ISSN 1364-0380 (on line) 1465-3060 (printed)

Geometry & Topology
Volume 8 (2004) 1189–1226
Published: 10 September 2004

A field theory for symplectic fibrations over surfaces

François Lalonde

Department of Mathematics and Statistics, University of Montreal
Montreal H3C 3J7, Quebec, Canada

Email: `lalonde@dms.umontreal.ca`

Abstract

We introduce in this paper a field theory on symplectic manifolds that are fibered over a real surface with interior marked points and cylindrical ends. We assign to each such object a morphism between certain tensor products of quantum and Floer homologies that are canonically attached to the fibration. We prove a composition theorem in the spirit of QFT, and show that this field theory applies naturally to the problem of minimising geodesics in Hofer's geometry. This work can be considered as a natural framework that incorporates both the Piunikhin–Salamon–Schwarz morphisms and the Seidel isomorphism.

AMS Classification numbers Primary: 53D45

Secondary: 53D40, 81T40, 37J50

Keywords: Symplectic fibration, field theory, quantum cohomology, Floer homology, Hofer's geometry, commutator length

Proposed: Leonid Polterovich — Received: 20 September 2003
Seconded: Yasha Eliashberg, Robion Kirby — Revised: 22 August 2004

1 Introduction

We describe in this paper a field theory for Hamiltonian fibrations over oriented surfaces with boundaries and marked points. In this context, a Hamiltonian fibration $(M, \omega) \hookrightarrow P \to \Sigma$ is a symplectic manifold (P, Ω) whose form restricts to a symplectic form on each fiber, endowed with interior marked points on Σ. The restriction of Ω to the 1–skeleton of Σ is assumed to be fiberwise symplectically trivialisable, although such a trivialisation is not part of the data. To each boundary component of Σ there corresponds a geometric version of Floer homology and to each marked point the quantum homology of the fiber at that point. Any partition in two subsets of the boundary components and marked points gives rise to a morphism between the corresponding tensor product of homologies. The gluing of two such surfaces (using gluing on the boundary components and index zero surgeries on a pair of marked points) and the compatible gluing of the fibrations give rise to the composition of morphisms. This setting is a geometric generalisation of the Piunikhin–Salamon–Schwartz [17] and the Seidel homomorphisms [22].

We will show below that such a generalisation is natural in the study of Hofer's geometry: we will use our field theory to give a new proof that a symplectic isotopy $\phi_t \in \mathrm{Diff}_{\mathrm{Ham}}(M)$ generated by a Hamiltonian H_t is length minimising in the Hofer metric in its homotopy class with fixed endpoints if two special classes corresponding to the fixed minimum and maximum of the generating Hamiltonian are essential in the Floer homology of H_t. The proof is geometric and works for all weakly monotone manifolds (actually, it very likely holds for all symplectic manifolds if one consider the virtual moduli space of pseudo-holomorphic curves, but we will restrict here to the case where (M, ω) is weakly monotone). See Corollary 3.3 for new results on length minimising geodesics in Hofer's geometry that this field theory yields. Here is a simple example of application of Corollary 3.3:

Example Given a symplectic manifold with minimal Chern number not in the range $[1, n]$, $\mathbb{CP}^n$ for instance, let $f\colon M \to \mathbb{R}$ be a Morse function. Let P be a global maximum of f and Q a global minimum of f such that the linearised flows have no nontrivial closed orbit in time ≤ 1. Let $\varepsilon > 0$ be sufficiently small so that the graph of the diffeomorphism ϕ_t^f induced by f be transversal to the diagonal in $M \times M$ for all $t \in (0, \varepsilon]$. Consider any Hamiltonian $H_{t\in[0,1]}$ such that:

(1) on some open neighbourhoods of P and Q, $H_t = f$ for all $t \in [0, 1]$

(2) $H_t = f$ for $t \leq \varepsilon$,

(3) the graph of ϕ_t^H remains transversal to the diagonal in $M \times M$ for all $t \in [\varepsilon, 1]$, and

(4) $H_t(x) \in [f(Q), f(P)]$ for all $x \in M$ and $t \in [0,1]$.

Then the path induced by $H_{t\in[0,1]}$ is length minimising rel endpoints in its homotopy class, in the Hofer metric. See the proof of this claim after Corollary 3.3 in Section 3.

This paper is a natural extension of the ideas presented in Lalonde–McDuff–Polterovich [9] and in Entov [2]. See the remark at the end of Section 2 for a generalisation of this setting to a category of symplectic fibrations slightly larger than the Hamiltonian ones.

We end this paper with a somewhat independent section treating a notion of "total norm" on Hamiltonian diffeomorphisms, defined as the sequence of its genus–g norms. The results and conjecture of that section turn out to be closely related to Entov's paper [3] on the relation of K–area and the commutator length of Hamiltonian diffeomorphisms.

Acknowledgements I wish to thank the referee, whose questions helped clarify some aspects of the first version. This research was partially supported by a CRC grant, NSERC grant OGP 0092913, and FCAR grant ER–1199.

2 A field theory for fibrations

2.1 Definition of the objects

Here is the definition of the basic objects. Let $\Sigma = \Sigma_{g,k,\ell}$ be a real compact oriented surface of genus g with interior marked points $p_1, \dots, p_k$ and oriented boundary components $S_1, \dots, S_\ell$ (here g is by definition the genus of the closed surface obtained by capping with discs all boundary components). Let $(M,\omega) \hookrightarrow P \xrightarrow{\pi} \Sigma$ be a Hamiltonian fibration with fiber *any* compact symplectic manifold (we will restrict ourselves to simpler manifolds later when we will define Φ_σ). By the characterisation of such fibrations in [12, 8], we may assume that each fiber M_b is equipped with a symplectic structure ω_b that admits an extension to a symplectic form Ω on P (note that the restriction of Ω to each M_b is equal to ω_b as forms, not only as cohomology classes). We may further assume that the restriction of P to the 1–skeleton Σ_1 of Σ (including its boundary components) is fiberwise symplectically trivialisable. We will denote by $\rho\colon \pi^{-1}(\Sigma_1) \to \Sigma_1 \times M$ such a trivialisation. Thus $\rho_*(\Omega)$ restricts to ω

on each fiber $\{b\} \times M$. Such a choice of trivialisation is not unique in general at the homotopy level. The restriction of Ω to $\pi^{-1}(S_j)$ has a one-dimensional kernel transversal to the fibers: the monodromy of this foliation followed in the direction of the boundary component corresponding to its orientation is a Hamiltonian diffeomorphism $\phi_j\colon M_{q_j} \to M_{q_j}$.

Definition 2.1 A *marked Hamiltonian symplectic fibration* P over Σ is given by the following data:

(1) a fibration $(M,\omega) \hookrightarrow (P,\Omega) \to \Sigma$ over a real compact oriented surface with ℓ oriented boundary components, each one with a base point $q_j \in S_j$, together with k interior marked points $p_1, \dots, p_k \in \Sigma$;

(2) symplectic identifications $\eta_i\colon (M_{p_i}, \omega_{p_i}) \to (M,\omega)$ for each $1 \leq i \leq k$ and $\eta_j\colon (M_{q_j}, \omega_{q_j}) \to (M,\omega)$ for each $1 \leq j \leq \ell$;

(3) a choice, for each j, of an element ξ_j of the universal cover $\widetilde{\mathrm{Diff}}_{\mathrm{Ham}}(M,\omega)$ which projects to $\psi_j =_{def} \eta_j \circ \phi_j \circ \eta_j^{-1} \in \mathrm{Diff}_{\mathrm{Ham}}(M,\omega)$;

(4) a partition in two subsets $\mathcal{A}_+, \mathcal{A}_-$ of the set $\{p_1, \dots, p_k, S_1, \dots, S_\ell\}$.

Here the orientation $\mathcal{O}_j$ of the "ingoing" S_js, ie those that belong to $\mathcal{A}_+$, is such that $\mathcal{O}_j$ followed by the inward normal orientation is equal to the orientation of the surface. The orientation $\mathcal{O}_j$ of the "outgoing" S_js, ie those that belong to $\mathcal{A}_-$, is such that $\mathcal{O}_j$ followed by the outward normal orientation is equal to the orientation of the surface.

Note that, although the identifications ηs are part of the data, the choice of a trivialisation ρ is not. In the sequel, any choice of a trivialisation ρ will be compatible with the η_js, $1 \leq j \leq \ell$.

The main goal of this section is to assign to each end of Σ a Floer homology generated by the one-turn closed leaves of the characteristic foliation of $\Omega|_{\pi^{-1}(S_j)}$ and to each marked Hamiltonian fibration $P \to \Sigma$ a morphism

$$\Phi_\sigma\colon \bigotimes_{p_i \in \mathcal{A}_+} QH_*(M_{p_i}) \otimes \bigotimes_{S_j \in \mathcal{A}_+} FH_*(S_j) \longrightarrow \bigotimes_{p_i \in \mathcal{A}_-} QH_*(M_{p_i}) \otimes \bigotimes_{S_j \in \mathcal{A}_-} FH_*(S_j)$$

depending on the choice of a homology class of sections σ, whose domain is the tensor product over the ingoing data of the Floer homologies (assigned to the ends) and the quantum homologies (assigned to the fibers over the interior marked points) and whose codomain is the corresponding tensor product over the outgoing data. This morphism is a generalisation to fibrations of the morphisms introduced in Piunikhin–Salamon–Schwarz: it is obtained by counting sections of P that obey a certain Cauchy–Riemann equation. However, in the counting of pseudo-holomorphic curves, we do not fix the conformal structure of the base. Note however that all the interesting cases that we have in mind involve marked surfaces Σ with cylindrical ends that admit, up to the data preserving diffeomorphism group, only one conformal structure; therefore, the variation of the conformal structure is not an issue in the applications that we have in mind.

Here are some special cases of this construction.

Example 1 When the base is S^2 with two marked points, the morphism:

$$\Phi_\sigma\colon\ QH_*(M_{p_1}) \to QH_*(M_{p_2})$$

is the Seidel homomorphism (see [22, 9]) that depends on the homotopy class of the loop of Hamiltonian diffeomorphisms used to define the fibration $P \to S^2$.

Example 2 When the base is the 2–disc with one interior marked point and the fibration is trivial away from the end, this is as in [17] the usual isomorphism between quantum and Floer homologies.

Example 3 When the base is the 2–sphere with three marked points and P is the trivial fibration (or three ends and the fibration is trivial away from the cylindrical ends), this is the quantum product (or Floer product).

It is not difficult to see that, when the base is the 2–sphere with an arbitrary number of ends and marked points in the ingoing data set and only one element (either a point or a boundary component) in the outgoing data set, our morphism is always a composition of the morphisms of the last three examples, in the sense of our Composition Theorem (see below). However, our theory makes sense in a wider context, where one may allow monodromies over interior loops of Σ to be non symplectically isotopic to the identity – see the last remark of this section.

The main result of this section is a Composition formula similar to Theorem 3.7 of [17]. It is useful even in the simplest case of the gluing of two fibrations $(M,\omega) \hookrightarrow P_i \to D^2$ (for $i = 1, 2$) over a disc with one interior marked point. We assume here that the fibrations P_1, P_2 have the same Hamiltonian monodromy round their end, and can therefore be glued together along the monodromies to yield a new (not necessarily trivial) fibration $P_1\#P_2$. The Composition Theorem then states that the Seidel homomorphism

$$\Phi_{\sigma=\sigma_1\#\sigma_2, P=P_1\#P_2}\colon\ QH_*(M)\to QH_*(M)$$

corresponding to $P_1\#P_2$ factorises through the morphisms $\Phi_{\sigma_1,P_1}\colon\ QH(M)\to FH(\partial P_1)$ and $\Phi_{\sigma_2,P_2}\colon\ FH(\partial P_2)\to QH(M)$. Thus, if one knows that the Seidel homomorphism is an isomorphism, one concludes that both Φ_{σ_1,P_1} and Φ_{σ_2,P_2} are isomorphisms too, and conversely. This will be used in our application to Hofer's geometry in Section 3.

Let us return to the construction of the main objects. Our first goal is to define a homology whose generators are the closed leaves of the characteristic foliation of the restriction of Ω to a given boundary component, together with filling discs to be defined below, that will be canonically identified with the Floer homology in a way that is independent of the chosen trivialisation ρ.

Let $H_2^S(M,\mathbb{Z})$ be the spherical part of the second homology group of M, ie the image of the Hurewicz homomorphism $\pi_2(M)\to H_2(M,\mathbb{Z})$, and let Λ be the usual rational Novikov ring of the group $\mathcal{H}=H_2^S(M,\mathbb{Z})/\!\sim$ where $B\sim B'$ if $\omega(B-B')=c(B-B')=0$. Thus the elements of Λ have the form

$$\sum_{B\in\mathcal{H}}\lambda_B e^B$$

where for each κ there are only finitely many nonzero $\lambda_B\in\mathbb{Q}$ with $\omega(B)<\kappa$.

Let us first quickly recall that the Floer homology $HF_*(H_t)$ can be considered as the Novikov homology of the classical action functional

$$\mathcal{A}_H\colon\ \widetilde{\mathcal{L}}\to\mathbb{R}$$

defined by $\mathcal{A}_H(x(t),v_x)=\int_{v_x}\omega-\int_0^1 H_t(x(t))dt$. Here $\widetilde{\mathcal{L}}$ is a covering of the loop space $\mathcal{L}$ of all contractible loops in M defined in the following way. Fix a constant loop $x_*\in\mathcal{L}$. Then the covering $\widetilde{\mathcal{L}}$ is defined by requiring that two paths from x_* to $x(t)$ are equivalent if the 2–sphere S obtained by gluing the corresponding discs has $\omega(S)=c(S)=0$. Thus the covering group of $\widetilde{\mathcal{L}}$ is $\mathcal{H}$, and the Floer homology is a Λ–module in a natural way.

