

Parametrized ring-spectra and the nearby Lagrangian conjecture

THOMAS KRAGH

APPENDIX BY MOHAMMED ABOUZAIID

Let L be an embedded closed connected exact Lagrangian submanifold in a connected cotangent bundle T^*N . In this paper we prove that such an embedding is, up to a finite covering space lift of T^*N , a homology equivalence. We prove this by constructing a fibrant parametrized family of ring spectra \mathcal{FL} parametrized by the manifold N . The homology of \mathcal{FL} will be the (twisted) symplectic cohomology of T^*L . The fibrancy property will imply that there is a Serre spectral sequence converging to the homology of \mathcal{FL} . The fiber-wise ring structure combined with the intersection product on N induces a product on this spectral sequence. This product structure and its relation to the intersection product on L is then used to obtain the result. Combining this result with work of Abouzaid we arrive at the conclusion that $L \rightarrow N$ is always a homotopy equivalence.

53D12; 55R70, 55T10

1 Introduction

Let N be any closed smooth connected manifold, and let $j: L \subset T^*N$ be a connected closed exact Lagrangian submanifold. Let $\pi: T^*N \rightarrow N$ be the obvious projection and define $p = \pi \circ j: L \rightarrow N$. The Nearby Lagrangian Conjecture states that L is Hamiltonian isotopic to the zero section. In [9], Fukaya, Seidel and Smith introduced new powerful methods, and there is a good summary of what was known at that time about this conjecture.

An easy to prove property of exact Lagrangians is that if $f: N' \rightarrow N$ is a smooth covering space then the associated pull-back

$$\begin{array}{ccc} L' & \xrightarrow{j'} & T^*N' \\ \downarrow & & \downarrow T^*f \\ L & \xrightarrow{j} & T^*N \end{array}$$

defines an exact Lagrangian $j': L' \subset T^*N'$. If L' is not connected, we may make a choice of connected component. Combining this with the result by Lalonde and Sikorav [13] that p is always surjective (which also works when N is not closed) one easily gets that $\pi_1(N)/p_*(\pi_1(L))$ is finite. Note that this surjectivity result works when L is assumed closed but not N , which means that this possibility is in fact excluded. So the assumption that N is closed is in fact superfluous, but we include it so as not to cause confusion.

Main Theorem *If in addition N is oriented and the induced map p_* on fundamental groups is surjective then the induced map p_* on homology is an isomorphism.*

To see just how restrictive this is in general we prove a conjecture by Arnold on the degree of p . This conjecture is stated without the additional assumptions of the Main Theorem.

Corollary 1.1 *The degree of p is nonzero.*

Proof Let N be any (possibly nonorientable) smooth connected manifold, and let $L \subset T^*N$ be as above. By lifting to the *finite* covering of N associated to the image $p_*(\pi_1(L))$ intersected with the kernel of the first Stiefel–Whitney class of N we get a lifted exact Lagrangian $j': L' \subset T^*N'$. The actual lift could have 2 connected components, but as above we let L' be one of these components. By the definition of degree (see Epstein [8]) we have that $p' = \pi' \circ j': L' \rightarrow N'$ has nonzero degree if and only if p has nonzero degree.

By construction this lift satisfies the assumptions of the Main Theorem and so p' has degree one. \square

We in fact see that the Main Theorem is much stronger because it tells us that up to a finite lift (and restricting to a component) the map p is a homology equivalence, and on all further finite lifts it is also a homology equivalence (this is not true for any homology equivalence so this strengthens the statement).

In the case of vanishing Maslov class and N simply connected Fukaya, Seidel and Smith proved in [9] that p induces a homology isomorphism. Later in [4], Abouzaid removed the assumption on N and strengthened this to a homotopy equivalence. However, the assumption on the Maslov class remained. Only a few, and not very general, results were known about the nonvanishing case. However, in Appendix E, Abouzaid uses Corollary 1.2 below to prove that in fact this assumption is always true. So as Theorem E.2 states we now know that $L \rightarrow N$ is a homotopy equivalence.

Corollary 1.2 *Lifting $L \rightarrow T^*N$ to the universal cover $N' \rightarrow N$, the Maslov class of a component $L' \rightarrow T^*N'$ vanishes.*

Note that the universal cover of N may not be closed and thus L' may not be closed either.

Proof Since we have established that some partial lift $L'' \rightarrow T^*N''$ is a homology equivalence, we have that in the diagram

$$\begin{array}{ccc}
 \pi_1(L'') & \xrightarrow{p_*} & \pi_1(N'') \\
 \downarrow & & \downarrow \\
 H_1(L'') & \xrightarrow[p_* \cong]{} & H_1(N'') \\
 \downarrow m_L & \swarrow \text{dotted} & \\
 \mathbb{Z} & \xleftarrow{k} &
 \end{array}$$

the dashed arrow exists, making it commutative. Here the top vertical arrows are the Abelianization maps, and m_L is given by evaluating the Maslov class. This proves that m_L vanishes on the kernel of $p_*: \pi_1(L'') \rightarrow \pi_1(N'')$. So since the Maslov index of any loop in the cover L' is the same as the Maslov index of its projection to L , the corollary follows. □

The proof of the Main Theorem uses parts of Viterbo’s transfer map construction in [21], which we constructed as a map of spectra in [12]. The idea is that the target spectrum, which represents the (twisted)¹ symplectic cohomology of T^*L , can be created as a fibered spectrum \mathcal{FL} over the base N and in fact turns out to be fibrant.² This means that the homology of the fibers form a local system. It also implies the existence of a Serre spectral sequence with page 2 given by the homology of N with coefficients in the homology of the fibers, strongly converging to the homology of \mathcal{FL} , ie the (twisted) symplectic cohomology of T^*L . There are also natural product structures on the fibers (continuous over the base) and when combined with the intersection product on N we get a product on the spectral sequence. The idea that such a spectral sequence might exist (and be important) was inspired by a description of a spectral sequence in [9] by Fukaya, Seidel and Smith.

To construct this fibered spectrum \mathcal{FL} we start in Section 2 by introducing finite-dimensional approximations of the action integral A , which is defined on paths

¹“Twisted” is explained in Remark 1.3 and Appendix D.

²Our notion of fibrant is defined in Section 9, but it is close to Serre fibrant.

$\gamma: I \rightarrow T^*N$ by

$$(1) \quad A(\gamma) = \int_{\gamma} \lambda - H dt.$$

Here λ is the canonical Liouville 1-form given by

$$\lambda_{q,p}(v) = p(\pi_*(v)), \quad q \in N, p \in T_q^*N,$$

and $H: T^*N \rightarrow \mathbb{R}$ is a Hamiltonian. We will restrict the class of Hamiltonians that we will consider. These finite-dimensional approximations are versions of Chaperon’s broken geodesic approach in [5], which is similar to the generating functions used by Viterbo in [21]. So, the finite-dimensional approximations will be functions defined on finite-dimensional manifolds depending on an $r \in \mathbb{N}$ with $r \gg 1$ depending on the Hamiltonian H and other structure. These functions will (when properly handled) have well-defined Morse homologies, which will equal the (twisted) Floer homology of A . To define symplectic cohomology one really only needs to consider closed loops in T^*N . However, the fiber over $q \in N$ of our spectrum is constructed by considering paths starting and ending on the Lagrangian T_q^*N , ie a different Lagrangian boundary condition.

There are some difficulties concerning orientations, which we summarize in the following remark. For concreteness (and because this is what we compute with our spectrum) we describe this in terms of symplectic cohomology. However, the same difference in orientations is there for the Morse homology of the action and the finite-dimensional approximations. Indeed, symplectic cohomology is a special case of the Floer (co)homology associated to a Hamiltonian quadratic at infinity (see Abbondandolo and Schwarz [1] or Salamon and Weber [18]).

Remark 1.3 When defining symplectic cohomology there are some orientation discrepancies (sign errors) in the literature. To explain this we assume that N and L are oriented. Until recently it was thought that this implied that the symplectic cohomology $SH^*(TN)$ was isomorphic to $H_*(\Lambda N)$. Indeed, this was established in [20], [1] and [18]. There is a slight error in this. The fact is that if the second Stiefel–Whitney class of N does not vanish on $\pi_2(N)$ this is false. Indeed, in [12] it was proven that Viterbo functoriality in the case of $L \subset T^*N$ can be realized as a map of Thom-spectra

$$(2) \quad (\Lambda N)^{TN} \rightarrow (\Lambda L)^{TL \oplus \eta},$$

where $(\cdot)^\beta$ denotes Thom-spectrum construction for a bundle β , and η is a virtual bundle, which is *not oriented* unless the relative Stiefel–Whitney class vanishes on $\pi_2(L)$ (this last fact is proved in Corollary 7.5 combined with Remark 7.6 and not in [12]). For the reader mostly concerned with (co)homology this simply means that the homology of

the target is not the homology of ΛL , but a twisted version of it (basically because the Thom isomorphism only works for oriented bundles). This is contradictory to previous results since Viterbo functoriality tells us that symplectic cohomology is natural with respect to restrictions (as Paul Seidel and Mohammed Abouzaid pointed out to us). Indeed, this bundle η was thought to *be* oriented due to the previous results. The only solution to this problem is that the transgression of the second Stiefel–Whitney class enters in when defining signs on the differentials of the Floer complex. As a test Paul Seidel produced a calculation that $SH^*(T^*\mathbb{C}P^2)$ in fact vanishes with rational coefficients, and this can be explained by it being 2-torsion with \mathbb{Z} coefficients, as one would not previously have expected, but with the correct signs this is the correct answer. These new signs in the differentials has further been verified by Abouzaid in [3, Appendix A].

For the purpose of this paper we mention here that for general $L \subset T^*N$ there are three potentially different orientation choices that seem important in this discussion. We will for simplicity also assume that L is orientable. The first two are only dependent on $L \subset T^*L$ and not how it is embedded into T^*N . As explained in Appendix D, different finite-dimensional approximations may give different results.

- As noticed by Viterbo in [21], in the finite-dimensional approximations for $L \subset T^*L$, the zero-section does in fact reproduce the loop space homology. Ie, there is no transgression of the second Stiefel–Whitney class.
- The symplectic cohomology of T^*L reproduces the loop space homology twisted by the canonical transgression of the second Stiefel–Whitney class to $\pi_1(\Lambda L)$, as described by Abouzaid in [3].
- The finite-dimensional approximation of the action provided by $DT^*L \subset T^*N$ gives a potentially third result, which is loop space homology twisted by the transgression of the *relative* second Stiefel–Whitney class of $L \rightarrow N$ transgressed to $\pi_1(\Lambda L)$. This is the homology of the target spectrum in (2) above, and is also the homology of \mathcal{FL} appearing in this paper, and the one used in the main argument.

The reason for these differences is that the orientation in symplectic homology does not generally agree with the one defined by finite-dimensional approximations, which for this purpose can be thought of as depending on a choice of Lagrangian foliation. This is explained in more detail in Appendix D.

Section 3 and Section 4 define parametrized pseudogradients and parametrized Conley index theory, which can be thought of as a version of parametrized Morse theory on noncompact manifolds, which do not suffer from the defect that one actually needs

the functions to be Morse. And for the finite-dimensional approximations we thus get parametrized Conley indices over the base N . The fact that we do not need Morse means that the fibers are defined even in the nongeneric cases.

In Section 5 we relate the finite-dimensional approximations for different r . The relation can heuristically be thought of as the idea that the fibers of the parametrized Conley index for $r + 1$ are given by a reduced suspension of the fibers of the parametrized Conley index for r . This is not completely true because the fibers may behave slightly bad due to the nongenericity that we cannot avoid in a parametrized setting. However, in the end these problems will disappear when considering \mathcal{FL} because it is fibrant.

In Section 6 we use the suspension result from Section 5 to construct fibered spectra over N whose homology (ie, the homology of the total space) will equal the (twisted) Floer homology associated to the action A . The homologies of the fibers will *not* form local systems, ie these fibered spectra will *not* be *fibrant*. In fact, we will only see this special feature for \mathcal{FL} because it is defined carefully as a certain limit (see below). This fibrancy can in fact be generalized to fibered spectra computing the (twisted) symplectic cohomology of any compact exact sub-Liouville domain in T^*N , but we will not consider this here.

In Section 7 we define \mathcal{FL} as a limit when $s \rightarrow \infty$ of the fibered spectra defined in Section 6 given by a family of Hamiltonians H^s , $s \in [0, \infty[$. This is a refinement of the idea that symplectic cohomology of DT^*L as a symplectic neighborhood of L can be defined by the limit of Floer homologies for some Hamiltonians that go to infinity on the complement of DT^*L and stay 0 on L . To get all the structure we need – most importantly fibrancy – we have to be very careful in defining these Hamiltonians, and this is done in Appendix A.

In Section 8 we construct parallel transport in parametrized Conley indices. We need this to properly define the continuation maps increasing s (ie, the maps over which we took the limit) in Section 7 as fibered maps (called *ex*-maps in the paper) over N . We also need this at several points later to prove fibrancy and construct other structures on \mathcal{FL} .

In Section 9 and Section 10 we use this parallel transport to construct (stable) homotopy lifts in \mathcal{FL} and prove that the homology of the fibers form a local system and that we have an associated Serre spectral sequence as mentioned above.

Section 11 through Section 14 concern structures such as products and a careful analysis of the constant loops on L , which is then related to the spectral sequence. The product is a generalization of the Chas–Sullivan product (and thus probably related to the pair of pants products). This product on the spectral sequence can be considered a

generalization of the product constructed in [6] by Cohen, Jones and Yan. In fact, when $L = N$ is the zero section we recover their spectral sequence.

All these structures are needed for the main argument, which is carried out in Section 15. Since this is great motivation for all the structure and is readable without understanding much of the paper we invite the reader to skip ahead and to read the main argument before continuing.

Acknowledgments

The author was Supported by Carlsberg-fondet, MIT, Topology in Norway, and the University of Oslo. I would like to thank John Rognes for many discussions on various subjects involved in the construction, and Mohammed Abouzaid for many discussions and suggestions, which lead me to rewrite the paper in a more transparent fashion. I would also like to thank the anonymous referee for several suggestions. Of course any nontransparency and errors are due solely to me.

2 Finite-dimensional approximation of the action

In this section we will define smooth functions on finite-dimensional manifolds approximating the action in (1) on the space of paths restricted to:

- (1) The free loop space ΛT^*N of *closed* paths in T^*N .
- (2) The path space $\Omega_q T^*N$ of paths in T^*N starting and ending in the fiber T_q^*N over a point $q \in N$.

In contrast to the infinite-dimensional cases each approximation in case (2) will simply be the restriction of an approximation in case (1), and so the approximations will be compatible with each other. Indeed, this is because to approximate the action on closed loops we in some sense insert discontinuities, which are similar to the one of restricting the start and end point to a fiber. These approximations will of course depend on a Hamiltonian H as the action does, but they will also depends on a subdivision α of the unit interval. In fact they will smoothly depend on the pair (H, α) . We will only need to consider a very restricted class of Hamiltonians. Indeed, we will only consider Hamiltonians linear at infinity with a slight positive slope.

We fix a Riemannian structure on N once and for all. Throughout the paper let δ_0 denote some fixed constant such that $2\delta_0$ is smaller than the injective radius of N (recall that we assume N to be closed).

Definition 2.1 The subspace $\mathcal{H} \subset C^\infty(T^*N, \mathbb{R})$ of *admissible* Hamiltonians is defined by $H \in \mathcal{H}$ iff

$$(3) \quad H(q, p) = \mu\|p\| + c \quad \text{when } \|p\| \geq 1$$

with $c \in \mathbb{R}$ and $\mu \in]0, \delta_0]$.

The slope μ is significant since this implies that the Hamiltonian flow outside DT^*N never starts and ends in the same fiber. Indeed, the Hamiltonian flow coincides with the geodesic flow with constant speed equal to this slope at infinity (see eg [12, Section 3]).

We let C_1^H denote the supremum of the norm of the covariant derivative (gradient) of H on the compact set DT^*N . Similarly let C_2^H denote the supremum of the norm of the second order covariant derivative on DT^*N . These are both continuous as functions of $H \in \mathcal{H}$. Note that C_1^H is a bound on the gradient on all of T^*N because of the linearity assumption outside of DT^*N .

The Riemannian structure on N provides a canonical vector bundle isomorphism between the tangent bundle TN and the cotangent bundle T^*N . We will without further warning suppress this isomorphism from the notation. However, we use the notation T^*N to denote the canonical symplectic manifold of dimension $2d$, and we use the notation TN to denote the vector bundle over N . The Riemannian structure on N induces a Riemannian structure on T^*N compatible with the symplectic structure (see eg [12, Section 3]). We also get induced Riemannian structures on the manifolds defined below, by considering them submanifolds of the r -fold products N^r and $(T^*N)^r = T^*(N^r)$, respectively. We define the finite-dimensional manifold approximation to free loops in N by

$$\Lambda_r N = \{\vec{q} = (q_j)_{j \in \mathbb{Z}_r} \in N^r \mid \text{dist}(q_j, q_{j+1}) < \delta_0\}.$$

This is a manifold of dimension rd and we see that its cotangent bundle is

$$T^* \Lambda_r N = \{(q_j, p_j)_{j \in \mathbb{Z}_r} \in (T^*N)^r \mid \text{dist}(q_j, q_{j+1}) < \delta_0\}.$$

We will use the notation $\vec{z} = (\vec{q}, \vec{p})$ for a point in this space. Compatible with this, we use the notation $z_j = (q_j, p_j)$ for a single “coordinate” in T^*N .

We define the evaluations $\text{Ev}_i: T^* \Lambda_r N \rightarrow N$ by $\text{Ev}_i(\vec{z}) = q_i$. We then define

$$T^\bullet \Omega_{q,r} N = \text{Ev}_0^{-1}(q).$$

This is the submanifold where q_0 is fixed. It is the cotangent bundle above restricted to the $(r - 1)d$ -dimensional submanifold

$$\Omega_{q,r} N = \{\vec{q} \in \Lambda_r N \mid q_0 = q\} \subset \Lambda_r N.$$

Obviously it is not the cotangent bundle because we have the extra cotangent vector p_0 , hence the \bullet notation.

The approximations will depend on a subdivision of the interval $I = [0, 1]$ into r subintervals. These are represented by points in the simplex Δ^{r-1} defined by

$$\Delta^{r-1} = \left\{ \alpha = (\alpha_0, \dots, \alpha_{r-1}) \in \mathbb{R}^r \mid \alpha_j \geq 0 \ \forall j, \sum_j \alpha_j = 1 \right\}.$$

The correspondence is such that the length of the j^{th} interval is α_{j-1} . We define the length l of a subdivision α as the longest interval in the subdivision given by $l(\alpha) = \max_j \alpha_j$.

Assume that $\alpha \in \Delta^{r-1}$ is a subdivision and $\gamma: I \rightarrow T^*N$ is a smooth path starting and ending in the same fiber T_q^*N . Also assume that $\|\gamma'(t)\| l(\alpha) < \delta_0$ for all $t \in I$. We may then define the α -dissection of γ to be the point

$$(4) \quad \vec{z} = (z_j)_{j \in \mathbb{Z}_r} \in T^* \Lambda_r N \quad \text{given by } z_j = \gamma \left(\sum_{i=0}^{j-1} \alpha_i \right)$$

for $j = 0, \dots, r - 1$. The assumption on the length makes the distance $\text{dist}(z_j, z_{j-1})$ less than δ_0 so that this is well-defined. Note that this is well-defined even though the end points of γ are far apart because their projection to N is the same. Indeed, this shows that

$$\begin{aligned} \text{dist}(q_{r-1}, q_0) &= \text{dist}(q_{r-1}, \pi(\gamma(0))) = \text{dist}(q_{r-1}, \pi(\gamma(1))) \\ &\leq \text{dist}(z_{r-1}, \gamma(1)) = \text{dist}(\gamma(1 - \alpha_{r-1}), \gamma(1)) < \delta_0. \end{aligned}$$

We will be α -dissecting time-1 flow curves of the Hamiltonian flow of H , so we assume that

$$(5) \quad l(\alpha) C_1^H < \delta_0.$$

The converse of dissecting a time-1 flow curve is to start with an ‘‘approximated curve’’ $\vec{z} \in T^* \Lambda_r N$ and create a piece wise flow curve by defining the j^{th} piece

$$(6) \quad \gamma_j: [0, \alpha_j] \rightarrow T^*N$$

to be the Hamiltonian flow curve $\gamma_j(t) = \varphi_t^H(z_j) = \varphi_t^H(q_j, p_j)$ for all $j \in \mathbb{Z}_r$. These may not fit together to a continuous curve, and the construction of the approximations of the action is all about fixing this *without* creating additional critical points, and retaining that the set of connecting orbits is compact.

Let $P_q^{q'}: T_q N \rightarrow T_{q'} N$ denote the parallel transport induced by the Riemannian structure of tangent vectors along the unique geodesic when $\text{dist}(q, q') < 2\delta_0$.

Because of our assumption in (5), each of the flow pieces γ_j are shorter than δ_0 . So, we may for all $j \in \mathbb{Z}_r$ define:




Figure 1: Broken “geodesics”

$$\begin{aligned}
 (q_j^-, p_j^-) &= \gamma_{j-1}(\alpha_j) & \tilde{p}_j^- &= P_{q_j^-}^{q_j}(p^-) \in T_{q_j^-}^* N \\
 \epsilon_{q_j} &= \exp_{q_j^-}^{-1}(q_j) \in T_{q_j^-} N & \tilde{\epsilon}_{q_j} &= P_{q_j^-}^{q_j-1}(\epsilon_{q_j}) \in T_{q_{j-1}^-} N \\
 \epsilon_{p_j} &= p_j - \tilde{p}_j^- & P &= \max_{j \in \mathbb{Z}_r} \|p_j\|
 \end{aligned}$$

These, with the exception of the two last in the right column, are visualized in Figure 1. Even though all of these depend on α, H and \vec{z} , we suppress this from the notation.

We can now define the finite-dimensional approximations. Since we will only be interested in these when the assumptions of Lemma 2.3 are satisfied, we only define them in such cases. The assumption in (5) may be obtained by simply assuming that $\delta < \delta_0$ in the definition below. However, to get a buffer around the critical points we assume that $\delta < \delta_0/5$. So, all the flow curves γ_j have length less than $\delta_0/5$.

Definition 2.2 For any subdivision $\alpha \in \Delta^{r-1}$ and any Hamiltonian $H \in \mathcal{H}$ such that

$$l(\alpha)(C_1^H + C_2^H) \leq \delta,$$

with δ the constant from Lemma 2.3, we define

$$S_r: T^* \Lambda_r N \rightarrow \mathbb{R}$$

by

$$(7) \quad S_r(\vec{z}) = \sum_{j \in \mathbb{Z}_r} \left(\int_{\gamma_j} (\lambda - H dt) + p_j^- \epsilon_{q_j} \right).$$

The integration term in each summand of (7) is the action of the flow curve piece. The second term is the cotangent vector p_j^- evaluated on the tangent vector ϵ_{q_j} and can be viewed as a way to compensate for the fact that we really should integrate λ over a closed curve. In fact this term is λ integrated over the horizontal geodesic from (q_j^-, p_j^-) to (q_j, \tilde{p}_j^-) (which is the dotted line in Figure 1). The integration of λ over the unique line segment connecting (q_j, \tilde{p}_j^-) to (q_j, p_j) (the vertical arrows in Figure 1) is zero. So, if we so wished we could rewrite S_r as the integration of λ over a closed curve depending on \vec{z} plus $\sum_j (\alpha_j H(z_j))$.

By definition we see that: if for each $j \in \mathbb{Z}_r$ the curve γ_j ends where γ_{j+1} starts then \vec{z} is an α -dissection of a closed periodic orbit for the Hamiltonian flow of H starting at $z_0 \in T^*N$ and the value of S_r is the action of this curve. The following lemma and corollary are taken from [12].

Lemma 2.3 *There are constants $K > 0$ and $0 < \delta < \delta_0/5$ independent of r , $H \in \mathcal{H}$, and $\alpha \in \Delta^{r-1}$ such that the assumption*

$$(8) \quad l(\alpha)(C_1^H + C_2^H) \leq \delta$$

in Definition 2.2 implies the estimates

$$(9) \quad \|\nabla_{p_j} S_r - \tilde{\epsilon}_{q_{j+1}}\| \leq \frac{1}{4} \|\tilde{\epsilon}_{q_{j+1}}\| = \frac{1}{4} \|\epsilon_{q_{j+1}}\| \quad \text{and}$$

$$(10) \quad \|\nabla_{q_j} S_r + \epsilon_{p_j}\| \leq K \max(1, P)(\|\epsilon_{q_j}\| + \|\epsilon_{q_{j+1}}\|)$$

for all $j \in \mathbb{Z}_r$.

Here $\nabla_{q_j} S_r \oplus \nabla_{p_j} S_r = \nabla_{(q_j, p_j)} S_r = \nabla_{z_j} S_r$ is the gradient with respect to the j^{th} factor split into horizontal and vertical vectors.

Corollary 2.4 *The critical points of S_r are precisely the α -dissections of the 1-periodic Hamiltonian flow curves of H , and the associated critical value is the action of this orbit.*

Using Lemma 2.3 we can deduce a similar corollary for the approximations restricted to $T^\bullet \Omega_{q,r} N$, which we denote by

$$S_{q,r} = S_r|_{T^\bullet \Omega_{q,r} N}.$$

The proof of this corollary for the restriction is almost identical to the proof of the above corollary.

Corollary 2.5 *The critical points of $S_{q,r}$ are precisely the α -dissections of the time-1 Hamiltonian flow curves of H that start and end in the fiber $T_q^* N$, and the associated critical value is the action of this flow curve.*

Proof The gradient estimates in Lemma 2.3 are the same for the restriction, except that there is no $\nabla_{q_0} S_{q,r}$ since q_0 is constantly equal to q . So if $\vec{z} \in T^\bullet \Omega_{q,r} N$ is a critical point we may first use (9) to conclude that $\epsilon_{q_j} = 0$ for all $j \in \mathbb{Z}_r$. Then we use (10) to conclude that all $\epsilon_{p_j} = 0$ for $j \in \mathbb{Z}_r - \{0\}$. This goes both ways so this is a necessary and sufficient condition to be a critical point. So each γ_j ends where γ_{j+1} begins except for γ_{r-1} , which ends in the fiber over $q = q_0$. On the other hand any time-1 flow curve starting and ending over q may be α -dissected and produce a critical point. The critical value is the action by the definition of S_r and the fact that all the ϵ_{q_j} are 0. \square

3 Pseudogradients and compactness/completeness

In this section we define pseudogradients X_r and $Y_{q,r}$ for the approximations S_r and $S_{q,r}$ from Section 2. These will smoothly depend on all parameters (H, α, q) . We will then prove two important compactness/completeness results for these pseudogradients:

- (C1) Their flows exists for all times (future *and* past).
- (C2) The rate of change of S_r when flowing with them is bounded from below by a positive constant outside a compact set.

These statements are short versions of Lemma 3.6 and Lemma 3.7 below, and combinations of these are very useful. The fact that the pseudogradients $Y_{q,r}$ are the restrictions of a globally defined vector field Y_r is very important, and this Y_r is an example of what we will define as a parametrized pseudogradient (Definition 3.3).

Recall that a pseudogradient for a smooth function $f: M \rightarrow \mathbb{R}$ is a smooth vector field such that $X(f) \geq 0$ and $X = 0$ only at critical points for f .

To define the pseudogradients we fix once and for all a smooth function $\chi: [0, \delta_0] \rightarrow I$ such that

$$\chi(t) = \begin{cases} 1 & t < \delta_0/5, \\ 0 & t > \delta_0/4. \end{cases}$$

Using this we define the smooth bump functions $\chi_r: T^* \Lambda_r N \rightarrow I$ by

$$\chi_r(\vec{z}) = \prod_{j=1}^r \chi(\|\epsilon_{q_j}\|).$$

Note that since ϵ_{q_j} depends on H and α so will this function, but smoothly. Assuming S_r is defined, we define the following vector fields.

Definition 3.1 Let X_r be the smooth vector field on $T^*\Lambda_r N$ given by

$$(X_r)_{\vec{z}} = (\chi_r(\vec{z})\nabla_{\vec{q}'}S_r, \nabla_{\vec{p}'}S_r)$$

using the splitting associated to coordinates $\vec{z} = (\vec{q}', \vec{p}')$.

Similarly we define the smooth vector field Y_r on $T^*\Lambda_r N$ by

$$(Y_r)_{\vec{z}} = (0, \chi_r(\vec{z})\nabla_{\vec{q}'}S_r, \nabla_{\vec{p}'}S_r),$$

where $\vec{q}' = (q_1, \dots, q_{r-1})$ does not have the q_0 -factor.

The coordinates in the definition of Y_r are those given by the splitting of the tangent bundle associated to $\vec{z} = (q_0, \vec{q}', \vec{p}')$. In both cases we are as in Lemma 2.3 implicitly using the Riemannian structure to properly define the horizontal directions associated to \vec{q}' .

Remark 3.2 The vector field Y_r is vertical with respect to the submersion

$$\text{Ev}_0: T^*\Lambda_r N \rightarrow N.$$

Indeed, Y_r is simply the orthogonal projection of X_r to the vertical vectors (the kernel of $D\text{Ev}_0$). These vertical vectors are canonically identified with the tangent spaces of $T^*\Omega_{q,r} N$, and so we may define

$$Y_{q,r} = Y_r|_{\text{Ev}_0^{-1}(q)}$$

as a vector field on $T^*\Omega_{q,r} N$.

