natural morphism

$$(A.7.2) \quad \text{hocolim}_{s \leq s' \in \mathcal{S}} A_{s,s'}^\bullet \otimes \text{hocolim}_{t \leq t' \in \mathcal{T}} B_{t,t'}^\bullet \rightarrow \text{hocolim}_{s \times t \leq s' \times t' \in \mathcal{S} \times \mathcal{T}} A_{s,s'}^\bullet \otimes B_{t,t'}^\bullet.$$

To define this morphism, we simply observe that the nerve of $\mathcal{S} \times \mathcal{T}$ is the standard simplicial subdivision of the product of the nerves of \mathcal{S} and \mathcal{T} , and this is covered by a morphism of coefficient systems.⁴⁶

A.8 Homotopy colimits of pure homotopy \mathcal{K} -sheaves

We introduce a gluing construction for pure homotopy \mathcal{K} -sheaves. The relevance of this Lemma A.8.2 is best understood in its (only) intended application, namely Proposition 4.3.3 (which is best understood in context).

Lemma A.8.1 (hocolim preserves homotopy \mathcal{K} -sheaves) *Let \mathcal{S} be a finite poset and let $\{\mathcal{F}_{s,t}^\bullet\}_{s \leq t \in \mathcal{S}}$ be a homotopy diagram of homotopy \mathcal{K} -sheaves. Then the resulting $\text{hocolim}_{s \leq t \in \mathcal{S}} \mathcal{F}_{s,t}^\bullet$ is a homotopy \mathcal{K} -sheaf.*

Proof The associated graded of the p -filtration on $\text{hocolim}_{s \leq t \in \mathcal{S}} \mathcal{F}_{s,t}^\bullet$ is a homotopy \mathcal{K} -sheaf by assumption; now use Lemma A.2.11. \square

Lemma A.8.2 (Gluing pure homotopy \mathcal{K} -sheaves) *Let A be a finite set. Let $\{U_I\}_{I \subseteq A}$ be an open cover of a space X satisfying:*

- (i) $U_I \cap U_K \subseteq U_J$ for $I \subseteq J \subseteq K$.
- (ii) $U_I \cap U_{I'} \subseteq U_{I \cup I'}$ for $I, I' \subseteq A$.

⁴⁶In fact, from this perspective one easily sees that (A.7.2) is always a quasi-isomorphism.

Let \mathcal{G} be a sheaf on X and let $\mathcal{G}_{IJ} := (j_{IJ})_!(j_{IJ})^*\mathcal{G}$, where $j_{IJ}: U_I \cap U_J \hookrightarrow X$. By property (i), this gives rise to a homotopy diagram $\{\mathcal{G}_{IJ}\}_{I \subseteq A}$ over 2^A of sheaves on X .

Let $\{\mathcal{F}_{IJ}^\bullet\}_{I \subseteq J \subseteq A}$ be a homotopy diagram over 2^A of pure homotopy \mathcal{K} -sheaves on X , and suppose we give a compatible system of isomorphisms $\mathcal{G}_{IJ} \xrightarrow{\sim} H^0\mathcal{F}_{IJ}^\bullet$. Then $\mathcal{F}^\bullet := \text{hocolim}_{I \subseteq J \subseteq A} \mathcal{F}_{IJ}^\bullet$ is a pure homotopy \mathcal{K} -sheaf and there is a canonical induced isomorphism $\mathcal{G} \xrightarrow{\sim} H^0\mathcal{F}^\bullet$.

Proof Certainly \mathcal{F}^\bullet is a homotopy \mathcal{K} -sheaf by Lemma A.8.1. Since A is finite, it is easy to see that $H^i\mathcal{F}_{IJ}^\bullet$ being bounded below implies that $H^i\mathcal{F}^\bullet$ is bounded below.

Now let us calculate \mathcal{F}_p^\bullet using the spectral sequence associated to the p -filtration. The E_1 term is concentrated along the $q = 0$ row since $(H^q\mathcal{F}_{IJ}^\bullet)_p$ is concentrated in degree zero; thus there are no further differentials after the E_1 page. On the E_1 page the differentials coincide precisely with the differentials in the definition of $\text{hocolim}_{I \subseteq J \subseteq A} (\mathcal{G}_{IJ})_p$ (regarding \mathcal{G}_{IJ} as complexes concentrated in degree zero). Hence we have an isomorphism

$$(A.8.1) \quad H^\bullet\mathcal{F}_p^\bullet = H^\bullet \text{hocolim}_{I \subseteq J \subseteq A} (\mathcal{G}_{IJ})_p.$$

To calculate the right-hand side of (A.8.1), let us start with the trivial observation that

$$(A.8.2) \quad (\mathcal{G}_{IJ})_p = \begin{cases} \mathcal{G}_p, & p \in U_I \cap U_J, \\ 0, & p \notin U_I \cap U_J. \end{cases}$$

Now consider $(2^A)_p := \{I \subseteq A : p \in U_I\}$, which satisfies the following properties:

- (i) $(2^A)_p$ has a maximal element (restatement of (ii)).
- (ii) $I, K \in (2^A)_p$ implies $J \in (2^A)_p$ for $I \subseteq J \subseteq K$ (restatement of (i)).

These two properties imply that $\text{hocolim}_{I \subseteq J \subseteq A} (\mathcal{G}_{IJ})_p$ is simply \mathcal{G}_p tensored with the simplicial chain complex of the nerve of $(2^A)_p$ (which is contractible). Hence (A.8.1) gives a canonical isomorphism of stalks $H^0\mathcal{F}_p^\bullet = \mathcal{G}_p$ (and thus in particular \mathcal{F}^\bullet is pure).

Now it remains to construct a canonical isomorphism of sheaves $\mathcal{G} \xrightarrow{\sim} H^0\mathcal{F}^\bullet$. Over the open set U_I , we define this isomorphism to be the composition

$$\mathcal{G}_{II} \xrightarrow{\sim} H^0\mathcal{F}_{II}^\bullet \rightarrow H^0\mathcal{F}^\bullet$$

(the second map being the inclusion of a $p = 0$ subcomplex in the homotopy colimit). At any point $p \in U_I$, this specializes to the isomorphism of stalks $\mathcal{G}_p = H^0\mathcal{F}_p^\bullet$ defined earlier. Hence since $\{U_I\}_{I \subseteq A}$ is an open cover of X , these isomorphisms patch together to give the desired global isomorphism. □

A.9 Steenrod homology

Steenrod homology \overline{H}_\bullet is a homology theory for compact Hausdorff spaces. It is characterized uniquely by a certain set of axioms due to Berikashvili [7; 8] (see also Goldfarb [35, page 355, Section 7.4] and Inassaridze [48]); a simpler axiomatic characterization of Steenrod homology on compact metrizable spaces is due to Milnor [64]. Note that it follows as usual from these axioms that Steenrod homology coincides with singular homology on finite CW-complexes. Steenrod homology is due to Steenrod for compact metrizable spaces [76] and was later generalized to compact Hausdorff spaces (some sources include Edwards and Hastings [18], Hastings [38], Carlsson and Pedersen [12] and Goldfarb [35]). Another reference is Mardešić’s book [59] (note that Mardešić studies the more general theory of strong homology, which coincides with Steenrod homology on compact spaces).

A.9.1 Čech cochains The following definition is due to Carlsson and Pedersen [12]. The idea of using open covers indexed by the points of the space being covered goes back at least to Godement [34, Chapter II, Section 5.8], and was also used by Friedlander [27] in the context of étale homotopy theory.

Definition A.9.1 (Rigid open cover) A *rigid open cover* of a compact space X consists of open sets $\{U_x \subseteq X\}_{x \in X}$ such that $x \in U_x$, $\{x : \overline{U_x} = U\} \subseteq U$, and $\#\{U : U = U_x \text{ for some } x\} < \infty$. Warning: even if $U_x = U_y$ for some $x \neq y$, they are still different elements of the cover; in particular, the nerve of a rigid cover is always infinite unless X is finite.

For doing Čech theory, the category of rigid open covers is technically more convenient than the usual category of open covers (as in Remark A.4.1). Specifically, the collection of rigid covers forms a set, and there is at most one morphism between any pair of rigid covers (we only consider refinements which act as the identity on the index set X).

Definition A.9.2 (\check{C}^\bullet and \check{C}_c^\bullet of sheaves; compare Definition A.4.3) Let \mathcal{F} be a sheaf. We define the Čech cochains

$$(A.9.1) \quad \check{C}^\bullet(X; \mathcal{F}) := \varinjlim_{\substack{\{U_x \subseteq X\}_{x \in X} \\ \text{rigid open cover}}} \bigoplus_{p \geq 0} \prod_{\substack{S \subseteq X \\ |S|=p+1}} \mathcal{F}\left(\bigcap_{x \in S} U_x\right)[-p]$$

with the standard Čech differential. For any compact $K \subseteq X$, define $\check{C}_K^\bullet(X; \mathcal{F})$ (Čech cochains with supports in K) via (A.9.1) except replacing every instance of $\mathcal{F}(U)$ with $\ker[\mathcal{F}(U) \rightarrow \mathcal{F}(U \setminus K)]$. We let $\check{C}_c^\bullet(X; \mathcal{F}) := \varinjlim_{K \subseteq X} \check{C}_K^\bullet(X; \mathcal{F})$ (Čech cochains with compact supports).

It is clear that the homologies of $\check{C}^\bullet(X; \mathcal{F})$, $\check{C}_K^\bullet(X; \mathcal{F})$ and $\check{C}_c^\bullet(X; \mathcal{F})$ are $\check{H}^\bullet(X; \mathcal{F})$, $\check{H}_K^\bullet(X; \mathcal{F})$ and $\check{H}_c^\bullet(X; \mathcal{F})$, respectively.

Definition A.9.3 (\check{C}^\bullet of complexes of \mathcal{K} -presheaves; compare Definition A.4.10)
 Let \mathcal{F}^\bullet be a complex of \mathcal{K} -presheaves on a compact space X . We define

$$(A.9.2) \quad \check{C}^\bullet(X; \mathcal{F}^\bullet) := \varinjlim_{\substack{\{U_x \subseteq X\}_{x \in X} \\ \text{rigid open cover}}} \bigoplus_{p \geq 0} \prod_{\substack{S \subseteq X \\ |S|=p+1}} \mathcal{F}^{\bullet-p} \left(\bigcap_{x \in S} \overline{U_x} \right)$$

with the standard Čech differential (plus the internal differential of \mathcal{F}^\bullet). We also define $\check{C}^\bullet(X; \mathcal{F})$ for any \mathcal{K} -presheaf \mathcal{F} by viewing it as a complex concentrated in degree zero.

The homology of $\check{C}^\bullet(X; \mathcal{F}^\bullet)$ is $\check{H}^\bullet(X; \mathcal{F}^\bullet)$, as can be seen by applying the arguments from the proof of Lemma A.4.11. Note that for a sheaf \mathcal{F} , the complexes (A.9.1) and (A.9.2) are canonically isomorphic.

Definition A.9.4 (Pullback on \check{C}^\bullet and \check{C}_c^\bullet ; compare Definitions A.4.6 and A.4.21)
 Let $f: X \rightarrow Y$ be a map of spaces. A rigid open cover $\{U_y\}_{y \in X}$ pulls back to a rigid open cover $\{f^{-1}(U_{f(x)})\}_{x \in X}$, and this gives natural maps:

- (i) $f^*: \check{C}^\bullet(Y; f_*\mathcal{F}) \rightarrow \check{C}^\bullet(X; \mathcal{F})$ for $\mathcal{F} \in \text{Prshv } X$.
- (ii) $f^*: \check{C}_c^\bullet(Y; f_*\mathcal{F}) \rightarrow \check{C}_c^\bullet(X; \mathcal{F})$ for $\mathcal{F} \in \text{Shv } X$ (if f is proper).
- (iii) $f^*: \check{C}^\bullet(Y; f_*\mathcal{F}^\bullet) \rightarrow \check{C}^\bullet(X; \mathcal{F}^\bullet)$ for $\mathcal{F}^i \in \text{Prshv}_{\mathcal{K}} X$.

The actions of these maps on homology coincide with the maps f^* defined earlier.

A.9.2 Derived inverse limits

Definition A.9.5 (Derived inverse limit $R\varprojlim$) Let $\{C_\lambda^\bullet\}_{\lambda \in \Lambda}$ be an inverse system of complexes. We define

$$(A.9.3) \quad R\varprojlim_{\lambda \in \Lambda} C_\lambda^\bullet := \prod_{q \geq 0} \prod_{\lambda_0 \leq \dots \leq \lambda_q} C_{\lambda_0}^{\bullet-q}$$

with differential obtained by viewing this as cochains on the nerve of Λ with a particular coefficient system. See Mardešić [59, Section 17] for more details on and basic properties of $R\varprojlim$.

Definition A.9.6 (Derived functors \varprojlim^i) Let $\{A_\lambda\}_{\lambda \in \Lambda}$ be an inverse system of abelian groups. We define

$$(A.9.4) \quad \varprojlim_{\lambda \in \Lambda}^i A_\lambda := H^i \left[R \varprojlim_{\lambda \in \Lambda} A_\lambda \right]$$

(viewing A_λ as an inverse system of complexes concentrated in degree zero). The inverse limit functor \varprojlim from inverse systems of abelian groups indexed by Λ to abelian groups is left exact, and \varprojlim^i are its right derived functors (see Mardešić [59, Corollary 11.47]). See Mardešić [59, Sections 11–15] for more details on and basic properties of \varprojlim^i .

Lemma A.9.7 (Cofinality for \varprojlim^i and $R \varprojlim$) Let $f: \Lambda' \rightarrow \Lambda$ be weakly increasing ($\lambda_1 \leq \lambda_2 \implies f(\lambda_1) \leq f(\lambda_2)$) and cofinal ($f(\Lambda') \subseteq \Lambda$ cofinal). Then the following natural map is an isomorphism:

$$(A.9.5) \quad \varprojlim_{\lambda \in \Lambda}^i A_\lambda \xrightarrow{\sim} \varprojlim_{\lambda' \in \Lambda'}^i A_{f(\lambda')}.$$

More generally, the following natural map is a quasi-isomorphism:

$$(A.9.6) \quad R \varprojlim_{\lambda \in \Lambda} C_\lambda^\bullet \xrightarrow{\sim} R \varprojlim_{\lambda' \in \Lambda'} C_{f(\lambda')}^\bullet.$$

Proof For \varprojlim^i , see [59, page 291, Theorem 14.9], and for $R \varprojlim$, see [59, page 349, Corollary 17.18]. □

Lemma A.9.8 If $C_\lambda^\bullet \xrightarrow{\sim} D_\lambda^\bullet$ is a morphism of inverse systems which is a level-wise quasi-isomorphism (ie for every λ), then the natural map

$$(A.9.7) \quad R \varprojlim_{\lambda \in \Lambda} C_\lambda^\bullet \xrightarrow{\sim} R \varprojlim_{\lambda \in \Lambda} D_\lambda^\bullet$$

is a quasi-isomorphism.

Proof See [59, page 348, Theorem 17.16]. □

A.9.3 Steenrod chains

Definition A.9.9 (Due to Chogoshvili [13]) A partition $X = \bigsqcup_{i=1}^n E_i$ shall mean an *unordered* partition into finitely many disjoint nonempty subsets $E_1, \dots, E_n \subseteq X$; there is an associated finite closed covering $X = \bigcup_{i=1}^n \overline{E}_i$, whose nerve is denoted $N(X; \{\overline{E}_i\}_{i=1}^n)$ (the simplicial complex with vertices $\{1, \dots, n\}$, where $S \subseteq \{1, \dots, n\}$ spans a simplex if and only if $\bigcap_{i \in S} \overline{E}_i \neq \emptyset$). A partition $\{F_j\}_{j=1}^m$ refines $\{E_i\}_{i=1}^n$ if and only if for all j there exists a (necessarily unique) i with $F_j \subseteq E_i$. For a refinement $\{E_i\}_{i=1}^n \rightarrow \{F_j\}_{j=1}^m$, there is an associated map of nerves $N(X; \{\overline{F}_j\}) \rightarrow N(X; \{\overline{E}_i\})$.

For doing Čech theory, Chogoshvili’s construction has exceptionally nice properties. The collection of partitions is a set, and there is at most one morphism between any pair of partitions. The poset of partitions is cofinite (a given partition refines only finitely many other partitions). Each nerve $N(X; \{\overline{E}_i\})$ is a finite simplicial complex, and the transition maps $N(X; \{\overline{F}_j\}) \rightarrow N(X; \{\overline{E}_i\})$ are all surjective. Also, any pair of partitions has a minimal common refinement. Given a map $f: X \rightarrow Y$, a partition $\{E_i\}$ of Y pulls back to a partition $\{f^{-1}(E_i)\}$ of X , and there is an associated map on nerves $N(X; \{f^{-1}(E_i)\}) \rightarrow N(Y, \{E_i\})$.

Lemma A.9.10 *Let X be compact. There is a natural map*

$$(A.9.8) \quad \varprojlim_{X = \bigsqcup_{i=1}^n E_i} N(X; \{\overline{E}_i\}) \rightarrow X$$

which induces an isomorphism on Čech cohomology.

Proof There is a natural correspondence $C \subseteq X \times N(X; \{\overline{E}_i\})$, whose fiber over $x \in X$ is the complete simplex on $\{i : x \in \overline{E}_i\}$. Now it is not hard to check that the inverse limit of these correspondences maps bijectively (and thus homeomorphically) to $\varprojlim_{X = \bigsqcup_{i=1}^n E_i} N(X; \{\overline{E}_i\})$, thus giving rise to the desired map (A.9.8).

Now the inverse image of $p \in X$ under (A.9.8) is an inverse limit of complete simplices with surjective transition maps (“the complete simplex on a profinite set”), and this has the Čech cohomology of a point. Indeed, for any inverse system of compact spaces $\{X_\alpha\}$, the natural map $\varprojlim_\alpha \check{H}^\bullet(X_\alpha) \rightarrow \check{H}^\bullet(\varprojlim_\alpha X_\alpha)$ is an isomorphism (since open covers of $\varprojlim_\alpha X_\alpha$ pulled back from some X_α are cofinal among all open covers). Thus it follows from the Leray spectral sequence that (A.9.8) induces an isomorphism on Čech cohomology. □

Definition A.9.11 (Steenrod chains and homology) *Let X be compact. We define*

$$(A.9.9) \quad \overline{C}_\bullet(X) := R\varprojlim_{X = \bigsqcup_{i=1}^n E_i} \bigoplus_{p \geq 0} \bigoplus_{\substack{S \subseteq \{1, \dots, n\} \\ |S| = p+1 \\ \bigcap_{i \in S} \overline{E}_i \neq \emptyset}} \mathbb{Z}[-p].$$

For a map $f: X \rightarrow Y$, there is an induced pushforward map $f_*: \overline{C}_\bullet(X) \rightarrow \overline{C}_\bullet(Y)$. We denote by $\overline{H}_\bullet(X; \mathcal{F})$ the homology of $\overline{C}_\bullet(X; \mathcal{F})$, and we define relative homology $\overline{H}_\bullet(X, Y; \mathcal{F})$ as the homology of the relevant mapping cone.

More generally, let \mathcal{F} be a locally constant sheaf of abelian groups on X . Fix a partition $X = \bigsqcup_{i=1}^n E_i^0$ and trivializations of $\mathcal{F}|_{E_i^0}$. This gives rise to a local system \mathcal{F} on $N(\{\overline{E}_i\}_{i=1}^n)$ for any partition $\{E_i\}$ refining $\{E_i^0\}$. We now define $\overline{C}_\bullet(X; \mathcal{F})$ as

in (A.9.9), restricted to partitions refining $\{E_i^0\}$ and using the coefficient system on the nerve induced by \mathcal{F} and the fixed trivializations of $\mathcal{F}|_{\overline{E}_i^0}$. Since any two choices of trivializations of $\mathcal{F}|_{\overline{E}_i}$ become isomorphic (in the sense that their sets of “constant sections” coincide) after pulling back to some common refinement, it follows that $\overline{C}_\bullet(X; \mathcal{F})$ is well-defined up to essentially unique quasi-isomorphism (by Lemma A.9.7). For a map $f: X \rightarrow Y$ and $\mathcal{F} \rightarrow f^*\mathcal{G}$, there is a pushforward map $\overline{C}_\bullet(X; \mathcal{F}) \rightarrow \overline{C}_\bullet(Y; \mathcal{G})$.

Although not logically necessary for our purposes, we now argue that Steenrod homology as given in Definition A.9.11 coincides with the definition from Mardešić [59] (at least for constant coefficient systems). First of all, note that (a special case of) Mardešić’s definition of strong homology is that if a compact space X is an inverse limit of compact polyhedra X_α over a cofinite index set, then

$$(A.9.10) \quad \overline{H}_\bullet(X) = H_\bullet R \varprojlim_\alpha C_\bullet(X_\alpha)$$

(see [59, page 379, Section 19.1], and note that such an inverse system is a “cofinite polyhedral resolution” of its inverse limit [59, page 103, Section 6]). In particular, $\overline{C}_\bullet(X)$ as defined in (A.9.9) computes the Steenrod homology of

$$\varprojlim_{X = \bigsqcup_{i=1}^n E_i} N(X; \{\overline{E}_i\})$$

(using the morphism of inverse systems

$$C_\bullet^{\text{simp}}(N(X; \{\overline{E}_i\})) \rightarrow C_\bullet^{\text{sing}}(N(X; \{\overline{E}_i\}))$$

given by barycentric subdivision, which is a level-wise quasi-isomorphism and thus induces a quasi-isomorphism on derived inverse limits by Lemma A.9.8). Finally, note that the map (A.9.8) induces an isomorphism on strong homology by Lemma A.9.10 and [59, page 446, Theorem 21.15].

Lemma A.9.12 (Steenrod homology is the derived dual of Čech cohomology) *Let \mathcal{F} be a locally constant sheaf whose stalks are finitely generated free R -modules, and let M be an R -module. Then there is a natural isomorphism*

$$(A.9.11) \quad \overline{C}_\bullet(X; \text{Hom}(\mathcal{F}, M)) = R \text{Hom}_{D(R)}(\check{C}^\bullet(X; \mathcal{F}), M)$$

in the derived category $D(R)$. The same holds for relative (co)chains of a closed subspace $Y \subseteq X$, and the isomorphisms are compatible with the relevant exact triangles relating Y , X and (X, Y) .

This result is very similar to Milnor [64, page 92, Definition, Lemma 5] and Mardešić [59, page 446, Theorem 21.15]. Note also that morally, the above result should be thought of as following trivially from some derived universal property of the form

$$R \varprojlim R \operatorname{Hom}(A_\alpha^\bullet, B^\bullet) = R \operatorname{Hom}(\varprojlim A_\alpha^\bullet, B^\bullet)$$

(hopefully this provides some motivation for the proof below).

Proof Let us relate $\check{C}^\bullet(X; \mathcal{F})$ as defined using rigid open covers to a version using partitions. Specifically, we consider the following quasi-isomorphisms:

$$(A.9.12) \quad \check{C}^\bullet(X; \mathcal{F}) = \varinjlim_{\substack{\{U_x \subseteq X\}_{x \in X} \\ \text{rigid open cover}}} \bigoplus_{p \geq 0} \prod_{\substack{S \subseteq X \\ |S|=p+1}} \mathcal{F}\left(\bigcap_{x \in S} U_x\right)[-p]$$

$$(A.9.13) \quad \varinjlim_{\substack{X = \bigsqcup_{i=1}^n E_i \\ \bar{E}_i \subseteq U_i}} \bigoplus_{p \geq 0} \prod_{\substack{S \subseteq \{1, \dots, n\} \\ |S|=p+1}} \mathcal{F}\left(\bigcap_{i \in S} U_i\right)[-p]$$

$$(A.9.14) \quad \varinjlim_{X = \bigsqcup_{i=1}^n E_i} \bigoplus_{p \geq 0} \prod_{\substack{S \subseteq \{1, \dots, n\} \\ |S|=p+1}} \mathcal{F}\left(\bigcap_{i \in S} \bar{E}_i\right)[-p]$$

$$(A.9.15) \quad \bigoplus_{q \geq 0} \bigoplus_{\{E_i^{(0)}\}_{i=1}^{n_0} < \dots < \{E_i^{(q)}\}_{i=1}^{n_q}} \bigoplus_{p \geq 0} \prod_{\substack{S \subseteq \{1, \dots, n_0\} \\ |S|=p+1}} \mathcal{F}\left(\bigcap_{i \in S} \overline{E_i^{(0)}}\right)[q-p]$$

$$(A.9.16) \quad \bigoplus_{q \geq 0} \bigoplus_{\{E_i^{(0)}\}_{i=1}^{n_0} \leq \dots \leq \{E_i^{(q)}\}_{i=1}^{n_q}} \bigoplus_{p \geq 0} \prod_{\substack{S \subseteq \{1, \dots, n_0\} \\ |S|=p+1}} \mathcal{F}\left(\bigcap_{i \in S} \overline{E_i^{(0)}}\right)[q-p]$$

The first map is induced by associating to a partition $\{E_i\}$ and $\bar{E}_i \subseteq U_i$ the rigid open cover assigning U_i to $x \in E_i$, and pulling back along the induced map on nerves; it can be seen to be a quasi-isomorphism by a filtration argument (basically the Leray spectral sequence). The second map is a quasi-isomorphism by Lemma A.1.7 and the definition of α^* . The third map is defined as the tautological map on $q = 0$ direct summands and zero for $q > 0$; it can be seen to be a quasi-isomorphism by expressing the domain (resp. codomain) as the direct limit over partitions $\{E_i\}$ of the corresponding complex (resp. direct limit) restricted to partitions refined by $\{E_i\}$ and observing that for any fixed $\{E_i\}$ the map is a quasi-isomorphism by Lemma A.7.3.

Finally, the last pair of maps consists of the natural inclusion (\downarrow) and the retraction annihilating any component with $\{E_i^{(j)}\}_{i=1}^{n_j} = \{E_i^{(j+1)}\}_{i=1}^{n_{j+1}}$ for some j (\uparrow); one can see that both are quasi-isomorphisms by filtering by $\{E_i\}$ (as before) and then by the number of distinct $\{E_i^{(j)}\}_{i=1}^{n_j}$.

For any partition $\{E_i\}$ of X and collection of trivializations of $\mathcal{F}|_{\overline{E_i}}$, there is a sub- \mathcal{K} -presheaf $\mathcal{F}^{\text{pre}} \subseteq \mathcal{F}$ whose sections over K are those sections of \mathcal{F} whose restrictions to $K \cap \overline{E_i}$ are constant with respect to the specified trivializations. The complex (A.9.14) calculates Čech cohomology (A.4.6) by a cofinality argument, and the inclusion $\mathcal{F}^{\text{pre}} \rightarrow \mathcal{F}$ induces an isomorphism on stalks and thus on Čech cohomology by Proposition A.4.18. The third map above remains a quasi-isomorphism for \mathcal{F}^{pre} (for the same reason), and furthermore we may restrict (A.9.16) to partitions refining the fixed $\{E_i\}$ by Lemma A.9.7.

The conclusion of the above discussion is that we have thus constructed a canonical (and functorial) quasi-isomorphism between $\check{C}^\bullet(X; \mathcal{F})$ and

$$(A.9.17) \quad \bigoplus_{q \geq 0} \bigoplus_{\{E_i\} \leq \{E_i^{(0)}\}_{i=1}^{n_0} \leq \dots \leq \{E_i^{(q)}\}_{i=1}^{n_q}} \bigoplus_{p \geq 0} \prod_{\substack{S \subseteq \{1, \dots, n_0\} \\ |S|=p+1}} \mathcal{F}^{\text{pre}} \left(\bigcap_{i \in S} \overline{E_i^{(0)}} \right) [q - p]$$

(for any fixed partition $\{E_i\}$ and trivializations of $\mathcal{F}|_{\overline{E_i}}$ giving rise to $\mathcal{F}^{\text{pre}} \subseteq \mathcal{F}$). Note that $\overline{C}_\bullet(X, \text{Hom}(\mathcal{F}, M))$ is precisely $\text{Hom}((A.9.17), M)$. Thus it suffices to show that this particular Hom is in fact the $R\text{Hom}$.

Recall that $R\text{Hom}$ in $D(R)$ may be computed using a \mathbf{K} -projective resolution of the first argument (see Spaltenstein [74]), where a complex P^\bullet is called \mathbf{K} -projective if and only if $\text{Hom}^\bullet(P^\bullet, M^\bullet)$ is acyclic for every acyclic complex M^\bullet (note that a bounded-above complex of projective modules is \mathbf{K} -projective). Thus it suffices to show that (A.9.17) is \mathbf{K} -projective.

We consider subsets \wp of the poset of partitions of X with the property that any partition refined by a partition in \wp also lies in \wp . Since the poset of partitions is cofinite, we may use Zorn’s lemma to choose a directed system $\{\wp_i\}_{i \in I}$ of such collections, indexed by a well-ordered set I , with the following properties:

- (i) $\wp_0 = \emptyset$, where $0 \in I$ is the least element.
- (ii) $|\wp_i \setminus \wp_{i-1}| = 1$ if $i \in I$ has a predecessor $i - 1$.
- (iii) $\wp_i = \bigcup_{i' < i} \wp_{i'}$ if $i \in I$ has no predecessor.

Now for $i \in I$, let $(A.9.17)_i \subseteq (A.9.17)$ denote the subcomplex generated by restricting to partitions in \wp_i . It follows that:

- (i) $(A.9.17)_0 = 0$, where $0 \in I$ is the least element.
- (ii) $(A.9.17)_{i-1} \rightarrow (A.9.17)_i$ is injective and component-wise split with K -projective cokernel if $i \in I$ has a predecessor $i - 1$ (this holds since the cokernel is a bounded-above complex of free modules).
- (iii) $\varinjlim_{i' < i} (A.9.17)_{i'} \rightarrow (A.9.17)_i$ is an isomorphism if $i \in I$ has no predecessor.

It follows that $(A.9.17) = \varinjlim_i (A.9.17)_i$ is K -projective by Spaltenstein [74, page 131, Corollary 2.8].

Since the construction of the isomorphism (A.9.11) was given by a functorial chain-level construction, it can be checked that it is natural, applies in the relative case, and is compatible with exact triangles. \square

Appendix B: Gluing for implicit atlases on Gromov–Witten moduli spaces

In this appendix, we provide the gluing analysis used in Section 9 to verify that the implicit atlases constructed there satisfy the openness and submersion axioms (specifically, we prove Proposition 9.2.6).

The gluing theorem we prove here is of a very standard sort which has been treated (in different related settings) many times over in the literature, and so we make no claim of originality in this appendix. The methods used here are based on our partial understanding of the treatments of gluing in Abouzaid [1] and McDuff and Salamon [62], as well as conversations with Abouzaid, Ekholm, Hofer and Mazzeo (we also thank the anonymous referee for their comments). We understand that there is work in progress of McDuff and Wehrheim proving a similar result (keeping track of more smoothness) in their (very similar) setting of Kuranishi atlases [63; 60]. We also refer the reader to Ekholm and Smith [19], Fukaya, Oh, Ohta and Ono [31] and Hofer, Wysocki and Zehnder [46] for related gluing results.

The essential content of our result is that, for a certain moduli space of holomorphic curves $\overline{\mathcal{M}}(X)$, the regular locus $\overline{\mathcal{M}}(X)^{\text{reg}} \subseteq \overline{\mathcal{M}}(X)$ (ie the locus where a certain linearized $\bar{\partial}$ -operator is surjective) is a topological manifold. We prove this by constructing local manifold charts covering $\overline{\mathcal{M}}(X)^{\text{reg}}$ (since being a topological manifold is a property rather than extra structure, we need not address any question of compatibility between different local charts or of their compatibility with any auxiliary group action).

B.1 Setup and main result

We use new notation, unrelated to (and simpler than) that from Section 9.

Fix a smooth almost complex manifold

$$(B.1.1) \quad (X, J).$$

Fix codimension-two submanifolds with boundary

$$(B.1.2) \quad D, D_1, \dots, D_r \subseteq X$$

(let $D^\circ := D \setminus \partial D$, and similarly for D_i). Fix a smooth manifold $\bar{\mathcal{M}}$ equipped with a smooth étale map $\bar{\mathcal{M}} \rightarrow \bar{\mathcal{M}}_{g,n+l}$ ⁴⁷ (in the orbifold sense). Denote by

$$\bar{\mathcal{C}}_{g,n+l} \rightarrow \bar{\mathcal{M}}_{g,n+l}$$

the universal family, and define a family of curves $\bar{\mathcal{C}} \rightarrow \bar{\mathcal{M}}$ via the following pullback square:

$$(B.1.3) \quad \begin{array}{ccc} \bar{\mathcal{C}} & \longrightarrow & \bar{\mathcal{C}}_{g,n+l} \\ \downarrow & & \downarrow \\ \bar{\mathcal{M}} & \longrightarrow & \bar{\mathcal{M}}_{g,n+l} \end{array}$$

Let $\bar{\mathcal{C}}^\circ \subseteq \bar{\mathcal{C}}$ denote the open subset where the fiber is smooth. Fix a finite-dimensional vector space E and a linear map

$$(B.1.4) \quad \lambda: E \rightarrow C^\infty(\overbrace{\bar{\mathcal{C}}^\circ \times \dots \times \bar{\mathcal{C}}^\circ}^{r \text{ times}} \times \bar{\mathcal{C}}^\circ \times X, \Omega_{\bar{\mathcal{C}}^\circ/\bar{\mathcal{M}}}^{0,1} \otimes_{\mathbb{C}} TX).$$

Here $\Omega_{\bar{\mathcal{C}}^\circ/\bar{\mathcal{M}}}^{0,1}$ is the $(0, 1)$ -part of the complexified vertical cotangent bundle

$$T_{\bar{\mathcal{C}}^\circ/\bar{\mathcal{M}}}^* \otimes_{\mathbb{R}} \mathbb{C}$$

of (the last factor of) $\bar{\mathcal{C}}^\circ \rightarrow \bar{\mathcal{M}}$, and C^∞ means smooth sections. We require that λ be zero in a neighborhood of the nodes $\bar{\mathcal{C}} \setminus \bar{\mathcal{C}}^\circ$ of the last $\bar{\mathcal{C}}^\circ$ factor.

For a nodal curve C , we denote by \tilde{C} its normalization, ie the unique smooth compact Riemann surface equipped with a map $\tilde{C} \rightarrow C$ which identifies points in pairs to form

⁴⁷We give $\bar{\mathcal{M}}_{g,n+l}$ the standard smooth structure from its structure as a complex analytic orbifold. Every point in a smooth orbifold M is in the image of a smooth étale map $M' \rightarrow M$. Indeed, M is covered by open sets each diffeomorphic to \mathbb{R}^n/G for some finite group $G \rightarrow \text{GL}_n(\mathbb{R})$, and the map $\mathbb{R}^n \rightarrow \mathbb{R}^n/G \hookrightarrow M$ is smooth étale.

the nodes of C . A function on C being smooth means (by definition) that its pullback to \tilde{C} is smooth.

Fix a homology class $\beta \in H_2(X; \mathbb{Z})$. We are interested in the following moduli space: (B.1.5)

$$\overline{\mathcal{M}}(X) := \left\{ \begin{array}{l} s \in \overline{\mathcal{M}} \\ u: C_s \rightarrow X \\ x_i \in C_s, \ 1 \leq i \leq r \\ e \in E \end{array} \left| \begin{array}{l} u \text{ smooth and } u_*[C_s] = \beta \\ u(p_{n+i}) \in D^\circ, \ 1 \leq i \leq \ell \\ u(x_i) \in D_i^\circ \text{ and } u \pitchfork D_i \text{ at } x_i, \ 1 \leq i \leq r \\ \bar{\partial}u + \lambda(e)(x_1, \dots, x_r, \cdot, u(\cdot)) = 0 \end{array} \right. \right\}.$$

We let C_s denote the fiber of $\bar{c} \rightarrow \overline{\mathcal{M}}$ over s , with marked points $p_1, \dots, p_{n+\ell} \in C_s$. Here $u \pitchfork D_i$ at x_i means (by definition) that x_i is not a node of C_s and the map

$$T_{x_i}C_s \oplus T_{u(x_i)}D_i \xrightarrow{du \oplus \text{id}} T_{u(x_i)}X$$

is surjective. Now $\lambda(e)(x_1, \dots, x_r, \cdot, u(\cdot))$ is a smooth section of $\Omega_{\tilde{C}_s}^{0,1} \otimes_{\mathbb{C}} u^*TX$ over C_s supported away from the nodes; hence the last equation makes sense. We give $\overline{\mathcal{M}}(X)$ the topology of uniform convergence (ie using the Hausdorff distance between graphs $\subseteq \bar{c} \times X$ to compare maps u).

We spend this appendix studying $\overline{\mathcal{M}}(X)$, though the only reason to care about $\overline{\mathcal{M}}(X)$ itself is as an intermediate tool for proving the desired result for $\overline{\mathcal{M}}_{g,n}^\beta(X)_I$. Note that $\overline{\mathcal{M}}(X)$ is not necessarily compact, however this will be irrelevant since we are only interested in its local properties.

Fix a subspace $E' \subseteq E$.

Let us now define the “ E' -regular locus” $\overline{\mathcal{M}}(X)^{\text{reg}} \subseteq \overline{\mathcal{M}}(X)$ (which for simplicity we will just call the “regular locus”). Fix $(s_0, u_0, \{x_i^0\}, e_0) \in \overline{\mathcal{M}}(X)$; we will describe when $(s_0, u_0, \{x_i^0\}, e_0) \in \overline{\mathcal{M}}(X)^{\text{reg}}$. We consider the smooth Banach manifold

$$(B.1.6) \quad \mathcal{B} := \{(u, e) \in W^{k,p}(C_{s_0}, X) \times E \mid u(p_{n+i}) \in D^\circ \ (1 \leq i \leq \ell)\}.$$

We consider the smooth Banach bundle \mathcal{E} over \mathcal{B} whose fiber over a map $u: C_{s_0} \rightarrow X$ is

$$W^{k-1,p}(\tilde{C}_{s_0}, \Omega_{\tilde{C}_{s_0}}^{0,1} \otimes_{\mathbb{C}} u^*TX) \oplus E/E'.$$

Now suppose k is large; then there are unique continuous functions

$$(B.1.7) \quad x_i: \mathcal{B} \rightarrow C_{s_0} \quad (1 \leq i \leq r)$$

defined in a neighborhood of $(u_0, e_0) \in \mathcal{B}$, which coincide with x_i^0 at (u_0, e_0) , and for which $u(x_i(u)) \in D_i^\circ$. Moreover, (B.1.7) are “highly differentiable”, by which we mean that for all $\ell < \infty$, the function (B.1.7) is of class C^ℓ provided $k \geq k_0(\ell)$. In the present situation, we can of course be more precise: x_i is C^ℓ as a function of $u \in C^\ell$

(by the implicit function theorem), so (B.1.7) is C^ℓ as long as $W^{k,p} \hookrightarrow C^\ell$ (which holds if and only if $(k - \ell)p > 2$). It follows that

$$(B.1.8) \quad (u, e) \mapsto [\bar{\partial}u + \lambda(e)(x_1, \dots, x_r, \cdot, u(\cdot))] \oplus e$$

is a highly differentiable section of \mathcal{E} over \mathcal{B} (the only nonsmoothness comes from how the x_i depend on u). Assume $1 < p < \infty$. We say $(s_0, u_0, \{x_i^0\}, e_0) \in \bar{\mathcal{M}}(X)^{\text{reg}}$ if and only if (B.1.8) is transverse to the zero section at (u_0, e_0) . It is an easy exercise using elliptic regularity to see that this notion is independent of (k, p) (as long as k is sufficiently large so that the condition makes sense).

For a topological manifold M , let \mathfrak{o}_M denote its orientation sheaf (whose fiber at a point $p \in M$ is canonically $H_{\dim M}(M, M \setminus p; \mathbb{Z})$). For a vector space E , let \mathfrak{o}_E denote its orientation module (canonically $H_{\dim E}(E, E \setminus 0; \mathbb{Z})$). The main result of this appendix is:

Theorem B.1.1 *In the above setup (from the beginning of Section B.1 until here), we have:*

- (i) $\bar{\mathcal{M}}(X)^{\text{reg}} \subseteq \bar{\mathcal{M}}(X)$ is open.
- (ii) $\bar{\mathcal{M}}(X)^{\text{reg}}$ is a topological manifold of dimension

$$\dim \bar{\mathcal{M}}_{g,n+\ell} + (1 - g) \dim X + 2\langle c_1(X), \beta \rangle + \dim E - 2\ell.$$

- (iii) The projection $\bar{\mathcal{M}}(X)^{\text{reg}} \rightarrow E/E'$ is a topological submersion, ie locally modeled on a projection $\mathbb{R}^n \times \mathbb{R}^m \rightarrow \mathbb{R}^n$.
- (iv) The orientation sheaf $\mathfrak{o}_{\bar{\mathcal{M}}(X)^{\text{reg}}}$ is canonically isomorphic to

$$\bigotimes_{i=1}^{\ell} \mathfrak{o}_{u^{(p_n+i)^*} N_{D/X}} \otimes \mathfrak{o}_E.$$

Let us now explain how Theorem B.1.1 implies Proposition 9.2.6. Proposition 9.2.6 is a local statement, so let us prove it in a neighborhood of a specific point

$$(C, u, \{\phi_\alpha\}_{\alpha \in I}, \{e_\alpha\}_{\alpha \in I}) \in \bar{\mathcal{M}}_{g,n}^\beta(X)_I^{\text{reg}}.$$

We will construct data $X, J, D, D_1, \dots, D_r, \bar{\mathcal{M}}, E, \lambda, E'$ as in the above setup and a point $(s, u, \{x_i\}, e) \in \bar{\mathcal{M}}(X)$. It will be clear from the construction that there is a natural homeomorphism between a small neighborhood of this point in $\bar{\mathcal{M}}(X)$ and a small neighborhood of the given point in $\bar{\mathcal{M}}_{g,n}^\beta(X)_I$, and using this we will infer Proposition 9.2.6 from Theorem B.1.1.

