

Symplectic Floer homology of area-preserving surface diffeomorphisms

ANDREW COTTON-CLAY

The symplectic Floer homology $HF_*(\phi)$ of a symplectomorphism $\phi: \Sigma \rightarrow \Sigma$ encodes data about the fixed points of ϕ using counts of holomorphic cylinders in $\mathbb{R} \times M_\phi$, where M_ϕ is the mapping torus of ϕ . We give an algorithm to compute $HF_*(\phi)$ for ϕ a surface symplectomorphism in a pseudo-Anosov or reducible mapping class, completing the computation of Seidel’s $HF_*(h)$ for h any orientation-preserving mapping class.

53D40, 37J10

1 Introduction

1.1 Statement of results

Let Σ be a compact, connected, oriented surface. Let ω be a symplectic form (here, an area form) on Σ . Let $\phi \in \text{Symp}(\Sigma, \omega)$ be a symplectomorphism (here, an area-preserving diffeomorphism) of Σ with nondegenerate fixed points. That is, the fixed points of ϕ are cut out transversally in the sense that $\det(1 - d\phi_x) \neq 0$ for any fixed point x . We then consider the symplectic Floer homology chain complex

$$CF_*(\phi) := \mathbb{Z}/2 \langle x \mid x \in \text{Fix}(\phi) \rangle,$$

a $\mathbb{Z}/2$ -graded chain complex over $\mathbb{Z}/2$. The grading of a generator x is given by the sign of $\det(1 - d\phi_x)$. Note that a fixed point x of ϕ corresponds to a constant section of the mapping torus M_ϕ , the Σ -bundle over S^1 with monodromy ϕ . The matrix coefficient $\langle \partial x, y \rangle$ of the differential ∂ on $CF_*(\phi)$ is then given by the mod 2 count of the index one components of the moduli space of holomorphic cylinders in $\mathbb{R} \times M_\phi$, with respect to a generic \mathbb{R} -invariant almost complex structure, which limit to the sections corresponding to x and y at $\pm\infty$, assuming these counts are finite. See Section 2 for details.

The standard condition, which in this form is due to Seidel [45], used to ensure that these counts are finite is *monotonicity*. Let ω_ϕ be the two-form on M_ϕ induced by ω on $\Sigma \times \mathbb{R}$ and let c_ϕ be the first Chern class of the vertical tangent bundle of $M_\phi \rightarrow S^1$.

Definition 1.1 A map $\phi: \Sigma \rightarrow \Sigma$ is *monotone* if $[\omega_\phi] = kc_\phi$ as elements of $H^2(M_\phi; \mathbb{R})$ for some $k \in \mathbb{R}$.

Because c_ϕ controls the index, or expected dimension, of moduli spaces of holomorphic curves under change of homology class and ω_ϕ controls their energy under change of homology class, this condition ensures that the energy is constant on the index one components of the moduli space, which implies compactness. See [Section 2.1](#) and [Section 2.2](#) for details.

The mapping class group Γ_Σ is given by $\pi_0(\text{Diff}^+(\Sigma))$, the connected components of the group of oriented diffeomorphisms (or, equivalently, homeomorphisms) of Σ . For $h \in \Gamma_\Sigma$, let $\text{Symp}_h(\Sigma)$ denote the space of symplectomorphisms in the mapping class h and $\text{Symp}_h^m(\Sigma)$ denote the space of monotone symplectomorphisms in the mapping class h . Seidel [\[45\]](#) has shown that the inclusions

$$\text{Symp}_h^m(\Sigma) \hookrightarrow \text{Symp}_h(\Sigma) \hookrightarrow \text{Diff}_h^+(\Sigma)$$

are homotopy equivalences; in particular, $\text{Symp}_h^m(\Sigma)$ is path connected. Furthermore, Seidel showed that $HF_*(\phi)$ is invariant as ϕ is deformed through monotone symplectomorphisms. These imply that we have a symplectic Floer homology invariant $HF_*(h)$ canonically assigned to each mapping class h given by $HF_*(\phi)$ for any monotone symplectomorphism.¹

In this paper, we give an algorithm to compute $HF_*(h)$ in all previously unknown cases on surfaces of negative Euler characteristic.² In addition, we define and give an algorithm to compute $HF_*(h)$ for all mapping classes on surfaces with boundary (see [Remark 1.2](#) for our conventions for mapping classes on surfaces with boundary). We furthermore compute the $H_*(\Sigma)$ -module structure. In a forthcoming paper [\[7\]](#), we extend the results of this paper to an algorithm to compute $HF_*(\phi)$ over a Novikov ring for any $\phi \in \text{Symp}_h(\Sigma)$ and use this to give a sharp lower bound on the number of fixed points of an area-preserving map with nondegenerate fixed points in a given mapping class, generalizing the Poincaré–Birkhoff fixed point theorem.

Thurston’s classification of surface diffeomorphisms [\[49; 12\]](#) states that every mapping class of Σ is precisely one of the following:

- Periodic (finite order): For some representative ϕ , we have $\phi^\ell = \text{id}$ for some $\ell \in \mathbb{Z}_{>0}$.
- Pseudo-Anosov: See [Section 3.2](#).

¹Furthermore, when Σ has negative Euler characteristic, these three groups, and in particular $\text{Symp}_h^m(\Sigma)$, are contractible, so this assignment really is canonical.

²The genus one case follows from Pozniak’s thesis [\[39, Sections 3.5.1–2\]](#).

- Reducible: Some representative ϕ fixes setwise a collection of curves C none of which are nullhomotopic or boundary parallel (and the mapping class is not periodic).

In the reducible case, cutting along a maximal collection of curves C gives a map on each component of $\Sigma \setminus C$ (given by the smallest power of ϕ which maps that component to itself) which is periodic or pseudo-Anosov.

Remark 1.2 When Σ has boundary, we instead consider the connected components of $\text{Diff}_\partial^+(\Sigma)$, the group of orientation-preserving diffeomorphisms of Σ with no fixed points on the boundary. We then additionally consider maps with full Dehn twists at the boundary to be reducible, and periodic and pseudo-Anosov maps with boundary come with the additional data of which direction each setwise fixed boundary component rotates. See [Section 4.2](#) for more details.

We briefly summarize previous work and which cases of $HF_*(h)$ have been computed. Dostoglou and Salamon [8] extended Floer's work [15] on symplectomorphisms Hamiltonian isotopic to the identity to other symplectomorphisms. Seidel [43] was the first to consider symplectic Floer homology specifically for surface symplectomorphisms, calculating it for arbitrary compositions of Dehn twists along a disjoint collection of curves. These are examples of reducible mapping classes, where the reducing collection of curves is the stated collection. Gautschi [17] calculated the symplectic Floer homology for all periodic mapping classes as well as reducible mapping classes in which the map on each component is periodic.

Eftekhary [9] generalized Seidel's work on Dehn twists in a different direction. Given two disjoint forests C_+ and C_- of embedded, homologically essential curves,³ he showed that, for ϕ any composition of single positive Dehn twists along each curve in C_+ and single negative Dehn twists along each curve in C_- , we have $HF_*([\phi]) = H_{*(\text{mod } 2)}(\Sigma \setminus C_-, C_+)$. He also showed that some of these maps are pseudo-Anosov.

We determine the symplectic Floer homology for any pseudo-Anosov mapping class and for any reducible mapping class. In particular, we allow pseudo-Anosov components as well as periodic components.

For pseudo-Anosov mapping classes, we deform the canonical singular representative ϕ_{sing} of the mapping class (see [Section 3.2](#)) near its singularities to a smooth, symplectic representative with nondegenerate fixed points ϕ_{sm} . For reducible mapping classes,

³That is, any two curves in C_+ intersect at most once, and there are no cycles in the graph whose vertices are curves in C_+ and whose edges are intersections, and similarly for C_- . In addition, no curve in C_+ intersects any curve in C_- .

we give a similar perturbation of the boundary behavior of the canonical singular representative of pseudo-Anosov maps in [Section 4.2](#) and combine this with standard representatives of maps on periodic components and of Dehn twists along the reducing curves to give standard-form-maps ϕ_{sm} for reducible mapping classes.

We give a notion of *weak monotonicity* satisfied by ϕ_{sm} such that $HF_*(\phi)$ is well-defined for and invariant among weakly monotone maps. Monotonicity implies weak monotonicity, and so $HF_*(h) = HF_*(\phi)$ for any weakly monotone ϕ . Furthermore we show that the $H_*(\Sigma)$ -module structure of $HF_*(\phi)$ is invariant among weakly monotone maps.

Definition 1.3 A map $\phi: \Sigma \rightarrow \Sigma$ is *weakly monotone* if the class $[\omega_\phi]$ vanishes on $\ker(c_\phi|_{T(M_\phi)})$, where $T(M_\phi) \subset H_2(M_\phi; \mathbb{R})$ is generated by tori T such that $\pi|_T: T \rightarrow S^1$ is a fibration with fiber S^1 , where the map $\pi: M_\phi \rightarrow S^1$ is the projection.

This brings us to our main results.

Theorem 1.4 For h a pseudo-Anosov mapping class, we have:

- (1) Every symplectomorphism $\phi: (\Sigma, \omega) \rightarrow (\Sigma, \omega)$ in the mapping class h is weakly monotone. Hence if ϕ has nondegenerate fixed points, then $HF_*(\phi)$ is well-defined and $HF_*(\phi) = HF_*(h) = HF_*(\phi_{\text{sm}})$.
- (2) Moreover, $HF_*(\phi_{\text{sm}}) = CF_*(\phi_{\text{sm}})$.
- (3) Only the fundamental class of $H_*(\Sigma)$ (the identity), acts nontrivially.

Remark 1.5 Furthermore, we show in [Section 5](#) that $CF_*(\phi_{\text{sm}})$ is combinatorially computable from an invariant train track representation. In particular, see [Theorem 5.1](#). An invariant train track for a mapping class can be found algorithmically from, for example, a representation of the mapping class as a composition of Dehn twists by work of Bestvina and Handel [\[1\]](#).

Theorem 1.6 For h a reducible mapping class, we have:

- (1) Any standard form map is weakly monotone.
- (2) $HF_*(h)$ splits into summands for each component⁴ of $\Sigma \setminus C$ (and annular region – see [Section 4.2](#)), with the periodic and annular components as in [Gautschi \[17\]](#)

⁴But note that the summand for a given component may be affected by the direction it is required to rotate near its boundary.

and the pseudo-Anosov components as in [Theorem 1.4](#), except when a pseudo-Anosov component with an unrotated boundary component (see [Section 4.2](#)) abuts a component on which the map is the identity, in which case there may be explicitly given additional flow lines between these two components. See [Theorem 4.16](#) for details and a complete statement and computation.

- (3) $H_*(\Sigma)$ acts by intersection products on components on which the map is the identity and which do not meet pseudo-Anosov components, and only the fundamental class (the identity) acts nontrivially on other components.

Corollary 1.7 *The total rank of $HF_*(h^n)$*

$$\begin{cases} \text{is bounded} & \text{if } \phi \text{ is periodic,} \\ \text{grows linearly} & \text{if } \phi \text{ is reducible with all components periodic,} \\ \text{grows exponentially} & \text{otherwise.} \end{cases}$$

In [Section 6](#) we explicitly work out $HF_*(h)$ for h a large collection of pseudo-Anosov mapping classes based on a construction of Penner [\[35\]](#) using [Theorem 1.4](#) and the train track computation in [Theorem 5.1](#):

Theorem 1.8 *Let $\{C_i\}$ and $\{D_j\}$ be collections of noncontractible simple closed curves on Σ with the C_i mutually disjoint, the D_j mutually disjoint, and C_i transverse to D_j . Suppose that every region of the complement $\Sigma - \bigsqcup_i C_i - \bigsqcup_j D_j$ is a disk with more than two edges.*

Let $\phi = \prod_j \tau_{D_j}^{-2} \prod_i \tau_{C_i}^2$ and let h be the mapping class of ϕ . Then $HF_(\phi) = HF_*(h)$ has total rank*

$$4 \sum_{i,j} |C_i \cap D_j|^2 + 2g - 2.$$

1.2 Relationships with other Floer theories

Symplectic Floer homology for surface symplectomorphisms is the $d = 1$ part of periodic Floer homology⁵ (PFH) of ϕ (see Hutchings and Sullivan [\[21, Section 2\]](#) and Hutchings [\[18\]](#)), a Floer homology theory whose chain complex is generated by certain multisets of periodic orbits of ϕ and whose differential counts certain embedded pseudoholomorphic curves in $\mathbb{R} \times M_\phi$. This theory was created to be a candidate for a 3-dimensional version of Taubes’ “SW = Gr” result [\[46\]](#). Since the first draft of this paper came out, Lee and Taubes [\[30\]](#) have shown this to be isomorphic to the

⁵Beware, however, that for some of the spin^c -structures, the monotonicity condition for PFH differs from that for symplectic Floer homology.

Seiberg–Witten Floer homology (SWFH) (see Kronheimer and Mrowka [25]) of M_ϕ in certain spin^c -structures when $d < g - 1$. This in turn is conjectured to agree with Ozsváth and Szabó's Heegaard Floer homology invariant HF^+ [34].

PFH is also related to embedded contact homology (ECH) (see Hutchings and Sullivan [22]), which is an analogous construction with $\mathbb{R} \times Y$ for Y a contact 3-manifold and is also conjectured to agree with SWFH and HF^+ . Taubes has recently announced a result relating ECH to SWFH [47; 48] and yet more recently Lee and Taubes have announced a result relating PFH to SWFH. This would imply that the results of this paper calculate the SWFH of M_ϕ in certain spin^c -structures.

PFH has been computed for a single Dehn twist and certain other cases [21], but no pseudo-Anosov cases are known for $g > 1$. ECH has been computed for T^3 [22] and for M_ϕ for ϕ an Anosov map on the torus (ie a map given by a hyperbolic element of $\text{SL}_2\mathbb{Z}$) by Lebow [26]. Jabuka and Mark [23] have computed $HF^+(M_\phi)$ for a subset of the mapping classes calculated by Seidel [43] and Eftekhary [9] for symplectic Floer homology, and the results agreed.

1.3 Relationship with vanishing cycles and mutation

Seidel has a conjecture for the symplectic Floer homology of a general composition of positive Dehn twists on a surface with boundary in terms of Hochschild homology coming from his work on vanishing cycles and mutation [44]. Due to the complicated relationship between Thurston's classification of surface diffeomorphisms, which our results are based on, and compositions of positive Dehn twists, comparing Seidel's conjecture with the results of this paper is nontrivial. Since the first draft of this paper, Perutz [37] has announced a proof, in progress, of Seidel's conjecture. It would be interesting to find applications of the equality of these two very different computations. It may also be interesting to understand this equality directly, perhaps by understanding how the Floer complex changes as the Bestvina–Handel algorithm is implemented.

1.4 Organization of the paper

In Section 2 we review symplectic Floer homology and give a detailed discussion of when it is well-defined. We introduce Nielsen classes and the direct sum decomposition of (CF_*, ∂) resulting therefrom. We use a simplified version of Yi-Jen Lee's bifurcation analysis to show that, in our setting, $HF_*(\phi)$ is invariant under deformations through weakly monotone maps. Finally we show invariance of the $H_*(\Sigma)$ -module structure using an algebraic interpretation of the $H_1(\Sigma)$ -action in terms of the twisted Floer homology complex.

In [Section 3.1](#) we show that every map in a pseudo-Anosov mapping class is weakly monotone.

In [Section 3.2](#), we describe the canonical singular representative of a pseudo-Anosov mapping class ϕ_{sing} in the closed case and give a Hamiltonian perturbation supported near its singularities which results in a smooth symplectomorphism ϕ_{sm} . In addition, we analyze the fixed points of ϕ_{sm} near the perturbed singularities and note that they are all of the same index.

In [Section 3.3](#), we use a result of Birman and Kidwell [2] and the conclusions of [Section 3.2](#) to show that all differentials vanish on $CF_*(\phi_{\text{sm}})$ and deduce the $H_*(\Sigma)$ -module structure.

In [Section 4.1](#), we study the topology of the space of weakly monotone maps in arbitrary mapping classes in order to prove a stronger invariance result for HF_* .

In [Section 4.2](#), we describe the standard form for reducible maps and discuss the behavior of pseudo-Anosov maps near the boundary. We give a Hamiltonian perturbation of ϕ_{sing} to ϕ_{sm} at the boundary. We then give in [Section 4.3](#) a classification of the Nielsen classes of fixed points of standard form reducible maps after Jiang and Guo [24], extending their work to the symplectic case, where one must also consider fixed regions between parallel Dehn twists (which can be isotoped away if one is not required to preserve area). We then use this to show in [Section 4.4](#) that standard form reducible maps ϕ_{sm} are weakly monotone in every Nielsen class, which requires a lengthy technical argument. This allows us to show that $HF_*(\phi_{\text{sm}}) = HF_*(h)$.

In [Section 4.5](#), we use a neck-stretching argument of Gautschi [17] in [Lemma 4.15](#) to show that $(CF_*(\phi_{\text{sm}}), \partial)$ splits into direct summands for each component (including regions between parallel Dehn twists) with one exception. This exception occurs when a pseudo-Anosov component abuts a component on which the map is the identity. We adapt Gautschi's neck stretching argument to this case and combine this with an argument of Salamon and Zehnder [41] in [Lemma 4.14](#) to compute the differential by relating holomorphic cylinders from fixed points in one of these regions to the other to Morse flow lines passing through the reducing curve between the regions. This allows us to give a computation in [Theorem 4.16](#) of $HF_*(\phi_{\text{sm}})$, together with its $H_*(\Sigma)$ -module structure, using our result for pseudo-Anosov components and Gautschi's [17] work on periodic components.