This leads us to the definition of parametrized pseudogradients. Let M be a smooth compact manifold possibly with boundary, corners, etc. A submersion of smooth manifolds $\pi: M' \rightarrow M$ is a smooth map with surjective differential such that $\partial M' = \pi^{-1}(\partial M)$, making each fiber a manifold of constant dimension *without* boundary.

Definition 3.3 A parametrized pseudogradient for $f: M' \rightarrow \mathbb{R}$ with respect to a submersion $\pi: M' \rightarrow M$ is a vector field on M' , which lies in the kernel of $D\pi$, and which in each fiber over $q \in M$ is a pseudogradient for f restricted to that fiber.

Remark 3.4 Since we assume that the base is compact, we only need to check (C1) and (C2) fiber-wise. This will come in handy later when we in fact will deal with compact smooth families of both X_r and Y_r .

We will call a critical point of $f|_{\pi^{-1}(q)}$ for any $q \in M$ a *fiber-critical point*. We will call the value of any fiber-critical point a *fiber-critical value*. Note that this is for all fibers simultaneously, ie we do not have an analogue of Sard's Theorem for these. Similarly we define *fiber-regular value*.

Lemma 3.5 *The vector fields X_r and $Y_{q,r}$ are pseudogradients for S_r and $S_{q,r}$, respectively. In particular Y_r is a parametrized pseudogradient for S_r with respect to $\text{Ev}_0: T^*\Lambda_r N \rightarrow N$.*

Proof The first gradient estimate (9) in Lemma 2.3 tells us that $\nabla_{\vec{p}} S_r$ is nonzero on the set where any of the ϵ_{q_j} are nonzero. So we conclude that on the set where we multiply the \vec{q} -component of ∇S_r with 0 in the definition of both X_r and Y_r , the \vec{p} -components of either gradient are nonzero. \square

Lemma 3.6 *The flows of $\pm X_r$ and $\pm Y_r$ are defined for all time.*

Proof These vector fields are constructed such that the positive or negative flows never reach the ‘‘boundary’’ where $\text{dist}(q_j, q_{j+1}) = \delta_0$. Indeed, this distance function is preserved by the flows when χ is zero, and this is the case if we assume $\text{dist}(q_j, q_{j+1}) \geq 9\delta_0/20$. Indeed, by the triangle inequality and this assumption we see that

$$\|\epsilon_{q_j}\| \geq \text{dist}(q_j, q_{j+1}) - \delta_0/5 \geq 9\delta_0/20 - \delta_0/5 = \delta_0/4.$$

The gradient estimates in (9) and the fact that $\nabla_{p_j} S_r = \nabla_{p_j} S_{q,r}$ are the p_j -components of both Y_r and X_r proves that

$$|\pm Y_r(\|p_j\|)| = |\pm X_r(\|p_j\|)| \leq \|\nabla_{p_j} S_r\| \|(\nabla\|p_j\|)\| \leq \|\nabla_{p_j} S_r\| \leq 5\delta_0/4$$

when $p_j \neq 0$. This implies that the rate of change of $\|p_j\|$ is less than $5\delta_0/4$ and the lemma follows. \square

Lemma 3.7 *There is an $\epsilon > 0$ such that $X_r(S_r) \geq Y_r(S_r) \geq \epsilon$ on the complement of the compact set*

$$B_r = \{\vec{z} \in T^*\Lambda_r N \mid \text{dist}(q_j, q_{j+1}) \leq \delta_0/2 \text{ and } P \leq 2\}.$$

Proof By definition $X_r(S_r) \geq Y_r(S_r)$, so we only need to consider Y_r . Also by definition of Y_r we see that $Y_r(S_r) \geq \|Y_r\|^2$. So getting a lower bound on $\|Y_r\|$ suffices.

We will thus assume that \vec{z} is a point in the set where $\|Y_r\| \leq \epsilon$ intersected with the complement of B_r , and arrive at a contradiction for ϵ small enough. However, for notational convenience we will often suppress the \vec{z} in the following.

Using the estimate (9) of $\nabla_{p_j} S_r$ in Lemma 2.3 we conclude that

$$(11) \quad \|\epsilon_{q_{j+1}}\| \leq \frac{4}{3} \|\nabla_{p_j} S_r\|, \quad j \in 0, \dots, r-1.$$

Since the \vec{p} -part of Y_r is the gradient $\nabla_{\vec{p}} S_r$ the above inequalities imply

$$(12) \quad \sum_{j=1}^r \|\epsilon_{q_j}\|^2 \leq \sum_{j=0}^{r-1} 2 \|\nabla_{p_j} S_r\|^2 \leq 2 \|Y_r\|^2 \leq 2\epsilon^2.$$

So with $2\epsilon^2 \leq \delta_0/5$ we see that

$$(13) \quad \|\epsilon_{q_j}\| \leq \delta_0/5 \quad \forall j = 1, \dots, r,$$

which implies

$$\text{dist}(q_j, q_{j+1}) \leq 2\delta_0/5 < \delta_0/2 \quad \forall j = 1, \dots, r.$$

Since we assumed that we are in the complement of B_r this implies the key observation that

$$P \geq 2.$$

Equation (13) also implies that $\chi_r = 1$ in Definition 3.1, which means that Y_r is the actual fiber-wise gradient and we thus get

$$(14) \quad \|Y_r\|^2 = \sum_{j \in \mathbb{Z}_r} \|\nabla_{p_j} S_r\|^2 + \sum_{j \in \mathbb{Z}_r - \{0\}} \|\nabla_{q_j} S_r\|^2.$$

The inequality $\|x - y\| \leq \|z\|$ implies that $\|y\|^2 \leq (\|x\| + \|z\|)^2$. Using (in order) this on the gradient estimates in (10), the inequality $(c + d + e)^2 \leq 3c^2 + 3d^2 + 3e^3$, (11) for all $j \in \mathbb{Z}_r$, and (14), we see that:

$$\begin{aligned} \sum_{j=1}^{r-1} \|\epsilon_{p_j}\|^2 &\leq \sum_{j=1}^{r-1} (\|\nabla_{q_j} S_r\| + KP(\|\epsilon_{q_j}\| + \|\epsilon_{q_{j+1}}\|))^2 \\ &\leq \sum_{j=1}^{r-1} (3\|\nabla_{q_j} S_r\|^2 + 3K^2P^2(\|\epsilon_{q_j}\|^2 + \|\epsilon_{q_{j+1}}\|^2)) \\ &\leq \sum_{j=1}^{r-1} (3\|\nabla_{q_j} S_r\|^2 + 6K^2P^2(\|\nabla_{p_j} S_r\|^2 + \|\nabla_{p_{j+1}} S_r\|^2)) \\ &\leq \max(3, 12K^2P^2) \|Y_r\|^2 \leq \max(3, 12K^2P^2) \epsilon^2 \end{aligned}$$

So for $\epsilon^2 < (4 \max(3, 12K^2)r)^{-1}$, we have

$$(15) \quad \sum_{j=1}^{r-1} \|\epsilon_{p_j}\|^2 < P^2/4r, \quad \text{which implies} \quad \sum_{j=1}^{r-1} \|\epsilon_{p_j}\| < P/2.$$

This is another key observation and Lemma 3.8 below thus implies that $\|p_j\| \geq P/2$ for all $j \in \mathbb{Z}_r$.

In particular we see that $\|p_j\| \geq 1$ for all $j \in \mathbb{Z}_r$, and on this set we see that S_r only depends on the constants μ and c (Definition 2.1) and the subdivision α . So in fact the gradient of S_r only depends on μ and α . Analyzing the Hamiltonians $H(q, p) = \mu\|p\|$ on $T^*N - DT^*N$ and the associated functions S_r (for any α) we have

$$(16) \quad S_r(\vec{q}, t\vec{p}) = tS_r(\vec{q}, \vec{p}),$$

as long as (\vec{q}, \vec{p}) has all $\|p_j\| \geq 1$ and $t \geq 1$. This implies that the norm of the gradient of S_r increases when \vec{p} is multiplied by t . Indeed, we have

$$(17) \quad \|\nabla_{\vec{q}} S_r(\vec{q}, t\vec{p})\| = t\|\nabla_{\vec{q}} S_r(\vec{q}, \vec{p})\|$$

and

$$(18) \quad \|\nabla_{t\vec{p}} S_r(\vec{q}, t\vec{p})\| = \|\nabla_{\vec{p}} S_r(\vec{q}, \vec{p})\|$$

on that same set.

The flow defined by these Hamiltonian has no time-1 flow curves starting and ending in the same fiber outside of DT^*N . Indeed, there is no geodesic starting and ending at the same point on N with length $\mu \in]0, \delta_0]$. This means that on the compact set given by the equations

- $1 \leq \min_j \|p_j\| \leq \max_j \|p_j\| \leq 2$ and
- $\text{dist}(q_j, q_{j+1}) \leq \delta_0/2$

the norm of ∇S_r has a lower bound. Using (18) and (17), this is a lower bound on the gradient of S_r on the noncompact set given by the equations

- $1 \leq P/2 \leq \min_j \|p_j\| \leq \max_j \|p_j\| = P$ and
- $\text{dist}(q_j, q_{j+1}) \leq \delta_0/2$.

This provides the contradiction if we finally make ϵ smaller than this lower bound. \square

Lemma 3.8 *If $P \geq 2$ and*

$$\sum_{j \in \mathbb{Z}_r - \{0\}} \|\epsilon_{p_j}\| \leq P/2$$

then $\|p_j\| \geq P/2$ for all $j \in \mathbb{Z}_r$.

Proof This relies on the fact that if either p_j or \tilde{p}_{j+1}^- lies outside DT^*N then so does the other and

$$(19) \quad \|\tilde{p}_{j+1}^-\| = \|p_j\|.$$

Indeed, these two are related by a composition of a parallel transport map (which preserves norms) and the Hamiltonian flow, which preserves the norm when outside DT^*N .

Since $0 \in \mathbb{Z}_r$ is special (omitted from the sum above) it is convenient in the following to identify \mathbb{Z}_r with $\{0, \dots, r-1\}$ and use the usual ordering on this set. We may pick i in this set such that $\|p_i\| = P \geq 2$. Now *assume for contradiction* that we have a $0 \leq j < i$ such that $\|p_j\| < P/2$; we may assume without loss of generality that j is the largest such $j < i$ with $\|p_j\| < P/2$. Then $\|p_{j+1}\| \geq P/2 \geq 1$ and using (19) we see that

$$\|\epsilon_{p_{j+1}}\| = \|p_{j+1} - \tilde{p}_{j+1}^-\| > \|p_{j+1}\| - P/2.$$

Using the same equation for all remaining p_k with $j < k \leq i$, which also lie outside DT^*N , we get

$$\|\epsilon_{p_k}\| = \|p_k - \tilde{p}_k^-\| \geq \|p_k\| - \|\tilde{p}_k^-\| = \|p_k\| - \|p_{k-1}\|.$$

Adding these for all $j < k \leq i$ we get that

$$\sum_{k=j+1}^i \|\epsilon_{p_k}\| > \|p_i\| - P/2 = P - P/2 \geq P/2,$$

which contradicts (15). The case with $i < j \leq r-1$ is similar, summing over k 's such that $i \leq k < j$. □

4 Parametrized Conley indices

In this section we define the Conley index (CI) associated to a smooth function f and a pseudogradient X satisfying the two compactness conditions (C1) and (C2) from Section 3. These are based compact Hausdorff spaces $I_a^b(f, X)$ depending on a

choice of interval $[a, b]$ and one additional contractible choice. These will come with a canonical inclusion

$$(20) \quad I_a^b(f, X) \subset f^{-1}([a, b])/f^{-1}(a),$$

which we will see is a homotopy equivalence. The only reason we do not use this “level-set quotient” as the actual definition is that later it will be very convenient to use the compact replacements that we describe in this section. We will also define a category of parametrized based spaces (called *ex-spaces*), and define the parametrized CI (PCI) as an object in such a category. A PCI will encode the fiber-wise CI, and Lemma 4.5 describes how these are related for different bases. In particular Corollary 4.6 relates the PCI of (S_r, Y_r) (whose fibers will be $I_a^b(S_{q,r}, Y_{q,r})$) with the *unparametrized* CI of (S_r, X_r) . For simplicity we will *always* assume that b is an upper bound on all fiber-critical values, and hence a fiber-regular value. However, since we will not be assuming such genericity for a , the fiber CI will be a slight generalization of the usual notion of CI’s. See eg Salamon [17] for a more conventional treatment of CI’s.

Since the unparametrized case is a special case of the parametrized case (the base is a single point) we formulate everything parametrized. We will need some basics about parametrized homotopy theory. As usual a map is a continuous function. In the following, W is any compact Hausdorff space.

Definition 4.1 The category \mathcal{C}_W^+ of compact *ex-spaces* over W has:

- Objects (or *ex-spaces*): $A = (A, p_A, s_A)$, with A a compact Hausdorff space, $p_A: A \rightarrow W$ a map, and $s_A: W \rightarrow A$ a map such that $p_A \circ s_A = \text{Id}_W$.
- Morphisms (or *ex-maps*): $f: A \rightarrow B$ a map such that $p_B \circ f = p_A$ and $s_B = f \circ s_A$.

We call p_A and s_A the projection of A and the section of A , respectively. For each $q \in W$ the *fiber* $A|_q = p_A^{-1}(q) \subset A$ is naturally a based space with base point $s_A(q)$. So intuitively this category is the category of based compact Hausdorff spaces continuously parametrized by W .

Definition 4.2 For any $A, B \in \mathcal{C}_W^+$ an *ex-homotopy* $G: A \times I \rightarrow B$ is a map such that $G|_{A \times \{t\}}$ is an *ex-map* for all $t \in I$.

In this definition we used the canonical identification of $A \times \{t\}$ with A . As for spaces, this provides the notions of *ex-homotopy equivalences*, *ex-homotopy type*, etc. In particular restricting any *ex-homotopy equivalence* $h: A \rightarrow B$ to the fibers over $x \in W$, we get a based homotopy equivalence of the fibers $h|_x: A|_x \rightarrow B|_x$.

By a pair of spaces (A, B) we mean an ordered pair of spaces (A, B) with an inclusion $B \hookrightarrow A$. For any pair of spaces (A, B) with a map $\pi: A \rightarrow W$ we define the *fiber-wise quotient* $A :_{\pi}^W B$ by the push-out diagram:

$$(21) \quad \begin{array}{ccc} B \hookrightarrow & A & \\ \downarrow \pi|_B & & \downarrow \\ W & \longrightarrow & A :_{\pi}^W B \end{array}$$

If this is compact Hausdorff it will define an ex-space in \mathcal{C}_W^+ . Indeed, the projection is the push-out of the map π and the identity on W , ie, the unique π' in the diagram:

$$(22) \quad \begin{array}{ccccc} B \hookrightarrow & A & & & \\ \downarrow \pi|_B & \downarrow & \searrow \pi|_A & & \\ W & \xrightarrow{s'} A :_{\pi}^W B & & & \\ & \downarrow \text{Id} & \searrow \pi' & & \\ & & & & W \end{array}$$

The section is the s' in the diagram. When unambiguous we omit the projection π and write $A :_W B$. The fiber of this construction at any point $q \in W$ is canonically identified with the based space given by the quotient $\pi^{-1}(q)/(\pi^{-1}(q) \cap B)$. To make the section more concrete one can add a W to both A and B . Ie, $A :_W B \cong (A \sqcup W) :_W (B \sqcup W)$. However, this is not actually necessary by the usual definition of push-out.

The fact that we will mostly use these notions for compact Hausdorff spaces makes many things easier, however, for a general treatment we refer to May and Sigurdsson [14]. We will at some points use noncompact ex-spaces. However, the properties we will use for these will be very basic.

Let $f: M' \rightarrow \mathbb{R}$ be a smooth function and let $Y: M' \rightarrow TM'$ be a parametrized pseudogradient (Definition 3.3) with respect to a submersion

$$\pi: M' \rightarrow M.$$

We assume that Y satisfies (C1) and (C2) from Section 3. So, there is an $\epsilon > 0$ such that $Y(f) > \epsilon$ outside some compact set $K \subset M'$. Let ψ_t denote the flow of $-Y$, which exists for all time. We may define

$$(23) \quad C_\tau = \psi_\tau(f^{-1}([a, b])) \cap f^{-1}([a, b])$$

and

$$(24) \quad \underline{C}_\tau = \psi_\tau(f^{-1}([a, b])) \cap f^{-1}(a) \subset C_\tau.$$

These sets are illustrated in Figure 2.




Figure 2: C_τ and \underline{C}_τ

Lemma 4.3 *The subspace $C_\tau \subset f^{-1}([a, b])$ is compact for $\tau \geq \epsilon^{-1}(b - a)$.*

Proof Let $t_0 = \epsilon^{-1}(b - a)$. Then define

$$K' = \bigcup_{t \in [0, t_0]} \psi_t(K).$$

This is compact, and any point $x \in f^{-1}([a, b])$ will either flow into and thus end in this set after time t_0 , or satisfy

$$\frac{\partial}{\partial t}(f \circ \gamma)(t) = -Y_{\gamma(t)}(f) < -\epsilon,$$

in which case we get $f(\gamma(t_0)) - f(\gamma(0)) < -\epsilon(b - a)\epsilon^{-1}$. This implies that $f(\gamma(t_0)) < a$, ultimately proving that

$$C_\tau \subset K'$$

for all $\tau \geq t_0$. Since C_τ is closed and K' is compact we have proved the lemma. \square

This leads us to define the parametrized Conley index (PCI) as

$$(25) \quad I_a^b(f, Y, \tau)_M = I_a^b(f, Y)_M = C_\tau :_M \underline{C}_\tau,$$

which for $\tau \geq t_0$ is compact and thus lies in \mathcal{C}_M^+ . The choice of τ is a contractible choice, and we will assume that C_τ and thus the PCI is compact if nothing else is mentioned. We will comment on the choice of τ when appropriate. The following lemma justifies this suppression.

Lemma 4.4 *The subspace C_τ is a subspace in $C_{\tau'}$ when $\tau \geq \tau' \geq 0$, and the inclusion induces an ex-homotopy equivalence*

$$I_a^b(f, Y, \tau)_M \rightarrow I_a^b(f, Y, \tau')_M$$

even when the indices are noncompact.

Proof Since Y is a pseudogradient it is clear that

$$\psi_{\tau'}(f^{-1}([\infty, b])) \subset \psi_{\tau}(f^{-1}([\infty, b])),$$

which is preserved when intersecting with $f^{-1}([a, b])$. So $C_{\tau} \subset C_{\tau'}$.

Furthermore, we may define what we will call a *stop-flow* $\underline{\psi}_t$ on $f^{-1}([a, b])$ by the formula

$$\underline{\psi}_t(x) = \psi_{s(x,t)}(x),$$

where $s(x, t) = t$ if $f(\psi_t(x)) \geq a$ and otherwise we put

$$s(x, t) = \max\{s \in [0, t] \mid f(\psi_s(x)) \geq a\}.$$

This is the flow ψ_s restricted to $f^{-1}([a, b])$ and stopped (abruptly) when reaching the set $f^{-1}(a)$. This is continuous, and since it preserves the set $f^{-1}(a)$, and the submersion to M it defines ex-maps

$$\underline{\psi}_{\tau-\tau'}: I_a^b(f, Y, \tau)_M \rightarrow I_a^b(f, Y, \tau')_M,$$

which is an ex-homotopy inverse to the inclusions. Indeed, in both cases the ex-homotopy to the identity is simply given by these maps $\underline{\psi}_t$ for $t \in [0, \tau' - \tau]$. \square

Let $\pi': M \rightarrow M''$ be another submersion and let X be another parametrized pseudogradient for f , this time with respect to $\pi' \circ \pi: M' \rightarrow M''$ satisfying (C1) and (C2). Then we have two parametrized Conley indices $I_a^b(f, Y)_M$ and $I_a^b(f, X)_{M''}$.

Lemma 4.5 *There is a canonical (up to contractible choices) ex-homotopy equivalence*

$$I_a^b(f, X)_{M''} \simeq_{M''} I_a^b(f, Y)_M :_{M''} M.$$

Corollary 4.6 *The Conley index $I_a^b(S_r, X_r)$ is homotopy equivalent to*

$$(I_a^b(S_r, Y_r)_N)/N,$$

where $(\cdot)/N$ is the functor from ex-spaces over N to based spaces given by collapsing the section.

Proof of Lemma 4.5 This follows from Lemma 4.4 and the fact that by definition the level-set quotients satisfies this, ie

$$(f^{-1}([a, b]) :_M f^{-1}(a)) :_{M''} M \cong f^{-1}([a, b]) :_{M''} f^{-1}(a).$$

This ‘‘canonical’’ ex-homotopy equivalences are thus defined as the maps induced (when taking quotients over M'') by the composition of:

- the inclusion $C_\tau \rightarrow f^{-1}([a, b])$ and
- the flow of $-X$ taking the image into a set C'_τ , defining $I_a^b(f, X)_{M''}$.

Here C_τ is defined using $-Y$ and C'_τ , is defined using $-X$. □

The fiber-wise quotient (used in Lemma 4.5) of the section for any map $h: W \rightarrow W'$ is called the push-forward of ex-spaces and is a functor denoted:

$$h_!: C_W^+ \rightarrow C_{W'}^+$$

The relation to sheaf-shriek push-forward is subtle.

5 The suspension maps

In this section we will, for a given Hamiltonian H , relate the Conley indices for S_r and $S_{q,r}$ defined in Section 4 for different choices of r . These turn out to be related by twisted suspensions, which essentially is why we can use these to define spectra. More specifically assume that H is a Hamiltonian and $\alpha \in \Delta^{r-1}$ is a subdivision such that S_r is defined. Then the top face map

$$(26) \quad d_r: \Delta^{r-1} \subset \Delta^r$$

given by appending a 0 defines S_{r+1} (and Y_{r+1} and X_{r+1}) using the same Hamiltonian H . Then for each $q \in N$ we construct a map

$$\sigma_{q,r}: \Sigma^d I_a^b(S_r, Y_r)_q \rightarrow I_a^b(S_{r+1}, Y_{r+1})_q.$$

Here $\Sigma^d(\cdot)$ denotes the d -fold reduced suspension of based topological spaces. We construct these such that they are defined continuously over N , and in the generic case they will induce homotopy equivalences of unparametrized indices (when a is regular for S_r). Here continuous in N means that we in fact define a map of ex-spaces

$$(27) \quad \sigma_r: I_a^b(S_r, Y_r)_{TN}^{TN} \rightarrow I_a^b(S_{r+1}, Y_{r+1})_N,$$

where $I_a^b(S_r, Y_r)_{TN}^{TN}$ denotes the twisted fiber-wise reduced suspension of the index $I_a^b(S_r, Y_r)_N$ using the one-point compactification of the fibers of TN . We start by defining this properly. We will refer to this procedure as TN -suspending.

The construction of the twisted fiber-wise reduced suspension does not depend on the fact that TN is the tangent bundle of N . So, assume that $\pi: \beta \rightarrow W$ is any metric vector bundle of dimension l over a compact Hausdorff space W . We have the associated disc bundle

$$D\beta = \{v \in \beta \mid \|v\| \leq 1\}$$

and sphere bundle

$$S\beta = \{v \in \beta \mid \|v\| = 1\}.$$

We then define

$$S^\beta = (D\beta) :_W (S\beta) \in \mathcal{C}_W^+,$$

where $:_W$ is the fiber-wise quotient defined in (21). This may be identified with the sphere bundle of $\beta \times \mathbb{R} \rightarrow W$, with its canonical section. Or, it may be identified with the fiber-wise one-point compactification of β .

For any two ex-spaces A and B over W we also define the fiber-wise product over W as the usual pull-back

$$A \times_W B = \{(x, y) \in A \times B \mid p_A(x) = p_B(y)\}.$$

This has the obvious section and projection and it is the categorical product. We also define the fiber-wise wedge product as

$$A \vee_W B = A \cup_W B$$

using the sections s_A and s_B . This is the categorical coproduct. We may then define the fiber-wise smash product as

$$A \wedge_W B = (A \times_W B) :_W (A \vee_W B).$$

These are symmetric monoidal products. The smash product \wedge_W distributes over the wedge \vee_W and defines a symmetric bimonoidal structure on the category of ex-spaces over W . Restricting to a fiber $q \in W$ defines a symmetric bimonoidal functor to based spaces (with the usual \wedge and \vee as products, so we are really just saying that these structures on ex-spaces behave similarly).

Any ex-space A over W may now be fiber-wise reduced suspended by the sphere bundle S^β . That is we define the β -suspension

$$(28) \quad A^\beta = A \wedge_W S^\beta.$$

A little more concretely, one may identify A^β with

$$D(p_A^*\beta) :_W (D(\beta) \cup S(p_A^*\beta)).$$

Here $D(\beta)$ is identified with the restriction of $D(p_A^*\beta)$ to the image of the section $s_A: W \rightarrow A$.

Remark 5.1 The symmetric bimonoidal structure on ex-spaces over W has:

- Zero: $W_W = (W, \text{Id}_W, \text{Id}_W)$ (to be read W over W).

- Unit: S_W^0 .

Here

$$S_W^n = (S^n \times W, \pi_W, * \times \text{Id}_W),$$

where $\pi_W: S^n \times W \rightarrow W$ is the projection and $* \times \text{Id}_W: W \rightarrow S^n \times W$ is the constant map to the base point of S^n times the identity.

The reason we denote the zero by W_W and not simply W is because this would be very ambiguous with the usual notation for Thom spaces. Ie, for $\beta \rightarrow W$ we would define its Thom space to be

$$W^\beta = D\beta/S\beta \cong ((S_W^0)^\beta)/W = S^\beta/W,$$

which is not $W_W^\beta (\cong W_W)$. So, the subscript emphasizes that we are in the category over W such that we know that $(\cdot)^\beta$ means the β -suspension.

So, $I_a^b(S_r, Y_r)_{TN}^{TN}$ is defined as above using the metric bundle $TN \rightarrow N$ and a choice of the parametrized Conley index from Section 4. So, we may identify the restriction of this to the fiber over $q \in N$ with the usual reduced suspension of the Conley index of that fiber:

$$\Sigma^d I_a^b(S_{q,r}, Y_{q,r}) = I_a^b(S_{q,r}, Y_{q,r}) \wedge S^d.$$

However, this is not canonical since it depends on a choice of isomorphism $T_q N \cong \mathbb{R}^d$.

Assume now that $(H, \alpha) \in \mathcal{H} \times \Delta^{r-1}$ is given and defines S_r . Then, we may define the stabilized subdivision $\alpha' \in \Delta^r$ by $\alpha'_j = \alpha_j$ for $j = 0, \dots, r-1$ and $\alpha'_r = 0$. Indeed, this is the usual top face map d_r from (26) evaluated on α . Using (H, α') we then define the function S_{r+1} and parametrized pseudogradient Y_{r+1} .

We will now define σ_r in a few steps. Let $\underline{C}_\tau \subset C_\tau \subset S_r^{-1}([a, b])$ be as defined in (23) and (24) using the function S_r and parametrized pseudogradient Y_r . Let τ be such that C_τ and hence

$$I_a^b(S_r, Y_r)_N = C_\tau :_N (C_\tau \cap S_r^{-1}(a))$$

is compact. Let $E \rightarrow C_\tau$ be the vector bundle given by $\text{Ev}_0^* TN$. Then, we define the map

$$(29) \quad f = f_r: E \rightarrow S_{r+1}^{-1}([a, b])$$

given by

$$(30) \quad f(z_0, \dots, z_{r-1}, v) = (z_0, \dots, z_{r-1}, (q_0, v)).$$

Here $\vec{z} = (z_0, \dots, z_{r-1}) \in T^* \Lambda_r N$ and $v \in T_{q_0} N$, and the new “coordinate” in $f(\vec{z})$ is simply given by (q_0, v) . So if we denote $\vec{z}' = f(\vec{z})$ we may write $z'_j = z_j$ for $j = 0, \dots, r-1$ and $z'_r = (q_0, v)$. This notation is a little abusive since we originally assumed that the indices j are in \mathbb{Z}_r and \mathbb{Z}_{r+1} , and now we are identifying these with $\{0, \dots, r-1\}$ and $\{0, \dots, r\}$, respectively.

Notice that this in fact does land in $S_{r+1}^{-1}([a, b])$ because the diagram

$$(31) \quad \begin{array}{ccc} E & \xrightarrow{f} & T^* \Lambda_{r+1} N \\ \downarrow & & \downarrow S_{r+1} \\ C_\tau & \xrightarrow{S_r} & \mathbb{R} \end{array}$$

commutes. Indeed, the two new terms in the sum are both zero independent of v . This is because $\epsilon_{q_r} = 0$ and the flow curve γ_r is the constant curve parametrized by the point $0 \in \mathbb{R}$, so the integral vanishes. In particular we see that S_{r+1} is constant on the image of f restricted to any fibers of the vector bundle.