Now let us define the geometric Floer homology attached to an end S_j. Note that any fiberwise symplectic trivialisation $\rho_j\colon\ \pi^{-1}(S_j)\to S_j\times M$ pushes the form Ω to a form Ω' on $S_j\times M$ whose characteristic flow round the boundary projects on the fiber (M,ω) to a Hamiltonian isotopy $\psi_{t\in[0,1]}$. Pick the trivialisation ρ_j so that this isotopy belongs to the class ξ_j. This is always possible by composing the trivialisation, if needed, with an appropriate loop of Hamiltonian diffeomorphisms. Let P_j be the canonical Hamiltonian fibration over the

2–disc whose monodromy round the boundary lies in the class ξ_j: this is given as the part $P^-(H)$ under the graph of the Hamiltonian $H_{t\in[0,1]}$ that generates $\psi_{t\in[0,1]}$ – see the precise definition in Section 3.1. Glue P_j to P via the map ρ_j, and denote by $\bar{P}$ the space obtained from P by gluing to P all the P_js, $1 \leq j \leq \ell$. We will call $\bar{P}$ the *closure of* P. It is endowed with a natural ruled symplectic form $\bar{\Omega}$ that extends Ω; it extends as well the symplectic forms on the fibers of $\bar{P}$. We will denote by the same symbol P_j its image in $\bar{P}$.

Definition 2.2 The *Floer complex* corresponding to S_j has generators $(x(t),v)$ where x is a parametrisation of a closed leaf of the characteristic foliation (induced by a fixed parametrisation of the base S_j), and where v is a 2–disc in $P_j \subset \bar{P}$ with boundary equal to x. Via the trivialisation ρ_j, these objects are in one-to-one correspondence with the generators of the usual Floer complex of H_t. Define the boundary operator as the pull-back through ρ_j of the usual Floer boundary operator.

Lemma 2.3 *The Floer complex corresponding to S_j depends only on the class ξ_j.*

Proof If ρ_j' is another trivialisation in class ξ_j, the two corresponding Hamiltonian isotopies ψ_j and ψ_j' are homotopic with fixed endpoints. There is however a delicate point here: the identification of these two isotopies (and therefore the identification of the corresponding cylinders P_j, P_j', capping discs v, v' and boundary operators) depends on a choice of a homotopy between the two paths ψ_j and ψ_j'. But two different homotopies differ by an element in $\pi_2(\mathrm{Diff}_{\mathrm{Ham}}(M))$ and it is shown in [8], Proposition 5.4, that such elements act trivially on the homology of M. Thus all these identifications act in the same way on capping discs. It is then easy to see that they preserve the boundary operator too. Note that this is equivalent to saying that the Seidel morphism induced by a contractible loop ϕ_t of Hamiltonian diffeomorphisms admits a unique lift to the universal cover of the loop space (independent of the choice of the homotopy between the loop ϕ_t and the trivial loop). □

We will denote by $FC(S_j)$ the Floer complex and by $FH(S_j)$ its homology (or by $FC(S_j,\xi_j), FH(S_j,\xi_j)$ if we wish to emphasise the dependence on ξ_j). It is a Λ–module, and there is a canonical Λ–module isomorphism between $FH(S_j,\xi_j)$ and the Floer homology of M. Recall that this complex is graded by the Conley–Zehnder index normalised so that it corresponds to the Morse index in the case of a C^1–small autonomous Hamiltonian.

One may define the quantum complex associated with a Smale–generic pair (κ, g_M) consisting of a Morse function $\kappa\colon M \to \mathbb{R}$ and a metric g_M on M: this is the usual Morse complex of (κ, g_M) tensored with Λ, ie $C_*(\kappa, g_M)\otimes\Lambda$. When there is no confusion, we will often omit the data κ, g_M and denote $C_*(\kappa, g_M)$ by $C_*(M)$. The quantum complex is defined by $QC_*(M) = C_*(M) \otimes \Lambda$. The quantum homology of M is defined as $QH_*(M) = H_*(M) \otimes \Lambda$, $\mathbb{Z}$–graded with $\deg(a \otimes e^B) = \deg(a) - 2c(B)$. Associate to each marked point p_i the quantum homology of the fiber at that point, $QH(M_{p_i})$. Each η_i identifies the quantum homology of the fiber at p_i with the quantum homology of M.

Lemma 2.4 *Every Hamiltonian fibration over a closed oriented real surface admits sections. Actually, the homology classes of sections form an affine space over $H_2^S(M)$.*

Proof This has been proved in [9] for Hamiltonian fibrations over the 2–sphere. Let $P \to \Sigma$ be a Hamiltonian fibration over an arbitrary Riemann surface with boundary, given with a symplectic fiberwise trivialisation over the 1–skeleton Σ_1 of Σ. Thus finding a section of P which is constant over Σ_1 is equivalent to finding sections of the quotient bundle over Σ/Σ_1. But the latter is a bouquet of 2–spheres, and the result therefore follows from the case proved in [9]. $\square$

2.2 Definition of the morphisms

To simplify the exposition and avoid dealing with virtual moduli spaces, assume from now on that (M, ω) is weakly monotone. Recall that this means that for every spherical homology class $B \in H_2(M)$

$$\omega(B) > 0, \;\; c_1(B) \geq 3 - n \quad \Longrightarrow \quad c_1(B) \geq 0.$$

This condition is satisfied if either $\dim M \leq 6$ or M is *semi-monotone*, ie there is a constant $\mu \geq 0$ such that, for all spherical homology classes $B \in H_2(M)$,

$$c_1(B) = \mu\omega(B).$$

Weak monotonicity is the condition under which the ordinary (non virtual) theory of J–holomorphic curves behaves well.

Define an equivalence relation on the set of homology classes of sections of $\bar{P}$ by identifying two such classes if their values under c_{vert} and $\bar{\Omega}$ are equal. Here c_{vert} is the first Chern class of the bundle $V \to \bar{P}$ whose fiber over $p \in \bar{P}$ is the tangent space $T_p(\pi^{-1}(\pi(p)))$ to the fiber of the M–bundle passing through

p. Note that this definition is independent of the extension $\bar{\Omega}$ of ω. Given an equivalence class of sections σ of $\bar{P}$, whose existence is guaranteed by Lemma 2.4, we are going to define a morphism of chain complexes

$$\Phi_\sigma\colon \bigotimes_{p_i\in\mathcal{A}_+} QC_*(M_{p_i})\otimes \bigotimes_{S_j\in\mathcal{A}_+} FC_*(S_j)\longrightarrow \bigotimes_{p_i\in\mathcal{A}_-} QC_*(M_{p_i})\otimes \bigotimes_{S_j\in\mathcal{A}_-} FC_*(S_j).$$

that depends on the auxiliary data given by the choices of the generic Smale–Morse pairs (κ_i,g_i) on the fibers M_{p_i}, but will turn out to be independent of these auxiliary data at the homology level. Note that the domain and the codomain are both Λ–modules in an obvious way. Here the total degree of the domain or codomain is the sum of the degrees of each factor, and we denote by d the degree of Φ_σ, ie the integer such that $\deg(\Phi_\sigma(w))=\deg(w)+d$ (see the computation of d below in Lemma 2.8). To simplify notations, assume that the set $\mathcal{A}_+$ contains the indices from $i=1$ to $i=i_+$ and from $j=1$ to $j=j_+$. Any chain in the domain of Φ_σ decomposes as a sum of elements of the form: $(a_1e^{\alpha_1}\otimes\ldots\otimes a_{i_+}e^{\alpha_{i_+}})\otimes((x_1,v_1)\otimes\ldots\otimes(x_{j_+},v_{j_+}))$ where $a_1,\ldots,a_{i_+}$ are unstable manifolds of critical points of κ_i. The image of such an element by Φ_σ has the form

$$\sum m\left(a_1,\ldots,a_{i_+},b_{i_++1},\ldots,b_k,(x_1,v_1),\ldots,(x_\ell,v_\ell),\sigma,\alpha_i\right)$$
$$\left(b_{i_++1}e^{\alpha_{i_++1}}\otimes\ldots\otimes b_ke^{\alpha_k}\right)\otimes\left((x_{j_++1},v_{j_++1})\otimes\ldots\otimes(x_\ell,v_\ell)\right)$$

where the summation runs through all generators $(x_{j_++1},v_{j_++1}),\ldots,(x_\ell,v_\ell)$, all sequences of 2–spheres $\alpha_{i_++1},\ldots,\alpha_k$ and sequences of stable manifolds b_is associated with the κ_is. The number m is defined to be equal to the cardinality, counted with signs, of the moduli space

$$\mathcal{M}_{g,k,\ell}(a_1,\ldots,a_{i_+},b_{i_++1},\ldots,b_k,(x_1,v_1),\ldots,(x_\ell,v_\ell),\tau)$$

where

$$\tau=\sigma-\sum_{1\le i\le i_+}\alpha_i+\sum_{i_++1\le i\le k}\alpha_i.$$

which is defined, roughly, by counting sections u of the bundle that are holomorphic away from the ends, satisfy a Floer-type equation on each end, converge on each end to the loops $x_1,\ldots,x_\ell$, and meet the a and b chains over the marked points of the surface. For this, one needs in general to deal with a moduli space of conformal structures on the base; however, one cannot quotient out the pairs (j,u) made of a conformal structure and of a section by the diffeomorphism group of Σ that preserves the ends and the marked points, since the action on the second factor does not yield a section. For this reason, we will introduce the following definition:

Definition 2.5 Let Σ be a real oriented surface of genus g with ℓ ends and k interior marked points. Let us denote by $\mathcal{J}_{g,\ell}$ the space of conformal structures on the interior of Σ that are standard on each end (ie biholomorphic to $\mathbb{R}/\mathbb{Z} \times [0,\infty)$) and by $\mathcal{D}_{g,k,\ell}$ the group of compactly supported orientation preserving diffeomorphisms of Σ that fix pointwise the ends and the marked points. Note that the quotient $\mathcal{J}_{g,\ell}/\mathcal{D}_{g,k,\ell}$ has a natural smooth (orbifold) structure. We say that Σ it is *normalisable* if the space of smooth sections of the quotient map $\mathcal{J}_{g,\ell} \to \mathcal{J}_{g,\ell}/\mathcal{D}_{g,k,\ell}$ is non-empty and connected. In such a case, we say that it is *normalised* if the choice of a section has been made.

Obviously, when the quotient space $\mathcal{J}_{g,\ell}/\mathcal{D}_{g,k,\ell}$ reduces to a point, which is the case when Σ has genus 0 and the total number of ends and marked points is less than or equal to 3, the marked surface is normalisable: there is then only one class of conformal structures to choose. In genus zero, with four marked points, this quotient is the cross-ratio of the four points, and a normalisation consists in sending say $0, 1, \infty$ to the first three points and keeping the fourth one free. In genus g with zero marked points, a normalisation is the choice of a particular structure for each class of biholomorphic ones (ie a section of the Teichmuller space quotiented out by the appropriate mapping class group). We will now assume that Σ is normalised and will denote by σ_{norm} the chosen section. It is required that any complex structure in the image of σ_{norm} be standard on the cylindrical ends where the bump functions are defined; it is allowed to vary everywhere else.

Here is a detailed description of this moduli space $\mathcal{M}_{g,k,\ell}$. In P, there is a inner collar neighbourhood V_j of the hypersurface $W_j = \pi^{-1}(S_j)$ of the form $\pi^{-1}(U_j)$ where $U_j = S^1 \times [0,\delta]$. On this inner collar neighbourhood, Ω has the form $\Omega_{W_j} + d(yd\theta) = \Omega_{W_j} + dy \wedge d\theta$ where y is the coordinate in $[0,\delta]$. (Here δ corresponds to the boundary S_j and 0 is in the interior of the surface.) Therefore, the characteristic foliation of the hypersurfaces $W_j(y)$ does not depend on y: W_j is *stable* and actually flat. Take a generic family of almost complex structures J_b on each fiber M_b of P, compatible with ω_b, and let $\mathcal{H}$ be the symplectic connection which at $p \in P$ is the Ω–orthogonal complement of the tangent space to the fiber at p. Consider the function $h_j\colon [0,\delta] \to \mathbb{R}^+$ that vanishes near 0, is non-decreasing and reaches the slope 1 at $y = \delta$, and let us denote by f_j its pull-back by the projection $S^1 \times [0,\delta] \to [0,\delta]$. Thus f_j has a periodic orbit of period 1 at δ with respect to the symplectic form $dy \wedge d\theta$. Endowing the jth end $S^1 \times [0,\delta]$ with the form $dy \wedge d\theta$ and the standard conformal structure $J_{st}\partial_y = \partial_\theta$, solves the equation

$$\partial_s v + J_{st}\partial_t v = -\nabla f_j$$

for a map $v\colon S^1\times[0,\infty)\to S^1\times[0,\delta)$ with initial condition $v(t,0)=(t,0)$ which converges to the periodic orbit of f_j at $y=\delta$. Push forward the conformal structure of the domain of v to its codomain: this gives a conformal structure J_j on the jth end, conformally equivalent to the standard structure on the semi-infinite cylinder. Now each $J_\Sigma\in\sigma_{norm}(\mathcal{J}_{g,\ell}/\mathcal{D}_{g,k,\ell})$ is required to coincide with J_j on each end. For any such structure, consider the sections with finite energy $u\colon\Sigma\to P$ which are solutions of the equation:

$$AC_{J_\Sigma,J_b}(\pi_{\mathcal{H}}\circ du)(b)=0 \tag{$*$}$$

for all $b\in\Sigma$, where $\pi_{\mathcal{H}}$ is the projection on the tangent space to the fiber defined by the connection $\mathcal{H}$ and AC is the anti-complex part of the derivative. Impose the following boundary conditions:

(1) $u(p_i)$ meets the chain $\iota(a_i)$ for $i\le i_+$ and $\iota(b_i)$ for $i>i_+$ (where ι denotes the map induced by inclusion of the fiber to P);

(2) $u(s,t)\to x_j(t)$ as $s\to\delta$ on each end, and

(3) the section class that u represents in $H_2(\bar{P};\mathbb{Z})$, once capped in the obvious way at each end by the v_js $1\le j\le\ell$, is equal to τ.