In [Section 5](#), we review the theory of train tracks and give a combinatorial formula computing $HF_*(\phi_{\text{sm}})$ for pseudo-Anosov mapping classes (and thus for reducible mapping classes as well, by the direct sum decomposition given in [Section 4.5](#)) from the action of ϕ_{sm} on an invariant train track. Note that Bestvina and Handel [1] give

an algorithm to find an invariant train track from eg the representation of a mapping class as a product of Dehn twists.

In Section 6, we apply the results of Section 5 to a large collection of examples of pseudo-Anosov maps due to Penner and obtain an explicit formula for the rank of $HF_*(h)$ in these cases.

1.5 Directions for further research

Cylindrical contact homology Symplectic field theory and contact homology invariants (see Eliashberg, Givental and Hofer [10]) exist for mapping tori due to the existence of a Hamiltonian structure (see Bourgeois et al [3], Cieliebak and Mohnke [5] and Fabert [11]). In this setting, the “cylindrical mapping torus contact homology” splits as a direct sum

$$\bigoplus_{n=1}^{\infty} H_*(C_*(\phi, n)),$$

where $C_*(\phi, n)$ is generated by good periodic orbits of degree n . This chain complex is similar to $CF_*(\phi^n)$, but the latter is generated instead by periodic *points*, with multiplicity, of degree n . This, and its more complicated “full” version, are a sort of intermediate step between symplectic Floer homology and (cylindrical) contact homology of open books. Mei-Lin Yau [51] has given a correspondence between (not necessarily cylindrical) holomorphic curves in mapping tori and in their 0–surgeries, which are contact 3–manifolds with a natural open book decomposition. Since the first draft of this paper came out, Colin and Honda [6] have computed certain examples of cylindrical contact homology for open books.

Product structures We have an understanding of the $H_*(\Sigma)$ –module structure, but lack an understanding of products $*$: $HF(\phi) \otimes HF(\psi) \rightarrow HF(\phi \circ \psi)$ unless one of the two maps is the identity. An important special case is an understanding of rings of the form

$$\bigoplus_{n=0}^{\infty} HF_*(\phi^n) \quad \text{or} \quad \bigoplus_{n=-\infty}^{\infty} HF_*(\phi^n).$$

Grading and spin^c–structures We use a $\mathbb{Z}/2$ –grading, but a relative $\mathbb{Z}/(2\text{div}(c_\phi)) = \mathbb{Z}/(4g - 4)$ grading is available for each Nielsen class⁶ or more generally for each homological Nielsen class, that is, among generators whose orbits agree in $H_1(M_\phi)$. A good understanding of the splitting not into Nielsen classes but into homological Nielsen classes, together with relative grading data, is lacking.

⁶In our calculations, we have more than one generator in a given Nielsen class only in situations in which we have a correspondence between Floer flow lines and Morse flow lines. Thus the relative grading for Nielsen classes is simply the Morse relative grading.

Acknowledgments I'd like to thank my advisor Michael Hutchings for his invaluable guidance and support and David Farris, Eli Lebow and Lee Mosher for helpful discussions. I'd also like to thank an anonymous referee for very thorough and insightful comments.

2 Floer homology, weak monotonicity, Nielsen classes, bifurcation analysis and module structure

2.1 Review of symplectic Floer homology

Let Σ be a compact, connected, oriented surface of negative Euler characteristic. Let ω be a symplectic form (ie an area form) on Σ . Let $\phi \in \text{Symp}(\Sigma, \omega)$ be a symplectomorphism (ie an area-preserving diffeomorphism) of Σ and consider

$$M_\phi = \frac{\mathbb{R} \times \Sigma}{(t + 1, x) \sim (t, \phi(x))},$$

the mapping torus of ϕ . Note that this is a Σ -bundle over S^1 and we have the projection $\pi: M_\phi \rightarrow \mathbb{R}/\mathbb{Z} = S^1$.

Suppose ϕ has nondegenerate fixed points; that is, at a fixed point, $d\phi$ does not have 1 as an eigenvalue (in particular, fixed points are isolated). Under certain conditions (see below), we can define the symplectic Floer homology $HF_*(\phi)$ associated to ϕ . This is the homology of a $\mathbb{Z}/2$ -graded⁷ chain complex $CF_*(\phi)$ (which depends on the choice of almost complex structure; see below) which is freely generated as a vector space over $\mathbb{Z}/2$ by the fixed points of ϕ . The grading $\epsilon(x)$ of a fixed point x is given by $(-1)^{\epsilon(x)} = \text{sign}(\det(1 - d\phi))$. The homology, but not necessarily the chain complex itself, is invariant under certain deformations of ϕ (and independent of the choice of almost complex structure). We give a summary of how to define $HF_*(\phi)$ and point out any relevant pitfalls. See also Dostoglou and Salamon [8] and Seidel [45].

Let $\Gamma(M_\phi)$ denote the space of smooth sections of $\pi: M_\phi \rightarrow S^1$, also known as the twisted loop space,⁸ consisting of smooth paths $\gamma: \mathbb{R} \rightarrow \Sigma$ such that $\gamma(t) = \phi(\gamma(t+1))$. Note that a fixed point $x \in \Sigma$ of ϕ can be interpreted as a constant path $\gamma(t) = x$ in $\Gamma(M_\phi)$. The tangent space of $\Gamma(M_\phi)$ at a path γ consists of sections ξ of $\gamma^*T\Sigma$ such that $\xi(t) = d\phi(\xi(t+1))$. We have an action one-form α_ϕ on $\Gamma(M_\phi)$ defined by

$$\alpha_\phi(\gamma)\xi = \int_0^1 \omega\left(\frac{d\gamma}{dt}, \xi\right) dt.$$

⁷There is a lift to a relative $\mathbb{Z}/\text{div}(2c_1(E)) = \mathbb{Z}/(4g - 4)$ grading, with a corresponding absolute grading by homotopy classes of 2-plane fields; see Perutz [38] or Hutchings [20].

⁸We will abuse notation and switch between sections of M_ϕ and paths in Σ at will.

Note that the zeroes of this one-form are precisely the constant paths γ , that is, the paths corresponding to fixed points of ϕ .

The differential on $CF_*(\phi)$ is defined by a certain count of “gradient flow lines” of α_ϕ . To get a metric (which doesn’t show up explicitly in the equations), we require a complex structure on Σ . Let \mathcal{J} be the space of smooth 1-periodic families $J = \{J_t\}$ (that is, $\phi_* J_{t+1} \phi_*^{-1} = J_t$) of ω -tame⁹ almost complex structures on Σ . Then we can define $\mathcal{M}(x, y, J)$ to be the set of maps $u: \mathbb{R}^2 \rightarrow \Sigma$ such that:

- (1)
$$u(s, t) = \phi(u(s, t + 1))$$
- (2)
$$\lim_{s \rightarrow +\infty} u(s, \cdot) = x \quad \text{and} \quad \lim_{s \rightarrow -\infty} u(s, \cdot) = y$$
- (3)
$$\partial_s u + J_t(u) \partial_t u = 0.$$

We call such a u a *flow line* from x to y . We may think of u as a section $u: \mathbb{R} \times S^1 \rightarrow \mathbb{R} \times M_\phi$ of $\text{id} \times \pi: \mathbb{R} \times M_\phi \rightarrow \mathbb{R} \times S^1$. We may also think of u as a path in $\Gamma(M_\phi)$ between γ_x and γ_y , the constant paths associated to the fixed points x and y . We will use these three interpretations interchangeably.

At each flow line u , we have a Fredholm linearized $\bar{\partial}$ operator D_u . We let $\mathcal{M}_k(x, y, J)$ be the subset consisting of those maps u for which the index of D_u is k . For generic $J \in \mathcal{J}$, D_u is onto and so by the implicit function theorem, all of these are smooth manifolds of dimension k (see Floer, Hofer and Salamon [16]). Additionally, these moduli spaces come with a natural action of \mathbb{R} by translation in the first coordinate, and so we consider their quotients $\mathcal{M}_k(x, y, J)/\mathbb{R}$.

We would like to define the differential¹⁰ by

$$(4) \quad \partial x = \sum_y c_{x,y} y$$

where $c_{x,y}$ is a count of the number of points (modulo 2) in $\mathcal{M}_1(x, y, J)/\mathbb{R}$, and we would also like to have $\partial^2 = 0$ so that we can take the homology $HF_*(\phi)$. In order to do this, we need compactness and gluing results (note that bubbling is not an issue as $\pi_2(\Sigma) = 0$), which require upper bounds on the energy

$$(5) \quad E(u) = \frac{1}{2} \int_{\mathbb{R} \times [0,1]} (|\partial_s u|^2 + |\partial_t u|^2) ds dt$$

⁹That is, $\omega(v, Jv) > 0$ for $v \neq 0$.

¹⁰We use symplectic Floer *homology*, which agrees with the conventions in periodic Floer homology [18; 21].

in each $\mathcal{M}_k(x, y, J)$ (at least for $k = 1$ or 2). For $u \in \mathcal{M}_k(x, y, J)$, we have

$$(6) \quad E(u) = \int_{\mathbb{R} \times [0,1]} u^* \omega.$$

These are locally constant. To ensure they are bounded, we impose:¹¹

Condition 2.1 $E(u)$ is constant on $\mathcal{M}_k(x, y, J)$.

Given this, the symplectic Floer homology groups $HF_*(\phi)$ are well defined. Additionally, they are independent of the choice of (generic) J (which follows, for example, by [Theorem 2.8](#)). We have yet to consider invariance under (certain) deformations of the map ϕ . More on that in [Section 2.4](#).

2.2 Index and weak monotonicity

In order to verify [Condition 2.1](#), we need to know more about the index $k = \text{ind}(D_u)$. We follow [\[18, Sections 2.1–2.3, 5.1\]](#). We first define the relative first Chern class and the Conley–Zehnder index and then use these to state a formula for $\text{ind}(D_u)$.

Let E denote the vertical tangent bundle of $M_\phi \rightarrow S^1$. This bundle carries a symplectic structure induced by ω . If $\gamma_x \in \Gamma(M_\phi)$ corresponds to a fixed point x (ie $\gamma_x(t) = x$), let $\mathcal{T}(\gamma_x)$ denote the set of homotopy classes of symplectic trivializations of $\gamma_x^* E$. This is an affine space over \mathbb{Z} .

Suppose we have a flow line $u: \mathbb{R} \times S^1 \rightarrow \mathbb{R} \times M_\phi$ which limits to x at $-\infty$ and y at $+\infty$. If we choose trivializations $\tau_x \in \mathcal{T}(\gamma_x)$ and $\tau_y \in \mathcal{T}(\gamma_y)$, we get a *relative first Chern class* $c_1(u^* E, \tau_x, \tau_y) \in \mathbb{Z}$ by counting the zeroes of a section of $u^* E$ which is constant (and nonzero) with respect to τ_x and τ_y over the two boundary circles. If v is another flow line with the same limits, let $[u - v]$ be the homology class in $H_2(\mathbb{R} \times M_\phi) = H_2(M_\phi)$ corresponding to the obvious (up to reparametrization) map $S^1 \times S^1 \rightarrow \mathbb{R} \times M_\phi$. Then we have

$$(7) \quad c_1(u^* E, \tau_x, \tau_y) - c_1(v^* E, \tau_x, \tau_y) = \langle c_1(E), [u - v] \rangle.$$

We also have an integer called the *Conley–Zehnder index* $\mu(\gamma_x, \tau_x)$ associated to a fixed point x and a trivialization $\tau_x \in \mathcal{T}(\gamma_x)$. We have a natural connection on the bundle $M_\phi \rightarrow S^1$ whose monodromy is given by ϕ . Linearized parallel transport induces a connection on $\gamma_x^* E$. Thus if we traverse γ_x once, then using τ_x this induces a path of symplectic matrices in $\text{Sp}(2, \mathbb{R})$ from 1 to $d\phi_x$. We let $\mu(\gamma_x, \tau_x)$ be the

¹¹We could make do without this if we were willing to work over a Novikov ring. This is the approach we take in [\[7\]](#).

Maslov index of this path of symplectic matrices (see eg Salamon and Zehnder [41]). Furthermore, $\mu(\gamma_x, \tau_x) \equiv \epsilon_x$ (modulo 2).

Theorem 2.2 (Index [40]) *For u a flow line from x to y ,*

$$\text{ind}(D_u) = 2c_1(u^* E, \tau_x, \tau_y) + \mu(\gamma_x, \tau_x) - \mu(\gamma_y, \tau_y).$$

*In particular, this is independent of the choice of trivializations.*¹² □

Corollary 2.3 (Change of homology class) *If u and v are flow lines with the same limits, then*

$$\text{ind}(D_u) - \text{ind}(D_v) = 2 \langle c_1(E), [u - v] \rangle. \quad \square$$

Following Seidel [45], we now turn this into a way to verify [Condition 2.1](#). If u and v are two flow lines from x to y , then by following u and then the reverse of v , we get a loop in $\Gamma(M_\phi)$ based at γ_x , ie an element $\psi_{uv} \in \pi_1(\Gamma(M_\phi), \gamma_x)$. We may think of an element $\psi \in \pi_1(\Gamma(M_\phi), \gamma_x)$ as a map $S^1 \times S^1 \rightarrow S^1 \times M_\phi$ which is a section of $\text{id} \times \pi: S^1 \times M_\phi \rightarrow S^1 \times S^1$. On $S^1 \times M_\phi$ as well as on M_ϕ , we have two elements of $H^2(M_\phi; \mathbb{R})$ (we will abuse notation and consider these on $S^1 \times M_\phi$ or M_ϕ interchangeably). First we have ω_ϕ induced by ω on Σ , and second we have $c_\phi = c_1(E)$. Both of these may be evaluated on elements of $\pi_1(\Gamma(M_\phi), \gamma_x)$. We denote these actions by $\omega_\phi(\psi) = \int_{S^1 \times S^1} \psi^* \omega_\phi$ and $c_\phi(\psi) = c_1(\psi^* E)[S^1 \times S^1]$.

Corollary 2.4 (Monotonicity for x ; cf [45, Lemma 9]) *Suppose that for all $\psi \in \pi_1(\Gamma(M_\phi), \gamma_x)$ such that $c_\phi(\psi) = 0$, we also have $\omega_\phi(\psi) = 0$. Then $E(u)$ is constant on $\mathcal{M}_k(x, y, J)$ for all k, y and J . That is, [Condition 2.1](#) is satisfied for flow lines starting at x . In this case, we say that ϕ satisfies monotonicity for x .*

Proof Suppose we have u and v in $\mathcal{M}_k(x, y, J)$. Then we have $\text{ind}(D_u) - \text{ind}(D_v) = k - k = 0$, so by [Corollary 2.3](#), $c_\phi(\psi_{uv}) = \langle c_1(E), [u - v] \rangle = 0$. Thus, using $\omega_\phi(\psi_{uv}) = 0$, we have:

$$0 = \omega_\phi(\psi_{uv}) = \int_{S^1 \times S^1} \psi^* \omega_\phi = \int_{\mathbb{R} \times [0, 1]} u^* \omega - \int_{\mathbb{R} \times [0, 1]} v^* \omega = E(u) - E(v). \quad \square$$

Thus we've reduced [Condition 2.1](#) to:

Condition 2.5 (Weak monotonicity) ω_ϕ vanishes on the kernel of

$$c_\phi: H_1(\Gamma(M_\phi)) \rightarrow \mathbb{R}.$$

¹²We use the notation of Hutchings [18, Theorem 5.1], who states this theorem for general Riemann surfaces, where it is due to Schwarz [42, Theorem 3.3.11].

2.3 Nielsen classes

Given a map ϕ , there is a topological separation of fixed points due to Nielsen. The chain complex $CF_*(\phi)$ splits into direct summands corresponding to each Nielsen class (defined below). That is, if two fixed points x and y are in different Nielsen classes, then the matrix coefficient $c_{x,y} = \langle \partial x, y \rangle$ is zero. This idea has been used eg by Pozniak [39, Sections 3.3.1, 3.5.1] and Seidel [43] and has also been noticed by Nielsen theorists, eg Fel'shtyn [13].

Definition 2.6 Two fixed points x and y of ϕ are in the same *Nielsen class*, or *Nielsen equivalent*, if there exists a path $\gamma: [0, 1] \rightarrow \Sigma$ with $\gamma(0) = x$, $\gamma(1) = y$ such that $\phi(\gamma)$ is homotopic to γ (rel. endpoints). Note that this is an equivalence relation.

There can be a flow line from x to y only if x and y are in the same Nielsen class: If $u(s, t)$ is a flow line from x to y , let the path $\gamma_t: [0, 1] \rightarrow \Sigma$ be a reparametrization of $u(\cdot, t)$ (for each t). Then the curve γ_1 exhibits the Nielsen equivalence between x and y : its limits are x and y , and it is homotopic (rel. endpoints) through γ_t to $\gamma_0 = \phi(\gamma_1)$.

Interpreting fixed points as sections of $M_\phi \rightarrow S^1$, that is, as elements of $\Gamma(M_\phi)$, we note that two fixed points are Nielsen equivalent if and only if their sections are homotopic. That is, $\pi_0(\Gamma(M_\phi))$ describes possible Nielsen classes.

For $\eta \in \pi_0(\Gamma(M_\phi))$, let $CF_*(\phi, \eta)$ be the chain complex generated by fixed points in Nielsen class corresponding to η . (Note that this may be the zero chain complex.) We have $CF_*(\phi) = \bigoplus_{\eta \in \pi_0(\Gamma(M_\phi))} CF_*(\phi, \eta)$ as chain complexes, and thus

$$HF_*(\phi) = \bigoplus_{\eta \in \pi_0(\Gamma(M_\phi))} HF_*(\phi, \eta).$$

Let $\Gamma(M_\phi)_\eta$ denote the η -component of $\Gamma(M_\phi)$. Then $HF_*(\phi, \eta)$ is well-defined if we have:

Condition 2.7 (η -Weak monotonicity) ω_ϕ vanishes on the kernel of

$$c_\phi: H_1(\Gamma(M_\phi)_\eta) \rightarrow \mathbb{R}.$$

2.4 Bifurcation analysis of Yi-Jen Lee

The standard argument to show invariance under certain deformations of ϕ is the continuation argument as in [15]. It gives invariance under Hamiltonian deformations

of ϕ . Seidel [45, Section 3] has a generalized continuation argument which shows invariance under deformations which preserve monotonicity. We have need of a more general version of monotonicity, and thus also require a more general invariance result.