The image of f is not compact, and it does not induce a map of the quotients that we need it to. However, composing with the flow $\psi'_{\tau'}$ of $-Y_{r+1}$ and intersecting with $S_{r+1}^{-1}([a, b])$ will by Lemma 4.3 be compact if τ' is large enough. That is, if we define the compact set

$$C'_{\tau'} = \psi'_{\tau'}(S_{r+1}^{-1}([a, b])) \cap S_{r+1}^{-1}([a, b])$$

as in (23), but for S_{r+1} . Then we get

$$(\psi'_{\tau'} \circ f)(E) \subset C'_{\tau'} \cup S_{r+1}^{-1}([-\infty, a]).$$

We may then (using compactness) find a $R \gg 0$ such that $S_R E$ (the sphere bundle of E with radius R) is sent by this map to $S_{r+1}^{-1}([-\infty, a])$. This means that it induces a map of quotients

$$D_R E / (S_R E \cup (D_R E)|_{C_\tau}) \rightarrow C'_{\tau'} \cup S_{r+1}^{-1}([a, b]) / S_{r+1}^{-1}([-\infty, a]).$$

However, all maps are over N , so it actually induces a map of the quotients over N . The source of this ex-map is identified as

$$D_R E :_N (S_R E \cup (D_R E)|_{C_\tau}) \cong I_a^b(S_r, Y_r)_{N}^{TN}$$

and the target is

$$C'_{\tau'} \cup S_{r+1}^{-1}([a, b]) :_N S_{r+1}^{-1}([-\infty, a]) \cong C'_{\tau'} :_N C'_{\tau'} = I_a^b(S_{r+1}, Y_{r+1})_N.$$

So we have defined the wanted map

$$\sigma_r : I_a^b(S_r, Y_r)_{N}^{TN} \rightarrow I_a^b(S_{r+1}, Y_{r+1})_N,$$

depending on the contractible choices of τ and τ' .

Since this construction generalizes our construction in [12] to a fiber-wise construction we immediately get the following corollary.

Corollary 5.2 *If a is a regular value for S_r (and thus also S_{r+1}) then the map*

$$\sigma_r/N: (I_a^b(S_r, Y_r)^{TN})/N \rightarrow I_a^b(S_{r+1}, X_{r+1})$$

is a homotopy equivalence.

Note that by assumption b is always assumed regular. One could try to prove more generally that this corollary is true for any a , and even prove a fiber-wise version. However, we will not need this, and it is thus convenient to avoid the technicalities of this by simply making sure later that when this corollary is needed a is, indeed, regular for $S_r: T^*\Lambda_r N \rightarrow \mathbb{R}$.

Proof Corollary 4.6 tells us that using the functor $(\cdot)/N$ we recover the unparametrized indices $I_a^b(S_r, X_r)$ and $I_a^b(S_{r+1}, X_{r+1})$. These were proved in [12] to be related by a relative Thom construction with TN relative to \underline{C}_τ , and inspecting the above-defined map we see that it induces this homotopy equivalence. □

6 The generating function spectra

In this section we assume that we are given:

- (1) A Hamiltonian $H: T^*N \rightarrow \mathbb{R}$ in \mathcal{H} .
- (2) A subdivision $\alpha \in \Delta^{r-1}$ such that S_r from Section 2 is defined using H and α .
- (3) A compact interval $[a, b] \subset \mathbb{R}$ such that
 - a is not the action of *closed* 1-periodic orbit associated to H , and
 - b is a strict upper bound on the action of the time-1 flow curves starting and ending in the same fiber of $T^*N \rightarrow N$.

After some additional choices (which we suppress from the notation) we will associate to this a fibered spectrum

$$\mathcal{GF}_a^b(H, \alpha)$$

over N . Here fibered³ simply means that for each $q \in N$ we have a spectrum, and these glue together continuously in a sense that we will make precise. This fibered

³Not to be confused with *fibrant*.

spectrum captures the Floer homology of the action integral restricted to those curves with action above a , encoding what takes place over each fiber. We warn that as defined in this section these are *not* parametrized spectra in the sense considered by May and Sigurdsson in [14], unless β in the following is trivial. Appendix B, however, describes a natural functor to a more standard notion of parametrized spectra.

We will then define the notion of homology of such fibered spectra, and prove that for $\mathcal{GF}_a^b(H, \alpha)$ this does not depend on anything but H and a , and this will recover the notion of generating function homology. In fact homology forgets the fibered structure, and only depends on the unparametrized index. If one wants information about fibers one should restrict to a fiber and *then* take homology.

Motivated by the construction in Section 5 we define the following very naïve version of a category of spectra over any compact Hausdorff space W . Let $\beta \rightarrow W$ be any metric vector bundle. We will use that $(\cdot)^\beta: C_W^+ \rightarrow C_W^+$ (defined in Section 5) is a functor. Indeed, we may extend the construction to ex-maps $f: A \rightarrow B$ by smashing with the identity on S^β , ie, we get

$$f^\beta = f \wedge_W \text{Id}_{S^\beta}: A^\beta \rightarrow B^\beta.$$

We note that there is an obvious natural isomorphism of functors between $((\cdot)^{\beta_1})^{\beta_2}$ and $(\cdot)^{\beta_1 \oplus \beta_2}$. We will use this implicitly on iterations of $(\cdot)^\beta$ and use the notation $(\cdot)^{k\beta}$. For the sake of notation we will (when there is no ambiguity) suppress these exponents on ex-maps, so we will often simply write f instead of $f^{k\beta}$. Sometimes we will write f^\oplus to emphasize that some (possible) exponent has been suppressed.

Definition 6.1 A β -spectrum \mathcal{A} is a sequence of pairs $\mathcal{A} = (A_r, \sigma_r)_{r \in \mathbb{N}_0}$ with each A_r an ex-space over W and each σ_r an ex-map

$$\sigma_r = \sigma_r^A: (A_r)^\beta \rightarrow A_{r+1}$$

over W .

We call the σ_r *structure maps*. We define $\sigma_{r,r'} = \sigma_{r,r'}^A$ as the composition

$$\sigma_{r'-1} \circ \sigma_{r'-2} \circ \dots \circ \sigma_r: A_r^{(r'-r)\beta} \rightarrow A_{r'},$$

so that $\sigma_{r,r+1} = \sigma_r$ and $\sigma_{r,r} = \text{Id}$. Here we have suppressed some β -exponents. Eg, we have

$$\sigma_{r,r+2} = \sigma_{r+1} \circ (\sigma_r)^\beta$$

as the map

$$A_r^{2\beta} \xrightarrow{(\sigma_r)^\beta} A_{r+1}^\beta \xrightarrow{\sigma_{r+1}} A_{r+2}.$$

We will later define morphisms of β -spectra and we denote the category of these by Sp_W^β .

If we take a β -spectrum and restrict it level-wise to the fiber over a single point $q \in W$ then we recover a naïve spectrum (this identification depends on a trivialization of β_q). That is, we get a sequence of based spaces A_r and a sequence of maps $\sigma_r: \Sigma^l(A_r) \rightarrow A_{r+1}$, where $\Sigma^l(\cdot)$ denotes reduced suspension. The “restriction” of the suppressed natural isomorphism of functors from $((\cdot)^\beta)^\beta$ to $(\cdot)^{2\beta}$ is then simply the natural isomorphisms of functors from $((\cdot) \wedge S^l) \wedge S^l$ to $(\cdot) \wedge S^{2l}$.

Now let H be any Hamiltonian as in Section 2 and $\alpha \in \Delta^{r_0-1}$ with $\alpha_{r_0} \neq 0$. We will think of this as a choice of subdivision for all $r \geq r_0$ by using the top face inclusion $\Delta^{r-1} \subset \Delta^r$ defined by appending any subdivision by a zero. This was already introduced when we defined the ex-map σ_r in Section 5. Using these maps we define the infinite simplex Δ^∞ by taking the topological union

$$(32) \quad \Delta^\infty = \bigcup_{r \in \mathbb{N}} \Delta^{r-1}$$

with the weak topology. For any point $\alpha \in \Delta^\infty$ we define $r_\alpha \in \mathbb{N}$ by

$$(33) \quad r_\alpha = \inf\{r \in \mathbb{N} \mid \alpha \in \Delta^{r-1}\}.$$

Now let $\alpha \in \Delta^\infty$ be such that S_{r_α} is defined using H and α . This implies that S_r is defined for all $r \geq r_\alpha$, which means that the Conley indices $I_a^b(s_r, Y_r)_N$ are defined. Also all the suspension ex-maps σ_r in Section 5 between these indices for different r are defined.

We thus define the TN -spectrum

$$(34) \quad \mathcal{GF}_a^b(H, \alpha)$$

by $\mathcal{GF}_a^b(H, \alpha)_r = N_N$ (from Remark 5.1) for $r < r_\alpha$. For each $r \geq r_\alpha$ we define

$$\mathcal{GF}_a^b(H, \alpha)_r = I_a^b(S_r, Y_r)_N.$$

We define σ_r by the canonical identification $N_N^{TN} \cong N_N$ when $r < r_\alpha - 1$, and we define $\sigma_{r_\alpha-1}$ by the section (or the fact that N_N is initial in C_N^+). For $r > r_\alpha$ we use the ex-maps

$$\sigma_{r-1}: \mathcal{GF}_a^b(H, \alpha)_{r-1}^{TN} \rightarrow \mathcal{GF}_a^b(H, \alpha)_r$$

constructed in Section 5 as the structure maps.

Remark 6.2 It may seem odd that this fibered spectrum is trivial on the first levels. However, this is a general fact about spectra. It does not really matter up to equivalence

if one replaces the first finite number of spaces by the trivial space. In fact it is *intuitively* convenient to think of spectra as a sort of colimit as $r \rightarrow \infty$ (which is not defined in the category of spaces). This is reflected in the following definition of homology.

Let \mathbb{F} be either \mathbb{Q} or \mathbb{F}_p for some prime p . Assume that β is \mathbb{F} orientable and that we have chosen an orientation. Using the functor in Appendix B we could define this for any general homology theory, but to make things as transparent as possible we won't. Recall that N is viewed as a subspace of A_r using its section. So in the following $H_*(A_r, N; \mathbb{F})$ denotes the usual relative singular homology groups.

Definition 6.3 The homology of a β -spectrum \mathcal{A} with coefficients in \mathbb{F} is defined to be the colimit

$$H_*(\mathcal{A}; \mathbb{F}) = \operatorname{colim}_{r \rightarrow \infty} H_{*+l_r}(A_r, N; \mathbb{F})$$

using the maps $H_*(A_r, N; \mathbb{F}) \rightarrow H_{*+l}(A_{r+1}, N; \mathbb{F})$ given by the composition of the maps:

- The Thom-map $H_*(A_r, N; \mathbb{F}) \rightarrow H_{*+l}(A_r^\beta, N; \mathbb{F})$.
- The map induced by the structure map $\sigma_r: A_r^\beta \rightarrow A_{r+1}$.

The *Thom-map* will, when the sections are cofibrant, be a Thom isomorphism. This is, indeed, the case for $\mathcal{GF}_a^b(H, \alpha)$ since a is regular (see the proof below). However, we want to talk about the homology of fibers as well. So we do not restrict the definition to this case, and in general we define the Thom-map by the following steps.

- First we use the Thom-isomorphism

$$H_*(A_r, N; \mathbb{F}) \cong H_{*+l}(Dp^*\beta, Sp^*\beta \cup D\beta; \mathbb{F})$$

of pairs, which depends on the orientation on β . In the proof of Proposition 10.5 a different perspective on this map is used in the case where β is trivial.

- Then, we compose with the map induced by the fiber-wise collapse map

$$(Dp^*\beta, Sp^*\beta \cup D\beta) \rightarrow (A_r^\beta, N)$$

defining A_r^β in (28).

Although we have not defined morphisms of TN -spectra yet, we note that homology is a functor.

Lemma 6.4 The homology $H_*(\mathcal{GF}_a^b(H, \alpha); \mathbb{F})$ is isomorphic to

$$H_{*+l_{r_\alpha}}(I_a^b(S_{r_\alpha}, Y_{r_\alpha})_N, N; \mathbb{F}) \cong \widetilde{H}_{*+l_{r_\alpha}}(I_a^b(S_{r_\alpha}, X_{r_\alpha}); \mathbb{F}).$$

Note that the notion of generating function homology (associated to the action on *closed* loops) discussed by, eg, Viterbo in [20] is recovered in this lemma. However, we will later only use the explicit computations in [12], and not use any relation to actual (twisted) Floer homology.

There is a notion of an ex-space being well-sectioned, which is more than the section being a standard cofibration of spaces. This notion does *not* appear in the following proof.

Proof If we can argue that all the sections are cofibrations then Corollary 4.6 tells us that $H_{*+lr_\alpha}(I_a^b(S_r, Y_r)_N, N; \mathbb{F}) \cong \widetilde{H}_{*+lr_\alpha}(I_a^b(S_r, X_r); \mathbb{F})$. Combining this with Corollary 5.2 we see that in this case the colimit in Definition 6.3 defining the homology is the colimit of the diagram

$$0 \cdots 0 \rightarrow H_*(I_a^b(S_{r_\alpha}, Y_{r_\alpha})_N, N; \mathbb{F}) \xrightarrow{\cong} H_*(I_a^b(S_{r_{\alpha+1}}, Y_{r_{\alpha+1}})_N, N; \mathbb{F}) \xrightarrow{\cong} \cdots$$

proving the corollary.

The sections *are* cofibrations. Indeed, since a is a regular value for S_r the inclusion $\underline{C}_\tau \subset C_\tau$ (defined in (23) and (24)) is a cofibration, and for any cofibrant pair (A, B) with a map to N the section of $A :_N B$ is cofibrant. □

Corollary 6.5 *The homology of $\mathcal{GF}_a^b(H, \alpha)$ does not depend on α .*

Proof In the proof above we related the homologies of the parametrized indices to the homologies of the unparametrized indices, and these satisfy homotopy invariance. So the corollary follows since the fact that a and b are regular values for S_r does not depend on the subdivision α . □

7 The TN –spectrum \mathcal{FL} associated to L exact in T^*N

Let $j: L \rightarrow T^*N$ be as in the introduction. In this section we construct the associated TN –spectrum \mathcal{FL} as a “colimit” of certain generating function spectra. We put colimit in quotes since we have not yet discussed morphisms of TN –spectra. We will instead construct \mathcal{FL} explicitly. The homology of this TN –spectrum will be the homology of the free loop space of L with possibly twisted coefficients. It will in fact represent the symplectic homology of DT^*L , but with possibly different twisted coefficients; we discussed these local coefficients in Remark 1.3 and Appendix D. The construction of \mathcal{FL} will depend on a lot of choices, which we suppress from the notation. Indeed, we will construct a smooth family

$$(H^s, \alpha^s) \quad \text{for } s \in [0, \infty[,$$

where this family consists of a Hamiltonian and a subdivision for each s . Then for each s we may pick r large enough to define associated finite-dimensional approximations S_r^s , and define associated (un)parametrized pseudogradients X_r^s and Y_r^s , all as in Section 2 and Section 3. This will define a family of generating function spectra, and we will explicitly describe the “colimit” of these as s tends to infinity. Since we will be tweaking these choices in later sections, and work with more than one set of choices simultaneously, we introduce the terminology that a set of choices defines an *instance* of \mathcal{FL} . However, we note that \mathcal{FL} is in the homotopy category of TN -spectra defined up to unique isomorphism, and we will discuss this in Appendix C.

By choosing an appropriate extension (and rescaling) of j and an appropriate Riemannian structure on L we may assume that we have a symplectic embedding of the closed unit disc bundle DT^*L into $D_{1/2}T^*N$. By abuse of notation we will also denote this embedding j .

Let λ_N and λ_L denote the canonical Liouville forms on T^*N and T^*L , respectively. Similarly we will let $\|\cdot\|_N: T^*N \rightarrow \mathbb{R}$ and $\|\cdot\|_L: T^*L \rightarrow \mathbb{R}$ denote the norm of the cotangent vectors in N and L , respectively.

Because j is exact symplectic we may choose a function $F: DT^*L \rightarrow \mathbb{R}$ such that $dF = j^*\lambda_N - \lambda_L$, and this is unique up to a constant because L is connected. Since DT^*L is compact F and its derivatives are bounded. Because of this we will assume that F takes the value 0, but is nonnegative, and hence $\|F\| = \|F\|_\infty$ is a bound on $|F(z_1) - F(z_2)|$.

If $\gamma: I \rightarrow DT^*L$ is any smooth path we can now consider its action in T^*L (using the 1-form λ_L) or the action of $j \circ \gamma$ in T^*N (using λ_N). We denote these two different action functionals by A_L and A_N , respectively. They satisfy the formula

$$(35) \quad A_N(j \circ \gamma) = A_L(\gamma) + (F(\gamma(1)) - F(\gamma(0))).$$

The extra terms vanish on closed loops. So on closed loops the two action integrals agree.

To create \mathcal{FL} we have to construct some families of Hamiltonians associated to j . We will describe some smooth functions and some of their properties in the following, and we postpone the actual construction of these to Appendix A. Indeed, we do so because we need them to satisfy a lot of additional properties later, which we do not want to mention here.

First we fix a smooth concave function $f: \mathbb{R}_+ \rightarrow \mathbb{R}$ (see Figure 3) such that:

- $f(t) \rightarrow -\infty$ when $t \rightarrow 0$.

- $f(t) = 0$ when $t \in [1, \infty[$.
- $f''(t) < 0$ when $t \in]0, 1[$, ie f is strictly concave on $]0, 1[$.

For each $s > 0$ we define t_s to be the unique point $t_s \in]0, 1[$ such that the tangent of f at t_s intersects the 2nd axis at $-s$ (see Figure 3). By assumptions on f this t_s will




Figure 3: The function f




Figure 4: The functions f_s and $f + s$

go monotonely to 0 as s goes to infinity.

For any smooth function $g: \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}$ we define the *action set* of g as the set of those $s \in \mathbb{R}$ such that there is a tangent to g which intersects the 2nd axis at $-s$ and has slope the length of a *closed* geodesic on L . We also define this for f even though f is not defined at 0. The action set of f consists of nonnegative real numbers. If we define the Hamiltonian $H(q, p) = f(\|p\|_L)$ away from the zero-section on T^*L the action of 1-periodic orbits would coincide with this action set of f , hence the name. Indeed, this is a standard calculation (see eg [12, Section 3]). Of course we cannot use this Hamiltonian as it is not extendable to $L \subset T^*L$. It will thus be convenient to translate the function f and cap it of depending on a parameter s .

So let $f_s: \mathbb{R} \rightarrow \mathbb{R}$ for $s \in [0, \infty[$ be the smooth family of smooth increasing functions constructed in Appendix A using f . These satisfy:

- (f2') For (t, s) in a neighborhood of $(0, s) \subset \mathbb{R}_{\geq 0}^2$ we have $f_s(t) = ct^2$ for some $c = c(s) \geq 0$.
- (f3) For $t \geq 1$, we have $f_s(t) = s$.
- (f4) For $s > 0$, the restriction of $f_s(t)$ to $]0, 1[$ is strictly increasing.

These force $f_0 = 0$. In this section we particularly need for $s > 5$:

- (f5) The tangent to f_s at any $t \in [0, t_s]$ intersects the 2nd axis in $] - 1, 0]$.
- (f6) We have $f_s(t) = f(t) + s$ when $t \geq t_s$, making the tangents of f_s at these points intersect the 2nd axis in $[0, s]$.

Here 5 is a rather arbitrary choice since we in this section really only care about the limit $s \rightarrow \infty$. The function f_s (for $s \geq 5$) is illustrated in Figure 3. Let $H_\infty: T^*N \rightarrow \mathbb{R}$ be a function in $\mathcal{H}_\infty \subset \mathcal{H}$ from Appendix A. In particular H_∞ satisfies:

- H_∞ is zero on $D_{2/3}T^*N$.
- Any time-1 Hamiltonian flow curve for H_∞ starting and ending in the same fiber T_q^*N is constant.

Now define the Hamiltonians $H^s: T^*N \rightarrow \mathbb{R}$ for all $s \geq 0$ by

$$H^s(z) = \begin{cases} H_\infty(z) + s & z \notin j(DT^*L), \\ f_s(\|j^{-1}(z)\|_L) & z \in j(DT^*L). \end{cases}$$




Figure 5: Hamiltonian H^{s_r}

The reason for introducing this H_∞ is simply to get a slight positive slope at infinity, so that H^s is in \mathcal{H} . These are smooth because $f_s(\|j^{-1}(z)\|_L)$ is smooth (by (f2') above), and it is smoothly extendable by the constant function s on a neighborhood of $j(DT^*L)$ and this agrees with $H_\infty(z) + s$. Since f_s is smooth in s we in fact have that the adjoint function

$$H^{\cdot}: T^*N \times [0, \infty[\rightarrow \mathbb{R}$$

is smooth, making H^s a smooth family of smooth Hamiltonians.

We wish to explicitly construct the TN -spectrum \mathcal{FL} by a limiting process as $s \rightarrow \infty$ of certain generating function spectra associated to H^s defined in Section 6. First we consider the critical set associated to the action on *closed* loops (corresponding to the unparametrized Conley indices).

Lemma 7.1 *For any $s \in [0, \infty]$ the actions of the closed 1-periodic orbits associated to the Hamiltonian H^s with negative action coincide with the action set of f translated down by s and intersected with $] - \infty, 0[$.*

Proof By (35) the actions A_L and A_N on closed loops agree in DT^*L and T^*N . By the explicit construction of H^s (and H_∞ in Appendix A) any closed 1-periodic orbit in T^*N is either

- constant in $D_{2/3}T^*N - j(DT^*L)$, or
- a 1-periodic orbit in $j(DT^*L)$.

The constant curves outside of $j(DT^*L)$ have critical value $-s$, and this corresponds to the unique horizontal tangent to f_s intersecting the 2nd axis at s . The orbits inside $j(DT^*L)$ can be calculated using the same precise formula that defines the action set of f_s . We conclude that there is a 1-1 correspondence between the action of 1-periodic orbits and the action set of f_s .

For the purpose of this section we now assume that $s \in [5, \infty]$. We postpone the rest of the proof to Appendix A.

The action set of $f + s$ is the action set of f shifted down by s . So we need to see that the negative part of the action set of f_s is the same as the negative part of the action set of $f + s$. This is indeed the case by construction of the family f_s for $s \geq 5$. Indeed, (f5) and (f6) show that we have been careful to cap off $f + s$ to get f_s without distorting the negative part of the action set. \square

The above lemma makes it rather easy to construct continuation maps if we did not care about the fiber-wise structure of the Conley indices. However, we need the fiber-wise structure, and thus we need similar control of fiber-critical values. So now we consider these.

Lemma 7.2 *For any $s \geq 5$ all time-1 Hamiltonian flow curves of H^s starting and ending in the same fiber have action less than $\|F\| + 1$, and if the action is less than or equal to $-\|F\|$ the flow curve will lie in the set $T^*N - j(D_{t_s}T^*L)$.*

The most important part of this construction is that we can control the *open* flow lines with action less than the fixed number $-\|F\|$ in such a way that their action is simply translated downwards when s is increased. This we fully prove in Lemma 7.4, but first we need this lemma.

Proof On the set $j(DT^*L)$ the Hamiltonian flow is constructed so that it preserves $\|j^{-1}(\cdot)\|_L$. So either all of a flow-curve is inside DT^*L (or $D_{t_s}T^*L$) or none of it is. Given any time-1 Hamiltonian flow curve $\gamma \in C^\infty(I, DT^*L)$ we consider the two action integrals in (35). The action A_L of γ in DT^*L is given by the usual intersection formula, which works on open paths in DT^*L , because here H^s is a

function of $\|-\|_L$. So the action A_N of $j \circ \gamma$ in T^*N is given by this *plus* the extra factors $F(\gamma(1)) - F(\gamma(0))$.

The first statement then follows by using the above on any time-1 flow curve γ in $j(DT^*L)$ to see that

$$A_N(\gamma) \leq 1 + F(\gamma(1)) - F(\gamma(0)) \leq \|F\| + 1.$$

Here 1 comes as minus the smallest possible intersection from (f5) (combined with (f6) this is a global bound on the intersection with the 2nd axis). By construction of H_∞ any flow-curve outside of $j(DT^*L)$ that starts and ends in the same fiber is constant and has action $-s \leq -5$.

The second statement follows from using it on any flow curve γ inside $j(D_{t_s}T^*L)$ to get:

$$A_N(\gamma) \geq 0 + F(\gamma(1)) - F(\gamma(0)) \geq -\|F\|$$

as needed. Here we used the upper bound 0 from (f5). □

The interval of action (energy) $[a, b]$ that we will be interested in as s tends to infinity is determined by the above lemma. That is, we will now assume that a and b are fixed such that

$$a < -\|F\| \quad \text{and} \quad b \geq \|F\| + 1.$$

Indeed, this makes b an upper bound on all fiber-critical values of the action integrals as we have been assuming in many constructions thus far. The reason for the bound on a is that the critical values below or equal to a behave in a very concrete manner (even fiber-wise), and Lemma 7.4 below is the precise description of this behavior. This idea and Lemma 7.4 will be essential in constructing the fiber-wise continuation maps and proving the fibranity property.

We need a suitable family of subdivisions $\alpha^s, s \in [0, \infty[$. This means that we need to choose a smooth map

$$\alpha^-: [0, \infty[\rightarrow \Delta^\infty$$

such that the finite-dimensional approximations $S_{r_{\alpha^s}}^s$ (see (33)) are defined using (H^s, α^s) for all $s \in [0, \infty[$. Here smoothness is well-defined because the top face inclusions $\Delta^{r-1} \subset \Delta^r$ are smooth and we used the weak topology on Δ^∞ . Such families are easy to construct because all we need α^s to satisfy is

$$l(\alpha^s)(C_1^{H^s} + C_2^{H^s}) \leq \delta,$$

with δ as in Definition 2.2. We may, in fact, choose this so that α^s is constantly equal to the unique subdivision in Δ^0 for s close to 0. Indeed, this follows from the fact that $f_0 = 0$ and property H4 of H_∞ described in Appendix A.

Lemma 7.3 *There exists a strictly increasing sequence $(s_r)_{r \in \mathbb{N}} \in [0, \infty[$ converging to infinity such that*

- S_r^s is defined for all $s \in [0, s_r]$, and
- the flow associated to H^{s_r} has no 1-periodic orbit with action a , ie a is regular for $S_r^{s_r}$.

Proof By construction r_{α^s} is defined and finite for all $s \in [0, s']$ and by compactness of $[0, s']$ it is bounded on such intervals (here we use the topology on Δ^∞). So we may indeed pick a sequence s_r (for all $r \geq 1$) such that S_r^s is defined for $s \in [0, s_r]$. Making each s_r smaller preserves this property and since r_{α^s} starts out being constantly equal to 1 at $s = 0$ we can always rechoose the sequence to make it strictly increasing. By the shift by s in Lemma 7.1 and Sard’s Theorem we see that since a is fixed and negative, the subset of $[0, \infty[$ where the last bullet point is satisfied is open and dense. Hence, we may perturb any sequence slightly downwards to satisfy this. \square

We now pick such a sequence $s_r \in [0, \infty[$ and call them *jump points*. As in Section 2 we may thus define the fiber-wise approximations

$$S_{q,r}^s: T^*\Omega_{q,r}N \rightarrow \mathbb{R}$$

as the restriction of S_r^s to $T^*\Omega_{q,r}N$ for each $s \in [0, s_r]$. We use the notation

$$S_r^-: T^*\Lambda_r N \times [0, s_r] \rightarrow \mathbb{R}$$

for the function defined by $S_r^-(\bar{z}, s) = S_r^s(\bar{z})$, and similarly for $S_{q,r}^-$. These functions are smooth because H^s and α^s are smooth in s .

We may also define the pseudogradients X_r^s and parametrized pseudogradients $Y_{q,r}^s$ as in Section 3. In fact we define X_r^- to be the unique vector field on $T^*\Lambda_r N \times [0, s_r]$ that restricts (and lies in the correct subspace) to X_r^s on $T^*\Lambda_r N \times \{s\}$ for each $s \in [0, s_r]$. This makes X_r^- a parametrized pseudogradient with respect to the projection to the second factor

$$T^*\Lambda_r N \times [0, s_r] \rightarrow [0, s_r].$$

Note that X_r^- is smooth because of how it is defined in Section 2 and because S_r^- is smooth.

Similarly we may define Y_r^- making it a parametrized pseudogradient with respect to the projection

$$(36) \quad \pi = \text{Ev}_0 \times \text{Id}: T^*\Lambda N \times [0, s_r] \rightarrow N \times [0, s_r].$$

Lemma 7.4 *If $\vec{z} \in T^* \Lambda_r N$ is a critical point for $S_{q,r}^{s'}$ with critical value less than $-\|F\|$, for some $s' \in]5, s_r]$, then*

$$\left(\frac{\partial}{\partial s} S_{q,r}^-(\vec{z})\right)(s') = -1.$$

Intuitively this lemma says that the critical values of $S_{q,r}^s$ below $-\|F\|$ are when $s \geq 5$ is increased, simply translated downwards with the same speed. Indeed, this should be considered a fiber-wise version of Lemma 7.1.

Proof By Corollary 2.5, \vec{z} is the $\alpha^{s'}$ -dissection of a time-1 flow curve for $H^{s'}$ with action less than $-\|F\|$. By Lemma 7.2 this flow curve lies in the set $T^*N - j(D_{t_s}, T^*L)$. By (f6) above and how we constructed H^s using f_s , this means that there is an $\epsilon > 0$ such that: for any $z \in T^*N$ sufficiently close to the time-1 flow curve we have

$$H^s(z) = H^{s'}(z) + (s - s'), \quad s \in [s' - \epsilon, s' + \epsilon].$$

This implies that this time-1 flow curve is a time-1 flow curve for all $s \in [s' - \epsilon, s' + \epsilon]$ with the action translated by $s' - s$.

This does not mean that \vec{z} is a critical point for all s in this interval. Indeed, the subdivision α^s is not necessarily constant so the dissection could change. So for each $s \in [s' - \epsilon, s' + \epsilon]$ we define $\vec{z}_s \in T^* \Omega_{q,r} N$ to be the α^s -dissection of this common time-1 flow curve. This is smooth in s because α^s and the flow curve is. By Corollary 2.5, \vec{z}_s is now a critical point for $S_{q,r}^s$ when $s \in [s' - \epsilon, s' + \epsilon]$.