Denote by

$$\mathcal{M}^{\Gamma}_{g,k,\ell,\sigma_{norm},\{J_b\},F}(a_1,\dots,a_{i_+},b_{i_++1},\dots,b_k,(x_1,v_1),\dots,(x_\ell,v_\ell),\tau)$$

the space of these solutions (Γ stands for sections), in which J_Σ is allowed to vary in the image of the chosen section of $\mathcal{J}_{g,\ell}\to\mathcal{J}_{g,\ell}/\mathcal{D}_{g,k,\ell}$. Here F is the sum of the $F_j\colon P\to\mathbb{R}$ defined to be zero away from the ends and given, on the above collar neighbourhood V_j, by $f_j\circ\pi$ where π is the projection of the fibration.

Thus this moduli space is the set of all pairs (u,J_Σ) where J_Σ belongs to the image of the generic section σ_{norm} and u is a section of the fibration $P\to\Sigma$ satisfying $(*)$ and the above three boundary conditions. Because the total space of the manifold is assumed to be weakly monotone, the standard transversality and orientation results can be used to show that such a moduli space is an oriented manifold for generic data $\sigma_{norm},\{J_b\},F$ – see for instance McDuff–Salamon [13] and Salamon [20] for closed curves and Floer flow lines, and Gatien–Lalonde [4] for a rigorous treatment of the transversality issues when the conformal structure of the source Σ is allowed to vary in a simple way.

Remark Observe that a family of almost-complex structures on the fibers $\{J_b\}$ and a complex structure J_Σ on Σ uniquely determine an almost-complex structure J on P, compatible with Ω, that restricts to $\{J_b\}$ on the fibers, is such that the projection π is (J_Σ,J)–holomorphic, and has the horizontal distribution $\mathcal{H}$ (ie the distribution of planes Ω–orthogonal to the fibers) as invariant subspaces. It is then obvious that each above solution corresponds to a section u which is a solution of the equation

$$\bar{\partial}_{J_\Sigma, J} u = -\nabla F \tag{$**$}$$

where ∇ is the Riemannian gradient corresponding to the metric $\Omega(\cdot, J\cdot)$ and F is defined as above. Note that we will not have to consider more general Js than the ones corresponding to a pair $(J_\Sigma, \{J_b\})$ in this way.

Now, define the number $m(a_1, \dots, a_{i_+}, b_{i_+ +1}, \dots, b_k, (x_1, v_1), \dots, (x_\ell, v_\ell), \tau)$ as zero if the above moduli space has dimension different from 0, and by the number of elements of that moduli space (counted with sign) if its dimension is zero.

Proposition 2.6 *The homomorphism Φ_σ is a morphism of chain complexes.*

Proof Since we work in a weakly monotone manifold, this proposition can be proved by standard techniques. Note first that the space P is a compact symplectic manifold with flat boundary. Thus the Gromov–Floer compactness theorem applies as in the case of the standard Floer theory.

Each element of the moduli space

$$\mathcal{M}^{\Gamma}_{g,k,\ell,\sigma_{norm},\{J_b\},F}(a_1, \dots, a_{i_+}, b_{i_+ +1}, \dots, b_k, (x_1, v_1), \dots, (x_\ell, v_\ell), \tau)$$

consists of three pieces:

(1) the ingoing flow lines from the a_is $(i \le i_+)$ to u,

(2) u itself, and

(3) the outgoing flow lines from u to the b_is $(i > i_+)$.

Now consider the same moduli space where exactly one of the terms in the tensor product of the elements in, say, the target complex is replaced by a term appearing in its differential (this amounts, by the Leibniz rule, to replacing the data in the target space

$$b_{i_+ +1}, \dots, b_k, (x_{j_+ +1}, v_{j_+ +1}), \dots, (x_\ell, v_\ell)$$

by a non-vanishing element appearing in its differential). The index is then increased by one, hence this moduli space has real dimension 1 and has boundaries that are due either to the boundary of the Morse chain complex associated to the quantum complexes, or to the boundary of the moduli space of holomorphic sections. The former consists of broken flow lines of the Morse–Smale complexes (κ_i, g_i) where the breaking point is either in a fiber corresponding to the ingoing data, or in a fiber corresponding to the outgoing data (ie they correspond to the boundary of the stable or unstable chains); and the latter consists of broken flow lines of the Floer equations over one of the ingoing or outgoing ends. By standard gluing techniques, either of these two events can be realised as the boundary of a moduli space of the above form. Thus each element of

$$\Phi_\sigma \partial \left((a_1 e^{\alpha_1} \otimes \ldots \otimes a_{i_+} e^{\alpha_{i_+}}) \otimes ((x_1, v_1) \otimes \ldots \otimes (x_{j_+}, v_{j_+})) \right)$$

can be paired, after cancellations, with an element in

$$\partial \Phi_\sigma \left((a_1 e^{\alpha_1} \otimes \ldots \otimes a_{i_+} e^{\alpha_{i_+}}) \otimes ((x_1, v_1) \otimes \ldots \otimes (x_{j_+}, v_{j_+})) \right).$$

By the compactness theorem, there are for each given energy level κ only finitely many homology classes D in P with $\Omega(D - \sigma) \leq \kappa$ that are represented by J–holomorphic curves in P with the fixed boundary conditions. Thus Φ_σ satisfies the finiteness condition for elements of the codomain. Finally, because the sum of a section with a sphere in the fiber gives a homology class which is obviously independent of the choice of the fiber where the sphere is taken (a fiber in the domain of Φ_σ or in its codomain), the map Φ_σ is Λ–linear. □

Thus the morphism Φ_σ descends to a homomorphism at the homology level. We omit the proof of the following proposition that follows by standard cobordism arguments as in Floer's theorem of invariance under change of auxiliary data:

Proposition 2.7 *The map*

$$\Phi_\sigma \colon \bigotimes_{p_i \in \mathcal{A}_+} QH_*(M_{p_i}) \otimes \bigotimes_{S_j \in \mathcal{A}_+} FH_*(S_j) \longrightarrow \bigotimes_{p_i \in \mathcal{A}_-} QH_*(M_{p_i}) \otimes \bigotimes_{S_j \in \mathcal{A}_-} FH_*(S_j).$$

is well-defined: it is independent of all auxiliary choices that have been made, ie it is independent of the choices of the fiberwise almost complex structures $\{J_b\}$ on M_b, of the bump function F and of the (generic) choice of the normalisation.

Lemma 2.8 *The degree d of Φ_σ is:*

(0) $d = 2c_{vert}(\sigma)$ *in the genus* 0 *case;*

(1) $d = 2c_{vert}(\sigma) - 2n + 2$ *in the genus* 1 *case;*

(2) $d = 2c_{vert}(\sigma) + g(6 - 2n) - 6$ *in the genus* $g \geq 2$ *case.*

Proof The index formula for closed parametrised J–holomorphic curves u of genus g in an almost-complex manifold V of dimension $2n$ is:

$$\text{index} = 2(c_1(u) + n(1 - g)) + \dim \mathcal{T}_g$$

where the dimensions are real, $\mathcal{T}_g$ is the Teichmuller space and c_1 denotes the first Chern class of the tangent bundle of V. Thus the dimension of the space of J–holomorphic sections of $\bar{P}$ is given by the same formula, where c_1 is replaced by c_{vert}. Hence the dimension of the moduli space

$$\mathcal{M}^{\Gamma}_{g,k,\ell,\sigma_{norm},\{J_b\},F}(a_1,\dots,a_{i_+},b_{i_++1},\dots,b_k,(x_1,v_1),\dots,(x_\ell,v_\ell),\tau)$$

is equal, for genus 0, to

$$\sum_{1\le j\le j_+}\mu(x_j,v_j) \;-\; \sum_{j_++1\le j\le \ell}\mu(x_j,v_j) \;+\; 2c_{vert}(\tau) \;+\; 2n$$
$$-\sum_{1\le i\le i_+}\operatorname{codim} a_i \;-\; \sum_{i_++1\le i\le k}\operatorname{codim} b_i.$$

Thus the dimension vanishes when

$$\sum_{i_++1\le i\le k}\operatorname{codim} b_i = \sum_{1\le j\le j_+}\mu(x_j,v_j) \;-\; \sum_{j_++1\le j\le \ell}\mu(x_j,v_j)$$
$$+\; 2c_{vert}(\tau) \;+\; 2n - \sum_{1\le i\le i_+}\operatorname{codim} a_i,$$

that is to say when

$$\sum_{i_++1\le i\le k}\dim a_i = \sum_{1\le j\le j_+}\mu(x_j,v_j) \;-\; \sum_{j_++1\le j\le \ell}\mu(x_j,v_j)$$
$$+\; 2c_{vert}(\tau) \;+\; 2n \;-\; \sum_{1\le i\le i_+}\operatorname{codim} a_i. \qquad (***)$$

But

$$\begin{aligned}
&\deg\Phi_\sigma = \deg\Psi_\sigma(a_1e^{\alpha_1}\otimes\ldots\otimes a_{i_+}e^{\alpha_{i_+}}\otimes(x_1,v_1)\otimes\ldots\otimes(x_{j_+},v_{j_+}))\\
&-\deg(a_1e^{\alpha_1}\otimes\ldots\otimes a_{i_+}e^{\alpha_{i_+}}\otimes(x_1,v_1)\otimes\ldots\otimes(x_{j_+},v_{j_+}))\\
&=(\textstyle\sum_{i_++1\le i\le k}\dim a_i-2\sum_{i_++1\le i\le k}c_{vert}(\alpha_i)+\sum_{j_++1\le j\le\ell}\mu(x_j,v_j))\\
&-(\textstyle\sum_{1\le i\le i_+}\dim a_i-2\sum_{1\le i\le i_+}c_{vert}(\alpha_i)+\sum_{1\le j\le j_+}\mu(x_j,v_j))\\
&=(\textstyle\sum_{1\le j\le j_+}\mu(x_j,v_j)-\sum_{j_++1\le j\le\ell}\mu(x_j,v_j)+2c_{vert}(\tau)+2n\\
&-\textstyle\sum_{1\le i\le i_+}\operatorname{codim} a_i-2\sum_{i_++1\le i\le k}c_{vert}(\alpha_i)+\sum_{j_++1\le j\le\ell}\mu(x_j,v_j))\\
&-(\textstyle\sum_{1\le i\le i_+}\dim a_i-2\sum_{1\le i\le i_+}c_{vert}(\alpha_i)+\sum_{1\le j\le j_+}\mu(x_j,v_j))\\
&=(\textstyle\sum_{1\le j\le j_+}\mu(x_j,v_j)-\sum_{j_++1\le j\le\ell}\mu(x_j,v_j)+2c_{vert}(\sigma)+2n\\
&-\textstyle\sum_{1\le i\le i_+}\operatorname{codim} a_i+\sum_{j_++1\le j\le\ell}\mu(x_j,v_j))-\sum_{1\le i\le i_+}\dim a_i\\
&-\textstyle\sum_{1\le j\le j_+}\mu(x_j,v_j))\\
&=2c_{vert}(\sigma)+2n-\textstyle\sum_{1\le i\le i_+}\operatorname{codim} a_i-\sum_{1\le i\le i_+}\dim a_i=2c_{vert}(\sigma).
\end{aligned}$$

where the first equality relies on the definition of the degree of a homomorphism, the second on the definition of the total degree of an element, the third on $(***)$, the fourth is obtained by replacing τ by its value $\sigma - \sum_{1\leq i\leq i_+}\alpha_i + \sum_{i_++1\leq i\leq k}\alpha_i$ and suppressing the terms that cancel out, and the last two are straightforward computations. The computations in higher genera are similar. □

Observe finally that

$$\Phi_{\sigma+B} = \Phi_\sigma \otimes e^{-B}.$$

2.3 The composition theorem

Definition 2.9 Let $P' \to \Sigma' = \Sigma'_{g',k',\ell'}$ and $P'' \to \Sigma'' = \Sigma''_{g'',k'',\ell''}$ be two marked Hamiltonian fibrations with same fiber (M,ω). Suppose that there is a bijection

$$\mu\colon \{i'_+ + 1, \dots, k'\} \to \{1, \dots, i''_+\}$$

between the "outgoing" p'_is and the "ingoing" p''_is and a bijection

$$\nu\colon \{j'_+ + 1, \dots, \ell'\} \to \{1, \dots, j''_+\}$$

between the "outgoing" S'_js and the "ingoing" S''_js. Assume moreover that the corresponding monodromies coincide, ie that $\xi'_j = \xi''_{\nu(j)}$.