Floer's original invariance argument [14] was not a continuation argument, but rather a bifurcation argument, in which one analyzes degeneracies that arise in generic one-parameter families and their effect on HF_* . He did not complete all of the analysis required, but with the advent of the continuation argument, this was no longer needed. Yi-Jen Lee [28; 29] has worked out a general bifurcation argument for what she calls Floer-type theories, including all of the necessary analysis. Michael Usher [50] has a nice summary of the invariance result this gives (which Lee conjectured in an earlier paper [27, Equation 3.2] but did not explicitly state as a Theorem in [28; 29]) and its algebraic aspects. We give a brief discussion of a simplified version sufficiently general for our needs.

Theorem 2.8 (Simplified version of [50, Theorem 3.6, due to Lee [28; 29]]) *Suppose (Σ, ω) is a symplectic manifold with $\pi_2(\Sigma) = 0$. Let $\phi_r: \Sigma \rightarrow \Sigma$ be a smooth family of symplectomorphisms and $J_r = \{J_t\}_r$ a smooth family (of 1-periodic families) of almost complex structures such that (ϕ_0, J_0) and (ϕ_1, J_1) are generic. Let $\eta \in \pi_0(\Gamma(M_{\phi_0})) \cong \pi_0(\Gamma(M_{\phi_r}))$ be a Nielsen class. Suppose that ω_{ϕ_r} vanishes on $\ker(c_\phi): \pi_1(\Gamma(M_{\phi_r}), \eta) \rightarrow \mathbb{R}$ (that is, ϕ_r is η -weakly monotone). Then $(CF_*(\phi_0, \eta, J_0), \partial_0)$ is chain homotopy equivalent to $(CF_*(\phi_1, \eta, J_1), \partial_1)$. \square*

The main idea of the proof is to study generic 1-parameter families and what happens to (CF_*, ∂) at times r_0 where the data fail to be generic. Because we have assumed that ω_ϕ vanishes on $\ker(c_\phi)$, we cannot have a flow line of index 0 from a fixed point x to itself (Lee calls these “type II handleslides”) because, by Corollary 2.3, it would have zero energy. The argument is then formally identical to bifurcation arguments in Morse theory; see for instance Hutchings [19] for a carefully done modern perspective in the more complicated circle-valued case (which includes type II handleslides). For Lee's account in the Floer theory setting, see [28, Section 4.4], noting that we have no type II handleslides.

In generic 1-parameter families, there are three types of degeneracies that can happen at a time r_0 :

- *Cancellation*: A degenerate flow line from x to y of index 1,
- *Birth-death*: A birth or death of two fixed points x and y at a degenerate fixed point,
- *Handleslide*: A flow line from x to y of index 0.

For a cancellation, at times $r_0 - \epsilon < r < r_0$ we have two flow lines (of opposite sign were we working with \mathbb{Z} -coefficients) which cancel at time r_0 (or the time-reversal of this situation). This does not change the chain complex (CF_*, ∂) .

For a birth-death, let us assume without loss of generality we have a death of a pair x and y , with a single flow line from x to y for r shortly before r_0 . Let z and w be two other fixed points. We compare $\langle \partial_{-z}, w \rangle$ and $\langle \partial_{+z}, w \rangle$, the matrix coefficients at times just before r_0 and at times just after r_0 , respectively. Because ω_ϕ vanishes on $\ker(c_\phi)$, we may have either flow lines of index 1 from z to y or flow lines of index 1 from y to z , but not both. Let us assume that the former is the case. Flow lines from z to y are glued to flow lines from x to w . Thus

$$\langle \partial_{+z}, w \rangle = \langle \partial_{-z}, w \rangle \pm \langle \partial_{-z}, y \rangle \langle \partial_{-x}, w \rangle.$$

This does not change the chain homotopy type of the chain complex (we need not worry about the sign as we are working with $\mathbb{Z}/2$ -coefficients, but the case with \mathbb{Z} -coefficients is fine as well); see eg Usher [50, Section 3.3].

For a handleslide, the generators of the chain complex do not change. We again compare ∂_- and ∂_+ . If we have a fixed point z with a flow line to x for r near r_0 , then, by gluing this with our flow line from x to y , we get a flow line from z to y on one side of r_0 . Similarly, if we have a fixed point w with a flow line from y for r near r_0 , then we get a flow line from x to w on one side of r_0 . Thus, letting $A: CF_* \rightarrow CF_*$ be given by $A(x) = x \pm y$ and $A(v) = v$ for all $v \neq x$, we have that $\partial_+ = A^{-1} \circ \partial_- \circ A$. Thus the isomorphism-type of (CF_*, ∂) does not change, and, in particular, the homology is unchanged.

The analysis required to make the above rigorous is the gluing described in the birth-death situation and in the handleslide situation. Lee does this in the second paper [29], where she must also glue tori due to type II handleslides.

2.5 Product structure and invariance

There is a product $*$: $HF(\phi) \otimes HF(\psi) \rightarrow HF(\phi \circ \psi)$, whose matrix coefficients count holomorphic sections of a hamiltonian symplectic fibration with fiber Σ over the pair of pants $\mathbb{C}P^1 \setminus \{0, 1, \infty\}$. The monodromies are given by ϕ , ψ and $\phi \circ \psi$.

In particular, this gives $HF(\phi)$ the structure of a module over the quantum homology $QH(\Sigma) = HF(\text{id})$ via $*$: $HF(\text{id}) \otimes HF(\phi) \rightarrow HF(\phi)$. By Liu and Tian [31], this intrinsic definition is equivalent to the following extrinsic definition: given $a \in HF(\text{id})$ and x a fixed point of ϕ , choose a cycle C representing $a \in HF(\text{id}) \cong H_*(\Sigma)$. Then the coefficient $\langle a * x, y \rangle$ for y a fixed point of ϕ is given by a count holomorphic

cylinders as in Section 2.1 with the extra condition that they pass through the cycle C at $s = t = 0$.

The action of the fundamental class on $HF(\phi)$ is the identity. The action of the point class has been studied by Seidel [45], who shows that it acts as zero unless the mapping class of ϕ is the identity mapping class. We calculate the action of $H_1(\Sigma)$ below. For this to have meaning, we have need of a proof of the invariance of the module structure under weakly monotone deformations, which we come to presently.

Proposition 2.9 *The action of $H_1(\Sigma)$ on $HF(\phi)$ is invariant under weakly monotone deformations of ϕ .*

Remark 2.10 This implies that the full module structure is invariant by the discussion in the previous paragraph, or by the observation that the action of $H_1(\Sigma)$ algebraically determines the full module structure. We expect the full module structure to be invariant for weakly monotone deformations of symplectomorphisms in higher dimensions, using more details of Lee's bifurcation analysis [28; 29].

Proof The key observation here is that the action of $H_1(\Sigma)$ has index one and thus, using the extrinsic definition, we may consider index one moduli spaces of cylinders, the same moduli spaces used in defining the differential, and then impose the constraint that they pass through C at $s = t = 0$ (whereas for the differential we mod out by the \mathbb{R} -action). We fix C missing the fixed points of ϕ . Then we have

$$(8) \quad \langle [C] * x, y \rangle = \sum_{p \in \mathcal{M}_1(x, y, J) / \mathbb{R}} [C] \cdot [p],$$

where $[p]$ is the relative homology class of the cylinder from x to y , in the $H_2(M_\phi)$ -torsor of 2-chains with boundary $\gamma_y - \gamma_x$, and $[C] \cdot [p]$ is the intersection product, where by $[C]$ we really mean $i_*[C]$ for $i: \Sigma \rightarrow M_\phi$ the inclusion to the fiber over $0 \in S^1$.

The extrinsic definition is equivalent to the following algebraic definition, as in [22, Section 12.1.3]. Let $H = H_2(M_\phi)$. We can define the fully twisted symplectic Floer homology $(CF^{\text{tw}}(\phi), \partial^{\text{tw}})$ as the free $\mathbb{Z}/2[H]$ -module generated by fixed points of ϕ , with differential given by

$$\langle \partial^{\text{tw}} x, y \rangle = \sum_{p \in \mathcal{M}_1(x, y, J) / \mathbb{R}} ([p] - b_{xy}),$$

where b_{xy} is a choice of an element in the H -torsor associated to x and y . We require the choices to be compatible in the sense that they form a cochain with respect

to the gluing maps and additionally require that $[C] \cdot b_{xy} = 0$ for all x, y . This is possible because we may assume that C is nonseparating (else $[C]$ is zero). The chain homotopy type of $(CF^{\text{tw}}(\phi), \partial^{\text{tw}})$ is invariant under weakly monotone deformations of ϕ . The proof is the same as that of [Theorem 2.8](#), with the additional comment that gluing is additive on relative homology classes, and so the same algebraic formulas given above for birth-deaths and handleslides carry over.¹³

Via the intersection product, the chain C gives a map $i_C: H \rightarrow \mathbb{Z}/2$. This induces a $\mathbb{Z}/2$ -linear map $\zeta_C: \mathbb{Z}/2[H] \rightarrow \mathbb{Z}/2[H]$ defined on elements $h \in H$ by $\zeta_C(h) = i_C(h)h$. Note that this is a derivation of $\mathbb{Z}/2[H]$ over $\mathbb{Z}/2$. We now define a $\mathbb{Z}/2[H]$ -linear map ∂_C^{tw} on $CF^{\text{tw}}(\phi)$ via

$$\partial_C^{\text{tw}}x = \sum_y \zeta_C(\langle \partial^{\text{tw}}x, y \rangle)y.$$

We have that $\partial^{\text{tw}} \circ \partial_C^{\text{tw}} = \partial_C^{\text{tw}} \circ \partial^{\text{tw}}$ (as noted in [\[22, Section 12.1.3\]](#), this follows from $(\partial^{\text{tw}})^2 = 0$ and $\zeta_C(xy) = \zeta_C(x)y + x\zeta_C(y)$). Setting all elements of H equal to one, we have ∂_C on $CF(\phi)$ which commutes with ∂ . This gives a map $\partial_C: HF(\phi) \rightarrow HF(\phi)$ which agrees with the map $[C] * -: HF(\phi) \rightarrow HF(\phi)$ by [Equation \(8\)](#). To complete the proof of the invariance of the latter under weakly monotone deformations, we have only to show that the former is invariant under chain homotopies of $CF^{\text{tw}}(\phi)$.

Consider a $\mathbb{Z}/2[H]$ -linear chain map $\Phi: CF^{\text{tw}}(\phi) \rightarrow CF^{\text{tw}}(\phi')$ (we're only interested in the case where this is a quasi-isomorphism, but it does not matter). Note that via the inclusion $i: CF \rightarrow CF^{\text{tw}}$ and the map $j: CF^{\text{tw}} \rightarrow CF$ sending all elements of H to one, we have a chain map $j \circ \Phi \circ i: CF(\phi) \rightarrow CF(\phi')$. We show $\Phi \partial_C^{\text{tw}}$ and $\partial_C^{\text{tw}} \Phi$ are chain homotopic, which implies $j \Phi i \partial_C$ and $\partial_C j \Phi i$ are chain homotopic. We compute:

$$\begin{aligned} \langle \Phi \partial_C^{\text{tw}}x_i, x_k \rangle &= \langle \Phi \langle \partial_C^{\text{tw}}x_i, x_j \rangle x_j, x_k \rangle \\ &= \langle \partial_C^{\text{tw}}x_i, x_j \rangle \langle \Phi x_j, x_k \rangle \\ &= \zeta_C(\langle \partial^{\text{tw}}x_i, x_j \rangle) \langle \Phi x_j, x_k \rangle \\ &= \zeta_C(\langle \partial^{\text{tw}}x_i, x_j \rangle \langle \Phi x_j, x_k \rangle) - \langle \partial^{\text{tw}}x_i, x_j \rangle \zeta_C(\langle \Phi x_j, x_k \rangle) \end{aligned}$$

Similarly,

$$\langle \partial_C^{\text{tw}} \Phi x_i, x_k \rangle = \zeta_C(\langle \Phi x_i, x_j \rangle \langle \partial^{\text{tw}}x_j, x_k \rangle) - \zeta_C(\langle \Phi x_i, x_j \rangle) \langle \partial^{\text{tw}}x_j, x_k \rangle.$$

The first terms on the right hand sides are equal as $\partial^{\text{tw}} \Phi = \Phi \partial^{\text{tw}}$. The others are of the form $\partial^{\text{tw}}K$ and $K \partial^{\text{tw}}$, so indeed the two left hand sides are chain homotopic. \square

¹³For the handleslide, we actually must use $A(x) = x \pm (h - b_{xy})y$, where h is the class of the index zero flow line. For the birth-death, we simply choose b_{xy} to be the class of the short flow line between x and y and the same formula carries over.

3 Pseudo-Anosov maps

3.1 Weak monotonicity and invariance in pseudo-Anosov mapping classes

The aim of this section is to show:

Proposition 3.1 *For any (not necessarily monotone) symplectomorphism $\phi: (\Sigma, \omega) \rightarrow (\Sigma, \omega)$ in a pseudo-Anosov mapping class h , if ϕ has nondegenerate fixed points, then $HF_*(\phi)$ is well-defined and $HF_*(\phi) = HF_*(h)$.*

In order to prove this we show that every symplectomorphism ϕ in a pseudo-Anosov mapping class is weakly monotone (ie satisfies [Condition 2.5](#)). To do this, we show that $H_1(\Gamma(M_\phi)) = 0$.

Lemma 3.2 *For ϕ pseudo-Anosov and $\gamma \in \Gamma(M_\phi)$, $\pi_1(\Gamma(M_\phi), \gamma) = 0$.*

Proof We consider γ as a map $\mathbb{R} \rightarrow \Sigma$ with $\gamma(t) = \phi(\gamma(t + 1))$. An element of $\pi_1(\Gamma(M_\phi), \gamma)$ is of the form $\gamma_s(t)$ for $s \in S^1 = \mathbb{R}/\mathbb{Z}$ with $\gamma_0(t) = \gamma(t)$. We consider $\alpha_0(s) = \gamma_s(0)$ and $\alpha_1(s) = \gamma_s(1)$. These are closed curves on Σ and $\phi(\alpha_1(s)) = \alpha_0(s)$. Furthermore, $\alpha_0(s)$ is homotopic to $\alpha_1(s)$ by the homotopy $\alpha_t(s) = \gamma_s(t)$. We now quote [Lemma 3.3](#) (below) to conclude that $\alpha_0(s)$ must be nullhomotopic.

To deal with this case, we consider the fibration:

$$\begin{array}{ccc} \Omega\Sigma & \longrightarrow & \Gamma(M_\phi) \\ & & \downarrow f \\ & & \Sigma \end{array}$$

Here f takes a section to its value at $t = 0$. The fiber over $\gamma(0) \in \Sigma$ is homeomorphic to $\Omega\Sigma$, which we think of as loops in Σ based at $\gamma(0)$, as follows: Let $\psi_t: \Sigma \rightarrow \Sigma$ be an isotopy from the identity to a map which takes $\gamma(0)$ to $\gamma(1)$. Then, given an element $x(t)$ of $\Omega\Sigma$ (ie a loop based at $\gamma(0)$), we consider $y(t) = \psi_t(x(t))$, an element of the fiber over $\gamma(0)$ of the map f . This gives the homeomorphism.

Thus we have a long exact sequence of homotopy groups, a piece of which is

$$\pi_2(\Sigma, \gamma(0)) \rightarrow \pi_1(\Gamma(M_\phi), \gamma) \rightarrow \pi_1(\Sigma, \gamma(0)).$$

The image in $\pi_1(\Sigma, \gamma(0))$ of the element of $\pi_1(\Gamma(M_\phi), \gamma)$ represented by the homotopy $\gamma_s(t)$ is represented by $\alpha_0(s)$, which is nullhomotopic. Thus it comes from $\pi_2(\Sigma, \gamma(0))$, which is trivial. Thus we have $\pi_1(\Gamma(M_\phi), \gamma) = 0$. □

Lemma 3.3 [4, Lemma 5.6] *If $\phi: \Sigma \rightarrow \Sigma$ is pseudo-Anosov, and C is a noncontractible curve on Σ , then $\phi^m(C) \simeq \phi^n(C)$ if and only if $m = n$. \square*

Proof of Proposition 3.1 By Lemma 3.2, Condition 2.5 is trivially satisfied. Thus $HF_*(\phi)$ is well-defined. To show that $HF_*(\phi) = HF_*(h)$, we use Theorem 2.8. We find an isotopy of symplectic maps ϕ_t such that $\phi_0 = \phi$ and ϕ_1 is monotone. This exists because $\text{Symp}^m \hookrightarrow \text{Symp}$ is a homotopy equivalence [45]. By Lemma 3.2, each ϕ_t satisfies the conditions of Theorem 2.8. \square

3.2 Symplectic representatives of pseudo-Anosov mapping classes

Pseudo-Anosov mapping classes have canonical representatives (up to topological conjugacy) which happen to be singular. The goal of this section is to describe how to symplectically smooth these singular representatives. We review the structure of the singular representatives and then spend the rest of the section symplectically smoothing their singularities.

A singular measured foliation \mathcal{F} on a surface Σ is, away from a finite number of points, a smooth foliation with a transverse measure, that is, a measure of how far a curve is travelling in the direction transverse to the foliation. Away from the singular points the foliation has charts which take the leaves to horizontal lines in \mathbb{R}^2 such that the transverse distance between leaves is simply the vertical distance in \mathbb{R}^2 . At a singular point, \mathcal{F} has a number of prongs $p \geq 3$ as in the left of Figure 1 for the case $p = 3$ (more on this below).