We claim that every unstable irreducible component of C has a point where du is injective. If $I = \emptyset$, this follows from the domain stabilization step Lemma 9.2.9, and if $I \neq \emptyset$ then, picking any $\alpha \in I$, this follows from the fact that $u \pitchfork D_\alpha$ and adding the intersections as marked points makes C stable. It follows from the claim that we may pick $D \subseteq X$ with $u \pitchfork D$ such that adding these intersections as extra marked points makes C stable, and none of these points are nodes or marked points. We take ℓ to be the *minimum* number of points in $u^{-1}(D)$ necessary to stabilize C , and we fix an ordered ℓ -tuple of such points, adding them as new marked points $p_{n+1}, \dots, p_{n+\ell} \in C$, so now C is a curve of type $(g, n + \ell)$. Now we let $r = \sum_{\alpha \in I} r_\alpha$, we let $D_1, \dots, D_r \subseteq X$ consist of r_α copies of D_α (union over $\alpha \in I$), and we let the x_i be the intersection points $u^{-1}(D_\alpha)$ (union over $\alpha \in I$). For the purposes of the present argument, let us reindex $\{x_i\}_{1 \leq i \leq r}$ as $\{x_i^\alpha\}_{\alpha \in I, 1 \leq i \leq r_\alpha}$. Now since C is a curve of type $(g, n + \ell)$, it corresponds to a point in $\overline{\mathcal{M}}_{g, n+\ell}$. Choose an étale map $\overline{\mathcal{M}} \rightarrow \overline{\mathcal{M}}_{g, n+\ell}$ covering this point, and choose a point $s \in \overline{\mathcal{M}}$ and an isomorphism $\iota: C \rightarrow C_s$. Now for every $\alpha \in I$, there is a unique map

$$(B.1.9) \quad \tilde{\phi}_\alpha: \overbrace{\overline{\mathcal{C}}^\circ \times \dots \times \overline{\mathcal{C}}^\circ}^{r_\alpha \text{ times}} \times \overline{\mathcal{C}}^\circ \rightarrow \overline{\mathcal{C}}_\alpha$$

defined in a neighborhood of $\{x_1^\alpha\} \times \dots \times \{x_{r_\alpha}^\alpha\} \times C_s$ so that

$$\tilde{\phi}_\alpha(x_1^\alpha, \dots, x_{r_\alpha}^\alpha, \iota(\cdot)) = \phi_\alpha(\cdot)$$

and $\tilde{\phi}_\alpha(y_1, \dots, y_{r_\alpha}, \cdot)$ classifies the curve in the last factor after forgetting the last ℓ marked points and adding y_1, \dots, y_{r_α} as marked points (this follows from the fact that $\overline{\mathcal{C}}_\alpha \rightarrow \overline{\mathcal{M}}_{g, n+r_\alpha}/S_{r_\alpha}$ is étale). We let $E := E_I = \bigoplus_{\alpha \in I} E_\alpha$, and we define

$$\lambda(\bigoplus_{\alpha \in I} e_\alpha)(\{y_i^\alpha\}_{\alpha \in I, 1 \leq i \leq r_\alpha}, \cdot, \cdot) = \sum_{\alpha \in I} \lambda_\alpha(e_\alpha)(\tilde{\phi}_\alpha(y_1^\alpha, \dots, y_{r_\alpha}^\alpha, \cdot), \cdot)$$

in a neighborhood of $\{x_i^\alpha\}_{\alpha \in I, 1 \leq i \leq r_\alpha} \times C_s \times X$ (and then we simply cut it off to be zero elsewhere). Finally, we observe that with this definition, a neighborhood of

$$(s, u \circ \iota^{-1}, \{x_i^\alpha\}_{\alpha \in I, 1 \leq i \leq r_\alpha}, \bigoplus_{\alpha \in I} e_\alpha) \in \overline{\mathcal{M}}(X)$$

coincides with a neighborhood of

$$(C, u, \{\phi_\alpha\}_{\alpha \in I}, \{e_\alpha\}_{\alpha \in I}) \in \overline{\mathcal{M}}_{g, n}^\beta(X)_I$$

(note that the point

$$(C, u, \{\phi_\alpha\}_{\alpha \in I}, \{e_\alpha\}_{\alpha \in I}) \in \overline{\mathcal{M}}_{g, n}^\beta(X)_I^{\text{reg}}$$

has trivial automorphism group by definition of $\overline{\mathcal{M}}_{g, n}^\beta(X)_I^{\text{reg}}$).

For $E' = E$, under this identification of a small open set in $\overline{\mathcal{M}}_{g,n}^\beta(X)_I$ and a small open set in $\overline{\mathcal{M}}(X)$, we have $\overline{\mathcal{M}}(X)^{\text{reg}} \subseteq \overline{\mathcal{M}}_{g,n}^\beta(X)_I^{\text{reg}}$ (for this it is important that we chose ℓ as the *minimum* number of points necessary to stabilize C). It follows that parts (i) and (iv) of Theorem B.1.1 imply parts (i) and (iii) of Proposition 9.2.6 (note that over this small open set, we have an identification

$$\mathfrak{o}_{u(p_{n+i})^* N_{D/X}} = \mathfrak{o}_{T_{p_{n+i}} C_S}$$

and that the latter is canonically trivial using the complex structure on C_S). Taking the above construction with J in place of I and setting $E' := E_I \subseteq E_J = E$, we get Proposition 9.2.6(ii) from Theorem B.1.1(ii)–(iii) as well.

B.2 Local model for resolution of a node

The rest of this appendix is devoted to the proof of Theorem B.1.1. We now fix $(s_0, u_0, \{x_i^0\}, e_0) \in \overline{\mathcal{M}}(X)^{\text{reg}}$, and we prove the desired statements (i)–(iv) in a neighborhood of this point.

Our first task is to fix a nice local coordinate system on $\overline{\mathcal{M}}$ near s_0 . Let d be the number of nodes of $C_0 := C_{s_0}$.

On each side of each node of C_0 , fix a “cylindrical end”, that is, a map

$$(B.2.1) \quad [0, \infty) \times S^1 \rightarrow C_0$$

which is a biholomorphism onto some small neighborhood $D^2 \setminus 0$ of the node. We use coordinates $(s, t) \in [0, \infty) \times S^1$, which is given the standard holomorphic structure $z = e^{s+it}$.

Let $\overline{\mathcal{M}}^d \subseteq \overline{\mathcal{M}}$ denote the locus of curves with exactly d nodes. Pick a smooth family of smooth almost complex structures j_y on C_0 parametrized by $y \in \mathbb{R}^{\dim \overline{\mathcal{M}}^d}$, where j_0 is the given almost complex structure on C_0 , which is constant over the cylindrical ends (B.2.1), and such that the induced map $\mathbb{R}^{\dim \overline{\mathcal{M}}^d} \rightarrow \overline{\mathcal{M}}^d$ is a diffeomorphism onto its image.

Consider the following procedure, which takes a “gluing parameter” $\alpha = e^{-6S+i\theta} \in \mathbb{C}$ and two copies of the standard end $[0, \infty) \times S^1 \sqcup [0, \infty) \times S^1$. We first truncate both ends, leaving just the subset $[0, 6S] \times S^1 \sqcup [0, 6S] \times S^1$. We then identify (s, t) in the first piece with $(6S - s, \theta - t) = (s', t')$ in the second piece. We call the resulting cylinder a “neck”.

Now given gluing parameters⁴⁸ $\alpha = (\alpha_1, \dots, \alpha_d) \in \mathbb{C}^d$, we may perform the gluing operation above on C_0 , using the chosen cylindrical ends (B.2.1). We call the resulting curve C_α , and it is equipped with cylindrical ends (corresponding to those $\alpha_i = 0$) and necks (corresponding to those $\alpha_i \neq 0$)

$$(B.2.2) \quad [0, \infty) \times S^1 \rightarrow C_\alpha,$$

$$(B.2.3) \quad [0, 6S_i] \times S^1 \rightarrow C_\alpha.$$

In each neck, we have coordinates $(s, t) \in [0, 6S_i] \times S^1$ and coordinates $(s', t') \in [0, 6S_i] \times S^1$, which satisfy $s + s' = 6S_i$ and $t + t' = \theta_i$. The curve C_α is also equipped with $n + \ell$ marked points $p_1, \dots, p_{n+\ell} \in C_\alpha$ coming from the given $p_1, \dots, p_{n+\ell} \in C_0$. Since j_y is constant over the cylindrical ends, it descends to give an almost complex structure on C_α .

Now since $\bar{\mathcal{M}} \rightarrow \bar{\mathcal{M}}_{g,n+\ell}$ is étale, there is an induced map

$$(B.2.4) \quad \mathbb{C}^d \times \mathbb{R}^{\dim \bar{\mathcal{M}}^d} \rightarrow \bar{\mathcal{M}},$$

$$(B.2.5) \quad (\alpha, y) \mapsto (C_\alpha, j_y, p_1, \dots, p_{n+\ell}),$$

which is a local diffeomorphism near zero. Using these local coordinates, we may alternatively define $\bar{\mathcal{M}}(X)$ as

$$(B.2.6) \quad \bar{\mathcal{M}}(X) := \left\{ \begin{array}{l} \alpha \in \mathbb{C}^d \\ y \in \mathbb{R}^{\dim \bar{\mathcal{M}}^d} \\ u: C_\alpha \rightarrow X \\ x_i \in C_\alpha, 1 \leq i \leq r \\ e \in E \end{array} \left| \begin{array}{l} u \text{ smooth and } u_*[C_\alpha] = \beta \\ u(p_{n+i}) \in D^\circ, 1 \leq i \leq \ell \\ u(x_i) \in D_i^\circ \text{ and } u \pitchfork D_i \text{ at } x_i, 1 \leq i \leq r \\ \bar{\partial}_y u + \lambda(e)(\alpha, y, x_1, \dots, x_r, \cdot, u(\cdot)) = 0 \end{array} \right. \right\},$$

and this coincides with the definition (B.1.5) in a neighborhood of $(s_0, u_0, \{x_i^0\}, e_0) = (0, 0, u_0, \{x_i^0\}, e_0)$. For the purposes of (B.2.6), λ denotes the function

$$\lambda: E \rightarrow C^\infty(\mathbb{C}^d \times \mathbb{R}^{\dim \bar{\mathcal{M}}^d} \times \overbrace{C_0 \times \dots \times C_0}^{r \text{ times}} \times C_0 \times X, \text{Hom}_{\mathbb{R}}(T\tilde{C}_0, TX))$$

⁴⁸In the present construction, and in many other constructions/definitions to come later, certain expressions, equalities, etc, only make sense or only hold if the gluing parameters $\alpha_i \in \mathbb{C}$ are sufficiently close to zero (ie $|\alpha_i| \leq \epsilon$ for some $\epsilon > 0$ depending only on data which we have previously fixed). We will often leave this assumption implicit, since we only care about what happens for α in a neighborhood of $0 \in \mathbb{C}^d$ anyway. The same goes for $y \in \mathbb{R}^{\dim \bar{\mathcal{M}}^d}$.

for which $\lambda(\cdot)(\cdot, y, \dots)$ lands in $\Omega_{\tilde{C}_0, j_y}^{0,1} \otimes_{\mathbb{C}} TX \subseteq \text{Hom}_{\mathbb{R}}(T\tilde{C}_0, TX)$, defined in terms of the old λ via

$$\begin{aligned} \lambda^{\text{new}}(e)(\alpha, y, x_1, \dots, x_r, p, x) \\ := \lambda^{\text{old}}(e)(x_1 \in (C_\alpha, j_y) \subseteq \bar{C}, \dots, x_r \in (C_\alpha, j_y) \subseteq \bar{C}, p \in (C_\alpha, j_y) \subseteq \bar{C}, x). \end{aligned}$$

We assume that the cylindrical ends were chosen disjoint from the support of λ^{old} , so we can make sense of λ^{new} as giving sections on C_α .

Now to prove the main result, it suffices to study the local structure of $\bar{M}(X)$ (defined as in (B.2.6)) near the basepoint

$$(B.2.7) \quad (0, 0, u_0: C_0 \rightarrow X, \{x_i^0\}, e_0),$$

which we are assuming lies in $\bar{M}(X)^{\text{reg}}$.

B.3 Pregluing

Let $\exp: TX \rightarrow X$ denote the exponential map of some Riemannian metric on X for which D is totally geodesic. Let ∇ denote any J -linear⁴⁹ connection on TX (equivalently, a connection for which $\nabla J = 0$). Let $\text{PT}_{x \rightarrow y}: T_x X \rightarrow T_y X$ denote parallel transport via ∇ along the shortest geodesic between x and y (we will only use this notation when it may be assumed that x and y are very close in X); note that $\text{PT}_{x \rightarrow y}$ is J -linear.

Fix a smooth function $\chi: \mathbb{R} \rightarrow [0, 1]$ satisfying

$$(B.3.1) \quad \chi(x) = \begin{cases} 0, & x \leq 0, \\ 1, & x \geq 1. \end{cases}$$

Definition B.3.1 (Flattening) For a gluing parameter $\alpha \in \mathbb{C}^d$, we define the “flattened” map $u_{0|\alpha}: C_0 \rightarrow X$ as follows. Away from the ends, $u_{0|\alpha}$ coincides with u_0 . Over an end $[0, \infty) \times S^1$, we define it as

$$(B.3.2) \quad u_{0|\alpha}(s, t) := \begin{cases} u_0(s, t), & s \leq S - 1, \\ \exp_{u_0(n)}[\chi(S - s) \cdot \exp_{u_0(n)}^{-1} u_0(s, t)], & S - 1 \leq s \leq S, \\ u_0(n), & S \leq s, \end{cases}$$

where $n \in C_0$ denotes the corresponding node.

⁴⁹Given any connection ∇^0 , the connection $\nabla_X Y := \frac{1}{2}(\nabla_X^0 Y - J(\nabla_X^0(JY)))$ is J -linear.

Definition B.3.2 (Pregluing) For $\alpha \in \mathbb{C}^d$, we define the “preglued” map $u_\alpha: C_\alpha \rightarrow X$ as follows. Away from the necks, u_α coincides with u_0 . Over a neck $[0, 6S] \times S^1$, we define it as

$$(B.3.3) \quad u_\alpha(s, t) := \begin{cases} u_0(s, t), & s \leq S - 1, \\ \exp_{u_0(n)}[\chi(S - s) \cdot \exp_{u_0(n)}^{-1} u_0(s, t)], & S - 1 \leq s \leq S, \\ u_0(n), & S \leq s \leq 5S, \\ \exp_{u_0(n)}[\chi(S - s') \cdot \exp_{u_0(n)}^{-1} u_0(s', t')], & 5S \leq s \leq 5S + 1, \\ u_0(s', t'), & 5S + 1 \leq s \end{cases}$$

(u_α should be thought of as the “descent” of $u_0|_\alpha$ from C_0 to C_α).

Definition B.3.3 (Pregluing vector fields) For $\xi \in C^\infty(C_0, u_0^*TX)$, we define

$$(B.3.4) \quad \xi_\alpha \in C^\infty(C_\alpha, u_\alpha^*TX)$$

as follows. Away from the necks, ξ_α coincides with ξ . Over a neck $[0, 6S] \times S^1$, we define it as

$$(B.3.5) \quad \xi_\alpha(s, t) := \begin{cases} \xi(s, t), & s \leq S - 1, \\ \text{PT}_{u_0(s,t) \rightarrow u_\alpha(s,t)} \xi(s, t), & S - 1 \leq s \leq S, \\ (1 - \chi(s - S)) \text{PT}_{u_0(s,t) \rightarrow u_\alpha(s,t)} \xi(s, t) + \chi(s - S) \xi(n), & S \leq s \leq S + 1, \\ \xi(n), & S + 1 \leq s \leq 5S - 1, \\ (1 - \chi(s' - S)) \text{PT}_{u_0(s',t') \rightarrow u_\alpha(s',t')} \xi(s', t') + \chi(s' - S) \xi(n), & 5S - 1 \leq s \leq 5S, \\ \text{PT}_{u_0(s',t') \rightarrow u_\alpha(s',t')} \xi(s', t'), & 5S \leq s \leq 5S + 1, \\ \xi(s', t'), & 5S + 1 \leq s. \end{cases}$$

B.4 Weighted Sobolev norms

Fix $k \in \mathbb{Z}_{\geq 6}$ and fix $\delta \in (0, 1)$. Fix a metric on TX for the purposes of defining Sobolev norms of sections; this could be the same as the metric used to define \exp , though there is no need for it to be.

We denote by $W^{k,2}(C_\alpha, u_\alpha^*TX)$ the space of sections of u_α^*TX over C_α whose pullback to \tilde{C}_α is $W^{k,2}$ (since $W^{k,2} \hookrightarrow C^0$ as $k \geq 6$, this is the same as $W^{k,2}$ sections over \tilde{C}_α which coincide on each pair of points identified in $\tilde{C}_\alpha \rightarrow C_\alpha$, where \tilde{C}_α denotes the normalization of C_α). Similarly we denote by

$$W^{k-1,2}(\tilde{C}_\alpha, \Omega_{\tilde{C}_\alpha, j_y}^{0,1} \otimes_{\mathbb{C}} u_\alpha^*TX)$$

the space of sections of class $W^{k-1,2}$. Note that both $W^{k,2}$ and $W^{k-1,2}$ refer to spaces of functions on the compact Riemann surface \tilde{C}_α . The specific choice of metric

on C_α used to define the norms $\|\cdot\|_{k,2}$ and $\|\cdot\|_{k-1,2}$ will be of little importance (they will matter only up to commensurability).

We now define certain weighted Sobolev spaces $W^{k,2,\delta}$ and $W^{k-1,2,\delta}$ on the possibly non-compact Riemann surface C_α minus the nodes. The specific choice of norms $\|\cdot\|_{k,2,\delta}$ and $\|\cdot\|_{k-1,2,\delta}$ (not just their commensurability classes) will be of great importance.

To make calculations in ends/necks easier, for each node $n \in C_0$ we fix the local trivialization of TX given by parallel transport

$$(B.4.1) \quad \text{PT}_{u_0(n) \rightarrow p}: T_{u_0(n)}X \rightarrow T_pX$$

over p contained in a small ball centered at $u_0(n)$. We assume that the ends (B.2.1) were chosen small enough so that u_0 maps each end into a compact subset of the small ball associated to the corresponding node.

Definition B.4.1 For $\alpha \in \mathbb{C}^d$, we define a weighted Sobolev space

$$(B.4.2) \quad W^{k,2,\delta}(C_0, u_{0|\alpha}^*TX)$$

using the following norm. We use the usual $(k, 2)$ norm away from the ends. Over an end $[0, \infty) \times S^1 \subseteq C_0$, we consider the local trivialization (B.4.1), in which a section of $u_{0|\alpha}^*TX$ simply becomes a function $\xi: [0, \infty) \times S^1 \rightarrow T_{u_0(n)}X$. In terms of this function, the contribution from the end to the norm squared is

$$(B.4.3) \quad |\xi(n)|^2 + \int_{[0,\infty) \times S^1} \left[|\xi(s, t) - \xi(n)|^2 + \sum_{j=1}^k |D^j \xi(s, t)|^2 \right] e^{2\delta s} ds dt.$$

The derivatives of ξ are measured with respect to the standard metric on $[0, \infty) \times S^1$.

By the Sobolev embedding of $W^{2,2}$ into C^0 in two dimensions, in any end we have a uniform bound on $|\xi(s, t) - \xi(n)|e^{\delta s}$ and $|D^j \xi(s, t)|e^{\delta s}$ ($1 \leq j \leq k - 2$) linear in $\|\xi\|_{k,2,\delta}$.

Definition B.4.2 For $\alpha \in \mathbb{C}^d$, we define a weighted Sobolev space

$$(B.4.4) \quad W^{k,2,\delta}(C_\alpha, u_\alpha^*TX)$$

using the following norm. We use the usual $(k, 2)$ norm away from the ends/necks. Over an end, the contribution to the norm squared is (B.4.3). Over a neck $[0, 6S] \times S^1 \subseteq C_\alpha$,

we again think of a section as a function $\xi: [0, 6S] \times S^1 \rightarrow T_{u_0(n)}X$, and the contribution to the norm squared is

$$\left| \int_{S^1} \xi(3S, t) dt \right|^2 + \int_{[0, 6S] \times S^1} \left[\left| \xi(s, t) - \int_{S^1} \xi(3S, t) dt \right|^2 + \sum_{j=1}^k |D^j \xi(s, t)|^2 \right] e^{2\delta \min(s, 6S-s)} ds dt.$$

Of course, for fixed $\alpha \in \mathbb{C}^d$, this norm is equivalent to not weighting the necks, though the two norms are not uniformly equivalent as $\alpha \rightarrow 0$. The weight in the necks is important since the key point is to establish certain estimates which are uniform as $\alpha \rightarrow 0$.

By the Sobolev embedding of $W^{2,2}$ into C^0 in two dimensions, in any neck we have a uniform bound on

$$\left| \xi(s, t) - \int_{S^1} \xi(3S, t) dt \right| e^{\delta \min(s, 6S-s)} \quad \text{and} \quad |D^j \xi(s, t)| e^{\delta \min(s, 6S-s)}$$

($1 \leq j \leq k - 2$) linear in $\|\xi\|_{k,2,\delta}$.

Definition B.4.3 For $\alpha \in \mathbb{C}^d$, we define a weighted Sobolev space

$$(B.4.5) \quad W^{k-1,2,\delta}(\tilde{C}_\alpha, \Omega_{\tilde{C}_\alpha, j_y}^{0,1} \otimes_{\mathbb{C}} u_\alpha^* TX)$$

as follows. We use the usual $(k - 1, 2)$ norm away from the ends/neck. Over an end or neck, we trivialize $(T\tilde{C}_\alpha, j_y)$ over \mathbb{C} by the basis vector $\partial/\partial s$, and we trivialize TX via (B.4.1), and hence the section is simply a function η from the end/neck to $T_{u_0(n)}X$. In terms of this function, the contribution to the norm squared is (for end/neck, respectively)

$$(B.4.6) \quad \int_{[0, \infty) \times S^1} \sum_{j=0}^{k-1} |D^j \eta(s, t)|^2 e^{2\delta s} ds dt,$$

$$(B.4.7) \quad \int_{[0, 6S] \times S^1} \sum_{j=0}^{k-1} |D^j \eta(s, t)|^2 e^{2\delta \min(s, 6S-s)} ds dt.$$

By the Sobolev embedding of $W^{2,2}$ into C^0 in two dimensions, in any end or neck we have a uniform bound on $|D^j \eta(s, t)| e^{\delta s}$ or $|D^j \eta(s, t)| e^{\delta \min(s, 6S-s)}$, respectively ($1 \leq j \leq k - 3$), linear in $\|\eta\|_{k,2,\delta}$.

B.5 Based $\bar{\partial}$ -section $\mathcal{F}_{\alpha,y}$ and linearized operator $D_{\alpha,y}$

Fix a norm on E . On direct sums of normed spaces we use the direct sum norm $\|a \oplus b\| := \|a\| + \|b\|$.

We consider the partially defined function

$$\mathcal{F}_{\alpha,y}: C^\infty(C_\alpha, u_\alpha^*TX)_D \oplus E \rightarrow C^\infty(\tilde{C}_\alpha, \Omega_{\tilde{C}_\alpha, j_y}^{0,1} \otimes_{\mathbb{C}} u_\alpha^*TX),$$

$$\xi \mapsto \text{PT}_{\exp_{u_\alpha} \xi \rightarrow u_\alpha} [\bar{\partial}_y \exp_{u_\alpha} \xi + \lambda(e_0 + \text{proj}_E \xi)(\alpha, y, x_1, \dots, x_r, \cdot, (\exp_{u_\alpha} \xi)(\cdot))].$$

This function $\mathcal{F}_{\alpha,y}$ is defined for ξ in a C^1 -neighborhood of zero; for these ξ we define $x_i = x_i(\xi)$ as in (B.1.7). The subscript D indicates restriction to sections which are tangent to D at p_{n+1}, \dots, p_{n+l} . It follows that for ξ contained in a C^0 -neighborhood of zero, $\exp_{u_\alpha} \xi$ sends p_{n+1}, \dots, p_{n+l} to D° .

Now we observe that $\mathcal{F}_{\alpha,y}$ extends continuously to a map

$$(B.5.1) \quad \mathcal{F}_{\alpha,y}: W^{k,2,\delta}(C_\alpha, u_\alpha^*TX)_D \oplus E \rightarrow W^{k-1,2,\delta}(\tilde{C}_\alpha, \Omega_{\tilde{C}_\alpha, j_y}^{0,1} \otimes_{\mathbb{C}} u_\alpha^*TX)$$

which is defined for $\|\xi\|_{k,2,\delta} \leq c'$ (some $c' > 0$). Moreover, this map is highly differentiable (see (B.1.7)–(B.1.8) and the surrounding discussion; recall we have fixed $k \geq 6$). We denote by

$$(B.5.2) \quad D_{\alpha,y}: W^{k,2,\delta}(C_\alpha, u_\alpha^*TX)_D \oplus E \rightarrow W^{k-1,2,\delta}(\tilde{C}_\alpha, \Omega_{\tilde{C}_\alpha, j_y}^{0,1} \otimes_{\mathbb{C}} u_\alpha^*TX)$$

the derivative of $\mathcal{F}_{\alpha,y}$ at zero.

Let $T_\nabla(X, Y) := \nabla_X Y - \nabla_Y X - [X, Y]$ denote the torsion of ∇ . Let $(\cdot)_y^{0,1}$ denote the projection

$$\text{Hom}_{\mathbb{R}}(T\tilde{C}_\alpha, u_\alpha^*TX) \rightarrow \Omega_{\tilde{C}_\alpha, j_y}^{0,1} \otimes_{\mathbb{C}} u_\alpha^*TX,$$

$$\text{so } (V)_y^{0,1} := \frac{1}{2}(V + J \circ V \circ j_y).$$

Lemma B.5.1 *The linearized operator $D_{\alpha,y}$ is given by*

$$(B.5.3) \quad D_{\alpha,y}\xi = (\nabla\xi + T_\nabla(\xi, du_\alpha))_y^{0,1} + \sum_{i=1}^r \frac{d[\lambda(e_0)]}{dx_i}(\alpha, y, x_1, \dots, x_r, \cdot, u_\alpha(\cdot))(-\text{proj}_{TC_\alpha} \xi(x_i)) + \nabla_\xi[\lambda(e_0)](\alpha, y, x_1, \dots, x_r, \cdot, u_\alpha(\cdot)) + \lambda(\text{proj}_E \xi)(\alpha, y, x_1, \dots, x_r, \cdot, u_\alpha(\cdot)),$$

where $\text{proj}: u_\alpha^*TX \rightarrow TC_\alpha$ denotes the projection associated to the splitting

$$T_{u_\alpha(x_i)}X = T_{u_\alpha(x_i)}D_i \oplus T_{x_i}C_\alpha,$$

and $\nabla_\xi[\lambda(e_0)]$ means covariant derivative in the direction of ξ along the X factor.

Proof The first term $(\nabla\xi + T_\nabla(\xi, du_\alpha))_y^{0,1}$ comes from differentiating

$$(B.5.4) \quad \text{PT}_{\exp_{u_\alpha} \xi \rightarrow u_\alpha}[\bar{\partial}_y \exp_{u_\alpha} \xi].$$

The second two terms come from differentiating

$$(B.5.5) \quad \text{PT}_{\exp_{u_\alpha} \xi \rightarrow u_\alpha}[\lambda(e_0)(\alpha, y, x_1, \dots, x_r, \cdot, (\exp_{u_\alpha} \xi)(\cdot))],$$

where we use the fact that $(d/dt)x_i(\exp_{u_\alpha}(t\xi)) = -\text{proj}_{TC_\alpha} \xi(x_i)$. The last term comes from differentiating

$$(B.5.6) \quad \text{PT}_{\exp_{u_\alpha} \xi \rightarrow u_\alpha}[\lambda(\text{proj}_E \xi)(\alpha, y, x_1, \dots, x_r, \cdot, (\exp_{u_\alpha} \xi)(\cdot))].$$

We leave the calculations to the reader. □

Lemma B.5.2 Consider the following commutative square:

$$(B.5.7) \quad \begin{array}{ccc} W^{k,2,\delta}(C_\alpha, u_\alpha^*TX)_D \oplus E & \xrightarrow{D_{\alpha,y}} & W^{k-1,2,\delta}(\tilde{C}_\alpha, \Omega_{\tilde{C}_\alpha, j_y}^{0,1} \otimes_{\mathbb{C}} u_\alpha^*TX) \\ \uparrow & & \uparrow \\ W^{k,2}(C_\alpha, u_\alpha^*TX)_D \oplus E & \xrightarrow{D_{\alpha,y}} & W^{k-1,2}(\tilde{C}_\alpha, \Omega_{\tilde{C}_\alpha, j_y}^{0,1} \otimes_{\mathbb{C}} u_\alpha^*TX) \end{array}$$

Both horizontal operators are Fredholm, and the induced maps from the kernel and cokernel of the bottom horizontal map to the kernel and cokernel of the top horizontal map, respectively, are isomorphisms.

Let us remark on the reason for the vertical inclusions in (B.5.7). For the leftmost vertical inclusion, one just needs to check that for a function $f: D^2 \rightarrow \mathbb{R}$, the $W^{k,2,\delta}$ norm of its pullback to $[0, \infty) \times S^1$ is bounded linearly by its $W^{k,2}$ norm on D^2 . For the rightmost vertical inclusion, one checks the same property for 1-forms. In both of these calculations we use the fact that $\delta < 1$.

Proof The operator $D_{\alpha,y}$ equals $(\nabla\xi)_y^{0,1}$ plus compact terms (by Lemma B.5.1), so it is equivalent to show that $\xi \mapsto (\nabla\xi)_y^{0,1}$ is Fredholm (and we may forget about E). Here ∇ is really the pullback of ∇ via u_α . Thus it suffices to show that if V is

a smooth complex vector bundle over a nodal Riemann surface C equipped with a smooth connection ∇ , then the horizontal maps

$$(B.5.8) \quad \begin{array}{ccc} W^{k,2,\delta}(C, V) & \xrightarrow{\xi \mapsto (\nabla\xi)^{0,1}} & W^{k-1,2,\delta}(\tilde{C}, \Omega_{\tilde{C}}^{0,1} \otimes_{\mathbb{C}} V) \\ \uparrow & & \uparrow \\ W^{k,2}(C, V) & \xrightarrow{\xi \mapsto (\nabla\xi)^{0,1}} & W^{k-1,2}(\tilde{C}, \Omega_{\tilde{C}}^{0,1} \otimes_{\mathbb{C}} V) \end{array}$$

are Fredholm. For the bottom map, the Fredholm property is standard. Indeed, $W^{k,2}(C, V)$ lies inside $W^{k,2}(\tilde{C}, V)$ as a closed subspace of finite codimension, and

$$\nabla^{0,1}: W^{k,2}(\tilde{C}, V) \rightarrow W^{k-1,2}(\tilde{C}, \Omega_{\tilde{C}}^{0,1} \otimes_{\mathbb{C}} V)$$

is Fredholm by elliptic regularity on the compact Riemann surface \tilde{C} . For the top map, the Fredholm property follows from results of Lockhart and McOwen [55] concerning the Fredholmness (in weighted Sobolev spaces) of elliptic operators on manifolds with cylindrical ends. Specifically, $W^{k,2,\delta}(C, V)$ lies inside $W^{k,2,\delta}(\tilde{C}, V)$ as a closed subspace of finite codimension, and the map

$$\nabla^{0,1}: W^{k,2,\delta}(\tilde{C}, V) \rightarrow W^{k-1,2,\delta}(\tilde{C}, \Omega_{\tilde{C}}^{0,1} \otimes_{\mathbb{C}} V)$$

is Fredholm by the theory of Lockhart and McOwen [55]. This uses the fact that the weight $\delta \in (0, 1)$ is not an eigenvalue of the asymptotic linearized operator in any end (the set of eigenvalues in any end is precisely \mathbb{Z} , corresponding to the powers $\{z^n\}_{n \in \mathbb{Z}}$ on $D^2 \setminus 0$). Let us also remark that both of these Fredholmness statements are easier at $p = 2$ than for general $p \in (1, \infty)$, since for $p = 2$ one can use Fourier analysis (after localizing the problem with a partition of unity).

Now it remains to show that the induced maps on kernel and cokernel are isomorphisms. The map on kernels is obviously injective, and the map on cokernels is easily seen to be surjective (use the fact that the image of the top horizontal map is closed and the fact that the rightmost vertical inclusion is dense).

For the surjectivity of the map on kernels and the injectivity of the map on cokernels, it suffices to show that if $D_{\alpha,y}\xi = \eta$ for $\xi \in W^{k,2,\delta}$ and $\eta \in W^{k-1,2}$, then $\xi \in W^{k,2}$. To show $\xi \in W^{k,2}$, it suffices to argue locally on one side of a given node.

Certainly (the vector field part of) ξ extends continuously to C_α (by definition of $W^{k,2,\delta}$), and this continuous extension satisfies $D_{\alpha,y}\xi = \eta + \epsilon$, where ϵ is a distribution on \tilde{C}_α valued in

$$\Omega_{\tilde{C}_\alpha, j_y}^{0,1} \otimes_{\mathbb{C}} u_\alpha^* TX$$

and supported inside the (inverse images of the) nodes. By elliptic regularity, it suffices to show that $\epsilon = 0$.

Now for smooth test functions φ supported inside a small neighborhood of the support of ϵ , we have

$$(B.5.9) \quad \langle \epsilon, \varphi \rangle = \langle D_{\alpha,y} \xi - \eta, \varphi \rangle = \langle \xi, D_{\alpha,y}^* \varphi \rangle - \langle \eta, \varphi \rangle.$$

We thus obtain the bound $|\langle \epsilon, \varphi \rangle| \leq c \|\varphi\|_{1,1}$ since $D_{\alpha,y}^*$ is a first-order operator and ξ, η are bounded (as they are continuous). On the other hand, ϵ is supported at a finite set of points, so it is a linear combination of δ -functions and their derivatives. Hence ϵ does not satisfy the bound $|\langle \epsilon, \varphi \rangle| \leq c \|\varphi\|_{1,1}$ (recall that $W^{1,1} \not\hookrightarrow C^0$ since we are in two dimensions) unless $\epsilon = 0$, as desired. \square

We denote the kernel of $D_{0,0}$ (whose meaning is unambiguous by Lemma B.5.2) by

$$(B.5.10) \quad K := \ker D_{0,0} \subseteq C^\infty(C_0, u_0^* TX)_D \oplus E.$$

Note that our assumption $(0, 0, u_0, \{x_i^0\}, e_0) \in \overline{\mathcal{M}}(X)^{\text{reg}}$ is equivalent to saying that $D_{0,0}$ is surjective and $K \rightarrow E/E'$ is surjective.

B.6 Pregluing estimates

Fix norms on K, \mathbb{C}^d , and $\mathbb{R}^{\dim \overline{\mathcal{M}}^d}$.

Lemma B.6.1 (Estimate for map pregluing) *We have*

$$(B.6.1) \quad \|\overline{\partial}_y u_\alpha + \lambda(e_0)(\alpha, y, x_1^0, \dots, x_r^0, \cdot, u_\alpha(\cdot))\|_{k-1,2,\delta} \leq c \cdot \left[|y| + \sum_{i=1}^d e^{-(1-\delta)S_i} \right]$$

uniformly over (α, y) in a neighborhood of zero, for $c < \infty$ depending on data which has been previously fixed.

Data which has been previously fixed includes $(s_0, u_0, \{x_i^0\}, e_0)$, the cylindrical ends, j_y , the metrics used to define exp and Sobolev norms, k, δ , the cutoff function χ , the norms on $E, K, \mathbb{C}^d, \mathbb{R}^{\dim \overline{\mathcal{M}}^d}$, etc.

Proof Away from the ends/necks of C_α , the expression is zero when $y = 0$ (since $(0, 0, u_0, \{x_i^0\}, e_0) \in \overline{\mathcal{M}}(X)$) and more generally it follows easily that its j^{th} derivative is bounded pointwise by $c_j |y|^j$. In the ends of C_α , the expression is zero.

In a given neck $[0, 6S] \times S^1 \subseteq C_\alpha$, argue as follows. Recall that λ vanishes in the ends/necks, so we just have to bound $\bar{\partial}_y u_\alpha$. Now $\bar{\partial}_y u_\alpha$ is only nonzero over $[S - 1, S] \cup [5S, 5S + 1]$; by symmetry we may just work over $[S - 1, S]$. Now since u_0 is smooth, it and all its derivatives decay exponentially in any cylindrical end $[0, \infty) \times S^1$ (since every end is holomorphically conjugate to the map $[0, \infty) \times S^1 \rightarrow D^2$ given by $z = e^{-s-it}$). Precisely, we have the following pointwise bound over $[S - 1, S]$:

$$(B.6.2) \quad |D^j u_0(s, t)| \leq c_j e^{-S}$$

(derivatives with respect to s and t). It follows from the definition of u_α that the same bound holds for u_α . Thus the contribution of $[S - 1, S]$ to the norm of $\bar{\partial}_y u_\alpha$ is bounded by a constant times $e^{S\delta} e^{-S}$. \square

Lemma B.6.2 (Estimate for kernel pregluing) *For all $\kappa \in K$, we have*

$$(B.6.3) \quad \|D_{\alpha,y} \kappa_\alpha\|_{k-1,2,\delta} \leq c \cdot \left[|y| + \sum_{i=1}^d e^{-(1-\delta)S_i} \right] \|\kappa\|$$

uniformly over (α, y) in a neighborhood of zero, for $c < \infty$ depending on data which has been previously fixed.

Proof Away from the ends/necks, the norm is bounded by $c \cdot |y| \|\kappa\|$. In the ends, the expression is zero. In the necks, argue as follows.

The expression is only nonzero over $[S - 1, S + 1] \cup [5S - 1, 5S + 1]$; by symmetry we may just work over $[S - 1, S + 1]$. Then since u_0 and κ are smooth, we have (by the same reasoning used to justify (B.6.2)) the following pointwise bound over $[S - 1, S + 1]$:

$$(B.6.4) \quad |D^j \kappa_\alpha(s, t)| \leq c_j e^{-S} \|\kappa\|$$

(derivatives with respect to s and t , in local coordinates (B.4.1)). Thus the overall contribution of $[S - 1, S + 1]$ to the total norm of the expression is bounded by a constant times $e^{S\delta} e^{-S} \|\kappa\|$. \square

B.7 Approximate right inverse

Recall that by assumption, the linearized operator

$$(B.7.1) \quad D_{0,0}: W^{k,2,\delta}(C_0, u_0^* TX)_D \oplus E \rightarrow W^{k-1,2,\delta}(\tilde{C}_0, \Omega_{\tilde{C}_0, j_0}^{0,1} \otimes_{\mathbb{C}} u_0^* TX)$$

is surjective (even if we replace E with E'). We now proceed to fix a bounded right inverse

$$(B.7.2) \quad Q_{0,0}: W^{k-1,2,\delta}(\tilde{C}_0, \Omega_{\tilde{C}_0,j_0}^{0,1} \otimes_{\mathbb{C}} u_0^*TX) \rightarrow W^{k,2,\delta}(C_0, u_0^*TX)_D \oplus E'$$

whose image admits a simple description (the description of the image will be important for the proof that the gluing map is continuous (Proposition B.10.4)). Fix a collection of points $q_i \in C_0$ ($1 \leq i \leq h$) not contained in any of the cylindrical ends, subspaces $V_i \subseteq T_{u_0(q_i)}X$, and a subspace $E'' \subseteq E'$ so that the natural evaluation map

$$(B.7.3) \quad L_0: K \xrightarrow{\simeq} \left(\bigoplus_{i=1}^h T_{u_0(q_i)}X / V_i \right) \oplus E/E''$$

is an isomorphism (such choices exist since $K \twoheadrightarrow E/E'$ is surjective and elements of K satisfy unique continuation). Now we can consider the same evaluation map on the larger space:

$$(B.7.4) \quad L_0: W^{k,2,\delta}(C_0, u_0^*TX)_D \oplus E \rightarrow W := \left(\bigoplus_{i=1}^h T_{u_0(q_i)}X / V_i \right) \oplus E/E''.$$

Since L_0 sends $K = \ker D_{0,0}$ isomorphically to W , it follows that the restriction of $D_{0,0}$ to $\ker L_0$ is an isomorphism of Banach spaces. Hence there is a unique right inverse

$$Q_{0,0}: W^{k-1,2,\delta}(\tilde{C}_0, \Omega_{\tilde{C}_0,j_0}^{0,1} \otimes_{\mathbb{C}} u_0^*TX) \rightarrow \ker L_0 \subseteq W^{k,2,\delta}(C_0, u_0^*TX)_D \oplus E,$$

and it is bounded. Since $E'' \subseteq E'$, $\ker L_0$ is in fact contained in the right-hand side of (B.7.2). We fix once and for all this $Q_{0,0}$.