The corollary also tells us that the critical value is the same as the action, which shows that

$$S_{q,r}^s(\vec{z}_s) = S_{q,r}^{s'}(\vec{z}) + (s' - s).$$

This proves that the directional derivative of $S_{q,r}^-$ along the tangent of the curve $s \mapsto (\vec{z}_s, s)$ is -1 , and the fact that \vec{z} is a critical point for $S_{q,r}^{s'}$ implies that the directional derivative along the curve $s \mapsto (\vec{z}, s)$ at the point (\vec{z}, s') is the same. Indeed, the tangent vectors of the two curves at the point $(\vec{z}, s') \in T^\bullet \Omega_{q,r} N \times [0, s_r]$ differ by a tangent that is 0 on the second factor. \square

The above lemma is the reason for the choice of a to be less than $-\|F\|$. Indeed, this gives us the control we need in Section 8 to define the maps q_r in (37) below. These maps are the usual quotient maps of Conley indices, but lifted to ex-maps over N . So, they are the fiber-wise continuation ex-maps that we are after, modulo the fact that we also need the TN -suspensions from Section 5.

Firstly, we pick $r_0 > 0$ such that $s_{r_0} > 5$, and consider any $r > r_0$. Put $W = N$ and define $G: W \times I \rightarrow N \times [0, s_r]$ by

$$G(q, t) = (q, (1 - t)s_{r-1} + ts_r).$$

Then Lemma 7.4 tells us that the adjoint of G , say $g: W \rightarrow C^\infty(I, N \times [0, s_r])$, is space-like with respect to S_r^- (see Definition 8.2), and so Lemma 8.3 provides us with an ex-map (contractible choice)

$$(37) \quad q_{r-1} = p_0^1: I_a^b(S_r^{s_{r-1}}, Y_r^{s_{r-1}})_N \rightarrow I_a^b(S_r^{s_r}, Y_r^{s_r})_N$$

over N . The map induced on the unparametrized indices by q_{r-1} is simply the usual quotient map of Conley indices combined with homotopy invariance when changing H and α (this uses the assumption that a is a regular value for the action).

We now define the TN -spectrum \mathcal{FL} by

$$\mathcal{FL}_r = \begin{cases} N_N & r < r_0, \\ \mathcal{GF}_a^b(H^{s_r}, \alpha^{s_r})_r = I_a^b(S_r^{s_r}, Y_r^{s_r})_N & r \geq r_0. \end{cases}$$

The first r_0 structure maps are trivial (N_N is initial in \mathcal{C}_N^+). For $r \geq r_0$ we define $\sigma_r^{\mathcal{FL}}: \mathcal{FL}_r^{TN} \rightarrow \mathcal{FL}_{r+1}$ as the composition of the two maps:

- The structure map of the generating function spectrum

$$\sigma_r: \mathcal{GF}_a^b(H^{s_r}, \alpha^{s_r})_r^{TN} \rightarrow \mathcal{GF}_a^b(H^{s_r}, \alpha^{s_r})_{r+1},$$

which were also the ex-maps σ_r defined in Section 5.

- The ex-map from (37) above:

$$q_r: \mathcal{GF}_a^b(H^{s_r}, \alpha^{s_r})_{r+1} \rightarrow \mathcal{GF}_a^b(H^{s_{r+1}}, \alpha^{s_{r+1}})_{r+1}.$$

Although we will not need it in the proof of the main theorem, we will in Appendix C discuss the fact that for all possible choices made we get canonically (up to homotopy) ex-homotopy equivalent instances of \mathcal{FL} , meaning that they define an object in the ex-homotopy category up to unique isomorphism. The only thing that we will need for the main theorem is the following computation of the homology, given some homological assumptions on $L \subset T^*N$.

Corollary 7.5 *Let \mathbb{F} be \mathbb{Q} or \mathbb{F}_p for p prime. If either*

- N and L are orientable and $j': L \rightarrow N$ is relative spin, or
- $\mathbb{F} = \mathbb{F}_2$,

then the homology $H_*(\mathcal{FL}, \mathbb{F})$ is isomorphic to the homology of the free loop space $H_*(\Lambda L, \mathbb{F})$ with the usual component-wise regrading by the Maslov class.

Proof Since we assumed that a is regular for $S_r^{s_r}$ for all $r \geq r_0$, we know from Lemma 6.4 that the homology of $\mathcal{GF}_a^b(H^{s_r}, \alpha^{s_r})$ is the reduced homology of the unparametrized index $I_a^b(S_r^{s_r}, X_r^{s_r})$. Since a is regular, the usual homotopy invariance of such Conley indices applies. This was used in [12] to prove that this Conley index is stably (ie, after some standard reduced suspensions) homotopy equivalent to a Thom space

$$(\Lambda^\mu L)^{V_\mu},$$

where $\Lambda^\mu L$ is the loop space of loops in L with length less than μ , $V_\mu \rightarrow \Lambda^\mu L$ is some vector bundle, and μ is the slope of the tangent of $f + s$ intersecting the second axis at $-a$. These identifications were also shown to be compatible with the inclusion of loops when the slope μ was increased to μ' , so that the structure map on reduced homology from $(\mathcal{FL}_r)/N$ to $(\mathcal{FL}_{r+1})/N$ is the reduced homology of the inclusion of $\Lambda^\mu L$ into $\Lambda^{\mu'} L$ with the twisted coefficients associated to $V_{\mu'}$.

This provides the isomorphism for $\mathbb{F} = \mathbb{F}_2$ because here the Thom isomorphism works for any vector bundle, except that we should note that the Maslov index shift occurs because the dimension of V_μ depends over each component on the Maslov index. For general \mathbb{F} we may use the assumptions in the corollary. Indeed, the virtual bundle over ΛL constructed in [12] (which is represented by V_μ minus a trivial bundle on $\Lambda^\mu L$) is identified with $j^*TN^{\oplus(r+1)} - TL + \eta$, where η is the Maslov bundle introduced there. This is orientable if η , TN and TL are all orientable, and thus V_μ is orientable in this case. So, by the assumptions in the corollary we may focus solely on the first Stiefel–Whitney class of the Maslov bundle η .

The Maslov bundle η was constructed in [12] by using a canonical lift of the map $f: L \rightarrow BO$, which classifies the orientable and spinable virtual bundle $TN - TL = j^*TN - TL$ to a map $f': L \rightarrow U/O$ (the infinite Lagrangian Grassmannian U/O maps to BO by forgetting that the vector spaces are Lagrangian). Then looping this we get a map

$$(38) \quad \Lambda L \rightarrow \Lambda(U/O) \rightarrow \Omega(U/O) \simeq \mathbb{Z} \times BO.$$

Here the middle map is the canonical projection using that U/O is a based loop space (in fact an infinite loop space), and the latter map is Bott periodicity as described by Milnor in [16].

Assume for contradiction that η is not orientable. Then there is a map $g: S^1 \rightarrow \Lambda L$ such that the first Stiefel–Whitney class of the composition with the map in (38) is

nontrivial. Using that the map $U/O \rightarrow BO$ (forgetting the Lagrangian structure) is an E_∞ map we get a homotopy commutative diagram:

$$\begin{array}{ccccccc}
 S^1 & \xrightarrow{g} & \Lambda L & \longrightarrow & \Lambda U/O & \longrightarrow & \Omega U/O = \mathbb{Z} \times BO & \longrightarrow & BO \\
 & & \downarrow & & \downarrow & & \downarrow & & \downarrow \\
 & & \Lambda(BO) & \longrightarrow & \Omega(BO) \simeq \mathbb{Z}/2 \times SO & \longrightarrow & SO & &
 \end{array}$$

Indeed, the first square commutes because the projection uses the loop space structure, and the second square because taking product with the constant map to $\{1\} \times BO \subset \mathbb{Z} \times BO$ proves that the composition $\mathbb{Z} \times BO \rightarrow \mathbb{Z}/2 \times SO \rightarrow SO$ is the same up to homotopy on each component. The first Stiefel–Whitney class is the pull-back of the first cohomology group $H^1(BO, \mathbb{Z}/2)$. One may check that the rightmost vertical map on first cohomology

$$(39) \quad \mathbb{Z}/2 \cong H^1(SO, \mathbb{Z}/2) \rightarrow H^1(BO, \mathbb{Z}/2) \cong \mathbb{Z}/2$$

is an isomorphism. So by assumption the induced map on cohomology from the lower right corner to $H^1(S^1, \mathbb{Z}/2)$ is nontrivial.

On the other hand, the map in the diagram

$$S^1 \rightarrow \Lambda L \xrightarrow{f} \Lambda BO$$

is adjoint to a map $S^1 \times S^1 \rightarrow L \rightarrow BO$, which by Lemma 7.7 and the assumptions is homotopy trivial. So, the map $S^1 \rightarrow \Lambda BO$ is seen to be null-homotopic by taking the adjoint of a null-homotopy. □

Remark 7.6 Note that the long exact sequence

4	→	0	→	0	→	0	→
3	→	\mathbb{Z}	$\xrightarrow{\cdot 2}$	\mathbb{Z}	\twoheadrightarrow	\mathbb{Z}_2	→
2	→	0	→	0	→	\mathbb{Z}_2	\xrightarrow{g}
1	\xrightarrow{g}	\mathbb{Z}_2	→	\mathbb{Z}	$\xrightarrow{\cdot 2}$	\mathbb{Z}	\twoheadrightarrow
0	\twoheadrightarrow	\mathbb{Z}_2	→	0	→	0	→
n		$\pi_n(O)$		$\pi_n(U)$		$\pi_n(U/O)$	

on homotopy groups shows that g is an isomorphism. Looping this isomorphism proves the isomorphism in (39). It also shows that if the second relative Stiefel–Whitney class does not vanish on π_2 then the bundle η will *not* be oriented.

Lemma 7.7 Any map $S^1 \times S^1 \rightarrow BO$ with vanishing first and second Stiefel–Whitney class is null-homotopic.

Proof Since the first Stiefel–Whitney class vanishes, the map lifts to a map

$$h: S^1 \times S^1 \rightarrow BSO.$$

One may construct the space BSO with a single 0–cell, no 1–cells and a single 2–cell. Indeed, it is simply connected with $\pi_2(BSO) \cong \mathbb{Z}/2$. So h is homotopic to a map that is cellular and thus factors through the quotient $S^1 \times S^1 / (S^1 \times \{*\} \cup \{*\} \times S^1) \cong S^2$. So, the homotopy type of the map is determined (not uniquely as we will see below) by a degree d . Indeed, it is homotopic to a composition

$$S^1 \times S^1 \xrightarrow{-/\sim} S^2 \xrightarrow{\cdot d} S^2 \subset BSO.$$

However, since the second Stiefel–Whitney class vanishes the degree d is even and since the map

$$S^2 \xrightarrow{\cdot 2} S^2 \subset BSO$$

is null-homotopic the same is true for h . □

8 Parallel transport in parametrized Conley indices

In this section we introduce the notion of parallel transport in parametrized Conley indices, which is the notion we use to construct the ex-maps q_r from (37). Heuristically these are a “continuous combination” of homotopy invariance and the natural quotient maps

$$I_a^b(f, Y)_M \rightarrow I_{a'}^b(f, Y)_M,$$

where $a < a' < b$. We could also simultaneously consider varying b and the natural inclusion/quotient maps:

$$I_a^b(f, Y)_M \rightarrow I_{a'}^{b'}(f, Y)_M$$

However, as we assume that b is always an upper bound on all fiber-critical values this will not be needed. So, in the following we assume that b is always a fiber-regular value. We will illustrate the idea by describing a simple case, and describe what assumptions we need and what parallel transport then is (for the most general description see Lemma 8.3). So assume that

- the base M is the unit interval I ,

- we have a smooth function $f: M' \rightarrow \mathbb{R}$, where M' is the product of a manifold without boundary and I , and
- we have a parametrized pseudogradient Y on M' , with respect to the projection to one factor $M' \rightarrow I$, which satisfies the two properties (C1) and (C2) described in the introduction to Section 3.

This is equivalent to having a smooth homotopy $(f_s, Y_s), s \in I$ of smooth functions and pseudogradients, satisfying (C1) and (C2) for all $s \in I$. The property we need to be able to define parallel transport forward in I is: any critical points x for f_s (ie, $(x, s) \in M'$ is a fiber-critical point for the fiber over $s \in I$) with critical value a satisfies

$$(40) \quad \left(\frac{\partial}{\partial s} f \right)(x, s) < 0.$$

The parallel transport (or homotopy lift if the reader prefers) are then maps

$$p_s^{s'}: I_a^b(f_s) \rightarrow I_a^b(f_{s'})$$

for $0 \leq s \leq s' \leq 1$ continuous in s and s' such that $p_s^s = \text{Id}$. The existence of these maps in both directions when a is regular for all fibers will describe mutual homotopy inverses and thus recovers the notion of homotopy invariance of Conley indices.

The condition in (40) is rather intuitive: if we encounter a critical point with value a when increasing s , it must cross downwards through a such that going forward corresponds to a usual collapse map of indices. So in generic cases this is homotopic to a composition of homotopy invariance when a is regular and collapse maps for the crossings. The trouble is nongeneric cases and the fact that we need these maps to exist in compact families, ie as ex-maps.

For the general setup we assume that (f, Y) is a smooth function and parametrized pseudogradient with respect to the submersion

$$(41) \quad \pi: M' \rightarrow M,$$

of smooth manifolds as in Definition 3.3. We as usual denote the flow of $-Y$ by ψ_t . In addition we are as mentioned above assuming the properties:

- (C1) The flow ψ_t of $-Y$ exists for all time (future and past).
- (C2) Outside a compact set, $Y(f)$ is bounded from below by a positive constant.

We will in this section use the abbreviation

$$I(f)_M = I_a^b(f, Y)_M.$$

To describe the general property we need to be able to construct the parallel transport map we need a few notions. Let $\gamma: I \rightarrow M$ be a smooth path. We may pull-back everything using γ . Ie, we get the pull-back diagram:

$$\begin{array}{ccc} M'_I = \gamma^* M' & \longrightarrow & M' \\ \downarrow \pi_I & & \downarrow \pi \\ I & \xrightarrow{\gamma} & M, \end{array}$$

where the two vertical maps are smooth submersions. The function and parametrized pseudogradient (f, Y) pulls back to a function and parametrized (with respect to π_I) pseudogradient $(f \circ \gamma, \gamma^* Y)$. The manifold M'_I is not in general a product, but there is still a canonical way to ask whether or not (40) is satisfied. Indeed, we may pick a lift $Z: M'_I \rightarrow TM'_I$ of $\partial/\partial s$ in the sense that

$$\begin{array}{ccc} M'_I & \xrightarrow{Z} & TM'_I \\ \downarrow \pi_I & & \downarrow D\pi_I \\ I & \xrightarrow{\partial/\partial s} & TI \end{array}$$

commutes. Then $Z_x(f \circ \gamma)$ for $x \in M'_I$ does not depend on the choice of lift provided x is a fiber-critical point. So we may define $(\partial/\partial s)^x(f \circ \gamma)$ as this unique value when x is a fiber-critical point. We can thus ask that

$$(42) \quad \left(\frac{\partial}{\partial s}\right)^x(f \circ \gamma) < 0$$

for all fiber-critical points x with critical value a .

Definition 8.1 Any smooth curve $\gamma: I \rightarrow M$ is called *space-like* (with respect to f, π and a) if the pull-back above satisfies (42) for the fiber-critical points with critical value a . We denote the space of these paths by $\Gamma^a(f, \pi)$ and topologize it as a subspace of $C^\infty(I, M)$ with the weak topology.

In the following W is compact Hausdorff. In particular any map $W \rightarrow \Gamma^a(f)$ has all derivatives continuous as maps from $W \times I$.

Definition 8.2 The adjoint $G: W \times I \rightarrow M$ of any map $W \rightarrow C^\infty(I, M)$ is called a *path-smooth* homotopy. Furthermore if the image of the adjoint is in $\Gamma^a(M, f)$ we call it a *space-like homotopy*.

For such homotopies we use the notation

$$G_s = G(\cdot, s): W \rightarrow M$$

for $s \in I$, and $C^\infty(I, M) \ni G(w) = G(w, \cdot): I \rightarrow M$ for $w \in W$.

Notice that since the flow of Y is fiber-wise the properties (C1) and (C2) are preserved under the pull-back. In fact the Conley indices pull back as ex-spaces. Indeed, if $h: W' \rightarrow W$ is any map of compact Hausdorff spaces then we have an induced pull-back functor

$$(43) \quad h^*: \mathcal{C}_W^+ \rightarrow \mathcal{C}_{W'}^+$$

of ex-spaces. Indeed, this is simply defined as the usual pull-back of spaces, which has canonical sections and projections.

Lemma 8.3 *Let $I(f)_M$ be a parametrized Conley index with all the assumptions above. Let $G: W \times I \rightarrow M$ be any space-like homotopy. Then there exists a contractible choice of a continuous family of parallel transport ex-maps*

$$p_s^{s'}: G_s^*(I(f)_M) \rightarrow G_{s'}^*(I(f)_M)$$

over W for $0 \leq s \leq s' \leq 1$ such that $p_s^{s''} = p_{s'}^{s''} \circ p_s^{s'}$ and p_s^s is the identity.

By contractible choice we simply mean that all the choices we make to construct these ex-maps form a contractible space, and that the family of course depends continuously on these choices.

Proof First we consider the case where W is a point, and thus G is a single smooth curve $G: I \rightarrow N$. As above we can use G to pull everything back and have everything parametrized by I , and thus in this first case we may assume that $M = I$, and the path is simply the identity on I .

Recall the subspaces

$$C_\tau = \psi_\tau(f^{-1}([a, b])) \cap f^{-1}([a, b])$$

and

$$\underline{C}_\tau = C_\tau \cap f^{-1}(a)$$

defined in Section 4. Assume, using Lemma 4.3, that τ is large enough for this to be compact. This is thus one of the possible definitions of $I(f)_I = C_\tau :_I \underline{C}_\tau$. We see that

$$C_\tau = \{x \in M' \mid f(x) \geq a, f(\psi_{-\tau}(x)) \leq b\},$$

which implies that the boundary satisfies

$$(44) \quad \partial C_\tau \subset f^{-1}(a) \cup (f \circ \psi_{-\tau})^{-1}(b).$$

So let $g = f \circ \psi_{-\tau}$. Notice that Y is also a parametrized pseudogradient for g and b is a fiber-regular value for g .

We define $D_s = C_\tau \cap \pi^{-1}(s)$ and $\underline{D}_s = \underline{C}_\tau \cap \pi^{-1}(s)$ for each $s \in I$ and we thus have

$$C_\tau = \bigcup_{s \in I} D_s \quad \text{and} \quad \underline{C}_\tau = \bigcup_{s \in I} \underline{D}_s.$$

Moreover, we have that the fiber of $I(f)_I$ at $s \in I$ is D_s/\underline{D}_s . The goal is thus to construct maps

$$p_s^{s'}: D_s/\underline{D}_s \rightarrow D_{s'}/\underline{D}_{s'}$$

satisfying the requirements.

The idea is to flow the set $D_s, s \in I$ to the fiber over $s' \in [s, 1]$, in such a way that we get maps induced on the quotients $D_s/\underline{D}_s \rightarrow D_{s'}/\underline{D}_{s'}$. In the following we make this idea precise.

Start by choosing any vector field $Z: M' \rightarrow TM'$ lifting $\partial/\partial s$. To accommodate the general situation later we will only assume that Z is continuous. The fact that the curve is space-like means that we have that

$$Z_x(f) = \left(\frac{\partial}{\partial s}\right)^x f < 0$$

for all $x \in M'$ such that $f(x) = a$ and x is a critical point for $f_s = f|_{\pi^{-1}(s)}$.

We will alter this Z to another lift Z' which has better properties. Indeed we will find a $c \gg 0$ such that the flow of the vector field $Z' = Z - cY: M' \rightarrow TM'$ defines a map from one fiber of the parametrized Conley index to the other. Notice that Z' is indeed another lift of $\partial/\partial s$ since Y is vertical. So the flow of Z' for time $s' - s$ takes the fiber $M'_s = \pi^{-1}(s)$ to the fiber $M'_{s'}$ (when it is defined). The point is that we may pick $c \gg 0$ such that

- $Z'_x(f) < 0$ for all $x \in C_\tau$ with $f(x) = a$, and
- $Z'_x(g) < 0$ for all $x \in C_\tau$ with $g(x) = b$.

Indeed, this is true by compactness of C_τ and in each separate case because:

- When $f(x) = a$ we have either $Y(f) < 0$ or by the space-like assumption $Y(f) = 0$ and $Z(f) < 0$.
- When $g = b$ then $Y(f) < 0$ since b is fiber-regular.

Now let $s \in I - \{1\}$ be given. Then since $D_s \subset C_\tau$ is compact there is an $\epsilon > 0$ such that the flow (which is C^1) of Z' is defined on D_s for time $[0, \epsilon]$. We may shrink this ϵ and get that any point in a neighborhood of $g = b$ in D_s is flown to the set where $g \leq b$. Similar for $f = a$. So by possibly shrinking ϵ again we may assume that the flow of Z' for time $t \in [0, \epsilon]$ maps D_s to

$$g^{-1}([-\infty, b]) \cup f^{-1}([-\infty, a]) = D_{s+t} \cup f^{-1}([-\infty, a])$$

and similarly \underline{D}_s to

$$f^{-1}([-\infty, a]),$$

and hence induces a map of quotients

$$p_s^{s+t}: D_s/\underline{D}_s \rightarrow (D_{s+t} \cup f^{-1}([-\infty, a]))/f^{-1}([-\infty, a]) \cong D_{s+t}/\underline{D}_{s+t}.$$

Obviously p_s^s is the identity, and dividing the interval $[0, \epsilon]$ into smaller bits and composing the resulting maps yields the same map, so we may uniquely extend this to all $s' \geq s$ and the maps will then satisfy

$$p_{s'}^{s''} \circ p_s^{s'} = p_s^{s''}.$$

There are no problems at $1 \in I$ since we can argue backwards in time and find the ϵ above such that the maps exists locally for any $s \in [1 - \epsilon, 1]$ and any $1 \geq s' \geq s$.

For any map $g: W \rightarrow \Gamma^a(f)$ the same argument extends. The pull-back M'' defined topologically by the pull-back diagram

$$\begin{array}{ccc} M'' & \longrightarrow & M' \\ \downarrow & & \downarrow \\ W \times I & \longrightarrow & M \end{array}$$

is, when restricted to $\{w\} \times I$ for each $w \in W$, naturally manifolds M''_w . Their tangent bundles define a vector bundle $E \rightarrow W \times I$ with a surjective map to the vector bundle $W \times TI \rightarrow W \times I$. We may lift $\partial/\partial s$ to a continuous section Z , and also the pull-back of Y defines a continuous section in E . By compactness we may now find $c \gg 0$ as above working for all $w \in W$ simultaneously. The flow of Z' is then defined fiber-wise, but glues together to continuous ex-maps over W .

Contractibility follows from the fact that lifts Z' satisfying the needed equations above are a contractible space. Also one could consider τ as part of the chosen structure, but this is also contractible, and even though the conditions on Z' depends on τ the combined set of choices is still contractible. □

9 3S-fibrancy of \mathcal{FL}

In this section we define the notion of a β -spectrum over a smooth compact manifold M being 3S-fibrant (smoothly and stably Serre fibrant). As we will see in Section 10 this will imply that the homology of the fibers is a local coefficient system and that there is a Serre spectral sequences, with second page isomorphic to the homology of N with coefficients in this local coefficient system. This will converge to the homology of the β -spectrum. It is convenient at this point to introduce the version of morphisms we will use in the categories Sp_W^β . The fact that we restricted our attention to spectra with levels given by compact Hausdorff ex-spaces makes everything a little easier than it would be in a more general setting.

The morphisms of β -spectra we will define can be considered analogous to the notion of maps of colimits over \mathbb{N} of finitely generated Abelian groups. Indeed, even though the homology groups of a β -spectrum are not always finitely generated at each level (eg the Hawaiian earring example) the morphisms of our spectra will behave as if they were, due to compactness. In particular this means that for any level, say r , of a β -spectrum any morphism has to be defined on the image of $\sigma_{r,r'}$ for some structure map with sufficiently large r' . Hence we use the following definition, which really only defines a map in the homotopy category.

A map of β -spectra $f = [(f_r, h_r, k_r)]: \mathcal{A} \rightarrow \mathcal{B}$, with $\mathcal{A}, \mathcal{B} \in \text{Sp}_W^\beta$, is an equivalence class of triples of

- a sequence of increasing integers $k_r \geq r$,
- a sequence of ex-maps

$$f_r: \mathcal{A}_r^{(k_r-r)\beta} \rightarrow \mathcal{B}_{k_r}$$

in \mathcal{C}_W^+ , and

- a sequence of ex-homotopies

$$h_r: \mathcal{A}_r^{(k_{r+1}-r)\beta} \times I \rightarrow \mathcal{B}_{k_{r+1}}$$

between the maps $f_{r+1} \circ \sigma_r^{\mathcal{A}}$ and $\sigma_{k_r, k_{r+1}}^{\mathcal{B}} \circ f_r$ (here we have suppressed some β -suspensions on ex-maps).

The equivalence relation is generated by:

- Ex-homotopy equivalence on both sequences f_r and h_r through ex-maps such that they satisfy the above.
- Replacing k_r with $k'_r \geq k_r$ and thus compatibly β -suspending f_r (and h_r) with $(k'_r - k_r)\beta$ and composing with $\sigma_{k_r, k'_r}^{\mathcal{B}}$.

For W a point this defines a morphisms in the usual homotopy category of spectra (only for spectra with compact Hausdorff levels, if not one has to be a little more careful).

Without the ex-homotopies h_r this would not make much sense as maps of colimits, ie homology would not be a functor. We could simply have asked that the maps were ex-homotopic and not included the ex-homotopies h_r in the structure. However, it is well-known that this produces an extra unwanted equivalence on the morphisms, and although this has little bearing on the main argument we include it so that we really do get morphisms in a more conventional homotopy category of β -spectra.

Let $g: W' \rightarrow W$ be any map between compact Hausdorff spaces. The pull-back g^*S^β of S^β is naturally identified with $S^{g^*\beta}$. Also the smash product over W is pulled back to the smash product over W' . So g^* naturally defines a functor from β -spectra over W to $(g^*\beta)$ -spectra over W' .

Let M be a closed neighborhood retract of some open manifold. Let \mathcal{A} be a β -spectrum in Sp_M^β , and let

$$G: W \times I \rightarrow M$$

with W compact Hausdorff. Then we say that G has a *stable lift* to \mathcal{A} if there is a map of $G^*\beta$ -spectra

$$(45) \quad \bar{G}: (G_0^*\mathcal{A}) \times I \rightarrow G^*\mathcal{A}$$

over $W \times I$, where

- $G_0 = G|_{W \times \{0\}}$,
- we have identified $G_0^*\beta \times I$ with $G^*\beta$ in a way compatible with the canonical choice over $0 \in I$, and
- \bar{G} is the canonical identification (or identity) over $0 \in I$.

The existence of such a stable lift has nothing to do with the choice in the second bullet point. Indeed, any two choices are related by an automorphism, which induces an automorphism of $G^*\mathcal{A}$, which is seen to be the identity because we only defined morphisms over $W \times I$ up to (ex-)homotopy. Also this choice is a contractible choice, since we required it to be the canonical choice over $0 \in I$. In fact, these identifications will in the following construction come for free since we are using the parallel transport from Lemma 8.3. Indeed, using the lemma on a parametrized Conley index that happens to be the sphere bundle $D\beta/S\beta$ (eg having a single nondegenerate critical point in each fiber) produces maps that are ex-homotopic to standard parallel transport in β descended to the sphere bundle $D\beta/S\beta$.

Definition 9.1 An ex-spectrum $\mathcal{A} \in \text{Sp}_M^\beta$ is 3S-fibrant if it has stable lifts of all path-smooth homotopies $G: W \times I \rightarrow M$ with W a compact CW-complex.

The TN -spectrum \mathcal{FL} has this property mainly because of the following lemma and Lemma 7.4. Let H^s, α^s for each $s \in [0, \infty[$ and $a, b, (s_r)_{r \geq r_0}$ be as in Section 8 defining the families of smooth finite-dimensional approximations S_r^s (and $S_{q,r}^s$) and subsequently defining an instance of \mathcal{FL} .

Lemma 9.2 For any $s \in [0, s_r]$ let \vec{z} be a fiber-critical point for S_r^s with respect to $\text{Ev}_0: T^* \Lambda_r N \rightarrow N$, ie a critical point for S_{r,q_0}^s . Then

$$\|(\nabla S_r^s)_{\vec{z}}\| \leq 2.$$

The proof of this is a simple reference to the gradient approximation from Lemma 2.3 of the finite-dimensional approximations. However, the action on the infinite-dimensional space of paths starting and ending in the same fibers behaves in the exact same way. Indeed, the gradient of an open time-1 path for the Hamiltonian flow only depends on the end-points and if we calculate the gradient keeping in mind that we are restricted to paths that start and end in the same fiber within DT^*N we get the exact same gradient and bound as in the proof below.

Proof Corollary 2.5 tells us that the fiber critical points over $q \in N$ are precisely those points for which $\epsilon_{q_j} = 0$ for all j and $\epsilon_{p_j} = 0$ for all $j \neq 0$. The gradient estimates in Lemma 2.3 then tell us that

$$\|\nabla S_r^s\| = \|\nabla_{q_0} S_r^s\| = \|\epsilon_{p_0}\| \leq 2,$$

the latter because all our Hamiltonians are such that S_r^s only has fiber-critical points with all z_j in DT^*N . In particular $p_0 \in DT^*N$ and $p_0^- \in DT^*N$. □

To prove the 3S-fibrancy structure of \mathcal{FL} we need to understand how to create ex-homotopies by “varying” the jump points s_r where we Thom-suspend using the map from Section 5 and combine this with the parallel transports from Section 8. The best way to understand this concept is proving the following concrete lemma.