A pseudo-Anosov map ϕ_{sing} (that is, a canonical singular representative) preserves two measured foliations \mathcal{F} and \mathcal{G} which are transverse everywhere except their singular points, which coincide and have the same number of prongs (as in the left of Figure 1). The map ϕ_{sing} expands the transverse measure on \mathcal{F} by a factor of $\lambda > 1$ and contracts the transverse measure on \mathcal{G} by the same factor. See Fathi, Laudenbach and Poenaru [12], Casson and Bleiler [4], Thurston [49], Mosher [33], Jiang and Guo [24] and Penner [35] for more on pseudo-Anosov maps.

As in [12, Exposé 9], we can find a Markov partition for ϕ_{sing} . That is, we can decompose our surface into a finite collection of homeomorphic images of rectangles, overlapping only on their boundaries, such that \mathcal{F} and \mathcal{G} pull back to the standard horizontal and vertical foliations. We equip these rectangles with the standard flat metric coming from the transverse measures. Upon gluing together the rectangles, we may obtain cone points of angle $p\pi$, which we call singular points with p prongs.

We change coordinates by a homeomorphism isotopic to the identity such that the foliations are smooth away from the singular points. Furthermore we require that in a

neighborhood of a singular point x with p prongs, we have a smooth chart such that one of our foliations \mathcal{F} is given by the subbundle of the tangent space to Σ on which the quadratic differential $z^{p-2}dz^2$ is positive real, and the other \mathcal{G} by the subbundle on which it is negative real. See the left of [Figure 1](#).

We choose our symplectic form to be the area form associated with the flat-metric-with-cone-points described above. This is smooth even at the singularities. We must find a smooth symplectic perturbation of the map ϕ near the singularities. We describe the perturbation near a fixed singularity. Near a singularity that maps to a different singularity, one can use the work below under an identification of neighborhoods of the singularities.




Figure 1: Left: The foliations \mathcal{F} and \mathcal{G} in the case $p = 3$. Right: The singular Hamiltonian vector field whose time-1 flow is ϕ when $\rho = \text{id}$ in the case $p = 3$.

If we divide up a neighborhood of this singular point by considering the components of the complement of the prongs of \mathcal{F} , each of these components is diffeomorphic to a rectangle on which the two foliations are the horizontal and vertical foliations. This diffeomorphism is given in radial coordinates by $(r, \theta) \mapsto (r, \frac{p}{2}\theta)$ (see [Figure 2](#)). This map multiplies area by the constant multiple $\frac{p}{2}$ and thus a map is symplectic on one of the components if and only if the associated map is symplectic on the associated rectangle.

If x is a fixed point, then ϕ cyclically permutes these rectangles. If we compose ϕ with an appropriate rotation ρ (rotation by some multiple of $2\pi/p$) so that each rectangle is taken to itself, then the map $\rho \circ \phi = \phi \circ \rho$ on the closure of one of these rectangles is given in coordinates by the linear map

$$A_\lambda = \begin{pmatrix} \lambda & 0 \\ 0 & \lambda^{-1} \end{pmatrix},$$




Figure 2: Angular dilation of one of the regions transforms it to a standard rectangle with horizontal and vertical foliations.

where the origin is the singularity. The map A_λ is the time-1 flow of the Hamiltonian vector field associated to the Hamiltonian $H = \mu xy$, where $\mu = \ln \lambda$.

We can obtain a singular Hamiltonian H_{sing} (singular at x only) on an entire neighborhood of x by piecing together the Hamiltonians on the various components. The Hamiltonians agree on the prongs, and the only issue is smoothness there. To see that we actually have a smooth Hamiltonian away from x , consider that if we use \mathcal{G} instead of \mathcal{F} and find Hamiltonians in the manner described above, they agree on the half-rectangles (given by the overlaps) with the Hamiltonians coming from \mathcal{F} and they are smooth on the prongs associated to \mathcal{F} . Alternatively, we can simply note that $H_{\text{sing}} = \mu r^2 \cos(p\theta)$ in polar coordinates, or $\mu \operatorname{Re}(z^p)/|z|^{p-2}$ in complex coordinates. See the right of Figure 1.

Our strategy now is to modify H_{sing} inside a small disk. We divide our task into two cases:

- $\rho \neq \text{id}$ (the “rotated” case)
- $\rho = \text{id}$ (the “unrotated” case)

Perturbing rotated singularities

Our map ϕ is given (in a neighborhood of x) by the time-1 flow of the singular Hamiltonian H_{sing} composed with ρ^{-1} . We modify ϕ by forming $H_{\text{sm}} = fH_{\text{sing}}$, where f is a smooth function whose zero set is a small ball and which is one outside a slightly larger one. Then we let ϕ_{sm} (in a neighborhood of x) be the time-1 flow of the Hamiltonian H_{sm} composed with ρ^{-1} .

The union of all of the prongs is still preserved setwise by ϕ_{sm} and the rectangles continue to be permuted in the same manner as before. Thus the only possible fixed point is at x , where we have an elliptic fixed point, for ϕ_{sm} is given simply by ρ^{-1} in a neighborhood of x .

Perturbing unrotated singularities

Here ϕ is simply given (in a neighborhood of x) by the time-1 flow of H_{sing} . We must do something more elaborate than in the previous section in order to avoid degenerate fixed points.

We construct H_{sm} in two steps. First we construct it on a small ball, where it has $p - 1$ nondegenerate saddle points, and then we join it to H_{sing} by an interpolation with no critical points. See Figure 3 for the basic idea.




Figure 3: Smoothing H_{sing} in the case $p = 3$. On the left we have the level sets of H_{sing} and on the right we have the level sets of the perturbed Hamiltonian H_{sm} with $p - 1 = 2$ saddle points.

Lemma 3.4 *There is a smooth function G on the closed disk $D = D(1)$ with $p - 1$ nondegenerate saddle points and no other critical points such that on the boundary of the disk, $G|_{\partial D}$ is Morse¹⁴ with p minima and p maxima. Additionally, consider the one-manifold C_{max} consisting of the points where $G|_{\partial D(r)}$ is maximum for r near 1. Likewise consider C_{min} . We require that the derivative of G along C_{max} (pointing outward) is positive, that the derivative along C_{min} is negative and that C_{max} and C_{min} are transverse to $\partial D(r)$ where they are defined.*

¹⁴We won't be doing Morse theory here; this is simply a convenient class of functions.

Proof The function $x \sin(\pi y)$ has a saddle point at $(0, k)$ for $k \in \mathbb{Z}$. We consider a small rectangle $[-\epsilon, \epsilon] \times [0, p - 2]$ with semicircular caps on each end. This is not smooth, but we may smooth it with a C^1 -small perturbation near the points where it fails to be C^2 . We claim the map on this domain B has the requisite properties (we then choose an appropriate map from $D(1)$ to B).

On B , except near the ends, the maxima, minima and zeroes near the boundary are (for an appropriate choice of the subdisks $D(r)$) horizontal lines and the map on the boundary has nondegenerate zeroes and extrema.

The C^1 -small perturbation from above can't produce extra critical points of $G|_{\partial D(r)}$. Additionally, for ϵ small enough, $x \sin(\pi y)$ is a C^2 -small perturbation of πxy , which has nondegenerate extrema and satisfies the derivative condition in the statement of the lemma, so $x \sin(\pi y)$ also has nondegenerate extrema. Finally, the map is seen to have the correct number of zeroes and extrema. □

Lemma 3.5 *Let G be as in Lemma 3.4. Then there exist constants $0 < c_1 \ll 1$ and $c_2 \gg 1$ such that if we define $H_{\text{sm}} = c_1 G(c_2 z)$ on a small disk of radius $1/c_2$, there exists an extension to a function H_{sm} which equals H_{sing} outside a somewhat larger disk and which has no critical points in the intermediate region.*

Proof We proceed in three steps (we will rescale at the end):

(1) Select r_1 slightly less than 1 and consider the disk $D(r_1) \subset D(1)$. We modify G so that the zeroes and extrema are in the correct places on $D(1) - D(r_1)$ (in the sense that their placement agrees with that of $H_{\text{sing}} = \mu r^2 \cos(p\theta)$).

Isotop the identity map on the disk (through maps which preserve $D(r)$ setwise) to one which takes the zeroes of $G|_{D(1)-D(r_1)}$ to the segments $\theta = (k/p)\pi/2$, the maxima of G on the boundary of disks intermediate to D_1 and D_2 to the segments $\theta = (2k/p)\pi$ and the minima to the segments $\theta = ((2k + 1)/p)\pi$. Note that these θ correspond with the placement of the zeroes, maxima and minima of the map $H_{\text{sing}} = \mu r^2 \cos(p\theta)$.

Let G_1 be the composition of the inverse of the result of this isotopy with G . Note that G_1 also satisfies the conditions of Lemma 3.4

(2) Now select r_2 such that $r_1 < r_2 < 1$ and consider the disk $D(r_2)$. We modify G_1 so that it is $f(r) \cos(p\theta)$ for some increasing function $f(r)$ on the region $D(1) - D(r_2)$.

To do this, select an increasing function $f(r)$ on $D(1)$ such that on $D(1) - D(r_1)$,

$$|f(r) \cos(p\theta)| \geq K|G_1|$$

for some large constant K .

Let G_2 be a modification of G_1 on $D(1) - D(r_1)$ which is increasing in absolute value (except where it is zero) on the rays $\theta = \text{const}$ as r increases from r_1 to 1 except possibly for a brief period initially (ie near r_1) and such that $G_2 = f(r) \cos(p\theta)$ on $D(1) - D(r_2)$. We claim that we do not create additional critical points. Because the maxima, minima and zeroes have the desired behavior in [Lemma 3.4](#), there are only issues in between them. We can make the initial period where G_2 is not increasing short enough so as not to create new critical points of G_2 in this region.

(3) Extend G_2 to a function G_3 on \mathbb{R}^2 by extending $f(r)$ to an increasing function which equals Cr^2 for $r \gg 0$ and some large constant C . Finally, let H_{sm} equal $c_1 G_3(c_2 z)$ for appropriate constants c_1, c_2 (defined on an appropriate region). \square

Thus our map ϕ_{sm} has $p - 1$ positive hyperbolic fixed points, and no others, in the perturbed region. There are no fixed points coming from closed orbits because there are no components of level sets of H_{sm} which are circles, for then there would necessarily be an extremum inside, and H_{sm} has none.

3.3 Floer homology of ϕ_{sm} for pseudo-Anosov maps

We now consider the symplectic Floer homology of the smooth representative ϕ_{sm} we have just constructed. By [Proposition 3.1](#), because the mapping class h of ϕ_{sm} is pseudo-Anosov, $HF_*(\phi_{\text{sm}}) = HF_*(h)$, which in turn is equal to $HF_*(\phi)$ for any other ϕ in the same mapping class.

Theorem 3.6 *In $CF_*(\phi_{\text{sm}})$, all differentials vanish. Thus $HF_*(h) = HF_*(\phi_{\text{sm}}) = CF_*(\phi_{\text{sm}})$.*

Proof In [\[2\]](#) it is shown that no two fixed points (this includes both singular and nonsingular fixed points) of a singular, standard form pseudo-Anosov map ϕ_{sing} are Nielsen equivalent. If we can show that two fixed points of ϕ_{sm} are Nielsen equivalent if and only if they are both associated to the same singularity of ϕ_{sing} , then we will be done, for a differential u gives a Nielsen equivalence as in [Section 2.3](#). Thus there are no differentials except between those associated to the same singularity. These, however, are all of the same index (mod 2) by [Section 3.2](#).

To see that it is indeed the case that two fixed points of ϕ_{sm} are Nielsen equivalent if and only if both are associated to the same singularity of ϕ_{sing} , we argue as follows:

(1) Whether two fixed points x and y of a map ψ are Nielsen equivalent is unchanged by modifying a map inside a disk away from either fixed point (and thus any number

of disks away from the fixed points), for any path passing through a disk is homotopic (rel boundary) to any other with the same endpoints.

(2) If two isolated fixed points x and y of a map ψ are Nielsen inequivalent and we modify our map in a small disk near x to a map ψ' , then any fixed point x' (of the modified map) inside this disk is Nielsen inequivalent to y , for if we take a path γ' exhibiting Nielsen equivalence of x' and y (under ψ') and append a short path from x to x' , then this new path γ exhibits Nielsen equivalence of x and y (under our ψ).

To see this, note that the path following γ from x to x' and then onward until it exits (for the last time) our small disk at a point x'' (call this path α) stays inside a slightly larger disk under ψ . Thus the homotopy from x'' to $\psi(x'') = \psi'(x'')$ given by the homotopy from γ' to $\psi'(\gamma')$ extends to a homotopy rel x from α to $\psi(\alpha)$. This implies that γ is homotopic to $\psi(\gamma)$.

Repeated application of these two points gives our conclusion. □

Corollary 3.7 $H_1(\Sigma)$ acts as zero on $HF_*(h)$ for h pseudo-Anosov.

Proof There are no flow lines between generators of $CF_*(\phi_{sm})$, so the extrinsic or algebraic definition of the action of $H_1(\Sigma)$ as in Section 2.5 is trivial. By Proposition 2.9, the same holds for $HF_*(h)$. □

4 Reducible maps

4.1 The space of η -weakly monotone maps

In the reducible case, not every map is weakly monotone. We need to understand the structure of the space of η -weakly monotone maps, and in particular show that it is path connected, in order to prove invariance.¹⁵ We first recall the definition:

Definition 4.1 The symplectomorphism $\phi: (\Sigma, \omega) \rightarrow (\Sigma, \omega)$ is *weakly monotone* for a Nielsen class $\eta \in \pi_0(\Gamma(M_\phi))$ (or η -weakly monotone) if $\omega_\phi|_N = 0$, where

$$N = \ker(c_\phi: H_1(\Gamma(M_\phi)_\eta) \rightarrow \mathbb{R}),$$

where $\Gamma(M_\phi)_\eta$ denotes the η -component of $\Gamma(M_\phi)$.

¹⁵The results of this section go through for weakly monotone maps as well with similar arguments, which is sufficient for our needs, but involving Nielsen classes is not much more difficult.

We would like to be able to define the space $\text{Symp}_h^\eta(\Sigma, \omega)$ of maps ϕ in a mapping class $h \in \Gamma$ which are η -weakly monotone. To do this we must show that Nielsen classes are well-defined not only for a single map but also for an entire mapping class. Suppose we have an isotopy ϕ_t from $\phi = \phi_0$ to ϕ_1 . Thinking of a Nielsen class as a homotopy class of maps $f: [0, 1] \rightarrow \Sigma$ with the restriction that $f(1) = \phi(f(0))$, we have a natural way to extend a given f (up to reparametrization) to a representative of a Nielsen class for ϕ_t by extending the path f from $f(0)$ to $\phi(f(0))$ to a path f_t from $f(0)$ to $\phi_t(f(0))$ in the obvious manner. We must show that if $\phi_1 = \phi$ then the Nielsen class of f_1 is the same as the Nielsen class of f .

Lemma 4.2 *Nielsen classes are well-defined on an entire mapping class if Σ has negative Euler characteristic. That is, in the above setup, the Nielsen class of f_1 is the same as the Nielsen class of f .*

Proof We consider the universal cover \mathbb{H} of Σ . We have a map $\Phi: \Sigma \times S^1 \rightarrow \Sigma$ given by $(x, t) \mapsto \phi_t(x)$. This extends to a map $\tilde{\Phi}: \mathbb{H} \times [0, 1] \subset \mathbb{H} \times \mathbb{R} \rightarrow \mathbb{H}$ on the universal covers. If $\tilde{\Phi}(\cdot, 0) = \tilde{\Phi}(\cdot, 1)$, then the result follows. That is, we must show the map $\Phi \circ p: \mathbb{H} \times S^1 \rightarrow \Sigma$ extends to a map to \mathbb{H} (where $p: \mathbb{H} \times S^1 \rightarrow \Sigma \times S^1$ is the covering map).

Recall the homotopy lifting property, which states that this map extends if and only if $(\Phi \circ p)_* \pi_1(\mathbb{H} \times S^1) = 0 \subset \pi_1(\Sigma)$. To see that this is the case, consider that $\pi_1(\Sigma \times S^1) = \pi_1(\Sigma) \times \mathbb{Z}$. The image that we're interested in is the image of the copy of \mathbb{Z} under the map $\Phi_*: \pi_1(\Sigma) \times \mathbb{Z} \rightarrow \pi_1(\Sigma)$. Because ϕ is a diffeomorphism, $\Phi_*(\cdot, 0): \pi_1(\Sigma) \rightarrow \pi_1(\Sigma)$ is an isomorphism. Thus the image of the copy of \mathbb{Z} is in the center of $\pi_1(\Sigma)$ because it commutes with every element of the image of $\Phi_*(\cdot, 0)$, which is all of $\pi_1(\Sigma)$.

Because $\chi(\Sigma) < 0$, the center of $\pi_1(\Sigma)$ is trivial,¹⁶ and so the image of the copy of \mathbb{Z} is indeed trivial. The result follows. \square

Thus we denote the space of maps in a mapping class $h \in \Gamma$ which are weakly monotone for η by $\text{Symp}_h^\eta(\Sigma, \omega)$. Similarly we use $\text{Symp}_h^m(\Sigma, \omega)$ for the monotone symplectic maps in the mapping class h and $\text{Symp}_h(\Sigma, \omega)$ for all symplectic maps in the mapping class h . We note that if ϕ is monotone, then it is η -weakly monotone (for any η): the constant k is simply the proportionality constant relating ω and $c_\phi = c_1(E)$.