(1) The *gluing of P' and P''*, denoted by $P = P' \# P''$, is by definition the marked Hamiltonian fibration obtained in the following way. First extend the symplectic trivialisations η'_i, $i \in \{i'_+ + 1, \dots, k'\}$ over small discs where the form Ω' can therefore be identified with $(D^2 \times M, \omega_{st} \oplus \omega)$ where ω_{st} is the area form on the disc. Do the same near η''_i, $i \in \{1, \dots, i''_+\}$. Use these identifications to perform the obvious surgery between $(\eta'_i)^{-1}(D^2 \times M)$ and $(\eta''_{\mu(i)})^{-1}(D^2 \times M)$ that covers the index 0 surgery between Σ' and Σ'' at the points p'_i and $p''_{\mu(i)}$. Similarly, consider a diffeomorphism $r_j\colon S'_j \to S''_{\nu(j)}$ preserving both the orientations and the base points $q'_j, q''_{\nu(j)}$, and define

$$\theta_j\colon W'_j \to W''_{\nu(j)}$$

as the symplectic diffeomorphism whose restriction to the fiber $M_{q'_j}$ is the map

$$(\eta'')^{-1}_{\nu(j)} \circ \eta'_j\colon M_{q'_j} \to M_{q''_{\nu(j)}}$$

and which sends each fiber M_b to its image $M_{r_j(b)}$ so that the characteristic foliations be preserved.

Then glue the hypersurface $W'_j \to S'_j$ to the hypersurface $W''_{\nu(j)} \to S''_{\nu(j)}$ via the attaching map θ_j. The resulting marked Hamiltonian fibration $P = P'\#P''$ over $\Sigma = \Sigma'\#\Sigma''$ has $\mathcal{A}'_+$ as ingoing data and $\mathcal{A}''_-$ as outgoing data.

(2) For each $i \in \{i'_+ + 1, \dots, \ell'\}$, there is an identification Φ_{QH} of $QC(M_{p'_i})$ with $QC(M_{p''_{\mu(i)}})$ induced by the identification of the fibers.

For each $j \in \{j'_+ + 1, \dots, \ell'\}$, there is an identification Φ_{FH} of $FC(S'_j)$ with $FC(S''_{\nu(j)})$ that maps (x', v') to an element (x'', v'') in the following way. The closed leaf $x'(t)$ is sent to $\theta(x'(t))$ and the disc $v' \subset P'_j$ is sent to the unique disc v'' in $P''_{\nu(j)}$ such that the gluing $v'\#v''$ in the fibration $P'_j\#P''_{\nu(j)} \to S^2$ is flat. To simplify notations, we will denote by Φ any tensor product of copies of Φ_{QH} and Φ_{FH}.

(3) Finally, if σ', σ'' are sections of $\bar{P}', \bar{P}''$, we will denote by $\sigma'\#\sigma''$ the section of $\overline{P'\#P''}$ such that in each $P'_j\#P''_{\nu(j)}$ the resulting sphere section be flat.

Here is the main result of this section.

Theorem 2.10 *Let P' and P'' be two marked Hamiltonian fibrations and $P = P'\#P''$ their gluing. Assume, to simplify the argument, that the number of outgoing data of P' is 1. Let σ', σ'' be section classes in P', P'' and $\sigma = \sigma'\#\sigma''$ their gluing. Then*

$$\Phi''_{\sigma''} \circ \Phi \circ \Phi'_{\sigma'} = \Phi_\sigma$$

Sketch of the proof

(A) To make the argument clearer, suppose first that each Morse chain and periodic orbit at the source and the target spaces are *cycles* and actually assume even that there is no element in their differential. We will examine the general case of the construction of a chain homotopy in (B) below.

First note that we may assume that the Morse–Smale data (κ'_i, g'_{M_i}) at the point p'_i coincides with the Morse–Smale data at the point $p''_{\mu(i)}$. So under these hypotheses, the statement reduces to establishing the formula

$$\begin{aligned}
&\sum m_{g',k',\ell',\sigma'_{norm},\{J'_b\},F'}\Big(a'_1,\dots,a'_{i'_+},b'_{i'_+ +1},\dots,b'_{k'},(x'_1,v'_1),\dots,(x'_{\ell'},v'_{\ell'}),\tau'\Big)\\
&\quad\times m_{g'',k'',\ell'',\sigma''_{norm},\{J''_b\},F''}\Big(\Phi_{QH}(b'_{i'_+ +1}),\dots,\Phi_{QH}(b'_{k'}),b''_{i''_+ +1},\dots,b''_{k''},\\
&\qquad\Phi_{FH}(x'_{j'_+ +1},v'_{j'_+ +1}),\dots,\Phi_{FH}(x'_{\ell'},v'_{\ell'}),(x''_{j''_+ +1},v''_{j''_+ +1}),\dots,(x''_{\ell''},v''_{\ell''}),\tau''\Big)\\
&= m_{g,k,\ell,\sigma_{norm},\{J_b\},F}\Big(a'_1,\dots,a'_{i'_+},b''_{i''_+ +1},\dots,b''_{k''},(x'_1,v'_1),\dots,(x'_{j'_+},v'_{j'_+}),\\
&\qquad\qquad(x''_{j''_+ +1},v''_{j''_+ +1}),\dots,(x''_{\ell''},v''_{\ell''}),\tau\Big) \quad (*)
\end{aligned}$$

where the sum is taken over all elements $b'_{i'_+ +1}, \ldots, b'_{k'}$ of a basis of $H_*(M) \otimes \ldots \otimes H_*(M)$, all tensor products $x'_{j'_+ +1} \otimes \ldots \otimes x'_{\ell'}$ of 1–turn closed characteristic leaves, and all classes τ', τ'' such that $\tau' \# \tau'' = \tau$.

The quantum part of the summation is taken over all elements $b'_{i'_+ +1}, \ldots, b'_{k'}$ of a basis of $H_*(M) \otimes \ldots \otimes H_*(M)$ corresponding to the outgoing data p'_i of the fibration P' (and their images by Φ_{QH} in the ingoing data of the second fibration P''). Since this is the basis of the Morse complex corresponding to (κ'_i, g'_{M_i}) at the M–fiber over the point p'_i, and because the boundary conditions of the defining equation of the moduli space tell us that a solution u' meets the *stable* submanifold of the critical points b'_i while a solution u'' meets the *unstable* submanifold of the corresponding points, the above summation corresponds to taking the inverse image $\mathcal{I}$ by the evaluation map

$$\begin{aligned}\mathcal{M}^{\Gamma}_{g',k',\ell',\sigma'_{norm},\{J'_b\},F'} &\left(a'_1, \ldots, a'_{i'_+}, (x'_1, v'_1), \ldots, (x'_{\ell'}, v'_{\ell'}), \tau'\right) \times \\ &\mathcal{M}^{\Gamma}_{g'',k'',\ell'',\sigma''_{norm},\{J''_b\},F''} \Big(b''_{i''_+ +1}, \ldots, b''_{k''}, \Phi_{FH}(x'_{j'_+ +1}, v'_{j'_+ +1}), \ldots, \\ &\qquad \Phi_{FH}(x'_{\ell'}, v'_{\ell'}), (x''_{j''_+ +1}, v''_{j''_+ +1}), \ldots, (x''_{\ell''}, v''_{\ell''}), \tau''\Big) \\ &\qquad\qquad \longrightarrow \Pi_{i \in \{i'_+ +1, \ldots,, k'\}} (M_i \times M_i)\end{aligned}$$

of the product over the $i \in \{i'_+ + 1, \ldots, k'\}$ of the cycles $\sum_{r \in \mathrm{Crit}(\kappa_i)} b'_{i,r} \otimes a'_{i,r}$ where $b'_{i,r}$ is the stable submanifold corresponding to a critical point q_r of the Morse function κ'_i and $a'_{i,r}$ is the unstable submanifold corresponding to the same critical point. But $\sum_{r \in \mathrm{Crit}(\kappa_i)} b'_{i,r} \otimes a'_{i,r}$ is a cycle in $M_i \times M_i$ homologous to the diagonal $\Delta_i \subset M_i \times M_i$. Hence the proof reduces to establish the above equation $(*)$ when the intermediate quantum data belong to the diagonal.

This boils down to a gluing theorem that gives a map from pairs of solutions in $\mathcal{I}$ to solutions in the moduli space

$$\begin{aligned}\mathcal{M}^{\Gamma}_{g,k,\ell,\sigma_{norm},\{J_b\},F} &\Big(a'_1, \ldots, a'_{i'_+}, b''_{i''_+ +1}, \ldots, b''_{k''}, (x'_1, v'_1), \ldots, (x'_{j'_+}, v'_{j'_+}), \\ &\qquad (x''_{j''_+ +1}, v''_{j''_+ +1}), \ldots, (x''_{\ell''}, v''_{\ell''}), \tau\Big).\end{aligned}$$

The well-known gluing techniques in weakly monotone manifolds give such a map: indeed, the gluing techniques for solutions near a fiber M_{p_i} are exactly those exposed in Lalonde–McDuff–Polterovich [9] and developed in McDuff [11] (see also Ruan–Tian [19]), while the gluing techniques for solutions on the asymptotic ends are Floer's theorems on the composition of homotopies. The essential point is to show that this gluing map is algebraically onto. This is what we now prove.

Observe first that if u is a solution in

$$\mathcal{M}^{\Gamma}_{g,k,\ell,\sigma_{norm},\{J_b\},F} = \mathcal{M}_{g,k,\ell,\sigma_{norm},\{J_b\},F}\Big(a'_1,\dots,a'_{i'_+},b''_{i''_+ +1},\dots,b''_{k''},$$
$$(x'_1,v'_1),\dots,(x'_{j'_+},v'_{j'_+}),(x''_{j''_+ +1},v''_{j''_+ +1}),\dots,(x''_{\ell''},v''_{\ell''}),\tau\Big)$$

then the Ω–area of u is determined by the boundary conditions and the homology class τ, ie it is fixed by the homology class of u relative to its boundary conditions.

Lemma 2.11 *Let u be a solution in $\mathcal{M}^{\Gamma}_{g,k,\ell,\sigma_{norm},\{J_b\},F}$. Then the Ω–area of each end is non-negative.*

Proof Recall that the restriction of a solution u near the jth end is a map

$$u_j\colon\ S^1\times[0,\infty)\to P$$

that satisfies the equation $(**)$. The computation of the s–energy of the restriction of the solution to that end leads to

$$0\le\int_{S^1\times[0,\infty]}\left\|\frac{\partial u_j}{\partial s}\right\|^2\le\left(\int_{S^1\times[0,\infty]}(u_j)^*(\Omega)\right)-T_j$$

where $T_j>0$ is the total variation of the function F_j. Thus the Ω–area of that end is bounded from below by $T=T_j$, and is therefore non-negative. □

Now let u_0 be a solution in $\mathcal{M}^{\Gamma}_{g,k,\ell,\sigma_{norm},\{J_b\},F}$ and γ_j be a loop on Σ in the class where the gluing of S'_j with $S''_{\nu(j)}$ took place – to simplify notation, we will now drop the indices. Denote by W the inverse image $\pi^{-1}(\gamma)$; note that the kernel of $\Omega_W=\Omega|_W$ is transverse to the fibers. Recall that in a neighbourhood $U=\gamma\times[-\delta,\delta]$ of γ, Ω has the form $\Omega'_W+d(s\alpha)$ where α is any 1–form on W that does not vanish on the kernel of Ω_W and s is the coordinate transverse to γ. Stretch the neck of U and deform the equation $\bar\partial_J u=0$ to some well-chosen non-homogeneous equation over the neck in order to force the solutions to approach closed leaves of W. More precisely, for arbitrarily large $\kappa\ge\delta$, let us denote by $J_{\Sigma,\kappa}$ any complex structure in the image of σ_{norm} for the surface Σ which, on $U=\gamma\times[-\delta,\delta]$ is conformally equivalent to the standard complex structure on $\gamma\times[-\kappa,\kappa]$ in such a way that the resulting structure be symmetric with respect to the S^1–action and anti-symmetric with respect to the $\mathbb{Z}_2$–action that changes the sign of the coordinate in $[-\delta,\delta]$. For an arbitrary positive real number $T\ge 0$, let $h_T\colon [-\delta,\delta]\to\mathbb{R}$ be a non-decreasing function which is constant and equal to $-T/2$ near $-\delta$ and is constant and equal to $T/2$ near the other end $+\delta$; thus its total variation is $T \in \mathbb{R}^+$. Denote by F_T its pull-back to $\pi^{-1}(U)$ by the projections $\pi^{-1}(U) \to \gamma \times [-\delta, \delta] \to [-\delta, \delta]$. The generic family $J_b, b \in \Sigma$, being fixed, for each $J_{\Sigma,\kappa}$, denote by $J_{P,\kappa}$ an almost-complex structures on P, compatible with Ω, which restricts to the family $J_b, b \in \Sigma$, on the fibers and which is such that the projection is pseudo-holomorphic with respect with the structures $J_{\Sigma,\kappa}$ on the base. For each κ and T, consider the sections $u = u_{\kappa,T} \colon \Sigma \to P$ that realise the class τ, which over $\gamma \times [-\delta, \delta]$ are solutions of the equation:

$$\bar{\partial}_{J_{\Sigma,\kappa}, J_\kappa} u = -\nabla F_T$$

and which over the rest of Σ are solutions of the same equation and boundary conditions as for the original $u_0 = u_{\kappa=\delta, T=0}$. Thus the moduli space of solutions is the set of all quadruples $(\kappa, J_{\Sigma,\kappa}, T, u_{\kappa,T})$ where $\kappa \geq \delta, T \geq 0$, $J_{\Sigma,\kappa}$ belongs to the image of σ_{norm} and $u_{\kappa,T}$ is a solution of the above system. Denote it by $\mathcal{M}_\kappa$. Because the solution u_0 is both generic and isolated (amongst *all* normalised conformal structures on Σ), the real dimension of $\mathcal{M}_\kappa$ is 1. This is because, basically, the moduli space $\mathcal{M}_\kappa$ is obtained from the original one by deforming the right hand side along a one-parameter family given by T. The computation of the s–energy over $\gamma \times [-\kappa, \kappa]$ leads to