Definition B.7.1 (Approximate right inverse $T_{\alpha,y}$) We define a map

$$(B.7.5) \quad T_{\alpha,y}: W^{k-1,2,\delta}(\tilde{C}_\alpha, \Omega_{\tilde{C}_\alpha,j_y}^{0,1} \otimes_{\mathbb{C}} u_\alpha^*TX) \rightarrow W^{k,2,\delta}(C_\alpha, u_\alpha^*TX)_D \oplus E$$

as the composition

$$(B.7.6) \quad T_{\alpha,y} := \text{glue} \circ \text{PT} \circ \text{id} \circ Q_{0,0} \circ I_y \circ \text{PT} \circ \text{break}$$

of maps in the following diagram, to be defined below:

$$\begin{array}{ccc}
 W^{k,2,\delta}(C_\alpha, u_\alpha^*TX)_D \oplus E & \xrightarrow{D_{\alpha,y}} & W^{k-1,2,\delta}(\tilde{C}_\alpha, \Omega_{\tilde{C}_\alpha, j_y}^{0,1} \otimes_{\mathbb{C}} u_\alpha^*TX) \\
 \uparrow \text{glue} & & \downarrow \text{break} \\
 W^{k,2,\delta}(C_0, u_{0|\alpha}^*TX)_D \oplus E & \xrightarrow{D_{0|\alpha,y}} & W^{k-1,2,\delta}(\tilde{C}_0, \Omega_{\tilde{C}_0, j_y}^{0,1} \otimes_{\mathbb{C}} u_{0|\alpha}^*TX) \\
 \uparrow \text{PT} & & \downarrow \text{PT} \\
 W^{k,2,\delta}(C_0, u_0^*TX)_D \oplus E & \xrightarrow{D_{0,y}} & W^{k-1,2,\delta}(\tilde{C}_0, \Omega_{\tilde{C}_0, j_y}^{0,1} \otimes_{\mathbb{C}} u_0^*TX) \\
 \parallel \text{id} & & \uparrow I_y \\
 W^{k,2,\delta}(C_0, u_0^*TX)_D \oplus E & \xleftrightarrow[\mathcal{Q}_{0,0}]{D_{0,0}} & W^{k-1,2,\delta}(\tilde{C}_0, \Omega_{\tilde{C}_0, j_0}^{0,1} \otimes_{\mathbb{C}} u_0^*TX)
 \end{array}$$

($D_{0|\alpha,y}$ denotes the linearized operator at $u_{0|\alpha}$).

We fix once and for all a smooth family $I_y: (T\tilde{C}_0, j_0) \rightarrow (T\tilde{C}_0, j_y)$ of (j_0, j_y) -linear identifications which are the identity over the ends/necks. This gives rise to the bottom right vertical map in (B.7.7).

We define the middle vertical maps in (B.7.7) using parallel transport (equivalently, using the local trivializations defined by (B.4.1)).

We define the map

$$\text{(B.7.8)} \quad W^{k-1,2,\delta}(\tilde{C}_\alpha, \Omega_{\tilde{C}_\alpha, j_y}^{0,1} \otimes_{\mathbb{C}} u_\alpha^*TX) \xrightarrow{\text{break}} W^{k-1,2,\delta}(\tilde{C}_0, \Omega_{\tilde{C}_0, j_0}^{0,1} \otimes_{\mathbb{C}} u_{0|\alpha}^*TX)$$

as follows. Fix a smooth function $\bar{\chi}: \mathbb{R} \rightarrow [0, 1]$ such that

$$\text{(B.7.9)} \quad \bar{\chi}(x) = \begin{cases} 1, & x \leq -1, \\ 0, & x \geq +1 \end{cases} \quad \text{and} \quad \bar{\chi}(x) + \bar{\chi}(-x) = 1.$$

Let

$$\eta \in W^{k-1,2,\delta}(\tilde{C}_\alpha, \Omega_{\tilde{C}_\alpha, j_y}^{0,1} \otimes_{\mathbb{C}} u_\alpha^*TX).$$

Away from the ends with $\alpha \neq 0$, $\text{break}(\eta)$ is simply η . In any particular end $[0, \infty) \times S^1 \subseteq C_0$ with $\alpha \neq 0$, we define

$$\text{(B.7.10)} \quad \text{break}(\eta)(s, t) := \begin{cases} \eta(s, t), & s \leq 3S - 1, \\ \bar{\chi}(s - 3S) \cdot \eta(s, t), & 3S - 1 \leq s \leq 3S + 1, \\ 0, & 3S + 1 \leq s \end{cases}$$

(noting the corresponding neck $[0, 6S] \times S^1 \subseteq C_\alpha$).

We define the map

$$(B.7.11) \quad W^{k,2,\delta}(C_0, u_{0|\alpha}^* TX)_D \xrightarrow{\text{glue}} W^{k,2,\delta}(C_\alpha, u_\alpha^* TX)_D.$$

Let $\xi \in W^{k,2,\delta}(C_0, u_{0|\alpha}^* TX)_D$. Away from the necks of C_α , $\text{glue}(\xi)$ is simply ξ . In any particular neck $[0, 6S] \times S^1 \subseteq C_\alpha$, we define

$$(B.7.12) \quad \text{glue}(\xi)(s, t) := \begin{cases} \xi(s, t), & s \leq 2S, \\ \xi(n) + \chi(4S - s) \cdot [\xi(s, t) - \xi(n)] \\ \quad + \chi(4S - s') \cdot [\xi(s', t') - \xi(n)], & 2S \leq s \leq 4S, \\ \xi(s', t'), & 4S \leq s, \end{cases}$$

noting the corresponding ends

$$(s, t) \in [0, \infty) \times S^1 \subseteq C_0 \quad \text{and} \quad (s', t') \in [0, \infty) \times S^1 \subseteq C_0.$$

The cutoff of $\xi(s, t)$ occurs around $4S \in [0, 6S]$, where the weight $e^{2\delta \min(s, 6S-s)}$ is much smaller than the weight $e^{2\delta s}$ at $4S \in [0, \infty)$. We will see in the proof of Lemma B.7.4 that this makes it easy to derive the desired estimates on $\bar{\partial} \text{glue}(\xi)$. This trick was explained to us by Abouzaid and attributed to Fukaya, Oh, Ohta and Ono [31].

Let us make the elementary observation that the definition of L_0 extends perfectly well to give an analogous bounded linear map

$$(B.7.13) \quad L_\alpha: W^{k,2,\delta}(C_\alpha, u_\alpha^* TX)_D \oplus E \rightarrow W.$$

Since $\text{im } Q_{0,0} \subseteq \ker L_0$, it follows from the definition of $T_{\alpha,y}$ that $\text{im } T_{\alpha,y} \subseteq \ker L_\alpha$ as well. This is a key ingredient in the proof that the gluing map is continuous (Proposition B.10.4).

Lemma B.7.2 *Let*

$$(B.7.14) \quad \begin{array}{ccc} X & \xrightarrow{D} & Y \\ G \uparrow & & \downarrow B \\ X' & \xrightarrow{D'} & Y' \end{array}$$

denote the bottom square in (B.7.7). Then for $\xi \in X'$ and $\eta \in Y$ with $D'\xi = B\eta$, we have

$$(B.7.15) \quad \|DG\xi - \eta\| \leq c \cdot \|\xi\| |y|$$

uniformly over (α, y) in a neighborhood of zero, for $c < \infty$ depending on data which has been previously fixed.

Proof In simpler terms, we have $\|D_{0,y} - I_y \circ D_{0,0}\| \leq c \cdot |y|$ (calculation left to the reader) and this trivially implies the claimed statement. \square

Lemma B.7.3 *Let*

$$(B.7.16) \quad \begin{array}{ccc} X & \xrightarrow{D} & Y \\ G \uparrow & & \downarrow B \\ X' & \xrightarrow{D'} & Y' \end{array}$$

denote the middle square in (B.7.7). Then for $\xi \in X'$ and $\eta \in Y$ with $D'\xi = B\eta$, we have

$$(B.7.17) \quad \|DG\xi - \eta\| \leq c \cdot \|\xi\| \sum_{i=1}^d e^{-S_i}$$

uniformly over (α, y) in a neighborhood of zero, for $c < \infty$ depending on data which has been previously fixed.

Proof In simpler terms, we bound the operator norm of the difference between the two diagonal compositions:

$$(B.7.18) \quad \|\text{PT} \circ D_{0,y} - D_{0|\alpha,y} \circ \text{PT}\| \leq c \cdot \sum_{i=1}^d e^{-S_i}$$

(this trivially implies the claimed statement). To show (B.7.18), observe that the two operators only differ over the $[S - 1, \infty)$ subset of each end $[0, \infty) \times S^1 \subseteq C_0$. Using estimates (B.6.2) arising from the smoothness of u_0 , we obtain the desired bound. \square

Lemma B.7.4 *Let*

$$(B.7.19) \quad \begin{array}{ccc} X & \xrightarrow{D} & Y \\ G \uparrow & & \downarrow B \\ X' & \xrightarrow{D'} & Y' \end{array}$$

denote the top square in (B.7.7). Then for $\xi \in X'$ and $\eta \in Y$ with $D'\xi = B\eta$, we have

$$(B.7.20) \quad \|DG\xi - \eta\| \leq c \cdot \|\xi\| \sum_{i=1}^d e^{-2\delta S_i}$$

uniformly over (α, y) in a neighborhood of zero, for $c < \infty$ depending on data which has been previously fixed.

Proof The difference $DG\xi - \eta$ is only nonzero in the necks over $s \in [2S, 2S + 1]$ and $s \in [4S - 1, 4S]$. By symmetry, we may just do the bound over $s \in [4S - 1, 4S]$. Over this region, one calculates that the norm of the difference is bounded as claimed. The factor of $e^{-2\delta S}$ comes as the ratio between the $e^{2\delta S}$ weight given to

$$[4S - 1, 4S] \times S^1 \subseteq [0, 6S] \times S^1 \subseteq C_\alpha$$

and the $e^{4\delta S}$ weight given to

$$[4S - 1, 4S] \times S^1 \subseteq [0, \infty) \times S^1 \subseteq C_0. \quad \square$$

Lemma B.7.5 *Let X and Y be Banach spaces, let $D: X \rightarrow Y$ and $T: Y \rightarrow X$ be bounded, and suppose $\epsilon := \|DT - \mathbf{1}_Y\| < 1$. Then*

$$(B.7.21) \quad Q := T \cdot \sum_{n \geq 0} (\mathbf{1}_Y - DT)^n$$

converges and satisfies $DQ = \mathbf{1}_Y$. In addition, we have the following estimates:

$$(B.7.22) \quad \|Q\| \leq \frac{1}{1-\epsilon} \|T\|, \quad \|Q - T\| \leq \frac{\epsilon}{1-\epsilon} \|T\|.$$

Proof Use the telescoping sum $(\mathbf{1} - A) \sum_{n \geq 0} A^n = \mathbf{1}$ for $\|A\| < 1$. □

Lemma B.7.6 *Suppose we have a diagram of Banach spaces as follows:*

$$(B.7.23) \quad \begin{array}{ccc} X_1 & \xrightarrow{D_1} & Y_1 \\ G_1 \uparrow & & \downarrow B_1 \\ \vdots & & \vdots \\ G_{n-2} \uparrow & & \downarrow B_{n-2} \\ X_{n-1} & \xrightarrow{D_{n-1}} & Y_{n-1} \\ G_{n-1} \uparrow & & \downarrow B_{n-1} \\ X_n & \xrightleftharpoons[Q_n]{} & Y_n \end{array}$$

where $\|D_i\|, \|G_i\|, \|B_i\|, \|Q_n\| \leq c$ and $D_n Q_n = \mathbf{1}$. Then for all $\delta > 0$ there exists $\epsilon = \epsilon(n, c, \delta) > 0$ such that if

$$(B.7.24) \quad D_i \xi = B_{i-1} \eta \implies \|D_{i-1} G_{i-1} \xi - \eta\| \leq \epsilon \cdot \|\xi\| \quad \text{for all } 1 < i \leq n$$

then we have

$$(B.7.25) \quad \|D_1 G_1 \cdots G_{n-1} Q_n B_{n-1} \cdots B_1 - \mathbf{1}_{Y_1}\| \leq \delta.$$

Proof We work by induction on n . The case $n = 1$ is clear: $\epsilon(1, c, \delta) = \infty$.

Now assume $n \geq 2$. Let us also assume without loss of generality that $c \geq 1$. Applying (B.7.24) to $i = n$ and $\xi = Q_n B_{n-1} \eta$, we see that

$$(B.7.26) \quad \|D_{n-1} G_{n-1} Q_n B_{n-1} - \mathbf{1}_{Y_{n-1}}\| \leq \epsilon \cdot \|Q_n\| \|B_{n-1}\| \leq \epsilon c^2.$$

Let us require that $\epsilon \leq \frac{1}{2} c^{-2}$, so the above bound is $\leq \frac{1}{2}$. Then by Lemma B.7.5 applied to $T = G_{n-1} Q_n B_{n-1}$, we see that there exists Q_{n-1} with $D_{n-1} Q_{n-1} = \mathbf{1}_{Y_{n-1}}$ and

$$(B.7.27) \quad \|Q_{n-1}\| \leq 2c^3, \quad \|Q_{n-1} - G_{n-1} Q_n B_{n-1}\| \leq 2\epsilon c^5.$$

Now we see that

$$(B.7.28) \quad \begin{aligned} & \|D_1 G_1 \cdots G_{n-2} G_{n-1} Q_n B_{n-1} B_{n-2} \cdots B_1 \\ & \quad - D_1 G_1 \cdots G_{n-2} Q_{n-1} B_{n-2} \cdots B_1\| \\ & \leq c^{2n-3} \|G_{n-1} Q_n B_{n-1} - Q_{n-1}\| \\ & \leq 2\epsilon c^{2n+2}. \end{aligned}$$

Let us require that $\epsilon \leq \frac{1}{4} \delta c^{-2n-2}$, so the above bound is $\leq \frac{1}{2} \delta$. Let us also require that $\epsilon \leq \epsilon(n-1, 2c^3, \frac{1}{2} \delta)$ (which exists by the induction hypothesis), so that

$$(B.7.29) \quad \|D_1 G_1 \cdots G_{n-2} Q_{n-1} B_{n-2} \cdots B_1 - \mathbf{1}_{Y_1}\| \leq \frac{1}{2} \delta.$$

Combining (B.7.28) and (B.7.29), we get the desired bound (B.7.25). □

Proposition B.7.7 (Approximate right inverse $T_{\alpha,y}$) We have

$$(B.7.30) \quad \|T_{\alpha,y}\| \leq c,$$

$$(B.7.31) \quad \|D_{\alpha,y} T_{\alpha,y} - \mathbf{1}\| \rightarrow 0,$$

$$(B.7.32) \quad \text{im } T_{\alpha,y} \subseteq \ker L_\alpha$$

as $(\alpha, y) \rightarrow 0$, for $c < \infty$ depending on data which has been previously fixed.

Proof It is easy to see that all the maps in (B.7.7) are uniformly bounded. Hence $\|T_{\alpha,y}\| \leq c$ as $(\alpha, y) \rightarrow 0$. Now Lemma B.7.6 combined with Lemmas B.7.2, B.7.3, B.7.4 show that for $(\alpha, y) \rightarrow 0$, we have $\|D_{\alpha,y} T_{\alpha,y} - \mathbf{1}\| \rightarrow 0$. We observed earlier that $\text{im } T_{\alpha,y} \subseteq \ker L_\alpha$. □

Definition B.7.8 (Right inverse $Q_{\alpha,y}$) We define a map

$$(B.7.33) \quad Q_{\alpha,y}: W^{k-1,2,\delta}(\tilde{C}_\alpha, \Omega_{\tilde{C}_\alpha, j_y}^{0,1} \otimes_{\mathbb{C}} u_\alpha^* TX) \rightarrow W^{k,2,\delta}(C_\alpha, u_\alpha^* TX)_D \oplus E$$

as the sum

$$(B.7.34) \quad Q_{\alpha,y} := T_{\alpha,y} \sum_{k=0}^{\infty} (\mathbf{1} - D_{\alpha,y} T_{\alpha,y})^k.$$

Proposition B.7.9 *We have*

$$(B.7.35) \quad \|Q_{\alpha,y}\| \leq c,$$

$$(B.7.36) \quad D_{\alpha,y} Q_{\alpha,y} = \mathbf{1},$$

$$(B.7.37) \quad \text{im } Q_{\alpha,y} \subseteq \ker L_{\alpha}$$

uniformly over (α, y) in a neighborhood of zero, for $c < \infty$ depending on data which has been previously fixed.

Proof Apply Lemma B.7.5 and Proposition B.7.7. □

B.8 Quadratic estimates

Proposition B.8.1 (Quadratic estimate) *There exist $c' > 0$ and $c < \infty$ (depending on data which has been previously fixed) such that for $\|\xi_1\|_{k,2,\delta}, \|\xi_2\|_{k,2,\delta} \leq c'$, we have*

$$(B.8.1) \quad \|D_{\alpha,y}(\xi_1 - \xi_2) - (\mathcal{F}_{\alpha,y}\xi_1 - \mathcal{F}_{\alpha,y}\xi_2)\|_{k-1,2,\delta} \leq c \cdot \|\xi_1 - \xi_2\|_{k,2,\delta} (\|\xi_1\|_{k,2,\delta} + \|\xi_2\|_{k,2,\delta})$$

(and $\mathcal{F}_{\alpha,y}\xi_1$ and $\mathcal{F}_{\alpha,y}\xi_2$ are both defined), uniformly over (α, y) in a neighborhood of zero.

Proof This is similar to McDuff and Salamon [62, page 68, Proposition 3.5.3].

We have already remarked that $\mathcal{F}_{\alpha,y}\xi$ is defined for $\|\xi\|_{k,2,\delta} \leq c'$.

Let $\mathcal{F}'_{\alpha,y}(\zeta, \xi)$ denote the derivative of $\mathcal{F}_{\alpha,y}$ at ζ applied to ξ . So, for instance, $D_{\alpha,y}(\xi) := \mathcal{F}'_{\alpha,y}(0, \xi)$. It suffices to show that

$$(B.8.2) \quad \|\mathcal{F}'_{\alpha,y}(0, \xi) - \mathcal{F}'_{\alpha,y}(\zeta, \xi)\|_{k-1,2,\delta} \leq c \cdot \|\zeta\|_{k,2,\delta} \|\xi\|_{k,2,\delta}$$

for $\|\zeta\|_{k,2,\delta} \leq c'$ uniformly as $(\alpha, y) \rightarrow 0$ (one recovers (B.8.1) by integrating $\int_{\xi_1}^{\xi_2} \mathcal{F}'_{\alpha,y}(0, d\zeta) - \mathcal{F}'_{\alpha,y}(\zeta, d\zeta)$).

For $\zeta \in W^{k,2,\delta}(C_{\alpha}, u_{\alpha}^*TX)_D \oplus E$, let

$$(B.8.3) \quad \mathcal{F}_{\alpha,y,\zeta}: W^{k,2,\delta}(C_{\alpha}, (\exp_{u_{\alpha}} \zeta)^*TX)_D \oplus E \rightarrow W^{k-1,2,\delta}(\tilde{C}_{\alpha}, \Omega_{\tilde{C}_{\alpha},j_y}^{0,1} \otimes_{\mathbb{C}} (\exp_{u_{\alpha}} \zeta)^*TX)$$

denote the $\bar{\partial}$ -section based at $\exp_{u_\alpha} \zeta: C_\alpha \rightarrow X$ and $e_0 + \text{proj}_E \zeta$ (so, for example, $\mathcal{F}_{\alpha,y} := \mathcal{F}_{\alpha,y,0}$). Let

$$(B.8.4) \quad D_{\alpha,y,\zeta}: W^{k,2,\delta}(C_\alpha, (\exp_{u_\alpha} \zeta)^*TX)_D \oplus E \rightarrow W^{k-1,2,\delta}(\tilde{C}_\alpha, \Omega_{\tilde{C}_\alpha,j_y}^{0,1} \otimes_{\mathbb{C}} (\exp_{u_\alpha} \zeta)^*TX)$$

denote the derivative of $\mathcal{F}_{\alpha,y,\zeta}$ at zero. Of course, $D_{\alpha,y,\zeta}$ may be calculated as in Lemma B.5.1, and the result is the same (ie we just substitute $(\exp_{u_\alpha} \zeta, e_0 + \text{proj}_E \zeta)$ in place of (u_α, e_0)).

Now the first step in proving (B.8.2) is to express $\mathcal{F}'_{\alpha,y}(\zeta, \xi)$ in terms of $D_{\alpha,y,\zeta}$. To do this, we observe that

$$(B.8.5) \quad \mathcal{F}_{\alpha,y}(a) = \left[\text{PT}_{\exp_{u_\alpha} a \rightarrow u_\alpha} \circ \text{PT}_{\exp_{u_\alpha} \zeta \rightarrow \exp_{u_\alpha} a} \right] \left[\mathcal{F}_{\alpha,y,\zeta}((\exp_{\exp_{u_\alpha} \zeta}^{-1} \exp_{u_\alpha} a) \oplus (\text{proj}_E a - \text{proj}_E \zeta)) \right].$$

We now differentiate with respect to a and evaluate at $a = \zeta$ and $\dot{a} = \xi$. We find

$$(B.8.6) \quad \mathcal{F}'_{\alpha,y}(\zeta, \xi) = \left[\frac{d}{da} \Big|_{\substack{a=\zeta \\ \dot{a}=\xi}} (\text{PT}_{\exp_{u_\alpha} a \rightarrow u_\alpha} \circ \text{PT}_{\exp_{u_\alpha} \zeta \rightarrow \exp_{u_\alpha} a}) \right] (\mathcal{F}_{\alpha,y,\zeta}(0)) + \text{PT}_{\exp_{u_\alpha} \zeta \rightarrow u_\alpha} \left[D_{\alpha,y,\zeta} \left(\frac{d}{da} \Big|_{\substack{a=\zeta \\ \dot{a}=\xi}} (\exp_{\exp_{u_\alpha} \zeta}^{-1} \exp_{u_\alpha} a) \oplus \text{proj}_E \xi \right) \right].$$

We rewrite the first term as

$$(B.8.7) \quad \left[\frac{d}{da} \Big|_{\substack{a=\zeta \\ \dot{a}=\xi}} (\text{PT}_{\exp_{u_\alpha} a \rightarrow u_\alpha} \circ \text{PT}_{\exp_{u_\alpha} \zeta \rightarrow \exp_{u_\alpha} a}) \right] \text{PT}_{u_\alpha \rightarrow \exp_{u_\alpha} \zeta} (\mathcal{F}_{\alpha,y}(\zeta)).$$

We know that $\|\mathcal{F}_{\alpha,y}(\zeta)\|_{k-1,2,\delta}$ is bounded uniformly for $(\alpha, y) \rightarrow 0$ and $\|\zeta\|_{k,2,\delta} \leq c'$ (Lemma B.6.1 implies that $\|\mathcal{F}_{\alpha,y}(0)\|_{k-1,2,\delta}$ is bounded as $(\alpha, y) \rightarrow 0$, and from this one may derive a bound on $\|\mathcal{F}_{\alpha,y}(\zeta)\|_{k-1,2,\delta}$ in terms of $\|\zeta\|_{k,2,\delta}$). The operator $[(d/da)(\text{PT} \circ \text{PT})]\text{PT}$ in front is of the form $H(\zeta, \xi)$ for a smooth (non-linear) bundle map $H: TX \oplus TX \rightarrow \text{End}(TX)$ (defined in a neighborhood of $\zeta = \xi = 0$). Since H satisfies $H(0, \cdot) = H(\cdot, 0) = 0$, it follows that $\|H(\zeta, \xi)\|_{k,2,\delta}$ is bounded by $c \cdot \|\zeta\|_{k,2,\delta} \|\xi\|_{k,2,\delta}$ for $\|\zeta\|_{k,2,\delta}, \|\xi\|_{k,2,\delta} \leq c'$. Hence the $\|\cdot\|_{k-1,2,\delta}$ -norm of the first term in (B.8.6) is bounded by $c \cdot \|\zeta\|_{k,2,\delta} \|\xi\|_{k,2,\delta}$, so for the purposes of proving (B.8.2) it may be ignored.

The second term in (B.8.6) is approximated by

$$\text{PT}_{\exp_{u_\alpha} \zeta \rightarrow u_\alpha} [D_{\alpha,y,\zeta} (\text{PT}_{u_\alpha \rightarrow \exp_{u_\alpha} \zeta} \xi)]$$

with error

$$(B.8.8) \quad \text{PT}_{\exp_{u_\alpha} \zeta \rightarrow u_\alpha} \left[D_{\alpha,y,\zeta} \left(\left[\text{PT}_{u_\alpha \rightarrow \exp_{u_\alpha} \zeta} \xi \xi - \frac{d}{da} \Big|_{\substack{a=\zeta \\ \dot{a}=\xi}} (\exp_{\exp_{u_\alpha} \zeta}^{-1} \exp_{u_\alpha} a) \right] \oplus 0 \right) \right],$$

which we may write as

$$(B.8.9) \quad \text{PT}_{\exp_{u_\alpha} \zeta \rightarrow u_\alpha} \left[D_{\alpha,y,\zeta} \left(\text{PT}_{u_\alpha \rightarrow \exp_{u_\alpha} \zeta} \left[\xi - \text{PT}_{\exp_{u_\alpha} \zeta \rightarrow u_\alpha} \frac{d}{da} \Big|_{\substack{a=\zeta \\ \dot{a}=\xi}} (\exp_{\exp_{u_\alpha} \zeta}^{-1} \exp_{u_\alpha} a) \right] \oplus 0 \right) \right].$$

Now the $\| \cdot \|_{(k-1,2,\delta) \rightarrow (k-1,2,\delta)}$ norm of the outer PT is bounded uniformly for $(\alpha, y) \rightarrow 0$ and $\|\zeta\|_{k,2,\delta} \leq c'$, as is the $\| \cdot \|_{(k,2,\delta) \rightarrow (k-1,2,\delta)}$ norm of $D_{\alpha,y,\zeta}$ and the $\| \cdot \|_{(k,2,\delta) \rightarrow (k,2,\delta)}$ norm of the following PT. The difference $\xi - \text{PT}(d/da)(\cdot)$ is of the form $H(\zeta, \xi)$ for a smooth (non-linear) bundle map $H: TX \oplus TX \rightarrow TX$ (defined in a neighborhood of $\zeta = \xi = 0$). Since H satisfies $H(0, \cdot) = H(\cdot, 0) = 0$, it follows that $\|H(\zeta, \xi)\|_{k,2,\delta}$ is bounded by $c \cdot \|\zeta\|_{k,2,\delta} \|\xi\|_{k,2,\delta}$ for $\|\zeta\|_{k,2,\delta}, \|\xi\|_{k,2,\delta} \leq c'$. Hence the error (B.8.8) has $\| \cdot \|_{k-1,2,\delta}$ bounded by $c \cdot \|\zeta\|_{k,2,\delta} \|\xi\|_{k,2,\delta}$.

Thus we have reduced the estimate (B.8.2) to proving

$$(B.8.10) \quad \|D_{\alpha,y} - \text{PT}_{\exp_{u_\alpha} \zeta \rightarrow u} \circ D_{\alpha,y,\zeta} \circ \text{PT}_{u_\alpha \rightarrow \exp_{u_\alpha} \zeta}\|_{(k,2,\delta) \rightarrow (k-1,2,\delta)} \leq c \cdot \|\zeta\|_{k,2,\delta}.$$

We calculated $D_{\alpha,y}$ in Lemma B.5.1, and $D_{\alpha,y,\zeta}$ may be expressed in exactly the same way (specifically, it is obtained by taking the expression for $D_{\alpha,y}$ and replacing every occurrence of u_α with $\exp_{u_\alpha} \zeta$ and e_0 by $e_0 + \text{proj}_E \zeta$). Now we compare term by term to prove (B.8.10). We omit the details of this calculation. \square

B.9 Newton–Picard iteration

Lemma B.9.1 *There exists $c' > 0$ (depending on data which has been previously fixed) such that for sufficiently small (α, y) :*

- (i) *The map $\mathcal{F}_{\alpha,y}$ is defined for $\|\xi\|_{k,2,\delta} \leq c'$.*
- (ii) *For $\xi_1 - \xi_2 \in \text{im } Q_{\alpha,y}$ and $\|\xi_1\|_{k,2,\delta}, \|\xi_2\|_{k,2,\delta} \leq c'$, we have*

$$(B.9.1) \quad \|(\xi_1 - \xi_2) - (Q_{\alpha,y} \mathcal{F}_{\alpha,y} \xi_1 - Q_{\alpha,y} \mathcal{F}_{\alpha,y} \xi_2)\|_{k,2,\delta} \leq \frac{1}{2} \|\xi_1 - \xi_2\|_{k,2,\delta}.$$

Proof The first assertion has been shown earlier. For the second, write

$$\begin{aligned}
 \text{(B.9.2)} \quad & \|(\xi_1 - \xi_2) - (Q_{\alpha,y}\mathcal{F}_{\alpha,y}\xi_1 - Q_{\alpha,y}\mathcal{F}_{\alpha,y}\xi_2)\|_{k,2,\delta} \\
 & = \|Q_{\alpha,y}D_{\alpha,y}(\xi_1 - \xi_2) - (Q_{\alpha,y}\mathcal{F}_{\alpha,y}\xi_1 - Q_{\alpha,y}\mathcal{F}_{\alpha,y}\xi_2)\|_{k,2,\delta} \\
 & \leq \|Q_{\alpha,y}\| \|D_{\alpha,y}(\xi_1 - \xi_2) - (\mathcal{F}_{\alpha,y}\xi_1 - \mathcal{F}_{\alpha,y}\xi_2)\|_{k-1,2,\delta} \\
 & \leq c \cdot \|Q_{\alpha,y}\| \|\xi_1 - \xi_2\|_{k,2,\delta} (\|\xi_1\|_{k,2,\delta} + \|\xi_2\|_{k,2,\delta})
 \end{aligned}$$

by Proposition B.8.1. Since $\|Q_{\alpha,y}\|$ is uniformly bounded, this is enough. □

Proposition B.9.2 (Newton–Picard iteration) *There exists $c' > 0$ (depending on data which has been previously fixed) so that for $(\alpha, y, \kappa \in K)$ sufficiently small, there exists a unique $\kappa_{\alpha,y} \in W^{k,2,\delta}(C_\alpha, u_\alpha^*TX)_D \oplus E$ satisfying*

$$\text{(B.9.3)} \quad \kappa_{\alpha,y} \in \kappa_\alpha + \text{im } Q_{\alpha,y},$$

$$\text{(B.9.4)} \quad \|\kappa_{\alpha,y}\|_{k,2,\delta} \leq c',$$

$$\text{(B.9.5)} \quad \mathcal{F}_{\alpha,y}\kappa_{\alpha,y} = 0.$$

Proof In fact, we will show that $\kappa_{\alpha,y}$ is given explicitly as the limit of the Newton iteration

$$\text{(B.9.6)} \quad \xi_0 := \kappa_\alpha,$$

$$\text{(B.9.7)} \quad \xi_n := \xi_{n-1} - Q_{\alpha,y}\mathcal{F}_{\alpha,y}\xi_{n-1}.$$

By Lemma B.9.1, the map $\xi \mapsto \xi - Q_{\alpha,y}\mathcal{F}_{\alpha,y}\xi$ is a $\frac{1}{2}$ -contraction mapping when restricted to

$$\text{(B.9.8)} \quad \{\xi \in \kappa_\alpha + \text{im } Q_{\alpha,y} : \|\xi\|_{k,2,\delta} \leq c'\}.$$

To finish the proof, it suffices to show that (for sufficiently small (α, y, κ)) (B.9.8) is nonempty and is mapped to itself by $\xi \mapsto \xi - Q_{\alpha,y}\mathcal{F}_{\alpha,y}\xi$.

We know that $\|\kappa_\alpha\|_{k,2,\delta} \rightarrow 0$ as $\kappa \rightarrow 0$ (uniformly in (α, y)), so (B.9.8) is nonempty. By using Proposition B.8.1 with $(\xi_1, \xi_2) = (0, \kappa_\alpha)$ and Lemmas B.6.1 and B.6.2, we conclude that

$$\text{(B.9.9)} \quad \|\mathcal{F}_{\alpha,y}\kappa_\alpha\|_{k-1,2,\delta} \rightarrow 0$$

as $(\alpha, y, \kappa) \rightarrow 0$. Since the operator norm of $Q_{\alpha,y}$ is bounded uniformly as $(\alpha, y) \rightarrow 0$, we see that κ_α is almost fixed by $\xi \mapsto \xi - Q_{\alpha,y}\mathcal{F}_{\alpha,y}\xi$ as $(\alpha, y, \kappa) \rightarrow 0$. It then follows from the contraction property that $\xi \mapsto \xi - Q_{\alpha,y}\mathcal{F}_{\alpha,y}\xi$ maps (B.9.8) to itself. □

B.10 Gluing

Definition B.10.1 (Gluing map) We define

$$(B.10.1) \quad u_{\alpha,y,\kappa} := \exp_{u_\alpha} \kappa_{\alpha,y},$$

$$(B.10.2) \quad e_{\alpha,y,\kappa} := e_0 + \text{proj}_E \kappa_{\alpha,y},$$

where $\kappa_{\alpha,y}$ is as in Proposition B.9.2, and we consider the map

$$(B.10.3) \quad \mathbb{C}^d \times \mathbb{R}^{\dim \bar{M}^d} \times K \rightarrow \bar{M}(X),$$

$$(B.10.4) \quad (\alpha, y, \kappa) \mapsto (\alpha, y, u_{\alpha,y,\kappa}, e_{\alpha,y,\kappa})$$

(with $\{x_i\}$ understood). It follows from the definition that (B.10.3) commutes with the projection from both sides to $\bar{M} \times E/E'$.

Lemma B.10.2 *The gluing map (B.10.3) maps sufficiently small (α, y, κ) to $\bar{M}(X)^{\text{reg}}$.*

Proof This is true since $Q_{\alpha,y}$ gives an approximate right inverse at $(u_{\alpha,y,\kappa}, e_{\alpha,y,\kappa})$ (use (B.8.2) with $\zeta = \kappa_{\alpha,y}$). □

Let $K_\alpha \subseteq C^\infty(C_\alpha, u_\alpha^*TX)_D \oplus E$ denote the image of $\kappa \mapsto \kappa_\alpha$. It is clear by definition that $K \rightarrow K_\alpha$ is an isomorphism, and the respective $W^{k,2,\delta}$ norms are uniformly commensurable. It is also clear that the diagram

$$(B.10.5) \quad \begin{array}{ccc} K & \xrightarrow{\kappa \mapsto \kappa_\alpha} & K_\alpha \\ & \searrow L_0 & \swarrow L_\alpha \\ & & W \end{array}$$

commutes (all maps being isomorphisms). Since $\text{im } Q_{\alpha,y} \subseteq \ker L_\alpha$, it follows in particular that $\text{im } Q_{\alpha,y} \cap K_\alpha = 0$. On the other hand, an index calculation shows that $\text{ind } D_{\alpha,y} = \text{ind } D_{0,0}$ (note that by Lemma B.5.2, it suffices to calculate their indices as operators $W^{k,2} \rightarrow W^{k-1,2}$ on \tilde{C}_α and \tilde{C}_0 respectively, and this is a standard calculation as in McDuff and Salamon [62]). Both are surjective, and hence we have $\dim \text{coker } Q_{\alpha,y} = \dim \ker D_{\alpha,y} = \dim \ker D_{0,0} = \dim K = \dim K_\alpha$. It follows that $\text{im } Q_{\alpha,y} = \ker L_\alpha$ and that

$$(B.10.6) \quad \text{im } Q_{\alpha,y} \oplus K_\alpha \xrightarrow{\sim} W^{k,2,\delta}(C_\alpha, u_\alpha^*TX)_D \oplus E$$

is an isomorphism of Banach spaces. We claim that in fact the two norms are uniformly commensurable as $(\alpha, y) \rightarrow 0$. The map written is clearly uniformly bounded, so we just need to show the same for its inverse. It suffices to show that the projection from the right-hand side to K_α is uniformly bounded, but this is nothing other than L_α

(clearly uniformly bounded) composed with the inverse of the isomorphism in (B.10.5) (also uniformly bounded).

Lemma B.10.3 *The map (B.10.3) is injective in a neighborhood of zero.*

Proof Suppose that

$$(B.10.7) \quad (\alpha, y, u_{\alpha,y,\kappa}, e_{\alpha,y,\kappa}) = (\alpha', y', u_{\alpha',y',\kappa'}, e_{\alpha',y',\kappa'}).$$

We see immediately that $(\alpha, y) = (\alpha', y')$. Now we see that

$$(B.10.8) \quad \exp_{u_\alpha} \kappa_{\alpha,y} = \exp_{u_{\alpha'}} (\kappa')_{\alpha,y},$$

$$(B.10.9) \quad \text{proj}_E \kappa_{\alpha,y} = \text{proj}_E (\kappa')_{\alpha,y}.$$

Since the norms of $\kappa_{\alpha,y}$ and $\kappa'_{\alpha,y}$ go to zero as $(\alpha, y, \kappa, \kappa') \rightarrow 0$ (in $W^{k,2,\delta}$, and hence in C^0) and the injectivity radius of the exponential map is fixed, we see that $\kappa_{\alpha,y} = (\kappa')_{\alpha,y}$. It follows that $\kappa_\alpha - (\kappa')_\alpha \in \text{im } Q_{\alpha,y}$, but since $K_\alpha \cap \text{im } Q_{\alpha,y} = 0$ we conclude $\kappa_\alpha = (\kappa')_\alpha$ and hence $\kappa = \kappa'$. \square

Proposition B.10.4 *The map (B.10.3) is continuous in a neighborhood of zero.*

Proof The key ingredient in this proof is our precise control of the image of the right inverse $Q_{\alpha,y}$ (specifically, that $\text{im } Q_{\alpha,y} = \ker L_\alpha$).

Suppose $(\alpha_i, y_i, \kappa_i) \rightarrow (\alpha, y, \kappa)$ is a convergent net.⁵⁰ We will show that

$$(B.10.10) \quad (u_{\alpha_i,y_i,\kappa_i}, e_{\alpha_i,y_i,\kappa_i}) \rightarrow (u_{\alpha,y,\kappa}, e_{\alpha,y,\kappa}).$$

First, we claim that $\|(\kappa_i)_{\alpha_i,y_i} - \kappa_{\alpha_i,y_i}\|_\infty \rightarrow 0$. In fact, we will show the stronger statement that $\|(\kappa_i)_{\alpha_i,y_i} - \kappa_{\alpha_i,y_i}\|_{k,2,\delta} \rightarrow 0$. Since the Newton iteration converges uniformly, it suffices to show that $\|(\kappa_i)_{\alpha_i,y_i}^n - \kappa_{\alpha_i,y_i}^n\|_{k,2,\delta} \rightarrow 0$ for all $n \geq 0$, where $\kappa_{\alpha,y}^n$ denotes the n^{th} step of the Newton iteration converging to $\kappa_{\alpha,y}$. It is easy to verify this for $n = 0$:

$$(B.10.11) \quad \|(\kappa_i)_{\alpha_i,y_i}^0 - \kappa_{\alpha_i,y_i}^0\|_{k,2,\delta} = \|(\kappa_i - \kappa)_{\alpha_i}\|_{k,2,\delta} \rightarrow 0.$$

The inductive step follows from the fact that Q_{α_i,y_i} is uniformly bounded and $\mathcal{F}_{\alpha_i,y_i}$ is uniformly Lipschitz (eg as a consequence of Proposition B.8.1). Now from the claim, we see that it suffices to show that

$$(B.10.12) \quad (u_{\alpha_i,y_i,\kappa}, e_{\alpha_i,y_i,\kappa}) \rightarrow (u_{\alpha,y,\kappa}, e_{\alpha,y,\kappa}).$$

⁵⁰We could restrict to sequences rather than nets since $\mathbb{C}^d \times \mathbb{R}^{\dim \overline{\mathcal{M}}^d} \times K$ is first countable. However, this would not make the argument any simpler.

Recall that, by definition,

$$(B.10.13) \quad u_{\alpha,y,\kappa} = \exp_{u_\alpha} \kappa_{\alpha,y}, \quad \kappa_{\alpha,y} = \kappa_\alpha + \xi \quad \text{for some } \xi \in \text{im } Q_{\alpha,y},$$

$$(B.10.14) \quad u_{\alpha_i,y_i,\kappa} = \exp_{u_{\alpha_i}} \kappa_{\alpha_i,y_i}, \quad \kappa_{\alpha_i,y_i} = \kappa_{\alpha_i} + \xi_i \quad \text{for some } \xi_i \in \text{im } Q_{\alpha_i,y_i}.$$

Now we define $\xi_{\alpha_i} \in W^{k,2,\delta}(C_{\alpha_i}, u_{\alpha_i}^* TX)_D \oplus E$ by “pregluing” ξ from C_α to C_{α_i} as follows. Note that we may assume without loss of generality that at the nodes where $\alpha \neq 0$ we also have $\alpha_i \neq 0$. Away from the ends/necks of C_{α_i} , we set $\xi_{\alpha_i} = \xi$. Over the ends of C_{α_i} , note there is a corresponding end of C_α , so we may also simply set $\xi_{\alpha_i} = \xi$ over the ends of C_{α_i} . Over the necks of C_{α_i} for which $\alpha = 0$, we define ξ_{α_i} via (B.3.5) (note that this is reasonable since ξ is smooth on C_α). Over the necks of C_{α_i} for which $\alpha \neq 0$ we define ξ_{α_i} as

$$(B.10.15) \quad \xi_{\alpha_i}(s, t) := \text{PT}_{u_\alpha(f_i(s),t) \rightarrow u_{\alpha_i}(s,t)}[\xi(f_i(s), t)],$$

where $f_i: [0, 6S_i] \rightarrow [0, 6S]$ is defined as

$$(B.10.16) \quad f_i(s) := \begin{cases} s, & s \leq S - 2, \\ s - 3S_i + 3S, & 3S_i - 2S + 2 \leq s \leq 3S_i + 2S - 2, \\ s - 6S_i + 6S, & 6S_i - S + 2 \leq s, \end{cases}$$

$$(B.10.17) \quad f_i([S - 2, 3S_i - 2S + 2]) \subseteq [S - 2, S + 2],$$

$$(B.10.18) \quad f_i([3S_i + 2S - 2, 6S_i - S + 2]) \subseteq [5S - 2, 5S + 2],$$

and is smooth with absolutely bounded derivatives of all orders. More informally, f_i is smooth and matches up $[0, S - 2]$, $[S + 2, 5S - 2]$, $[5S + 2, 6S] \subseteq [0, 6S]$ with corresponding intervals of the same length inside $[0, 6S_i]$, symmetrically.