¹⁶To see this, note that if $\pi_1(\Sigma)$ has a center, then we get a copy of $\mathbb{Z}/k \times \mathbb{Z}/\ell$ acting freely and properly discontinuously by deck transformations (and, in particular, hyperbolic isometries) on \mathbb{H} . The free hyperbolic isometries (ie the parabolic and hyperbolic ones) are infinite order, so $k = \ell = 0$. Thus the quotient is a 2-manifold with fundamental group $\mathbb{Z} \times \mathbb{Z}$. The only one of these is the torus. This gives a hyperbolic structure on the torus, which violates Gauss–Bonnet.

Proposition 4.3 (cf [45, Lemma 6 ff]) *The inclusions*

$$\text{Symp}_h^m(\Sigma, \omega) \hookrightarrow \text{Symp}_h^\eta(\Sigma, \omega) \hookrightarrow \text{Symp}_h(\Sigma, \omega)$$

are homotopy equivalences.

Proof For the duration of this proof, the Σ and ω will be implied and we will omit them when they are not needed.

Fix $\phi \in \text{Symp}_h^m$. We have the action $\text{Symp}_0 \times \text{Symp}_h \rightarrow \text{Symp}_h$. Using ϕ , this gives a homeomorphism $\text{Symp}_0 \times \{\phi\} \rightarrow \text{Symp}_h$. We ask which elements of Symp_0 correspond under this homeomorphism to elements of Symp_h^η .

Consider $\psi \in H_1(\Gamma(M_\phi)_\eta)$. We have a map $H_1(\Gamma(M_\phi)) \rightarrow H_2(M_\phi; \mathbb{Z})$ given by taking the image in $H_2(M_\phi; \mathbb{Z})$ of the fundamental class of the torus representing the element of $H_1(\Gamma(M_\phi))$. Denote the image of ψ by $A_\psi \in H_2(M_\phi; \mathbb{Z})$. Then if ϕ is η -weakly monotone, we have that $c_\phi(A_\psi) = 0$ implies $\omega_\phi(A_\psi) = 0$.

Suppose we perturb the map ϕ by an isotopy ϕ_t from $\phi = \phi_0$ to ϕ_1 . There is then a natural way to extend ψ (up to reparametrization of S^1): ψ is (as an element of $\Gamma(M_\phi)$) a path from $\gamma = \psi(0, \cdot)$ to $\phi^{-1}(\gamma)$. We let ψ_t be the extension of this to a path from γ to $\phi_t^{-1}(\gamma)$ in the obvious manner. We wish to understand how $\omega_{\phi_t}(A_{\psi_t})$ and $c_{\phi_t}(A_{\psi_t})$ vary with t .

The latter is constant: we are simply taking the Euler numbers of isomorphic bundles.

For the former, we briefly recall the Flux homomorphism (see McDuff and Salamon [32, Section 10.2] for details). This is a map $\text{Flux}: \widetilde{\text{Symp}}_0(\Sigma, \omega) \rightarrow H^1(\Sigma; \mathbb{R})$ (where $\widetilde{\text{Symp}}_0(\Sigma, \omega)$ is the universal cover of the identity component $\text{Symp}_0(\Sigma, \omega)$, ie paths in $\text{Symp}(\Sigma, \omega)$ starting at the identity up to homotopy). Its value, when paired with an element $\gamma \in H_1(\Sigma, \mathbb{Z})$ represented by the smooth image of an S^1 , is the area swept out by γ under the path of symplectomorphisms. In our situation with $\chi(\Sigma) < 0$, we actually have $\text{Flux}: \text{Symp}_0(\Sigma, \omega) \rightarrow H^1(\Sigma; \mathbb{R})$ because maps $S^1 \times S^1 \rightarrow \Sigma$ are nullhomologous, and so the area swept out by these is zero.

Thus the difference between $\omega_{\phi_0}(A_{\psi_0})$ and $\omega_{\phi_1}(A_{\psi_1})$ is

$$\text{Flux}(\{\phi_t^{-1} \circ \phi_0\}_{t \in [0,1]}[\phi_0^{-1}(\gamma)]),$$

ie the area swept out (on Σ) by $\phi_t^{-1}(\gamma)$ as t goes from 0 to 1. By the last line of the previous paragraph, we may simply write this as $\text{Flux}(\phi_1^{-1} \circ \phi_0)[\phi_0^{-1}(\gamma)]$.

Let N be the subgroup of $H_1(\Sigma; \mathbb{R})$ generated by $[\phi^{-1}(\gamma)]$ for γ the smooth image of an S^1 such that there exists a $\psi \in H_1(\Gamma(M_\phi)_\eta)$ with $\psi(0, \cdot) = \gamma$. Note that for each such γ , the class A_ψ is well defined, for if we have two different ψ 's, putting

them together we get a map $S^1 \times S^1 \rightarrow \Sigma$, which then must be nullhomologous. Let $N_0 \subset N$ be generated by $[\phi^{-1}(\gamma)]$ for those γ which additionally satisfy $c_\phi(A_\psi) = 0$.

We get that $\phi_1 \in \text{Symp}_h^\eta$ if and only if

$$\text{Flux}(\phi_1^{-1} \circ \phi)(\phi_*^{-1}(\cdot))|_{N_0} = 0.$$

That is, Symp_h^η is the kernel of the antihomomorphism

$$\text{Symp}_h \xrightarrow{\cdot^{-1} \circ \phi} \text{Symp}_0 \xrightarrow{\text{Flux}} H^1(\Sigma; \mathbb{R}) \xrightarrow{\cdot|_{N_0}} \text{Hom}(N_0, \mathbb{R})$$

This antihomomorphism is surjective and continuous and thus all the fibers are homeomorphic (and homotopic in Symp_h) and, seeing as the range is contractible, each fiber is also homotopic to all of Symp_h . Thus $\text{Symp}_h^\eta \simeq \text{Symp}_h$ by the inclusion. As asserted in [45, Lemma 6ff], $\text{Symp}_h^m \simeq \text{Symp}_h$ by the inclusion (the proof is similar to what we've done here but is slightly simpler), and so $\text{Symp}_h^m \simeq \text{Symp}_h^\eta$ by the inclusion as well. □

Proposition 4.4 *For any η -weakly-monotone symplectomorphism $\phi: (\Sigma, \omega) \rightarrow (\Sigma, \omega)$ in a mapping class h , $HF_*(\phi; \eta)$ is well-defined and $HF_*(\phi; \eta) = HF_*(h; \eta)$.*

Proof We apply Proposition 4.3: there is an isotopy ϕ_t with $\phi_0 = \phi$ and ϕ_1 monotone such that ϕ_t is η -weakly monotone for all t . Now well-definedness follows by Condition 2.7. The final equality follows from Theorem 2.8. □

Corollary 4.5 *If $\phi: (\Sigma, \omega) \rightarrow (\Sigma, \omega)$ is a symplectomorphism in mapping class h and is η -weakly monotone for all η , then $HF_*(\phi) = HF_*(h)$.* □

4.2 Structure of reducible maps

By Thurston's classification (see Thurston [49] and Fathi, Laudenbach and Poenaru [12]; also cf Gautschi [17, Definition 8]), in a reducible mapping class $h \in \Gamma$, there is a (not necessarily smooth) map ϕ which satisfies the following:

Definition 4.6 A reducible map ϕ is in *standard form* if there is a ϕ - and ϕ^{-1} -invariant finite union of disjoint noncontractible (closed) annuli $N \subset \Sigma$ such that:

- (1) For A a component of N and ℓ the smallest positive integer such that ϕ^ℓ maps A to itself, the map $\phi^\ell|_A$ is either a *twist map* or a *flip-twist map*. That is, with respect to coordinates $(q, p) \in [0, 1] \times S^1$, we have one of

$$\begin{aligned} (q, p) &\mapsto (q, p - f(q)) && \text{(twist map)} \\ (q, p) &\mapsto (1 - q, -p + f(q)) && \text{(flip-twist map),} \end{aligned}$$

where $f: [0, 1] \rightarrow \mathbb{R}$ is a strictly monotonic smooth map. We call the (flip-)twist map *positive* or *negative* if f is increasing or decreasing, respectively. Note that these maps are area-preserving.

- (2) Let A and ℓ be as in (1). If $\ell = 1$ and $\phi|_N$ is a twist map, then $\text{Im}(f) \subset [0, 1]$. That is, $\phi|_{\text{int}(A)}$ has no fixed points. (If we want to twist multiple times, we separate the twisting region into parallel annuli separated by regions on which the map is the identity.) We further require that parallel twisting regions twist in the same direction.
- (3) For S a component of $\Sigma \setminus N$ and ℓ the smallest integer such that ϕ^ℓ maps A to itself, the map $\phi^\ell|_S$ is area-preserving and is either *periodic* (ie some power is the identity) or *pseudo-Anosov*. If periodic, we require it to have the form of an isometry of a hyperbolic surface with geodesic boundary. We define pseudo-Anosov on a surface with boundary and standard form for such below. If a periodic component S has $\ell = 1$ and $\phi|_S = \text{id}_S$ we will furthermore call it *fixed*.

Remark 4.7 When a pseudo-Anosov component abuts a fixed component with no twist region in between, we'll need to deal with this case differently throughout. This is the major interaction between components. (There is also the influence of the twisting direction in neighboring twist regions on fixed components.) When two pseudo-Anosov components meet, we perturb both as below and there is no interaction.

A model for a pseudo-Anosov map on a surface with boundary is given in [24, Section 2.1]. The idea is to define a pseudo-Anosov map on a surface with punctures and then “blow up” the punctures to recover boundary components. To define a pseudo-Anosov map on a surface with punctures, we simply consider the punctures as (yet more) singularities and let our foliations \mathcal{F} and \mathcal{G} be given by the subbundles of the tangent space of Σ on which the quadratic differential $z^{p-2}dz^2$ is positive (resp. negative) real, but instead of requiring $p \geq 3$, we only require $p \geq 1$. The case $p = 2$ corresponds to a smooth fixed point (which we have punctured at the fixed point). As in Section 3.2, the map ϕ is given in polar coordinates in a neighborhood of the puncture by the time-1 flow of the (singular at the puncture) Hamiltonian $H_{\text{punc}} = \mu r^2 \cos(p\theta)$.

We must now blow up each puncture to a boundary circle such that the map on the boundary S^1 is a rotation by some angle α (to match up with the maps in item (1) in Definition 4.6). This can be done in various ways (see Jiang and Guo [24, Section 2.1] for example), but we must do it in such a way that the resulting map is area-preserving and all fixed points are nondegenerate. Note that we are free to choose α as the difference can be made up by twisting (or by undoing twisting).

As in [Section 3.2](#), we again have our rotation ρ (which is rotation by some multiple of $2\pi/p$) and this again gives us two cases:

- $\rho \neq \text{id}$ (the “rotated” case)
- $\rho = \text{id}$ (the “unrotated” case)

Perturbing rotated punctures

In the rotated case, we choose α to be the same as the rotation angle for ρ^{-1} . The same argument as in [Section 3.2](#) allows us to modify the map so that it has no fixed points except inside a disk which is rotated with angle α . Excising a subdisk, we have our boundary component with rotation angle α .

Perturbing unrotated punctures

The unrotated case is again more complicated. In order for the fixed points to be nondegenerate, we’ll need the rotation angle α to be a nontrivial one. The desired result will be a Hamiltonian whose flow is a rotation inside a small ball. We’ll find one which has p hyperbolic fixed points around the rotating region and agrees with the flow of H_{punc} outside a slightly larger ball. See [Figure 4](#) for the basic idea.




Figure 4: Modifying H_{punc} in the case $p = 3$. On the left we have the level sets of H_{punc} and on the right we have the level sets of the perturbed Hamiltonian H_{sm} whose flow is rotation inside a small ball (with an elliptic fixed point at the center, but we’re excising the ball so it won’t count) and has $p = 3$ hyperbolic fixed points surrounding the rotating region.

We proceed in four steps. First we add a positive bump of the form $f(r, \theta) = Ae^{-kr^2}$ to H_{punc} to produce the p equally spaced saddle points. Then we smoothly cut off the tail of f without modifying the critical points of $H_{\text{punc}} + f$. Next we modify

$H_{\text{punc}} + f$ near the origin so that its flow rotates at constant angular velocity. Finally we ensure there are no fixed points of the time-1 flow of the Hamiltonian vector field except for those corresponding to our p critical points.

Step 1 We consider $H_1 = H_{\text{punc}} + f(r, \theta) = \mu r^2 \cos(p\theta) + Ae^{-kr^2}$ with A and k positive constants to be chosen later. We're interested in the critical points of this function (away from the origin, where it's singular), so we calculate its partial derivatives:

$$\begin{aligned} \frac{\partial H_1}{\partial r} &= 2\mu r \cos(p\theta) - 2krAe^{-kr^2} \\ \frac{\partial H_1}{\partial \theta} &= -p\mu r^2 \sin(p\theta). \end{aligned}$$

Thus at the critical points of H_1 we must have $\sin(p\theta) = 0$. In this case, $\cos(p\theta) = \pm 1$. If it's -1 , $\partial H_1 / \partial r$ cannot be zero (because both terms are negative). When it's 1 , $\partial H_1 / \partial r = 2\mu r - 2krAe^{-kr^2}$. For $r > 0$, this is zero when $e^{kr^2} = kA/\mu$, ie when $r = r_c := \sqrt{(1/k) \ln(kA/\mu)}$. We impose the restriction that $kA > \mu$. Note that by making k large, we can make r_c arbitrarily small.

In summary, we get critical points at one value of r at the values of θ when $\cos(p\theta) = 1$, that is, for p values of θ . To check that these are all saddle points, we compute the Hessian at these points:

$$\begin{aligned} \frac{\partial^2 H_1}{\partial r^2} &= 2\mu + 2\mu k r^3 > 0 \\ \frac{\partial^2 H_1}{\partial \theta^2} &= -p^2 \mu r^2 < 0 \\ \frac{\partial^2 H_1}{\partial r \partial \theta} &= 0. \end{aligned}$$

Thus the critical points are all nondegenerate and of index one, so we have p saddle points.

Step 2 Keeping f solely a function of r , we cut it off smoothly starting at some point past r_c to give a Hamiltonian H_2 which agrees with H_{punc} outside a ball. As long as we keep $\partial f / \partial r > -2\mu r$, we create no new critical points.

Note that $f(r_c) = \mu/k$. Keeping $\partial f / \partial r$ near $-\mu r_c$ (which, using eg $A = e\mu/k$, is $1/\sqrt{k}$), we can bring f to zero in a radial distance of a constant times $1/\sqrt{k}$; ie for k large we can make H_2 agree with H_{punc} outside an arbitrarily small ball.

Step 3 Now we modify H_2 near the origin to give us H_3 which is $-Br^2 + C$ near the origin (for B positive), which corresponds to the Hamiltonian flow rotating at a

constant angular rate. Since $\partial H_2/\partial r$ is negative for $r < r_c$, we can patch together $-Br^2 + C$ near the origin with H_2 outside a small ball (of radius less than r_c) in a radially symmetric manner to get H_3 such that $\partial H_3/\partial r$ is negative for $r < r_c$ (we do this by choosing C sufficiently large).

Step 4 Finally, to ensure no fixed points of the time-1 flow of H_{sm} , we let H_{sm} be H_3 multiplied by a radially symmetric function which is ϵ for $r < R$ (for ϵ sufficiently small that the only fixed points of the time-1 flow inside radius R are the critical points and for R large enough that H_3 agrees with H_{punc} for $r > R$) and 1 for $r > 2R$ (we may have to make r_c smaller so we can fit a ball of radius $2R$, but this is no problem). This creates no new fixed points in the region $R < r < 2R$ because H_3 and $\partial H_3/\partial r$ have the same sign there. Now there are no fixed points of ϕ_{sm} , the time-1 flow of the Hamiltonian vector field of H_{sm} , except for the p critical points of H_{sm} because outside radius R there are no compact flow lines.

Thus our map ϕ_{sm} has p positive hyperbolic fixed points, and no others, in the perturbed region. Note that everything in this section works in exactly the same manner if we add on a small negative bump instead of a small positive bump, which we may do in order to twist in the opposite direction.

4.3 Nielsen classes for reducible maps

We now study the Nielsen classes of fixed points of standard form reducible maps. Let us briefly describe the fixed points of our standard form reducible maps. We have:

- (Type Ia) The entire component of fixed components S of $\Sigma \setminus N$ with $\chi(S) < 0$.
- (Type Ib) The entire component of fixed components S of $\Sigma \setminus N$ with $\chi(S) = 0$. These are annuli and only occur when we have multiple parallel Dehn twists.
- (Type IIa) Fixed points x of periodic components S of $\Sigma \setminus N$ with $\chi(S) < 0$ which are setwise fixed by ϕ . (These can be understood by considering the map on S as a hyperbolic isometry, from which we see that x must be an elliptic fixed point).
- (Type IIb) Fixed points x of flip-twist regions. (These are elliptic. Note that each flip-twist region has two fixed points.)
- (Type III) Fixed points x of pseudo-Anosov components S of $\Sigma \setminus N$ which are setwise fixed by ϕ . These come in 4 types (note that there are no fixed points associated to a rotated puncture):
 - (Type IIIa) Fixed points which are not associated with any singularity or puncture (ie boundary component) of the presmoothed map. These may be positive or negative hyperbolic.

- (Type IIIb- p) Fixed points which come from an unrotated singular point with p prongs. There are $p-1$ of these for each such, all positive hyperbolic.
- (Type IIIc) Fixed points which come from a rotated singular point. There is one for each such and it is elliptic.
- (Type IIId- p) Fixed points which come from an unrotated puncture (boundary component) with p prongs. There are p for each such, all positive hyperbolic.