Lemma 9.3 If we replace s_{r+1} with $s'_{r+1} \in [s_r, s_{r+1}]$ in the definition of \mathcal{FL} we get the same composed structure map $\sigma_{r,r+2}^{\mathcal{FL}}$ up to ex-homotopy.

When we defined \mathcal{FL} we did not actually allow $s_r = s_{r+1}$. However, this was only to exclude the possibility of having to do parallel transport along constant paths (see definition of q_r from (37)) and because it will be convenient later when considering products. It is convenient to allow the case of $s_r = s_{r+1}$ in this section and the argument in the following proof makes it very clear how this can be handled in more generality.

Proof The definition of the composed structure map $\sigma_{r,r+2}^{\mathcal{FL}}$ depends on contractible choices. We will make these choices now such that the similar structure maps, say

$$\sigma_{r,r+2}^{s'_{r+1}},$$

defined using s'_{r+1} instead of s_{r+1} is continuous in s'_{r+1} . Indeed, let

$$P_s^{s'} : I_a^b(S_{r+2}^s, Y_{r+2}^s)_N \rightarrow I_a^b(S_{r+2}^{s'}, Y_{r+2}^{s'})_N$$

be parallel transport maps for $s \leq s' \in [s_r, s_{r+2}]$ defined by the obvious space-like homotopy $N \times I \rightarrow N \times [s_r, s_{r+2}]$ (similar to those defining q_{r+1} in (37)) and let

$$p_s^{s'} : I_a^b(S_{r+1}^s, Y_{r+1}^s)_N \rightarrow I_a^b(S_{r+1}^{s'}, Y_{r+1}^{s'})_N$$

be defined similarly for $s \leq s' \in [s_r, s_{r+1}]$. These are illustrated in Figure 6.




Figure 6: The horizontal lines represent parallel transport forward in s and the vertical dotted line represents when we choose to use the suspension from Section 5.

Then the structure maps may be defined by

$$\sigma_{r,r+2}^{s'_{r+1}} = P_{s'_{r+1}}^{s_{r+2}} \circ \sigma_{r+1}^{s'_{r+1}} \circ p_{s_r}^{s'_{r+1}} \circ \sigma_r,$$

where $\sigma_{r+1}^{s'_{r+1}}$ is the suspension ex-map

$$I_a^b(S_{r+1}^{s'_{r+1}}, Y_{r+1}^{s'_{r+1}})_{TN} \rightarrow I_a^b(S_{r+2}^{s'_{r+1}}, Y_{r+2}^{s'_{r+1}})_N$$

easily defined to be continuous in s'_{r+1} . Indeed, the suspension maps from Section 5 can be defined continuously over any compact smooth family of Hamiltonians and subdivisions. □

Proposition 9.4 Any instance of \mathcal{FL} is 3S-fibrant.

The intuitive idea of the following proof is: Lemma 9.2 tells us that the movement of the fiber-critical points for fixed s is bounded by 2 when following a unit length path in N . Lemma 7.4 tells us that the relevant critical points moves down with speed

1 when s increases by unit length. This means that if we lift any smooth path in N to $N \times [s_r, \infty]$ such that the speed in the second factor is more than two times the speed in the first factor then we get a space-like path. This is also the reason for the suggestive name space-like; indeed, we are lifting the path to the cone where the rate of change of these critical points is guaranteed to be negative by Lemma 7.4.

Proof Let $G: W \times I \rightarrow N$ be any path-smooth homotopy, and let $r \in \mathbb{N}_0$ be given. Then let $c \gg 0$ be a bound on the norm of the first derivative of all the paths $G(w) \in C^\infty(I, N)$. Pick an increasing sequence $k_r \geq r$ such that

$$(46) \quad s_{k_r} - s_r > 2c.$$

We claim that there is a stable lift $\bar{G} = (f_r, h_r, k_r)$ as in (45) with this sequence k_r .

The structure maps of the generating function spectrum $\mathcal{GF}(H^{s_r}, \alpha^{s_r})$ (the ex-map from Section 5) are

$$(47) \quad \sigma_{r,k_r}: \mathcal{FL}_r^{\beta(k_r-r)} = I_a^b(S_r^{s_r}, Y_r^{s_r})_N^{(k_r-r)\beta} \rightarrow I_a^b(S_{k_r}^{s_r}, Y_{k_r}^{s_r})_N.$$

This does not map to the correct ex-space. Indeed, we need to map to the indices defined at s_{k_r} (not s_r). But, moreover we need it to be defined on the pull-back using G_0 and “lift” the homotopy G in the sense of (45).

We will generalize the construction of the structure maps using the bound $s_{k_r} - s_r > 2c$ to construct space-like homotopies, which are equal to restrictions (and reparametrization) of G in the N factor. Indeed, for each $r \geq r_0$ define

$$F = F^r: W \times I \times I \rightarrow N \times [s_r, s_{k_r}]$$

by

$$F(w, t, \tau) = (G(w, \tau t), (s_{k_r} - s_r)\tau + s_r).$$

Then this is space-like (using the τ -factor) with respect to $S_{k_r}^-$ and a . Indeed, with $s = (s_{k_r} - s_r)\tau + s_r$ we see that at a fiber critical point (ie, critical for $S_{q,r}^s$) we have

$$\left(\frac{\partial}{\partial \tau} F\right)(w, t, \tau) \leq 2ct + (s_{k_r} - s_r) \left(\frac{\partial}{\partial S} S_{k_r}^-(s)\right).$$

Indeed, this uses Lemma 9.2, the assumed bound on G , and the chain rule. Then using Lemma 7.4 (remembering that $a < -\|F\|$ and $s_r \geq 5$) and (46) we may bound this to be strictly less than $2ct - 2c$ which is nonpositive for all $t \in I$.

Lemma 8.3 thus gives us parallel transport ex-maps

$$\begin{aligned} p_0^1: F_0^* I_a^b(S_{k_r}^-, Y_{k_r}^-)_N &= G_0^* I_a^b(S_{k_r}^{s_r}, Y_{k_r}^{s_r})_N \times I \\ &\rightarrow F_1^* I_a^b(S_{k_r}^-, Y_{k_r}^-)_N = G^* I_a^b(S_{k_r}^{s_{k_r}}, Y_{k_r}^{s_{k_r}})_N, \end{aligned}$$

over $W \times I$, which when precomposed with the map from (47) define the wanted lifts f_r for each r . However, we also need to construct the homotopies h_r and argue that the morphisms over $0 \in I$ are (homotopic to) the identity.

Firstly we notice that the restriction of each f_r to $0 \in I$ is ex-homotopic to the structure maps (pulled back by G_0). Indeed, the homotopy for $t = 0$ is a homotopy that is constant in the N factor. This means that the map is defined precisely as in the lemma above, except that we are considering more than two suspensions and the longer interval $[s_r, s_{k_r}]$. The natural generalization of the above lemma to several jump points $s'_{r+1} \leq \dots \leq s'_{k_r-1}$ with each $s'_m \in [s_r, s_m]$ shows that f_r restricted to $0 \in I$ is defined precisely as $G_0^* \sigma_{r,k_r}$ if we have $s_r = s'_{r+1} = \dots = s'_{k_r-1}$. Intuitively this means that we get an ex-homotopy between the two maps by sliding down the jump points in the interval to all be equal to s_r .

A generalization of the same idea can be used to construct h_r . For this it is convenient to introduce some notions:

- We say that two homotopies $G', G'' : W \times I \rightarrow N$ are concatenable if $G'_1 = G''_0$.
- When two space-like homotopies G', G'' are concatenable we define the concatenated parallel transport

$$p_0^s = \begin{cases} (p')_0^{2s} & s \in [1, \frac{1}{2}], \\ (p'')_0^{2s-1} \circ (p')_0^1 & s \in [\frac{1}{2}, 1], \end{cases}$$

where p' and p'' are parallel transports associated to G' and G'' , respectively.

- We may generalize this to any number of concatenable homotopies, and in fact in general get parallel transport for any piecewise space-like homotopy.
- The choices of such parallel transports are contractible even when considering different choices of division points. Indeed, the choice of a parallel transport is contractible on each piece, and introducing a redundant division point we may consider the usual parallel transport as defined by the concatenation above.

We now realize that the two compositions $f_{r+1} \circ \sigma_r$ and $\sigma_{k_r, k_{r+1}} \circ f_r$ are in fact both defined in the following general way:

- We have given piece-wise space-like homotopies in $N \times [5, \infty[$, which are space-like because their derivatives lie in the cone given by $\|\gamma'_N(t)\| \leq 2\gamma'_s(t)$. Here the subscript denotes the components in TN and $T[5, \infty[$, respectively.
- We then parallel transport along these homotopies and at several given values of s we stop the parallel transport and suspend to a higher level and continue the parallel transport on this higher level.

The generalization of the above lemma tells us that we can ignore where we do the suspensions. Indeed, we may homotope them all to happen at the beginning, ie at s_r . The space-like homotopies used to define these two different maps are homotopies between the same two maps, ie the maps $W \times I \rightarrow N$ defined using G_0 (constant in the I factor) and all of G . Figure 7 illustrates the two different paths taken for some




Figure 7: Paths in $N \times [0, \infty[$ defined by homotopies for a fixed $(w, t) \in W \times I$. The paths are labeled by which ex-map they induce, and $q_{k_r, k_r+1} = q_{k_{r+1}-1} \circ \dots \circ q_{k_r}$.

fixed $(w, t) \in W \times I$ (by construction we see that all these paths are constant in N when $t = 0$). We can interpolate between these two homotopies and stay within the space-like cone, and we may use parallel transport on the family of such an interpolation to construct the ex-homotopies h_r . □

10 Serre’s spectral sequence and 3S–fibrancy

In Section 9 we introduced the notion of a β –spectrum $\mathcal{A} \in \text{Sp}_N^\beta$ being 3S–fibrant over N . In this section we describe the Serre spectral sequence converging to the homology of such a β –spectrum. In particular we show that page 2 is the homology of N with coefficients in the local coefficient system given by the homology of the fibers. Some may want to call this an Atiyah–Hirzebruch spectral sequence. However, since we are calculating standard singular homology with coefficients \mathbb{F} of a fibered spectrum, and not thinking of this as homology of the base with coefficient in a twisted homology theory, we have kept the Serre terminology.

Firstly – using Appendix B – we will assume that β comes with a trivialization $\beta \cong \mathbb{R}^l \times N$. Secondly, the following lemma tells us that we need not worry that we can only lift path-smooth homotopies.

Lemma 10.1 *Let W be compact, any homotopy $H: W \times I \rightarrow N$ is homotopic rel $W \times \partial I$ to a homotopy that is path-smooth. If the original homotopy was constant on a subspace $A \subset W$ we may assume that the new path-smooth replacement is also constant on A .*

Proof One can compose smooth paths and get smooth paths if one is willing to reparametrize them to make them constant close to the endpoints. So by subdividing the homotopy using compactness we may assume that each path $H(x, \cdot)$ is contained within a δ_0 neighborhood of $H(x, 0)$ for all $x \in W$. Now the family of paths $\gamma(q, q', t)$ given by the unique shortest geodesic between two points q and q' in N with distance less than ϵ_0 are smooth in all three variables. This means that all the derivatives with respect to t (defined locally in charts) are continuous as functions of (q, q', t) . So, let $\pi_W: W \times I \rightarrow W$ and $\pi_I: W \times I \rightarrow I$ be the projections; then, the homotopy defined by $\gamma(H_0 \circ \pi_W, H_1 \circ \pi_W, \pi_I): W \times I \rightarrow N$ is adjoint to a map $W \rightarrow C^\infty(I, N)$ and is homotopic to $H \text{ rel } W \times \partial I$.

The last statement in the lemma follows by construction. \square

As in Appendix B we let $l \in \mathbb{N}$ denote the trivial l -dimensional bundle over any space.

Lemma 10.2 *All fibers of a 3S-fibrant l -spectrum $\mathcal{A} \in \text{Sp}_N^l$ are homology equivalent and their homologies define a local coefficient system on N .*

Proof This is standard, except that from the definition of 3S-fibrant we only have a lifting property for smooth paths and path-smooth homotopies of these smooth paths. However, we simply note that defining the fundamental groupoid using smooth paths and path-smooth homotopies yields the same groupoid as the continuous construction, essentially because of Lemma 10.1. \square

We will need to get a good hold on products and suspensions on the level of chains. So, we need the Eilenberg–Zilber operators. Let A and B be spaces. Define the Eilenberg–Zilber operator

$$P^{n,m}: C_n(A) \otimes C_m(B) \rightarrow C_{n+m}(A \times B)$$

by subdividing $\Delta^n \times \Delta^m$ in the following way: let $\Delta^n = [x_0, \dots, x_n]$ and let $\Delta^m = [y_0, \dots, y_m]$ then each subset $S \subset \{1, \dots, n+m\}$ with $|S| = n$ defines a sequence of cross-products of 0-simplices in $\Delta^n \times \Delta^m$ by starting with (x_0, y_0) , and then if 1 is in S the next point is (x_1, y_0) , but if 1 is not in S it is (x_0, y_1) . Continuing like this, the k^{th} point is $(x_{|S \cap \{1, \dots, k\}|}, y_{k - |S \cap \{1, \dots, k\}|})$. This defines a nondegenerate linear $n+m$ simplex in $\Delta^n \times \Delta^m$, and these subdivide the product into n choose $n+m$ simplices. We then define $P^{n,m}$ on the tensor product of generators $\alpha: \Delta^n \rightarrow A$ and $\beta: \Delta^m \rightarrow B$ as the sum over $\alpha \times \beta: \Delta^n \times \Delta^m \rightarrow A \times B$ precomposed with each of these $(n+m)$ -simplices times the sign on each that corresponds to preserving orientation.

The Eilenberg–Zilber operators are strictly associative since we may index similar nondegenerate simplices in the triple product $\Delta^n \times \Delta^m \times \Delta^k$ by two disjoint subsets

$S_1, S_2 \subset \{1, \dots, n + m + k\}$ of size $|S_1| = n$ and $|S_2| = m$, and the sign is still the orientation-preserving sign. So this defines

$$P^{n,m,k}: C_n(A) \otimes C_m(B) \otimes C_k(C) \rightarrow C_{n+m+k}(A \times B \times C),$$

and similar for more than 3 factors.

Remark 10.3 By keeping track of the orientations one may check that the Eilenberg–Zilber operator satisfies the derivation property

$$(48) \quad \partial P^{n,m}(\alpha \otimes \beta) = P^{n-1,m}(\partial\alpha \otimes \beta) + (-1)^{nm} P^{n,m-1}(\alpha \otimes \partial\beta),$$

which implies that it induces maps on homology

$$H_*(A) \otimes H_*(B) \rightarrow H_*(A \times B).$$

More generally, for pairs (X, A) and (Y, B) it induces maps

$$H_*(X, A) \otimes H_*(Y, B) \rightarrow H_*(X \times Y, A \times Y \cup X \times B).$$

We will use these to define products, and also to systematically treat the suspensions.

Remark 10.4 We will also need the fact that the Eilenberg–Zilber operators are strictly (graded) commutative, that is, the diagram

$$(49) \quad \begin{array}{ccc} C_n(A) \otimes C_m(B) & \xrightarrow{P^{n,m}} & C_{n+m}(A \times B) \\ \downarrow (-1)^{nm} \tau & & \downarrow T_* \\ C_m(B) \otimes C_n(A) & \xrightarrow{P^{m,n}} & C_{m+n}(B \times A) \end{array}$$

commutes, where τ and T are the obvious twists. The sign is obvious by considering that the twist on the product $\Delta^n \times \Delta^m$ is not always orientation-preserving.

All of the above can be done with coefficients in any ring, and we may now formulate the following stable (in the fiber direction) version of the Serre spectral sequence.

Proposition 10.5 *Let \mathbb{F} be any coefficient ring. For a 3S-fibrant $\mathcal{A} \in \text{Sp}_N^\beta$, there is a 1st and 4th quadrant spectral sequence $\{E_{n,m}^r, d_r\}$ such that:*

- *The spectral sequence strongly converges to a filtered quotient of $H_*(\mathcal{A}; \mathbb{F})$.*
- *$E_{n,m}^2 \cong H_n(N; H_m(\mathcal{A}_{|\bullet}; \mathbb{F}))$, where $H_*(\mathcal{A}_{|\bullet}; \mathbb{F})$ is the graded local coefficient system defined by \mathcal{A} from Lemma 10.2. Since N is d -dimensional this implies that the spectral sequence collapses on the $(d + 1)^{\text{st}}$ page.*

Remark 10.6 The proof follows several ideas from Cohen, Jones and Yan [6], Hatcher [10], Hatcher [11] and McCleary [15].

Proof For notational purposes we will only consider $\mathbb{F} = \mathbb{Z}$. However, the general case is similar.

We start by using the Eilenberg–Zilber operators to define the stabilization maps, used in Definition 6.3, on the level of chains. Define the chain map on singular chains

$$(50) \quad \Sigma_* = \Sigma_{*,r}: C_*(\mathcal{A}_r, N) \rightarrow C_{*+l}(\mathcal{A}_{r+1}, N)$$

as the composition of

- The chain map

$$C_*(\mathcal{A}_r, N) \rightarrow C_{*+l}(\mathcal{A}_r \times S^l, N \times S^l \cup \mathcal{A}_r \times \{s_0\}),$$

is given by sending α to $P^{*,l}(\alpha \otimes \beta)$, where $\beta \in C_l(S^l)$ is some fixed representative of the generator of $H_l(S^l, \{s_0\})$. This is a chain map due to the derivation property (48) and the fact that β is closed.

- The chain map induced by the quotient map

$$(\mathcal{A}_r \times S^l, N \times S^l \cup \mathcal{A}_r \times \{s_0\}) \rightarrow (\mathcal{A}_r \wedge_N S^l, N).$$

- The chain map induced by the structure map σ_r of \mathcal{A} .

Now define

$$C_* = C_*(\mathcal{A}) = \operatorname{colim}_{r \rightarrow \infty} (C_{*+lr}(\mathcal{A}_r, N))$$

in the category of chain complexes, using the chain maps Σ_* . Since limits of chain complexes commute with homology we have $H_*(C_*) \cong H_*(\mathcal{A})$. Indeed, Σ_* represents on chains the limit used in Definition 6.3 to define $H_*(\mathcal{A})$.

We call a singular m -simplex $\alpha: \Delta^m \rightarrow \mathcal{A}_r$ n -degenerate over N if the projection $p \circ \alpha$ is a singular simplex in N that is a degeneration of a simplex of dimension less than or equal to n . Ie, there is a commutative diagram

$$\begin{array}{ccc} \Delta^m & \longrightarrow & \mathcal{A}_r \\ \downarrow L & & \downarrow p_{\mathcal{A}_r} \\ \Delta^n & \longrightarrow & N, \end{array}$$

where L is a linear degeneration. So L is given by an order-preserving map from $\{0, \dots, m\}$ to $\{0, \dots, n\}$. Let $F_n C_m(\mathcal{A}_r, N)$ be the free Abelian group generated by the n -degenerate m -simplices in \mathcal{A}_r quotiented by the free subgroup generated

by those in N (or more precisely in the image of the section of \mathcal{A}_r). Since the degeneracy property is preserved when taking boundary we see that $F_n C_*(\mathcal{A}_r, N)$ is a subchain complex. We also see that the chain map Σ_* preserves this property since the suspension is in fiber direction. Indeed in general if $\alpha \in C_n(\mathcal{A}_r)$ and $\beta \in C_m(Y)$ then $P^{n,m}(\alpha, \beta)$ is n -degenerate when projected to \mathcal{A}_r . So we may define the n^{th} filtration of C_* by

$$F_n C_* = \operatorname{colim}_{r \rightarrow \infty} (F_n C_{*+lr}(\mathcal{A}_r, N)) \subset C_*.$$

Since $F_n C_* = 0$ for $n < 0$ and $\bigcup_n F_n C_* = C_*$ we get a filtration of $H_*(C_*)$ as $0 \subset H_*(F_0 C_*) \subset H_*(F_1 C_*) \subset \dots$, giving a spectral sequence $(E_{n,m}^*, d_*)$ converging as described in the first bullet point. Here the initial exact triangle $A^1 \rightarrow A^1 \rightarrow E^1 \rightarrow A^1$ is given by $A_{n,m}^1 = H_{n+m}(F_n C_*)$ and $E_{n,m}^1 = H_{n+m}(F_n C_*/F_{n-1} C_*)$, with the usual maps.

To get the second bullet point we need to inspect the first page and its differential $d_1: E_{*,*}^1 \rightarrow E_{*,-1,*}^1$. First notice that any n -degenerate simplex that is not $(n-1)$ -degenerate has a unique nondegenerate n -simplex in N over which it is nondegenerate. This implies that the quotient $F_n C_*/F_{n-1} C_*$ is given by a direct sum as chain complexes over all the nondegenerate n -simplices in N . Let $\alpha: \Delta^n \rightarrow N$ be such a nondegenerate simplex, and let $F_s^\alpha C_*$ denote the subspace in $F_s C_*$ spanned by those s -degenerate simplices which if projected to N are a composition of face maps and degeneracies of α . Then, the direct summand of the quotient $F_n C_*/F_{n-1} C_*$ corresponding to α may be canonically identified with $F_n^\alpha C_*/F_{n-1}^\alpha C_*$.

Let $h: \Delta^n \times I \rightarrow N$ be a homotopy from the constant simplex mapping to $q = \alpha(1, 0, \dots, 0)$ and α relative to $\Delta^0 = \{(1, 0, \dots, 0)\} \subset \Delta^n$. We may assume by Lemma 10.1 that h is path-smooth and thus invoke Definition 9.1 and get a map (when evaluating at $1 \in I$) of ex-spectra


$$\Delta^n \times \mathcal{A}|_q \rightarrow \alpha^* \mathcal{A}$$

over Δ^n . By using the homotopy in the opposite direction we may get a map of ex-spectra in the other direction

$$\alpha^* \mathcal{A} \rightarrow \Delta^n \times \mathcal{A}|_q.$$

By usual arguments (using a null-homotopy of concatenated homotopies) these two maps are homotopy inverses as $(\alpha^* \beta)$ -spectra over Δ^n (since we are only considering maps up to homotopy this means it is an isomorphism in the category). So they induce a homology isomorphism, but that is not precisely what we are after. Looking at the

universal property for pull-backs



we realize that the chain complex $F_s^\alpha C_*$ is isomorphic to the chain complex given by the span of those simplices in $\alpha^* \mathcal{A}$, which projects to actual simplices in the simplicial structure of Δ^n that are s -degenerate. However, the maps also induce a chain homotopy equivalence of such restricted chain complexes simply because the maps and homotopies are over Δ^n . This means that we may identify the homology of the quotient complex $F_n^\alpha C_*/F_{n-1}^\alpha C_*$ with $H_n(\Delta^n, \partial\Delta^n; H_*(\mathcal{A}_q))$. In fact we may do this such that d_1 is compatible with the usual boundary operator on α . We have thus argued that $(E_{*,m}^1, d_1)$ is chain homotopy equivalent to the normed chain complex $C_*(N; H_m(\mathcal{A}_\bullet))$, which establishes the second bullet point. \square

11 The inclusion of constant loops

The identification of the homology in Corollary 7.5 tells us that there is an inclusion of constant loops $H_*(L, \mathbb{F}) \rightarrow H_*(\mathcal{FL}, \mathbb{F})$. The goal of this section is to prove that this inclusion exists as a map of spectral sequences from the usual Serre spectral sequence for a fibrant replacement of the map $L \rightarrow N$ into the Serre spectral sequence associated to \mathcal{FL} in Proposition 10.5.

The main part of the section is to construct a map of pairs

$$(51) \quad i_k: (DV, SV) \rightarrow (\mathcal{FL}_k, N)$$

for each k sufficiently large (depending on the instance of \mathcal{FL}). Here $V \rightarrow L$ will (under the assumptions of Corollary 7.5) be an \mathbb{F} -oriented kd -dimensional vector bundle. The induced map on homology

$$H_*(L; \mathbb{F}) \cong H_{*+kd}(DV, SV; \mathbb{F}) \xrightarrow{(i_k)_*} H_{*+kd}(\mathcal{FL}_k, N; \mathbb{F}) \rightarrow H_*(\mathcal{FL}; \mathbb{F}) \cong H_{*,*}(\Lambda L; \mathbb{F})$$

will be the inclusion. Note that the latter isomorphism is graded component-wise by the Maslov index. However, the grading on the component containing constant loops is,

indeed, 0. Of course since we want to involve the Serre spectral sequence associated to the fibrant replacement of $L \rightarrow N$ we need i_k to somehow relate to this map. Indeed, we will construct i_k such that after restricting it to the zero-section $L \subset DV$ and projecting to N (using the projection of $\mathcal{F}\mathcal{L}_k$) it is homotopic to the composition $p: L \rightarrow T^*N \rightarrow N$.

Lemma 11.1 *A map i_k as in (51) with the above properties will induce a map from the Serre spectral sequence associated to the fibrant replacement of $p: L \rightarrow N$ into the Serre spectral sequence associated to $\mathcal{F}\mathcal{L}$.*

In the following proof we use several noncompact ex-spaces. However, no subtleties arise from this and we refer to May and Sigurdsson [14] for a thorough discussion.

Proof Define $A = DV :_N^P SV$ by using the map $P: DV \xrightarrow{i_k} \mathcal{F}\mathcal{L}_k \rightarrow N$. This is an ex-space over N . However, it is *not* in general a sphere bundle and not in general fibrant in any way. The map from $A \rightarrow \mathcal{F}\mathcal{L}_k$ induced by i_k is by definition of A an ex-map over N . So if we replace the map with a map of ex-fibrations (see [14]) over N it will induce a map of the Serre spectral sequences associated to the ex-fibrant replacements of A and $\mathcal{F}\mathcal{L}_k$. Since the spectrum $\mathcal{F}\mathcal{L}$ is already 3S-fibrant there is a map from the target spectral sequence into the spectral sequence associated to $\mathcal{F}\mathcal{L}$ in Proposition 10.5. So to prove the lemma we are left with identifying the Serre spectral sequence associated to the ex-fibrant replacement of A as the one associated with the fibrant replacement of the map $p: L \rightarrow N$ (shifted upwards in degree by the dimension of V).

When constructing A and taking ex-fibrant replacement we would get ex-homotopy equivalent results if we replaced the map to N by a homotopy equivalent map. So since P restricted to the zero section is homotopic to $p: L \rightarrow N$ and the inclusion $L \rightarrow DV$ is a homotopy equivalence we may pick a homotopy from P to $P': DV \rightarrow L \xrightarrow{p} N$, where the first map is the projection to the base of the vector bundle. So, the ex-fibrant replacement of A is ex-homotopy equivalent to the ex-fibrant replacement of $A' = DV :_N^{P'} SV$, and hence has the same Serre spectral sequence.

Now let $\pi: P_N L \rightarrow N$ be a fibrant replacement of $L \rightarrow N$ (not an ex-space so standard path-space construction will do). Let $V_P \rightarrow P_N L$ denote the pull-back of V to $P_N L$. One may check that

$$PA' = DV_P :_N SV_P$$

is an ex-fibrant replacement of the ex-space A' .

Recall that the Serre spectral sequence for a fibration $P_N L \rightarrow N$ uses nonrelative homology groups H_* , while the Serre spectral sequence for an ex -fibration uses relative homology groups $H_*(\cdot; B)$, where B is the image of the section. This is compatible with the functor

$$(\text{spaces with a map to } N) \rightarrow \mathcal{C}_N^+$$

given by disjoint union with N , and defining the section and projection as the obvious ones. Indeed, by applying this functor to $\pi: P_N L \rightarrow N$ we see that for all $B \subset N$ the Thom-isomorphism takes the form

$$\begin{aligned} H_*(\pi^{-1}(B)) &\cong H_*(\pi^{-1}(B) \sqcup B, B) \cong H_{*+kd}((DV_P)_{|\pi^{-1}(B)}, (SV_P)_{|\pi^{-1}(B)}) \\ &\cong H_{*+kd}((DV_P)_{|\pi^{-1}(B)} :_B (SV_P)_{|\pi^{-1}(B)}, B) = H_{*+kd}(PA'_{|B}, B). \end{aligned}$$

Note that we used here that the section $B \rightarrow \pi^{-1}(B) \sqcup B$ is a cofibration. Using the fact that the Thom-isomorphism is natural with respect to restriction we now see that the Serre spectral sequence associated to PA' is isomorphic (only from page 1) to the one associated to $P_N L \rightarrow N$ shifted by the dimension of V . □

The rest of the section thus concerns the construction of i_k . Let \mathcal{FL} be defined as in Section 7. In particular $a < -\|F\|$ and $s_r \in [0, \infty]$ is an increasing sequence. We will need k to satisfy

$$(52) \quad s_k > -a \quad \text{and as usual } s_k > 5.$$

So fix k for the rest of this section such that this is true. The level \mathcal{FL}_k is defined using the Hamiltonian H^{s_k} and the subdivision $\alpha^{s_k} \in \Delta^{k-1}$, and so we need to consider these more carefully.