Now the C^0 -distance between

$$(B.10.19) \quad u_{\alpha,y,\kappa} = \exp_{u_\alpha}(\kappa_\alpha + \xi): C_\alpha \rightarrow X \quad \text{and} \quad \exp_{u_{\alpha_i}}(\kappa_{\alpha_i} + \xi_{\alpha_i}): C_{\alpha_i} \rightarrow X$$

goes to zero (this is easy to see from (B.3.5) and (B.10.15)). Hence it suffices to show that $\|\xi_i - \xi_{\alpha_i}\|_\infty \rightarrow 0$. In fact, we will prove the stronger statement that

$$(B.10.20) \quad \|\xi_i - \xi_{\alpha_i}\|_{k,2,\delta} \rightarrow 0.$$

First, we claim that we may assume that

$$(B.10.21) \quad \limsup_i \|\xi_{\alpha_i}\|_{k,2,\delta}$$

is arbitrarily small by taking (α, y, κ) sufficiently close to zero. We know that $\|F_{\alpha,y,\kappa}\|_{k-1,2,\delta} \rightarrow 0$ as $(\alpha, y, \kappa) \rightarrow 0$ by (B.9.9). Now since the Newton iteration is uniformly convergent and $Q_{\alpha,y}$ is uniformly bounded, it follows that $\|\xi\|_{k,2,\delta} \rightarrow 0$ as

$(\alpha, y, \kappa) \rightarrow 0$. We may easily bound $\|\xi_{\alpha_i}\|_{k,2,\delta}$ in terms of $\|\xi\|_{k,2,\delta}$ and the desired claim follows.

Now since the \limsup (B.10.21) can be assumed to be arbitrarily small, we may, in particular, assume that for sufficiently large i , $\|\kappa_{\alpha_i} + \xi_{\alpha_i}\|_{k,2,\delta} \leq c'$ for the constant $c' > 0$ from Lemma B.9.1 for which (B.9.8) is a domain of contraction. By construction, the fact that $\xi \in \text{im } Q_{\alpha,y} = \ker L_{\alpha}$ implies that $\xi_{\alpha_i} \in \ker L_{\alpha_i} = \text{im } Q_{\alpha_i,y_i}$. Hence $\kappa_{\alpha_i} + \xi_{\alpha_i}$ lies in the domain of contraction (B.9.8) for (α_i, y_i, κ) where the Newton iteration applies. By definition, we have $\mathcal{F}_{\alpha_i,y_i}(\kappa_{\alpha_i} + \xi_i) = 0$ and $\kappa_{\alpha_i} + \xi_i$ is in the same domain of contraction. Since the Newton iteration is a $\frac{1}{2}$ -contraction on this domain and Q_{α_i,y_i} is uniformly bounded, to show (B.10.20), it suffices to show that

$$(B.10.22) \quad \|\mathcal{F}_{\alpha_i,y_i}(\kappa_{\alpha_i} + \xi_{\alpha_i})\|_{k-1,2,\delta} \rightarrow 0.$$

To prove (B.10.22), first recall that $\mathcal{F}_{\alpha,y}(\kappa_{\alpha} + \xi) = 0$.

Away from the ends/necks of C_{α_i} , the expression $\mathcal{F}_{\alpha_i,y_i}(\kappa_{\alpha_i} + \xi_{\alpha_i})$ differs from $\mathcal{F}_{\alpha,y}(\kappa_{\alpha} + \xi) = 0$ only in terms of the complex structure j_{y_i} in place of j_y . Clearly this difference goes to zero as $y_i \rightarrow y$.

Over the ends of C_{α_i} , the expression $\mathcal{F}_{\alpha_i,y_i}(\kappa_{\alpha_i} + \xi_{\alpha_i})$ coincides with $\mathcal{F}_{\alpha,y}(\kappa_{\alpha} + \xi) = 0$.

Over the necks of C_{α_i} , we bound $\mathcal{F}_{\alpha_i,y_i}(\kappa_{\alpha_i} + \xi_{\alpha_i})$ as follows. Fix a neck $[0, 6S_i] \subseteq C_{\alpha_i}$. If $\alpha \neq 0$ for this neck, then the desired estimate follows easily from the definition of ξ_{α_i} , κ_{α_i} , and u_{α_i} , and the fact that $\mathcal{F}_{\alpha,y}(\kappa_{\alpha} + \xi) = 0$. If $\alpha = 0$ for this neck, then we argue as follows. The expression $\mathcal{F}_{\alpha_i,y_i}(\kappa_{\alpha_i} + \xi_{\alpha_i})$ is only nonzero over $[S_i-1, S_i+1] \cup [5S_i-1, 5S_i+1]$; by symmetry we may just bound it over $[S_i-1, S_i+1]$. As in the proof of Lemma B.6.1, we win because of the weights: the expression in question has derivatives bounded by e^{-S} and is weighted by $e^{\delta S}$. \square

B.11 Surjectivity of gluing

We identify $\mathbb{R}^{2n} = \mathbb{C}^n$ via $z_j = x_j + iy_j$. Let g_{std} denote the standard metric on \mathbb{R}^{2n} , and let J_{std} denote the standard almost complex structure on \mathbb{C}^n .

The following proposition is well known, and is essentially contained in McDuff and Salamon [62].

Proposition B.11.1 (A priori estimate on long pseudo-holomorphic necks) *For all $\mu < 1$, there exists $\epsilon > 0$ with the following property. Let J be a smooth almost complex structure on $B^{2n}(1) \subseteq \mathbb{R}^{2n}$ satisfying $\|J - J_{\text{std}}\|_{C^2} \leq \epsilon$ (measured with*

respect to g_{std}). There exists $c_k < \infty$ depending only on $\|J\|_{C^k}$ with the following property. Let $u: [-R, R] \times S^1 \rightarrow B^{2n}(1)$ be J -holomorphic. If

$$(B.11.1) \quad \int_{\partial[-R,R] \times S^1} \left| \frac{du}{dt} \right|^2 dt < \epsilon$$

then

$$(B.11.2) \quad |(D^k u)(s, t)| \leq c_k \cdot e^{\mu \cdot (|s| - R)} \left(\int_{\partial[-R,R] \times S^1} \left| \frac{du}{dt} \right|^2 dt \right)^{1/2}$$

for $|s| \leq R - 1$.

Proof First, apply Lemma B.11.2 to get an exponential decay bound on the $W^{1,2}$ norm of u restricted to $[-r, r] \times S^1$. Second, apply Lemma B.11.3 to conclude that the bound on the $W^{1,2}$ norm implies a similar bound on the $W^{1,\infty}$ norm. Finally, apply elliptic bootstrapping Lemma B.11.4 to conclude that the bound on the $W^{1,\infty}$ norm implies a similar bound on the C^k norm for all $k < \infty$. \square

Lemma B.11.2 ($W^{1,2}$ norm decays exponentially) *For all $\mu < 1$, there exists $\epsilon > 0$ with the following property. Let $u: [-R, R] \times S^1 \rightarrow B^{2n}(1)$ be J -holomorphic for an almost complex structure J on $B^{2n}(1) \subseteq \mathbb{R}^{2n}$ satisfying $\|J - J_{\text{std}}\| \leq \epsilon$ (measured with respect to g_{std}). Then*

$$\left(\int_{[-r,r] \times S^1} \left[\left| \frac{du}{ds} \right|^2 + \left| \frac{du}{dt} \right|^2 \right] ds dt \right)^{1/2} \leq c \cdot e^{\mu \cdot (r - R)} \left(\int_{\partial[-R,R] \times S^1} \left| \frac{du}{dt} \right|^2 dt \right)^{1/2}$$

for $0 \leq r \leq R$.

Proof This proof is essentially lifted from McDuff and Salamon [62, page 99, Lemma 4.7.3].

Let $\lambda_{\text{std}} := \frac{1}{2} \sum_j x_j dy_j - y_j dx_j$. Note the identity

$$(B.11.3) \quad \sum_j z_j d\bar{z}_j = \frac{1}{2} \sum_j d|z_j|^2 - 2i \lambda_{\text{std}}.$$

Let $\gamma: S^1 \rightarrow \mathbb{R}^{2n}$ be a smooth loop. Write the Fourier series $\gamma(t) = \sum_k a_k e^{ikt} \in \mathbb{C}^n$, where $a_k \in \mathbb{C}^n$. Now using (B.11.3), we see that

$$\int_{S^1} \gamma^* \lambda_{\text{std}} = \frac{i}{2} \int_{S^1} \sum_j \gamma^*(z_j d\bar{z}_j) = \frac{i}{2} \int_{S^1} \sum_{k,\ell} a_k \bar{a}_\ell (-i\ell) e^{i(k-\ell)t} dt = \pi \sum_k k |a_k|^2.$$

We may also calculate

$$(B.11.4) \quad \int_{S^1} \left| \frac{d\gamma}{dt} \right|^2 dt = \int_{S^1} \sum_{k,\ell} a_k \bar{a}_\ell (ik)(-i\ell) e^{i(k-\ell)t} dt = 2\pi \sum_k k^2 |a_k|^2.$$

Hence we conclude that

$$(B.11.5) \quad \left| \int_{S^1} \gamma^* \lambda_{\text{std}} \right| \leq \frac{1}{2} \int_{S^1} \left| \frac{d\gamma}{dt} \right|^2 dt.$$

Let

$$(B.11.6) \quad E(r) := \int_{[-r,r] \times S^1} \left| \frac{du}{dt} \right|^2 ds dt \approx \int_{[-r,r] \times S^1} u^* \omega_{\text{std}} = \int_{\partial[-r,r] \times S^1} u^* \lambda_{\text{std}}$$

(\approx means equality up to a factor which can be made arbitrarily close to 1 by taking $\epsilon > 0$ sufficiently small). Applying (B.11.5), we conclude that

$$(B.11.7) \quad E(r) \lesssim \frac{1}{2} \int_{\partial[-r,r] \times S^1} \left| \frac{du}{dt} \right|^2 dt$$

(\lesssim means inequality up to a factor which can be made arbitrarily close to 1 by taking $\epsilon > 0$ sufficiently small). The right-hand side above equals $\frac{1}{2} E'(r)$, so we have $E'(r) \gtrsim 2E(r)$, and hence $E'(r) \geq 2\mu E(r)$ (for $\epsilon > 0$ sufficiently small), from which we conclude that $E(r) \leq e^{2\mu \cdot (r-R)} E(R)$. Using (B.11.7) to bound $E(R)$, we see that

$$(B.11.8) \quad E(r) \lesssim e^{2\mu \cdot (r-R)} \frac{1}{2} \int_{\partial[-R,R] \times S^1} \left| \frac{du}{dt} \right|^2 dt.$$

We have

$$(B.11.9) \quad E(r) \approx \frac{1}{2} \int_{[-r,r] \times S^1} \left[\left| \frac{du}{ds} \right|^2 + \left| \frac{du}{dt} \right|^2 \right] ds dt.$$

Hence we conclude that

$$(B.11.10) \quad \int_{[-r,r] \times S^1} \left[\left| \frac{du}{ds} \right|^2 + \left| \frac{du}{dt} \right|^2 \right] ds dt \lesssim e^{2\mu \cdot (r-R)} \int_{\partial[-R,R] \times S^1} \left| \frac{du}{dt} \right|^2 dt,$$

which is the desired estimate. □

Lemma B.11.3 ($W^{1,2}$ controls $W^{1,\infty}$) *Let J be an almost complex structure on $B^{2n}(1) \subseteq \mathbb{R}^{2n}$; there exist $\epsilon > 0$ and $c < \infty$, depending only on $\|J\|_{C^2}$ (measured with respect to g_{std}) with the following property. Let $u: [0, 1] \times [0, 1] \rightarrow B^{2n}(1)$ be J -holomorphic, and suppose that*

$$(B.11.11) \quad \int_{[0,1]^2} \left[\left| \frac{du}{dx} \right|^2 + \left| \frac{du}{dy} \right|^2 \right] dx dy < \epsilon.$$

Then we have

$$(B.11.12) \quad \left| \frac{du}{dx} \left(\frac{1}{2}, \frac{1}{2} \right) \right| + \left| \frac{du}{dy} \left(\frac{1}{2}, \frac{1}{2} \right) \right| \leq c \cdot \left(\int_{[0,1]^2} \left[\left| \frac{du}{dx} \right|^2 + \left| \frac{du}{dy} \right|^2 \right] dx dy \right)^{1/2}.$$

Proof This proof is essentially lifted from McDuff and Salamon [62, page 80, Lemma 4.3.1].

Let $w := \frac{1}{2}u_x^2 + \frac{1}{2}u_y^2$ (using the standard inner product g_{std}), so $w: [0, 1]^2 \rightarrow \mathbb{R}_{\geq 0}$. Now we calculate

$$(B.11.13) \quad w_{xx} + w_{yy} = [u_{xx}^2 + 2u_{xy}^2 + u_{yy}^2] + [u_x \cdot (u_{xxx} + u_{xyy}) + u_y \cdot (u_{xxy} + u_{yyy})].$$

Now differentiating (B.11.20) yields

$$\begin{aligned} u_{xxx} + u_{xyy} &= \ddot{J}(u, u_x, u_y) \circ u_x + \dot{J}(u, u_{xy}) \circ u_x + \dot{J}(u, u_y) \circ u_{xx} \\ &\quad - \ddot{J}(u, u_x, u_x) \circ u_y - \dot{J}(u, u_{xx}) \circ u_y - \dot{J}(u, u_x) \circ u_{xy}, \\ u_{xxy} + u_{yyy} &= \ddot{J}(u, u_y, u_y) \circ u_x + \dot{J}(u, u_{yy}) \circ u_x + \dot{J}(u, u_y) \circ u_{xy} \\ &\quad - \ddot{J}(u, u_x, u_y) \circ u_y - \dot{J}(u, u_{xy}) \circ u_y - \dot{J}(u, u_x) \circ u_{yy}. \end{aligned}$$

Now we conclude that

$$\begin{aligned} &|u_x \cdot (u_{xxx} + u_{xyy}) + u_y \cdot (u_{xxy} + u_{yyy})| \\ &\leq 2\|J\|_{C^2} \cdot (|u_x|^3|u_y| + |u_x||u_y|^3) \\ &\quad + 2\|J\|_{C^1} \cdot (|u_x|^2|u_{xy}| + |u_x||u_y||u_{xx}| + |u_y|^2|u_{xy}| + |u_x||u_y||u_{yy}|) \\ &\leq c\|J\|_{C^2}w^2 + c\|J\|_{C^1}w \cdot (|u_{xx}| + |u_{xy}| + |u_{yy}|) \\ &\leq c\|J\|_{C^2}w^2 + c\|J\|_{C^1}^2w^2 + \frac{1}{100}(u_{xx}^2 + 2u_{xy}^2 + u_{yy}^2). \end{aligned}$$

Plugging this into (B.11.13), we conclude that

$$(B.11.14) \quad \Delta w = w_{xx} + w_{yy} \geq -c \cdot (\|J\|_{C^2} + \|J\|_{C^1}^2)w^2.$$

Now we may apply a mean value inequality (see McDuff and Salamon [62, page 81, Lemma 4.3.2] or Wehrheim [79, page 306, Theorem 1.1]) to see that there exist $\epsilon > 0$ and $c < \infty$ such that if $\int_{[0,1]^2} w dx dy < \epsilon$, then $w(\frac{1}{2}, \frac{1}{2}) \leq c \cdot \int_{[0,1]^2} w dx dy$. Thus we are done. \square

Lemma B.11.4 ($W^{1,\infty}$ controls $W^{k,p}$) Let $u: [0, 1] \times [0, 1] \rightarrow B^{2n}(1)$ be J -holomorphic for an almost complex structure J on $B^{2n}(1) \subseteq \mathbb{R}^{2n}$. For all $k \geq 1$, there

exists $c_k < \infty$ depending only on $\|J\|_{C^k}$ such that if

$$(B.11.15) \quad \sup_{(x,y) \in [0,1]^2} \left[\left| \frac{du}{dx} \right| + \left| \frac{du}{dy} \right| \right] \leq 1$$

then

$$(B.11.16) \quad |(D^k u)\left(\frac{1}{2}, \frac{1}{2}\right)| \leq c_k \cdot \sup_{(x,y) \in [0,1]^2} \left[\left| \frac{du}{dx} \right| + \left| \frac{du}{dy} \right| \right].$$

Proof For more details see McDuff and Salamon [62, page 533, Section B.4].

The $\bar{\partial}$ -equation for u may be written as

$$(B.11.17) \quad u_y = J(u) \circ u_x.$$

Differentiating with respect to x and to y , we conclude that

$$(B.11.18) \quad u_{xy} = \dot{J}(u, u_x) \circ u_x + J(u) \circ u_{xx},$$

$$(B.11.19) \quad u_{xx} = \dot{J}(u, u_y) \circ u_x + J(u) \circ u_{xy}.$$

Combining these and using the fact that \dot{J} and J anticommute, we conclude that

$$(B.11.20) \quad \Delta u = u_{xx} + u_{yy} = \dot{J}(u, u_y) \circ u_x - \dot{J}(u, u_x) \circ u_y.$$

Now a standard elliptic bootstrapping argument based on (B.11.20) gives the desired result. By hypothesis, we have an L^∞ bound in terms of $\|J\|_{C^1}$ on the right-hand side of (B.11.20), which gives a $W^{2,p}$ bound ($2 < p < \infty$) on u in terms of $\|J\|_{C^1}$. Now we have a $W^{1,p}$ bound on the right-hand side of (B.11.20) in terms of $\|J\|_{C^2}$, which gives a $W^{3,p}$ bound on u in terms of $\|J\|_{C^2}$. Iterating, we get a $W^{k,p}$ bound on u in terms of $\|J\|_{C^{k-1}}$, which is enough. There is no need to worry about elliptic regularity estimates near the boundary since we can shrink the domain slightly after each iteration. □

Proposition B.11.5 *The restriction of (B.10.3) to any neighborhood of zero is surjective onto a neighborhood of $(0, 0, u_0, e_0) \in \bar{\mathcal{M}}(X)$.*

Proof Let $(\alpha_i, y_i, u_i, \{x_j^i\}, e_i) \in \bar{\mathcal{M}}(X)$ be a sequence converging to $(0, 0, u_0, \{x_j^0\}, e_0)$. We must show that for i sufficiently large, $(\alpha_i, y_i, u_i, \{x_j^i\}, e_i)$ is contained in the image of the map (B.10.3). We may restrict to sequences rather than nets since the topology on $\bar{\mathcal{M}}(X)$ is first countable (recall the definition of the topology following (B.1.5); it is even a metric topology). This is convenient when we apply Arzelà–Ascoli.

Let us define $\xi_i \in C^\infty(C_{\alpha_i}, u_{\alpha_i}^* TX)_D \oplus E$ by the property

$$(B.11.21) \quad u_i = \exp_{u_{\alpha_i}} \xi_i,$$

$$(B.11.22) \quad e_i = e_0 + \text{proj}_E \xi_i$$

(and the exponential follows the shortest geodesic). Obviously $\text{proj}_E \xi_i \rightarrow 0$ and $\|\xi_i\|_\infty \rightarrow 0$.

Now we claim that $u_i \rightarrow u$ in the C^∞ topology away from the nodes, or equivalently, that $\xi_i \rightarrow 0$ in the C^∞ topology away from the nodes of C_0 . To see this, argue as follows. The thickened holomorphic curve equation from (B.2.6) is equivalent to an honest holomorphic curve equation for the graph $u: C \rightarrow C \times X$, where the almost complex structure on $C \times X$ is defined in terms of $\lambda(e)$. Hence the Gromov–Schwarz lemma (see [36, page 316, 1.3.A] or [65, page 223, Corollary 4.1.4]) applies and we conclude that $\|du_i\|_\infty$ is uniformly bounded on compact sets away from the nodes (equivalently, the same for $\|d\xi_i\|_\infty$). Now elliptic bootstrapping (as in the proof of Lemma B.11.4) implies that all derivatives are bounded uniformly on compact sets away from the nodes. Using Arzelà–Ascoli and diagonalization, we conclude that there exists a subsequence of u_i which is convergent in the C^∞ topology away from the nodes. Since we know that $u_i \rightarrow u$ in the C^0 topology, the limit of this C^∞ convergent subsequence must be u . This argument can be applied to any subsequence of u_i , so we conclude that every subsequence of u_i has a subsequence which converges to u in C^∞ away from the nodes. It follows that in fact u_i itself converges to u in C^∞ away from the nodes.

In the ends/necks of C_{α_i} , note that u_i is genuinely J -holomorphic (the λ term vanishes), so we may apply Proposition B.11.1 to conclude that $\|\xi_i\|_{k,2,\delta} \rightarrow 0$ (here using the fact that $\delta < 1$). We should remark that to apply Proposition B.11.1, we need to be able to choose some $M < \infty$ such that for any neck $[0, 6S] \times S^2 \subseteq C_{\alpha_i}$, the derivatives of u_i restricted to $\partial[M, 6S - M] \subseteq C_{\alpha_i}$ are bounded by δ (the constant in Proposition B.11.1). Such an M exists because of our earlier observation that $u_i \rightarrow u$ in C^∞ on compact sets away from the nodes, and we can certainly choose a large M such that the derivatives of u over $\{M\} \times S^1 \subseteq [0, \infty) \times S^1 \subseteq C_0$ (any end) are arbitrarily small.

Let us define $\kappa_i \in K$ by the property that $\xi_i \in (\kappa_i)_{\alpha_i} + \text{im } Q_{\alpha_i, y_i}$. Now since $\|\xi_i\|_{k,2,\delta} \rightarrow 0$ and the norms on the left and right of (B.10.6) are uniformly commensurable, we see that $\kappa_i \rightarrow 0$. Now the uniqueness in Proposition B.9.2 shows that $\xi_i = (\kappa_i)_{\alpha_i, y_i}$ for i sufficiently large. Thus we are done. □

B.12 Conclusion of the proof

Lemma B.12.1 *Let X be a Hausdorff topological space and let $f: \mathbb{R}^n \rightarrow X$ be continuous. Suppose that:*

- (i) *f is injective.*
- (ii) *The restriction of f to any neighborhood of zero is surjective onto a neighborhood of $f(0)$.*

Then there is an open set $0 \in U \subseteq \mathbb{R}^n$ such that $f(U) \subseteq X$ is open and $f: U \xrightarrow{\sim} f(U)$ is a homeomorphism.

Proof Let $B(1) \subseteq \mathbb{R}^n$ denote the closed unit ball and $B(1)^\circ$ its interior. We know that $f(B(1))$ is compact (since f is continuous) and Hausdorff (since X is Hausdorff). We know that $f: B(1) \rightarrow f(B(1))$ is bijective (since f is injective), so it is in fact a homeomorphism.

Now choose an open subset $V \subseteq X$ with $f(0) \in V \subseteq f(B(1)^\circ)$ (which exists by assumption). Set $U = f^{-1}(V) \cap B(1)^\circ$, which is clearly open. Now $f: U \rightarrow f(U)$ is a homeomorphism onto its image (since $U \subseteq B(1)$), and $f(U) = V$ is open. \square

Proof of Theorem B.1.1(i)–(iii) We have shown that the map $g := (B.10.3)$ is continuous, injective, and that its restriction to any neighborhood of zero is surjective onto a neighborhood of the image of zero. The target $\overline{\mathcal{M}}(X)$ is Hausdorff, and thus it follows from Lemma B.12.1 that for some open neighborhood of zero $U \subseteq \mathbb{C}^d \times \mathbb{R}^{\dim \overline{\mathcal{M}}^d} \times K$, we have that $g(U)$ is open and $g: U \xrightarrow{\sim} g(U)$ is a homeomorphism. Since g respects the projection map from both sides to $\overline{\mathcal{M}} \times E/E'$, we obtain the desired conclusions (i), (ii) and (iii). \square

B.13 Gluing orientations

To show Theorem B.1.1(iv) (the statement about orientations), observe that since $\overline{\mathcal{M}}(X)^{\text{reg}} \rightarrow \overline{\mathcal{M}}$ is a submersion, and $\overline{\mathcal{M}}$ is canonically oriented as a complex manifold, it follows that we can canonically identify $\mathfrak{o}_{\overline{\mathcal{M}}(X)^{\text{reg}}}$ with the orientation sheaf of the fibers of $\overline{\mathcal{M}}(X)^{\text{reg}} \rightarrow \overline{\mathcal{M}}$. Now orienting a fiber is the same as orienting the kernel K . It is standard to see that

$$\mathfrak{o}_K = \mathfrak{o}_E \otimes \bigotimes_{i=1}^{\ell} \mathfrak{o}_{u(p_{n+i})}^\vee \otimes N_{D/X}^*$$

This argument gives an identification of $\mathfrak{o}_{\overline{\mathcal{M}}(X)^{\text{reg}}}$ with

$$\mathfrak{o}_E \otimes \bigotimes_{i=1}^{\ell} \mathfrak{o}_{u(p_{n+i})}^\vee \otimes N_{D/X}^*$$

at every point in $\overline{\mathcal{M}}(X)^{\text{reg}}$. Via the gluing map, we get an identification of $\mathfrak{o}_{\overline{\mathcal{M}}(X)^{\text{reg}}}$ with

$$\mathfrak{o}_E \otimes \bigotimes_{i=1}^{\ell} \mathfrak{o}_{u(p_{n+i})^* N_{D/X}}^{\vee}$$

over a small neighborhood of every point in $\overline{\mathcal{M}}(X)^{\text{reg}}$. It is probably straightforward to check that these identifications agree on overlaps (this is a concrete question about whether the projection map $K \rightarrow K_{\alpha} \rightarrow \ker D_{\alpha,y}$ is orientation-preserving; a similar question is dealt with in Floer and Hofer [25]).

This is the best way to prove the desired result, but we can actually get away with a less technical argument. Namely, it is easy to see that the identifications induced by points with smooth domain curve agree on overlaps (since then there is no gluing and we have a nice smooth Banach bundle picture). Since the nodal locus is codimension-two inside $\overline{\mathcal{M}}(X)^{\text{reg}}$ (this follows from the gluing map constructed previously), the identification

$$\mathfrak{o}_{\overline{\mathcal{M}}(X)^{\text{reg}}} = \mathfrak{o}_E \otimes \bigotimes_{i=1}^{\ell} \mathfrak{o}_{u(p_{n+i})^* N_{D/X}}^{\vee}$$

away from the nodal locus extends uniquely to all of $\overline{\mathcal{M}}(X)^{\text{reg}}$.

Appendix C: Gluing for implicit atlases on Hamiltonian Floer moduli spaces

In this appendix, we supply the gluing analysis which was quoted in Section 10 to justify our assertions that the implicit atlases constructed there satisfy the openness and submersion axioms and to identify their orientation local systems (specifically, we prove Propositions 10.3.3 and 10.6.2). Our arguments here are very similar to those used in Appendix B to prove the analogous results for the implicit atlases on Gromov–Witten moduli spaces, and Appendix B should be read first. As in Appendix B, our work here is little more than an appropriate combination of already existing techniques.

C.1 Setup for gluing

The goal of this subsection is to reduce Propositions 10.3.3 and 10.6.2⁵¹ to a single concrete gluing statement concerning a new moduli space $\overline{\mathcal{M}}(M)$. We then spend the rest of the appendix proving this statement.

⁵¹The moduli spaces in Proposition 10.6.2 carry a natural S^1 -action, and following our arguments carefully may lead to the construction of an S^1 -equivariant gluing map. However we do not need such an S^1 -equivariant gluing map, so henceforth we will ignore this S^1 -action.

Fix a symplectic manifold (M, ω) , an integer $n \geq 0$, and a simplex $\sigma \in JH_n(M)$, ie maps $H: \Delta^n \rightarrow C^\infty(M \times S^1)$ and $J: \Delta^n \rightarrow J(M)$ as in Definition 10.1.3. Also fix $\gamma^- \in \mathcal{P}_{H(0)}$ and $\gamma^+ \in \mathcal{P}_{H(n)}$.

Since Propositions 10.3.3 and 10.6.2 are local statements, it suffices to prove them in a neighborhood of a given point. Thus fix a basepoint $(C_0, \ell_0, u_0, \{\phi_0^\beta\}_{\beta \in I}, \{e_0^\beta\}_{\beta \in I}) \in \overline{\mathcal{M}}(\sigma, \gamma^-, \gamma^+)_I^{\leq 5}$ in the setup of either Proposition 10.3.3 or 10.6.2. This point consists (in particular) of the following data:

- (i) C_0 , a Riemann surface of genus zero with marked points $q^-, q^+ \in C_0$. Note that C_0 may be written uniquely as $\coprod_{i=1}^k C_0^{(i)} / \sim$, where each $C_0^{(i)}$ is a nodal Riemann surface with marked points q_{i-1}, q_i (contained in the same irreducible component), the equivalence relation \sim identifies $q_i \in C_0^{(i)}$ with $q_i \in C_0^{(i+1)}$, and $q^- = q_0, q^+ = q_k$. We call any irreducible component containing some q_i a *main component*, and we call all other irreducible components *bubble components*. The nodes other than $q_1, \dots, q_{k-1} \in C_0$ will be called *bubble nodes*.
- (ii) $\ell_0 = \{\ell_0^i: \mathbb{R} \rightarrow \Delta^n\}_{1 \leq i \leq k}$, where ℓ_0^i is a (possibly constant) Morse flow line (for the flow from Definition 10.1.4) from vertex v_{i-1} to vertex v_i , for vertices $0 = v_0 \leq \dots \leq v_k = n$.
- (iii) $A_0 = \{A_0^i: C_0^{(i)} \setminus \{q_{i-1}, q_i\} \rightarrow S^1 \times \mathbb{R}\}_{1 \leq i \leq k}$, where A_0^i is holomorphic, sends q_{i-1} and q_i to $-\infty$ and ∞ , respectively, and restricts to a biholomorphism on the main component of $C_0^{(i)}$. Recall that we always use $(t, s) \in S^1 \times \mathbb{R}$ as coordinates, and that we equip $S^1 \times \mathbb{R}$ (and any subset thereof) with the standard complex structure $z = e^{s+it}$. We may identify each main component of C_0 with $S^1 \times \mathbb{R}$ via A_0 , and thus we have coordinates (t, s) on each such component.
- (iv) $u_0 = \{u_0^i: C_0^{(i)} \setminus \{q_{i-1}, q_i\} \rightarrow M\}_{1 \leq i \leq k}$, where u_0^i is a smooth map with finite energy, converging to $u_0^i(t, -\infty) = \gamma_{i-1}(t)$ and $u_0^i(t, \infty) = \gamma_i(t)$, for periodic orbits $\gamma_i \in \mathcal{P}_{H(v_i)}$ with $\gamma_0 = \gamma^-$ and $\gamma_k = \gamma^+$.
- (v) $e_0 := \bigoplus_{\beta \in I} e_0^\beta \in \bigoplus_{\beta \in I} E_\beta = E_I =: E$.

We require that

$$(C.1.1) \quad \left(du_0 + 2d(\text{proj}_{S^1} A_0) \otimes X_{H((\ell_0 \times \text{id}_{S^1})(A_0(\cdot)))}(u_0) + \sum_{\beta \in I} \lambda_\beta (e_0^\beta)(\phi_0^\beta(\cdot), u_0(\cdot)) \right)_{J(\ell_0(A_0(\cdot)))}^{0,1} = 0.$$

There are a few differences (of an essentially self-explanatory nature) between our notation here and the notation from Section 10. Note that A_0 is holomorphic even

when $e_0 \neq 0$, since λ_β lands in

$$\Omega_{\bar{C}_{0,2+r_\beta}/\bar{M}_{0,2+r_\beta}}^{0,1} \otimes_{\mathbb{R}} TM,$$

as opposed to

$$\Omega_{\bar{C}_{0,2+r_\beta}/\bar{M}_{0,2+r_\beta}}^{0,1} \otimes_{\mathbb{R}} T(M \times S^1 \times \mathbb{R}).$$

We now proceed to formulate an alternative description of the moduli space of interest, $\bar{M}(\sigma, \gamma^-, \gamma^+)_{\bar{I}}^{\leq s}$, in a neighborhood of $(C_0, \ell_0, u_0, \{\phi_0^\beta\}_{\beta \in I}, \{e_0^\beta\}_{\beta \in I})$. This description will be tailored specifically for the present goal of proving a gluing theorem, and we will give a more manageable repackaging of the set of thickening datums $\beta \in I$.

C.1.1 Points $\{x_0^i\}_{1 \leq i \leq r}$ and submanifolds $\{D_i\}_{1 \leq i \leq r} \subseteq M \times S^1 \times \Delta^n$ Let us consider the intersections $\{x_0^{\beta,i}\}_{\beta \in I, 1 \leq i \leq r_\beta}$ of $(\text{id}_{M \times S^1} \times \ell_0)(u_0 \times A_0)|_{(C_0)_\beta}$ with D_β (recall from Definition 10.3.2(iii) that $(C_0)_\beta \subseteq C_0$ denotes the union of the $C_0^{(i)}$ corresponding to β , and that there are exactly r_β such intersections, all of which are transverse). We immediately reindex these intersection points as $\{x_0^i\}_{1 \leq i \leq r}$ (defining $r := \sum_{\beta \in I} r_\beta$). Now

$$D_\beta \subseteq M \times S^1 \times \Delta^{[i_0 \dots i_m]}$$

(for $[i_0 \dots i_m]$ dictated by β), and a neighborhood of $\Delta^{[i_0 \dots i_m]} \subseteq \Delta^n$ may be naturally parametrized by

$$\Delta^{[i_0 \dots i_m]} \times [0, \epsilon)^{\{1, \dots, n\} \setminus \{i_0, \dots, i_m\}}.$$

Let

$$\tilde{D}_\beta := (D_\beta \setminus \partial^{\text{ess}} D_\beta) \times [0, \epsilon)^{\{1, \dots, n\} \setminus \{i_0, \dots, i_m\}} \subseteq M \times S^1 \times \Delta^n,$$

and define $D_1, \dots, D_r \subseteq M \times S^1 \times \Delta^n$ as the reindexing of $\{\tilde{D}_\beta\}_{\beta \in I, 1 \leq i \leq r_\beta}$ corresponding to our earlier reindexing of $\{x_0^{\beta,i}\}_{\beta \in I, 1 \leq i \leq r_\beta}$ to $\{x_0^i\}_{1 \leq i \leq r}$ (thus D_1, \dots, D_r contains r_β copies of \tilde{D}_β). In particular, $(\text{id}_{M \times S^1} \times \ell_0)((u_0 \times A_0)(x_0^i)) \in D_i$ with transverse intersection (this intersection is transverse, rather than merely “transverse when viewed inside $M \times S^1 \times \Delta^{[i_0 \dots i_m]}$ ”, since we replaced D_β with \tilde{D}_β ; we made this replacement purely for the sake of this linguistic convenience).

C.1.2 Points $\{p_i\}_{1 \leq i \leq L}, \{p'_i\}_{1 \leq i \leq L'} \in C_0$ and submanifolds $D, H \subseteq M$ We claim that every unstable bubble component of C_0 has a point where du_0 is injective. Indeed, on any unstable bubble component which is not contained in $(C_0)_\beta$ for any $\beta \in I$, the map u_0 is a (nonconstant!) $J(\ell_0(A_0(\cdot)))$ -holomorphic sphere in M , and thus has such a point of injectivity. Any unstable bubble component contained in $(C_0)_\beta$ for some $\beta \in I$ is stabilized by its intersections with the corresponding D_β . Since these intersections are transverse (and A_0 is constant on the bubble), it follows that du_0

has a point of injectivity on such bubble components as well. Hence the claim is valid. Let us now choose points $p_1, \dots, p_L \in C_0$ (a minimal set of distinct non-nodal points stabilizing all bubble components of C_0) and a codimension-two submanifold with boundary $D \subseteq M$ such that $u_0(p_i) \in D$ with transverse intersection for $1 \leq i \leq L$.

We claim that every unstable main component of C_0 over which ℓ_0 is constant has a point where $(\partial/\partial s)u_0$ is nonzero. Indeed, this holds because such an unstable main component cannot be a trivial cylinder (otherwise the trajectory would be unstable). Let us now choose points $p'_1, \dots, p'_{L'} \in C_0$ (consisting of exactly one non-nodal point in every unstable main component over which ℓ_0 is constant) and a codimension-one submanifold with boundary $H \subseteq M$ such that $u_0(p'_i) \in H$ with $(\partial/\partial s)u_0$ transverse to H for $1 \leq i \leq L'$. We assume that D and H are disjoint: this can be achieved by first perturbing D (and correspondingly p_1, \dots, p_L) so that on each of the main components under consideration here, there exists a point where $(\partial/\partial s)u_0$ is nonzero and which is not mapped to D .

C.1.3 Gluing $C_0, \coprod_{i=1}^k \mathbb{R}$ and varying j_0, A_0, ℓ_0 A subset of the Riemann sphere $\mathbb{C} \cup \{\infty\}$ homeomorphic to S^1 is called a *circle* if and only if its intersection with \mathbb{C} is either a straight line or a circle (in the usual sense). It is well-known that this notion is invariant under biholomorphisms of the Riemann sphere. For any Riemann surface C biholomorphic to the Riemann sphere minus finitely many points, a subset of C homeomorphic to S^1 is called a *circle* if and only if its image under such a biholomorphism is a circle.

On both sides of every bubble node of C_0 , fix cylindrical ends

$$(C.1.2) \quad S^1 \times [0, \infty) \rightarrow C_0$$

which are *circular*, in the sense that every (equivalently, some) cross-section $S^1 \times \{s\}$ is a circle (inside the corresponding irreducible component of C_0). We will call the ends (C.1.2) the *bubble ends*. We also fix some large $N < \infty$, and we call the subsets $S^1 \times [N, \infty)$ and $S^1 \times (-\infty, -N]$ of the main components of C_0 the *main ends* (which come in two types: *positive* and *negative*). At a few later points in the gluing argument, we will need to assume that the bubble ends were chosen sufficiently small and that N was chosen sufficiently large.

We now fix a smooth family of (necessarily integrable) almost complex structures j_y on C_0 and a smooth family of j_y -holomorphic maps

$$(C.1.3) \quad A_y^i: C_0^{(i)} \setminus \{q_{i-1}, q_i\} \rightarrow S^1 \times \mathbb{R}$$

parametrized by $y \in \mathbb{R}^*$ (for an integer $*$ ≥ 0 to be specified shortly), specializing to (j_0, A_0) at $y = 0$ (where j_0 denotes the given almost complex structure on C_0). We

require that this family $(j_y, A_y)_{y \in \mathbb{R}^*}$ satisfy the following conditions. The family j_y must be constant over the bubble ends, and A_y (and hence j_y) must be constant over the main ends. The bubble ends (C.1.2) must also be circular with respect to every j_y . Now, for each stable bubble component of C_0 (say, with $\nu \geq 3$ special points), the family j_y induces a map $\mathbb{R}^* \rightarrow \mathcal{M}_{0,\nu}$. For each main component of C_0 (say, with $\nu \geq 0$ bubble nodes), the family A_y induces a map $\mathbb{R}^* \rightarrow (S^1 \times \mathbb{R})^\nu$ (namely taking the images of the bubble nodes under A_y). We require that the resulting map

$$\mathbb{R}^* \rightarrow \prod_{\substack{\text{bubble components} \\ \text{with } \nu \geq 3}} \mathcal{M}_{0,\nu} \times \prod_{\substack{\text{main components} \\ \text{with } \ell_0 \text{ non-constant}}} (S^1 \times \mathbb{R})^\nu \times \prod_{\substack{\text{main components} \\ \text{with } \ell_0 \text{ constant} \\ \text{and } \nu \geq 1}} (S^1 \times \mathbb{R})^\nu / \mathbb{R}$$

be a diffeomorphism onto its image (in particular, this determines $* \in \mathbb{Z}_{\geq 0}$). We fix, once and for all, a family $(j_y, A_y)_{y \in \mathbb{R}^*}$ satisfying these properties (such a family always exists).