In the smooth case (as opposed to the area-preserving case, which we're dealing with), Type Ib fixed points can be isotoped away. To see why, we introduce the concept of a *multiple twist region* which is a (maximal) annulus which is the union of a collection of twist annuli and fixed annuli (ie those of type Ib); note that every fixed annulus is between two twist annuli. In other words, this is a region in which we're performing multiple parallel Dehn twists. If we're allowed non-area-preserving maps, we can get rid of all fixed points in such a region by performing an isotopy which moves every interior point closer to one boundary component.

With this setup, Jiang and Guo [24] have studied the Nielsen classes of fixed points of maps which are in standard form with the following modifications:

- All Type Ib fixed points have been eliminated (in the above manner).
- Fixed points of Type IIIb associated to the same singular point are left as one fixed point.
- Fixed points of Type IIId associated to the same puncture are left as a full circle (ie the blown-up puncture).

Jiang and Guo show that in this case, Nielsen classes of fixed points are all connected – that is, two fixed points are in the same Nielsen class if and only if there is a path between them through fixed points. That is, we have a separate Nielsen class for every component of Type Ia, for every single fixed point of Type IIa, IIb, IIIa or IIIc and for every unrotated singular point of the presmoothed map for Type IIIb (ie the collection of fixed points associated to a single unrotated singular point are all in the same Nielsen class).

Type IIId is special, as mentioned in [Remark 4.7](#). If a pseudo-Anosov component meets a fixed component, then notice that the full circle of fixed points of Type IIId (in the above non-area-preserving model) coincides with fixed points on the boundary of the Type Ia fixed points (note that a pseudo-Anosov component never meets a Type Ib component, as these only occur between twist regions). In this case there will be some interaction, and, while we will perturb the pseudo-Anosov side as in [Section 4.2](#) and

the fixed side with a corresponding small rotation near the boundary, the resulting p fixed points will be in the same Nielsen class as the Type Ia fixed points in the fixed component. If two pseudo-Anosov components meet, the two full circles coincide and, after perturbing as in [Section 4.2](#), we will have $p + q$ positive hyperbolic fixed points all in the same Nielsen class, where p is the number of prongs on one side and q is the number of prongs on the other side. Note that in this case we will have to twist in a coherent manner, but this is not an issue. In all other cases, we have a separate Nielsen class associated to the boundary component of the pseudo-Anosov component containing p positive hyperbolic fixed points.

Jiang and Guo's argument actually extends to imply that, if we don't eliminate Type Ib fixed points (which we won't be able to do in an area-preserving manner), there is a separate Nielsen class (separate from all the others just mentioned) for each component of Type Ib. We briefly explain this.¹⁷ We note that the argument will crucially use the fact that in multiple twist region, all twisting happens in the same direction (ie all parallel Dehn twists have the same sign).

Recall that two fixed points x and y are Nielsen equivalent if there exists a path $\gamma(t)$ with $\gamma(0) = x$, $\gamma(1) = y$, and $\phi(\gamma(t))$ is homotopic rel boundary to $\gamma(t)$. This notion extends to ϕ -invariant sets (ie setwise fixed). We say that two ϕ -invariant sets A and B (either of which may be a single fixed point) are ϕ -related if there exists a path $\gamma: ([0, 1], 0, 1) \rightarrow (\Sigma, A, B)$ such that $\phi(\gamma) \simeq \gamma$ through maps $([0, 1], 0, 1) \rightarrow (\Sigma, A, B)$. We say that such a γ is a ϕ -relation between A and B .

For a map ϕ in standard form, we consider a collection of ϕ -invariant reducing curves $\{C_i\}$ given by ∂N for N our collection of annuli, excepting that when we have parallel annuli, we only use the two outermost ϕ -invariant curves; that is, we consider a multiple twist region as a single annulus. The following is (part of what is) proved in [\[24, Section 3.2-3\]](#):

Lemma 4.8 [\[24, Sections 3.2–3\]](#) *Suppose x and y are two fixed points of a map ϕ which is in standard form which are Nielsen equivalent. Let $\tilde{\gamma}$ be a ϕ -relation between x and y . Homotop $\tilde{\gamma}$ to a ϕ -relation γ with minimal (geometric) intersection number with the C_i in its homotopy class. Break this path into segments at its crossings with the C_i . Let γ_j be the j -th segment and let S_j be the (closure of the) component of $\Sigma \setminus \{C_i\}$ on which γ_j lives (S_j may be an annulus). If γ_j is an interior segment, let A_j and B_j be the two reducing curves (ie elements of $\{C_i\}$) the segment γ_j intersects. If γ_j is an initial segment, instead let $A_j = x$ and if γ_j is a terminal segment, let $B_j = y$.*

¹⁷Gautschi explains how to do this as well [\[17, Proposition 20\]](#), but his argument is specific to the case in which there are no pseudo-Anosov components.

Then γ_j is a ϕ -relation between A_j and B_j on the subsurface S_j . Furthermore, A_j and B_j must be pointwise fixed. Furthermore, we may homotop γ_j inside the fixed point locus unless we're dealing with a fixed point (ie $A_j = x$ or $B_j = y$) and the fixed point is in a multiple twist region. (The same is true here – see the Corollary below.) \square

Corollary 4.9 For a map ϕ in standard form, each component of Type Ib fixed points is in its own Nielsen class, distinct from the Nielsen class of any other fixed point of ϕ . Furthermore, if γ is a ϕ -relation between two fixed points lying inside a multiple twist region, we may homotop it to lie inside the fixed point locus.

Proof By Lemma 4.8, we need only show that on the closed annulus on which we've performed multiple parallel Dehn twists all of the same sign, Nielsen classes of fixed points are connected (ie they are simply the fixed subannuli and, if fixed, boundary circles). Note that we allow the boundary to twist or remain fixed.

Consider a path γ between two fixed points x and y in this situation. Consider 1-chains with fixed endpoints x and y up to homology. This is an affine space over the first homology of the annulus, which is \mathbb{Z} . The difference between $\phi(\gamma)$ and γ in terms of this relative homology is the element of the homology of the annulus corresponding to the number of twisting regions between x and y , with sign. Because all twisting regions have the same sign, this number is not zero unless x and y are in the same component of the fixed point locus of ϕ , and so they cannot be homotopic unless x and y are in the same component.

If x and y are in the same component of the fixed point locus, then we can easily homotop our ϕ -relation γ to lie inside the component (because the annulus deformation retracts onto the circle), proving the second statement as well. \square

4.4 Weak monotonicity of standard form reducible maps

We now show that $\phi = \phi_{sm}$ in standard form is weakly monotone.

Lemma 4.10 Let $\phi: \Sigma \rightarrow \Sigma$ be in standard form. Then $\omega_\phi: H_1(\Gamma(M_\phi)) \rightarrow \mathbb{R}$ is the zero map.

Proof Consider a section ψ of

$$\begin{array}{ccc} \Sigma & \longrightarrow & S^1 \times M_\phi \\ & & \downarrow \\ & & S^1 \times S^1 \end{array}$$

such that $\psi(0, \cdot)$ is in Nielsen class η and $[\psi] \in H_1(\Gamma(M_\phi)_\eta)$ is the class we wish to test ω_ϕ on.

As in [Lemma 3.2](#), let $\alpha_t(s) = \psi(s, t)$. Each α_t should be thought of as a closed curve on Σ by identifying all fibers of $M_\phi = (\mathbb{R} \times \Sigma) / \sim \rightarrow S^1$ with the fiber over $t = 0$ by projecting $\mathbb{R} \times \Sigma$ to Σ (with $t = 0$ and $t = 1$ thought of as separate times). Then $\phi(\alpha_0) = \alpha_1$ is homotopic to α_0 through α_t . Let A_ψ denote the 2-chain with boundary given by this homotopy. Note that this is equivalent to giving the homology class $[\psi(S^1 \times S^1)]$ (the two are related by appending the tube $S^1 \times \alpha_0$ to the 2-chain A_ψ , thought of as living in the fiber Σ over zero). We have $\omega_\phi(A_\psi) = \int_{S^1 \times S^1} \psi^* \omega_\phi$.

We desire to show that $\omega_\phi(A_\psi) = 0$. Note that continuously varying the homotopy α_t has no effect on $\omega_\phi(A_\psi)$. Nor does continuously varying α_0 (which requires that we vary $\alpha_1 = \phi(\alpha_0)$ appropriately) as this does not change the cohomology class $[\psi(S^1 \times S^1)]$.

Claim *There is a choice of ψ such that A_ψ either lies entirely inside a single pseudo-Anosov component or avoids all pseudo-Anosov components.*

Given this, we have two cases.

Case 1 A_ψ lies entirely in S , a pseudo-Anosov component. We show that $\omega_\phi(A_\psi) = 0$ (ie the constant k is zero). In this case, by [Lemma 3.3](#), we must have α_0 either nullhomologous or boundary parallel.¹⁸ If it's nullhomologous, we may argue as in [Lemma 3.2](#) and conclude that $\omega_\phi(A_\psi) = 0$.

If α_0 is boundary parallel, we homotop it near the boundary component so that the map there is given by the local model of [Section 4.2](#) (see also [Figure 4](#)). This map is then Hamiltonian, and so, choosing our homotopy to be one also lying in this region (which, as noted above, does not change the class of A_ψ) we see that $\omega_\phi(A_\psi) = 0$.

Case 2 A_ψ avoids all pseudo-Anosov components. In this case, we may as well be considering a map in standard form which has no pseudo-Anosov components (for example, by replacing pseudo-Anosov components with periodic components with the same boundary behavior). This is the situation Gauduchi considers, and he shows in [[17](#), Proposition 13] that $\omega_\phi(A_\psi) = 0$ for all ψ in this case.

We now prove the above claim. We consider two cases.

Case A ϕ has a fixed point in class η .

¹⁸Technically, that result was for pseudo-Anosov maps on closed surfaces. Jiang and Guo give an argument [[24](#), Proof of Lemma 2.2] which includes pseudo-Anosov maps with boundary.

We show that we may assume A_ψ lies entirely inside the component S of $\Sigma \setminus N$ in which x lies. This requires showing both that we can restrict to $\alpha_0 \subset S$ (which implies $\alpha_1 = \phi(\alpha_0) \subset S$) and that we can bring the homotopy α_t inside S .

Notice that α_0 is a ϕ -relation between x and itself. As in the statement of [Lemma 4.8](#), we homotop α_0 so that it has minimal geometric intersection with a reducing set of curves. Then by [Lemma 4.8](#) and the fact that the collection of fixed points in a single Nielsen class is a subset of a single component of $\Sigma \setminus N$ (see [Section 4.3](#)), α_0 must lie in S .

To see that we can bring the homotopy α_t inside S , notice that the inclusion $S \hookrightarrow \Sigma$ is π_1 -injective. Thus there exists a possibly different homotopy α'_t between α_0 and α_1 which lies in S . Piecing these together we get a map $S^1 \times S^1 \rightarrow \Sigma$ which is necessarily nullhomologous because $\chi(\Sigma) < 0$. Thus the class A_ψ is the same for our new homotopy.

Case B ϕ has no fixed point in class η .

We argue in 3 steps that A_ψ may be homotoped either entirely inside a pseudo-Anosov component or to avoid all pseudo-Anosov components:

Step 1 Let $\{C_i\}$ be the subset of the reducing curves given by the boundaries of the pseudo-Anosov components (including those abutting Type Ia components). Inspired by Jiang and Guo [[24](#), Sections 3.2–3] (cf [Lemma 4.8](#)), we homotop α_0 so it has minimal geometric intersection with the collection of the C_i and work in the universal cover \mathbb{H} of Σ . Each C_i is noncontractible, so each lift in \mathbb{H} is a copy of \mathbb{R} , and each component S_k of $\Sigma \setminus \bigsqcup_i C_i$ is π_1 -injective, so each lift of S_k is also a universal cover of S_k . Furthermore, each lift of a C_i separates \mathbb{H} into two components. Thus $\tilde{\alpha}_0$ (a chosen lift of α_0) intersects each lift of C_i at most once (else we could perform a finger move to reduce intersections).

We lift ϕ to $\tilde{\phi}: \mathbb{H} \rightarrow \mathbb{H}$ and note that $\tilde{\phi}$ maps lifts of C_i to lifts of some C_j (for each i and some j). Thus $\tilde{\phi}(\tilde{\alpha}_0)$ also has minimal intersection with the lifts of the C_i .

Step 2 Choose some basepoint p on α_0 and let \tilde{p} be a lift compatible with our chosen lift of $\tilde{\alpha}_0$. Consider $\beta(t) = \psi(0, t)$, a curve from p to $\phi(p)$. Choose a lift $\tilde{\beta}$ of β through \tilde{p} . Let $\tilde{\alpha}$ be a segment inside $\tilde{\alpha}_0$ from \tilde{p} to \tilde{p}' (another lift of p) which is injective on its interior (ie its projection wraps around α once).

Thus we have a rectangle in the universal cover with corners \tilde{p} , \tilde{p}' , $\tilde{\phi}(\tilde{p})$ and $\tilde{\phi}(\tilde{p}')$ and edges $\tilde{\alpha}$, $\tilde{\phi}(\tilde{\alpha})$, $\tilde{\beta}$ and $\tilde{\beta}'$ (the lift of $\tilde{\beta}$ through \tilde{p}'). We call its interior $\tilde{\psi}$.

We will be homotoping these objects into easier to work with positions but will not change their names as we do so: homotop β (and thus also $\tilde{\beta}$ and $\tilde{\beta}'$) and thus also α_0 ,

so it has minimal intersection with the C_i , so the lifts each intersect each lift of a C_i at most once. We can do this without disturbing this property for α_0 .

Step 3 Homotop $\tilde{\psi}$ rel boundary so that it is transverse to each lift of a C_i . Let \tilde{C} be the intersection of $\tilde{\psi}$ with the lifts of the C_i . This is a compact one-manifold with boundary in the rectangle $\tilde{\psi}$.

Note that \tilde{C} has no circle components because all of the lifts of the C_i are copies of \mathbb{R} . Also note that no component of \tilde{C} has both boundary components on the same edge of the rectangle $\tilde{\psi}$ because each of the edges intersects each lift of a C_i at most once.

Step 4 Each of the edges of ψ has a certain number of boundary points of \tilde{C} on it. Note that opposite edges have the same number of such (which are in corresponding positions) because one pair is a lift of the same path β on Σ , and the other pair are lifts of curves α_0 and $\phi(\alpha_0)$ and ϕ preserves the collection $\{C_i\}$.

Using this and Step 3, we have three cases:

Case 1 There are no boundary points of \tilde{C} on $\tilde{\alpha}$ (and thus also on $\tilde{\phi}(\tilde{\alpha})$).

Thus \tilde{C} consists of some number of (nonintersecting) paths from $\tilde{\beta}$ to $\tilde{\beta}'$ which each go from a point on $\tilde{\beta}$ to the corresponding point on $\tilde{\beta}'$. If we project to Σ , what we see is the annulus given by A_ψ with boundary α_0 and $\phi(\alpha_0) = \alpha_1$ with some number of intermediate C_i 's. If there is at least one intermediate C_i , then α_0 is homotopic to a C_i . If we homotop it so that its image is C_i , we find that $\omega_\phi(A_\psi) = 0$, as desired.

If there are no intermediate C_i 's, then A_ψ is contained in one component of $\Sigma \setminus \coprod_i C_i$, as desired.

Case 2 There are no boundary points of \tilde{C} on $\tilde{\beta}$ (and thus also on $\tilde{\beta}'$).

In this case, \tilde{C} consists of some number of (nonintersecting) paths from $\tilde{\alpha}$ to $\tilde{\phi}(\tilde{\alpha})$ which each go from a point on $\tilde{\alpha}$ to a point on $\tilde{\phi}(\tilde{\alpha})$. If we project to Σ , what we see is the annulus given by A_ψ with boundary α_0 and $\phi(\alpha_0) = \alpha_1$ with some number of C_i 's (which are nonintersecting) running from a point on α_0 to the corresponding point on α_1 . If there are no such C_i 's, then A_ψ is contained in one component of $\Sigma \setminus \coprod_i C_i$, as desired. Similarly if there is only one C_i .

Suppose there are at least two such C_i . We desire to show that A_ψ avoids pseudo-Anosov components. Suppose that it does not. Then there is some rectangle in A_ψ with boundary edges given by portions of some C_i and some C_j as well as portions of α_0 and α_1 which is contained in a pseudo-Anosov component. Note that ϕ applied to this portion of α_0 is the portion of α_1 .

Work only in the pseudo-Anosov component as a surface with boundary. If C_i (resp. C_j) corresponds to a boundary component with $p > 1$ prongs, consider a surface in which

we've collapsed C_i (resp. C_j) and have a pseudo-Anosov map on this surface with fewer boundary components. Near any boundary component with one prong, we have a (hyperbolic) fixed point nearby. After collapsing any with $p > 1$, we have a fixed point of some kind nearby (we may perturb to our standard map if desired; there will still be a fixed point of some kind nearby). By the argument in the proof of [Theorem 3.6](#), we may modify α_0 slightly to exhibit a Nielsen equivalence between a fixed point near C_i and a fixed point near C_j (recall that $\alpha_1 = \phi(\alpha_0)$). Thus by the results of [Section 4.3](#) (or by [\[2\]](#), as in the proof of [Theorem 3.6](#), if our surface is now closed), $C_i = C_j$. Now if C_i has $p > 1$ prongs so that we've collapsed it, then we have a curve γ on the surface such that $\phi(\gamma) = \gamma$ (using the induced map ϕ). Thus γ is nullhomotopic by [\[24, Lemma 2.2\]](#). Similarly, if C_i is a $p = 1$ boundary component and again by [\[24, Lemma 2.2\]](#) we find that it is contractible. But this is impossible – we can reduce intersections of α_0 with $C_i = C_j$ by a finger move. Thus A_ψ avoids pseudo-Anosov components, as desired.

Case 3 There are boundary points of \tilde{C} on both $\tilde{\alpha}$ and $\tilde{\beta}$.