$$0 \leq \int_{\gamma \times [-\kappa,\kappa]} \left\| \frac{\partial u_{\kappa,T}}{\partial s} \right\|^2 \leq \left(\int_{\gamma \times [-\kappa,\kappa]} (u_{\kappa,T})^*(\Omega) \right) - T \leq \operatorname{area}(u) - T$$

because, by Lemma 2.11, the restriction of Ω to the rest of the image of $u_{\kappa,T}$ is non-negative. Here $\operatorname{area}(u)$ is the Ω–area of the solution $u_{\kappa,T}$ over all of Σ, which depends only on the homology class of the solution and the boundary conditions – it is a constant attached to the moduli space. Thus there is no solution when $T > \operatorname{area}(u)$. By Gromov–Floer's compactness theorems, the conformal structure of $J_{\Sigma,\kappa}$ must degenerate as T approaches its upper bound. But any degeneracy that would occur away from the set consisting of the union of the ends and of $\gamma \times [-\kappa, \kappa]$, is of real codimension 2 and can therefore be avoided in $\mathcal{M}_\kappa$. Because by our hypothesis on cycles, there is no broken Floer flow line near the ends (ingoing or outgoing), then the degeneracy must occur over $\gamma \times [-\kappa, \kappa]$, ie there is a sequence of triples (κ_i, T_i, u_i) such that $T_i \to T_\infty \leq \operatorname{area}(u)$ and $\kappa_i \to \infty$. But because the s–energies of the u_is are bounded, there is a sequence s_i of levels for which the integrals

$$\int_{s=s_i} \left\| \frac{\partial u_i}{\partial s} \right\|^2 dt$$

converge to zero, which means that $J_{\kappa_i} \frac{\partial u_i}{\partial t} + \nabla F_i$ is arbitrarily L^2–small: thus the loops $u_i|_{s_i}$ must converge to a periodic orbit of the characteristic foliation of the function $F_{T_\infty}\colon P \to \mathbb{R}$. But on U, this foliation corresponds to the characteristic flow of a copy of W (lying over s_i). This shows that the space of solutions in the moduli space $\mathcal{M}^{\Gamma}_{g,k,\ell,\sigma_{norm},\{J_b\},F}$ can be deformed to a space of solutions that has the same counting and for which each solution decomposes near γ_j.

This proves the claim if the surgery takes place on a cylindrical end. But if the surgery takes place on a marked point, one can use exactly the same "stretch the neck" argument with the difference that one applies it instead to the symplectically trivial cylinder $W \times I$ where I is an interval and W is the trivial product $S^1 \times (M, \omega)$. This produces a *flat* holomorphic cylinder over the infinite cylindrical end, because the closed orbits of a monotone function on $S^1 \times I$ that depends only on the the variable $t \in I$, are the circles $S^1 \times \{t\}$ over which the monodromy in $S^1 \times M \times I$ is trivial. By choosing the family J_b over $S^1 \times I$ independent of b, and projecting solutions to the fiber, this yields a solution of an elliptic Cauchy–Riemann type equation over $S^1 \times [0, \infty)$ and over $S^1 \times (-\infty, 0]$ as well. But each one is conformally equivalent to the unit disc D^2 with the origin removed. Thus, since the solution converges to a constant in the fiber M, one may extend it over all D^2 (in other words, we have reproved here the "removal of singularity" theorem for maps with finite energy). This yields two solutions, one on each unit disc, that meet transversally at a common point when the two M–fibers over the origins of the two discs are identified.

(B) It is now easy to generalise the previous argument to the case when the chains are not assumed to be cycles. Clearly, one can still consider the same deformed moduli space $\mathcal{M}_\kappa$. It is non-empty and can only degenerate in three ways:

(1) at either a Morse or a Floer broken flow line at one of the ingoing data,

(2) over some $\gamma \times [-\kappa, \kappa]$, or

(3) at a Morse or a Floer broken flow line at one of the outgoing data.

But this, counted with signs, leads immediately to the formula of a homotopy operator:

$$\partial H = H\partial + \Phi''_{\sigma''} \circ \Phi \circ \Phi'_{\sigma'} - \Phi_\sigma$$

where the three terms on the right hand side correspond to the three above possibilities in the same order. □

Lemma 2.12 *Let P be a topologically trivial M–fibration over S^2 with two marked points $p_1 \in \mathcal{A}_+$ and $p_2 \in \mathcal{A}_-$, equipped with a ruled symplectic form Ω deformation equivalent to the split form, and let σ_0 be the flat section $pt \times S^2$ of*

P. Then the map $\Phi_{\sigma_0}\colon QH_*(M) \to QH_*(M)$ corresponding to these marked points is the identity map.

The same is true if P is a topologically trivial M–fibration over the cylinder, equipped with a ruled symplectic form Ω deformation equivalent to the split form, and if σ_0 be the flat section $pt \times (S^1 \times I)$ of P. Then the map $\Phi_{\sigma_0}\colon FH_*(M) \to FH_*(M)$ is the identity map.

The proof of the first part was sketched in Lemma 4.A of Lalonde–McDuff–Polterovich [9] and was written down for general symplectic manifolds by McDuff in [11]. The second part is proved similarly – it is relatively easy in the weakly monotone case and left to the reader.

We say that a pair (P, σ) is *deformation equivalent to the split pair* if P has a smooth trivialisation $\phi\colon P \to \Sigma \times M$ that sends σ to the flat section and the form Ω to a form which is deformation equivalent to the split form.

Corollary 2.13 *Let $(M, \omega) \hookrightarrow P \to S^2$ be a Hamiltonian fibration with two marked points $p_1 \in \mathcal{A}_+$ and $p_2 \in \mathcal{A}_-$ and σ a homology class of sections up to equivalence. If there is another pair (P', σ') with two marked points $p_1' \in \mathcal{A}_+'$ and $p_2' \in \mathcal{A}_-'$ such that the pair $(P \# P', \sigma \# \sigma')$ is deformation equivalent to the split pair, then $\Phi_{P,\sigma}\colon QH(M_{p_1}) \to QH(M_{p_2})$ is an isomorphism.*

Let $(M, \omega) \hookrightarrow P \to D^2$ be a Hamiltonian fibration with one marked point $p \in \mathcal{A}_+$ and with $\partial D^2 \in \mathcal{A}_-$ and let σ a homology class of sections of $\bar{P}$ up to equivalence. If there is another similar pair (P', σ') with one marked point $p' \in \mathcal{A}_-'$ and $\partial D^2 \in \mathcal{A}_+$, such that the monodromies α and α' of their boundaries coincide, and if the pair $(P \# P', \sigma \# \sigma')$ is deformation equivalent to the split pair, then $\Phi_\sigma\colon QH(M_p) \to FH(\partial P)$ is an isomorphism.

Proof The proof is a direct consequence of the Composition theorem 2.10 and of Lemma 2.12. □

Remark Note that the first part of the corollary says that the Seidel map is an isomorphism, while the second part implies that the PSS map from quantum to Floer homology is an isomorphism. Note that in the second part of the last corollary, one could permute the sets $\mathcal{A}_+$ and $\mathcal{A}_-$ so that the gluing of P, P' would give a cylinder – using the second part of Lemma 2.12 instead, we would be led to the conclusion that $\Phi_\sigma\colon FH(M_p) \to QH(\partial P)$ is an isomorphism, which is then of course the inverse of $\Phi_\sigma\colon QH(M_p) \to FH(\partial P)$.

Remark All results of this section obviously generalise to ruled symplectic fibrations, ie fibrations $(M,\omega) \to P \to \Sigma$ having $\mathrm{Symp}(M,\omega)$ as structural group instead of $\mathrm{Diff}_{Ham}(M,\omega)$, if the following three conditions are satisfied:

(1) there is a closed extension of the fiberwise symplectic forms,

(2) the monodromy round each end is a Hamiltonian diffeomorphism, and

(3) there is at least one section of the fibration $\bar{P}$.

Thus, for instance, the monodromy round some closed loop in the base Σ could be a symplectic diffeomorphism not isotopic to the identity. In this case, the field theory developed in this section cannot be reduced to a composition of PSS and Seidel homomorphisms.

3 How essentiality implies minimality

The following definition is due to Polterovich [18] and Schwartz [21]: a generator $(x,v) \in FC_*(H_t)$ is *essential* if there is a class $a \in FH$ such that any cycle α representing a must contains the element (x,v) (ie it appears with non-vanishing coefficient in the cycle). An equivalent way of expressing this is: the inclusion $L \to FC(H_t)$ of the subcomplex generated by all the generators of $FC_*(H_t)$ except (x,v) induces a morphism between the corresponding homologies which is not onto.

Let $\mathcal{H}$ be the space $C^\infty([0,1] \times M, \mathbb{R})$. It is well-known that the Hamiltonian paths $\phi^H_{t\in[0,1]}$ generated by these functions can as well be generated by functions in $C^\infty(S^1 \times M, \mathbb{R})$.

Here and after, we use the following definition of Hofer's length of an element $H \in \mathcal{H}$: $\int_0^1 (\max_M H_t - \min_M H_t)dt$. A function in $\mathcal{H}$ is called *quasi-autonomous* if there are two points $P, Q \in M$ such that P is a global maximum of each $H_t, t \in [0,1]$, and Q is a global minimum of each $H_t, t \in [0,1]$. We will refer to these points as *fixed maximum, fixed minimum* respectively.

For any $H \in \mathcal{H}$, denote by $\bar{H}$ the *opposite* Hamiltonian, ie the function defined by

$$\bar{H}(t,x) = -H(t, \phi^H_t(x)).$$

This Hamiltonian generates the path $(\phi^H_t)^{-1}$. Obviously, the roles played by P and Q are reversed while the Hofer length remains the same.

We will show:

Theorem 3.1 *Let (M, ω) be a symplectic manifold, that we assume to be weakly monotone for simplicity. Let $H = H_{t\in[0,1]}$ be any Floer-generic quasi-autonomous Hamiltonian for which the class (x_{max}, v_{cst}) is essential in $FC_*(H)$ and (x_{min}, v_{cst}) is essential in $FC_*(\bar{H})$ (here v_{cst} denotes the constant disc). Then the Hofer length of the path ϕ_t generated by H is minimal among all paths joining the identity to ϕ_1 which are homotopic to ϕ_t with fixed endpoints* id, ϕ_1. *If the manifold is symplectically aspherical, the path ϕ_t is length minimising among all paths joining* id *to* ϕ_1.

Recall that a *symplectically aspherical* means that the integral of ω over any 2–sphere vanishes.

In all results on the minimality of geodesics in Hofer's geometry, there are essentially two steps: in the first one, one shows that some conditions concerning the dynamics of the Hamiltonian path (inexistence of periodic orbits of period one, etc) implies that a more abstract property is satisfied, like the essentiality of the two classes (x_{min}, v_{cst}) and (x_{max}, v_{cst}) or some energy-capacity inequality. The second step then consists in proving that this latter property implies length-minimality.

The first step was carried out in this setting in [5] for aspherical manifolds. The goal of this section is to prove that the second step is true for all weakly monotone manifolds (we restrict ourselves to weakly monotone manifolds only for simplicity). There was an attempt by Oh to complete this scheme in full generality in [14, 16], but a recent erratum [15] restricts the scope of validity of his results (however the second step in Oh's papers, by which he proves that the energy-capacity inequality implies length-minimality, does not seem at first sight affected by his erratum). In any case, our aim is to show that the homological essentiality implies minimality, using only simple geometric methods and our field theory. The key ingredient of this section is Proposition 3.10: it is there where our field theory intervenes. Inasmuch as one is willing to consider that our field theory can be generalised to all manifolds using perturbations and virtual cycles, then Theorem 3.1 would hold for all symplectic manifolds and therefore the second step of the above scheme would be completed.

Recall that a fixed point x of the flow $\phi^H_{t\in[0,1]}$ is *under-twisted* if, given any value $T \in [0,1]$, the linearised flow $D\phi^H_{t\in[0,T]}(x)\colon T_xM \to T_xM$ has no non-trivial closed orbit. It is *generically under-twisted* if, given any value $T \in [0,1]$, the linearised flow $D\phi^H_{t\in[0,T]}(x)\colon T_xM \to T_xM$ has only 0 as fixed point (ie we also exclude all trivial closed orbits, except the orbit at 0).

Definition 3.2 Let H_0 is a Morse function on a symplectic manifold and let $H_{t\in[0,1]}$ be a quasi-autonomous Hamiltonian starting with H_0. We say that $H_{t\in[0,1]}$ *has no index jump* if

(1) each closed orbit at time 1 is the endpoint of a continuous family, as T goes from 0 to 1, of closed orbits of $\phi^H_{t\in[0,T]}$ for which the capping discs can be chosen continuously so that, at time $T=0$, it is the constant one, and

(2) with respect to these capping discs, as T goes from 0 to 1 the Conley–Zehnder indices remain the same as in the $t=0$ Morse case.

We say that $H_{t\in[0,1]}$ is *restrained* if the condition (1) holds but the second one is replaced by

(2′) with respect to these capping discs, the action functional remains in the range $[\mathcal{A}^H_T(Q), \mathcal{A}^H_T(P)]$ for all $T\in[0,1]$.

Note that the hypothesis of "under-twisted" simply means that the condition (2) holds for the fixed minimum and maximum. Note also that these conditions do not mean that new closed orbits or index jumps cannot appear in the time interval $T\in[0,1]$; what they say is that, if they appear, they must disappear before time 1.