Given a set of gluing parameters⁵² $\alpha \in \mathbb{C}^d \times \mathbb{R}_{\geq 0}^{k-1}$ (where d is the number of bubble nodes of C_0), we may glue C_0 to obtain a curve C_α as follows. At a bubble node with gluing parameter $\alpha = e^{-6S+i\theta} \in \mathbb{C}$, we truncate both bubble ends (C.1.2) from $S^1 \times [0, \infty)$ to $S^1 \times [0, 6S]$, and we identify them via the map $s' = 6S - s$ and $t' = \theta - t$ (if $\alpha = 0$, then we do nothing). At each main node q_i with gluing parameter $\alpha = e^{-6S} \in \mathbb{R}_{\geq 0}$, we truncate the main ends $S^1 \times [N, \infty) \subseteq C_0^{(i)}$ and $S^1 \times (-\infty, -N] \subseteq C_0^{(i+1)}$ incident at q_i to $S^1 \times [N, 6S - N]$ and $S^1 \times [-6S + N, -N]$, and we identify them via $s' = s - 6S$ and $t' = t$ (if $\alpha = 0$, then we do nothing). The curve C_α now has *main ends* and *bubble ends* (those where no gluing was performed), as well as *bubble necks* $S^1 \times [0, 6S]$ and *main necks* $S^1 \times [N, 6S - N]$.

We perform a similar gluing operation to $\coprod_{i=1}^k \mathbb{R}$, and we denote the result by $(\coprod_{i=1}^k \mathbb{R})_\alpha$. Namely, for the gluing parameter $\alpha = e^{-6S} \in \mathbb{R}_{\geq 0}$ associated to q_i , we truncate the i^{th} copy of \mathbb{R} to $(-\infty, 6S]$, we truncate the $(i + 1)^{\text{st}}$ copy of \mathbb{R} to $[-6S, \infty)$, and we identify $[0, 6S]$ in the i^{th} copy with $[-6S, 0]$ in the $(i + 1)^{\text{st}}$ copy via $s' = s - 6S$ (if $\alpha = 0$, then we do nothing).

The almost complex structure j_y clearly descends to C_α , since it is constant over the ends of C_0 . Furthermore, the maps (C.1.3) induce j_y -holomorphic maps

$$(C.1.4) \quad A_y: C_\alpha \setminus \{q_0, \dots, q_k\} \rightarrow S^1 \times \left(\prod_{i=1}^k \mathbb{R} \right)_\alpha$$

⁵²We only care about what happens for gluing parameters α in a neighborhood of zero, and many constructions will only make sense (and many statements will only be true) for sufficiently small gluing parameters α , even though this assumption is not always explicitly stated. The same goes for the parameters w (introduced below) and y .

characterized uniquely by the property that they agree with (C.1.3) over the images of the truncated main components of C_0 (the existence of such a map A_y follows from our assumption that the bubble ends (C.1.2) are circular with respect to every j_y). The points $p_1, \dots, p_L, p'_1, \dots, p'_{L'} \in C_0$ clearly descend to points $p_1, \dots, p_L, p'_1, \dots, p'_{L'} \in C_\alpha$, and there is a node $q_i \in C_\alpha$ whenever the corresponding gluing parameter is zero.

Eventually, only a subset of $\mathbb{C}^d \times \mathbb{R}_{\geq 0}^{k-1}$ will be relevant for us, namely the subset

$$\mathbb{C}^d \times (\mathbb{R}_{\geq 0}^{k-1})^{\leq s} \subseteq \mathbb{C}^d \times \mathbb{R}_{\geq 0}^{k-1}$$

cut out by the requirement that if the break in the trajectory at (v_i, γ_i) is dictated by the stratum \mathfrak{s} , then the gluing parameter at q_i is zero.

For every nonconstant ℓ_0^i , choose a local hypersurface $H^i \subseteq \Delta^{[v_{i-1} \dots v_i]}$ transverse to ℓ_0^i at $\ell_0^i(0)$. For $w \in \prod_{i=1}^{k'} H^i$ (where $\prod_{i=1}^{k'}$ denotes the product over those i for which ℓ_0^i is nonconstant), let $\ell_w: \prod_{i=1}^k \mathbb{R} \rightarrow \Delta^n$ denote the unique trajectory satisfying $\prod_{i=1}^{k'} \ell_w^i(0) = w$ and broken at the same sequence of vertices $0 = v_0 \leq \dots \leq v_k = n$. We let $0 := \prod_{i=1}^{k'} \ell_0^i(0) \in \prod_{i=1}^{k'} H^i$, so in this notation $\ell_w = \ell_0$ for $w = 0$.

Now choose a neighborhood parametrization

$$\Delta^{[v_{i-1} \dots v_i]} \times [0, \epsilon)^{\{1, \dots, n\} \setminus \{v_{i-1}, \dots, v_i\}} \rightarrow \Delta^n,$$

and define a local hypersurface

$$\bar{H}^i := H^i \times [0, \epsilon)^{\{1, \dots, n\} \setminus \{v_{i-1}, \dots, v_i\}} \subseteq \Delta^n$$

and a projection map $\bar{H}^i \rightarrow H^i$. For $\alpha \in \mathbb{R}_{\geq 0}^{k-1}$, let $\ell_{\alpha, w}: (\prod_{i=1}^k \mathbb{R})_\alpha \rightarrow \Delta^n$ denote the trajectory characterized uniquely by the property that $\prod_{i=1}^{k'} \ell_{\alpha, w}(0_i) \in \prod_{i=1}^{k'} \bar{H}^i$ projects to $w \in \prod_{i=1}^{k'} H^i$, where $0_i \in (\prod_{i=1}^k \mathbb{R})_\alpha$ denotes the image of 0 in the i^{th} copy of \mathbb{R} . In this notation, $\ell_{\alpha, w} = \ell_0$ for $\alpha = 0$ and $w = 0$. The existence and uniqueness of $\ell_{\alpha, w}$ for (α, w) close to $(0, 0)$ follows simply by explicitly integrating the Morse flow on Δ^n from Definition 10.1.4.

Eventually, only a subset of $\prod_{i=1}^{k'} H^i$ will be relevant for us, namely the subset $\prod_{i=1}^{k'} (H^i)^{\leq s} \subseteq \prod_{i=1}^{k'} H^i$, where $(H^i)^{\leq s} \subseteq H^i$ is the intersection of H^i with

$$\Delta^{[i_0 \dots i_m]} \cap \Delta^{[v_{i-1} \dots v_i]} = \Delta^{[i_0 \dots i_m] \cap [v_{i-1} \dots v_i]},$$

where $(\sigma[[i_0 \dots i_m], \gamma_a, \gamma_b)$ is the part of the stratum \mathfrak{s} containing ℓ_0^i .

C.1.4 Linear map λ Recall that for all $\beta \in I$, we have a linear map

$$(C.1.5) \quad \lambda_\beta: E_\beta \rightarrow C^\infty(\bar{\mathbb{C}}_{0,2+r_\beta} \times M, \Omega_{\bar{\mathbb{C}}_{0,2+r_\beta}/\bar{\mathcal{M}}_{0,2+r_\beta}}^{0,1} \otimes_{\mathbb{R}} TM).$$

We now repackage these λ_β into a single linear map

$$(C.1.6) \quad \lambda: E \rightarrow C^\infty(\mathbb{C}^d \times (\mathbb{R}_{\geq 0}^{k-1})^{\leq s} \times \mathbb{R}^* \times C_0^r \times C_0 \times M, \Omega_{C_0}^1 \otimes_{\mathbb{R}} TM)$$

which vanishes over the ends of (the last factor of) C_0 . We will actually only define λ in a small neighborhood of $\{0\} \times \{0\} \times \{x_0^1\} \times \dots \times \{x_0^r\} \times C_0 \times M$, as this is the only part of the domain which will be relevant.

Suppose we are given $e \in E = \bigoplus_{\beta \in I} E_\beta$, $\alpha \in \mathbb{C}^d \times (\mathbb{R}_{\geq 0}^{k-1})^{\leq s}$ and $y \in \mathbb{R}^*$, along with points $x^1, \dots, x^r \in C_0$, each in a small neighborhood of the corresponding x_0^i ; let us define $\lambda(e)(\alpha, y, x^1, \dots, x^r, \cdot, \cdot)$ as a section

$$C_0 \times M \rightarrow \Omega_{C_0, j_y}^{0,1} \otimes_{\mathbb{R}} TM \subseteq \Omega_{C_0}^1 \otimes_{\mathbb{R}} TM.$$

Consider the glued curve C_α equipped with the almost complex structure j_y and the marked points $x^1, \dots, x^r \in C_\alpha$ (descended from C_0). Since $\alpha \in \mathbb{C}^d \times (\mathbb{R}_{\geq 0}^{k-1})^{\leq s}$, there is a subcurve $(C_\alpha)_\beta \subseteq C_\alpha$ corresponding to any given $\beta \in I$. If we equip it with the r_β marked points from x^1, \dots, x^r corresponding to β (with respect to the reindexing from Section C.1.1), then this induces a unique map $\phi^\beta: (C_\alpha)_\beta \rightarrow \bar{C}_{0,2+r_\beta}$ (isomorphism onto a fiber) which is close to the given map $\phi_0^\beta: (C_0)_\beta \rightarrow \bar{C}_{0,2+r_\beta}$ (this assumes we are in the setting of Proposition 10.3.3; in the setting of Proposition 10.6.2, we instead have a map $\phi^\beta: ((C_\alpha)_\beta)^{st} \rightarrow \bar{C}_{0,2+r_\beta}$ close to the given map $\phi_0^\beta: ((C_0)_\beta)^{st} \rightarrow \bar{C}_{0,2+r_\beta}$, where $(C_\alpha)_\beta \rightarrow ((C_\alpha)_\beta)^{st}$ contracts those components all of whose constituent components of $(C_0)_\beta$ were contracted by $(C_0)_\beta \rightarrow ((C_0)_\beta)^{st}$). Now the pullback of $\lambda_\beta(\text{proj}_{E_\beta} e)$ under ϕ^β gives us a section

$$C_\alpha \times M \rightarrow \Omega_{C_\alpha, j_y}^{0,1} \otimes_{\mathbb{R}} TM$$

(defined to be zero at those points not contained in $(C_\alpha)_\beta$). We may assume without loss of generality that the ends of C_0 were chosen small enough so that this section vanishes over the ends/necks of C_α , and hence gives rise to a well-defined lift to a section

$$C_0 \times M \rightarrow \Omega_{C_0, j_y}^{0,1} \otimes_{\mathbb{R}} TM.$$

We declare $\lambda(e)(\alpha, y, x^1, \dots, x^r, \cdot, \cdot)$ to be the sum of these sections over all $\beta \in I$. This defines the function λ , which is indeed of class C^∞ since the partially defined composition $C_0 \rightarrow C_\alpha \rightarrow (C_\alpha)_\beta \rightarrow \bar{C}_{0,2+r_\beta}$ depends smoothly on $\alpha, y, x^1, \dots, x^r$.

C.1.5 A new moduli space $\bar{\mathcal{M}}(M)$ Let $\bar{\mathcal{M}}(M)$ denote the moduli space of tuples $(\alpha, w, y, u, e, \{x^i\}_{1 \leq i \leq r})$, where:

- (i) $\alpha \in \mathbb{C}^d \times (\mathbb{R}_{\geq 0}^{k-1})^{\leq s}$.
- (ii) $w \in \prod_{i=1}^{k'} (H^i)^{\leq s}$.
- (iii) $y \in \mathbb{R}^*$.
- (iv) $u: C_\alpha \setminus \{q_0, \dots, q_k\} \rightarrow M$ is smooth with finite energy, is asymptotic to γ_i at $q_i \in C_\alpha$ (whenever the gluing parameter α at $q_i \in C_0$ is zero), and $u(p_i) \in D^\circ$ ($1 \leq i \leq L$) and $u(p'_i) \in H^\circ$ ($1 \leq i \leq L'$).
- (v) $e \in E$.
- (vi) $x^i \in C_\alpha$ (not nodes) satisfy $(\text{id}_{M \times S^1} \times \ell_{\alpha,w})(u \times A_y)(x^i) \in D_i$ with transverse intersection ($1 \leq i \leq r$).
- (vii) We require that

$$(C.1.7) \quad (du + 2d(\text{proj}_{S^1} A_y) \otimes X_{H((\ell_{\alpha,w} \times \text{id}_{S^1})(A_y(\cdot)))}) + \lambda(e)(\alpha, y, x^1, \dots, x^r, \cdot, u(\cdot)))_{j_y, J(\ell_{\alpha,w}(A_y(\cdot)))}^{0,1} = 0.$$

We equip $\overline{\mathcal{M}}(M)$ with the topology of uniform convergence. More precisely, let $\overline{\mathcal{M}} := \mathbb{C}^d \times (\mathbb{R}_{\geq 0}^{k-1})^{\leq s}$, and let $\overline{\mathcal{C}}^S \rightarrow \overline{\mathcal{M}}$ be the bundle whose fiber over α is C_α^S , obtained from C_α by replacing each q_i ($0 \leq i \leq k$) with a copy of S^1 (thus u as above extends continuously to C_α^S and equals $\gamma_i(t)$ on the S^1 over q_i whenever the corresponding α is zero). We equip $\overline{\mathcal{M}}(M)$ with the obvious topology on $w \in \prod_{i=1}^{k'} (H^i)^{\leq s}$, $y \in \mathbb{R}^*$, $e \in E$, $\{x^i\}_{1 \leq i \leq r} \in \overline{\mathcal{C}}^r$, and the Hausdorff topology on the graph of u inside $\overline{\mathcal{C}}^S \times M$.

Note that there is a distinguished basepoint $(0, 0, 0, u_0, e_0, \{x_0^i\}_{1 \leq i \leq r}) \in \overline{\mathcal{M}}(M)$. Furthermore, a neighborhood of this basepoint in $\overline{\mathcal{M}}(M)$ is canonically homeomorphic to a neighborhood of the given basepoint in $\overline{\mathcal{M}}(\sigma, \gamma^-, \gamma^+)^{\leq s}$ (as long as the given basepoint in $\overline{\mathcal{M}}(\sigma, \gamma^-, \gamma^+)^{\leq s}$ has trivial automorphism group). Note that to justify this statement in the setting of Proposition 10.6.2, we must appeal to Lemma 10.5.4.

C.1.6 The regular locus $\overline{\mathcal{M}}(M)^{\text{reg}}$ We now define a subset $\overline{\mathcal{M}}(M)^{\text{reg}} \subseteq \overline{\mathcal{M}}(M)$, depending on a choice of subspace $E' \subseteq E$. Fix a point $(\alpha, w, y, u_1, e_1, \{x_1^i\}_{1 \leq i \leq r}) \in \overline{\mathcal{M}}(M)$, and we will describe when it lies in $\overline{\mathcal{M}}(M)^{\text{reg}}$.

Fix an integer $k \in \mathbb{Z}_{\geq 6}$ (not to be confused with the number k of main components of C_0) and a small real number $\delta \in (0, 1)$ (we will be precise about how small δ must be shortly). Let $W^{k,2,\delta}(C_\alpha, M)$ denote the smooth Banach manifold consisting of maps $u: C_\alpha \setminus \{q_0, \dots, q_k\} \rightarrow M$ which are of locally of class $W^{k,2}$ such that for all positive main ends of C_α with corresponding node q_i , we have

$$(C.1.8) \quad \int_{S^1 \times [N', \infty)} \sum_{j=0}^k |D^j[\exp_{\gamma_i}^{-1} u(t, s)]|^2 e^{2s\delta} dt ds < \infty$$

for sufficiently large $N' < \infty$ (along with the analogous condition over negative main ends); this definition is independent of the choice of metric and connection used in (C.1.8). Over $W^{k,2,\delta}(C_\alpha, M)$, we consider the smooth Banach bundle \mathcal{E} whose fiber over $u: C_\alpha \rightarrow M$ is

$$W^{k-1,2,\delta}(\tilde{C}_\alpha, \Omega_{\tilde{C}_\alpha, j_y}^{0,1} \otimes_{\mathbb{C}} u^* TM_{J(\ell_{\alpha,w}(A_y(\cdot)))}),$$

namely the space of sections

$$\eta: \tilde{C}_\alpha \rightarrow \Omega_{\tilde{C}_\alpha, j_y}^{0,1} \otimes_{\mathbb{C}} u^* TM_{J(\ell_{\alpha,w}(A_y(\cdot)))}$$

which are locally of class $W^{k-1,2}$ and which satisfy

$$(C.1.9) \quad \int_{S^1 \times [N', \infty)} \sum_{j=0}^{k-1} |D^j [\text{PT}_{u(t,s) \rightarrow \gamma_i(t)} \eta(t, s)]|^2 e^{2s\delta} dt ds < \infty$$

over positive main ends (along with the analogous condition over negative main ends); this definition is independent of the choice of metric and connection used in (C.1.9).

It is well-known (see Proposition C.11.1 below) that for every non-degenerate periodic orbit γ of a smooth Hamiltonian $H: M \times S^1 \rightarrow \mathbb{R}$ and every smooth ω -compatible almost complex structure J on M , there exists $\delta_{H,J,\gamma} > 0$ such that for every $u: S^1 \times [0, \infty) \rightarrow M$ of finite energy satisfying $(du + 2dt \otimes X_H)_J^{0,1} = 0$ and asymptotic to $u(t, \infty) = \gamma(t)$, all derivatives of u decay like $e^{-\delta s}$ for any $\delta < \delta_{H,J,\gamma}$ (precisely, $\delta_{H,J,\gamma}$ is the smallest magnitude of any eigenvalue of the asymptotic linearized operator, which depends only on ω, H, J , and γ), and likewise for every $u: S^1 \times (-\infty, 0] \rightarrow M$ satisfying $(du + 2dt \otimes X_H)_J^{0,1} = 0$ and asymptotic to $u(t, -\infty) = \gamma(t)$. In particular, for $0 \leq i \leq k$, there is a corresponding $\delta_i := \delta_{H(v_i), J(v_i), \gamma_i} > 0$. We fix $\delta \in (0, 1)$ with $\delta < \delta_i$ for all i ; thus $u_1 \in W^{k,2,\delta}(C_\alpha, M)$.

Now since $k \geq 6$, there are unique continuous functions

$$(C.1.10) \quad x^i: W^{k,2,\delta}(C_\alpha, M) \rightarrow C_\alpha \quad (1 \leq i \leq r)$$

defined for u in a neighborhood of $u_1 \in W^{k,2,\delta}(C_\alpha, M)$, coinciding with the given $x_1^i \in C_\alpha$ at u_1 , and which satisfy $(\text{id}_{M \times S^1} \times \ell_{\alpha,y})(u \times A_y)(x^i(u)) \in D^i$ (the intersection is automatically transverse for u close to u_1). Moreover, (C.1.10) are of class C^{k-2} (see the discussion following (B.1.7)). It follows that the left-hand side of (C.1.7) is a C^{k-2} section of \mathcal{E} over $W^{k,2,\delta}(C_\alpha, M) \times E$. By results of Lockhart and McOwen [55], this section is Fredholm for $\delta > 0$ as above (note that over each end, the almost complex structure on M is constant and ω -compatible). Let $W^{k,2,\delta}(C_\alpha, M)_{D,H}$ denote the subspace cut out by the requirements $u(p_i) \in D$ and $u(p'_i) \in H$. We say that the given point $(\alpha, w, y, u_1, e_1, \{x_1^i\}_{1 \leq i \leq r}) \in \overline{\mathcal{M}}(M)$ lies in $\overline{\mathcal{M}}(M)^{\text{reg}}$ if and

only if the section $(u, e) \mapsto (C.1.7) \oplus e$ of $\mathcal{E} \oplus E/E'$ over $W^{k,2,\delta}(C_\alpha, M)_{D,H} \times E$ is transverse to the zero section at (u_1, e_1) (it follows from elliptic regularity theory that this condition is independent of the choice of k, δ as above).

Now suppose that we take $E' := E_{I'}$ for $I' \subseteq I$ and that the given basepoint in $\overline{\mathcal{M}}(\sigma, \gamma^-, \gamma^+)_{I'}^{\leq s}$ was chosen inside the inverse image of $(\overline{\mathcal{M}}(\sigma, \gamma^-, \gamma^+)_{I'}^{\leq s})^{\text{reg}}$. Then, the basepoint of $\overline{\mathcal{M}}(M)$ lies in $\overline{\mathcal{M}}(M)^{\text{reg}}$ (this uses the fact that L, L' were chosen to be minimal), and furthermore $\overline{\mathcal{M}}(M)^{\text{reg}} \cap \text{proj}_{E/E'}^{-1}(0)$ is contained inside the inverse image of $(\overline{\mathcal{M}}(\sigma, \gamma^-, \gamma^+)_{I'}^{\leq s})^{\text{reg}}$.

C.2 Our goal: the gluing map

We henceforth assume that the basepoint in $\overline{\mathcal{M}}(M)$ lies in $\overline{\mathcal{M}}(M)^{\text{reg}}$. The remainder of this appendix is devoted to the construction of a germ of a local chart

$$(C.2.1) \quad \left(\mathbb{C}^d \times (\mathbb{R}_{\geq 0}^{k-1})^{\leq s} \times \mathbb{R}^* \times \prod_{i=1}^{k'} (H^i)^{\leq s} \times K, (0, 0, 0, 0) \right) \rightarrow (\overline{\mathcal{M}}(M), (0, 0, 0, u_0, e_0, \{x_0^i\}_{1 \leq i \leq r}))$$

which respects the natural projection from both sides to

$$\mathbb{C}^d \times (\mathbb{R}_{\geq 0}^{k-1})^{\leq s} \times \mathbb{R}^* \times \prod_{i=1}^{k'} (H^i)^{\leq s} \times E/E'$$

and whose image is contained in $\overline{\mathcal{M}}(M)^{\text{reg}}$ (K denotes the kernel of the linearization of (C.1.7) at the distinguished basepoint of $\overline{\mathcal{M}}(M)$; see (C.5.9)). We will also discuss the compatibility of gluing with orientations, and more generally we will discuss how to define coherent orientations on the moduli spaces of Floer trajectories. Propositions 10.3.3 and 10.6.2 follow from the existence of such a chart (C.2.1) and our earlier observations relating $\overline{\mathcal{M}}(M)$ to $\overline{\mathcal{M}}(\sigma, \gamma^-, \gamma^+)_{I'}^{\leq s}$.

C.3 Pregluing

Let $\text{exp}: TM \rightarrow M$ denote the exponential map of some Riemannian metric on M for which D and H are totally geodesic (such a metric exists since D and H are disjoint). Recall that J is a smooth family of almost complex structures on M parametrized by Δ^n which is constant near the vertices. Let ∇ denote a smooth family of J -linear connections on M parametrized by Δ^n which is constant near the vertices (for instance, we could take $\nabla_X Y := \frac{1}{2}(\nabla_X^0 Y - J(\nabla_X^0(JY)))$ for any fixed connection ∇^0). Let $\text{PT}_{x \rightarrow y}: T_x M \rightarrow T_y M$ denote parallel transport via ∇ along the shortest geodesic between x and y (we will only use this notation when it may be assumed that x and y

are very close in M); note that $\text{PT}_{x \rightarrow y}$ is J -linear, and that $\text{PT}_{x \rightarrow y}$ is a family of maps parametrized by Δ^n .

Fix a smooth function $\chi: \mathbb{R} \rightarrow [0, 1]$ satisfying

$$(C.3.1) \quad \chi(x) = \begin{cases} 0, & x \leq 0, \\ 1, & x \geq 1. \end{cases}$$

Definition C.3.1 (Flattening) For $\alpha \in \mathbb{C}^d \times \mathbb{R}_{\geq 0}^{k-1}$, we define the “flattened” map $u_{0|\alpha}: C_0 \rightarrow M$ as follows. Away from the ends, $u_{0|\alpha}$ coincides with u_0 . Over a bubble end $S^1 \times [0, \infty)$, we define $u_{0|\alpha}$ as

$$(C.3.2) \quad u_{0|\alpha}(t, s) := \begin{cases} u_0(t, s), & s \leq S - 1, \\ \exp_{u_0(n)}[\chi(S - s) \cdot \exp_{u_0(n)}^{-1} u_0(t, s)], & S - 1 \leq s \leq S, \\ u_0(n), & S \leq s, \end{cases}$$

where $n \in C_0$ denotes the corresponding node. Over a positive main end $S^1 \times [N, \infty)$, we define $u_{0|\alpha}$ as

$$(C.3.3) \quad u_{0|\alpha}(t, s) := \begin{cases} u_0(t, s), & s \leq S - 1, \\ \exp_{\gamma(t)}[\chi(S - s) \cdot \exp_{\gamma(t)}^{-1} u_0(t, s)], & S - 1 \leq s \leq S, \\ \gamma(t), & S \leq s, \end{cases}$$

where $\gamma(t) = u_0(t, \infty)$ denotes the corresponding periodic orbit; an analogous definition applies over the negative main ends.

Definition C.3.2 (Pregluing) For $\alpha \in \mathbb{C}^d \times \mathbb{R}_{\geq 0}^{k-1}$, we define the “preglued” map $u_\alpha: C_\alpha \rightarrow M$ as follows. Away from the necks, u_α coincides with u_0 . Over a bubble neck $S^1 \times [0, 6S]$, we define u_α as

$$(C.3.4) \quad u_\alpha(t, s) := \begin{cases} u_0(t, s), & s \leq S - 1, \\ \exp_{u_0(n)}[\chi(S - s) \cdot \exp_{u_0(n)}^{-1} u_0(t, s)], & S - 1 \leq s \leq S, \\ u_0(n), & S \leq s \leq 5S, \\ \exp_{u_0(n)}[\chi(S - s') \cdot \exp_{u_0(n)}^{-1} u_0(t', s')], & 5S \leq s \leq 5S + 1, \\ u_0(t', s'), & 5S + 1 \leq s, \end{cases}$$

and over a main neck $S^1 \times [N, 6S - N]$, we define u_α as

$$(C.3.5) \quad u_\alpha(t, s) := \begin{cases} u_0(t, s), & s \leq S - 1, \\ \exp_{\gamma(t)}[\chi(S - s) \cdot \exp_{\gamma(t)}^{-1} u_0(t, s)], & S - 1 \leq s \leq S, \\ \gamma(t), & S \leq s \leq 5S, \\ \exp_{\gamma(t)}[\chi(S + s') \cdot \exp_{\gamma(t)}^{-1} u_0(t', s')], & 5S \leq s \leq 5S + 1, \\ u_0(t', s'), & 5S + 1 \leq s \end{cases}$$

(u_α should be thought of as the “descent” of $u_{0|\alpha}$ from C_0 to C_α).

C.4 Weighted Sobolev norms

Recall that we have fixed $k \in \mathbb{Z}_{\geq 6}$ and $\delta \in (0, 1)$ smaller than $\delta_i > 0$ for $0 \leq i \leq k$.

We now introduce new weighted Sobolev spaces $W^{k,2,\delta,\delta}$ (with weights over all ends and necks) which we will work with from now on. The specific choice of norms (not just their commensurability classes) on these $W^{k,2,\delta,\delta}$ spaces is crucial.

Definition C.4.1 We define the weighted Sobolev space $W^{k,2,\delta,\delta}(C_\alpha, u_\alpha^*TM)$ using the usual $(k, 2)$ -norm away from the ends/neck, and using the following weighted $(k, 2)$ -norms over the bubble ends/neck and main ends/neck, respectively (we will write the contribution to the norm squared):

$$(C.4.1) \quad |\xi(n)|^2 + \int_{S^1 \times [0, \infty)} \left[|\xi(t, s) - \xi(n)|^2 + \sum_{j=1}^k |D^j \xi(t, s)|^2 \right] e^{2\delta s} dt ds,$$

$$(C.4.2) \quad \left| \int_{S^1} \xi(t, 3S) dt \right|^2 + \int_{S^1 \times [0, 6S]} \left[\left| \xi(t, s) - \int_{S^1} \xi(t, 3S) dt \right|^2 + \sum_{j=1}^k |D^j \xi(t, s)|^2 \right] e^{2\delta \min(s, 6S-s)} dt ds,$$

$$(C.4.3) \quad \int_{S^1 \times [N, \infty)} \sum_{j=0}^k |D^j \xi(t, s)|^2 e^{2\delta(s-N)} dt ds,$$

$$(C.4.4) \quad \int_{S^1 \times [N, 6S-N]} \sum_{j=0}^k |D^j \xi(t, s)|^2 e^{2\delta \min(s-N, 6S-N-s)} dt ds.$$

These are to be interpreted as follows. For each bubble node $n \in C_0$, we fix, once and for all, a small trivialization of TM in a neighborhood of $u_0(n)$; this allows us to view ξ as a function (rather than a section) for the purposes of the integrals over the bubble ends/neck. For each $i = 0, \dots, k$, we fix, once and for all, a smooth family of trivializations of TM near $\gamma_i(t)$ (parametrized by $t \in S^1$); this allows us to view ξ as a section of γ_i^*TM for the purposes of the integrals over the main ends/neck. We also fix, once and for all, a connection on each such bundle γ_i^*TM over S^1 . The derivatives in the integrals above are measured with respect to the standard metric on $S^1 \times \mathbb{R}$. The case of negative main ends is completely analogous to that of positive main ends.

By the Sobolev embedding of $W^{2,2}$ into C^0 in two dimensions, we get uniform bounds linear in $\|\xi\|_{k,2,\delta,\delta}$ on $|\xi(t, s) - \xi(n)|e^{\delta s}$ in the bubble ends, $|\xi(t, s)|e^{\delta s}$ in the main ends, and $|D^j \xi(t, s)|e^{\delta s}$ ($1 \leq j \leq k - 2$) in all ends (as well as similar estimates in the necks).

We will occasionally use other weighted Sobolev spaces (eg $W^{k,2,\delta,\delta}(C_0, u_{0|\alpha}^* TM)$), and we leave it to the reader to make the necessary adjustments to the definition (which is essentially identical to the above).

Remark C.4.2 The particular choice of trivializations and connections in the definition above is not crucial: any other (fixed) choice would lead to a uniformly commensurable norm (this holds because u_0 satisfies the exponential decay estimates (C.6.1), (C.6.3), and because $\delta < 1$ and $\delta < \delta_i$).

Definition C.4.3 We define the weighted Sobolev space

$$W^{k-1,2,\delta,\delta}(\tilde{C}_\alpha, T^*\tilde{C}_\alpha \otimes_{\mathbb{R}} u_\alpha^* TM)$$

using the usual $(k - 1, 2)$ -norm away from the ends/necks, and using the following weighted $(k - 1, 2)$ -norms over the bubble ends/necks (we will write the contribution to the norm squared):

$$(C.4.5) \quad \int_{S^1 \times [0, \infty)} \sum_{j=0}^{k-1} |D^j \eta(t, s)|^2 e^{2\delta s} dt ds,$$

$$(C.4.6) \quad \int_{S^1 \times [0, 6S]} \sum_{j=0}^{k-1} |D^j \eta(t, s)|^2 e^{2\delta \min(s, 6S-s)} dt ds$$

(for the main ends/necks, simply make the obvious replacement of $[0, \infty)$ with $[N, \infty)$ and $e^{2\delta s}$ with $e^{2\delta(s-N)}$, etc.). These are to be interpreted as follows. We trivialize $T\tilde{C}_\alpha$ over any end/neck with the basis vectors $\partial/\partial t, \partial/\partial s$. We trivialize TM as in Definition C.4.1, and hence the section η is simply a pair of functions $\eta = (\eta_1, \eta_2)$.

By the Sobolev embedding of $W^{2,2}$ into C^0 in two dimensions, we get uniform exponential decay bounds on η up to $k - 3$ derivatives in any end/neck, linear in $\|\eta\|_{k-1,2,\delta,\delta}$.

We are actually interested in certain closed subspaces of

$$W^{k-1,2,\delta,\delta}(\tilde{C}_\alpha, T^*\tilde{C}_\alpha \otimes_{\mathbb{R}} u_\alpha^* TM),$$

for example

$$W^{k-1,2,\delta,\delta}(\tilde{C}_\alpha, \Omega_{\tilde{C}_\alpha, j}^{0,1} \otimes_{\mathbb{C}} u_\alpha^* TM_J)$$

for certain almost complex structures j, J on C_α, M respectively, which we equip with the restriction of the norm defined above. We will occasionally use other very similar weighted Sobolev spaces, and we leave it to the reader to make the necessary adjustments to the definition (which is essentially identical to the above).

Henceforth, we will work exclusively with the weighted Sobolev spaces defined above, rather than those from Section C.1.6. The Fredholm index and the kernel/cokernel of the relevant linearized operators are unchanged by the placement of weights in the bubble ends/necks (the argument from Lemma B.5.2 applies without modification).

C.5 Based section $\mathcal{F}_{\alpha,w,y}$ and linearized operator $D_{\alpha,w,y}$

We consider the partially defined function

$$\mathcal{F}_{\alpha,w,y}: C^\infty(C_\alpha, u_\alpha^*TM)_{D,H} \oplus E \rightarrow C^\infty(\tilde{C}_\alpha, \Omega_{\tilde{C}_\alpha}^{0,1} \otimes_{\mathbb{C}} u_\alpha^*TM_{J(\ell_{\alpha,w}(A_y(\cdot)))})$$

given by

$$(C.5.1) \quad \mathcal{F}_{\alpha,w,y}(\xi) := \text{PT}_{\exp u_\alpha \xi \rightarrow u_\alpha}^{\ell_{\alpha,w}(A_y(\cdot))} (d \exp_{u_\alpha} \xi + 2d(\text{proj}_{S^1} A_y) \otimes X_{H((\ell_{\alpha,w} \times \text{id}_{S^1})(A_y(\cdot)))}(\exp_{u_\alpha} \xi)) + \lambda(e_0 + \text{proj}_E \xi)(\alpha, y, x^1, \dots, x^r, \cdot, (\exp_{u_\alpha} \xi)(\cdot))^{0,1}_{j_y, J(\ell_{\alpha,w}(A_y(\cdot)))}$$

(recall that PT and $(\cdot)^{0,1}$ commute). This function $\mathcal{F}_{\alpha,w,y}$ is defined for ξ in a C^1 -neighborhood of zero; for these ξ we define $x^i = x^i(\xi)$ as the unique intersection of $(\text{id}_{M \times S^1} \times \ell_{\alpha,w}) \circ (\exp_{u_\alpha} \xi \times A_y)$ with D_i close to the image of $x_0^i \in C_0$ in C_α (note, however, that even $x^i(0)$ may not coincide exactly with the image of $x_0^i \in C_0$ in C_α); as before, these functions x^i are of class C^{k-2} . The subscript D,H indicates restriction to sections which are tangent to D at p_1, \dots, p_L and tangent to H at $p'_1, \dots, p'_{L'}$. Thus for ξ contained in a C^0 -neighborhood of zero, $\exp_{u_\alpha} \xi$ sends p_1, \dots, p_L to D° and sends $p'_1, \dots, p'_{L'}$ to H° .

Now we observe that $\mathcal{F}_{\alpha,w,y}$ induces a continuous map

$$(C.5.2) \quad \mathcal{F}_{\alpha,w,y}: W^{k,2,\delta,\delta}(C_\alpha, u_\alpha^*TM)_{D,H} \oplus E \rightarrow W^{k-1,2,\delta,\delta}(\tilde{C}_\alpha, \Omega_{\tilde{C}_\alpha}^{0,1} \otimes_{\mathbb{C}} u_\alpha^*TM_{J(\ell_{\alpha,w}(A_y(\cdot)))})$$

which is defined for $\|\xi\|_{k,2,\delta,\delta} \leq c'$ (some $c' > 0$) and small α, w . Moreover, this map is of class C^{k-2} . We denote by

$$(C.5.3) \quad D_{\alpha,w,y}: W^{k,2,\delta,\delta}(C_\alpha, u_\alpha^*TM)_{D,H} \oplus E \rightarrow W^{k-1,2,\delta,\delta}(\tilde{C}_\alpha, \Omega_{\tilde{C}_\alpha}^{0,1} \otimes_{\mathbb{C}} u_\alpha^*TM_{J(\ell_{\alpha,w}(A_y(\cdot)))})$$

the derivative of $\mathcal{F}_{\alpha,w,y}$ at zero.

Let $T_\nabla(X, Y) := \nabla_X Y - \nabla_Y X - [X, Y]$ denote the torsion of ∇ .

Lemma C.5.1 *The linearized operator $D_{\alpha,w,y}$ is given by*

$$(C.5.4) \quad D_{\alpha,w,y}\xi = \left(\nabla^{\ell_{\alpha,w}(A_y(\cdot))}\xi + T_{\nabla}^{\ell_{\alpha,w}(A_y(\cdot))}(\xi, du_{\alpha}) \right)$$

$$(C.5.5) \quad + 2d(\text{proj}_{S^1} A_y) \otimes \nabla_{\xi}^{\ell_{\alpha,w}(A_y(\cdot))} X_{H((\ell_{\alpha,w} \times \text{id}_{S^1})(A_y(\cdot)))}$$

$$(C.5.6) \quad + \sum_{i=1}^r \frac{d[\lambda(e_0)]}{dx^i}(\alpha, y, x^1, \dots, x^r, \cdot, u_{\alpha}(\cdot))(-\text{proj}_{TC_{\alpha}} \xi(x^i))$$

$$(C.5.7) \quad + \nabla_{\xi}^{\ell_{\alpha,w}(A_y(\cdot))}[\lambda(e_0)](\alpha, y, x^1, \dots, x^r, \cdot, u_{\alpha}(\cdot))$$

$$(C.5.8) \quad + \lambda(\text{proj}_E \xi)(\alpha, y, x^1, \dots, x^r, \cdot, u_{\alpha}(\cdot)) \Big)_{j_y, J(\ell_{\alpha,w}(A_y(\cdot)))}^{0,1},$$

where

$$\text{proj}_{TC_{\alpha}}: T(M \times S^1 \times \Delta^n) \rightarrow TC_{\alpha}$$

denotes the projection associated to the splitting

$$T(M \times S^1 \times \Delta^n) = TD_i \oplus TC_{\alpha}$$

at the point

$$x^i \in C_{\alpha} \quad \text{and} \quad (\text{id}_{M \times S^1} \times \ell_{\alpha,w})((u_{\alpha} \times A_y)(x^i)) \in M \times S^1 \times \Delta^n,$$

and $\nabla_{\xi}[\lambda(e_0)]$ denotes the derivative in the direction of ξ along the M factor with respect to the connection ∇ .

Proof Calculation as in Lemma B.5.1. Note that PT and $(\cdot)^{0,1}$ in (C.5.1) commute, since PT is J -linear. □

We denote the kernel of $D_{0,0,0}$ by

$$(C.5.9) \quad K := \ker D_{0,0,0} \subseteq C^{\infty}(C_0, u_0^*TM)_{D,H} \oplus E.$$

Note that our assumption that $(0, 0, 0, u_0, e_0, \{x_0^i\}) \in \overline{\mathcal{M}}(M)^{\text{reg}}$ is equivalent to the statement that $D_{0,0,0}$ is surjective and $K \rightarrow E/E'$ is surjective.

Definition C.5.2 (Kernel pregluing) For $\kappa \in K \subseteq C^{\infty}(C_0, u_0^*TM)$, we define $\kappa_{\alpha} \in C^{\infty}(C_{\alpha}, u_{\alpha}^*TM)$ as follows. Away from the necks, κ_{α} coincides with κ . Over a

bubble neck $S^1 \times [0, 6S]$, we define κ_α as

$$(C.5.10) \quad \kappa_\alpha(t, s) := \begin{cases} \kappa(t, s), & s \leq S - 1, \\ \text{PT}_{u_0(t,s) \rightarrow u_\alpha(t,s)} \kappa(t, s), & S - 1 \leq s \leq S, \\ (1 - \chi(s - S)) \text{PT}_{u_0(t,s) \rightarrow u_\alpha(t,s)} \kappa(t, s) \\ \quad + \chi(s - S) \kappa(n), & S \leq s \leq S + 1, \\ \kappa(n), & S + 1 \leq s \leq 5S - 1, \\ (1 - \chi(s' - S)) \text{PT}_{u_0(t',s') \rightarrow u_\alpha(t',s')} \kappa(t', s') \\ \quad + \chi(s' - S) \kappa(n), & 5S - 1 \leq s \leq 5S, \\ \text{PT}_{u_0(t',s') \rightarrow u_\alpha(t',s')} \kappa(t', s'), & 5S \leq s \leq 5S + 1, \\ \kappa(t', s'), & 5S + 1 \leq s, \end{cases}$$

and over a main neck $S^1 \times [N, 6S - N]$, we define κ_α as

$$(C.5.11) \quad \kappa_\alpha(t, s) := \begin{cases} \kappa(t, s), & s \leq S - 1, \\ \text{PT}_{u_0(t,s) \rightarrow u_\alpha(t,s)} \kappa(t, s), & S - 1 \leq s \leq S, \\ (1 - \chi(s - S)) \text{PT}_{u_0(t,s) \rightarrow u_\alpha(t,s)} \kappa(t, s), & S \leq s \leq S + 1, \\ 0, & S + 1 \leq s \leq 5S - 1, \\ (1 - \chi(-s' - S)) \text{PT}_{u_0(t',s') \rightarrow u_\alpha(t',s')} \kappa(t', s'), & 5S - 1 \leq s \leq 5S, \\ \text{PT}_{u_0(t',s') \rightarrow u_\alpha(t',s')} \kappa(t', s'), & 5S \leq s \leq 5S + 1, \\ \kappa(t', s'), & 5S + 1 \leq s. \end{cases}$$

C.6 Pregluing estimates

Fix norms on E and K , and equip $\mathbb{C}^d \times \mathbb{R}^{k-1}$ and \mathbb{R}^* with their standard norms. Equip each H^i (and hence $\prod_{i=1}^{k'} H^i$) with the pullback of the standard norm under a fixed choice of local diffeomorphism to $\mathbb{R}^n \times \mathbb{R}_{\geq 0}^m$ sending 0 to 0.