In this case, the projection to Σ looks much like it did in Case 2, but now the portions of the C_i 's running from α_0 to α_1 needn't start and end at corresponding points; that is, they twist in the annulus. Inspired by [\[17, Lemma 7\]](#), we consider an appropriate power of ϕ , ϕ^k . We get $A_\psi, \phi(A_\psi), \dots, \phi^{k-1}(A_\psi)$ which patch together to give an annulus kA_ψ between α_0 and $\phi^k(\alpha_0)$ of area $k\omega_\phi(A_\psi)$. If we choose k properly, then the C_i running through kA_ψ start and end at corresponding points of α_0 and $\phi^k(\alpha_0)$. Thus we're in Case 2 (possibly with an irrelevant Dehn twist in the annulus), and we see that $k\omega_\phi(A_\psi) = 0$, and so $\omega_\phi(A_\psi) = 0$. □

Corollary 4.11 *Let $\phi: \Sigma \rightarrow \Sigma$ be in standard form. Then ϕ is η -weakly monotone for all $\eta \in \pi_0(\Gamma(M_\phi))$.* □

Corollary 4.12 *Let $\phi: (\Sigma, \omega) \rightarrow (\Sigma, \omega)$ be a reducible symplectomorphism in standard form in mapping class $h \in \Gamma$. Then $HF_*(\phi) = HF_*(h)$.*

Proof This follows from [Corollary 4.11](#) and [Corollary 4.5](#). □

4.5 Floer homology of ϕ_{sm} for reducible maps

In this section we show how to compute $HF_*(\phi_{sm})$ by showing that $(CF_*(\phi_{sm}), \partial)$ splits into a contribution from each component (and annular region between parallel Dehn twists), though what this contribution is may be affected by the behavior of the neighboring regions. The argument follows Gaultsch [\[17, Sections 4–5\]](#), with the

pseudo-Anosov components and their contribution “coming along for the ride,” with the exception of pseudo-Anosov components directly abutting fixed components (cf [Remark 4.7](#) and [Section 4.3](#)).

The basic idea is to perturb the fixed components by the Hamiltonian flow of a Morse function (actually a small Morse–Smale function). This Morse function should agree near the boundary with the direction of twisting of a twist component, if any. That is, the Morse function achieves a local maximum at boundaries that twist positively and a local minimum at boundaries that twist negatively. Let $\hat{\phi}$ denote the perturbed map. At boundary components of the fixed components where there’s a positive twist joining up, a local model for $\hat{\phi}$ in a neighborhood of the boundary component is $(q, p) \mapsto (q, p - \delta)$ for $\delta > 0$ small. Similarly, when joining up with a negative twist, the map in a neighborhood of the boundary component is $(q, p) \mapsto (q, p + \delta)$ for $\delta > 0$ small. We let Σ_0 denote the collection of fixed components, and $\partial_{\pm}\Sigma_0$ the collection of components of $\partial\Sigma_0$ on which we’ve joined up with a positive (resp. negative) twist.

Additionally, the Morse function should agree near the boundary of a pseudo-Anosov component with the Hamiltonian perturbation we chose in [Section 4.2](#), which we may have to modify slightly by shrinking the size of the perturbation further (ie the same process we did already in [Section 4.2](#), but with a smaller ϵ). We may choose which direction to twist on each side when a pseudo-Anosov component meets a fixed component so long as we do so coherently. The collection of components of the boundary of Σ_0 which meet a pseudo-Anosov component will be denoted $\partial_p\Sigma_0$.

With this setup, we call $\hat{\phi}$ a perturbed standard form map. Note that we have left some ambiguity with respect to which direction we’re twisting when a fixed component meets a pseudo-Anosov component.

Now Gautschi [[17](#), Proposition 14] has a neck-stretching argument (actually, he enlarges the symplectic form) to show that different components of $\Sigma \setminus N$ don’t interact because flow lines originating at fixed points of Type Ia and Type Ib can’t approach the boundary of a component except where a Type IIIId boundary abuts a Type Ia component (see the proof of [Lemma 4.15](#), below, for other fixed points). This uses the fact that $\omega_{\phi}: H_1(\Gamma(M_{\phi})) \rightarrow \mathbb{R}$ is the zero map, ie our [Lemma 4.10](#). We present a modification of Gautschi’s argument in order to show that flow lines between Type Ia and Type IIIId fixed points in abutting components don’t reach past the Hamiltonian portion of the pseudo-Anosov component corresponding to the Type IIIId fixed points. We combine this with an argument of Salamon and Zehnder [[41](#)] to determine these flow lines.

Remark 4.13 Hutchings and Sullivan have an argument involving intersection positivity with the J -holomorphic foliation coming from $\mathbb{R} \times T$ for T a torus in M_{ϕ}

which is invariant under the mapping torus flow [21, Lemma 3.11]. They consider only such tori coming from fixed circles $C \subset \Sigma$, but the argument extends to collections of circles which are permuted. This can also be used, together with an argument as in [24, Sections 3.2–3] or Case B of the proof of Lemma 4.10, to show that flow lines originating at fixed points of Type Ia and Type Ib can't approach the boundary of a component except where a Type IIIId boundary abuts a Type Ia component. Their argument seems to be morally equivalent to Gautschi's, with both essentially getting at a local energy contribution. We don't have an argument along the lines of Hutchings and Sullivan's for Lemma 4.14, below, but one may be possible.

Lemma 4.14 *Consider the subset Σ'_0 of Σ consisting of Type Ia regions which abut Type IIIId fixed points together with neighborhoods of the corresponding Type IIIId boundaries (including the fixed points) in the appropriate pseudo-Anosov components. Then there exists a modification of the Hamiltonian on the interior of Σ'_0 such that the differentials between fixed points in Σ'_0 are given by gradient flow lines of the Hamiltonian.*

Proof We perform the same analysis at each Type IIIId boundary.

We assume the setup of Step 4 of Section 4.2. Our plan is to modify what is done in that step. We multiply by ϵ_1 for $r < R$. For $\frac{5R}{4} < r < \frac{3R}{2}$, we multiply by ϵ_2 , and we multiply by 1 for $r > 2R$. We assume that $0 \leq \epsilon_i \leq 1$. We furthermore multiply our Hamiltonian on Σ_0 by ϵ_1 . The plan is to obtain a local energy contribution whenever a differential passes through the region $\frac{5R}{4} < r < \frac{3R}{2}$ and then to use Gautschi's trick of varying the symplectic form to increase the local energy contribution without bound while keeping constant the energy of differentials between fixed points in Σ'_0 .

Part 1 We derive bounds on $|\nabla H|$. Fix R . We use the standard Euclidean metric. We perform a brief calculation to bound $|\nabla H|$ from above and below for $R < r < 2R$. Suppose H is of the form $\kappa\mu r^2 \cos(p\theta)$ (for some κ) in a subset of this region. Then

$$|\nabla H| = \kappa\mu \sqrt{4r^2 \cos^2(p\theta) + p^2 r^6 \sin^2(p\theta)}.$$

Using $\frac{1}{2} \leq \min(\cos^2(p\theta), \sin^2(p\theta)) \leq \max(\cos^2(p\theta), \sin^2(p\theta)) \leq 1$,

we have:
$$\kappa\mu \sqrt{\frac{1}{2} \min(4r^2, p^2 r^6)} \leq |\nabla H| \leq \kappa\mu \sqrt{4r^2 + p^2 r^6}$$

$$\implies \kappa\mu \frac{1}{\sqrt{2}} \min(2r, pr^3) \leq |\nabla H| \leq \kappa\mu(2r + pr^3)$$

$$\implies \kappa\mu \frac{1}{\sqrt{2}} \min(2R, pR^3) \leq |\nabla H| \leq \kappa\mu(4R + 8pR^3).$$

We must also bound the derivative in the region $R < r < \frac{5R}{4}$. When $r = \frac{5R}{4}$, we're supposed to have $H(r) = \epsilon_2 \mu r^2 \cos(p\theta)$. No matter what ϵ_1 is, we'll have to increase H by a quantity of at most $\epsilon_2 \mu \frac{25}{16} R^2$ over an interval of length $\frac{R}{4}$. We can do this keeping $|\nabla H|$ less than twice the ratio, namely $|\nabla H| \leq \epsilon_2 \mu \frac{25}{4} R$.

Part 2 We seek the local energy contribution. Consider our differential $u(s, t)$ and suppose that $u(s, t)$ achieves the r -value $\frac{3R}{2}$. We obtain an energy estimate for such u and then use Gautschi's trick of varying the symplectic form in a neighborhood of the annulus $\frac{5R}{4} < r < \frac{3R}{2}$ to show that a differential between two fixed points in Σ'_0 cannot reach beyond $\frac{3R}{2}$ at any point.

For s in the interval between s_1 and s_2 , we have that $u(s, 0)$ is between r -values $\frac{5R}{4} + \frac{R}{16}$ and $\frac{3R}{2} - \frac{R}{16}$. Thus we have bounds of the form $\epsilon_2 K_1 \geq |\nabla H| \geq \epsilon_2 K_2$ (where K_1 and K_2 are some constants, because μ and R are fixed).

To find a lower bound on $d(\phi(x), x)$ for $x \in \Sigma$ with an r -value of between $\frac{5R}{4} + \frac{R}{16}$ and $\frac{3R}{2} - \frac{R}{16}$, we bound the second derivative of H (that is, the Hessian) from above. It is not hard to see that this is bounded in terms of ϵ_2, μ and R , and the bound in terms of ϵ_2 is linear. Thus we have a bound of the form $|\text{Hess}(H)| \leq \epsilon_2 L$.

The maximum displacement of ϕ at time 1 is bounded by $\epsilon_2 K_1$ (the upper bound on $|\nabla H|$). Thus the maximum change in ∇H is $(\epsilon_2)^2 K_1 L$. Thus for sufficiently small ϵ_2 , this is smaller than $\frac{1}{2} \epsilon_2 K_2$. Thus the displacement of the flow at time 1 is at least $\frac{1}{2} \epsilon_2 K_2$ (by the above lower bound on $|\nabla H|$).

Thus
$$\int_0^1 |\partial_t u(s, t)| dt \geq \frac{1}{2} \epsilon_2 K_2$$

for s between s_1 and s_2 . By Hölder's inequality, we have

$$\int_0^1 |\partial_t u(s, t)|^2 dt \geq \frac{1}{4} (\epsilon_2)^2 (K_2)^2$$

for s between s_1 and s_2 . Thus

$$\int_{-\infty}^{\infty} \int_0^1 |\partial_t u(s, t)|^2 dt ds \geq \frac{1}{4} (\epsilon_2)^2 (K_2)^2 |s_1 - s_2|.$$

We now use Gautschi's trick. If we multiply the symplectic form and the metric by $M \gg 0$ (while keeping the almost complex structure J constant) on the entire pseudo-Anosov component minus a region near the boundary. On the region between the compact pieces of the level set through the p saddle points and the boundary, we multiply the symplectic form and metric by a function $M(s)$ constant on level sets $H = s$ which quickly reaches one (and similarly at any other boundary component). In

particular, we've multiplied the symplectic form and metric by M in a neighborhood of the annulus $\frac{5R}{4} < r < \frac{3R}{2}$. The holomorphic curves all remain exactly the same because we haven't changed the complex structure, but we get a better energy estimate in terms of the new metric, with a factor of M showing up:

$$E(u) = \int_{-\infty}^{\infty} \int_0^1 |\partial_t u(s, t)|^2 dt ds \geq \frac{1}{4} M(\epsilon_2)^2 (K_2)^2 |s_1 - s_2|.$$

In addition, by Lemma 4.10, our action functional is exact and so the energy of a holomorphic curve is given by the difference in the action between its two endpoints. If we temper $M(s)$ from M to 1 quickly enough, we can make the difference in action between any two fixed points in Σ'_0 arbitrarily close to the difference in action before we expanded the symplectic form and metric because the region in which we have modified the symplectic form does not separate the fixed points from each other in the sense that they can be connected by a topological cylinder which avoids the region. The action of this topological cylinder is then nearly unchanged and is equal to the energy of any holomorphic curve u between the two fixed points. Thus as we vary M , $E(u)$ must remain (nearly) constant, contradicting the above inequality. Thus differentials between fixed points in Σ'_0 do not go past $r = \frac{3R}{2}$.

Part 3 Now we use the argument of Salamon and Zehnder [41, Theorem 7.3] to show that for ϵ small enough, the only differentials in the region Σ_0 together with everything within radius $\frac{3R}{2}$ (we call this combined region Σ'_0) are those corresponding to gradient flow lines of our Hamiltonian (which we make Morse–Smale¹⁹).

The main point is that making ϵ_1 and ϵ_2 small makes $|\nabla H|$ small in the region governed by ϵ_1 and ϵ_2 – that is, Σ'_0 (recall that $|\nabla H| \leq \epsilon_2 \mu \frac{25}{4} R$ in the region $R \leq r \leq \frac{5R}{4}$, so $|\nabla H|$ is governed linearly in this entire region by ϵ_1 and ϵ_2). We note that on Σ'_0 we may assume that ∇H is determined up to a scalar multiple (namely, ϵ_2) by keeping the ratio between ϵ_1 and ϵ_2 constant.

We briefly recap Salamon and Zehnder's argument in order to show that it applies in our situation. They show linearization of the $\bar{\partial}$ operator is onto for every differential which corresponds to a gradient flow line [41, Theorem 7.3 (1)] for ∇H sufficiently small. This is a local condition and thus carries over to our setting.

They then argue as follows: if there's a differential for τH for $\tau > 0$ arbitrarily close to zero, we can take a sequence of these with $\tau \rightarrow 0$. The energy of such goes like τE , where E is the maximum energy possibly when using H .

¹⁹Morse–Smale functions are dense in the space of Morse functions, and so we can do this without disturbing any of our inequalities or critical points qualitatively.

By Gromov compactness we get a sequence converging to a broken differential. In our setting, we should add that this broken differential must stay in Σ'_0 by the reasoning in Part 2 and the fact that there are no fixed points in the same Nielsen class as fixed points in Σ'_0 which are outside Σ'_0 . The rest of the argument can now be repeated verbatim. \square

Lemma 4.15 (cf [17, Lemma 16]) *For any generic path of complex structures J_t on Σ , the Floer homology chain complex $(CF_*(\phi), \partial_{J_t})$ splits into a sum of chain complexes (C_i, ∂_i) for each component of $\Sigma \setminus N$. These chain complexes are the Floer homology chain complex of the component S with boundary, except that when a pseudo-Anosov component abuts a fixed component, we include the boundary fixed points as part of the fixed component and not as part of the pseudo-Anosov component.*

Proof Fixed points in different components of $\Sigma \setminus N$ are in different Nielsen classes by Section 4.3 (except where a Type Ia region meets a Type III d boundary). The moduli spaces of flow lines between different Nielsen classes are empty, as discussed in the proof of Theorem 3.6, so indeed the chain complex splits. The neck stretching argument of Gaultsch [17, Proposition 14] mentioned earlier implies that the flow lines don't reach the boundary of the component for components of Type Ia and Type Ib (unless a Type Ia region meets a Type III d boundary). For all other components, there are no possible differentials because fixed points of Type IIa, Type IIb, Type IIIa and Type IIIc are all the only fixed point in their Nielsen class by Section 4.3, and fixed points of Type IIIb and fixed points of Type III d not abutting a component of Type Ia are in a Nielsen class with only points of the same index by Sections 3.2, 4.2 and 4.3. For Type Ia components abutting Type III d boundaries, we use Lemma 4.14. \square

This splitting together with the fact that we've identified the summands with the Floer homology of the components S means that we can cite the rest of Gaultsch's calculation and add on the contributions from pseudo-Anosov components and from Type Ia components which abut Type III d fixed points.

First, some notation. Recall the notation of Σ_0 for the collection of fixed components as well as the three types of boundary $\partial_{\pm}\Sigma_0$ and $\partial_p\Sigma_0$. Additionally let Σ_1 be the collection of periodic components and let Σ_2 the collection of pseudo-Anosov components with punctures (ie before any perturbation in Section 4.2) instead of boundary components wherever there is a boundary component that meets a fixed component.

We further subdivide Σ_0 . Let Σ_a be the collection of fixed components which don't meet any pseudo-Anosov components. Let $\Sigma_{b,p}$ be the collection of fixed components

which meet one pseudo-Anosov component at a boundary with p prongs. In this case, we assign the boundary components to $\partial_- \Sigma_0$ (this is an arbitrary choice). Let $\Sigma_{b,p}^\circ$ be the collection of the $\Sigma_{b,p}$ with each component punctured once. Let $\Sigma_{c,q}$ be the collection of fixed components which meet at least two pseudo-Anosov components such that the total number of prongs over all the boundaries is q . In this case, we assign at least one boundary component to $\partial_+ \Sigma_0$ and at least one to $\partial_- \Sigma_0$ (and beyond that, it does not matter).

Theorem 4.16 *Let $\hat{\phi}$ be a perturbed standard form map $\Sigma \rightarrow \Sigma$ in a reducible mapping class $h \in \Gamma$ with choices of the signs of components of $\partial_p \Sigma_0$ and notation as in the preceding paragraphs. Then*

$$\begin{aligned}
 HF_*(h) = HF_*(\hat{\phi}) \cong & H_{*(\text{mod } 2)}(\Sigma_a, \partial_- \Sigma_a; \mathbb{Z}/2) \\
 & \oplus \bigoplus_p (H_{*(\text{mod } 2)}(\Sigma_{b,p}^\circ, \partial_- \Sigma_{b,p}; \mathbb{Z}/2) \oplus (\mathbb{Z}/2)^{(p-1)|\pi_0(\Sigma_{b,p})|}) \\
 & \oplus \bigoplus_q (H_{*(\text{mod } 2)}(\Sigma_{c,q}, \partial_- \Sigma_{c,q}; \mathbb{Z}/2) \oplus (\mathbb{Z}/2)^{q|\pi_0(\Sigma_{c,q})|}) \\
 & \oplus \mathbb{Z}/2^{\Lambda(\hat{\phi}|_{\Sigma_1})} \oplus CF_*(\hat{\phi}|_{\Sigma_2}),
 \end{aligned}$$

where $\Lambda(\hat{\phi}|_{\Sigma_1})$ is the Lefschetz number of $\hat{\phi}|_{\Sigma_1}$, the $\mathbb{Z}/2^{\Lambda(\hat{\phi}|_{\Sigma_1})}$ summand is all in even degree, the other two $\mathbb{Z}/2$ summands are all in odd degree, and $CF_*(\hat{\phi}|_{\Sigma_2})$ denotes the chain complex for $\hat{\phi}$ on the components Σ_2 (where, as mentioned above, we've left off the Type IIIId fixed points which abut a fixed component).