It is an interesting question to know whether or not each of these two sets of conditions ((1) and (2), or (1) and (2')) implies Hofer's minimality of the path $\phi^H_{t\in[0,1]}$. We will not examine in this paper the "restrained" set of conditions. But we will show that the first set of conditions (ie (1) and (2)) is sufficient if the minimal Chern number is either 0 or large enough:

Corollary 3.3 *Let (M^{2n},ω) be a symplectic manifold with minimal Chern number not in the range $[1,n]$ (a projective space for instance). Suppose that H_0 is a Morse function and let $H_{t\in[0,1]}$ be an undertwisted quasi-autonomous Hamiltonian without index jump. Then $H_{t\in[0,1]}$ induces a Hamiltonian path which is minimal in Hofer's length in its homotopy class with fixed endpoints.*

Proof First, it is an easy exercise to check that the proof in Section 8 of Lalonde–Kerman [5] applies here as well, with minor changes, so that we may assume that the Hamiltonian $H_{t\in[0,1]}$ in the statement of the corollary has the following additional generic properties:

(a) for all $t\in[0,1]$, P is the unique global maximum of H_t and is non-degenerate,

(b) for all $t \in [0,1]$, Q is the unique global minimum of H_t and is non-degenerate, and

(c) H is Floer-generic.

Now, the generators of the Floer complex $FC_*(H)$ are pairs consisting of a periodic orbit $x(t)$ and a capping disc v. By definition of "no index jump", each closed orbit x is the endpoint of a continuous family starting at some critical point of H_0 with continuously evolving capping discs with constant Conley–Zehnder index. Thus a generator (x,v) has index equal to the Morse index of a critical point of H_0 (in $[0,2n]$) plus or minus some multiple of twice the Chern class of some 2–sphere in M. By hypothesis, this Chern number cannot be in $[1,n]$, so the index of a generator cannot be equal to $2n+1$. This shows that two different cycles of index $2n$ cannot be homologous. Hence, to establish the essentiality of (P, v_{cst}) in $FC_*(H)$ (here v_{cst} denotes the constant disc at P), there only remains to show that there is a cycle in $FC_{2n}(H)$ that contains the element (P, v_{cst}). This is done exactly as in the proof of Proposition 5.2 of [5]. (It is for that proof that one needs the above additional generic properties (a) and (b).) Actually, the proof in [5] is written down for an aspherical manifold, but it works as well for a general manifold; one simply needs to replace P by (P, v_{cst}). This establishes the essentiality of (P, v_{cst}) in $FC_*(H)$; the essentiality of (Q, v_{cst}) in $FC_*(\bar{H})$ is proved similarly. By Theorem 3.1 above, we get minimality. □

Here is a simple example of application of that Corollary. Given a symplectic manifold with minimal Chern number not in the range $[1,n]$, $\mathbb{CP}^n$ for instance, let $f\colon M \to \mathbb{R}$ be a Morse function, for instance a function on $\mathbb{CP}^n$ which once pulled-back on $\mathbb{C}^{n+1} - \{0\}$, is of the form

$$\frac{\sum_{i,j}(a_{i,j}x_ix_j + b_{i,j}x_iy_j + c_{i,j}y_iy_j)}{\|z\|^2}.$$

Let P be a global maximum of f and Q a global minimum of f such that the linearised flows have no nontrivial closed orbit in time ≤ 1. Note that they are then automatically generically undertwisted by the Morse condition. Let $\varepsilon > 0$ be sufficiently small so that the flow ϕ_t^f induced by f be transversal to the diagonal in $M \times M$ for all $t \in (0,\varepsilon]$. Consider any Hamiltonian $H_{t\in[0,1]}$ such that:

(1) on some open neighbourhoods of P and Q, $H_t = f$ for all $t \in [0,1]$,

(2) $H_t = f$ for $t \leq \varepsilon$,

(3) the graph of ϕ_t^H remains transversal to the diagonal in $M \times M$ for all $t \in [\varepsilon, 1]$, and

(4) $H_t(x) \in [f(Q), f(P)]$ for all $x \in M$ and $t \in [0,1]$.

Then, clearly, each fixed point of ϕ_1^H is the endpoint of a continuous family, as T goes from 0 to 1, of closed orbits of $\phi^H_{t\in[0,T]}$. Thus the capping discs can be chosen continuously so that, at time $T = 0$, they are the constant ones. By the preceding corollary, one concludes that this path is length minimising rel endpoints in its homotopy class.

Let us go back to Theorem 3.1. Its proof takes the next two paragraphs: in the first one, we introduce the geometric tools that will be necessary: gluing of monodromies, Hamiltonian cylinders and fibrations, and area estimates. In the second paragraph, we conclude the argument in two cases:

(1) M is monotone and the paths are in the same homotopy class rel endpoints in $\mathrm{Ham}(M,\omega)$;

(2) M is aspherical

and we prove the stronger statement of Theorem 3.1, ie minimality with respect to *all* paths with fixed endpoints.

3.1 Gluing Hamiltonian fibrations along monodromies

Let us first recall from [6, 7] that if $H_{t\in[0,1]}, K_{t\in[0,1]}$ generate paths in the group of Hamiltonian diffeomorphisms of (M,ω) joining the identity to the same Hamiltonian diffeomorphism $\phi = \phi^H_{t=1} = \phi^K_{t=1}$, one can construct a symplectic manifold $R_{H,K}$ by gluing the region under the graph of H with the region above the graph of K. We will introduce a variant of this construction.

First, after reparametrisation of the Hamiltonian paths $\phi_{t\in[0,1}$, we may assume that all the generating Hamiltonians functions $G_{t\in[0,1]}$ are normalised so that

(1) they are the restrictions to the time interval $[0,1]$ of a time-periodic Hamiltonian of period 1 and

(2) the minimum over M of G_t equals 0 for every t and the maximum over M of G_t is independent of t (it is therefore equal to $\mathcal{L}(G)$).

Then G may be considered as a map $G\colon M \times S^1 \to [0,\infty)$. The restriction of the form $\Omega = \omega \oplus ds \wedge dt \in \Omega^2(M \times S^1 \times \mathbb{R})$ to the graph of a normalised Hamiltonian G gives rise to a characteristic foliation. It is well known that the monodromy of that foliation is equal to ϕ_1^G. The "region over the graph" of G is the subset of $M \times S^1 \times \mathbb{R}$ defined by

$$R^+(G) = \left\{ (x,t,s) : G_t(x) \le s \le \max_M G_t = \mathcal{L}(G) \right\}$$

and the "region under the graph" of G is the subset of $M \times S^1 \times \mathbb{R}$

$$R^-(G) = \left\{(x,t,s) : \min_M G_t = 0 \le s \le G_t(x)\right\}.$$

In order that all constructions in the sequel be smooth, we will have to consider spaces $R^-_\varepsilon(G)$ for arbitrarily small $\varepsilon > 0$:

$$R^-_\varepsilon(G) = \{(x,t,s) : 0 \le s \le G_t(x) + \varepsilon\}.$$

so that the lower and upper components of the boundary of $R^-_\varepsilon(G)$ are disjoint. Similarly:

$$R^+_\varepsilon(G) = \{(x,t,s) : G_t(x) \le s \le \mathcal{L}(G) + \varepsilon\}.$$

Note that $R^-_\varepsilon(G)$ is symplectomorphic to $\{(x,t,s) : -\varepsilon \le s \le G_t(x)\}$, so both spaces are obtained by gluing to $R^-(G), R^+(G)$ the product of (M,ω) with the annulus $S^1 \times I$, where I is an interval of length ε.

Let $q_-\colon M \times S^1 \times [0,\infty) \to M \times \mathbb{R}^2$ be the map that is the identity on M, sends t to the angle coordinate of $\mathbb{R}^2$ that we will still denote by t, and maps $s \in [0,\infty)$ to the action coordinate $c = \pi r^2 \in [0,\infty)$ of $\mathbb{R}^2$. Thus q_- maps the form $\omega + ds \wedge dt$ to the form $\omega + dc \wedge dt$ and sends $R^-_\varepsilon(G)$ to

$$P^-_\varepsilon(G) = \{(x,t,c) : 0 \le c \le G_t(x) + \varepsilon\} \subset M \times \mathbb{R}^2.$$

Similarly, let $q_+\colon M \times S^1 \times (-\infty, \mathcal{L}(G) + \varepsilon] \to M \times \mathbb{R}^2$ be the map that is the identity on M, sends t to the angle coordinate of $\mathbb{R}^2$ and maps s to $c(s) = \mathcal{L}(G) + \varepsilon - s$. Thus q_+ maps the form $\omega + ds \wedge dt$ to $\omega - dc \wedge dt$ and sends $R^+_\varepsilon(G)$ to

$$P^+_\varepsilon(G) = \{(x,t,c) : 0 \le c \le \mathcal{L}(G) + \varepsilon - G_t(x)\} \subset M \times \mathbb{R}^2.$$

This means that, with respect to the orientation $\mathcal{O}^+$ induced by $-dc \wedge dt$ on the boundary of the base of $P^+_\varepsilon(G)$, the monodromy of the characteristic flow of $\partial P^+_\varepsilon(G)$, followed in the direction of $\mathcal{O}^+$, is the inverse of the monodromy of $\partial P^-_\varepsilon(G)$, followed in the direction $\mathcal{O}^-$ induced by $dc \wedge dt$ on the boundary of the base of $P^-_\varepsilon(G)$.

Hence, up to diffeomorphism, $P^\pm_\varepsilon(G)$ may be considered as a topologically trivial fibration $M \overset{\pi}{\hookrightarrow} P^\pm_\varepsilon(G) \to D^2$ equipped with a ruled symplectic form Ω, ie a symplectic form that restricts to a non-degenerate form on each fiber, such that the monodromy round the boundary of the base covered in the positive direction is ϕ^G_1 for $P^-_\varepsilon(G)$ and $(\phi^G_1)^{-1}$ for $P^+_\varepsilon(G)$. Here the positive direction is intrinsically defined by following the symplectic gradient of a function whose level set is $\partial P^\pm_\varepsilon(G)$ and whose ordinary gradient points outward. The diffeomorphism that does this will be denoted by $\alpha^-_\varepsilon(G)$: it preserves t and maps the sets $(x, t, G_t(x)+\varepsilon)$, for each given t, to the M fiber over the point of ∂D^2 with angle t (and similarly for $\alpha_\varepsilon^+(G)$).

Note finally that any ruled symplectic form Ω on a fibration $M \hookrightarrow P \to D^2$ has a fiberwise symplectic trivialisation, ie a map $\phi\colon P \to D^2 \times (M, \omega)$ that commutes with the projection to D^2 and maps Ω to a form that restricts precisely to ω on each fiber (this trivialisation can be constructed by using the connexion induced by Ω over the rays in D^2 emanating from a base point). However, if $\pi_1(\mathrm{Ham}(M, \omega))$ is not $\{0\}$, this trivialisation is not unique up to homotopy. Our construction associates to each normalised Hamiltonian G a ruled symplectic fibration $P_\varepsilon^\pm(G)$ equipped with such a trivialisation.

Let P_1, P_2 be two M–fibrations over D^2 equipped with ruled symplectic structures Ω_i and with the fibers over the base point $1 \in \partial D^2$ symplectically identified. Suppose that the monodromies, via this identification, are inverses of each other. This yields a ruled symplectic structure on a M–fibration $P_1 \# P_2$ over S^2 by gluing the monodromies in the obvious way. This applies in particular to the pair $P_\varepsilon^-(H), P_\varepsilon^+(K)$ if H_t, K_t generate Hamiltonian flows with the same time-one maps $\phi_{t=1}^H = \phi_{t=1}^K$. Denote by $P_\varepsilon(H, K)$ the resulting fibration over S^2. Similarly, denote by $P_\varepsilon(K, H)$ the fibration over S^2 obtained by gluing $P_\varepsilon^-(K)$ to $P_\varepsilon^+(H)$. We will call them the *mixed fibrations* associated to H and K. Note that $P_\varepsilon(G, G)$ is the symplectically trivial fibration over a sphere of area $\mathcal{L}(G) + 2\varepsilon$.

Definition 3.4

(i) Let $(M, \omega) \hookrightarrow (P, \Omega) \to B$ be a ruled symplectic form over a compact surface Σ. The *area* of P is by definition the quotient of the Ω–volume of P by the ω–volume of the fiber.

(ii) When $M \to P \to D^2$ is equipped with a fiberwise symplectic trivialisation, the characteristic flow round ∂P gives rise to a Hamiltonian flow ϕ_t on M. Then each pair $(x(t), v)$ of a periodic orbit of ϕ_t and of a disc v of M (defined up to homotopy) with boundary equal to $x(t)$ corresponds via the trivialisation to a section of P, that we will denote $\sigma(x(t), v)$, with boundary lying on the characteristic leaf of ∂P given by $x(t)$. If P_1, P_2 are two fibrations with fiberwise symplectic trivialisations and inverse monodromies, a periodic orbit $x(t)$ of the flow $\phi_t^{P_1}$, a bounding disc $v_1 \subset M$, and a bounding disc $v_2 \subset M$ of the flow $\phi_t^{P_2}(x(0))$ gives rise to a section $\sigma(x(t), v_1, v_2)$ of $P_1 \# P_2$.

Remark Note that if $H_{t\in[0,1]}$ has a fixed maximum p_{max} and if $v_{p_{max}}$ denotes the constant disc, the Hofer length of $H_{t\in[0,1]}$ is of course equal, up to ε, to the Ω–area of the section $\sigma(p_{max}, v_{p_{max}})$ in $P_\varepsilon^-(H)$. The same applies for p_{min} and $P_\varepsilon^+(H)$.