Note that we have the following estimates on u_0 and $\kappa \in K$:

$$(C.6.1) \quad \text{in bubble ends:} \quad |D^j \exp_{u_0(n)}^{-1} u_0(t, s)| \leq c_j e^{-s},$$

$$(C.6.2) \quad \text{in bubble ends:} \quad |D^j [\kappa(t, s) - \kappa(n)]| \leq c_j e^{-s} \|\kappa\|,$$

$$(C.6.3) \quad \text{in main ends:} \quad |D^j \exp_{\gamma_i(t)}^{-1} u_0(t, s)| \leq c_j e^{-\delta' s} \quad \text{for } \delta' < \delta_i,$$

$$(C.6.4) \quad \text{in main ends:} \quad |D^j \kappa(t, s)| \leq c_j e^{-\delta' s} \|\kappa\| \quad \text{for } \delta' < \delta_i.$$

The estimates in the bubble ends hold simply because u_0 and κ are smooth on $C_0 \setminus \{q_0, \dots, q_k\}$. The estimates in the main ends hold for u_0 by Proposition C.11.1 and

for κ since $K = \ker D_{0,0,0}$ remains the same for any choice of $k \geq 6$ and any collection of end weights δ , each of which is less than the corresponding δ_i .

Lemma C.6.1 (Estimate for map pregluing) *We have the following estimate on $\mathcal{F}_{\alpha,w,y}(0)$:*

$$(C.6.5) \quad \left\| \left(du_\alpha + 2d(\text{proj}_{S^1} A_y) \otimes X_{H((\ell_{\alpha,w} \times \text{id}_{S^1})(A_y(\cdot)))}(u_\alpha) + \lambda(e_0)(\alpha, y, x^1, \dots, x^r, \cdot, u_\alpha(\cdot)) \right)_{j_y, J(\ell_{\alpha,w}(A_y(\cdot)))}^{0,1} \right\|_{k-1,2,\delta,\delta} \leq c \cdot (|\alpha|^\epsilon + |w| + |y|)$$

uniformly over (α, w, y) in a neighborhood of zero, for $c < \infty$ and $\epsilon > 0$ depending on data which has been previously fixed.

Proof Recall that

$$(C.6.6) \quad \left(du_0 + 2d(\text{proj}_{S^1} A_0) \otimes X_{H((\ell_{0,0} \times \text{id}_{S^1})(A_0(\cdot)))}(u_0) + \lambda(e_0)(0, 0, x_0^1, \dots, x_0^r, \cdot, u_0(\cdot)) \right)_{j_0, J(\ell_{0,0}(A_0(\cdot)))}^{0,1} = 0.$$

We estimate $\mathcal{F}_{\alpha,w,y}(0)$ over the main ends/necks. Note that over this region, the λ term vanishes, $j_y = j_0$,

$$J(\ell_{\alpha,w}(A_y(\cdot))) = J(\ell_{0,0}(A_0(\cdot))), \\ H((\ell_{\alpha,w} \times \text{id}_{S^1})(A_y(\cdot))) = H((\ell_{0,0} \times \text{id}_{S^1})(A_0(\cdot))).$$

Now due to (C.6.6), it follows that in this region, $\mathcal{F}_{\alpha,w,y}(0)$ is supported inside the subsets $S^1 \times ([S - 1, S] \cup [5S, 5S + 1])$ of the main necks $S^1 \times [N, 6S - N]$. The contribution of such a region to the norm of $\mathcal{F}_{\alpha,w,y}(0)$ is bounded above by a constant times $|\alpha|^\epsilon$, as follows from the estimate (C.6.3) on the derivatives of u_0 and the fact that $\delta < \delta_i$.

We estimate $\mathcal{F}_{\alpha,w,y}(0)$ away from the ends/necks. Note that $\ell_{\alpha,w}$ is close to $\ell_{0,0}$, in the sense that they differ over $\coprod_{i=1}^k [-N, N]$ (which may be viewed as a subset of $\coprod_{i=1}^k \mathbb{R}$ and of $(\coprod_{i=1}^k \mathbb{R})_\alpha$) by a constant times $|\alpha|^\epsilon + |w|$ in C^ℓ for some fixed $\epsilon > 0$ (depending on the flow on Δ^n from Definition 10.1.4) and any $\ell < \infty$. It follows that away from the ends/necks, the C^ℓ distance between $J(\ell_{\alpha,w}(A_y(\cdot)))$ on C_α and $J(\ell_{0,0}(A_0(\cdot)))$ on C_0 is bounded by a constant (depending on ℓ) times $|\alpha|^\epsilon + |w| + |y|$, where we identify C_0 and C_α away from the ends/necks in the canonical way (any $\ell < \infty$). The same holds for $H((\ell_{\alpha,w} \times \text{id}_{S^1})(A_y(\cdot)))$ and $H((\ell_{0,0} \times \text{id}_{S^1})(A_0(\cdot)))$. It also follows that the distance between $x^i \in C_\alpha$ and (the image in C_α of) $x_0^i \in C_0$ is bounded by a constant times $|\alpha|^\epsilon + |w|$. Hence we conclude that away from the ends/necks, $\mathcal{F}_{\alpha,w,y}(0)$ differs from the left-hand side of (C.6.6) in C^ℓ by a constant

times $|\alpha|^\epsilon + |w| + |y|$ for any $\ell < \infty$. It follows that the contribution of this region to the norm of $\mathcal{F}_{\alpha,w,y}(0)$ is bounded by a constant times $|\alpha|^\epsilon + |w| + |y|$.

We estimate $\mathcal{F}_{\alpha,w,y}(0)$ over the bubble ends. The reasoning above applies as written to imply that $\mathcal{F}_{\alpha,w,y}(0)$ is bounded in C^ℓ (with respect to the usual metric on \tilde{C}_α) by a constant times $|\alpha|^\epsilon + |w| + |y|$ for any $\ell < \infty$. Since $\delta < 1$, the weighted Sobolev norm over the bubble ends is also bounded by a constant times $|\alpha|^\epsilon + |w| + |y|$.

We estimate $\mathcal{F}_{\alpha,w,y}(0)$ over the bubble necks $S^1 \times [0, 6S]$. The argument for the bubble ends applies to show that the contribution outside $S^1 \times [S - 1, 5S + 1]$ is bounded by a constant times $|\alpha|^\epsilon + |w| + |y|$. Over $S^1 \times ([S - 1, S] \cup [5S, 5S + 1])$, we bound the expression termwise: both u_α and A_y are $O(e^{-s})$ in all derivatives, so the contribution of this region is bounded by $|\alpha|^\epsilon$. Over $S^1 \times [S, 5S]$, only the term involving X_H is nonzero, and since $A_y = O(e^{-s})$ in all derivatives, its contribution to the norm is bounded by a constant times $|\alpha|^\epsilon$ since $\delta < 1$. \square

Lemma C.6.2 (Estimate for kernel pregluing) *For all $\kappa \in K$, we have*

$$(C.6.7) \quad \|D_{\alpha,w,y}\kappa_\alpha\|_{k-1,2,\delta,\delta} \leq c \cdot (|\alpha|^\epsilon + |w| + |y|)\|\kappa\|$$

uniformly over (α, w, y) in a neighborhood of zero, for $c < \infty$ and $\epsilon > 0$ depending on data which has been previously fixed.

Proof Recall that $D_{0,0,0}\kappa = 0$; we will use this to estimate $D_{\alpha,w,y}\kappa_\alpha$ via the explicit expression for $D_{\alpha,w,y}$ from Lemma C.5.1.

We estimate $D_{\alpha,w,y}\kappa_\alpha$ over the main ends/neck. Over this region, the λ terms vanish, and J, H are the same for (α, w, y) as they are for $(0, 0, 0)$. Thus $D_{\alpha,w,y}\kappa_\alpha$ is supported inside $S^1 \times ([S - 1, S + 1] \cup [5S - 1, 5S + 1])$ in the main necks, and vanishes in the main ends. From the exponential decay estimates (C.6.3)–(C.6.4), we obtain that the contribution to the norm of $D_{\alpha,w,y}\kappa_\alpha$ over the main ends/neck is bounded by a constant times $|\alpha|^\epsilon\|\kappa\|$.

We estimate $D_{\alpha,w,y}\kappa_\alpha$ away from the ends/neck. As in the proof of Lemma C.6.1, the difference between J, H for (α, w, y) and for $(0, 0, 0)$ is bounded in C^ℓ by a constant times $|\alpha|^\epsilon + |w| + |y|$ for any $\ell < \infty$; similarly for the distance between x^i and x_0^i . It thus follows from the explicit form in Lemma C.5.1 that the contribution to the norm of $D_{\alpha,w,y}\kappa_\alpha$ over this region is bounded by a constant times $(|\alpha|^\epsilon + |w| + |y|)\|\kappa\|$.

Over the bubble ends, the same reasoning applies, and we obtain the desired bound since $\delta < 1$.

We estimate $D_{\alpha,w,y}\kappa_\alpha$ over the bubble necks. Outside $S^1 \times [S - 1, 5S + 1]$, the reasoning for the bubble ends applies to show that the contribution is bounded as

desired. Over $S^1 \times ([S - 1, S + 1] \cup [5S - 1, 5S + 1])$, the exponential decay estimates on u_0 and κ , along with similar estimates on A_y , show that the contribution of this region is bounded by $|\alpha|^\epsilon \|\kappa\|$ for some $\epsilon > 0$ since $\delta < 1$. Over $S^1 \times [S + 1, 5S - 1]$, only the term involving X_H is nonzero, and its contribution is bounded by a constant times $|\alpha|^\epsilon \|\kappa\|$ since $\delta < 1$ and A_y decays like $O(e^{-s})$ in all derivatives. \square

C.7 Approximate right inverse

Recall that by assumption, the linearized operator

$$(C.7.1) \quad D_{0,0,0}: W^{k,2,\delta,\delta}(C_0, u_0^*TM)_{D,H} \oplus E \rightarrow W^{k-1,2,\delta,\delta}(\tilde{C}_0, \Omega_{\tilde{C}_0,j_0}^{0,1} \otimes_{\mathbb{C}} u_0^*TM_{J(\ell_{0,0}(A_0(\cdot)))})$$

is surjective (even if we replace E with E'). We now proceed to fix a bounded right inverse

$$(C.7.2) \quad Q_{0,0,0}: W^{k-1,2,\delta,\delta}(\tilde{C}_0, \Omega_{\tilde{C}_0,j_0}^{0,1} \otimes_{\mathbb{C}} u_0^*TM_{J(\ell_{0,0}(A_0(\cdot)))}) \rightarrow W^{k,2,\delta,\delta}(C_0, u_0^*TM)_{D,H} \oplus E'$$

whose image admits a simple description. Fix a collection of points $z_i \in C_0$ ($1 \leq i \leq h$) not nodes and not contained in any of the ends, subspaces $V_i \subseteq T_{u_0(z_i)}M$, and a subspace $E'' \subseteq E'$ so that the natural evaluation map

$$(C.7.3) \quad L_0: K \xrightarrow{\simeq} \left(\bigoplus_{i=1}^h T_{u_0(z_i)}M / V_i \right) \oplus E / E''$$

is an isomorphism (such choices exist since $K \twoheadrightarrow E/E'$ is surjective and we may shrink the ends without loss of generality). Now we can consider the same evaluation map on the larger space:

$$(C.7.4) \quad L_0: W^{k,2,\delta,\delta}(C_0, u_0^*TM)_{D,H} \oplus E \rightarrow W := \left(\bigoplus_{i=1}^h T_{u_0(z_i)}M / V_i \right) \oplus E / E''.$$

Since L_0 sends $K = \ker D_{0,0,0}$ isomorphically to W , it follows that the restriction of $D_{0,0,0}$ to $\ker L_0$ is an isomorphism of Banach spaces. Hence there is a unique right inverse

$$(C.7.5) \quad Q_{0,0,0}: W^{k-1,2,\delta,\delta}(\tilde{C}_0, \Omega_{\tilde{C}_0,j_0}^{0,1} \otimes_{\mathbb{C}} u_0^*TM_{J(\ell_{0,0}(A_0(\cdot)))}) \rightarrow W^{k,2,\delta,\delta}(C_0, u_0^*TM)_{D,H} \oplus E$$

with image $\ker L_0$, and it is bounded. Since $E'' \subseteq E'$, $\text{im } Q_{0,0,0} = \ker L_0$ is in fact contained in the right-hand side of (C.7.2). We fix once and for all this $Q_{0,0,0}$.

Definition C.7.1 (Approximate right inverse $T_{\alpha,w,y}$) We define a map

$$(C.7.6) \quad T_{\alpha,w,y}: W^{k-1,2,\delta,\delta}(\tilde{C}_\alpha, \Omega_{\tilde{C}_\alpha}^{0,1} \otimes_{\mathbb{C}} u_\alpha^* TM_{J(\ell_{\alpha,w}(A_y(\cdot)))}) \rightarrow W^{k,2,\delta,\delta}(C_\alpha, u_\alpha^* TM)_{D,H} \oplus E$$

as the composition

$$(C.7.7) \quad T_{\alpha,w,y} := \text{glue} \circ \text{PT} \circ Q_{0,0,0} \circ \text{PT} \circ (\text{id}_*^{1,0} \otimes \text{id}^{1,0})^{-1} \circ \text{break}$$

of maps in the following diagram, to be defined below:

$$(C.7.8) \quad \begin{array}{ccc} W^{k,2,\delta,\delta}(C_\alpha, u_\alpha^* TM)_{D,H} \oplus E & \xrightarrow{D_{\alpha,w,y}} & W^{k-1,2,\delta,\delta}(\tilde{C}_\alpha, \Omega_{\tilde{C}_\alpha}^{0,1} \otimes_{\mathbb{C}} u_\alpha^* TM_{J(\ell_{\alpha,w}(A_y(\cdot)))}) \\ \uparrow \text{glue} & & \downarrow \text{break} \\ W^{k,2,\delta,\delta}(C_0, u_{0|\alpha}^* TM)_{D,H} \oplus E & \xrightarrow{D_{0|\alpha,0,y}} & W^{k-1,2,\delta,\delta}(\tilde{C}_0, \Omega_{\tilde{C}_0}^{0,1} \otimes_{\mathbb{C}} u_{0|\alpha}^* TM_{J(\ell_{0,0}(A_y(\cdot)))}) \\ \parallel & & \uparrow \text{id}_*^{1,0} \otimes \text{id}^{1,0} \\ W^{k,2,\delta,\delta}(C_0, u_{0|\alpha}^* TM)_{D,H} \oplus E & \xrightarrow{D_{0|\alpha,0,0}} & W^{k-1,2,\delta,\delta}(\tilde{C}_0, \Omega_{\tilde{C}_0}^{0,1} \otimes_{\mathbb{C}} u_{0|\alpha}^* TM_{J(\ell_{0,0}(A_0(\cdot)))}) \\ \uparrow \text{PT} & & \uparrow \text{PT} \\ W^{k,2,\delta,\delta}(C_0, u_0^* TM)_{D,H} \oplus E & \xrightleftharpoons[Q_{0,0,0}]{D_{0,0,0}} & W^{k-1,2,\delta,\delta}(\tilde{C}_0, \Omega_{\tilde{C}_0}^{0,1} \otimes_{\mathbb{C}} u_0^* TM_{J(\ell_{0,0}(A_0(\cdot)))}) \end{array}$$

The top and bottom horizontal maps $D_{\alpha,w,y}$ and $D_{0,0,0}$ are the linearized operators defined earlier. The third horizontal map $D_{0|\alpha,0,0}$ is the linearized operator at the flattened map $u_{0|\alpha}$ (its definition is identical to that of $D_{0,0,0}$ except for using $u_{0|\alpha}$ in place of u_0). Similarly, the second horizontal map $D_{0|\alpha,0,y}$ is the linearized operator at the flattened map $u_{0|\alpha}$, using (j_y, A_y) in place of (j_0, A_0) .

The vertical maps PT are simply parallel transport $\text{PT}^{\ell_{0,0}(A_0(\cdot))}$; this is valid since PT is J -linear.

The vertical map $\text{id}_*^{1,0} \otimes \text{id}^{1,0}$ denotes the tensor product of

$$\text{id}^{1,0}: TM_{J(a)} \rightarrow TM_{J(b)}$$

(the $(1, 0)$ -component of the identity map) and

$$\text{id}_*^{1,0}: \Omega_{C_0, j_0}^{0,1} \rightarrow \Omega_{C_0, j_y}^{0,1}$$

(the map induced by $\text{id}^{1,0}: (TC_0, j_0) \rightarrow (TC_0, j_y)$).

We define the map

$$(C.7.9) \quad W^{k,2,\delta,\delta}(C_0, u_{0|\alpha}^* TM)_{D,H} \xrightarrow{\text{glue}} W^{k,2,\delta,\delta}(C_\alpha, u_\alpha^* TM)_{D,H}.$$

Let $\xi \in W^{k,2,\delta,\delta}(C_0, u_{0|\alpha}^* TM)_{D,H}$. Away from the necks of C_α , $\text{glue}(\xi)$ is simply ξ . In any particular bubble neck $S^1 \times [0, 6S] \subseteq C_\alpha$, we define

$$(C.7.10) \quad \text{glue}(\xi)(s, t) := \begin{cases} \xi(s, t), & s \leq 2S, \\ \xi(n) + \chi(4S - s) \cdot [\xi(s, t) - \xi(n)] \\ \quad + \chi(4S - s') \cdot [\xi(s', t') - \xi(n)], & 2S \leq s \leq 4S, \\ \xi(s', t'), & 4S \leq s \end{cases}$$

(noting the corresponding ends $(t, s) \in S^1 \times [0, \infty) \subseteq C_0$ and $(t', s') \in S^1 \times [0, \infty) \subseteq C_0$); this definition also applies over the main necks with the obvious adjustment (and no $\xi(n)$ terms).

We define the map

$$(C.7.11) \quad W^{k-1,2,\delta,\delta}(\tilde{C}_\alpha, \Omega_{\tilde{C}_\alpha, j_y}^{0,1} \otimes_{\mathbb{C}} u_\alpha^* TM_{J(\ell_{\alpha,w}(A_y(\cdot)))}) \xrightarrow{\text{break}} W^{k-1,2,\delta,\delta}(\tilde{C}_0, \Omega_{\tilde{C}_0, j_y}^{0,1} \otimes_{\mathbb{C}} u_{0|\alpha}^* TM_{J(\ell_{0,0}(A_y(\cdot)))})$$

as follows. Fix a smooth function $\bar{\chi}: \mathbb{R} \rightarrow [0, 1]$ such that

$$(C.7.12) \quad \bar{\chi}(x) = \begin{cases} 1, & x \leq -1, \\ 0, & x \geq +1, \end{cases} \quad \bar{\chi}(x) + \bar{\chi}(-x) = 1.$$

Let

$$\eta \in W^{k-1,2,\delta,\delta}(\tilde{C}_\alpha, \Omega_{\tilde{C}_\alpha, j_y}^{0,1} \otimes_{\mathbb{C}} u_\alpha^* TM_{J(\ell_{\alpha,w}(A_y(\cdot)))}).$$

Away from the ends with $\alpha \neq 0$, $\text{break}(\eta)$ is given by $(\text{id}^{1,0})^{-1}(\eta)$, where

$$\text{id}^{1,0}: TM_{J(\ell_{0,0}(A_y(\cdot)))} \rightarrow TM_{J(\ell_{\alpha,w}(A_y(\cdot)))}$$

denotes the $(1, 0)$ -component of the identity map. In any particular bubble end $[0, \infty) \times S^1 \subseteq C_0$ with $\alpha \neq 0$, we define

$$(C.7.13) \quad \text{break}(\eta)(t, s) := \begin{cases} (\text{id}^{1,0})^{-1}(\eta(t, s)), & s \leq 3S - 1, \\ \bar{\chi}(s - 3S) \cdot (\text{id}^{1,0})^{-1}(\eta(t, s)), & 3S - 1 \leq s \leq 3S + 1, \\ 0, & 3S + 1 \leq s \end{cases}$$

(noting the corresponding neck $[0, 6S] \times S^1 \subseteq C_\alpha$); this definition also applies over the main ends, with the obvious adjustment for negative main ends.

Let us make the elementary observation that the definition of L_0 extends perfectly well to give an analogous bounded linear map

$$(C.7.14) \quad L_\alpha: W^{k,2,\delta,\delta}(C_\alpha, u_\alpha^* TM)_{D,H} \oplus E \rightarrow W.$$

Since $\text{im } Q_{0,0,0} \subseteq \ker L_0$, it follows from the definition of $T_{\alpha,w,y}$ that $\text{im } T_{\alpha,w,y} \subseteq \ker L_\alpha$ as well.

Lemma C.7.2 *Let*

$$(C.7.15) \quad \begin{array}{ccc} X & \xrightarrow{D} & Y \\ G \uparrow & & \downarrow B \\ X' & \xrightarrow{D'} & Y' \end{array}$$

denote the bottom square in (C.7.8). Then for $\xi \in X'$ and $\eta \in Y$ with $D'\xi = B\eta$, we have

$$(C.7.16) \quad \|DG\xi - \eta\| \leq c \cdot |\alpha|^\epsilon \|\xi\|$$

uniformly over (α, w, y) in a neighborhood of zero, for $c < \infty$ and $\epsilon > 0$ depending on data which has been previously fixed.

Proof In simpler terms, we bound the operator norm of the difference between the two diagonal compositions:

$$(C.7.17) \quad \|PT \circ D_{0,0,0} - D_{0|\alpha,0,0} \circ PT\| \leq c|\alpha|^\epsilon$$

(this trivially implies the claimed statement). To show (C.7.17), observe that the two operators only differ over the $S^1 \times [S - 1, \infty)$ subset of each end. Now it follows from Lemma C.5.1 that the contribution to the operator norm over $S^1 \times [S - 1, \infty)$ is bounded by a constant times the C^k -distance between u_0 and $u_{0|\alpha}$ over $S^1 \times [0, \infty)$. The estimates (C.6.1), (C.6.3) imply that this distance is bounded by a constant times $|\alpha|^\epsilon$ (to be precise, the bubble ends contribute $|\alpha|^{1/6-\rho}$ and each main end contributes $|\alpha|^{\delta_i/6-\rho}$ for any $\rho > 0$). \square

Lemma C.7.3 *Let*

$$(C.7.18) \quad \begin{array}{ccc} X & \xrightarrow{D} & Y \\ G \uparrow & & \downarrow B \\ X' & \xrightarrow{D'} & Y' \end{array}$$

denote the middle square in (C.7.8). Then for $\xi \in X'$ and $\eta \in Y$ with $D'\xi = B\eta$, we have

$$(C.7.19) \quad \|DG\xi - \eta\| \leq c \cdot |y| \|\xi\|$$

uniformly over (α, w, y) in a neighborhood of zero, for $c < \infty$ and $\epsilon > 0$ depending on data which has been previously fixed.

Proof In simpler terms, we have

$$\|D_{0|\alpha,0,y} - (\text{id}_*^{1,0} \otimes \text{id}^{1,0}) \circ D_{0|\alpha,0,0}\| \leq c \cdot |y|,$$

which trivially implies the claimed statement. To prove this, argue as follows.

We are comparing two first-order differential operators. Appealing to their explicit form from Lemma C.5.1, we see that their coefficients differ by a constant times $|y|$ in C^ℓ for any $\ell < \infty$ (measuring with respect to the cylindrical coordinates $S^1 \times [0, \infty)$ in the ends). It follows that we have the desired estimate. \square

Lemma C.7.4 *Let*

$$(C.7.20) \quad \begin{array}{ccc} X & \xrightarrow{D} & Y \\ G \uparrow & & \downarrow B \\ X' & \xrightarrow{D'} & Y' \end{array}$$

denote the top square in (C.7.8). Then for $\xi \in X'$ and $\eta \in Y$ with $D'\xi = B\eta$, we have

$$(C.7.21) \quad \|DG\xi - \eta\| \leq c \cdot (|\alpha|^\epsilon + |w|)\|\xi\|$$

uniformly over (α, w, y) in a neighborhood of zero, for $c < \infty$ and $\epsilon > 0$ depending on data which has been previously fixed.

Proof We consider the following diagram, which is a copy of the top square in (C.7.8) with one extra vector space and some extra maps (to be defined below).

$$(C.7.22) \quad \begin{array}{ccccc} & & W^{k-1,2,\delta,\delta}(\tilde{C}_\alpha, \Omega_{\tilde{C}_\alpha, j_y}^{0,1} \otimes_{\mathbb{C}} u_\alpha^* TM_{J(\ell_{\alpha,w}(A_y(\cdot)))}) & & \\ & D_{\alpha,w,y} \rightarrow & & \begin{array}{c} \text{glue} \\ \updownarrow \\ \text{break}' \end{array} & \text{break} \\ W^{k,2,\delta,\delta}(C_\alpha, u_\alpha^* TM)_{D,H} \oplus E & & & & \\ \uparrow \text{glue} & & W^{k-1,2,\delta,\delta}(\tilde{C}_0, \Omega_{\tilde{C}_0, j_y}^{0,1} \otimes_{\mathbb{C}} u_{0|\alpha}^* TM_{J(\ell_{\alpha,w}(A_y(\cdot)))}) & & \\ W^{k,2,\delta,\delta}(C_0, u_{0|\alpha}^* TM)_{D,H} \oplus E & D_{\alpha,w,y} \rightarrow & & \begin{array}{c} \text{id}^{1,0} \\ \updownarrow \\ (\text{id}^{1,0})^{-1} \end{array} & \\ & D_{0|\alpha,0,y} \rightarrow & W^{k-1,2,\delta,\delta}(\tilde{C}_0, \Omega_{\tilde{C}_0, j_y}^{0,1} \otimes_{\mathbb{C}} u_{0|\alpha}^* TM_{J(\ell_{0,0}(A_y(\cdot)))}) & & \end{array}$$

Since the almost complex structure $J(\ell_{\alpha,w}(A_y(\cdot)))$ is a function of a point in C_α , we should remark immediately on what we mean by

$$W^{k-1,2,\delta,\delta}(\tilde{C}_0, \Omega_{\tilde{C}_0, j_y}^{0,1} \otimes_{\mathbb{C}} u_{0|\alpha}^* TM_{J(\ell_{\alpha,w}(A_y(\cdot)))}).$$

Away from the ends, the curve C_0 is identified canonically with C_α , and this identification extends holomorphically to a (non-injective) map from C_0 to C_α defined outside

the $S^1 \times [4S, \infty)$ subsets of the ends. Thus there is a well-defined such $W^{k-1,2,\delta,\delta}$ space of sections on C_0 defined outside the $S^1 \times [4S, \infty)$ subsets of the ends, valued in

$$\Omega_{\tilde{C}_0, j_y}^{0,1} \otimes_{\mathbb{C}} u_{0|\alpha}^* TM_{J(\ell_{\alpha,w}(A_y(\cdot)))}.$$

Now, for instance, the middle horizontal map $D_{\alpha,w,y}$ can be seen as giving a section in this $W^{k-1,2,\delta,\delta}$ space defined outside $S^1 \times [4S, \infty)$. In the proof below, it is convenient to use this $W^{k-1,2,\delta,\delta}$ space of sections defined outside $S^1 \times [4S, \infty)$, though we must be careful that the expressions we write are well-defined.

Let us define the rest of the maps in (C.7.22). The vertical map break has been factored as $(\text{id}^{1,0})^{-1} \circ \text{break}'$ in the obvious way. Finally, let us define the map

$$(C.7.23) \quad W^{k-1,2,\delta,\delta}(\tilde{C}_0, \Omega_{\tilde{C}_0, j_y}^{0,1} \otimes_{\mathbb{C}} u_{0|\alpha}^* TM_{J(\ell_{\alpha,w}(A_y(\cdot)))}) \xrightarrow{\text{glue}} W^{k-1,2,\delta,\delta}(\tilde{C}_\alpha, \Omega_{\tilde{C}_\alpha, j_y}^{0,1} \otimes_{\mathbb{C}} u_\alpha^* TM_{J(\ell_{\alpha,w}(A_y(\cdot)))}).$$

Let

$$\eta \in W^{k-1,2,\delta,\delta}(\tilde{C}_0, \Omega_{\tilde{C}_0, j_y}^{0,1} \otimes_{\mathbb{C}} u_{0|\alpha}^* TM_{J(\ell_{\alpha,w}(A_y(\cdot)))}).$$

Away from the necks of C_α , $\text{glue}(\eta)$ is simply η . In any particular bubble neck $S^1 \times [0, 6S] \subseteq C_\alpha$, we define

$$(C.7.24) \quad \text{glue}(\eta)(t, s) := \begin{cases} \eta(t, s), & s \leq 2S, \\ \chi(4S - s)\eta(t, s) + \chi(4S - s')\eta(t', s'), & 2S \leq s \leq 4S, \\ \eta(t', s'), & 4S \leq s \end{cases}$$

(this definition also applies over the main necks with the obvious adjustment). Note that $\text{glue} \circ \text{break}'$ is the identity map.

Now suppose that $D_{0|\alpha,0,y}\xi = \text{break}(\eta)$; we must show that

$$(C.7.25) \quad \|D_{\alpha,w,y}(\text{glue}(\xi)) - \eta\|_{k-1,2,\delta,\delta} \leq c \cdot (|\alpha|^\epsilon + |w|) \|\xi\|_{k,2,\delta,\delta}.$$

Using the fact that $\eta = \text{glue}(\text{break}'(\eta)) = \text{glue}(\text{id}^{1,0}(D_{0|\alpha,0,y}\xi))$ and the triangle inequality, we conclude that $\|D_{\alpha,w,y}(\text{glue}(\xi)) - \eta\|$ is bounded above by

$$\|D_{\alpha,w,y}(\text{glue}(\xi)) - \text{glue}(D_{\alpha,w,y}(\xi))\| + \|\text{glue}[D_{\alpha,w,y}(\xi) - \text{id}^{1,0}(D_{0|\alpha,0,y}(\xi))]\|.$$

We now estimate each term separately.

To estimate $\|D_{\alpha,w,y}(\text{glue}(\xi)) - \text{glue}(D_{\alpha,w,y}(\xi))\|$, argue as follows. The difference is only nonzero over the regions $S^1 \times ([2S, 2S + 1] \cup [4S - 1, 4S])$ of each neck. Now the norm is bounded by $\|\xi\|e^{-2S\delta}$ (calculation left to the reader), where the factor of $e^{-2S\delta}$ comes as the ratio of the weight given to $S^1 \times [4S - 1, 4S]$ inside a neck vs inside an end; this gives the desired bound since $\delta > 0$.

To estimate $\|\text{glue}[D_{\alpha,w,y}(\xi) - \text{id}^{1,0}(D_{0|\alpha,0,y}\xi)]\|$, argue as follows. This is bounded by a constant times the $(k - 1, 2, \delta, \delta)$ -norm of $D_{\alpha,w,y}(\xi) - \text{id}^{1,0}(D_{0|\alpha,0,y}\xi)$ over the complement of the subsets $S^1 \times [4S, \infty)$ of the ends. Now this is bounded by $(|\alpha|^\epsilon + |w|)\|\xi\|_{k,2,\delta,\delta}$ using the reasoning from Lemma C.7.3. \square

Proposition C.7.5 (Approximate right inverse $T_{\alpha,w,y}$) We have

$$(C.7.26) \quad \|T_{\alpha,w,y}\| \leq c,$$

$$(C.7.27) \quad \|D_{\alpha,w,y}T_{\alpha,w,y} - \mathbf{1}\| \rightarrow 0,$$

$$(C.7.28) \quad \text{im } T_{\alpha,w,y} \subseteq \ker L_\alpha$$

as $(\alpha, w, y) \rightarrow 0$, for $c < \infty$ depending on data which has been previously fixed.

Proof It is easy to see that all the maps in (C.7.8) are uniformly bounded. Hence $\|T_{\alpha,w,y}\| \leq c$ as $(\alpha, w, y) \rightarrow 0$. Now Lemma B.7.6 combined with Lemmas C.7.2–C.7.4 show that for $(\alpha, w, y) \rightarrow 0$, we have $\|D_{\alpha,w,y}T_{\alpha,w,y} - \mathbf{1}\| \rightarrow 0$. We observed earlier that $\text{im } T_{\alpha,w,y} \subseteq \ker L_\alpha$. \square

Definition C.7.6 (Right inverse $Q_{\alpha,w,y}$) We define a map

$$(C.7.29) \quad Q_{\alpha,w,y}: W^{k-1,2,\delta,\delta}(\tilde{C}_\alpha, \Omega_{\tilde{C}_\alpha,j_y}^{0,1} \otimes_{\mathbb{C}} u_\alpha^* TM_{J(\ell_{\alpha,w}(A_y(\cdot)))}) \rightarrow W^{k,2,\delta,\delta}(C_\alpha, u_\alpha^* TM)_{D,H} \oplus E$$

as the sum

$$(C.7.30) \quad Q_{\alpha,w,y} := T_{\alpha,w,y} \sum_{k=0}^{\infty} (\mathbf{1} - D_{\alpha,w,y}T_{\alpha,w,y})^k.$$

Proposition C.7.7 We have

$$(C.7.31) \quad \|Q_{\alpha,w,y}\| \leq c,$$

$$(C.7.32) \quad D_{\alpha,w,y}Q_{\alpha,w,y} = \mathbf{1},$$

$$(C.7.33) \quad \text{im } Q_{\alpha,w,y} \subseteq \ker L_\alpha$$

uniformly over (α, w, y) in a neighborhood of zero, for $c < \infty$ depending on data which has been previously fixed.

Proof Apply Lemma B.7.5 and Proposition C.7.5. \square

C.8 Quadratic estimates

Proposition C.8.1 (Quadratic estimate) *There exist $c' > 0$ and $c < \infty$ (depending on data which has been previously fixed) such that for $\|\xi_1\|_{k,2,\delta,\delta}, \|\xi_2\|_{k,2,\delta,\delta} \leq c'$, we have*

$$(C.8.1) \quad \|D_{\alpha,w,y}(\xi_1 - \xi_2) - (\mathcal{F}_{\alpha,w,y}\xi_1 - \mathcal{F}_{\alpha,w,y}\xi_2)\|_{k-1,2,\delta,\delta} \leq c \cdot \|\xi_1 - \xi_2\|_{k,2,\delta,\delta} (\|\xi_1\|_{k,2,\delta,\delta} + \|\xi_2\|_{k,2,\delta,\delta})$$

(and $\mathcal{F}_{\alpha,w,y}\xi_1$ and $\mathcal{F}_{\alpha,w,y}\xi_2$ are both defined), uniformly over (α, w, y) in a neighborhood of zero.

Proof The proof is identical to the proof of Proposition B.8.1, with the evident notational differences, $\mathcal{F}_{\alpha,w,y}$ for $\mathcal{F}_{\alpha,y}$, and Lemma C.5.1 for Lemma B.5.1. As with Proposition B.8.1, we actually prove the stronger statement that

$$(C.8.2) \quad \|\mathcal{F}'_{\alpha,w,y}(0, \xi) - \mathcal{F}'_{\alpha,w,y}(\zeta, \xi)\|_{k-1,2,\delta,\delta} \leq c \cdot \|\zeta\|_{k,2,\delta,\delta} \|\xi\|_{k,2,\delta,\delta}. \quad \square$$

C.9 Newton–Picard iteration

Lemma C.9.1 *There exists $c' > 0$ (depending on data which has been previously fixed) such that for sufficiently small (α, w, y) :*

- (i) *The map $\mathcal{F}_{\alpha,w,y}$ is defined for $\|\xi\|_{k,2,\delta,\delta} \leq c'$.*
- (ii) *For $\xi_1 - \xi_2 \in \text{im } Q_{\alpha,w,y}$ and $\|\xi_1\|_{k,2,\delta,\delta}, \|\xi_2\|_{k,2,\delta,\delta} \leq c'$, we have*

$$(C.9.1) \quad \|(\xi_1 - \xi_2) - (Q_{\alpha,w,y}\mathcal{F}_{\alpha,w,y}\xi_1 - Q_{\alpha,w,y}\mathcal{F}_{\alpha,w,y}\xi_2)\|_{k,2,\delta,\delta} \leq \frac{1}{2} \|\xi_1 - \xi_2\|_{k,2,\delta,\delta}.$$

Proof The proof is identical to the proof of Lemma B.9.1, using Proposition C.8.1 in place of Proposition B.8.1. □

Proposition C.9.2 (Newton–Picard iteration) *There exists $c' > 0$ (depending on data which has been previously fixed) so that for $(\alpha, w, y, \kappa \in K)$ sufficiently small, there exists a unique $\kappa_{\alpha,w,y} \in W^{k,2,\delta,\delta}(C_{\alpha}, u_{\alpha}^*TM)_{D,H} \oplus E$ satisfying*

$$(C.9.2) \quad \kappa_{\alpha,w,y} \in \kappa_{\alpha} + \text{im } Q_{\alpha,w,y},$$

$$(C.9.3) \quad \|\kappa_{\alpha,w,y}\|_{k,2,\delta,\delta} \leq c',$$

$$(C.9.4) \quad \mathcal{F}_{\alpha,w,y}\kappa_{\alpha,w,y} = 0.$$

Proof The proof is essentially identical to the proof of Proposition B.9.2; we write it out anyway.

In fact, we will show that $\kappa_{\alpha,w,y}$ is given explicitly as the limit of the Newton iteration

$$(C.9.5) \quad \xi_0 := \kappa_\alpha,$$

$$(C.9.6) \quad \xi_n := \xi_{n-1} - Q_{\alpha,w,y} \mathcal{F}_{\alpha,w,y} \xi_{n-1}.$$

By Lemma C.9.1, the map $\xi \mapsto \xi - Q_{\alpha,w,y} \mathcal{F}_{\alpha,w,y} \xi$ is a $\frac{1}{2}$ -contraction mapping when restricted to

$$(C.9.7) \quad \{\xi \in \kappa_\alpha + \text{im } Q_{\alpha,w,y} : \|\xi\|_{k,2,\delta,\delta} \leq c'\}$$

To finish the proof, it suffices to show that (for sufficiently small (α, w, y, κ)) (C.9.7) is nonempty and is mapped to itself by $\xi \mapsto \xi - Q_{\alpha,w,y} \mathcal{F}_{\alpha,w,y} \xi$.

We know that $\|\kappa_\alpha\|_{k,2,\delta,\delta} \rightarrow 0$ as $\kappa \rightarrow 0$ (uniformly in (α, w, y)), so (C.9.7) is nonempty. By using Proposition C.8.1 with $(\xi_1, \xi_2) = (0, \kappa_\alpha)$ and Lemmas C.6.1 and C.6.2, we conclude that

$$(C.9.8) \quad \|\mathcal{F}_{\alpha,w,y} \kappa_\alpha\|_{k-1,2,\delta,\delta} \rightarrow 0$$

as $(\alpha, w, y, \kappa) \rightarrow 0$. Since the operator norm of $Q_{\alpha,w,y}$ is bounded uniformly as $(\alpha, w, y) \rightarrow 0$, we see that κ_α is almost fixed by $\xi \mapsto \xi - Q_{\alpha,w,y} \mathcal{F}_{\alpha,w,y} \xi$ as $(\alpha, w, y, \kappa) \rightarrow 0$. It then follows from the contraction property this map takes (C.9.7) to itself. □

C.10 Gluing

Definition C.10.1 (Gluing map) We define

$$(C.10.1) \quad u_{\alpha,w,y,\kappa} := \exp_{u_\alpha} \kappa_{\alpha,w,y},$$

$$(C.10.2) \quad e_{\alpha,w,y,\kappa} := e_0 + \text{proj}_E \kappa_{\alpha,w,y},$$

where $\kappa_{\alpha,w,y}$ is as in Proposition C.9.2, and we consider the *gluing map*

$$(C.10.3) \quad \mathbb{C}^d \times (\mathbb{R}_{\geq 0}^{k-1})^{\leq s} \times \mathbb{R}^* \times \prod_{i=1}^{k'} (H^i)^{\leq s} \times K \rightarrow \overline{\mathcal{M}}(M),$$

$$(C.10.4) \quad (\alpha, w, y, \kappa) \mapsto (\alpha, w, y, u_{\alpha,w,y,\kappa}, e_{\alpha,w,y,\kappa})$$

($\{x^i\}_{1 \leq i \leq r}$ are determined uniquely by $\alpha, w, y, u_{\alpha,w,y,\kappa}$, so we omit them from the notation). It follows from the definition that the gluing map commutes with the

projection from both sides to

$$\mathbb{C}^d \times (\mathbb{R}_{\geq 0}^{k-1})^{\leq s} \times \mathbb{R}^* \times \prod_{i=1}^{k'} (H^i)^{\leq s} \times E/E'.$$

Lemma C.10.2 *The gluing map (C.10.3) maps sufficiently small (α, w, y, κ) to $\overline{\mathcal{M}}(M)^{\text{reg}}$.*

Proof The statement holds since $Q_{\alpha,w,y}$ gives an approximate right inverse at $(u_{\alpha,w,y,\kappa}, e_{\alpha,w,y,\kappa})$ (use (C.8.2) with $\zeta = \kappa_{\alpha,w,y}$). \square

Let $K_\alpha \subseteq C^\infty(C_\alpha, u_\alpha^* TM)_{D,H} \oplus E$ denote the image of $\kappa \mapsto \kappa_\alpha$. It is clear by definition that $K \rightarrow K_\alpha$ is an isomorphism and that the respective $W^{k,2,\delta,\delta}$ norms are uniformly commensurable. It is also clear that the diagram

$$(C.10.5) \quad \begin{array}{ccc} K & \xrightarrow{\kappa \mapsto \kappa_\alpha} & K_\alpha \\ & \searrow L_0 & \swarrow L_\alpha \\ & & W \end{array}$$

commutes (all maps being isomorphisms). Since $\text{im } Q_{\alpha,w,y} \subseteq \ker L_\alpha$, it follows in particular that $\text{im } Q_{\alpha,w,y} \cap K_\alpha = 0$. On the other hand, an index calculation shows that $\text{ind } D_{\alpha,w,y} = \text{ind } D_{0,0,0}$ (note that by the argument for Lemma B.5.2, it suffices to calculate their indices as operators $W^{k,2,\delta} \rightarrow W^{k-1,2,\delta}$, ie with no weights at the bubble nodes, on \tilde{C}_α and \tilde{C}_0 respectively; the calculation of such indices is originally due to Floer [24; 22] and is by now standard). Both are surjective, and hence we have $\dim \text{coker } Q_{\alpha,w,y} = \dim \ker D_{\alpha,w,y} = \dim \ker D_{0,0,0} = \dim K = \dim K_\alpha$. It follows that $\text{im } Q_{\alpha,w,y} = \ker L_\alpha$ and that

$$(C.10.6) \quad \text{im } Q_{\alpha,w,y} \oplus K_\alpha \xrightarrow{\sim} W^{k,2,\delta,\delta}(C_\alpha, u_\alpha^* TX)_D \oplus E$$

is an isomorphism of Banach spaces (an alternative justification for this is to use an argument similar to the proof of Propositions B.11.5/C.11.3 to show that the natural projection

$$K_\alpha \xrightarrow{\kappa_\alpha \mapsto \kappa_\alpha - Q_{\alpha,w,y} D_{\alpha,w,y} \kappa_\alpha} \ker D_{\alpha,w,y}$$

is surjective, and thus bijective). We claim that in fact the two norms are uniformly commensurable as $(\alpha, w, y) \rightarrow 0$. The map written is clearly uniformly bounded, so we just need to show the same for its inverse. It suffices to show that the projection from the right-hand side to K_α is uniformly bounded, but this is nothing other than L_α (clearly uniformly bounded) composed with the inverse of the isomorphism in (C.10.5) (also uniformly bounded).