To construct the map i_k we use that the family $H^s, s \in [\epsilon, s_k]$ defines a homotopy from H^{s_k} to an almost constant Hamiltonian H^ϵ . We pick this $\epsilon > 0$ small enough such that:

- ($\epsilon 1$) H^ϵ is sufficiently C^2 -close to H_∞ for S_1^ϵ to be defined.
- ($\epsilon 2$) H^ϵ has only constant periodic orbits of the Hamiltonian flow, hence by (f4) in Appendix A the critical set of S_1^ϵ is $\{-\epsilon, 0\}$.
- ($\epsilon 3$) The interval $]-s, -s + \epsilon[$ is regular for S_k^s for all $s \in [\epsilon, s_k]$.
- ($\epsilon 4$) $-s_k + \epsilon/2 < a$.
- ($\epsilon 5$) Lemma 11.2 holds.

The first is simply (H3) in Appendix A and the fact that $f_0 = 0$. The second follows from (f4) in Appendix A and by making ϵ so small that the slope of f_s is smaller

than the length of any nonconstant closed geodesic on L . The third follows from $(\epsilon 2)$ and Lemma 7.1. Indeed, the action set of f does not depend on s . The fourth follows easily for small ϵ from (52). The last one will be taken care of in the proof of Lemma 11.2. We then define the varying lower bound




Figure 8: Hamiltonian H^ϵ

$$a_v(s) = -s + \epsilon/2 \quad (\text{v for varying}).$$

By $(\epsilon 3)$ above $a_v(s)$ is regular for S_k^s and we thus get by homotopy invariance of Conley indices that

$$(53) \quad I_{-\epsilon/2}^b(S_k^\epsilon, X_k^\epsilon) = I_{a_v(\epsilon)}^b(S_k^\epsilon, X_k^\epsilon) \simeq I_{a_v(s_k)}^b(S_k^{s_k}, X_k^{s_k}) = I_{-s_k+\epsilon/2}^b(S_k^{s_k}, X_k^{s_k}).$$

As usual b is a fixed upper bound on all critical values for any s . We may then compose with the quotient map

$$I_{a_v(s_k)}^b(S_k^{s_k}, X_k^{s_k}) \rightarrow I_a^b(S_k^{s_k}, X_k^{s_k}) = \mathcal{FL}_k/N.$$

Indeed, by $(\epsilon 4)$ above we have $a_v(s_k) < a$. As we will see below the index

$$I_{-\epsilon/2}^b(S_k^\epsilon, X_k^\epsilon)$$

is homotopy equivalent to a Thom-suspension of L and it is well-known that the quotient above recovers the inclusion of constant curves into the loop space when passing to homology (see eg [12]). So this map is the map we need to lift to a map of pairs $(DV, SV) \rightarrow (\mathcal{FL}_k, N)$. The plan for the rest of the section is:

- First we construct a similar lift for $s = \epsilon$ and $r = 1$.
- Then we use the TN -suspensions in Section 5 to get lifts for larger r (but keeping $s = \epsilon$). This is where the vector bundle $V = V_k$ appears for the specific $r = k$ that we chose above.
- Then we use the homotopy invariance in (53) (with some extra structure) to “transport” the lift to work for $s = s_k$.

The first part is to explicitly describe a map that induces a homotopy equivalence on quotients for S_1^ϵ .

Lemma 11.2 *The embedding $j: DT^*L \rightarrow T^*N$ induces a homotopy equivalence*

$$DT^*L/ST^*L \rightarrow I_{a_v(\epsilon)}^b(S_1^\epsilon, X_1^\epsilon).$$

Note that inducing a map to a Conley index means: The embedding j satisfies

$$(54) \quad j(DT^*L) \subset C_\tau \cup (S_1^\epsilon)^{-1}([-\infty, a_v(\epsilon)]) \quad \text{and}$$

$$(55) \quad j(ST^*L) \subset (S_1^\epsilon)^{-1}([-\infty, a_v(\epsilon)]).$$

for some index pair $(C_\tau, \underline{C}_\tau)$. Indeed, this is how we have produced most maps to Conley indices so far.

Proof By construction S_1^ϵ has critical set precisely L disjoint union with

$$(D_{2/3}T^*N - j(DT^*L)),$$

and the critical values are 0 on L and $-\epsilon$ on the rest. Since L is isolated from infinity by the other set this implies that L consists entirely of local maxima. So since $a_v(\epsilon) = -\epsilon/2$ we see that (54) and (55) are satisfied for *any* index pair $(C_\tau, \underline{C}_\tau)$ (in fact the index pair does not depend on τ and is hence unique). We will show that in a neighborhood of L , S_1^ϵ is Morse–Bott. This will imply that the index pairs $(C_\tau, \underline{C}_\tau)$ are in fact diffeomorphic to the disc bundle and sphere bundle of the normal bundle, which is isomorphic to $T^*L \rightarrow L$. So, indeed, the map j induces a homotopy equivalence.

To see that S_1^ϵ is Morse–Bott we consider points $(p_0, q_0) \in \Lambda_1 T^*N = T^*N$ close to L . Let $d = d(q_0, p_0)$ denote the distance from (q_0, p_0) to L in the Riemannian structure induced by L on DT^*L . We may express S_1^ϵ as a sum of two terms,

$$S_1^\epsilon(q_0, p_0) = -H^\epsilon(q_0, p_0) + \int_\gamma \lambda,$$

where $\gamma: S^1 \rightarrow T^*N$ is the small contractible curve defined by the concatenation of the pieces:

- The time-1 Hamiltonian flow line associated to H^ϵ starting at (q_0, p_0)
- The horizontal geodesic in T^*N (using the Riemannian structure induced by the one on N) back to a point in the fiber over $q_0 \in N$
- The vertical geodesic back to the point (q_0, p_0)

Indeed, this is how S_r is defined when $r = 1$ in Section 2. We may bound the symplectic area of this small “right angled triangle” by the length of the flow path squared times a constant. By (f2) in Appendix A this length is $2\epsilon c_0 d$, for some $c_0 > 0$,

in the Riemannian structure induced by L . The length is thus bounded by a constant times $2\epsilon c_0 d$ in the structure induced from N on T^*N (a compactness argument on DT^*L). So we in fact see that close to L we have

$$S_1^\epsilon(q_0, p_0) \leq -\epsilon c_0 d^2 + K(2\epsilon c_0 d)^2.$$

Indeed, H^ϵ is the distance d squared times ϵc_0 by (f2) in Appendix A. For small ϵ this proves that the Hessian of S_1^ϵ is negative definite in the normal bundle to L . \square

The next step is to use the iterations of the suspensions from Section 5. So, we define the vector bundle $V_1 = TL$ over L and inductively define $V_r = V_{r-1} \oplus p^*TN$ over L for $r > 1$. Also we abbreviate

$$V = V_k = TL \oplus (p^*TN)^{\oplus(k-1)}.$$

Lemma 11.3 *There is a map $DV_r \rightarrow T^*\Lambda_r N$ (homotopic to p when restricted to L and projected to N) that induces a homotopy equivalence*

$$(56) \quad DV_r/SV_r \rightarrow I_{a_v(\epsilon)}^b(S_r^\epsilon, X_r^\epsilon).$$

Proof Since $DV_1 = DTL \cong DT^*L$ the lemma is true for $r = 1$ by Lemma 11.2. So we proceed by induction and assume that the lemma is true for $r - 1$ with a map $g_{r-1}: DV_{r-1} \rightarrow T^*\Lambda_{r-1}N$.

The fact that the map induces a map like in (56) means that we may fix an index pair $(C_\tau, \underline{C}_\tau)$ such that

$$(57) \quad g_{r-1}(DV_{r-1}) \subset C_\tau \cup (S_{r-1}^\epsilon)^{-1}]-\infty, a[\quad \text{and}$$

$$(58) \quad g_{r-1}(SV_{r-1}) \subset (S_{r-1}^\epsilon)^{-1}]-\infty, a[.$$

Recall the construction of the suspension maps in Section 5. We may pull the bundle called E in that section back to DV_{r-1} . We thus get a map

$$D_R(g_{r-1}^*E) \xrightarrow{\bar{g}_{r-1}} D_R E \xrightarrow{f_r} T^*\Lambda_r N.$$

Here f_r is the map from (29) and $R \gg 0$ was chosen in that section as well. After applying the same flow as we did in Section 5 this will produce a map, which sends both $S_R(g_{r-1}^*E)$ and $D_R(g_{r-1}^*E)|_{SV_{r-1}}$ to the set $(S_r^\epsilon)^{-1}]-\infty, a[$. Indeed, the latter is already sent there by (58) and (31) and the former was the whole point of making R large and using the flow in Section 5. The disc bundle $D_R(g_{r-1}^*E)$ over the disc bundle DV_{r-1} is homeomorphic to a disc bundle over L . Indeed, it is homeomorphic to the disc bundle of

$$V'_r = V_{r-1} \oplus (g_{r-1}^*E)|_L.$$

Furthermore, the sphere bundle $S(V'_r)$ corresponds under this homeomorphism to the union of the two sets $S_R(g_{r-1}^*E)$ and $D_R(g_{r-1}^*E)|_{SV_{r-1}}$. So we have a map

$$D(V'_r) \rightarrow T^* \Lambda_r N$$

inducing a map

$$D(V'_r)/S(V'_r) \rightarrow I_{a_v(\epsilon)}^b(S_r^\epsilon, X_r^\epsilon).$$

Notice that $D(V'_r)/S(V'_r)$ is homeomorphic to the relative Thom-construction of E on the pair $(D(V_{r-1}), S(V_{r-1}))$. So, by Corollary 5.2 and the fact that g_{r-1} induced a homotopy equivalence we see that this later map is a homotopy equivalences.

Finally, we see that $(g_{r-1}^*E)|_L$ is isomorphic to p^*TN . Indeed, by induction assumption g_{r-1} is homotopic to p when restricted to L and projected to N . Also, we used the projection to N given by q_0 to define E in Section 5. So $V_r \cong V'_r$ and the new g_r is homotopic to p when restricted to L and projected to N . □

The next (and most subtle) step is to use the homotopy (H^s, α^s) and homotope the map from $DV = DV_k$ to induce a similar homotopy equivalence to the index defined at $s = s_k$. Here we need the parallel transport from Lemma 8.3, and since this assumes the lower bound $a_v(s)$ to be constant it is convenient to shift the Hamiltonians and approximations. This will also make the proof of Lemma 11.5 easier. Ie, we define:

$$\begin{aligned} \bar{H}^s &= H^s + c - s \\ \bar{S}_k^s &= S_k^s - c + s \\ \bar{a} &= a_v(s) - c + s = -\epsilon/2 - c \\ \bar{b}(s) &= b - c + s \end{aligned}$$

Here c is chosen such that the asymptotic tangent of \bar{H}^s goes through 0 for one s , hence all s , because we used the same H_∞ for all s . Notice that \bar{S}_k^s is indeed the finite-dimensional approximation associated to (\bar{H}^s, α^s) and X_k^s is still a pseudogradient. We may ignore the fact that $\bar{b}(s)$ depends on s because we may simply pick the common maximum $\bar{b} = \bar{b}(s_k)$ for $s \in [\epsilon, s_k]$. Indeed, all we need is an upper bound on all critical values.

The following lemma may be considered a lift of the parallel transport construction in Lemma 8.3 to the pair (DV, SV) .

Lemma 11.4 *There is a homotopy $h_s: DV \rightarrow T^* \Lambda_k N$, $s \in [\epsilon, s_k]$ which induces homotopy equivalences*

$$h_s/\sim: DV/SV \rightarrow I_{\bar{a}}^{\bar{b}}(\bar{S}_k^s, X_k^s) = I_{a_v(s)}^b(S_k^s, X_k^s).$$

Proof We use the map from Lemma 11.3 for $r = k$ as h_ϵ . As mentioned earlier, since \bar{a} and \bar{b} are regular for all $s \in [\epsilon, s_k]$ we get by homotopy invariance of Conley indices that all these indices are, indeed, homotopy equivalent. The parallel transport maps from Lemma 8.3 using the projection

$$T^* \Lambda_k N \times [\epsilon, s_k] \rightarrow [\epsilon, s_k]$$

realizes the homotopy equivalences of the fibers of the index parametrized over $[\epsilon, s_k]$. However, to get the maps defined on DV before we take the quotients we need an elaboration on these parallel transport maps. Indeed, the maps were constructed by creating a vector field whose induced flow locally defined maps of the quotients. We can use the idea from the proof of Lemma 8.3 to flow the map h_ϵ forward in the s variable. However, we need to be able to control it so that the image of $SV \subset DV$ is contained in the sets defined by

$$(59) \quad (\bar{S}_k^s)^{-1} (]-\infty, \bar{a}])$$

and the image of DV is contained in a set

$$(60) \quad (\bar{S}_k^s)^{-1} (]-\infty, \bar{a}]) \cup C_\tau,$$

where $(C_\tau, \underline{C}_\tau)$ is an index pair for the specific s that we are considering; and this is a little tricky. Indeed, we will need some global bounds that were unnecessary in the proof of Lemma 8.3 (due to compactness and the fact that we only needed the flow to work locally in s).

We thus elaborate on the proof of Lemma 8.3 by constructing the vector field Z' appearing there with certain additional properties. In that proof, Y was the parametrized pseudogradient over the base, which in our case is simply the interval $[\epsilon, s_k]$, and so, consistent with this, we use the notation $Y = X_k^-$ viewed as a parametrized pseudogradient for \bar{S}_k^- over $[\epsilon, s_k]$, not to be confused with the previously defined parametrized pseudogradient Y_k^- , which is parametrized by $N \times [\epsilon, s_k]$.

We see that $Y(\bar{S}_k^-)$ is bounded from below by a positive constant $c' > 0$ on the set

$$W_a = (\bar{S}_k^-)^{-1}(\bar{a}) \subset T^* \Lambda_k N \times [\epsilon, s_k].$$

Indeed, \bar{a} is a regular value for all s and Lemma 3.7 tells us that outside a compact set $Y(\bar{S}_k^-)$ is bounded from below by a nonzero constant. The set W_a is the set we wish to keep SV (and parts of DV) below for all s . We thus define

$$Z' = \frac{\partial}{\partial s} - K(1 + P)Y,$$

for some large $K \gg 0$. In the proof of Lemma 8.3 $\partial/\partial s$ was denoted Z because we did not have a canonical lift of $\partial/\partial s$ to the total space. Also the interval was I not $[\epsilon, s_k]$. Furthermore, we have put an extra factor of $1 + P$ in front of Y , which did not appear in the proof. Recall that $P = \max_j \|p_j\|$ is continuous. By Lemma 11.5 we may pick this $K \gg 0$ such that

$$(61) \quad Z'_{\vec{z},s}(\bar{S}_k^-) < 0 \quad \text{for all } (\vec{z}, s) \in W_a.$$

Indeed, the lemma says that the factor $\partial/\partial s \bar{S}_k^-$ can be bounded by $K(1 + P)Y(\bar{S}_k^-) \geq k(1 + P)c'$. We may also pick this K large enough for the top part of the index pairs to flow into the index pairs, precisely as in the proof of Lemma 8.3, where the top part is the latter part of the union in (44).

It is now clear that all points in W_a are immediately flowed below W_a in positive time, and thus the set in (59) is preserved under positive time flow. So, the image of SV stays below W_a for all $s \in [\epsilon, s_k]$. It is also clear that the positive time flow preserves the set in (60).

The only subtlety left is that we need to argue that the flow of Z' exists until we reach $s = s_k$. This is an easy extension of Lemma 3.6, and requires the fact that any solution to $P' \leq c_1 P + c_2$ does not go to infinity in finite time. Indeed, in Lemma 3.6 we basically used $|Y(P)| \leq 5\delta_0/4$. In this case we see that

$$|Z'(P)| = K(1 + P)|Y(P)| \leq K(1 + P)5\delta_0/4.$$

So the change of P when flowing with Z' is indeed bounded by $c_1 P + c_2$ for some $c_1, c_2 \in \mathbb{R}_+$. Here we are hiding the fact that P is not everywhere smooth. However, fixing a flow line γ for Z' and looking at $P(t) = P(\gamma(t))$ we see that

$$\left| \frac{\partial}{\partial t} \right| P = \limsup_{t \rightarrow t_0+} \left| \frac{P(t) - P(t_0)}{t - t_0} \right| \leq K(1 + P)5\delta_0/4,$$

which is good enough.

So define h_s as the map into $T^* \Lambda_k N \times \{s\} \subset T^* \Lambda_k N \times [\epsilon, s_k]$ given by h_ϵ composed with the flow of Z' for time $s - \epsilon$. □

The next part of the construction is realizing that h_{s_k} from the lemma above actually induces a map to the *parametrized* index. Indeed, we simply do not take the full quotient but the quotient over N and get a map of pairs

$$h_{s_k}: (DV, SV) \rightarrow (I_{\bar{a}}^{\bar{b}}(\bar{S}_k^{s_k}, Y_k^{s_k})_N, N).$$

Here the sphere is sent to $(\bar{S}_k^1)^{-1}([\infty, \bar{a}])$, which is collapsed to N (instead of a point) when defining the parametrized index. We may then compose h_{s_k} with the quotient (over N)

$$I_{\bar{a}}^b(\bar{S}_k^{s_k}, Y_k^{s_k})_N = I_{a_v(s_k)}^b(S_k^{s_k}, Y_k^{s_k})_N \rightarrow I_a^b(S_k^{s_k}, Y_k^{s_k})_N$$

to define

$$i_k: (DV, SV) \rightarrow (I_a^b(S_k^{s_k}, Y_k^{s_k})_N, N) = (\mathcal{FL}_k, N),$$

which then realizes the inclusion of the constant curves on the level of homology.

We used the following technical lemma above.

Lemma 11.5 *There is a $K > 0$ such that: for any $\vec{z} \in T^* \Lambda_k N$ and any $s \in [\epsilon, s_k]$ we have*

$$\left| \frac{\partial}{\partial s} (\bar{S}_k^-(\vec{z}))(s) \right| \leq K(P + 1),$$

where $P(\vec{z}) = \max_j \|p_j\|$.

Proof First notice that \bar{S}_k^s splits as a finite sum

$$\bar{S}_k^s(\vec{z}) = \sum_{j \in \mathbb{Z}_k} f_j(q_j, p_j, q_{j+1}, s),$$

where

$$f_j(q, p, q', s) = \int_{\gamma_j^s} \lambda_N - \bar{H}^s dt + p^-(\exp_q^{-1}(q')),$$

and γ_j^s is the smooth homotopy of curves given by (6) and the smooth homotopies α^s and \bar{H}^s , and $(q^-, p^-) = \gamma^s(\alpha_j^s)$. We suppress that (q^-, p^-) depends on s , but keep it in mind. We define this f_j when

$$(q, p) \in T^*N \quad \text{and} \quad \text{dist}(q', q) \leq \delta_0 \quad \text{and} \quad t \in I.$$

This is slightly more than we need because \bar{S}_k^s is a sum of these with $\text{dist}(q', q) < \delta_0$, but we need the extra boundary for a compactness argument. Since $P \geq \|p_j\|$ we may assume without loss of generality when bounding the derivative $\partial/\partial s f_j$ that $P = \|p\|$.

The rate of change $\partial/\partial s f_j$ of f_j is bounded by compactness on the set given by $\|p\| \leq 1$. Outside this set the integration term vanishes because of the adjustment we did to H^s to define \bar{H}^s . Indeed, the action of any flow curve outside DT^*N is 0 for \bar{H}^s . This implies that we need to bound the s derivative of the function

$$g_j = f_{j|\{\|p\| \geq 1\}}(q, p, q', s) = p^-(\exp_q^{-1}(q'))$$

by a constant times $P + 1$. We, in fact, can bound it proportionally to p . Indeed, this is easy since it scales in p . That is, for $\tau \geq 1$ we have

$$g_j(q, \tau p, q', s) = \tau p^{-1} (\exp_q^{-1}(q'))$$

because the Hamiltonian flow outside of DT^*N is scale equivariant. So any bound on the set $\|p\| = 1$ scales with τ . \square

12 The fiber-wise product

In this section we will construct a fiber-wise product on (some instances of) the TN -spectrum \mathcal{FL} from Section 7. That is, we will for some (not all) instances of \mathcal{FL} construct a product map of spectra

$$(62) \quad \mu: \mathcal{FL} \wedge_N \mathcal{FL} \rightarrow \widetilde{\mathcal{FL}},$$

where $\widetilde{\mathcal{FL}}$ is another instance of \mathcal{FL} , ie, $\widetilde{\mathcal{FL}}$ is defined using another set of input data (see below). There are a couple of very good reasons for using different instances of \mathcal{FL} at this point:

- (1) This is how the maps naturally occur.
- (2) Had we insisted on using the same \mathcal{FL} as the target we would have to “stabilize” the map using ideas from Appendix C. Doing it this way we can actually get a morphism of $2TN$ -spectra that has $k_r = r$ in the definition in Section 9.
- (3) All we really need for the main argument is that it realizes the intersection product on the constant loops on L , and this we only need homologically using the identifications in Corollary 7.5.

Smash products of spectra are rather delicate, and we will only define a very naïve version of these, and this will suffice for our purpose.

Remark 12.1 We will not explicitly need the Chas–Sullivan product in this construction. However, we note that the construction here is based on it. In fact for $L = N$ (the zero section) we recover the spectral sequence by Cohen, Jones and Yan in [6], which converges to the Chas–Sullivan homology ring. This is not that surprising since below we are concatenating the loops in each fiber.

We define the smash product $\mathcal{FL} \wedge_N \mathcal{FL}$ as the $2TN$ -spectrum defined by

$$(\mathcal{FL} \wedge_N \mathcal{FL})_r = \mathcal{FL}_r \wedge \mathcal{FL}_r,$$

with structure maps

$$\sigma_r^\wedge: (\mathcal{F}\mathcal{L}_r \wedge_N \mathcal{F}\mathcal{L}_r)^{2TN} \cong \mathcal{F}\mathcal{L}_r^{TN} \wedge_N \mathcal{F}\mathcal{L}_r^{TN} \xrightarrow{\sigma_r \wedge_N \sigma_r} \mathcal{F}\mathcal{L}_{r+1} \wedge_N \mathcal{F}\mathcal{L}_{r+1}.$$

We can easily consider any other instance $\widetilde{\mathcal{F}\mathcal{L}}$ as a $2TN$ -spectrum by simply forgetting every other ex-space in the sequence $\widetilde{\mathcal{F}\mathcal{L}}_r$, and making the new structure maps the old $\sigma_{2r,2r+2}$. This does not change its homology, and in fact defines an equivalence of categories, but we will not need the latter. We will thus define μ in the category Sp_N^{2TN} .

We start by defining the fiber-wise product on the level of the underlying smooth manifolds as maps

$$G_r: (T^* \Lambda_r N) \times_N (T^* \Lambda_r N) \rightarrow T^* \Lambda_{2r} N$$

given by simply concatenating. Ie, given two points $\vec{z}_1 = (q_0^1, p_0^1, \dots, q_{r-1}^1, p_{r-1}^1)$ and $\vec{z}_2 = (q_0^2, p_0^2, \dots, q_{r-1}^2, p_{r-1}^2)$ with $q_0^1 = q_0^2$ (the definition of \times_N provides precisely this equality) we can define

$$(63) \quad G_r(\vec{z}_1, \vec{z}_2) = (q_0^1, p_0^1, \dots, q_{r-1}^1, p_{r-1}^1, q_0^2, p_0^2, \dots, q_{r-1}^2, p_{r-1}^2).$$

This is well-defined because $\text{dist}(q_{r-1}^1, q_0^2) = \text{dist}(q_{r-1}^1, q_0^1) < \delta_0$ and similar for q_0^1 and q_{r-1}^2 .

Now assume we are given a Hamiltonian H and a subdivision $\alpha \in \Delta^{r-1}$ such that S_r is defined for the pair (H, α) . We may then define \widetilde{S}_{2r} as the finite-dimensional approximation using the Hamiltonian $2H$ and the subdivision $2_r\alpha \in \Delta^{2r-1}$ given by

$$(64) \quad (2_r\alpha)_j = (\alpha_{(j \bmod r)})/2.$$

The twiddle is put on \widetilde{S}_{2r} because we wish to emphasize that we are not using the same Hamiltonian and subdivision. In fact twiddles will signify that we are dealing with structure related to the target $\widetilde{\mathcal{F}\mathcal{L}}$ (which will be constructed during the course of this section). Using that $q_0^1 = q_0^2$ in the definition of G_r it is easy to check that these functions satisfy

$$(65) \quad \widetilde{S}_{2r}(G_r(\vec{z}_1, \vec{z}_2)) = S_r(\vec{z}_1) + S_r(\vec{z}_2).$$

The factor 2 in the Hamiltonian is compensated by the $\frac{1}{2}$ in the subdivision, and is explained by the fact that the action is additive under strict concatenation. Here strict concatenation takes two curves parametrized by an interval of length 1 and spits out a curve parametrized by an interval of length 2. However, since this is not the concatenation we use (we identify the result as parametrized by an interval of length 1 using the unique affine reparametrization) we effectively need to scale up the

Hamiltonian and scale down the length of the pieces in the subdivision to get the same result.

The point of these maps and (65) is that G_r now induces a map of Parametrized Conley indices. Indeed, for any pair $(C_\tau, \underline{C}_\tau)$ in $T^*\Lambda_r N$ as in Section 4, (65) shows that

$$G_r(C_\tau, C_\tau) \subset (\tilde{S}_{2r})^{-1}([2a, 2b]) \quad \text{and}$$

$$G_r(\underline{C}_\tau, C_\tau) \cup G_r(C_\tau, \underline{C}_\tau) \subset (\tilde{S}_{2r})^{-1}([2a, a + b]),$$

and thus as usual by composing with the flow of $-\tilde{Y}_{2r}$ (to get the image inside some \tilde{C}_t defining the target parametrized index) we get an induced ex-map

$$\mu_r: I_a^b(S_r, Y_r)_N \wedge_N I_a^b(S_r, Y_r)_N \rightarrow I_{a+b}^{2b}(\tilde{S}_{2r}, \tilde{Y}_{2r})_N.$$

Note that this, indeed, is an ex-map over N because G_r preserves the projections to N and so does the flow. We would like to say that this defines a map of generating function spectra. However, strictly speaking it does not. Indeed, the doubling of the subdivision in (64) is not compatible with the top face inclusions used to define the suspensions in the generating function spectra. Ie, the diagram

$$(66) \quad \begin{array}{ccc} \Delta^{r-1} & \xrightarrow{2_r(-)} & \Delta^{2r-1} \\ \downarrow d_r & & \downarrow d_{2r+1} \circ d_{2r} \\ \Delta^r & \xrightarrow{2_{r+1}(-)} & \Delta^{2r+1} \end{array}$$

does not commute. The difference is precisely a reordering of the numbers in the sequence (moving one of the introduced zeros around).

To be as precise as possible we now assume that we have

- a smooth family of Hamiltonians $H^s \in \mathcal{H}$, $s \in [0, \infty[$, defined as in Section 7 but with $H_\infty \in \mathcal{H}_\infty$ such that $2H_\infty \in \mathcal{H}_\infty$,
- a smooth family $\alpha^s \in \Delta^\infty$, $s \in [0, \infty[$ of subdivisions,
- constants $a < -2\|F\| - 1$ and $b = \|F\| + 1$ ($\|F\|$ as in Section 7), and
- a strictly increasing sequence $s_r \in [0, \infty[$, $r \geq r_0$,

defining \mathcal{FL} (the source) as in Section 7. Here we have put slightly more restrictions on H_∞ and the bounds a and b than we did in that section, and we will need this to define the product.

We will need the target spectra $\widetilde{\mathcal{FL}}$ to be defined using a smooth family of subdivisions $\tilde{\alpha}^s \in \Delta^\infty$, $s \in [0, \infty[$ such that $\tilde{\alpha}^{s_r} = 2_r(\alpha^{s_r})$. Since

$$l(2_r(\alpha^s)) = \frac{1}{2}l(\alpha^s)$$

there are no problems in choosing such a family such that

$$l(\tilde{\alpha}^s) \leq \frac{1}{2}(\alpha^s).$$

Basically this family could be constructed as a twisted α^s , where “twisting” takes place on each open interval $]s_r, s_{r+1}[$ and compensates for the noncommutativity of the diagram in (66). We thus use the following data to define $\widetilde{\mathcal{FL}}$:

- The smooth family of Hamiltonians $2H^s$.
- A strictly increasing sequence \tilde{s}_r with $\tilde{s}_{2r} = s_r$.
- A smooth family $\tilde{\alpha}^s$ such that $\tilde{\alpha}^{s_{2r}} = 2_r \alpha^{s_r}$ satisfying the above.
- The constants $\tilde{a} = a + b < -\|F\|$ and $\tilde{b} = 2b$.

By the above we have made sure that \tilde{S}_{2r}^s is defined for $s \in [0, \tilde{s}_{2r}]$. It is no problem to extend the sequence \tilde{s}_{2r} to include odd numberings \tilde{s}_{2r+1} . However, we are throwing the odd levels away anyway so in light of Lemma 9.3 it really does not matter. We thus define for each $r \in \mathbb{N}$ such that $s_r > 5$ the ex-map

$$(67) \quad \mu_r: \mathcal{FL}_r \wedge_N \mathcal{FL}_r \rightarrow \widetilde{\mathcal{FL}}_{2r}$$

by the map induced by G_r above.

Lemma 12.2 *The ex-maps μ_r fit together to define a morphism $\mu = (\mu_r, h_r, r)$ in Sp_N^{2TN} as in (62).*

See Section 9 for definition of morphism in Sp_N^{2TN} .

Proof Since the proof is highly technical and not very deep (except the usual issues about reordering suspension for spectra), we only provide a sketch of the argument. Also we do not *really* need this because as an alternative one can prove the main theorem using a compactness argument and a single μ_r , provided r is chosen large enough. However, this way of thinking makes notation much more compact.