Proposition 3.5 (Basic inequalities) *Let H_t, K_t generate Hamiltonian flows with the same time-one maps $\phi_{t=1}^H = \phi_{t=1}^K$. Then $\mathcal{L}(H_t) \leq \mathcal{L}(K_t)$ if and only if*

$$2\mathcal{L}(H_t) \leq \mathrm{area}(P_\varepsilon(H,K)) + \mathrm{area}(P_\varepsilon(K,H))$$

for all $\varepsilon > 0$. Therefore $\mathcal{L}(H_t) \leq \mathcal{L}(K_t)$ if each of the following inequalities holds for all $\varepsilon > 0$:

$$\mathcal{L}(H_t) \leq \mathrm{area}(P_\varepsilon(H,K)) \tag{1}$$

$$\mathcal{L}(H_t) \leq \mathrm{area}(P_\varepsilon(K,H)) \tag{2}$$

Moreover, if the space $P_\varepsilon(H,K)$ has a homology class of sections $\sigma_{flat}^{H,K}$ whose $\Omega_{H,K}$–area is equal to $\mathrm{area}(P_\varepsilon(H,K))$ and which decomposes as

$$\sigma_{flat}^{H,K} = \sigma(p_{max}, v_{p_{max}}) + \sigma'$$

over $P_\varepsilon^-(H)$ and $P_\varepsilon^+(K)$ respectively, then the inequality (1) amounts to

$$\Omega_{P^+(K)}\text{–area of } (\sigma') \ \geq \ 0.$$

Similarly, if the space $P_\varepsilon(K,H)$ has a homology class of sections $\sigma_{flat}^{K,H}$ whose $\Omega_{K,H}$–area is equal to $\mathrm{area}(P_\varepsilon(K,H))$ and which decomposes as

$$\sigma_{flat}^{K,H} = \sigma' + \sigma(p_{min}, v_{p_{min}})$$

over $P_\varepsilon^-(K)$ and $P_\varepsilon^+(H)$ respectively, then the inequality (2) amounts to

$$\Omega_{P^-(K)}\text{–area of } (\sigma') \ \geq \ 0.$$

In that statement, of course, the various indices affecting Ω denote the spaces in which the form lives (we deleted the ε in the indices).

Proof We have: $\mathcal{L}(H_t) \leq \mathcal{L}(K_t)$ if and only if

$$2\mathrm{vol}(P_\varepsilon(H,H)) \leq \mathrm{vol}(P_\varepsilon(K,K)) + \mathrm{vol}(P_\varepsilon(H,H)).$$

But the latter is equal to $\mathrm{vol}(P_\varepsilon(H,K)) + \mathrm{vol}(P_\varepsilon(K,H))$, and after dividing out by the volume of M, we find $2\mathcal{L}(H_t) + 4\varepsilon \leq \mathrm{area}(P_\varepsilon(H,K)) + \mathrm{area}(P_\varepsilon(K,H))$ for all $\varepsilon > 0$, which means that

$$2\mathcal{L}(H_t) \leq \mathrm{area}(P_\varepsilon(H,K)) + \mathrm{area}(P_\varepsilon(K,H))$$

for all $\varepsilon > 0$.

The rest of the proposition is a direct consequence of the remark preceding the proposition. ∎

This establishes the strategy for the proof of minimality: we will find sections of the ruled symplectic fibrations $P_\varepsilon(H,K), P_\varepsilon(K,H)$ which decompose as above and for which the σ'–part is positive.

We need a last result on the relation of the areas of the fibrations $P_\varepsilon(H,K))$, $P_\varepsilon(K,H)$ and the areas of *flat sections* in these spaces. We will consider two cases:

(1) the path ϕ_t^K is homotopic to ϕ_t^H, and

(2) the condition on the homotopy of the paths is released but M is aspherical.

Let (P,Ω) be a (M,ω)–ruled symplectic manifold over the 2–sphere. By the characterisation of Hamiltonian bundles in [12], Theorem 6.36 (or [8], Theorem 1.1), any ruled symplectic manifold over a simply connected base is a Hamiltonian fibration. Now suppose that there is a fiberwise symplectic trivialisation $\psi\colon P \to S^2 \times M$. Then any other such trivialisation ψ' will differ by a map $\Psi\colon S^2 \to \mathrm{Symp}_0(M)$. Because $\pi_2(\mathrm{Symp}_0(M)) = \pi_2(\mathrm{Ham}(M))$, Ψ is homotopic to a map $S^2 \to \mathrm{Ham}(M)$. But these act trivially on the homology of $S^2 \times M$ by Proposition 5.4 in Lalonde–McDuff [8]. This means that the homology class of the section $\psi^{-1}(S^2 \times \{pt\})$ is independent of the chosen trivialisation of P. Such a section will be called *flat*.

Now, if ϕ^K is homotopic to ϕ^H with fixed endpoints, the fibrations $P_\varepsilon(H,K)$ and $P_\varepsilon(K,H)$ are symplectically trivial. The fiberwise symplectic trivialisations are induced by a choice of a homotopy $G_{s,t}$ between the Hamiltonian paths H_t and K_t. Thus there is a well-defined class of flat sections. Here is a description of that class: if $G_{s,t}$ is a homotopy of paths with fixed endpoints between H and K, with $G_{0,t} = H, G_{1,t} = K$, and if $(x(t), v_x)$ is any pair of a closed orbit of H and a bounding disc, then the homotopy $G_{s,t}$ gives rise to a homotopy between (x, v_x) and $(\phi_t^K(x), v_x^K)$, which defines a bounding disc v_x^K up to homotopy (ie as an element in $\pi_2(M, \phi_t^K(x))$). Both (x, v_x) and $(\phi_t^K(x), v_x^K)$ may be interpreted as sections of the corresponding cylinders $P_\varepsilon^-(H), P_\varepsilon^+(K)$, which after gluing, give the flat section σ_{flat} of P.

Lemma 3.6 *Let $(M,\omega) \hookrightarrow (P,\Omega) \to S^2$ be a ruled symplectic manifold with compact fiber, with $P = P_\varepsilon(H,K)$ or $P_\varepsilon(K,H)$ where K is a path homotopic to H with fixed endpoints, and H has a fixed minimum and fixed maximum. Then P is a Hamiltonian fibration and there is a unique homology class $\sigma_{flat} \in H^2(P)$ which represents the flat section. In the case of $P_\varepsilon(H,K)$, that class can be decomposed as the union (along the common boundary loop) of two homologically well-defined sections of the spaces $R^-(H)$ and $R^+(K)$, the first* *one given by* $(p_{max}, v_{p_{max}})$*, with* $v_{p_{max}}$ *the constant disc, and the second one given by the corresponding pair* $(\phi_t^K(p_{max}), v_{p_{max}}^K)$*. Moreover the area of* P *is equal to the* Ω*–area of* σ_{flat}*:*

$$\mathrm{vol}(P, \Omega) = (\mathrm{vol}(M, \omega)) \int_{\sigma_{flat}} \Omega.$$

A similar statement applies to $P_\varepsilon(K, H)$ *with* p_{max} *replaced by* p_{min}*.*

This lemma produces the σ' required in the last proposition, and there will remain to show later (see Section 3.2) that its area is non-negative.

Proof There remains only to establish the last equality. Here is a proof using Moser's argument. Take any homotopy $G_{s,t}$ and consider the corresponding spaces P_s, Ω_s obtained by gluing the cylinders $P_\varepsilon^{\pm}(G_0) = P_\varepsilon^{\pm}(H)$ and $P_\varepsilon^{\mp}(G_s)$. This gives a one-parameter family of forms Ω_s on $S^2 \times M$ for which some fixed fiber, say M_0, is symplectic for all $0 \leq s \leq 1$. Since $\mathrm{vol}(P, \Omega_s) - (\mathrm{vol}(M, \omega)) \int_{\sigma_{flat}} \Omega_s$ is a continuous function of s, the inverse image of 0 is closed and therefore, if this equality fails to be true for all s, there is smallest s_0 for which it does not hold. But, at this value s_0, one may transform the deformation Ω_s into a genuine isotopy on some small interval $[s_0, s_0 + \varepsilon)$ by either adding or subtracting a small multiple of the Thom class of the normal bundle of the fiber M_0. This boils down to add or remove a small symplectically split tube $D^2 \times M_0$. These new forms are isotopic to Ω_{s_0} and therefore the equality holds. But, it holds for the small (added or removed) split tube too, so it must hold for the undeformed form as well. □

Let us consider now the second case, when M is aspherical but the path ϕ_t^K need not be homotopic to ϕ_t^H. The mixed fibration $P = P_\varepsilon(H, K)$ or $P_\varepsilon(K, H)$ has therefore a single homology class of sections. The unicity is clear from the aspherical condition. The existence is a consequence of the fact that in this case P is a Hamiltonian fibration corresponding to the loop $(\phi_t^K)^{-1} \star \phi_t^H \in \pi_1(\mathrm{Ham}(M))$ or its inverse, and it is a consequence of the Arnold conjecture that the orbit of such a loop on a point of M is contractible (see [9] for instance).

Proposition 3.7 *Let* (M, ω) *be an aspherical manifold and* H_t, K_t *two Hamiltonians generating paths with the same endpoints, with* H *quasi-autonomous. Then each of the spaces* $P = P_\varepsilon(H, K)$ *or* $P_\varepsilon(K, H)$ *has a single homology class of sections. Moreover, denoting by* $\sigma_{H,K}, \sigma_{K,H}$ *these sections, we have:*

$$\mathrm{area}(P_\varepsilon(H, K)) + \mathrm{area}(P_\varepsilon(K, H)) = \int_{\sigma_{H,K}} \Omega_{H,K} + \int_{\sigma_{K,H}} \Omega_{K,H}.$$

Proof Consider an arc $\gamma \subset M$ joining a fixed minimum p_{min} of H_t to a fixed maximum p_{max} of H_t. Its trace under the flow of ϕ^H defines a subset of the graph of H (as a subset of $M \times \mathbb{R}^2$, with action-angle c, t–coordinates):

$$\begin{array}{rcc} f\colon\, [0,1] \times [0,1] & \longrightarrow & M \times \mathbb{R}^2 \\ (z,t) & \longmapsto & \left(\phi_t^H\left(\gamma\left(z\right)\right), t, H_t\left(\phi_t^H\left(\gamma\left(z\right)\right)\right)\right) \end{array}$$

Because M is aspherical, there is up to homotopy a unique way to extend this map to a cylinder

$$C_H\colon\, [0,1] \times S^1 \to M \times \mathbb{R}^2$$

which is obtained by gluing f with a map

$$f'\colon\, [0,1] \times [0,1] \to M \times \{(c,t) : t = 0\}$$

in such a way that the projection of C_H on the first factor is a homotopically trivial 2–sphere. Note that, when considered as a subset of $P_\varepsilon(H,H) = M \times S^2$, this cylinder can be extended in the obvious way to a 2–sphere in the flat section class by capping small discs of area ε.The $\Omega_{H,H}$–area of C_H is the length of $H_{t\in[0,1]}$. Because the graph of H is mapped symplectically to the graph of K along the characteristic foliations via the diffeomorphisms $\alpha_\varepsilon^\pm(H), \alpha_\varepsilon^\pm(K)$ (see the construction described in the three paragraphs before Definition 3.4), the cylinder C_H is mapped to a cylinder C with same area in the space graph$(K) \subset P_\varepsilon(K,K)$. Now C can be extended in a unique way, up to homotopy, to a 2–sphere C_K of $P_\varepsilon(K,K) = (M \times S^2, \Omega = \omega \oplus \sigma)$ where σ is the standard area form on S^2 whose area is the Ω–area of the fibration ie is $\mathcal{L}(K) + 2\varepsilon$. It is obtained by gluing two discs to C:

$$g^+, g^-\colon\, D^2 \to P_\varepsilon^+(K), P_\varepsilon^-(K).$$

Here g^+ is uniquely determined up to homotopy by requiring that its boundary is mapped to $(\phi_t^K(p_{max}), t, K_t(\phi_t^K(p_{max}))) \subset$ graph(K). Similarly, g^- has its image in $P_\varepsilon^-(K)$ and its boundary is sent to $(\phi_t^K(p_{min}), t, K_t(\phi_t^K(p_{min}))) \subset$ graph(K).

This construction shows that (we omit the epsilons for simplicity)

$$\begin{aligned} \mathcal{L}(K) = \text{area}(C_K) = \text{area}(C) + \text{area}(g^+) + \text{area}(g^-) \\ = \mathcal{L}(H) + \text{area}(g^+) + \text{area}(g^-). \end{aligned}$$

On the other hand, the decomposition of the flat section of $P_\varepsilon(H,K)$ described in the last lemma leads to:

$$\int_{\sigma_{H,K}} \Omega_{H,K} = \mathcal{L}(H) + \int_{g^+} \Omega_{H,K}$$

and the decomposition of the flat section of $P_\varepsilon(K,H)$ to:

$$\int_{\sigma_{K,H}} \Omega_{K,H} = \int_{g^-} \Omega_{H,K} + \mathcal{L}(H).$$

Putting this together, we have:

$$\begin{aligned}
\int_{\sigma_{H,K}} \Omega_{H,K} + \int_{\sigma_{K,H}} \Omega_{K,H} &= \mathcal{L}(H) + \left(\mathcal{L}(H) + \int_{g^+} \Omega_{H,K} + \int_{g^-} \Omega_{H,K}\right) \\
&= \mathcal{L}(H) + \mathcal{L}(K) = \frac{\mathrm{vol}(P_{H,H}) + \mathrm{vol}(P_{K,K})}{\mathrm{vol}(M)} \\
&= \frac{\mathrm{vol}(P_{H,K}) + \mathrm{vol}(P_{K,H})}{\mathrm{vol}(M)} = \mathrm{area}(P_{H,K}) + \mathrm{area}(P_{K,H}).
\end{aligned}$$

□

As a consequence of that proposition and Proposition 3.5, we then have:

Corollary 3.8 *In the aspherical case, the path $\phi^H_{t\in[0,1]}$ is minimal amongst all paths with same endpoints if*

$$\int_{\sigma_{H,K}} \Omega_{H,K} \geq \mathcal{L}(H) \quad \textit{and} \quad \int_{\sigma_{K,H}} \Omega_{K,H} \geq \mathcal{L}(H)$$

for all $\varepsilon > 0$.