Lemma C.10.3 *The gluing map (C.10.3) is injective in a neighborhood of zero.*

Proof The proof is identical to the proof of Lemma B.10.3, with the obvious notational differences. □

Proposition C.10.4 *The gluing map (C.10.3) is continuous in a neighborhood of zero.*

Proof The proof follows the same basic outline as the proof of Proposition B.10.4; some parts of the proof are identical, and we will omit these. The key ingredient is our precise control of the image of the right inverse $Q_{\alpha,w,y}$ (specifically, that $\text{im } Q_{\alpha,w,y} = \ker L_{\alpha}$).

Suppose $(\alpha_i, w_i, y_i, \kappa_i) \rightarrow (\alpha, w, y, \kappa)$ is a convergent net.⁵³ We will show that

$$(C.10.7) \quad (u_{\alpha_i, w_i, y_i, \kappa_i}, e_{\alpha_i, w_i, y_i, \kappa_i}) \rightarrow (u_{\alpha, w, y, \kappa}, e_{\alpha, w, y, \kappa}).$$

First, observe that the argument from the proof of Proposition B.10.4 applies verbatim to give that in fact it suffices to show that

$$(C.10.8) \quad (u_{\alpha_i, w_i, y_i, \kappa}, e_{\alpha_i, w_i, y_i, \kappa}) \rightarrow (u_{\alpha, w, y, \kappa}, e_{\alpha, w, y, \kappa}).$$

Now recall that, by definition,

$$(C.10.9) \quad u_{\alpha, w, y, \kappa} = \exp_{u_{\alpha}} \kappa_{\alpha, w, y}, \quad \kappa_{\alpha, w, y} = \kappa_{\alpha} + \xi \quad \text{for some } \xi \in \text{im } Q_{\alpha, w, y}.$$

Now we define $\xi_{\alpha_i} \in W^{k, 2, \delta, \delta}(C_{\alpha_i}, u_{\alpha_i}^* TX)_D \oplus E$ by “pregluing” ξ from C_{α} to C_{α_i} as follows. Note that we may assume without loss of generality that at the nodes where $\alpha \neq 0$, we also have $\alpha_i \neq 0$. Away from the ends/necks of C_{α_i} , we set $\xi_{\alpha_i} = \xi$. Note that for every end of C_{α_i} , there is a corresponding end of C_{α} , so we may also simply set $\xi_{\alpha_i} = \xi$ over the ends of C_{α_i} . Over the necks of C_{α_i} for which $\alpha = 0$, we define ξ_{α_i} via (C.5.10)–(C.5.11) (note that this is reasonable since ξ satisfies the estimates (C.6.2), (C.6.4) as a consequence of Proposition C.11.1). Over the necks of C_{α_i} for which $\alpha \neq 0$, we define ξ_{α_i} as

$$(C.10.10) \quad \xi_{\alpha_i}(s, t) := \text{PT}_{u_{\alpha}(t, f_i(s)) \rightarrow u_{\alpha_i}(t, s)}[\xi(t, f_i(s))],$$

⁵³We could restrict to sequences rather than nets since $\mathbb{C}^d \times (\mathbb{R}_{\geq 0}^{k-1})^{\leq s} \times \mathbb{R}^* \times \prod_{i=1}^{k'} (H^i)^{\leq s} \times K$ is first countable. However, this would not make the argument any simpler.

where $f_i: [0, 6S_i] \rightarrow [0, 6S]$ is smooth and satisfies

$$(C.10.11) \quad f_i(s) := \begin{cases} s, & s \leq S - 2, \\ s - 3S_i + 3S, & 3S_i - 2S + 2 \leq s \leq 3S_i + 2S - 2, \\ s - 6S_i + 6S, & 6S_i - S + 2 \leq s, \end{cases}$$

$$(C.10.12) \quad f_i([S - 2, 3S_i - 2S + 2]) \subseteq [S - 2, S + 2],$$

$$(C.10.13) \quad f_i([3S_i + 2S - 2, 6S_i - S + 2]) \subseteq [5S - 2, 5S + 2],$$

so that $f_i \rightarrow \text{id}$ uniformly in all derivatives as $S_i \rightarrow S$. More informally, f_i is smooth and matches up $[0, S - 2], [S + 2, 5S - 2], [5S + 2, 6S] \subseteq [0, 6S]$ with corresponding intervals of the same length inside $[0, 6S_i]$, symmetrically.

Now at this point we appeal to the corresponding arguments from the proof of Proposition B.10.4, which apply as written (using the fact that $\text{im } Q_{\alpha,w,y} = \ker L_\alpha$) and imply that it suffices to show that

$$(C.10.14) \quad \|\mathcal{F}_{\alpha_i,w_i,y_i}(\kappa_{\alpha_i} + \xi_{\alpha_i})\|_{k-1,2,\delta,\delta} \rightarrow 0.$$

Recall that $\mathcal{F}_{\alpha,w,y}(\kappa_\alpha + \xi) = 0$; we will use this to estimate $\mathcal{F}_{\alpha_i,w_i,y_i}(\kappa_{\alpha_i} + \xi_{\alpha_i})$.

Away from the ends/necks of C_{α_i} , it is clear by definition that $\mathcal{F}_{\alpha_i,w_i,y_i}(\kappa_{\alpha_i} + \xi_{\alpha_i}) \rightarrow \mathcal{F}_{\alpha,w,y}(\kappa_\alpha + \xi)$ uniformly in all derivatives. Thus the contribution to the norm of $\mathcal{F}_{\alpha_i,w_i,y_i}(\kappa_{\alpha_i} + \xi_{\alpha_i})$ over this region approaches zero.

Over the necks of C_{α_i} which correspond to necks of C_α , we again have uniform convergence in all derivatives, so the contribution of these regions approaches zero as well.

We estimate the contribution over the ends of C_{α_i} (recall that these necessarily correspond to ends of C_α). Over main ends, we have $\mathcal{F}_{\alpha_i,w_i,y_i}(\kappa_{\alpha_i} + \xi_{\alpha_i}) = \mathcal{F}_{\alpha,w,y}(\kappa_\alpha + \xi)$. Now the convergence $\mathcal{F}_{\alpha_i,w_i,y_i}(\kappa_{\alpha_i} + \xi_{\alpha_i}) \rightarrow \mathcal{F}_{\alpha,w,y}(\kappa_\alpha + \xi)$ uniformly in all derivatives is valid near the bubble nodes in the usual metric on $C_{\alpha_i} = C_\alpha$. It follows that we also have convergence in the relevant δ -weighted Sobolev norm since $\delta < 1$. Thus the contribution of this region to the norm approaches zero.

Finally, let us estimate the contribution to the norm over the necks of C_{α_i} which correspond to ends of C_α . We treat main necks and bubble necks separately. Over main necks, $\mathcal{F}_{\alpha_i,w_i,y_i}(\kappa_{\alpha_i} + \xi_{\alpha_i})$ is supported inside $S^1 \times ([S - 1, S + 1], [5S - 1, 5S + 1])$. The contribution of this region to its norm goes to zero, as follows using the exponential decrease on κ and ξ from Proposition C.11.1 and the fact that $\delta < \delta_i$. Over bubble ends, $\mathcal{F}_{\alpha_i,w_i,y_i}(\kappa_{\alpha_i} + \xi_{\alpha_i})$ converges to $\mathcal{F}_{\alpha,w,y}(\kappa_\alpha + \xi)$ uniformly in all derivatives over the complement of $S^1 \times [S - 1, 5S + 1]$ (measured with respect to the usual metric on C_0), and hence the contribution of this region to the weighted norm goes to zero since $\delta < 1$. Over $S^1 \times ([S - 1, S + 1] \cup [5S - 1, 5S + 1])$, the smoothness

of κ and ξ over C_α and the fact that $\delta < 1$ shows that the contribution of this region to the norm goes to zero (since $S \rightarrow \infty$). Finally, over $S^1 \times [S + 1, 5S - 1]$, both u_{α_i} and $\kappa_{\alpha_i} + \xi_{\alpha_i}$ are constant, so $\mathcal{F}_{\alpha_i, w_i, y_i}(\kappa_{\alpha_i} + \xi_{\alpha_i})$ is simply

$$\text{PT}_{\exp_{u_0(n)}(\kappa(n) + \xi(n)) \rightarrow u_0(n)}^{\ell_{\alpha_i, w_i}(A_{y_i}(\cdot))} (2d(\text{proj}_{S^1} A_{y_i}) \otimes X_{H((\ell_{\alpha_i, w_i} \times \text{id}_{S^1})(A_{y_i}(\cdot))))}_{j, J(\ell_{\alpha_i, w_i}(A_{y_i}(\cdot)))}^{0,1}.$$

The norm of this expression over $S^1 \times [S + 1, 5S - 1]$ approaches zero since $A_{y_i} = O(e^{-s})$ in all derivatives and $\delta < 1$. □

C.11 Surjectivity of gluing

We recall a well-known principle of exponential decay for Floer trajectories converging to non-degenerate periodic orbits; this has appeared in many forms in the literature, for example in Floer [23, pages 801–802].

Proposition C.11.1 (A priori estimate on decay of connecting cylinders) *Let (M, ω) be a symplectic manifold, J an ω -compatible almost complex structure on M , $H: M \times S^1 \rightarrow \mathbb{R}$ a Hamiltonian, and $\gamma: S^1 \rightarrow M$ a non-degenerate periodic orbit of H . We consider the (unbounded) asymptotic linearized operator*

$$(C.11.1) \quad L^2(S^1, \gamma^*TM) \rightarrow L^2(S^1, \gamma^*TM),$$

$$(C.11.2) \quad \xi \mapsto J \mathcal{L}_{X_H} \xi,$$

where \mathcal{L}_{X_H} denotes the symplectic connection on γ^*TM induced by X_H , and we use the standard inner product $g(\cdot, \cdot) := \omega(\cdot, J\cdot)$ on γ^*TM for the inner product on $L^2(S^1, \gamma^*TM)$. This operator is self-adjoint; we denote by $\delta > 0$ the smallest magnitude of any of its eigenvalues (δ is positive since the orbit is non-degenerate).

We consider (partially defined) sections $\xi: S^1 \times \mathbb{R} \rightarrow \gamma^*TM$ with $|\xi(t, s)| < \epsilon$ such that $u := \exp_\gamma \xi: S^1 \times \mathbb{R} \rightarrow M$ satisfies $(du + 2dt \otimes X_H(u))_J^{0,1} = 0$. Now for all $\mu < 1$, there exists $\epsilon > 0$ such that we have the following estimates.

Suppose that ξ as above is defined on $S^1 \times [0, \infty)$. Then

$$(C.11.3) \quad |D^k \xi| \leq c_k \cdot e^{-\mu \delta s} \left(\int_{S^1} |\xi(t, 0)|^2 dt \right)^{1/2}$$

for $s \geq 1$ all $k \geq 0$. A symmetric statement holds for u and ξ defined over $S^1 \times (-\infty, 0]$.

Suppose that ξ as above is defined on $S^1 \times [0, N]$. Then

$$|D^k \xi| \leq c_k \cdot \left[e^{-\mu \delta s} \left(\int_{S^1} |\xi(t, 0)|^2 dt \right)^{1/2} + e^{-\mu \delta (N-s)} \left(\int_{S^1} |\xi(t, N)|^2 dt \right)^{1/2} \right]$$

for $1 \leq s \leq N - 1$ and all $k \geq 0$.

Proof This proof is adapted from Salamon [72, page 170, Lemma 2.11].

We have by assumption that

$$\partial_s u + J(u)(\partial_t u - X_H(u)) = 0,$$

where $u(t, s) = \exp_{\gamma(t)} \xi(t, s)$. Now we may rewrite this equation in exponential coordinates in terms of ξ as follows. Denote by ∂_t the connection \mathcal{L}_{X_H} on γ^*TM , which is given by the Lie derivative with respect to X_H . Now the equation for u is equivalent to

$$(C.11.4) \quad \partial_s \xi + J \partial_t \xi + A(\xi) \partial_t \xi + Q(\xi) = 0$$

for certain smooth (non-linear) bundle maps

$$A: \gamma^*TM \rightarrow \text{End}(\gamma^*TM) \quad \text{and} \quad Q: \gamma^*TM \rightarrow \gamma^*TM.$$

We have $A(0) = 0$, $Q(0) = 0$ and $Q'(0, \cdot) = 0$ (we denote by A' and Q' their ‘‘vertical’’ derivatives).

Now we let

$$(C.11.5) \quad f(s) := \int_{S^1} |\xi(t, s)|^2 dt = \|\xi\|_2^2,$$

and we will show that $f''(s) \geq (1 - o(1))4\delta^2 f(s)$, where $o(1)$ denotes a quantity which can be made arbitrarily small by choosing $\epsilon > 0$ sufficiently small. We have

$$(C.11.6) \quad f''(s) = 2 \int_{S^1} |\partial_s \xi|^2 dt + 2 \int_{S^1} \langle \xi, \partial_s \partial_s \xi \rangle dt = 2\|\partial_s \xi\|_2^2 + 2\langle \xi, \partial_s \partial_s \xi \rangle,$$

and we will now prove the estimates

$$\|\partial_s \xi\|_2^2 \geq (1 - o(1))\delta^2 f(s) \quad \text{and} \quad \langle \xi, \partial_s \partial_s \xi \rangle \geq (1 - o(1))\delta^2 f(s).$$

Recall that $J \partial_t$ is self-adjoint and that $\|J \partial_t \xi\|_2 \geq \delta \|\xi\|_2$.

To bound $\|\partial_s \xi\|_2^2$, we write

$$(C.11.7) \quad \begin{aligned} \|\partial_s \xi\|_2 &= \|J \partial_t \xi + A(\xi) \partial_t \xi + Q(\xi)\|_2 \\ &\geq \|J \partial_t \xi\|_2 - \|A(\xi) \partial_t \xi\|_2 - \|Q(\xi)\|_2 \\ &\geq \|J \partial_t \xi\|_2 - \|A(\xi)\|_\infty \|\partial_t \xi\|_2 - c \|\xi\|_\infty \|\xi\|_2 \\ &= \|J \partial_t \xi\|_2 - o(1) \|J \partial_t \xi\|_2 - o(1) \|\xi\|_2 \geq (1 - o(1))\delta \|\xi\|_2, \end{aligned}$$

and thus $\|\partial_s \xi\|_2^2 \geq (1 - o(1))\delta^2 f(s)$.

To bound $\langle \xi, \partial_s \partial_s \xi \rangle$, we write

$$\begin{aligned}
 \text{(C.11.8)} \quad \langle \xi, \partial_s \partial_s \xi \rangle &= -\langle \xi, \partial_s [J \partial_t \xi + A(\xi) \partial_t \xi + Q(\xi)] \rangle \\
 &= -\langle J \partial_t \xi, \partial_s \xi \rangle - \langle \xi, A'(\xi, \partial_s \xi) \partial_t \xi \rangle \\
 &\quad - \langle \xi, A(\xi) \partial_t \partial_s \xi \rangle - \langle \xi, Q'(\xi, \partial_s \xi) \rangle.
 \end{aligned}$$

Now substituting in $\partial_s \xi = -J \partial_t \xi - A(\xi) \partial_t \xi - Q(\xi)$, we obtain

$$\begin{aligned}
 \text{(C.11.9)} \quad \langle \xi, \partial_s \partial_s \xi \rangle &= \|J \partial_t \xi\|_2^2 + \langle J \partial_t \xi, A(\xi) \partial_t \xi \rangle \\
 &\quad + \langle J \partial_t \xi, Q(\xi) \rangle \\
 &\quad + \langle \xi, A'(\xi, [J + A(\xi)] \partial_t \xi) \partial_t \xi \rangle \\
 &\quad + \langle \xi, A'(\xi, Q(\xi)) \partial_t \xi \rangle \\
 &\quad + \langle \xi, A(\xi) \partial_t ([J + A(\xi)] \partial_t \xi) \rangle \\
 &\quad + \langle \xi, A(\xi) \partial_t (Q(\xi)) \rangle \\
 &\quad + \langle \xi, Q'(\xi, [J + A(\xi)] \partial_t \xi) \rangle \\
 &\quad + \langle \xi, Q'(\xi, Q(\xi)) \rangle.
 \end{aligned}$$

The first term $\|J \partial_t \xi\|_2^2$ is the main term; let us estimate the remaining error terms. In order, they are bounded by:

- (i) $\|\xi\|_\infty \|J \partial_t \xi\|_2^2$.
- (ii) $\|\xi\|_\infty \|\xi\|_2 \|J \partial_t \xi\|_2$.
- (iii) $\|\xi\|_\infty \|J \partial_t \xi\|_2^2$.
- (iv) $\|\xi\|_\infty^2 \|\xi\|_2 \|J \partial_t \xi\|_2$.
- (v) $\|\xi\|_\infty \|J \partial_t \xi\|_2^2 + \|\xi\|_\infty \|\xi\|_2 \|J \partial_t \xi\|_2$ (integrate by parts to move the outermost ∂_t onto ξ , $A(\xi)$, and $\langle \cdot, \cdot \rangle$).
- (vi) $\|\xi\|_\infty^2 \|\xi\|_2 \|J \partial_t \xi\|_2$.
- (vii) $\|\xi\|_\infty \|\xi\|_2 \|J \partial_t \xi\|_2$.
- (viii) $\|\xi\|_\infty^2 \|\xi\|_2^2$.

All of these are $o(1) \|J \partial_t \xi\|_2^2$, so we conclude that $\langle \xi, \partial_s \partial_s \xi \rangle \geq (1 - o(1)) \delta^2 f(s)$. Combining the above estimates, we obtain the desired inequality $f''(s) \geq (1 - o(1)) 4\delta^2 f(s)$.

Now we have the following maximum principle for the differential inequality $g''(s) \geq r^2 g(s)$. Namely, suppose that $g: [a, b] \rightarrow \mathbb{R}_{\geq 0}$ satisfies $g''(s) \geq r^2 g(s)$ and that $G: [a, b] \rightarrow \mathbb{R}_{\geq 0}$ satisfies $G''(s) = r^2 G(s)$; if $g \leq G$ at the endpoints a, b , then it follows that $g \leq G$ over the whole interval $[a, b]$ (one may easily derive a contradiction

by assuming that, on the contrary, $g = G$ at a_1, b_1 and $g > G$ on (a_1, b_1) for $a \leq a_1 < b_1 \leq b$). Specializing to the case at hand, we have

$$\begin{aligned}
 f(s) &\leq e^{-(1-o(1))2\delta s} f(0) && \text{in the case of } S^1 \times [0, \infty), \\
 f(s) &\leq e^{-(1-o(1))2\delta s} f(0) + e^{-(1-o(1))2\delta(N-s)} f(N) && \text{in the case of } S^1 \times [0, N].
 \end{aligned}$$

Now Lemma C.11.2 below asserts that for a J -holomorphic curve u_0 and a sufficiently small perturbation thereof $u = \exp_{u_0} \xi$ which is also J -holomorphic, we have uniform bounds on all derivatives of ξ (away from the boundary) in terms of the L^2 -norm of ξ . Applying this where u_0 is the trivial cylinder over $\gamma(t)$, we conclude that our bounds on $f(s)$ imply the desired result. \square

Lemma C.11.2 *Let $u: [0, 1] \times [0, 1] \rightarrow B^{2n}(1)$ be J -holomorphic with respect to some smooth almost complex structure J on the unit ball $B^{2n}(1) \subseteq \mathbb{R}^{2n}$. There exists $\epsilon > 0$ (depending only on upper bounds on $\|J\|_{C^\ell}$ and $\|u\|_{C^\ell}$ for some absolute $\ell < \infty$) with the following property. Suppose that $\xi: [0, 1] \times [0, 1] \rightarrow \mathbb{R}^{2n}$ with $|\xi| < \epsilon$ pointwise is such that $u + \xi$ has image contained in $B^{2n}(1)$ and is J -holomorphic. Then*

$$(C.11.10) \quad |D^k \xi(0, 0)| \leq c_k \left(\int_{[0,1] \times [0,1]} |\xi(x, y)|^2 dx dy \right)^{1/2}$$

for all $k < \infty$ for constants $c_k < \infty$ depending only on upper bounds on $\|J\|_{C^\ell}$ and $\|u\|_{C^\ell}$ for some $\ell = \ell(k) < \infty$.

Proof By assumption, we have

$$(C.11.11) \quad u_x + J(u)u_y = 0,$$

$$(C.11.12) \quad (u + \xi)_x + J(u + \xi)(u + \xi)_y = 0,$$

which together imply that

$$(C.11.13) \quad \xi_x + J(u + \xi)\xi_y = [J(u) - J(u + \xi)]u_y,$$

which we prefer to write as

$$(C.11.14) \quad (\partial_x + B\partial_y)\xi = A(\xi),$$

where $B: [0, 1]^2 \rightarrow \text{End}(\mathbb{R}^{2n})$ denotes $J(u + \xi)$ and $A: [0, 1]^2 \times \mathbb{R}^{2n} \rightarrow \mathbb{R}^{2n}$ denotes $A(\xi) = [J(u) - J(u + \xi)]u_y$. By definition, we have $A(0) = 0$ and

$$(C.11.15) \quad |D^k A| \leq c_k$$

for some $c_k < \infty$ depending only on upper bounds on $\|J\|_{C^\ell}$ and $\|u\|_{C^\ell}$ for some $\ell = \ell(k) < \infty$. By definition, we have $B^2 = -1$; let us now argue that we also have

$$(C.11.16) \quad |D^k B| \leq c_k$$

for sufficiently small $\epsilon > 0$ and $c_k < \infty$ as before. Indeed, for sufficiently small $\epsilon > 0$, we may apply the Gromov–Schwarz lemma to $u + \xi$ and conclude that $|D\xi| \leq c < \infty$ depending only on upper bounds on $\|J\|_{C^\ell}$ and $\|u\|_{C^\ell}$ for some $\ell < \infty$. Now Lemma B.11.4 applied to $u + \xi$ gives $|D^k \xi| \leq c_k < \infty$, which implies the desired bound on the derivatives of B .

Now we apply $(\partial_x - B\partial_y)$ to both sides of (C.11.14) to obtain

$$(C.11.17) \quad \Delta\xi = (\partial_x - B\partial_y)A(\xi) - ((\partial_x - B\partial_y)B)(\xi_y).$$

For any smooth function $\phi: [0, 1]^2 \rightarrow \mathbb{R}_{\geq 0}$ supported in the interior of $[0, 1]^2$, we thus have

$$\Delta(\phi\xi) = \xi\Delta\phi + 2\phi_x\xi_x + 2\phi_y\xi_y + \phi \cdot (\partial_x - B\partial_y)A(\xi) - \phi \cdot ((\partial_x - B\partial_y)B)(\xi_y).$$

Now the desired result follows from the usual bootstrapping of the elliptic regularity estimates for the Laplacian on \mathbb{R}^2 , namely $\|f\|_{H^{s+2}} \leq c_s\|f + \Delta f\|_{H^s}$ (the first step being the case $s = -1$). We may shrink the support of ϕ slightly at each step, so there is no need to worry about regularity near the boundary. \square

Proposition C.11.3 *The restriction of the gluing map (C.10.3) to any neighborhood of zero is surjective onto a neighborhood of $(0, 0, 0, u_0, e_0, \{x_i^i\}_{1 \leq i \leq r}) \in \bar{\mathcal{M}}(M)$.*

Proof The proof is identical to the proof of Proposition B.11.5, except for appealing to Proposition C.11.1 in addition to Proposition B.11.1 in the appropriate places. \square

C.12 Conclusion of the proof

We have shown that the map $g := (C.10.3)$ is continuous, injective, and that its restriction to any neighborhood of zero is surjective onto a neighborhood of the image of zero. The target $\bar{\mathcal{M}}(M)$ is Hausdorff, and thus it follows from Lemma B.12.1 that for some open neighborhood of zero

$$U \subseteq \mathbb{C}^d \times (\mathbb{R}_{\geq 0}^{k-1})^{\leq s} \times \mathbb{R}^* \times \prod_{i=1}^{k'} (H^i)^{\leq s} \times K,$$

we have that $g(U)$ is open and $g: U \xrightarrow{\sim} g(U)$ is a homeomorphism. Thus the gluing map (C.10.3) satisfies the properties desired for the map (C.2.1).

C.13 Gluing orientations

We now show the how to endow the moduli spaces $\overline{\mathcal{M}}(\sigma, p, q)$ with coherent orientations using the results of Floer and Hofer [25].

C.13.1 Orientations on spaces of flow lines on Δ^n For any simplex σ , define the orientation line⁵⁴

$$(C.13.1) \quad \mathfrak{o}_\sigma := \begin{cases} \bigotimes_{i=1}^{\dim \sigma - 1} \mathfrak{o}_{\mathbb{R}}, & \dim \sigma > 0, \\ \mathfrak{o}_{\mathbb{R}}^\vee, & \dim \sigma = 0. \end{cases}$$

Let us construct an identification between \mathfrak{o}_σ and the orientation sheaf of the space of flow lines on σ (from Definition 10.1.4) for $n = \dim \sigma > 0$. Let $f: \mathbb{R} \rightarrow [0, 1]$ be the unique solution to the initial value problem $f(0) = \frac{1}{2}$ and $f'(x) = \pi \sin(\pi f(x))$. Now every flow line $\ell: \mathbb{R} \rightarrow \Delta^n$ is of the form $\ell(t) = (f(t + a_1), \dots, f(t + a_n))$ for some $a_1 \leq \dots \leq a_n$ which are unique up to the addition of an overall constant. Thus the space of flow lines is parametrized by $(b_1, \dots, b_{n-1}) \in [0, \infty)^{n-1}$, where $b_i = a_{i+1} - a_i$; moreover, this parametrization extends continuously to a homeomorphism between $[0, \infty)^{n-1}$ and the space of broken flow lines on Δ^n . In these coordinates, $b_k = \infty$ if and only if the flow line is broken at vertex k , and $b_k = 0$ if and only if the flow line factors through $\Delta^{[0 \dots \hat{k} \dots n]} \subseteq \Delta^n$. Now the coordinates b_1, \dots, b_{n-1} determine an identification between $\mathfrak{o}_\sigma = \bigotimes_{i=1}^{n-1} \mathfrak{o}_{\mathbb{R}}$ and the orientation sheaf of the space of flow lines on σ , when $\dim \sigma > 0$.

For $\dim \sigma = 0$, let us simply remark that tensoring with $\mathfrak{o}_\sigma = \mathfrak{o}_{\mathbb{R}}^\vee$ is the effect on orientation sheaves of quotienting by an \mathbb{R} -action, and it will follow from this fact that this is the correct definition of \mathfrak{o}_σ when $\dim \sigma = 0$ (morally speaking, we may think of the space of flow lines on σ as the stacky quotient pt/\mathbb{R}).

Now let us observe that there are natural (odd) “boundary” maps

$$(C.13.2) \quad \mathfrak{o}_\sigma \rightarrow \mathfrak{o}_{\sigma|[0 \dots k]} \otimes \mathfrak{o}_{\sigma|[k \dots n]},$$

$$(C.13.3) \quad \mathfrak{o}_\sigma \rightarrow \mathfrak{o}_{\sigma|[0 \dots \hat{k} \dots n]}$$

induced by the geometric inclusions of boundary strata (10.1.3)–(10.1.4). Specifically, the first map is induced by the inclusion of the space of pairs of flow lines on $\sigma|[0 \dots k]$ and $\sigma|[k \dots n]$ into the space of flow lines on σ ($0 < k < n$), and the second by the inclusion of flow lines on $\sigma|[0 \dots \hat{k} \dots n]$ into flow lines on σ ($0 < k < n$). In fact, the first map (C.13.2) admits the following alternative description, which shows that it is in fact defined for $0 \leq k \leq n$. Let $\mathcal{F}^{(i)}(\sigma)$ denote the moduli space of stable broken Morse flow lines on a simplex σ with i ordered marked points appearing in order

⁵⁴An orientation line is a $\mathbb{Z}/2$ -graded free \mathbb{Z} -module of rank one.

(this is defined in the expected way, allowing constant flow lines as long as they are stabilized by the presence of at least one marked point). Denote the orientation module of $\mathcal{F}^{(i)}(\sigma)$ by $\mathfrak{o}_\sigma^{(i)}$, and note that there is a natural isomorphism $\mathfrak{o}_\sigma^{(i)} = \mathfrak{o}_\sigma \otimes \mathfrak{o}_\mathbb{R}^{\otimes i}$ (even for $\dim \sigma = 0$; this provides another justification of our definition of \mathfrak{o}_σ in this case). Now there is a natural concatenation map

$$\mathcal{F}^{(i)}(\sigma|[0 \dots k]) \times \mathcal{F}^{(j)}(\sigma|[k \dots n]) \rightarrow \mathcal{F}^{(i+j)}(\sigma)$$

which is the inclusion of a codimension-one boundary stratum, thus giving rise to a boundary map

$$\mathfrak{o}_\sigma^{(i+j)} \rightarrow \mathfrak{o}_{\sigma|[0 \dots k]}^{(i)} \otimes \mathfrak{o}_{\sigma|[k \dots n]}^{(j)}.$$

After factoring out $\mathfrak{o}_\mathbb{R}^{\otimes(i+j)}$, this map coincides with (C.13.2) for $0 < k < n$, and thus may be used to define (C.13.2) for $0 \leq k \leq n$.

Remark C.13.1 The standard orientation of \mathbb{R} gives a standard generator

$$[\mathbb{R}]_1 \otimes \dots \otimes [\mathbb{R}]_{n-1} \in \mathfrak{o}_\sigma$$

for $\dim \sigma > 0$, and $[\mathbb{R}]^\vee \in \mathfrak{o}_\sigma$ for $\dim \sigma = 0$, where $[\mathbb{R}] \cdot [\mathbb{R}]^\vee = 1$ (ie $[\mathbb{R}]^\vee \in \mathfrak{o}_\mathbb{R}^\vee$ is the “right dual” of $[\mathbb{R}] \in \mathfrak{o}_\mathbb{R}$). Using these generators to trivialize \mathfrak{o}_σ determines a sign convention in which (C.13.2) is given by $(-1)^{k+1}$ and (C.13.3) is given by $(-1)^k$.

C.13.2 Orientations of linearized operators of Floer equations Fix two Hamiltonians $H_0, H_1: M \times S^1 \rightarrow \mathbb{R}$ and two non-degenerate periodic orbits $p, q: S^1 \rightarrow M$ of H_0, H_1 , respectively. Now for any path of ω -compatible almost complex structures $J: \mathbb{R} \rightarrow J(M)$ (constant near $s = \pm\infty$), any path of Hamiltonians $H: M \times S^1 \times \mathbb{R} \rightarrow \mathbb{R}$ (constant near $s = \pm\infty$) with $H(s = -\infty) = H_0$ and $H(s = \infty) = H_1$, and any map $u: S^1 \times \mathbb{R} \rightarrow M$ in $W^{k,2,\delta}(S^1 \times \mathbb{R}, M)$ (converging to q at $-\infty$ and to p at ∞), there is a natural linearized operator

$$(C.13.4) \quad D_{\text{Floer}}: W^{k,2,\delta}(S^1 \times \mathbb{R}, u^*TM) \rightarrow W^{k-1,2,\delta}(S^1 \times \mathbb{R}, \Omega_{S^1 \times \mathbb{R}}^{0,1} \otimes_{\mathbb{C}} u^*TM_J)$$

(we assume that $\delta > 0$ is less than the smallest positive eigenvalues of the asymptotic linearized operators). Let us agree to define this linearized operator using the \mathbb{R} -family of J -linear connections on TM given by $\nabla_X Y := \frac{1}{2}(\nabla_X^0 Y - J(\nabla_X^0(JY)))$, where ∇^0 denotes the Levi-Civita connection of the metric associated to the compatible pair (ω, J) . The results of Floer and Hofer [25] (see the remark after Theorem 2) imply that the Fredholm orientation line of D_{Floer} is a trivial bundle over the space of paths H, J , and maps u . Thus if we fix p, q , and a homotopy class of maps u , there is a well-defined orientation line $\mathfrak{o}_{p,q}$, canonically isomorphic to $\mathfrak{o}_{D_{\text{Floer}}}$ for any choice of H, J , and u .

Floer and Hofer [25, Theorem 10] also construct natural associative (even) isomorphisms

$$(C.13.5) \quad \mathfrak{o}_{p,q} \otimes \mathfrak{o}_{q,r} \rightarrow \mathfrak{o}_{p,r}$$

by a certain kernel gluing procedure. We refer the reader to [25] for the details of this construction. Suffice it to say here that, after adding a finite-dimensional piece to the domains of each of two linearized operators (C.13.4) (from p to q and from q to r , respectively) so that they become surjective with kernels K_1 and K_2 , there is a natural kernel pregluing map from $K_1 \oplus K_2$ to the domain of certain glued operator (from p to r); the L^2 -orthogonal projection onto the kernel K_3 of the glued operator gives an isomorphism $K_1 \oplus K_2 \xrightarrow{\sim} K_3$, and this induces the map (C.13.5), which may be shown to be independent of the choices used to define it.

The coherent trivializations of $\mathfrak{o}_{p,q}$ resulting from the maps (C.13.5) are known to coincide with the usual coherent orientations of Morse theory, when restricted to S^1 -invariant Hamiltonians and their S^1 -invariant Floer trajectories; see eg Floer [24].

C.13.3 Orientations on thickened moduli spaces $\overline{\mathcal{M}}(\sigma, p, q)_I^{\text{reg}}$ The existence of the desired coherent orientations for $\overline{\mathcal{M}}(\sigma, p, q)$ follows easily from the following result, which we spend the rest of this appendix proving.

Proposition C.13.2 *The orientation sheaf (in the sense of Definition 4.1.3) of every moduli space $\overline{\mathcal{M}}(\sigma, p, q)$ can be canonically identified with $\mathfrak{o}_\sigma \otimes \mathfrak{o}_{p,q}$. Moreover, under these identifications, the boundary maps on orientation sheaves induced (as in (4.4.2)) by the structure maps*

$$(C.13.6) \quad \overline{\mathcal{M}}(\sigma|[0 \dots k], p, q) \times \overline{\mathcal{M}}(\sigma|[k \dots n], q, r) \rightarrow \partial \overline{\mathcal{M}}(\sigma, p, r),$$

$$(C.13.7) \quad \overline{\mathcal{M}}(\sigma|[0 \dots \hat{k} \dots n], p, q) \rightarrow \partial \overline{\mathcal{M}}(\sigma, p, q)$$

coincide with the maps

$$(C.13.8) \quad \mathfrak{o}_\sigma \otimes \mathfrak{o}_{p,r} \rightarrow \mathfrak{o}_{\sigma|[0 \dots k]} \otimes \mathfrak{o}_{p,q} \otimes \mathfrak{o}_{\sigma|[k \dots n]} \otimes \mathfrak{o}_{q,r},$$

$$(C.13.9) \quad \mathfrak{o}_\sigma \otimes \mathfrak{o}_{p,q} \rightarrow \mathfrak{o}_{\sigma|[0 \dots \hat{k} \dots n]} \otimes \mathfrak{o}_{p,q}$$

induced by (C.13.2)–(C.13.3) and (C.13.5).

Let us introduce various linearized operators which will play a role in the proof of Proposition C.13.2 below. At any point in a thickened moduli space $\overline{\mathcal{M}}(\sigma, p, q)_I$, we have a linearized operator

$$(C.13.10) \quad D: E_I \oplus W^{k,2,\delta}(C, u^*TM) \rightarrow W^{k-1,2,\delta}(C, \Omega_C^{0,1} \otimes_{\mathbb{C}} u^*TM_{J(\ell(\cdot))})$$

More precisely, the operator D denotes the usual linearization of the I -thickened holomorphic equation, eg as calculated in Lemma C.5.1 (corresponding to variations of the map u and of the element $e \in E_I$; we always keep ℓ fixed when defining linearized operators). We will need to make use of this and other linearized operators at maps u which do not necessarily satisfy the relevant pseudo-holomorphic curve equation, and for such u , the linearized operator D depends on a choice of (a family of) J -linear connections on M . Let us fix the convention of always using the J -linear connection $\nabla_X Y := \frac{1}{2}(\nabla_X^0 Y - J(\nabla_X^0(JY)))$, where ∇^0 is the Levi-Civita connection of the metric associated to the compatible pair (ω, J) (more precisely, this is a family of connections parametrized by $\sigma = \Delta^n$).

We will make use of another linearized operator

$$(C.13.11) \quad D_{\text{hol}}: E_I \oplus W^{k,2,\delta}(C, u^*TM) \rightarrow W^{k-1,2,\delta}(C, \Omega_C^{0,1} \otimes_C u^*TM_{J(\ell(\cdot))}).$$

The operator D_{hol} denotes the linearization of the usual holomorphic curve equation (ie without the thickening terms $\lambda_\alpha(e_\alpha)$); thus it is given by the expression in Lemma C.5.1 without the terms (C.5.6)–(C.5.8). Clearly there is a natural isomorphism $\mathfrak{o}_D = \mathfrak{o}_{D_{\text{hol}}}$ since the terms (C.5.6)–(C.5.8) are compact.

Proof of Proposition C.13.2 To identify the orientation sheaf of $\overline{\mathcal{M}}(\sigma, p, q)$ with $\mathfrak{o}_\sigma \otimes \mathfrak{o}_{p,q}$, it suffices to identify the orientation sheaf of every $\overline{\mathcal{M}}(\sigma, p, q)_I^{\text{reg}}$ with $\mathfrak{o}_{E_I} \otimes \mathfrak{o}_\sigma \otimes \mathfrak{o}_{p,q}$ in a compatible way (for all finite subsets $I \subseteq \mathcal{A}(\sigma, p, q)^{\geq \varepsilon^{\text{top}}}$). Moreover, it suffices to make this identification over the open subset

$$\overline{\mathcal{M}}(\sigma, p, q)_I^{\text{reg } \circ} \subseteq \overline{\mathcal{M}}(\sigma, p, q)_I^{\text{reg}}$$

where the domain curve is smooth (it then automatically extends uniquely to all of $\overline{\mathcal{M}}(\sigma, p, q)_I^{\text{reg}}$, by virtue of the local topological description of $\overline{\mathcal{M}}(\sigma, p, q)_I^{\text{reg}}$ given by the gluing map); we may also check compatibility with the inclusions $I \subseteq J$ over $\overline{\mathcal{M}}(\sigma, p, q)_I^{\text{reg } \circ}$. Now the kernel $\tilde{K} = \ker D$ of the linearization (C.13.10) forms a vector bundle over $\overline{\mathcal{M}}(\sigma, p, q)_I^{\text{reg } \circ}$, and the orientation sheaf of $\overline{\mathcal{M}}(\sigma, p, q)_I^{\text{reg } \circ}$ is isomorphic to $\mathfrak{o}_\sigma \otimes \mathfrak{o}_{\tilde{K}}$ (to see this, one must distinguish the two cases $\dim \sigma > 0$ and $\dim \sigma = 0$). Now we have

$$\mathfrak{o}_{\tilde{K}} = \mathfrak{o}_D = \mathfrak{o}_{D_{\text{hol}}} = \mathfrak{o}_{E_I} \otimes \mathfrak{o}_{p,q};$$

this defines a fiberwise isomorphism of $\mathfrak{o}_{\tilde{K}}$ with $\mathfrak{o}_{E_I} \otimes \mathfrak{o}_{p,q}$, and since the operators D and D_{hol} vary nicely over the base $\overline{\mathcal{M}}(\sigma, p, q)_I^{\text{reg } \circ}$, it is easy to see that this is in fact an isomorphism of sheaves. Thus we have the desired identification.