Firstly it is convenient to apply the functor from Appendix B to trivialize the bundles, and to do so in a way that makes the trivial bundles even-dimensional (this avoids some sign issues).

The noncommutativity in (66) is taken care of by the fact that the family $\tilde{\alpha}^s$ interpolates between the different $2_r(\alpha^{s_r})$ and $2_{r+1}(\alpha^{s_{r+1}})$. However, inspecting the Thom-suspensions from Section 5 we see that when comparing the two sides we are in fact putting in the vector from the first copy of TN over different q_j , and in fact shifting the identification of the p_j coming from the second factor by one. Indeed, this corresponds to the usual problem with nonassociativity of smash-products of spectra and the rearranging needed to define a map from the product of two spectra to a third spectrum. □

13 The global product

The products μ_r in Section 12 were defined fiber-wise (as ex-maps) over N . In this section we define associated products on the “total spaces” by combining the fiber-wise products with the intersection product on N , and we prove that it will be compatible with the intersection product on L . We will use the description of the intersection product as the map induced by a Pontryagin–Thom collapse map $N \times N \rightarrow N^{TN}$. The goal is thus to extend μ_r from Section 12 to induce a map

$$(\mu_r)_!: \mathcal{FL}_r/N \wedge \mathcal{FL}_r/N \rightarrow (\widetilde{\mathcal{FL}}_{2r}^{TN})/N,$$

and then prove that for large r this extends the intersection product on constant loops in L (compare with Corollary 7.5). Note that we will use the fact that \mathcal{FL} is a TN –spectrum below and use a suspension from Section 5 to get the TN factor on the target of the above map. However, in general it is not important that such a fiber-wise product is defined on a TN –spectrum to induce a global product as above. We use this not so general construction because it is the most convenient way of proving that the product extends the intersection product on L .

The additivity rule in (65) of G_r does not a priori extend to a neighborhood of

$$T^* \Lambda_r N \times_N T^* \Lambda_r N \subset T^* \Lambda_r N \times T^* \Lambda_r N.$$

However, if we insert an extra point in the target $(q, p) \in T^* N$ we can actually accomplish this. So extend G_r to points

$$\begin{aligned} \vec{z}_1 &= (q_0^1, p_0^1, \dots, q_{r-1}^1, p_{r-1}^1) \quad \text{and} \\ \vec{z}_2 &= (q_0^2, p_0^2, \dots, q_{r-1}^2, p_{r-1}^2) \end{aligned}$$

such that all three distances $\text{dist}(q_0^1, q_{r-1}^2)$, $\text{dist}(q_0^2, q_{r-1}^1)$ and $\text{dist}(q_0^1, q_0^2)$ are less than δ_0 . Then

$$G_r(\vec{z}_1, \vec{z}_2) = (q_0^1, p_0^1, \dots, q_{r-1}^1, p_{r-1}^1, q_0^2, p_0^2, \dots, q_{r-1}^2, p_{r-1}^2)$$

lies in $T^* \Lambda_{2r} N$. In the following we are implicitly using Hamiltonians H and $2H$ and subdivisions α and $\tilde{\alpha} = 2_r(\alpha)$ as assumed in (65). We wish to construct a smooth function

$$p = F(\vec{z}_1, \vec{z}_2) \in T_{q_0}^* N$$

such that if we define

$$(68) \quad G_r^+(\vec{z}_1, \vec{z}_2) = (q_0^1, p_0^1, \dots, q_{r-1}^1, p_{r-1}^1, q_0^2, p_0^2, \dots, q_{r-1}^2, p_{r-1}^2, q_0^2, p),$$

which by assumption on \vec{z}_1, \vec{z}_2 above lies in $T^*\Lambda_{2r+1}N$, we get the subadditivity rule

$$(69) \quad S_r(\vec{z}_1) + S_r(\vec{z}_2) \leq \tilde{S}_{2r+1}(G_r^+(\vec{z}_1, \vec{z}_2)),$$

with equality only on the diagonal where $q_0^2 = q_0^1$. We will construct it so that the equality on the diagonal is nondegenerate in the sense that if we move all terms to one side of the equation the diagonal is a nondegenerate critical manifold in the sense of Morse–Bott. Here we used the usual top face inclusion $\tilde{\alpha} \in \Delta^{2r} \subset \Delta^{2r+1}$ to define \tilde{S}_{2r+1} as we did when defining the suspension maps in Section 5. Indeed, as mentioned in the introduction we will use this suspension to get the extra TN factor.

Using Definition 2.2 we calculate (keeping in mind the reindexing from (68)) the difference for arbitrary $p \in T_{q_0^2}^*N$ to be

$$(70) \quad \tilde{S}_{2r+1}(\vec{z}) - S_r(\vec{z}_1) + S_r(\vec{z}_2)$$

$$(71) \quad = (p_{r-1}^1)^-(\exp_{(q_{r-1}^1)^-}^{-1}(q_0^2) - \exp_{(q_{r-1}^1)^-}^{-1}(q_0^1)) + p \exp_{q_0^2}^{-1}(q_0^1).$$

Indeed, since $\tilde{\alpha}_{2r} = 0$ we do not need the minuses on q_0^2 and p in the last term. If $q_0^2 = q_0^1$ (the diagonal on which we worked in Section 12) both terms vanish and we can set p equal to anything. However, to solve when $q_0^1 \neq q_0^2$ we need that the gradient of this thing on the diagonal with respect to both q_0^2 and q_0^1 vanishes. This is easily accomplished by noticing that there is a unique p where this is true. Indeed, the gradient on the diagonal is antisymmetric in q_0^2 and q_0^1 , and since the term

$$p \exp_{q_0^2}^{-1}(q_0^1)$$

has gradient with respect to (q_0^1, q_0^2) equal to $(p, -p)$ there is such a unique p . Then, it is simply a matter of defining the general

$$p \in T_{q_0^2}^*N$$

by choosing some extension of this unique p defined on the diagonal and adding a positive smooth function times

$$-\exp_{q_0^2}^{-1}(q_0^1)$$

to make the quantity in (70) nondegenerately negative for all \vec{z}_1 and \vec{z}_2 . In fact by picking this function we can get the value of \tilde{S}_{2r+1} as small as needed on any compact set disjoint from the diagonal, which we will use in the proof below.

For any $\delta > 0$ and any two maps $p_i: A_i \rightarrow N$ for $i = 1, 2$ we define

$$A_1 \times_N^\delta A_2 = \{(x, y) \in A_1 \times A_2 \mid \text{dist}(p_1(x), p_2(y)) \leq \delta\}.$$

Let $C_\tau \subset T^*\Lambda_r N$ be as in Section 4 defining the parametrized Conley index \mathcal{FL}_r .

Lemma 13.1 *There is a $\delta > 0$ such that the map*

$$G_r^+ : C_\tau \times_N^\delta C_\tau \rightarrow T^*\Lambda_{2r+1}N$$

is well-defined, and being careful about the choices above it will induce a map

$$(\mu_r)! : \mathcal{FL}_r/N \wedge \mathcal{FL}_r/N \rightarrow \widetilde{\mathcal{FL}}_{2r}^{TN}/N,$$

which realizes the Pontryagin–Thom collapse map to the new factor TN .

Proof Since C_τ is compact and satisfies the inequalities assumed on \vec{z}_1 and \vec{z}_2 above, G_r^+ must be well-defined on a small neighborhood of $C_\tau \times_N C_\tau$. This small neighborhood will contain $C_\tau \times_N^\delta C_\tau$ for small δ . By using the subadditivity (and possibly composing with the negative gradient flow) we get a map

$$C_\tau \times_N^\delta C_\tau \rightarrow \widetilde{\mathcal{FL}}_{2r+1}^{TN}/N$$

induced by G_r^+ , which sends $\underline{C}_\tau \times_N^\delta C_\tau \cup C_\tau \times_N^\delta \underline{C}_\tau$ to the base point. Inspecting $C_\tau \times_N^\delta C_\tau$ we see that it is neither open nor closed. It will be a closed set if we add the compact set of $(\vec{z}_1, \vec{z}_2) \in C_\tau \times C_\tau$, which has $\text{dist}(q_0^1, q_0^2) = \delta$. This compact set is isolated from the diagonal, and by being careful with the choices above we can make sure that G_r^+ sends this compact set to points with value less than \tilde{a} . Ie, the lower bound used to define $\widetilde{\mathcal{FL}}_{2r}$. This means that the induced map extends to all of $C_\tau \times C_\tau$ by sending everything outside $C_\tau \times_N^\delta C_\tau$ to the base point. This is the usual idea of Pontryagin–Thom collapse maps, except we now have good control over what G_r^+ does to the action values.

For the last statement notice that when defining the suspension in Section 5 the vectors of TN correspond to the p defined above and with a large factor in front of

$$\exp_{q_0^2}^{-1}(q_0^1)$$

we see that this p really measures the distance of the two points q_0^1 and q_0^2 , which is how one constructs the Pontryagin–Thom collapse map. □

The rest of this section is devoted to proving that for large r this product extends the intersection product on constant loops in L . This is done by combining this construction with the construction in Section 11. We will not need the subtle two last lemmas in that section. Indeed, the statement we need here is a purely homological statement, without any fiber-wise concerns.

Lemma 13.2 With \mathbb{F} as in Corollary 7.5 the diagram

$$\begin{CD}
 H_*(L; \mathbb{F}) \otimes H_*(L; \mathbb{F}) @>\cap>> H_*(L; \mathbb{F}) \\
 @V(i_r)_* \otimes (i_r)_*VV @VV(\tilde{i}_{2r})_*V \\
 H_*(\mathcal{FL}_r, N; \mathbb{F}) \otimes H_*(\mathcal{FL}_r, N; \mathbb{F}) @>(\mu_r)_!>> H_*(\widetilde{\mathcal{FL}}_{2r+1}^{TN}, N; \mathbb{F})
 \end{CD}$$

commutes up to a possible graded sign. Here \cap is the intersection product and the vertical maps are the inclusions of constant loops described in Section 11.

Proof We assume that r is as large as needed in the construction in Section 11, and the notation from that section is used in the following. We need to identify the product on the inclusion constructed there, and we start by taking a closer look at the homotopy $(S_r^s, X_r^s), s \in [\epsilon, s_r]$.

In Section 11 we considered a varying lower bound $a: [\epsilon, s_r] \rightarrow \mathbb{R}$, but we got rid of this s -dependence by translating the functions based on the parameter s (which we shall generalize in Lemma 13.3). We will also need to vary the upper bound b in this proof. Indeed, we will need the varying bounds:

$$\begin{aligned}
 a_v(s) &= -s + \epsilon/2 \\
 b_v(s) &= s/3
 \end{aligned}$$

for $s \in [\epsilon, s_r]$. This is the same $a_v(s)$ as in Section 11, which is regular for S_r^s . The construction of the product is such that we need to define the bounds associated to the target $\widetilde{\mathcal{FL}}$ as:

$$\begin{aligned}
 \tilde{a}_v(s) &= a_v(s) + b_v(s) = -2s/3 + \epsilon/2, \\
 \tilde{b}_v(s) &= 2b_v(s) = 2s/3.
 \end{aligned}$$

Varying bounds is not really problematic when dealing with parametrized Conley indices. However, we postpone the technical details of this to the following lemma, and simply suppress the s from the notation.

We may construct the products $(\mu_r)_!^s$ above for each $s \in [\epsilon, s_r]$. In fact using a compactness argument we may define the family of these as an ex-map:

$$(\mu_r)_!^s: I_{a_v}^{b_v}(S_r^-, X_r^-)_{[\epsilon, s_r]} \wedge_{[\epsilon, s_r]} I_{a_v}^{b_v}(S_r^-, X_r^-)_{[\epsilon, s_r]} \rightarrow I_{\tilde{a}_v}^{\tilde{b}_v}(\tilde{S}_{2r+1}^-, \tilde{X}_{2r+1}^-)_{[\epsilon, s_r]}$$

Lemma 13.3 tells us that there are parallel transports

$$\begin{aligned}
 (72) \quad p_s^{s'} &: I_a^b(S_r^s, X_r^s) \rightarrow I_a^b(S_r^{s'}, X_r^{s'}), \quad \text{and} \\
 \tilde{p}_s^{s'} &: I_{\tilde{a}}^{\tilde{b}}(\tilde{S}_{2r+1}^s, \tilde{X}_{2r+1}^s) \rightarrow I_{\tilde{a}}^{\tilde{b}}(\tilde{S}_{2r+1}^{s'}, \tilde{X}_{2r+1}^{s'}),
 \end{aligned}$$

both for $s \leq s' \in [\epsilon, s_r]$.

This provides a homotopy

$$\tilde{p}_t^{s_r} \circ (\mu_r)_! \circ (p_\epsilon^t \wedge p_\epsilon^t): I_a^b(S_r^\epsilon, X_r^\epsilon) \wedge I_a^b(S_r^\epsilon, X_r^\epsilon) \rightarrow I_{\tilde{a}}^{\tilde{b}}(\tilde{S}_{2r}^{s_r}, \tilde{X}_{2r}^{s_r})$$

for $t \in [\epsilon, s_r]$, which shows that the diagram

$$\begin{array}{ccc} I_{a_v}^{b_v}(S_r^\epsilon, X_r^\epsilon) \wedge I_{a_v}^{b_v}(S_r^\epsilon, X_r^\epsilon) & \xrightarrow{(\mu_r)_!^\epsilon} & I_{\tilde{a}_v}^{\tilde{b}_v}(\tilde{S}_{2r+1}^\epsilon, \tilde{X}_{2r+1}^\epsilon) \\ \downarrow p_\epsilon^{s_r} \wedge_N p_\epsilon^{s_r} & & \downarrow \tilde{p}_\epsilon^{s_r} \\ I_{a_v}^{b_v}(S_r^{s_r}, X_r^{s_r}) \wedge I_{a_v}^{b_v}(S_r^{s_r}, X_r^{s_r}) & \xrightarrow{(\mu_r)_!^{s_r}} & I_{\tilde{a}_v}^{\tilde{b}_v}(\tilde{S}_{2r+1}^{s_r}, \tilde{X}_{2r+1}^{s_r}) \end{array}$$

homotopy commutes. The $a_v(s_r), b_v(s_r), \tilde{a}_v(s_r), \tilde{b}_v(s_r)$ we have used as bounds in the lower horizontal map (s_r were suppressed) are as in Section 11 all smaller than the actual chosen $a, b, \tilde{a}, \tilde{b}$ when we defined the product $(\mu_r)_!$ in Lemma 13.1. Moreover, the construction of $(\mu_r)_!$ is obviously compatible with quotients to larger such bounds if of course these satisfy $\tilde{a} = a + b$ and $\tilde{b} = 2b$ before and after. So we get another commutative diagram

$$\begin{array}{ccc} I_{a_v}^{b_v}(S_r^{s_r}, X_r^{s_r}) \wedge I_{a_v}^{b_v}(S_r^{s_r}, X_r^{s_r}) & \xrightarrow{(\mu_r)_!^{s_r}} & I_{\tilde{a}_v}^{\tilde{b}_v}(\tilde{S}_{2r+1}^{s_r}, \tilde{X}_{2r+1}^{s_r}) \\ \downarrow & & \downarrow \\ I_a^b(S_r^{s_r}, X_r^{s_r}) \wedge I_a^b(S_r^{s_r}, X_r^{s_r}) & \xrightarrow{(\mu_r)_!} & I_{\tilde{a}}^{\tilde{b}}(\tilde{S}_{2r+1}^{s_r}, \tilde{X}_{2r+1}^{s_r}) \end{array}$$

which fits directly below the previous one. Putting them on top of each other we see that the inclusion of constant curves corresponds to going from top to bottom, and so we only need to see that $(\mu_r)_!^\epsilon$ realizes the intersection product. So we need to consider

$$G_r^+: U = T^* \Lambda_r N \times_N^\delta T^* \Lambda_r N \rightarrow T^* \Lambda_{2r+1} N$$

a little more carefully in the case of $(H^\epsilon, \alpha^\epsilon)$.

Since ϵ is small we have from Lemma 11.2 and the ensuing construction that both Conley indices used in the definition of $(\mu_r)_!^\epsilon$ are Thom-spaces over L . In fact the nondegeneracy proved for small s in Lemma 11.2 close to L can easily be extended to S_r^ϵ . Indeed, consider \mathbb{R}^{2n} where the action on closed loops for $H = 0$ is well understood and the constant loops (a copy of \mathbb{R}^{2n}) is a nondegenerate critical manifold, if we change H by a sufficiently small perturbation that has a nondegenerate critical manifold then the action will have this same nondegenerate critical manifold. So let

E_- be a maximal subspace of $T(T^*\Lambda_r N)|_L$ on which the Hessian of S_r^ϵ is negative definite. Also let ν_L be the normal bundle of the diagonal $L \subset L \times L$. We have the commutative diagram:

$$\begin{array}{ccc} U & \xrightarrow{G_r^+} & T^*\Lambda_{2r+1}N \\ \uparrow & & \uparrow \\ L \times L & \xleftarrow{\Delta} & L \end{array}$$

Now $E_- \times E_-$ is a subbundle in the tangent bundle of U along $L \times L$. Restricting it to the diagonal we get $E_- \oplus E_- \subset TU|_L$. The bundle ν_L may also be considered such a subbundle by $\nu_L \subset T(L \times L)|_L \subset TU|_L$, and these two subbundles intersect trivially. Indeed, the Hessian of $S_r^\epsilon(\vec{z}_1) + S_r^\epsilon(\vec{z}_2)$ is zero on $L \times L$ so ν_L is part of the kernel and cannot be part of any negative bundle.

Now G_r^+ is an embedding and by the subadditivity and its nondegeneracy on the diagonal we see that $E_- \oplus E_- \oplus \nu_L$ is sent to a negative subspace, which by dimension counting is maximal. This means that the map realizes the Pontryagin–Thom collapse map from $L \times L$ to $L^{\nu L}$, and the extra negative bundle of $E_- \oplus E_-$ is by assumption oriented, so on the level of homology the lemma follows. \square

Lemma 13.3 *The parallel transport maps in (72) exist.*

Proof The varying bounds are easy to get rid of by simply replacing S_r^s with the function

$$\frac{S_r^s - a_v(s)}{b_v(s) - a_v(s)}.$$

This does not change the fact that X_r^s is a pseudogradient, and now the bounds $a_v(s)$ and $b_v(s)$ are replaced by the constants 0 and 1. We can do the same for the target approximations \tilde{S}_{2r+1}^s using $\tilde{b}_v(s)$ and $\tilde{a}_v(s)$.

The trouble is that when defining \tilde{p}_s^s the value \tilde{a}_v probably has to cross several critical values because it is far from $\tilde{s}_{2r} = s_r$. However, this is where the parallel transport in Lemma 8.3 comes in handy again.

Indeed, for there to be a parallel transport forward in s for \tilde{S}_{2r+1}^s we need that

$$\left(\frac{\partial}{\partial s}\right)^{\vec{z}} \left(\frac{\tilde{S}_{2r+1}^- - \tilde{a}_v}{\tilde{b}_v - \tilde{a}_v}\right) < 0$$

for any $\vec{z} \in T^*\Lambda_r N$ regular for \tilde{S}_{2r+1}^s and such that $\tilde{S}_{2r+1}^s(\vec{z}) = \tilde{a}_v$. This translates into the need for

$$(73) \quad \left(\frac{\partial}{\partial s}\right)^{\vec{z}} \tilde{S}_{2r+1}^- < \frac{\partial}{\partial s} \tilde{a}_v$$

reflecting that \tilde{a}_v is not constant. By Lemma 7.1 the left hand side is -1 and by the choice of \tilde{a}_v the right hand side is $-\frac{2}{3}$. The intuition is that even though the lower bound \tilde{a}_v goes down with speed $-\frac{2}{3}$, the critical values of \tilde{S}_r moves down faster, ie with speed -1 . So, indeed, there is a parallel transport

$$\tilde{p}_s^{s'} : I_{\tilde{a}}^{\tilde{b}}(\tilde{S}_{2r+1}^s, \tilde{X}_{2r+1}^s) \rightarrow I_{\tilde{a}}^{\tilde{b}}(\tilde{S}_{2r+1}^{s'}, \tilde{X}_{2r+1}^{s'}).$$

The source is much easier to handle since a_v is always regular. In fact it is a limit case of the above idea where the bound a_v moves with same speed -1 as s , but then we still have a parallel transport because it is always regular (it moves parallel to the critical values without being one). □

14 The product on the Serre spectral sequence

In this section we describe how the products define a product on the Serre spectral sequence. First we consider an external smash product and the associated map on spectral sequences, which will be an appropriate derivation on all pages inducing the map on subsequent pages. Then we consider the Pontryagin–Thom collapse map, and finally we use the map μ defined in Section 12.

As in the proof of Proposition 10.5 this will work with any coefficient ring \mathbb{F} , but for notational purposes we suppress this and consider only $\mathbb{F} = \mathbb{Z}$. We have borrowed some ideas from Cohen, Jones and Yan [6].

First we define the exterior smash-product \wedge^e as a functor from $\text{Sp}_W^\beta \times \text{Sp}_{W'}^{\beta'}$ to $\text{Sp}_{W \times W'}^{\beta \times \beta'}$ by

$$(\mathcal{A} \wedge^e \mathcal{B})_r = (\mathcal{A}_r \times \mathcal{B}_r) :_{W \times W'} (\mathcal{A}_r \times W' \cup W \times \mathcal{B}_r),$$

so that the fiber at $(x, x') \in W \times W'$ is $\mathcal{A}_r|_x \wedge \mathcal{B}_r|_{x'}$. The structure maps are the obvious ones defined using the identification

$$(\mathcal{A} \wedge^e \mathcal{B})_r^{\beta \times \beta'} \cong (\mathcal{A}_r^\beta \times \mathcal{B}_r^{\beta'}) :_{W \times W'} (\mathcal{A}_r^\beta \times W' \cup W \times \mathcal{B}_r^{\beta'}).$$

The Eilenberg–Zilber operators defined in Section 10 induce maps on the chains relative to the sections and we thus have maps

$$C_*(\mathcal{A}_r, W) \otimes C_*(\mathcal{B}_r, W') \rightarrow C_*((\mathcal{A} \wedge^e \mathcal{B})_r, W \times W').$$

When both β and β' are trivialized of respective dimension l and l' (which we may assume) and one of them is even-dimensional the suspensions Σ_* defined on the chain

level in (50) are compatible with these maps. Ie, the diagram

$$\begin{array}{ccc}
 C_*(\mathcal{A}_r, W) \otimes C_*(\mathcal{B}_r, W') & \longrightarrow & C_*((\mathcal{A} \wedge^e \mathcal{B})_r, W \times W') \\
 \downarrow \Sigma_* \otimes \Sigma'_* & & \downarrow \Sigma''_* \\
 C_{*+l}(\mathcal{A}_{r+1}, W) \otimes C_{*+l'}(\mathcal{B}_{r+1}, W') & \longrightarrow & C_{*+l+l'}((\mathcal{A} \wedge^e \mathcal{B})_{r+1}, W \times W')
 \end{array}$$

commutes. Here Σ''_* is the suspension on chains induced by the structure maps of $\mathcal{A} \wedge^e \mathcal{B}$. Indeed, this follows from the commutativity mentioned in Remark 10.4 and the way the suspension maps on chains were defined. Furthermore, this is graded from the bigraded tensor complex to the graded complex. In addition it adds the filtration degrees. This implies that it induces a map from the tensored spectral sequence, which on each page is a derivation inducing the product on the next page (see eg McCleary [15]).

The next step is the spectral sequence version of the Pontryagin–Thom collapse map. For this we need a relative version of the spectral sequence. So assume $B' \subset B \subset M$ are closed neighborhood retracts of any compact manifold M and \mathcal{A} is a 3S–fibrant l –spectrum over M . Then we have a relative Serre spectral sequence given by the relative chain complex

$$C_*(\mathcal{A}|_B, \mathcal{A}|_{B'}) = \operatorname{colim}_{r \rightarrow \infty} C_{*+lr}(\mathcal{A}_r|_B, \mathcal{A}_r|_{B'} \cup B),$$

which on page two can be identified with $H_*(B, B'; H_*(\mathcal{A}|_\bullet))$. Indeed, this is a relative version of Proposition 10.5.

For any closed proper submanifold $M' \subset M$ with oriented (in general with respect to \mathbb{F}) normal bundle ν of dimension k we would like to create a spectral sequence version of the “Pontryagin–Thom collapse” map

$$(74) \quad H_*(\mathcal{A}) \rightarrow H_{*-k}(\mathcal{A}|_{M'}),$$

where \mathcal{A} is any 3S–fibrant l –spectrum over M . We will start by replacing the spectral sequence with a spectral sequence which is isomorphic from page 2 and onwards.

We may identify $D\nu$ with a closed tubular neighborhood of M . Let $U \supset D\nu$ be a neighborhood whose closure deformation retracts onto $D\nu$, and let $D\nu^c$ denote the open complement of $D\nu$. We may use Hatcher [10, Proposition 2.21] together with the identification of page 2 in our proof of Proposition 10.5, and conclude that the inclusion

$$(75) \quad C_*^{U, (D\nu)^c}(\mathcal{A}) := \operatorname{colim}_{r \rightarrow \infty} C_{*+lr}^{\mathcal{A}_r|_U, \mathcal{A}_r|(D\nu)^c}(\mathcal{A}_r, N) \rightarrow C_*(\mathcal{A}),$$

induces an isomorphism on page 2 and thus also on higher pages (this subcomplex inherits its filtration from $C_*(\mathcal{A})$). So we may replace the spectral sequence for \mathcal{A} with this.

Now pick a deformation retraction of the closure of U onto $D\nu$, where we can as usual use Lemma 10.1 to make sure that it is path-smooth. We may also make sure that everything that starts outside $D\nu$ never enters the interior of $D\nu$, and thus everything outside $D\nu$ must be deformed onto the sphere bundle $S\nu \subset D\nu$. Using this with Definition 9.1 and restricting to $1 \in I$ we get a map

$$p: \mathcal{A}|_{\bar{U}} \rightarrow \mathcal{A}|_{D\nu},$$

where all the fibers outside $D\nu$ are mapped to fibers over points in $S\nu$. We may use this on the chain complex in (75) by sending all simplices α living over U to $p_*(\alpha)$ and the rest to 0. This is a chain map into $C_*(\mathcal{A}|_{D\nu}, \mathcal{A}|_{S\nu})$ because the new boundary we create by “ignoring” some simplices in $D\nu^c$ is all sent to $C_*(\mathcal{A}|_{S\nu})$. It is also filtration-preserving because the projections of simplices are governed by the deformation retraction.

To continue this map into $C_{*-k}(\mathcal{A}|_M)$ we simply do as in [6] and pick a Thom class τ in $C^*(D\nu, S\nu)$ that vanishes on simplices contained in $S\nu$ and on degenerate simplices. Denoting τ pulled back to $\mathcal{A}_r|_{D\nu}$ by τ' we get

$$\tau' \cap (-): C_{*+lr}(\mathcal{A}_r|_{D\nu}, \mathcal{A}_r|_{S\nu}) \rightarrow C_{*-k+lr}(\mathcal{A}_r|_{D\nu}),$$

which is a chain map for the usual reasons and because τ' is zero on simplices in $\mathcal{A}_r|_{S\nu}$. This lowers the filtration by k for the same reasons as in [6] although we above made the choice to cap away the first c indices. This choice makes it commute with our suspensions (which adds to the other end of the indices defining the simplices). So, this is compatible with the limit as r tends to infinity, and thus induces a map of spectral sequences which lowers the degree by $(k, 0)$. To end up in $C_*(\mathcal{A}|_{M'})$ we again simply use a path-smooth deformation retraction of $D\nu$ onto M' with Definition 9.1 and restrict to $1 \in I$ to get a morphism

$$\mathcal{A}|_{D\nu} \rightarrow \mathcal{A}|_M,$$

which preserves filtration.

Using the above steps on $\mathcal{F}\mathcal{L} \wedge^e \mathcal{F}\mathcal{L}$ and the diagonal $N \subset N \times N$ we get a map from the tensor product of the spectral sequence defined using $\mathcal{F}\mathcal{L}$ with itself to the spectral sequence defined by $\mathcal{F}\mathcal{L} \wedge_N \mathcal{F}\mathcal{L}$. Indeed, $\mathcal{F}\mathcal{L} \wedge_N \mathcal{F}\mathcal{L}$ is the restriction of $\mathcal{F}\mathcal{L} \wedge^e \mathcal{F}\mathcal{L}$ to the diagonal. Composing this with the product μ from Section 12 we get a product on the spectral sequence associated to $\mathcal{F}\mathcal{L}$. It lowers degree by $(d, 0)$ due to the Thom

isomorphism used at one of the steps, so if unital the unit would have degree $(d, 0)$ (the fundamental class for L is in fact a unit, but we do not prove this at any point). Since Section 13 and the construction here both realize the Pontryagin–Thom construction combined with the fiber-wise product we have the following lemma.

Lemma 14.1 *The product above induces on nonfiltered homology the same map as the product in Section 13 does up to the shift by d coming from the Thom-isomorphism.*

15 Proof of the main theorem

As usual N is a closed d -dimensional manifold with a Riemannian structure and $L \subset T^*N$ is an exact Lagrangian.