3.2 End of the proof

Proposition 3.9 *If $H_{t\in[0,1]}, K_{t\in[0,1]}$ induce paths with same endpoints, σ is any section class of $P = P_{H,K}$, and $p_1 \in \mathcal{A}_+$ and $p_2 \in \mathcal{A}_-$ are two marked points in P, then*

$$\Phi_{P,\sigma}\colon\ QH(M_{p_1}) \longrightarrow QH(M_{p_2})$$

is an isomorphism. The same statement holds for $P = P_{K,H}$.

Proof This is a consequence of Corollary 2.13. Indeed, compose $P_{H,K}$ with $P_{K,H}$. We then get, inside $M \times S^2$, the graphs of two Hamiltonian paths that are the inverse of each other. Isotope the union of the two graphs to the graph of the trivial loop in $\mathrm{Ham}(M,\omega)$. Finally choose the section of $P_{H,K}$ so that its connected union with σ is the flat section. □

The next proposition is the key result of this section:

Proposition 3.10 *Let H_t, K_t be two Hamiltonian paths with same endpoints on a symplectic manifold. Assume either that they are homotopic rel endpoints or that M is aspherical. If a pair $(x(t), v_x)$ is essential in the Floer homology of H_t, then there is a solution u in $P_\varepsilon^\pm(K)$ to the equation $(**)$, which converges to $\phi_t^K(x(0))$ in the obvious class v_K.*

Proof To fix notations, consider $P_\varepsilon^+(K)$. If the statement in the proposition did not hold, the isomorphism

$$\Phi_\sigma\colon\ QH_*(M_{p_1}) \to QH_*(M_{p_2})$$

in the bundle $P = P_\varepsilon(H, K)$ over S^2 would factorise by the Composition Theorem 2.10 through

$$QH_*(M_{p_1}) \xrightarrow{\Phi'_{\sigma'_{flat}}} FH_*(\partial P_\varepsilon^-(H)) \xrightarrow{\Phi} FH_*(\partial P_\varepsilon^+(K)) \xrightarrow{\Phi''_{\sigma''_{flat}}} QH_*(M_{p_2})$$

where it would be enough to consider a subcomplex (L, ∂) of $FC_*(\partial P_\varepsilon^-(H) \subset P_\varepsilon^-(H))$ that does not contain $(x(t), v_x)$. Thus Φ_σ would factorise through $H_*(L)$. Since both the domain and codomain of Φ_σ can be canonically identified with $HF_*(H_t)$, this would mean that the inclusion of L in $FC_*(H_t)$ would induce an isomorphism, a contradiction with the hypothesis of essentiality. $\square$

We proved the last proposition using the Composition Theorem. It is clear that, unwrapping the proof of the Composition Theorem, the last proposition is actually established by a "stretch-the-neck" argument.

End of the proof of Theorem 3.1 We will apply the previous proposition to the pairs (x_{min}, v_{cst}) and (x_{max}, v_{cst}) (where v_{cst} is the constant map), which are essential in the Floer homology of H_t. By Proposition 3.5 and Corollary 3.8, it is enough to show:

(1) the symplectic area of the solution u_{max} of the preceding proposition to the equation $(**)$ in $P_\varepsilon^+(K)$ that converges to $\phi_t^K(x_{max})$ in class v_K is non-negative, and

(2) the symplectic area of the solution u_{min} of the preceding proposition to the equation $(**)$ in $P_\varepsilon^-(K)$ that converges to $\phi_t^K(x_{min})$ in class v_K is non-negative.

To prove (1), consider the s–energy of u_{max}:

$$0 \le \int_\Sigma \left\| \frac{\partial u_{max}}{\partial s} \right\|^2 = \int_\Sigma \left\langle \frac{\partial u_{max}}{\partial s}, -J\frac{\partial u}{\partial t} - \nabla F \right\rangle$$
$$\le \Omega_{P_\varepsilon^+(K)} \operatorname{area}(u_{max}) - \operatorname{Totvar}(f)$$

Here, on the end (identified with the semi-infinite cylinder), $\frac{\partial}{\partial s}$ and $\frac{\partial}{\partial t}$ are the standard vector fields while, inside the unit disc, they are given as rdr and $d\theta$ respectively. Because $\mathrm{Totvar}(f)$ is non-negative, we get:

$$\Omega_{P_\varepsilon^+(K)} \operatorname{area}(u_{max}) \geq 0.$$

A similar argument applies to (x_{min}, v_{cst}). □

4 Higher-genus norms

As an addendum to this paper, we briefly discuss a semi-norm of higher genus on the group of Hamiltonian diffeomorphisms and state a conjecture on its relation with the Hofer norm.

Let $(M,\omega) \hookrightarrow P \to \Sigma_{g,1}$ be a Hamiltonian fibration over a compact oriented surface obtained by removing an open disc from a closed surface of genus g. Let q be a base point on $S = \partial\Sigma$, and $\phi \in \mathrm{Diff}_{\mathrm{Ham}}(M,\omega)$ the monodromy of the characteristic foliation on $W = \partial P$.

Definition 4.1 Let $\phi \in \mathrm{Diff}_{\mathrm{Ham}}(M,\omega)$ be given. Its *genus g norm*, $\|\phi\|_g$ is by definition the infimum, over all Hamiltonian fibrations $(M,\omega) \hookrightarrow P \to \Sigma_{g,1}$ with monodromy equal to ϕ, of the area of P.

Let $\mathcal{R}$ be the semigroup of all non-increasing sequences of non-negative real numbers, which contain only finitely non-zero terms, with the operation:

$$(a_0, a_1, a_2, \ldots) +_T (b_0, b_1, b_2, \ldots) = (c_0, c_1, c_2, \ldots)$$

where the c_ks are defined by:

$$c_k = \min_{0 \leq j \leq k} (a_j + b_{k-j}).$$

This operation is associative and commutative. Because the sequences are non-increasing, the zero sequence is a neutral element. Defining the *total norm* of a Hamiltonian diffeomorphism as $\|\phi\|_T = (\|\phi\|_0, \|\phi\|_1, \|\phi\|_2, \ldots)$, one can show easily:

Proposition 4.2 *The total norm $\|\cdot\|_T \in \mathcal{R}$ satisfies the triangle inequality with respect to the total sum defined above, ie*

$$\|\phi \circ \psi\|_T \leq \|\phi\|_T +_T \|\psi\|_T.$$

Proposition 4.3 *Let $\phi \in \mathrm{Diff}_{\mathrm{Ham}}(M,\omega)$, different from the identity. Then*

$$\|\phi\|_g = 0$$

if and only if $g \leq \ell_c$.

Proof To simplify notation, we will first prove this when $\ell_c = 1$. In this case, $\phi = [f,g]$. Endow $([0,\varepsilon]\times[0,\varepsilon])\times M$ with the split symplectic structure, identify $[0,\varepsilon]\times\{0\}\times M$ to $[0,\varepsilon]\times\{\varepsilon\}\times M$ via the map f to get a Hamiltonian fibration over a cylinder. Identify $\{0\}\times[\varepsilon/3, 2\varepsilon/3]\times M$ with $\{\varepsilon\}\times[\varepsilon/3, 2\varepsilon/3]\times M$ via the map g. This gives a Hamiltonian fibration over the punctured 2–torus whose monodromy is $[f,g]$. Its area is arbitrarily small, thus $\|\phi\|_1 = 0$. The same argument shows that if ϕ can be written as a product of k commutators, one can construct a Hamiltonian fibration $P \to \Sigma_{k,1}$ of arbitrarily small area whose monodromy round the boundary is ϕ. Thus $\|\phi\|_k = 0$. Now, to prove that all other semi-norms $\|\phi\|_{k'}$ also vanish for $k' > k$, one can add as many small handles as one wants over an arbitrarily small disc where the above fibration $P \to \Sigma_{k,1}$ is trivialised.

The converse statement saying that $\|\phi\|_k \neq 0$ if ϕ has commutator length larger than k is a consequence of the non-degeneracy of the usual Hofer norm and will be proved elsewhere (see [10]). □

Because $\mathrm{Diff}_{\mathrm{Ham}}(M,\omega)$ is a simple group, the group generated by products of commutators must be the whole of $\mathrm{Diff}_{\mathrm{Ham}}(M,\omega)$. Thus any element $\phi \in \mathrm{Diff}_{\mathrm{Ham}}(M,\omega)$ can be written in the form:

$$\phi = [\phi_1, \phi_1'] \circ \ldots \circ [\phi_k, \phi_k'].$$

Definition 4.4 The least such integer k is called the *commutator length* of ϕ, $\ell_c(\phi)$.

Let $\mathrm{Diff}^k_{\mathrm{Ham}}(M)$ be the subspace of all diffeomorphisms with commutator length less or equal to k. The following conjecture is studied in [10].

Conjecture 4.5 Let $\phi \in \mathrm{Diff}_{\mathrm{Ham}}(M)$. Then the distance, in the usual Hofer norm, between ϕ and $\mathrm{Diff}^k_{\mathrm{Ham}}(M)$ is equal to $\|\phi\|_k$.

The last Proposition and Conjecture are closely related to Entov's paper on the relation between K–area and commutator length. Indeed, Proposition 4.3 is the analogue of Entov's result that the *size* of P (the inverse of our g–norm) is larger or equal to the K–area of P. However, in Entov's paper, the Hofer norm is taken to be $\max_M |H|$ with H normalised with ω–integral zero, and it is not obvious to pass from results concerning one norm to results concerning the other.

References

[1] **Y Eliashberg**, **A Givental**, **H Hofer**, *Introduction to Symplectic Field Theory*, from: "GAFA 2000 (Tel Aviv, 1999)", Geom. Funct. Anal. Special volume, part II (2000) 560–673

[2] **M Entov**, *K–area, Hofer metric and geometry of conjugacy classes in Lie groups*, Invent. Math. 146 (2001) 93–141

[3] **M Entov**, *Commutator length of symplectomorphisms*, Comm. Math. Helv. 79 (2004) 58–104

[4] **D Gatien**, **F Lalonde**, *Holomorphic cylinders with Lagrange boundary conditions and Hamiltonian dynamics*, Duke Math. J. 102 (2000) 485–511

[5] **E Kerman**, **F Lalonde**, *Length minimising paths for symplectically aspherical manifolds*, Ann. Inst. Fourier (Grenoble) 53 (2003) 1503–1526

[6] **F Lalonde**, **D McDuff**, *Hofer's L^∞–geometry: energy and stability of Hamiltonian flows, part II*, Invent. Math. 122 (1995) 34–69

[7] **F Lalonde**, **D McDuff**, *Local non-squeezing theorems and stability*, Geom. Funct. Anal. 5 (1995) 364–386

[8] **F Lalonde**, **D McDuff**, *Symplectic structures on fiber bundles*, Topology 42 (2003) 309–347

[9] **F Lalonde**, **D McDuff**, **L Polterovich**, *Topological rigidity of Hamiltonian loops and quantum homology*, Invent. Math. 135 (1999) 369–385

[10] **P Larchevêque**, in preparation

[11] **D McDuff**, *Quantum homology of fibrations over S^2*, Internat. J. Math. 11 (2000) 665–721

[12] **D McDuff**, **D Salamon**, *Introduction to Symplectic Topology*, Oxford Mathematical Monographs, Oxford University Press (2000)

[13] **D McDuff**, **D Salamon**, *J–holomorphic curves and Quantum cohomology*, AMS University Lecture Series 6, Amer. Math. Soc. Providence, RI (2004)

[14] **Y-G Oh**, *Chain level Floer theory and Hofer's geometry of the Hamiltonian diffeomorphism group*, Asian J. Math. 6 (2002) 579–624

[15] **Y-G Oh**, *Erratum to "Chain level Floer theory and Hofer's geometry of the Hamiltonian diffeomorphism group"*, to appear in Asian J. Math.

[16] **Y-G Oh**, *Spectral invariants and length minimizing property of Hamiltonian paths*, `arXiv:math.SG/0212337`

[17] **P Piunikhin**, **D Salamon**, **M Schwarz**, *Symplectic Floer–Donaldson theory and Quantum Cohomology*, from: "Contact and Symplectic Geometry (Cambridge, 1994)", (C Thomas, editor), Cambridge University Press, Cambridge (1996) 171–200

[18] **L Polterovich**, *The geometry of the group of symplectic diffeomorphisms*, Lectures Math. ETH Zürich, Birkhäuser, Basel (2001)

[19] **Y Ruan**, **G Tian**, *A mathematical theory of quantum cohomology*, J. Diff. Geom. 42 (1995) 259–367

[20] **D Salamon**, *Lectures on Floer theory*, from: "Symplectic geometry and topology (Park City, UT, 1997)", (Y Eliashberg and L Traynor, editors), IAS/Park City Math. Series 7, Amer. Math. Soc. Providence, RI (1999) 143–230

[21] **M Schwarz**, *On the action spectrum for closed symplectically aspherical manifolds*, Pacific J. Math. 193 (2000) 419–461

[22] **P Seidel**, π_1 *of symplectic automorphism groups and invertibles in quantum cohomology rings*, Geom. Funct. Anal. 7 (1997) 1046–1095