Now let us show that the boundary map induced by (C.13.7) coincides with the tautological map (C.13.9) (this just amounts to chasing definitions). Let $\varepsilon \in \mathcal{S}_{\overline{\mathcal{M}}}(\sigma, p, q)$ denote

the stratum $(\sigma|[0 \dots \hat{k} \dots n], p, q)$ (ie the stratum consisting of trajectories over σ which factor through $\sigma|[0 \dots \hat{k} \dots n]$). Now on the space $\overline{\mathcal{M}}(\sigma|[0 \dots \hat{k} \dots n], p, q) = \overline{\mathcal{M}}(\sigma, p, q)^{\leq s}$, we have three implicit atlases:

$$(C.13.12) \quad \mathcal{A}(\sigma|[0 \dots \hat{k} \dots n], p, q)^{\geq s^{\text{top}}} \subseteq \mathcal{A}(\sigma, p, q)^{\geq s} \supseteq \mathcal{A}(\sigma, p, q)^{\geq s^{\text{top}}}$$

(note that the two occurrences of s^{top} refer to the top elements of two different strata posets); in fact (the index set of) the middle atlas is the disjoint union of (the index sets of) the atlases on the right and on the left. The reasoning used above to identify the orientation sheaf of $\overline{\mathcal{M}}(\sigma|[0 \dots \hat{k} \dots n], p, q) = \overline{\mathcal{M}}(\sigma, p, q)^{\leq s}$ under the leftmost implicit atlas applies equally well for all three atlases (and, in particular, the resulting identifications coincide). Hence, for the purpose of identifying the boundary map on orientation sheaves, it suffices to consider the single atlas $\mathcal{A}(\sigma, p, q)^{\geq s^{\text{top}}}$ on $\overline{\mathcal{M}}(\sigma, p, q)$ and the inclusion of the boundary stratum $\overline{\mathcal{M}}(\sigma, p, q)^{\leq s} \hookrightarrow \overline{\mathcal{M}}(\sigma, p, q)$. Now it suffices to check that the desired compatibility holds for the inclusion $(\overline{\mathcal{M}}(\sigma, p, q)_I^{\leq s})^{\text{reg}} \hookrightarrow \overline{\mathcal{M}}(\sigma, p, q)_I^{\text{reg}}$ (for all finite subsets $I \subseteq \mathcal{A}(\sigma, p, q)^{\geq s^{\text{top}}}$), and this follows by definition.

Now let us show that the boundary map induced by (C.13.6) coincides with the tautological map (C.13.8). First, we chase definitions as above. Let $s \in \mathcal{S}_{\overline{\mathcal{M}}}(\sigma, p, r)$ denote the stratum $(\sigma|[0 \dots k], p, q) - (\sigma|[k \dots n], q, r)$ (ie the stratum consisting of Floer trajectories over σ which are broken at vertex k and periodic orbit $q \in \mathcal{P}_H(k)$). Now on the space

$$\overline{\mathcal{M}}(\sigma|[0 \dots k], p, q) \times \overline{\mathcal{M}}(\sigma|[k \dots n], q, r) = \overline{\mathcal{M}}(\sigma, p, r)^{\leq s},$$

we have three implicit atlases:

$$\mathcal{A}(\sigma|[0 \dots k], p, q)^{\geq s^{\text{top}}} \sqcup \mathcal{A}(\sigma|[k \dots n], q, r)^{\geq s^{\text{top}}} \subseteq \mathcal{A}(\sigma, p, r)^{\geq s} \supseteq \mathcal{A}(\sigma, p, r)^{\geq s^{\text{top}}}$$

and in fact (the index set of) the middle atlas is the disjoint union of (the index sets of) the atlases on the right and on the left. The reasoning used above to identify the orientation sheaves of $\overline{\mathcal{M}}(\sigma|[0 \dots k], p, q)$ and $\overline{\mathcal{M}}(\sigma|[k \dots n], q, r)$ applies equally well to identify the orientation sheaf of their product

$$\overline{\mathcal{M}}(\sigma|[0 \dots k], p, q) \times \overline{\mathcal{M}}(\sigma|[k \dots n], q, r) = \overline{\mathcal{M}}(\sigma, p, r)^{\leq s}$$

under all three atlases above (the fact that the thickening datums in the latter two atlases do not treat the portions of the trajectories from p to q and from q to r “independently” does not cause any problems). Hence for the purposes of identifying the boundary map on orientation sheaves, it suffices to consider the single atlas $\mathcal{A}(\sigma, p, r)^{\geq s^{\text{top}}}$ on $\overline{\mathcal{M}}(\sigma, p, r)$ and the inclusion of the boundary stratum $\overline{\mathcal{M}}(\sigma, p, r)^{\leq s} \hookrightarrow \overline{\mathcal{M}}(\sigma, p, r)$. As before, it suffices to check that the desired compatibility holds for the inclusion $(\overline{\mathcal{M}}(\sigma, p, r)_I^{\leq s})^{\text{reg}} \hookrightarrow \overline{\mathcal{M}}(\sigma, p, r)_I^{\text{reg}}$ (for all finite subsets $I \subseteq \mathcal{A}(\sigma, p, r)^{\geq s^{\text{top}}}$), and

moreover, this may be checked on the locus where the domain curve is smooth except for the required node over k asymptotic to $q \in \mathcal{P}_{H(k)}$.

It will be convenient to assume that $0 < k < n$, so let us first deduce the case of general $0 \leq k \leq n$ from the case $0 < k < n$. Fix σ and $0 \leq k \leq n = \dim \sigma$. We consider the degenerate simplex $\sigma' \rightarrow \sigma$ given by “doubling vertex 0” if $k = 0$ and “doubling vertex n ” if $k = n$ (ie if $0 = k < n$, we consider $\Delta^{n+1} \rightarrow \Delta^n$ given by superimposing the first two vertices of Δ^{n+1} , if $0 < k = n$, we consider $\Delta^{n+1} \rightarrow \Delta^n$ given by superimposing the last two vertices of Δ^{n+1} , and if $0 = k = n$ we consider $\Delta^2 \rightarrow \Delta^0$); let $k' = 1$ if $k = 0$ and let $k' = n' - 1$ if $k = n$, where $n' = \dim \sigma'$, and note that $0 < k' < n'$. Now there is a natural map $\mathcal{A}(\sigma, p, q)^{\geq s^{\text{top}}} \rightarrow \mathcal{A}(\sigma', p, q)^{\geq s^{\text{top}}}$ given by pulling back along $\sigma' \rightarrow \sigma$, and $\sigma' \rightarrow \sigma$ maps flow lines to flow lines (for the flow from Definition 10.1.4). Hence, given any point $x \in \overline{\mathcal{M}}(\sigma, p, r)_I^{\text{reg}}$ with a single node over k asymptotic to $q \in \mathcal{P}_{H(k)}$, we may lift it (canonically, up to translating the part(s) of the trajectory over $0 = k$ and/or over $k = n$) to a point $x' \in \overline{\mathcal{M}}(\sigma', p, r)_I^{\text{reg}}$ with a single node over k' asymptotic to $q \in \mathcal{P}_{H'(k')} = \mathcal{P}_{H(k)}$. Moreover, there is a germ of a homeomorphism

$$\overline{\mathcal{M}}(\sigma', p, r)_I^{\text{reg}} = \overline{\mathcal{M}}(\sigma, p, r)_I^{\text{reg}} \times \mathbb{R}^{\mathbf{1}(k=0)+\mathbf{1}(k=n)}$$

between neighborhoods of x' and $x \times 0$. Now the desired compatibility of orientations for $(\overline{\mathcal{M}}(\sigma, p, r)_I^{\leq s})^{\text{reg}} \hookrightarrow \overline{\mathcal{M}}(\sigma, p, r)_I^{\text{reg}}$ follows from the compatibility for $(\overline{\mathcal{M}}(\sigma', p, r)_I^{\leq s'})^{\text{reg}} \hookrightarrow \overline{\mathcal{M}}(\sigma', p, r)_I^{\text{reg}}$ (which has from $0 < k' < n'$). Hence we may assume without loss of generality that $0 < k < n$.

Now we have come to the heart of the matter, where we will need to analyze the gluing map. To review: we have identified the orientation sheaf of $\overline{\mathcal{M}}(\sigma, p, r)_I^{\text{reg}}$ with $\sigma_{E_I} \otimes \sigma_{\sigma} \otimes \sigma_{p,r}$, and we have identified the orientation sheaf of $(\overline{\mathcal{M}}(\sigma, p, r)_I^{\leq s})^{\text{reg}}$ with $\sigma_{E_I} \otimes \sigma_{\sigma[[0\dots k]]} \otimes \sigma_{p,q} \otimes \sigma_{\sigma[[k\dots n]]} \otimes \sigma_{q,r}$ (recall that s denotes the stratum consisting of trajectories broken at vertex k and periodic orbit $q \in \mathcal{P}_{H(k)}$). We must show that the boundary map on orientation sheaves induced by the inclusion $(\overline{\mathcal{M}}(\sigma, p, r)_I^{\leq s})^{\text{reg}} \hookrightarrow \overline{\mathcal{M}}(\sigma, p, r)_I^{\text{reg}}$ is the expected map (C.13.8), and it suffices to check this over the locus where the domain curve is smooth except for the required node over k asymptotic to $q \in \mathcal{P}_{H(k)}$. We may further assume that $0 < k < n$.

We consider the gluing setup for $\overline{\mathcal{M}}(\sigma, p, r)_I^{\text{reg}}$ at a point in $(\overline{\mathcal{M}}(\sigma, p, r)_I^{\leq s})^{\text{reg}}$ where C_0 is smooth except for the required node over vertex k asymptotic to $q \in \mathcal{P}_{H(k)}$. In other words,

$$C_0 \setminus \{q_0, q_1, q_2\} = S^1 \times \mathbb{R} \sqcup S^1 \times \mathbb{R},$$

and

$$u_0: S^1 \times \mathbb{R} \sqcup S^1 \times \mathbb{R} \rightarrow M$$

is a trajectory from p to q and from q to r . From Section C.1.2, we take $L = L' = 0$ (ie no points p_i or p'_i), and we may take $D = H = \emptyset$. From Section C.1.3, there is one gluing parameter $\alpha \in \mathbb{R}_{\geq 0}$, there is no nontrivial variation in (j_0, A_0) (ie $y \in \mathbb{R}^* = \mathbb{R}^0$), and there is some variation in ℓ_0 parametrized by $w \in \prod_{i=1}^{2'} H^i$. As usual, we denote by $K \subseteq W^{k,2,\delta}(C_0, u_0^*TM)$ the kernel of the linearized operator (we shall omit the subscript D, H since $L = L' = 0$). Now the gluing construction gives rise to a gluing map

$$(C.13.13) \quad K \times \mathbb{R}_{\geq 0} \times \prod_{i=1}^{2'} H^i \rightarrow \overline{\mathcal{M}}(\sigma, p, r)_I^{\text{reg}},$$

$$(C.13.14) \quad (\kappa, \alpha, w) \mapsto (u_{\alpha,w,y,\kappa}, e_{\alpha,w,y,\kappa})$$

($y = 0$). Now if we restrict to $\alpha = 0$, this map realizes the identification of the orientation sheaf of $(\overline{\mathcal{M}}(\sigma, p, r)_I^{\leq s})^{\text{reg}}$ with $\mathfrak{o}_K \otimes \mathfrak{o}_{\sigma|[0\dots k]} \otimes \mathfrak{o}_{\sigma|[k\dots n]}$. Thus it suffices to show that for fixed (sufficiently small) $\alpha > 0$ and fixed $w = 0$, the gluing map is differentiable in the K direction, and that its derivative (a map from K to the kernel of the linearized operator at the glued map $u_{\alpha,w,y,\kappa}$) agrees (on orientation lines) with the Floer–Hofer map (C.13.5) (recall that \mathfrak{o}_K is identified with $\mathfrak{o}_{E_I} \otimes \mathfrak{o}_{p,q} \otimes \mathfrak{o}_{q,r}$ and that the orientation line of the kernel at the glued map is identified with $\mathfrak{o}_{E_I} \otimes \mathfrak{o}_{p,r}$).

For fixed (α, w, y) , the gluing map is given by

$$(C.13.15) \quad K \xrightarrow{\kappa \mapsto \kappa_\alpha} K_\alpha \xrightarrow{\kappa_\alpha \mapsto \kappa_{\alpha,w,y}} \mathcal{F}_{\alpha,w,y}^{-1}(0).$$

The second map is defined *a priori* by a Newton iteration, however a more natural description *a posteriori* is that

$$\{\kappa_{\alpha,w,y}\} = \{\kappa_\alpha + \text{im } Q_{\alpha,w,y}\} \pitchfork \mathcal{F}_{\alpha,w,y}^{-1}(0)$$

is the unique (necessarily transverse) intersection in a $(k, 2, \delta, \delta)$ -neighborhood of zero of fixed size. Note that in this neighborhood of fixed size, $\mathcal{F}_{\alpha,w,y}^{-1}(0)$ is a (highly differentiable) submanifold, since $\mathcal{F}_{\alpha,w,y}$ is highly differentiable and $D_{\kappa_{\alpha,w,y}} \mathcal{F}_{\alpha,w,y}$ is surjective, as observed in the proof of Lemma C.10.2. From this description of the second map, it is clearly differentiable. Thus the derivative of the restricted gluing map (C.13.15) at a given κ is given by

$$(C.13.16) \quad T_\kappa K \xrightarrow{\dot{\kappa} \mapsto \dot{\kappa}_\alpha} T_{\kappa_\alpha} K_\alpha \xrightarrow{\text{proj } Q_{\alpha,w,y}} T_{\kappa_{\alpha,w,y}} \mathcal{F}_{\alpha,w,y}^{-1}(0) = \ker D_{\kappa_{\alpha,w,y}} \mathcal{F}_{\alpha,w,y}.$$

The second map is “projection with respect to $Q_{\alpha,w,y}$ ”, ie the map induced by identifying both the domain and codomain with $W^{k,2,\delta,\delta}(C_\alpha, u_\alpha^*TM)/\text{im } Q_{\alpha,w,y}$. It suffices to study the derivative at $\kappa = 0$, namely the first line of the following commuting

diagram:

$$(C.13.17) \quad \begin{array}{ccc} T_0 K & \xrightarrow{\dot{\kappa} \mapsto \dot{\kappa}_\alpha} & T_0 K_\alpha & \xrightarrow{\text{proj}_{\mathcal{Q}_{\alpha,w,y}}} & \ker D_{(0)\alpha,w,y} \mathcal{F}_{\alpha,w,y} \\ & & \searrow & & \swarrow \\ & & \text{proj}_{\mathcal{Q}_{\alpha,w,y}} & & \text{proj}_{\mathcal{Q}_{\alpha,w,y}} \\ & & & & \ker D_{\alpha,w,y} \end{array}$$

The rightmost diagonal map is orientation-preserving (ie it commutes with the identification of the orientation lines of the domain and codomain with $\mathfrak{o}_{E_I} \otimes \mathfrak{o}_{p,r}$), because $\text{proj}_{\mathcal{Q}_{\alpha,w,y}}$ gives a local trivialization of the bundle $\ker D \mathcal{F}_{\alpha,w,y}$ with respect to which the orientation is constant by definition. Now the leftmost diagonal map is just $\dot{\kappa}_\alpha \mapsto \dot{\kappa}_\alpha - \mathcal{Q}_{\alpha,w,y} D_{\alpha,w,y} \dot{\kappa}_\alpha$. Hence it suffices to show that the composition (note that we have changed notation from $\dot{\kappa}$ back to κ)

$$(C.13.18) \quad \ker D_{0,0,0} = K \xrightarrow{\kappa \mapsto \kappa_\alpha} K_\alpha \xrightarrow{\kappa_\alpha \mapsto \kappa_\alpha - \mathcal{Q}_{\alpha,w,y} D_{\alpha,w,y} \kappa_\alpha} \ker D_{\alpha,w,y}$$

acts as (C.13.5) on orientations (with respect to the previously defined isomorphisms $\mathfrak{o}_K = \mathfrak{o}_{E_I} \otimes \mathfrak{o}_{p,q} \otimes \mathfrak{o}_{q,r}$ and $\mathfrak{o}_{\ker D_{\alpha,w,y}} = \mathfrak{o}_{E_I} \otimes \mathfrak{o}_{p,r}$).

We analyze (C.13.18) as follows. For convenience, we may as well assume that $w = 0$, and recall that necessarily $y = 0$. Now let us consider the following diagram:

$$(C.13.19) \quad \begin{array}{ccccc} K_{\alpha,0,0} & \xleftrightarrow{\text{proj}_{T_{\alpha,0,0}}} & W^{k,2,\delta,\delta}(C_\alpha, u_\alpha^* TM) \oplus E_I & \xrightarrow{D_{\alpha,0,0}} & W^{k-1,2,\delta,\delta}(\tilde{C}_\alpha, \Omega_{\tilde{C}_\alpha, j_0}^{0,1} \otimes_{\mathbb{C}} u_\alpha^* TM_{J(\ell_{\alpha,0}(A_0(\cdot)))}) \\ & & \uparrow \text{glue} & & \downarrow \text{break} \\ K_{0|\alpha,0,0} & \xleftrightarrow{\text{proj}_{T_{0|\alpha,0,0}}} & W^{k,2,\delta,\delta}(C_0, u_{0|\alpha}^* TM) \oplus E_I & \xrightarrow{D_{0|\alpha,0,0}} & W^{k-1,2,\delta,\delta}(\tilde{C}_0, \Omega_{\tilde{C}_0, j_0}^{0,1} \otimes_{\mathbb{C}} u_{0|\alpha}^* TM_{J(\ell_{0,0}(A_0(\cdot)))}) \\ & \nearrow \kappa \mapsto \kappa_{0|\alpha} & \uparrow \text{PT} & & \downarrow \text{PT} \\ K_{0,0,0} & \xleftrightarrow{\text{proj}_{\mathcal{Q}_{0,0,0}}} & W^{k,2,\delta,\delta}(C_0, u_0^* TM) \oplus E_I & \xrightarrow{D_{0,0,0}} & W^{k-1,2,\delta,\delta}(\tilde{C}_0, \Omega_{\tilde{C}_0, j_0}^{0,1} \otimes_{\mathbb{C}} u_0^* TM_{J(\ell_{0,0}(A_0(\cdot)))}) \\ & & \nwarrow \mathcal{Q}_{0,0,0} & & \end{array}$$

The right half is just a copy of (C.7.8) (with the middle square collapsed since $y = 0$). The left half consists of the inclusions of the kernels of the operators on the right half,

as well as the projection maps associated to the images of the (approximate) right inverses $Q_{0,0,0}$,

$$T_{0|\alpha,0,0} := \text{PT} \circ Q_{0,0,0} \circ \text{PT},$$

and

$$T_{\alpha,0,0} = \text{glue} \circ \text{PT} \circ Q_{0,0,0} \circ \text{PT} \circ \text{break}$$

(recall that $\text{im } T_{\alpha,0,0} = \text{im } Q_{\alpha,0,0}$). The diagonal map $\kappa \mapsto \kappa_{0|\alpha}$ is defined by cutting off as in Definition C.5.2; thus $\kappa_{\alpha} = \text{glue}(\kappa_{0|\alpha})$.

Now the map (C.13.18) which we would like to analyze may be written as the composition $\kappa \mapsto \text{proj}_{T_{\alpha,0,0}}(\text{glue}(\kappa_{0|\alpha}))$ from (C.13.19). Now we know that

$$\|D_{0|\alpha,0,0}\kappa_{0|\alpha}\|_{k-1,2,\delta,\delta} = \|D_{\alpha,0,0}\kappa_{\alpha}\|_{k-1,2,\delta,\delta},$$

which is small by Lemma C.6.2. It follows that the map we would like to analyze is well-approximated (as $\alpha \rightarrow 0$) by the map

$$\kappa \mapsto \text{proj}_{T_{\alpha,0,0}}(\text{glue}(\text{proj}_{T_{0|\alpha,0,0}} \kappa_{0|\alpha})).$$

In particular, since we are only interested in its action on orientations, it suffices to consider the latter map

$$\kappa \mapsto \text{proj}_{T_{\alpha,0,0}}(\text{glue}(\text{proj}_{T_{0|\alpha,0,0}} \kappa_{0|\alpha})).$$

Now we claim that the map $K_{0,0,0} \rightarrow K_{0|\alpha,0,0}$ given by $\kappa \mapsto \text{proj}_{T_{0|\alpha,0,0}} \kappa_{0|\alpha}$ preserves orientation (the orientation lines of the domain and codomain are both identified with $\mathfrak{o}_{E_I} \otimes \mathfrak{o}_{p,q} \otimes \mathfrak{o}_{q,r}$). To see this, simply observe the maps (as well as their domains and codomains and their orientations) vary continuously over $\alpha \in [0, \epsilon]$, and that the statement is true for $\alpha = 0$ because then the map is the identity map. Hence we have reduced the problem to showing that the map

$$(C.13.20) \quad W^{k,2,\delta,\delta}(C_0, u_{0|\alpha}^* TM) \supseteq K_{0|\alpha,0,0} \xrightarrow{\text{glue}} W^{k,2,\delta,\delta}(C_{\alpha}, u_{\alpha}^* TM) \oplus E_I \xrightarrow{\text{proj}_{T_{\alpha,0,0}}} K_{\alpha,0,0}$$

acts by (C.13.5).

To analyze the map


$$\text{proj}_{T_{\alpha,0,0}} \circ \text{glue}: K_{0|\alpha,0,0} \rightarrow K_{\alpha,0,0},$$

argue as follows. For any finite-dimensional vector space F and a map

$$\Lambda: F \rightarrow W^{k-1,2,\delta,\delta}(\tilde{C}_0, \Omega_{\tilde{C}_0, j_0}^{0,1} \otimes_{\mathbb{C}} u_0^* TM_{J(\ell_{0,0}(A_0(\cdot)))})$$

supported away from the ends, we may consider the following modified version of (C.13.19):

(C.13.21)



Here $t \in [0, 1]$ indicates that the terms (C.5.6)–(C.5.8) carry a factor of t . The estimates from Lemmas C.7.2 and C.7.4 apply to this modified diagram as well. Thus as long as we fix a bounded right inverse $Q_{0,0,0}^{\Lambda,t}$ of $D_{0,0,0}^{\Lambda,t}$, the rest of the diagram makes sense (and, in particular, $T_{0|\alpha,0,0}^{\Lambda,t}$ and $T_{\alpha,0,0}^{\Lambda,t}$ are approximate right inverses) for sufficiently small $\alpha > 0$. Furthermore, in any family of $(\Lambda, t, Q^{\Lambda,t})$, the kernels $K_{0,0,0}, K_{0|\alpha,0,0}, K_{\alpha,0,0}$ form vector bundles, and the identifications of their orientation lines with $\mathfrak{o}_F \otimes \mathfrak{o}_{p,q} \otimes \mathfrak{o}_{q,r} \otimes \mathfrak{o}_{E_I}$ and $\mathfrak{o}_F \otimes \mathfrak{o}_{p,r} \otimes \mathfrak{o}_{E_I}$ (respectively) vary continuously.

Now the map

$$(C.13.22) \quad K_{0|\alpha,0,0}^{\Lambda,t} \xrightarrow{\text{proj}_{T_{\alpha,0,0}^{\Lambda,t}} \circ \text{glue}} K_{\alpha,0,0}^{\Lambda,t}$$

is exactly the map we wish to analyze when $F = 0, \Lambda = 0, t = 1$, and $Q_{0,0,0}^{\Lambda,t} = 0 \oplus Q_{0,0,0}$. More generally, if we allow F nonzero (but still $\Lambda = 0$), this map is simply our desired map plus the identity map on F . Since the space of acceptable maps Λ is contractible, it suffices to show that (C.13.22) has the desired action on orientations for any single pair $(F, \Lambda), t = 1$, and $Q_{0,0,0}^{\Lambda,t} = 0 \oplus Q_{0,0,0}$.

Now by compactness of $[0, 1]$, there exists a pair (F, Λ) so that $\Lambda \oplus D_{0,0,0}^t$ is surjective for all $t \in [0, 1]$. Fix such a pair (F, Λ) , and also fix a continuously varying family

of bounded right inverses $Q_{0,0,0}^{\Lambda,t}$ with $Q_{0,0,0}^{\Lambda,1} = 0 \oplus Q_{0,0,0}$. Since the kernels form a bundle over the base $[0, 1]$ and their orientations vary continuously, it suffices to analyze (C.13.22) for this (F, Λ) , $t = 0$, and this $Q_{0,0,0}^{\Lambda,0}$.

Now the Floer–Hofer map (C.13.5) is *defined* (see [25, Proposition 9]) by the property that it is induced by a certain map $K_{0|\alpha,0,0}^{\Lambda,0} \rightarrow K_{\alpha,0,0}^{\Lambda,0}$ closely related to (C.13.22); clearly it suffices to show that the difference between the Floer–Hofer map and (C.13.22) is very small. Consider the map

$$(C.13.23) \quad K_{0|\alpha,0,0}^{\Lambda,t} \xrightarrow{\text{proj}_{L^2} \circ \text{glue}} K_{\alpha,0,0}^{\Lambda,t}$$

in the setting of the above choices of (F, Λ) , $t = 0$, and $Q_{0,0,0}^{\Lambda,t}$, where instead of projecting off of $\text{im } T_{\alpha,0,0}^{\Lambda,0}$, we use the L^2 -orthogonal projection. The difference between (C.13.23) and (C.13.22) is exactly

$$\text{proj}_{L^2} \circ Q_{\alpha,0,0}^{\Lambda,0} \circ (\Lambda \oplus D_{\alpha,0,0}^0) \circ \text{glue},$$

which has small norm by the estimate in Lemma C.7.4. Thus it suffices to compare (C.13.23) to the Floer–Hofer map. Now the Floer–Hofer map is given by

$$(C.13.24) \quad K_{0|\alpha,0,0}^{\Lambda,0} \xrightarrow{\text{proj}_{L^2} \circ \text{glue}' } K_{\alpha,0,0}^{\Lambda,0}$$

for a certain map glue' (see [25, Proposition 9]). However, the norm of the difference $\text{glue} - \text{glue}'$ is very small due to the exponential decay of elements of the kernel (here, we may use the explicit description of $K_{0|\alpha,0,0}$ as the image of $\kappa \mapsto \kappa_{0|\alpha} - Q_{0|\alpha,0,0} D_{0|\alpha,0,0} \kappa_{0|\alpha}$). Thus we are done. \square

References

- [1] **M Abouzaid**, *Framed bordism and Lagrangian embeddings of exotic spheres*, Ann. of Math. 175 (2012) 71–185 MR2874640
- [2] **M Abouzaid**, *Nearby Lagrangians with vanishing Maslov class are homotopy equivalent*, Invent. Math. 189 (2012) 251–313 MR2947545
- [3] **J F Adams**, *On the cobar construction*, Proc. Nat. Acad. Sci. U.S.A. 42 (1956) 409–412 MR0079266
- [4] **F J Almgren, Jr**, *The homotopy groups of the integral cycle groups*, Topology 1 (1962) 257–299 MR0146835
- [5] **M F Atiyah, I G Macdonald**, *Introduction to commutative algebra*, Addison-Wesley, Reading, MA (1969) MR0242802
- [6] **J-F Barraud, O Cornea**, *Lagrangian intersections and the Serre spectral sequence*, Ann. of Math. 166 (2007) 657–722 MR2373371

- [7] **NA Berikashvili**, *The Steenrod–Sitnikov homology theory on the category of compact spaces*, Dokl. Akad. Nauk SSSR 254 (1980) 1289–1291 MR592491
- [8] **NA Berikashvili**, *Axiomatics of the Steenrod–Sitnikov homology theory on the category of compact Hausdorff spaces*, Trudy Mat. Inst. Steklov. 154 (1983) 24–37 MR733823
In Russian; translated in Proc. Steklov Inst. Math. 154 (1984) 25–39
- [9] **R H Bing**, *The cartesian product of a certain nonmanifold and a line is E^4* , Ann. of Math. 70 (1959) 399–412 MR0107228
- [10] **D Borisov, J Noel**, *Simplicial approach to derived differential manifolds*, preprint (2011) arXiv:1112.0033
- [11] **J W Cannon**, *Shrinking cell-like decompositions of manifolds: codimension three*, Ann. of Math. 110 (1979) 83–112 MR541330
- [12] **G Carlsson, E K Pedersen**, *Čech homology and the Novikov conjectures for K - and L -theory*, Math. Scand. 82 (1998) 5–47 MR1634649
- [13] **G S Chogoshvili**, *On the equivalence of the functional and spectral theory of homology*, Izvestiya Akad. Nauk SSSR. Ser. Mat. 15 (1951) 421–438 MR0046649
- [14] **R L Cohen, J D S Jones, G B Segal**, *Floer’s infinite-dimensional Morse theory and homotopy theory*, from: “The Floer memorial volume”, (H Hofer, C H Taubes, A Weinstein, E Zehnder, editors), Progr. Math. 133, Birkhäuser, Basel (1995) 297–325 MR1362832
- [15] **C C Conley, E Zehnder**, *The Birkhoff–Lewis fixed point theorem and a conjecture of VI Arnol’d*, Invent. Math. 73 (1983) 33–49 MR707347
- [16] **M Damian**, *On the stable Morse number of a closed manifold*, Bull. London Math. Soc. 34 (2002) 420–430 MR1897421
- [17] **A Dress**, *Zur Spectralsequenz von Faserungen*, Invent. Math. 3 (1967) 172–178 MR0267585
- [18] **D A Edwards, H M Hastings**, *Čech and Steenrod homotopy theories with applications to geometric topology*, Lecture Notes in Mathematics 542, Springer, Berlin (1976) MR0428322
- [19] **T Ekholm, I Smith**, *Exact Lagrangian immersions with one double point revisited*, Math. Ann. 358 (2014) 195–240 MR3157996
- [20] **T Ekholm, I Smith**, *Exact Lagrangian immersions with a single double point*, J. Amer. Math. Soc. 29 (2016) 1–59 MR3402693
- [21] **Y Eliashberg**, *Estimates on the number of fixed points of area preserving transformations*, preprint, Syktyvar University (1979)
- [22] **A Floer**, *A relative Morse index for the symplectic action*, Comm. Pure Appl. Math. 41 (1988) 393–407 MR933228

- [23] **A Floer**, *The unregularized gradient flow of the symplectic action*, Comm. Pure Appl. Math. 41 (1988) 775–813 MR948771
- [24] **A Floer**, *Symplectic fixed points and holomorphic spheres*, Comm. Math. Phys. 120 (1989) 575–611 MR987770
- [25] **A Floer, H Hofer**, *Coherent orientations for periodic orbit problems in symplectic geometry*, Math. Z. 212 (1993) 13–38 MR1200162
- [26] **A Floer, H Hofer, K Wysocki**, *Applications of symplectic homology, I*, Math. Z. 217 (1994) 577–606 MR1306027
- [27] **E M Friedlander**, *Étale homotopy of simplicial schemes*, Annals of Mathematics Studies 104, Princeton Univ. Press (1982) MR676809
- [28] **K Fukaya**, *The symplectic s -cobordism conjecture: a summary*, from: “Geometry and physics”, (J E Andersen, J Dupont, H Pedersen, A Swann, editors), Lecture Notes in Pure and Appl. Math. 184, Dekker, New York (1997) 209–219 MR1423167
- [29] **K Fukaya, Y-G Oh, H Ohta, K Ono**, *Lagrangian intersection Floer theory: Anomaly and obstruction, I*, AMS/IP Studies in Advanced Mathematics 46, Amer. Math. Soc., Providence, RI (2009) MR2553465
- [30] **K Fukaya, Y-G Oh, H Ohta, K Ono**, *Lagrangian intersection Floer theory: Anomaly and obstruction, II*, AMS/IP Studies in Advanced Mathematics 46, Amer. Math. Soc., Providence, RI (2009) MR2548482
- [31] **K Fukaya, Y-G Oh, H Ohta, K Ono**, *Technical details on Kuranishi structure and virtual fundamental chain*, preprint (2012) arXiv:1209.4410
- [32] **K Fukaya, Y-G Oh, H Ohta, K Ono**, *Kuranishi structure, pseudo-holomorphic curve, and virtual fundamental chain, Part 1*, preprint (2015) arXiv:1503.07631
- [33] **K Fukaya, K Ono**, *Arnold conjecture and Gromov–Witten invariant*, Topology 38 (1999) 933–1048 MR1688434
- [34] **R Godement**, *Topologie algébrique et théorie des faisceaux*, Actualités Sci. Ind. 1252, Hermann, Paris (1958) MR0102797
- [35] **B Goldfarb**, *Novikov conjectures for arithmetic groups with large actions at infinity*, K-Theory 11 (1997) 319–372 MR1451760
- [36] **M Gromov**, *Pseudoholomorphic curves in symplectic manifolds*, Invent. Math. 82 (1985) 307–347 MR809718
- [37] **M Gromov**, *Singularities, expanders and topology of maps, Part 2: From combinatorics to topology via algebraic isoperimetry*, Geom. Funct. Anal. 20 (2010) 416–526 MR2671284
- [38] **H M Hastings**, *Steenrod homotopy theory, homotopy idempotents, and homotopy limits*, from: “Proceedings of the 1977 Topology Conference, II”, (W Kuperberg, G M Reed, P Zenor, editors), Topology Proc. 2 (1977) 461–477 MR540623

- [39] **A Hatcher**, *Algebraic topology*, Cambridge Univ. Press (2002) MR1867354
- [40] **H Hofer, D A Salamon**, *Floer homology and Novikov rings*, from: “The Floer memorial volume”, (H Hofer, C H Taubes, A Weinstein, E Zehnder, editors), Progr. Math. 133, Birkhäuser, Basel (1995) 483–524 MR1362838
- [41] **H Hofer, K Wysocki, E Zehnder**, *A general Fredholm theory, I: A splicing-based differential geometry*, J. Eur. Math. Soc. (JEMS) 9 (2007) 841–876 MR2341834
- [42] **H Hofer, K Wysocki, E Zehnder**, *A general Fredholm theory, II: Implicit function theorems*, Geom. Funct. Anal. 19 (2009) 206–293 MR2507223
- [43] **H Hofer, K Wysocki, E Zehnder**, *A general Fredholm theory, III: Fredholm functors and polyfolds*, Geom. Topol. 13 (2009) 2279–2387 MR2515707
- [44] **H Hofer, K Wysocki, E Zehnder**, *Integration theory on the zero sets of polyfold Fredholm sections*, Math. Ann. 346 (2010) 139–198 MR2558891
- [45] **H Hofer, K Wysocki, E Zehnder**, *sc-smoothness, retractions and new models for smooth spaces*, Discrete Contin. Dyn. Syst. 28 (2010) 665–788 MR2644764
- [46] **H Hofer, K Wysocki, E Zehnder**, *Applications of polyfold theory, I: The polyfolds of Gromov–Witten theory*, preprint (2014) arXiv:1107.2097
- [47] **H Hofer, K Wysocki, E Zehnder**, *Polyfold and Fredholm theory, I: Basic theory in M -polyfolds*, preprint (2014) arXiv:1407.3185
- [48] **H Inassaridze**, *On the Steenrod homology theory of compact spaces*, Michigan Math. J. 38 (1991) 323–338 MR1116492
- [49] **D Joyce**, *D-manifolds and d-orbifolds: a theory of derived differential geometry*, Book manuscript (2012) Available at <https://people.maths.ox.ac.uk/joyce/dmbook.pdf>
- [50] **D Joyce**, *An introduction to d-manifolds and derived differential geometry*, from: “Moduli spaces”, (Cambridge, editor), London Math. Soc. Lecture Note Ser. 411, Cambridge Univ. Press (2014) 230–281 MR3221297
- [51] **M Kontsevich**, *Enumeration of rational curves via torus actions*, from: “The moduli space of curves”, (R Dijkgraaf, C Faber, G van der Geer, editors), Progr. Math. 129, Birkhäuser, Boston (1995) 335–368 MR1363062
- [52] **M Kontsevich, Y Manin**, *Gromov–Witten classes, quantum cohomology, and enumerative geometry*, Comm. Math. Phys. 164 (1994) 525–562 MR1291244
- [53] **J Li, G Tian**, *Virtual moduli cycles and Gromov–Witten invariants of general symplectic manifolds*, from: “Topics in symplectic 4-manifolds”, (R J Stern, editor), First Int. Press Lect. Ser. 1, Int. Press, Cambridge, MA (1998) 47–83 MR1635695
- [54] **G Liu, G Tian**, *Floer homology and Arnold conjecture*, J. Differential Geom. 49 (1998) 1–74 MR1642105

- [55] **R B Lockhart, R C McOwen**, *Elliptic differential operators on noncompact manifolds*, Ann. Scuola Norm. Sup. Pisa Cl. Sci. 12 (1985) 409–447 MR837256
- [56] **G Lu, G Tian**, *Constructing virtual Euler cycles and classes*, Int. Math. Res. Surv. 2007 (2007) MR2369606
- [57] **J Lurie**, *Higher topos theory* (2012) Updated online version of *Higher topos theory*, Annals of Mathematics Studies 170, Princeton Univ. Press, 2009 Available at <http://www.math.harvard.edu/~lurie/papers/croppedtopoi.pdf>
- [58] **J Lurie**, *Higher algebra* (2014) Available at <http://www.math.harvard.edu/~lurie/papers/higheralgebra.pdf>
- [59] **S Mardešić**, *Strong shape and homology*, Springer, Berlin (2000) MR1740831
- [60] **D McDuff**, *Notes on Kuranishi atlases*, preprint (2015) arXiv:1411.4306v2
- [61] **D McDuff, D Salamon**, *J-holomorphic curves and quantum cohomology*, University Lecture Series 6, Amer. Math. Soc., Providence, RI (1994) MR1286255
- [62] **D McDuff, D Salamon**, *J-holomorphic curves and symplectic topology*, Amer. Math. Soc. Colloq. Publ. 52, Amer. Math. Soc., Providence, RI (2004) MR2045629
- [63] **D McDuff, K Wehrheim**, *Kuranishi atlases with trivial isotropy*, preprint (2015) arXiv:1208.1340v6
- [64] **J Milnor**, *On the Steenrod homology theory*, from: “Novikov conjectures, index theorems and rigidity, Vol. 1”, (S C Ferry, A Ranicki, J Rosenberg, editors), London Math. Soc. Lecture Note Ser. 226, Cambridge Univ. Press (1995) 79–96 MR1388297
- [65] **M-P Muller**, *Gromov’s Schwarz lemma as an estimate of the gradient for holomorphic curves*, from: “Holomorphic curves in symplectic geometry”, (M Audin, J Lafontaine, editors), Progr. Math. 117, Birkhäuser, Basel (1994) 217–231 MR1274931
- [66] **K Ono**, *On the Arnol’d conjecture for weakly monotone symplectic manifolds*, Invent. Math. 119 (1995) 519–537 MR1317649
- [67] **M S Osborne**, *Basic homological algebra*, Graduate Texts in Mathematics 196, Springer, New York (2000) MR1757274
- [68] **J Pardon**, *The Hilbert–Smith conjecture for three-manifolds*, J. Amer. Math. Soc. 26 (2013) 879–899 MR3037790
- [69] **S Piunikhin, D Salamon, M Schwarz**, *Symplectic Floer–Donaldson theory and quantum cohomology*, from: “Contact and symplectic geometry”, (C B Thomas, editor), Publ. Newton Inst. 8, Cambridge Univ. Press (1996) 171–200 MR1432464
- [70] **J J Rotman**, *An introduction to homological algebra*, 2nd edition, Springer, New York (2009)
- [71] **Y Ruan**, *Virtual neighborhoods and pseudo-holomorphic curves*, from: “Proceedings of 6th Gökova Geometry-Topology Conference”, Turkish J. Math. 23 (1999) 161–231 MR1701645

- [72] **D Salamon**, *Lectures on Floer homology*, from: “Symplectic geometry and topology”, (Y Eliashberg, L Traynor, editors), IAS/Park City Math. Ser. 7, Amer. Math. Soc., Providence, RI (1999) 143–229 MR1702944
- [73] **D Salamon, E Zehnder**, *Morse theory for periodic solutions of Hamiltonian systems and the Maslov index*, Comm. Pure Appl. Math. 45 (1992) 1303–1360 MR1181727
- [74] **N Spaltenstein**, *Resolutions of unbounded complexes*, Compositio Math. 65 (1988) 121–154 MR932640
- [75] **D I Spivak**, *Derived smooth manifolds*, Duke Math. J. 153 (2010) 55–128 MR2641940
- [76] **N E Steenrod**, *Regular cycles of compact metric spaces*, Ann. of Math. 41 (1940) 833–851 MR0002544
- [77] **M G Sullivan**, *K-theoretic invariants for Floer homology*, Geom. Funct. Anal. 12 (2002) 810–872 MR1935550
- [78] **W P Thurston**, *The geometry and topology of three-manifolds*, lecture notes, Princeton University (1979) Available at <http://msri.org/publications/books/gt3m>
- [79] **K Wehrheim**, *Energy quantization and mean value inequalities for nonlinear boundary value problems*, J. Eur. Math. Soc. 7 (2005) 305–318 MR2156603
- [80] **K Wehrheim**, *Smooth structures on Morse trajectory spaces, featuring finite ends and associative gluing*, from: “Proceedings of the Freedman Fest”, (R Kirby, V Krushkal, Z Wang, editors), Geom. Topol. Monogr. 18 (2012) 369–450 MR3084244

Department of Mathematics, Stanford University
450 Serra Mall, Building 380, Stanford, CA 94305, USA

`pardon@math.stanford.edu`

<http://math.stanford.edu/~pardon/>

Proposed: Leonid Polterovich
 Seconded: Richard Thomas, Ronald Stern

Received: 26 May 2014
 Revised: 20 May 2015