Proposition 15.1 *Let \mathbb{F} be either \mathbb{Q} or \mathbb{F}_p for some p such that the requirements in Corollary 7.5 are satisfied. Assume also that $\pi_1(L) \rightarrow \pi_1(N)$ is surjective. Then $L \rightarrow N$ is a homology equivalence with coefficients in \mathbb{F} .*

Proof Under the assumptions in Corollary 7.5 the induced map

$$(i_r)_*: H_*(L) \rightarrow H_*(\mathcal{FL}) \cong H_{**}(\Lambda L, \mathbb{F})$$

from Section 11 is an injection on the level of homology. Indeed, evaluation at base-point defines a splitting. In Lemma 11.1 we proved that this gives a map of Serre spectral sequences, one coming from fibrantly replacing $L \rightarrow N$ and the one from \mathcal{FL} from Proposition 10.5. More precisely this defines two (at this point possibly different) filtrations on $H_p(L)$, one coming from the Serre filtration over N and one coming from applying $(i_r)_*(\alpha)$ and restricting the filtration associated to \mathcal{FL} . Denote the Serre-filtration $F_n^S H_*(L)$ and the \mathcal{FL} -filtration $F_n^{\mathcal{FL}} H_*(L)$. The fact that Lemma 11.1 says that we have a map of the spectral sequences proves that

$$F_n^S H_*(L) \subset F_n^{\mathcal{FL}} H_*(L).$$

In particular

$$(76) \quad H_n(L) \subset F_n^{\mathcal{FL}} H_*(L),$$

which was not a priori clear because the spectral sequence associated to \mathcal{FL} can be nontrivial in the 4th quadrant.

Claim: $F_{n-1}^S H_n(L) = 0$. Proof of claim: Assume for contradiction that $0 \neq \alpha \in H_n(L)$ is in Serre-filtration $n - 1$. This implies by the above that it is also in \mathcal{FL} -filtration $n - 1$. By assumption \mathbb{F} is a field for which the intersection product on $H_*(L)$, which

is Poincaré dual (up to a sign) to the cup product, is a perfect pairing. So there exists a $\beta \in H_{d-n}(L)$ such that $\beta \cdot \alpha = 1_L \in H_0(L)$. By (76), β is in filtration $d - n$. The construction of the product in Section 14, which by Lemma 13.2 and Lemma 14.1 realizes the intersection product, tells us that the intersection product preserves the \mathcal{FL} -filtration up to a shift by d . This implies that 1_L is in filtration -1 in the \mathcal{FL} -filtration. This is a contradiction, since the entire spectral sequence is supported in the 1st and 4th quadrant.

The above claim implies that the Serre spectral sequence associated to the map $L \rightarrow N$ has support entirely on the 1st axis in the abutment.

The fact that $\pi_1(L) \rightarrow \pi_1(N)$ is surjective implies that the homotopy fiber of $L \rightarrow N$ is connected, and combining this with the fact that the E_∞ term of the Serre spectral sequence for $L \rightarrow N$ is concentrated on the first axis we get that $H_*(L)$ injects into $H_*(N)$. Since L and N are both d -dimensional manifolds this implies that the degree is invertible and thus the map is also surjective on homology. \square

This now implies the main theorem rather easily.

Main Theorem *If N is oriented and the induced map p_* on fundamental groups is surjective then the induced map p_* on homology is an isomorphism.*

Proof of Main Theorem Since the assumptions in Corollary 7.5 are satisfied for $\mathbb{F} = \mathbb{F}_2$ the above proposition proves that p_* is a \mathbb{F}_2 -homology equivalence. It is a fact that any \mathbb{F}_2 -cohomology equivalence of manifolds preserves the Stiefel–Whitney classes of the tangent bundles (this is also mentioned by Fukaya, Seidel and Smith in [9] and used by Abouzaid in [4] in the exact same context). This implies that L is also oriented. It also implies that $p: L \rightarrow N$ is relative spin. This means that the assumptions in Corollary 7.5 are satisfied for any \mathbb{F} .

The above proposition thus works for all $\mathbb{F} = \mathbb{F}_p$ and $\mathbb{F} = \mathbb{Q}$, and it is well-known that a map which is a homology equivalence for all these coefficients is a homology equivalence with \mathbb{Z} coefficients. \square

Appendix A: Construction of Hamiltonians

Asymptotic behavior and \mathcal{H}_∞

In the constructions throughout the paper we use functions $H_\infty: T^*N \rightarrow \mathbb{R}$ with certain properties. We now construct the space $\mathcal{H}_\infty \subset \mathcal{H}$ in which these should lie.

We need this space to satisfy several assumptions, and the easiest way to get this is to explicitly construct it and then prove all of these.

So fix any $h: \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}$, which satisfies

- h is convex,
- $h(t)$ is zero for $t \in [0, 2/3]$, and
- $h(t) = \mu t + c$ for $t \in [1, \infty[$ for some constants $\mu, -c \in \mathbb{R}_+$.

Define H_∞ to be in \mathcal{H}_∞ if

$$H_\infty(q, p) = \epsilon h(\|p\|)$$

for some $0 < \epsilon < \beta$ with $\beta > 0$ small enough for (H1) through (H4) below to hold.

For $\beta > 0$ small enough we have for any $H_\infty \in \mathcal{H}_\infty$ that:

- (H1) H_∞ is 0 on $D_{2/3}T^*N$.
- (H2) $tH_\infty \in \mathcal{H}_\infty$ for all $t \in]0, 1]$.
- (H3) Any time-1 Hamiltonian flow line for H_∞ starting and ending in the same fiber T_q^*N is constant.
- (H4) The finite-dimensional approximation S_1 from Section 2 is defined using the unique subdivision in Δ^0 and any Hamiltonian in \mathcal{H} sufficiently C^2 -close to H_∞ .

Obviously (H1) and (H2) hold by construction. (H3) follows because the Hamiltonian flow of H_∞ is a reparametrization of the geodesic flow with constant speed $\|\nabla H_\infty\|$. So if $\|\nabla H_\infty\| \leq \delta_0$ (smaller than the injective radius) then it cannot return to the same starting point, unless it had speed 0, in which case it is constant. (H4) is a simple matter of making sure that

$$C_1^H + C_2^H < \delta,$$

with δ from Definition 2.2.

Constructing f_s from f

In this section we construct the smooth family of smooth increasing functions

$$f_s: \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}, \quad s \in [0, \infty[,$$

used in Section 7 and later. Throughout the paper we need this to satisfy several properties.

We assume that we are given a smooth function $f: \mathbb{R}_+ \rightarrow \mathbb{R}$ such that:

- $f(t) \rightarrow -\infty$ when $t \rightarrow 0$.
- $f(t) = 0$ when $t \in [1, \infty]$.
- $f''(t) < 0$ when $t \in]0, 1[$, ie f is strictly concave on $]0, 1[$.

For each $s > 0$ we define t_s to be the unique point $t_s \in]0, 1[$ such that the tangent of f at t_s intersects the 2nd axis at $-s$ (see Figure 3).

The first properties we want (see Figure 3) for all f_s are:

- (f1) $f_s: \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}$ is a smooth family for $s \in [0, \infty[$ of smooth *increasing* functions.
- (f2) In a neighborhood of the set $(0, s) \subset \mathbb{R}_{\geq 0}^2$ we have $f_s(t) = c(s)t^2$ for some $c(s) \geq 0$, and $c(s) = sc_0$ for s close to 0 and some constant $c_0 \in \mathbb{R}_+$.
- (f3) For $t \geq 1$ we have $f_s(t) = s$.
- (f4) For $s > 0$ the restriction of $f_s(t)$ to $]0, 1[$ is strictly increasing.

Of course these force $f_0 = 0$. In the construction of \mathcal{FL} we need the following for $s \geq 5$ (here 5 is a rather arbitrary choice).

- (f5) The tangent to f_s at any $t \in [0, t_s]$ intersects the 2nd axis in $] -1, 0]$.
- (f6) For $s \geq 5$ we have $f_s(t) = f(t) + s$ when $t \geq t_s$.

At other points in the paper we will need for all $s \in [0, \infty[$ that:

- (f7) The tangent to f_s at any $t \in [0, 1]$ intersects the 2nd axis above $-s/4$.
- (f8) Lemma 7.1 is true for all $s \in [0, \infty[$.

The last one uses the definition of the Hamiltonians H^s in Section 7, and depends on the fixed choice of Riemannian structure on L from that section, which we will thus assume is given.

As Figure 9 suggests it is not difficult to use bump functions and cut-off functions together with f' to construct a smooth family $h_s, s > 0$ satisfying (f1)–(f7) for all $s > \epsilon$ for any $\epsilon > 0$ (the ϵ makes the last part of (f2) void). Indeed, pick h_s as the unique antiderivative to the pictured h'_s (which is equal to f' for $t \geq t_s$) going through $(0, 0)$. By making this h'_s depend smoothly on s we get (f1). We get the first part of (f2) because h'_s is linear close to 0. We get (f4) trivially, and by making sure that the two areas marked with a are the same, this will automatically satisfy (f6) and thus also (f3). Properties (f5) and (f7) follow if we make h'_s have a unique critical point (a maximum) close to t_s , as suggested in the figure. Indeed this controls the inflection point of h_s . We can make this maximum value as close to $f'(t_s)$ as we would like by increasing c and thus decreasing the area a . We may thus assume that




Figure 9: Constructing h'_s

- h'_s tends to 0 when s tends to 0.

The proof of Lemma 7.1 then works as written for this family for all $s > 0$. The problem is having (f6) satisfied close to 0, and retaining smoothness (C^2 would be enough, but this seems no easier).

We thus use that h'_s tends to zero when s does. Indeed, pick some $s_0 > 0$ small enough such that $h'_s(t)$ is less than the shortest nonzero geodesic length of L for all t and all $s \in [0, s_0]$. Then the family

$$\frac{s}{s_0} h_{s_0}: \mathbb{R} \rightarrow \mathbb{R}$$

satisfies all the requirements for $s < s_0$ ((f6) is now considered void because $s < 5$). Indeed, the family satisfies all but the last part of (f2), which is now taken care of by the explicit construction. All properties (f1)–(f7) are for fixed s preserved by convex combinations. So we may choose any smooth function $\psi: \mathbb{R}_{\geq 0} \rightarrow [0, 1]$ with $\psi(s) = 0$ for $s \leq s_0/2$ and $\psi(s) = 1$ for $s \geq s_0$ and define

$$f_s(t) = (1 - \psi(s)) \frac{s}{s_0} h_{s_0}(t) + \psi(s) h_s(t).$$

This then satisfies all the requirements. Indeed, it is smooth (f1), satisfies the last part of (f2) by construction, and first part of (f2) and (f3)–(f7) are handled by the convex combination argument. Even (f8) is true simply because f_s for $s \in [0, s_0]$ has f'_s bounded by the length of any nonzero geodesic on L , and so only the constant geodesics contribute, and by (f4) these lie in the set defined by $\|p\|_L = 0$ or $\|p_L\| \geq 1$, which were already considered in the proof of Lemma 7.2.

Appendix B: A functor from β -spectra to parametrized spectra

In Section 6 we defined β -spectra and their homology. In this appendix we describe a functor F from β -spectra over W to l -spectra, where $l \in \mathbb{N}$ denotes the trivial bundle $W \times \mathbb{R}^l$. This functor preserves the homology in the cases that we care about, and we used it to simplify the construction of the spectral sequence in Section 10.

Let $i: \beta \rightarrow W \times \mathbb{R}^l$ be an embedding of the bundle β over W . Then we get a canonical isomorphism $\beta^\perp \oplus \beta \cong l = W \times \mathbb{R}^l$, where β^\perp is the orthogonal complement bundle of β in $W \times \mathbb{R}^l$. We then define the functor F on an object $\mathcal{A} \in \text{Sp}_W^\beta$ by

$$(F\mathcal{A})_r = \mathcal{A}_r^{r\beta^\perp}.$$

The new structure maps are defined by

$$(F\sigma_r) = \sigma_r^{\beta^\perp}: \mathcal{A}_r^{r\beta^\perp \oplus l} \cong \mathcal{A}_r^{r\beta^\perp \oplus \beta^\perp \oplus \beta} \cong \mathcal{A}_r^{\beta \oplus (r+1)\beta^\perp} \rightarrow \mathcal{A}_{r+1}^{(r+1)\beta^\perp}.$$

This is easily extended to morphisms since the fiber-wise Thom suspension is functorial on ex-spaces as described in Section 6.

Global homology is easily seen to be preserved in all cases of interest. Indeed, we made sure that the sections are cofibrations and that β (and hence β^\perp) is oriented. We could also argue that for a 3S-fibrant β -spectrum the homologies of the fibers are preserved, ie, the nongenericity of the fibers goes away. However, we will not actually need this since we never use what exactly the homology of the fibers are in any of the arguments.

Lemma B.1 *The functor F preserves 3S-fibrancy.*

Proof This follows from the fact that any stable lift for \mathcal{A} can, by using parallel transport in (copies of) β^\perp , be turned into a stable lift for $F(\mathcal{A})$. \square

Appendix C: Homotopy invariance of \mathcal{FL}

In this appendix we discuss why the construction of \mathcal{FL} is unique up to unique isomorphism in the homotopy category of parametrized spectra over N . A similar discussion can be done for the product, but we omit this.

Following notation from Section 7 we assume for the moment that f and the capping off family f_s is fixed. The extra choices of $(a, b, (s_r, q_r, \sigma_r)_{r \in \mathbb{N}})$ in defining an instance

of \mathcal{FL} is contractible, that is, if we forget the assumption that a should be a regular value for each $S_r^{S^r}$. We only assumed this so that the sections would be cofibrations making subsequent calculations easier, so we will simply not assume this here. The argument in Section 9 that \mathcal{FL} is 3S-fibrant can be extended to any compact smooth family of such a set of choices. That is, if K is a compact subset of a smooth manifold that smoothly parametrizes choices as above then we get a 3S-fibrant TN -spectrum over $N \times K$. So, in fact the object in the homotopy category does not depend on these choices up to unique isomorphism.

This can be extended to families of (f, f_s) if we ask that the family of f 's agree close to 0, and that the same is true for f_s for large $s > 0$ (on the same given neighborhood of 0). Indeed, with this we can see that the differences in the functions are pushed away from what happens at a critical values a if s is large enough, and so the fibrancy still holds at least for large s .

For two different choices $(f^i, f_s^i), i = 1, 2$ that do not agree like above one can create interpolations and construct a zig-zagging argument. Indeed, we can construct a pair (F, F_s) such that away from a small neighborhood of 0, $F = f^1$, but on an even smaller neighborhood we have $F = f^2$; we may thus construct the sequence F_s such that for small s it looks like f_s^1 and for large s it looks like f_s^2 close to 0. Sequences of these that agree with f^1 on smaller and smaller sets can be used to construct a map from the instance defined by (f^1, f_s^1) to the one defined by (f^2, f_s^2) . To prove that this is an equivalence we create interpolations from (F, F_s) to (f^1, f_s^1) , which on smaller and smaller sets agree with (F, F_s) , but which look like (f^2, f_s^2) on a set (ϵ, ϵ') (here one needs to choose ϵ dependent on ϵ' for the last part to work) and for some range of s such that we can factor the identity on the instance defined by (f^1, f_s^1) through the map defined by the first sequence.

Appendix D: Coherent orientations and finite-dimensional approximations

To understand the issue in Remark 1.3 better we will in this appendix describe why coherent orientations can be subtle from the point of view of finite-dimensional approximations. We will also relate this to stable homotopy types and discuss why in some sense Viterbo functoriality is not natural for the stable homotopy types, unless extra structure is considered as in [12] where we realize it as a map of spectra.

Consider the following abstract situation: assume that $i: M \rightarrow M'$ is a finite-dimensional manifold M embedded into an infinite-dimensional manifold M' . Also assume

that we have a Morse–Smale function $f: M' \rightarrow \mathbb{R}$ for which some sort of infinite-dimensional Morse homology with \mathbb{Z} coefficients can be defined, even though Morse indices happen to be infinite. Furthermore, assume that all critical points of f and all flow lines between such (defined by a PDE and not a flow) are contained in the image of i as an isolated invariant set $S \subset M$ (see Conley [7] for definition of isolated invariant set) for $f \circ i$. To define this we use some pseudogradient for $f \circ i$.

In such a situation we cannot conclude that the homology of the Conley index of S equals the Morse homology defined for f . Indeed, to define Morse homology of f with \mathbb{Z} coefficients we need some sort choice of coherent orientations that depends on more than just the isolated invariant set. To see this take any finite-dimensional vector bundle $E \rightarrow M$, whose restriction to S is not orientable. Then because the Morse indices of f are infinite there will often (if not always) be “room” to do the following: we may extend i to an embedding $i': E \rightarrow M'$ in such a way that $f \circ i'$ is strictly concave on the fibers with unique maximum on the zero-section. With a little precision one can prove that the Conley index of $f \circ i'$ is the Thom space of E on the Conley index of $f \circ i$ (relative to the base point). This means that the homologies are not the same because E is not orientable on S .

The above discussion seems to indicate that it is difficult to capture the “correct” orientation on the homology theories when doing finite-dimensional approximation. However, in cotangent bundles there is a canonical choice of Lagrangian foliations which from the stable homotopy theory point of view settles the above problem in more generality than just \mathbb{Z} -coefficients. Indeed, one may heuristically think of this choice of Lagrangian at each point in T^*N as canonically identifying a subset of the finite-dimensional approximations as above, which are only related by an E (also as above) if this E is stably parallizable. So in fact there is a canonical way of constructing a spectrum. However, it turns out that the orientation coming from this choice is not the one usual used in Floer theory, and, indeed, this is why the symplectic cohomology does not agree with the homology of this spectrum. Even more, when one considers Viterbo functoriality from T^*N to DT^*L the Lagrangian foliations induced by the cotangent bundle structures do not agree, and the different choices induce the bundle η in (2).

Appendix E: Nearby Lagrangians are homotopy equivalent by Mohammed Abouzaid

As in the introduction, let N be a closed smooth manifold, L a closed exact Lagrangian in T^*N , N' the universal cover of N , and L' the inverse image of L in T^*N' . In this short appendix, we prove the following result:

Proposition E.1 *If L' has vanishing Maslov class, then so does L .*

As a corollary, we conclude:

Theorem E.2 *If L is a closed exact Lagrangian in T^*N , then L is homotopy equivalent to N .*

Proof By Corollary 1.2, every closed exact Lagrangian satisfies the hypothesis of Proposition E.1, so L has vanishing Maslov index. The result therefore follows from [4], which proves that every closed exact Lagrangian in T^*N with trivial Maslov class is homotopy equivalent to the base. \square

The starting point for the proof of Proposition E.1 is the main result of [3] that the wrapped Fukaya category of T^*N is split-generated by a fibre T_q^*N . In particular, the Floer cohomology of L with T_q^*N does not vanish. To be more precise, [3, Theorem 1.1] shows that the Fukaya category consisting of Lagrangians of vanishing Maslov index is generated by a fibre. As an intermediate step, one proves a split-generation result, and that proof applies, without modification, to the more general case in which the Maslov index may not necessarily vanish. To see this, we use the fact that the proof of [3, Theorem 1.1] itself relies on [2, Theorem 1.1], which asserts that the Fukaya category of a symplectic manifold is split-generated by a Lagrangian L whenever a certain geometric criterion involving the Hochschild homology of the Floer complex $CF^*(L, L)$ is satisfied. The fact that the cotangent fibre in a cotangent bundle satisfies this property is verified in [3, Proposition 1.6]. Once this criterion is satisfied, the annulus degeneration argument described in [2, Section 6] shows that the self-Floer cohomology of any Lagrangian in T^*N vanishes if and only if its Floer cohomology with a fibre vanishes. This result in no way uses integral gradings in Floer cohomology. If L is an exact Lagrangian, its self-Floer cohomology is isomorphic to its ordinary cohomology, so we conclude:

Proposition E.3 *If L is an exact Lagrangian in T^*N , then the Floer cohomology $HF(L, T_q^*N)$ does not vanish.*

In order to derive Proposition E.1, we need to discuss the construction of gradings in Lagrangian Floer cohomology. Starting with a Riemannian metric on any smooth manifold N (not necessarily closed), we obtain a density which assigns to a basis of tangent vectors at a point the square of the volume of the corresponding parallelepiped. Whenever N is oriented, this density is the square of the volume form determined by the metric and the orientation. By complexifying this density, we obtain a *quadratic*

complex volume form η on T^*N ; in local volume-preserving coordinates (q_1, \dots, q_n) with dual coordinates (p_1, \dots, p_n) we have the expression

$$(77) \quad \eta = ((dp_1 + \sqrt{-1}dq_1) \wedge \dots \wedge (dp_n + \sqrt{-1}dq_n))^{\otimes 2},$$

where the symbol $\otimes 2$ means that we assign to an n -tuple of tangent vectors on T^*N the square of the complex number obtained by applying the complex-valued n -form written in coordinates inside the parentheses.

Given any Lagrangian $L \subset T^*N$, we may evaluate η on a basis of tangent vectors to obtain a complex number. The Lagrangian condition implies that this number does not vanish, so we obtain a *complex phase map*

$$(78) \quad \theta_L: L \rightarrow \mathbb{R}/2\pi\mathbb{Z},$$

$$(79) \quad e^{\sqrt{-1}\theta_L(x)} = \frac{\eta(v_1 \wedge \dots \wedge v_n)}{|\eta(v_1 \wedge \dots \wedge v_n)|},$$

where (v_1, \dots, v_n) is an arbitrary basis of tangent vectors for L at x .

Definition E.4 The *Maslov class* of L is the integral first cohomology class represented by θ_L . If L has vanishing Maslov class then a *grading* is a choice $\tilde{\theta}_L$ of an \mathbb{R} -valued lift of θ_L :

$$(80) \quad \begin{array}{ccc} L & \xrightarrow{\quad} & \mathbb{R} \\ & \searrow & \downarrow \\ & & S^1 = \mathbb{R}/2\pi\mathbb{Z} \end{array}$$

Vanishing of the Maslov class is the condition needed in order to be able to assign an integer (up to global shift) to each intersection point between a pair of Lagrangians, and the choice of a grading determines such an integer and hence equips Lagrangian Floer complexes and cohomology groups with \mathbb{Z} -gradings. We shall only need the following case: Assume that we are given a Lagrangian L that intersects T_q^*N transversely as a point x . Since L is Lagrangian, one may choose local coordinates (q_1, \dots, q_n) on N so that the tangent space of L is spanned by products of n lines, each lying in a plane spanned by ∂_{p_i} and ∂_{q_i} , with phase $\alpha_i(x) \in [0, \pi)$.

Note that $\theta_{T_q^*N}$ is independent of the point on T_q^*N , so that we can choose $\tilde{\theta}_{T_q^*N}$ to be identically 0. In this case, the following formula for the Maslov index is given in [19, Section 3.2]:

$$(81) \quad \mu(x) = \frac{1}{\pi} \left(\sum_{i=1}^n \alpha_i(x) - \theta_L(x) \right).$$

With this formula at hand, we can finish the proof of the result announced at the beginning of this note.

Proof of Proposition E.1 Assume that L' has vanishing Maslov class. In [4], a wrapped Fukaya category was assigned to the universal cover of T^*N , which is also the cotangent bundle T^*N' of the universal cover of N . By comparing holomorphic discs in T^*N and T^*N' , we find that the Floer cohomology of L with a fibre T_q^*N is isomorphic to that of L' with any cotangent fibre $T_{q'}^*N'$, and in particular has finite rank. Upon choosing a grading on L' , we conclude from Proposition E.3 that

$$(82) \quad HF^*(L', T_{q'}^*N') \text{ is nonzero, and is supported in finitely many cohomological degrees.}$$

Since all cotangent fibres are Hamiltonian isotopic via an isotopy that preserves the vanishing of the phase, and Floer cohomology together with its grading is invariant under such isotopies, the group $HF^*(L', T_{q'}^*N')$ is moreover independent of q' .

We now shall use (81) to derive a contradiction to the bounded support of the Floer cohomology group $HF^*(L', T_{q'}^*N')$ if L does not also have vanishing Maslov class.

Choose a quadratic volume form on T^*N' that is pulled back from T^*N . In particular, a grading $\tilde{\theta}_{L'}$ on L' fits in a commutative diagram:

$$(83) \quad \begin{array}{ccc} L' & \xrightarrow{\tilde{\theta}_{L'}} & \mathbb{R} \\ \downarrow & & \downarrow \\ L & \xrightarrow{\theta_L} & S^1 \end{array}$$

Given any point $x' \in L'$, lying over a point x , the commutativity of the above diagram implies that for any loop $\gamma \in \pi_1(L, x)$, with image $p_*(\gamma) \in \pi_1(N, p(x))$, we have

$$(84) \quad \theta_L(p_*(\gamma) \cdot x') = \theta_L(x') + \langle [\theta_L], [\gamma] \rangle,$$

where $p_*(\gamma) \cdot x'$ is the image of x' under the action of $p_*(\gamma)$ on T^*N' by deck transformations, $[\gamma]$ is the class of γ in $H_1(N)$, and $[\theta_L]$ is the Maslov class of L .

Let us now assume that L intersects T_q^*N transversely, which may be achieved after a small perturbation. Applying (81), we conclude that:

$$(85) \quad \mu(p_*(\gamma) \cdot x') = \mu(x') - \langle [\theta_L], [\gamma] \rangle$$

Since $\gamma \cdot (L' \cap T_{q'}^*N') = L' \cap T_{\gamma \cdot q'}^*N'$, with a corresponding identification of holomorphic curves for appropriate complex structures, we conclude that

$$(86) \quad HF^*(T_{\gamma \cdot q'}^*N', L') = HF^{* - \langle [\theta_L], [\gamma] \rangle}(T_{q'}^*N', L'),$$

which, because of the finite-dimensionality and invariance of this group, implies that it vanishes, contradicting the output of Proposition E.3. \square

References

- [1] **A Abbondandolo, M Schwarz**, *On the Floer homology of cotangent bundles*, Comm. Pure Appl. Math. 59 (2006) 254–316 MR2190223
- [2] **M Abouzaid**, *A geometric criterion for generating the Fukaya category*, Publ. Math. Inst. Hautes Études Sci. (2010) 191–240 MR2737980
- [3] **M Abouzaid**, *A cotangent fibre generates the Fukaya category*, Adv. Math. 228 (2011) 894–939 MR2822213
- [4] **M Abouzaid**, *Nearby Lagrangians with vanishing Maslov class are homotopy equivalent*, Invent. Math. 189 (2012) 251–313 MR2947545
- [5] **M Chaperon**, *Une idée du type “géodésiques brisées” pour les systèmes hamiltoniens*, C. R. Acad. Sci. Paris Sér. I Math. 298 (1984) 293–296 MR765426
- [6] **R L Cohen, J D S Jones, J Yan**, *The loop homology algebra of spheres and projective spaces*, from: “Categorical decomposition techniques in algebraic topology”, (G Arone, J Hubbuck, R Levi, M Weiss, editors), Progr. Math. 215, Birkhäuser, Basel (2004) 77–92 MR2039760
- [7] **C Conley**, *Isolated invariant sets and the Morse index*, CBMS Regional Conference Series in Mathematics 38, American Mathematical Society (1978) MR511133
- [8] **D B A Epstein**, *The degree of a map*, Proc. London Math. Soc. 16 (1966) 369–383 MR0192475
- [9] **K Fukaya, P Seidel, I Smith**, *Exact Lagrangian submanifolds in simply-connected cotangent bundles*, Invent. Math. 172 (2008) 1–27 MR2385665
- [10] **A Hatcher**, *Algebraic topology*, Cambridge Univ. Press (2002) MR1867354
- [11] **A Hatcher**, *Spectral sequences in algebraic topology* (2004) Available at www.math.cornell.edu/~hatcher/SSAT/SSATpage.html
- [12] **T Kragh**, *The Viterbo transfer as a map of spectra* (2009) arXiv:0712.2533v2
- [13] **F Lalonde, J-C Sikorav**, *Sous-variétés lagrangiennes et lagrangiennes exactes des fibrés cotangents*, Comment. Math. Helv. 66 (1991) 18–33 MR1090163
- [14] **J P May, J Sigurdsson**, *Parametrized homotopy theory*, Mathematical Surveys and Monographs 132, American Mathematical Society (2006) MR2271789
- [15] **J McCleary**, *A user’s guide to spectral sequences*, 2nd edition, Cambridge Studies in Advanced Mathematics 58, Cambridge Univ. Press (2001) MR1793722
- [16] **J Milnor**, *Morse theory*, Annals of Mathematics Studies 51, Princeton Univ. Press (1963) MR0163331

- [17] **D Salamon**, *Connected simple systems and the Conley index of isolated invariant sets*, Trans. Amer. Math. Soc. 291 (1985) 1–41 MR797044
- [18] **D A Salamon, J Weber**, *Floer homology and the heat flow*, Geom. Funct. Anal. 16 (2006) 1050–1138 MR2276534
- [19] **R P Thomas, S-T Yau**, *Special Lagrangians, stable bundles and mean curvature flow*, Comm. Anal. Geom. 10 (2002) 1075–1113 MR1957663
- [20] **C Viterbo**, *Generating functions, symplectic geometry, and applications*, from: “Proceedings of the International Congress of Mathematicians”, (S D Chatterji, editor), volume 1, Birkhäuser, Basel (1995) 537–547 MR1403954
- [21] **C Viterbo**, *Exact Lagrange submanifolds, periodic orbits and the cohomology of free loop spaces*, J. Differential Geom. 47 (1997) 420–468 MR1617648

TK: *Department of Mathematics, Uppsala University*
PO Box 480, SE-751 06 Uppsala, Sweden

MA: *Department of Mathematics, Columbia University*
2990 Broadway, New York, NY 10027, USA

and

Simons Center for Geometry and Physics, State University of New York
Stony Brook, NY 11794-3636, USA

thomas.kragh@math.uu.se, abouzaid@math.columbia.edu

Proposed: Yasha Eliashberg

Received: 6 September 2011

Seconded: Leonid Polterovich, Ralph Cohen

Revised: 19 June 2012