Furthermore, $H_1(\Sigma)$ acts as zero except on the first summand, where it acts by the intersection product. That is, given a component S of Σ_a with $\partial_+ S = 0$, the action of $H_1(\Sigma)$ is by intersection products for its image under the map to $H_1(\Sigma, \Sigma \setminus S) = H_1(S, \partial_- S)$, and given a component S with $\partial_- S = 0$, the action of $H_1(\Sigma)$ is by intersection products with the inclusion of $H_1(S)$ to $H_1(\Sigma)$, noting that this action preserves $H_1(S)$.

Proof The first summand and the $\mathbb{Z}/2^{\Lambda(\hat{\phi}|_{\Sigma_1})}$ summand are as in [17, Theorem 1]. The last summand comes from the fact that there are no differentials in the Floer chain complex on the pseudo-Anosov components as in the proof of Theorem 3.6. (We've left off the Type IIIId fixed points which meet fixed components and will deal with them elsewhere.)

The sum over q arises in the same manner as the first summand with the extra note that if we choose a Morse function on Σ_c which has no maxima or minima, which we can do because every component of Σ_c has both a positive and a negative boundary

component, then there are no differentials between the fixed points in each of the Σ_c components and the various Type III_d fixed points (which correspond to the second term in the sum over q).

The sum over p arises in this same manner, except that we choose a Morse function on each component of Σ_b with one maximum and no minima (which we can do due to its one positive boundary component). Then one of the Type III_d fixed points cancels with the maximum (they form a cancelling pair), resulting in the homology of the punctured component and in one fewer fixed point in the second half of the sum over p .

For the action of $H_1(\Sigma)$, we use the extrinsic or algebraic formulation given in [Section 2.5](#) for $HF_*(\hat{\phi})$, using [Proposition 2.9](#) to see that our calculation holds for $HF_*(h)$. The action preserves Nielsen classes because it is defined in terms of flow lines so it acts on the summands in the statement of the theorem. It acts as zero on the last two because there are no flow lines between generators of those summands. Furthermore, it can only act nontrivially on a summand which is nonzero in both degrees mod two, which now leaves only the first. There the correspondence of Morse flow lines and Floer flow lines shows that the action is the same as the action of $H_1(\Sigma)$ on Morse homology, which is the same as the intersection product. \square

Remark 4.17 In comparing this result with [\[43\]](#) or [\[17\]](#), note that we use symplectic Floer homology instead of symplectic Floer cohomology.

5 Train tracks and combinatorially computing $HF_*(h)$

In order to compute $HF_*(h)$ for a pseudo-Anosov or reducible mapping class from [Theorem 3.6](#) or [Theorem 4.16](#), we need to understand the canonical singular representative ϕ_{sing} of the pseudo-Anosov mapping class, or of the pseudo-Anosov pieces of the reducible map. In particular, we must understand its fixed points, its singular points, and how each singular fixed point is mapped to itself (ie whether it is rotated). All this data is combinatorially computable from the action of a representative of the mapping class on an invariant train track. It is the purpose of this section to describe how this computation works. We note that by work of Bestvina and Handel [\[1\]](#), one can algorithmically produce an invariant train track from a representation of ϕ_{sing} as a collection of Dehn twists (or, if ϕ_{sing} is reducible, produce reducing curves from the Dehn twists, and then repeat with the pseudo-Anosov pieces). Additionally, there is a large collection of examples due to Penner [\[35\]](#) for which producing the invariant train track from the Dehn twists is trivial. We review Penner's construction and compute some examples based on it in the next section. In this section as well, we follow

Penner [35]. See also Bestvina and Handel [1], Penner and Harer [36, Epilogue], Mosher [33, Section 9], Fathi, Laudenbach and Poenaru [12] and Thurston [49].

A *train track* τ on a surface Σ is a collection of simple closed curves and graphs in which every vertex has valence at least three embedded in the interior of Σ . It is required to be a smooth embedding on each closed edge (and curve), both of which we call a *branch*, and at the vertices, called *switches*, the tangent vectors for each edge must be collinear and in addition they must not be all of the same sign (that is, the track must not “stop”). The train track separates the surface into various faces, which are components of $\Sigma \setminus \tau$. These faces have straight vertices (which we do not count as vertices of the face) and cusp vertices (which we do). We require each face to have negative Euler characteristic, with cusps counting as $-\frac{1}{2}$. Note that the sum of the Euler characteristics of the faces (with this cusp contribution) equals the Euler characteristic of the surface. We say that a track τ *fills* Σ if every component of $\Sigma \setminus \tau$ is either a disk or a boundary parallel annulus (with some number of cusps).

For a train track τ , define the *branch space* V_τ to be the vector space generated by the branches b_i of τ . We say that a track τ *carries* a track τ' if we can isotop τ' such that it lies entirely inside τ . We may assume that switches of τ' are taken to switches of τ (but not every switch of τ will have a switch of τ' on it, even if τ' passes through that switch). In this case, we get an *incidence matrix* $M: V_\tau \rightarrow V_{\tau'}$ given by $Mb_i = \sum_j M_{ij}c_j$ for branches b_i of τ and c_j of τ' , where M_{ij} is the count of “how many times c_j passes over b_i .”

A track τ is *recurrent* if for each branch b of τ , there is a curve γ which is carried by τ and which runs over b . A track τ is *transversely recurrent* if for each branch b of τ , there is a simple closed curve γ which intersects b nontrivially and creates no bigons with τ (a bigon with cusps on any of its edges doesn't count). If both of these are true, we say the track τ is *birecurrent*.

Following [36, Epilogue], we consider a map $\phi: \Sigma \rightarrow \Sigma$ and a birecurrent train track τ such that $\phi(\tau)$ is carried by τ (this is what is meant by an *invariant train track* for ϕ). In this case we get an incidence matrix $M: V_\tau \rightarrow V_{\phi(\tau)} = V_\tau$ (where the equality is given by identifying $\phi(b)$ with b). For ϕ to be pseudo-Anosov, it will be necessary that τ fills Σ and also that (possibly after modifying the track by moves given in [1], such as collapsing invariant trees) M is *Perron–Frobenius*, which means that some (positive) power of M has all positive entries. It then follows that M has a unique (up to scaling) eigenvector $v = (v_i)$ with all positive entries. Furthermore, its (positive real) eigenvalue $\lambda > 1$ is the eigenvalue of M which is maximum in absolute value. M^T is also Perron–Frobenius, and we let $w = (w_i)$ denote its unique (up to scaling) eigenvector with all entries positive. Its eigenvalue must also be λ .

From this data we can construct the measured foliations for ϕ_{sing} . We fatten each branch and let its width be v_i and its length be w_i . These automatically satisfy the *switch condition* that the sum of the widths of branches abutting a switch from one direction equals the sum of the widths of the branches abutting the switch from the opposite direction. Additionally, they satisfy a dual switch condition on the w 's (see Penner [35]). It then turns out that we can “zip up” the fattened track, collapsing the faces and getting a singular point for each face [35].

Computing the number of fixed points

Now we come to the crux of the matter, which is how to compute the number of fixed points based on this description. To warm up we consider a toy model: suppose R is a rectangle in the plane (parallel to the x - and y -axes) and we apply an affine transformation Φ which stretches by a factor of $\lambda > 1$ in the x -direction and shrinks by a factor of λ in the y -direction. Suppose furthermore that the vertical extent of $\Phi(R)$ is contained within the vertical extent of R and the horizontal extent of R is contained within the horizontal extent of $\Phi(R)$, where in both cases we allow their boundaries to coincide. Then Φ has precisely one fixed point in R , as can be readily deduced from the contraction mapping principle in the x and y coordinates separately. Furthermore, if their boundaries coincide either horizontally or vertically, then the fixed point will be on this boundary (and if two boundaries coincide, the fixed point will be at the common vertex). Note also that we again get a unique fixed point if the affine transformation stretches in the x -direction, contracts in the y -direction and is then rotated by 180° (ie is negative hyperbolic instead of positive hyperbolic), subject to the same conditions on the vertical/horizontal extents.

What this means is that there is one fixed point of ϕ every time the branch $\phi(b_i)$ runs over the branch b_i . This count is the trace of M . Unfortunately, we have overcounted, counted some singular fixed points together with the nonsingular fixed points and additionally undercounted a specific type of nonsingular fixed point, so we have to make some corrections.

In order to deal with the overcounting and undercounting, we formalize two further pieces of data. These are the tangential and transverse lists. Arbitrarily orient each branch b_i . We let σ_i denote the (tangential) ordered list given by the b_j 's which $\phi(b_i)$ runs over (we also record whether the orientations agree when $\phi(b_i)$ runs over b_j). Note that the number of times b_j appears in σ_i equals M_{ij} . Also, let v_i denote the (transverse) ordered list given by the b_j 's for which $\phi(b_j)$ runs over b_i , in order from left to right (with b_i pointing up; again we record whether their orientations agree). Note that the number of times b_j appears in v_i is given by M_{ji} .

Let u_p denote the number of unrotated fixed singular points with p prongs and let r_p denote the number of rotated fixed singular points with p prongs. Both of these can be read off from the track τ together with the lists σ_i and ν_i . Rotated fixed singular points do not contribute any to $\text{tr}(M)$, but unrotated fixed singular points contribute p to $\text{tr}(M)$, corresponding to p points which are identified when the fattened track is zipped up (ignoring under- and over-counting that will be dealt with below). Note that there can be no other fixed points that are on the boundary of the fattened track (and which would then be identified when this boundary is zipped up) because the track is stretched by λ .

Next we have overcounting of nonsingular fixed points when there is a fixed point at the end of a (fattened) branch, ie living over a switch, for then it is counted twice, once in each of the two branches it lies in. Let s_i equal one-half the number of times that b_i is the first entry in σ_i plus one-half the number of times it's the last (in both cases we require that the orientation of $\phi(b_i)$ agree with the orientation of b_i). The number $\sum_i s_i$ represents how many of these sorts of overcounting errors there are.

Finally we have undercounting of nonsingular fixed points when there is a "flip." That is when b_i and b_j share a common switch and, supposing without loss of generality that this is the initial switch for each of b_i and b_j , the list σ_i begins with b_j and σ_j begins with b_i (both with compatible orientations). Let f denote the number of flips.

Theorem 5.1 *For ϕ pseudo-Anosov on a closed surface Σ , the rank of $HF_*(\phi)$ in the above setup is $\text{tr}(M) + f - \sum_i s_i + \sum_p (r_p - u_p)$.*

Proof By [Theorem 3.6](#), we get a contribution of $\sum_p (p - 1)u_p$ from the unrotated singular points, a contribution of $\sum_p r_p$ from the rotated singular points, and the above discussion shows that the number of normal fixed points is $\text{tr}(M) + f - \sum_i s_i - \sum_p p u_p$. Adding together these three quantities, we get the desired result. \square

We could get the number of even generators and the number of odd generators from the lists σ_i and ν_i with orientations attached, if desired. Alternatively, we can simply use the Lefschetz fixed point theorem to get the difference between the number of odd and the number of even generators and combine that with the above theorem.

6 Penner's construction and computation of examples

For a simple closed curve C on Σ , let τ_C be the positive Dehn twist about C . Consider a collection $\{C_i\}$ of disjoint embedded curves on Σ and $\{D_j\}$ another such collection. Let ϕ be any composition of positive Dehn twists along the C_i 's and negative Dehn

twists along the D_j 's such that every curve listed shows up at least once. Penner [35] shows that if every region of the complement $\Sigma - \coprod_i C_i - \coprod_j D_j$ is a disk with more than two edges, then the mapping class $[\phi]$ is pseudo-Anosov. Furthermore, if every curve listed shows up at least twice, then we have an explicit invariant train track for ϕ whose incidence matrix M is Perron–Frobenius.²⁰

This invariant train track τ is simply take the union of all the curves C_i and D_j with their intersection points modified to be switches: at an intersection point of C_i with D_j , have the switch curve to the right as viewed along C_i (and thus to the left as viewed along D_j). This convention guarantees that the train track is invariant under positive Dehn twists along the C_i 's and under negative Dehn twists along the D_j 's.

At an intersection point of C_i with D_j , there are two ends of D_j meeting C_i , and these (possibly) correspond to separate branches of the train track τ . If we perform a Dehn twist along a curve just to one side of C_i , then one of these branches is modified to one which goes all the way around C_i , and if we Dehn twist on the other side of C_i , the other is. We will be twisting along C_i at least twice, so we may ensure that we twist on each side of C_i at least once. This will ensure that the incidence matrix M is Perron–Frobenius.

The incidence matrix M is simply the product of the incidence matrices for each curve we twist along (the matrix depends on the side of the curve we twist on). Thus $\text{tr}(M)$ can be computed as the trace of this product of matrices. The other terms in Theorem 5.1 are also combinatorial. We have:

Theorem 6.1 *Let $\{C_i\}$ and $\{D_j\}$ be collections of noncontractible simple closed curves on Σ with the C_i mutually disjoint and the D_j mutually disjoint such that every region of the complement $\Sigma - \coprod_i C_i - \coprod_j D_j$ is a disk with more than two edges.*

Let $\phi = \prod_j \tau_{D_j}^{-2} \prod_i \tau_{C_i}^2$ and let h be the mapping class of ϕ . Then $HF_(\phi) = HF_*(h)$ has rank*

$$4 \sum_{i,j} |C_i \cap D_j|^2 + 2g - 2.$$

Proof We calculate each of the terms in the formula given in Theorem 5.1. In order to find $\text{tr}(M)$, we express M as a product of the incidence matrix for $\prod_j \tau_{D_j}^{-2}$ with the incidence matrix for $\prod_i \tau_{C_i}^2$.

First we fix some notation. We have curves C_i and D_j and branches c_α and d_β which lie on some C_i and some D_j , respectively. We let $C_{[\alpha]}$ denote the curve C_i such that

²⁰This train track will still be invariant if some curves show up only once, but the incidence matrix might not be Perron–Frobenius and we might need to perform some moves as in [1].

c_α lies on C_i and likewise for $D_{[\beta]}$. We will write our incidence matrices in block form with the entries associated to the c_α and the d_β separated.

The incidence matrix for $\prod_i \tau_{C_i}^2$ is of the form

$$M_C = \begin{bmatrix} I & 0 \\ A & I \end{bmatrix}$$

where A is a matrix with entries $A_{\beta\alpha}$ equal to the number of times the branch d_β passes over c_α . This equals the number of times d_β passes over the curve $C_{[\alpha]}$, which equals the number of times the branch d_β intersects the curve $C_{[\alpha]}$, that is, how many of its ends lie on $C_{[\alpha]}$, with each end counting separately. We denote this as $\#(d_\beta \cap C_{[\alpha]})$, though this may not equal their set-theoretic intersection number when the two ends of d_β coincide. We get a contribution from each end of d_β because we've twisted on either side of each C_i . Likewise, the incidence matrix for $\prod_j \tau_{D_j}^{-2}$ is of the form

$$M_D = \begin{bmatrix} I & B \\ 0 & I \end{bmatrix}$$

where B is a matrix with entries $B_{\alpha\beta}$ equal to the number of ends of the branch c_α which lie on $D_{[\beta]}$.

We have
$$M = M_D M_C = \begin{bmatrix} I & B \\ 0 & I \end{bmatrix} \begin{bmatrix} I & 0 \\ A & I \end{bmatrix} = \begin{bmatrix} I + BA & B \\ A & I \end{bmatrix}.$$

Thus

$$(9) \quad \text{tr}(M) = b + \sum_{\alpha,\beta} B_{\alpha\beta} A_{\beta\alpha},$$

where b is the total number of branches.

We claim that

$$(10) \quad \sum_{\alpha\beta} B_{\alpha\beta} A_{\beta\alpha} = 4 \sum_{i,j} |C_i \cap D_j|^2.$$

To see this, note that

$$\begin{aligned} \sum_{\beta \in [\beta]} B_{\alpha\beta} A_{\beta\alpha} &= \sum_{\beta \in [\beta]} \#(c_\alpha \cap D_{[\beta]}) \#(d_\beta \cap C_{[\alpha]}) \\ &= \#(c_\alpha \cap D_{[\beta]}) 2\#(D_{[\beta]} \cap C_{[\alpha]}), \end{aligned}$$

where the factor of two comes from the fact that each intersection point of $D_{[\beta]}$ with $C_{[\alpha]}$ is associated to two ends of branches in $D_{[\beta]}$. Now summing over $\alpha \in [\alpha]$ gives Equation (10). We furthermore note that the number of branches b is twice the number

of intersections, ie $b = 2 \sum_{i,j} |C_i \cap D_j|$, because each intersection involves four branch ends (of which there are two per branch).

It remains to compute the remaining terms in [Theorem 5.1](#). First we consider the start and end of each branch (from its ordered list σ). Consider the branch c_α . It is unchanged by twisting along the C_i 's. When we twist along the D_j 's, c_α covers the D -curve at its initial end, then itself, then the D -curve at its final end. In particular, s_α is zero and there are no contributions to flips from the c_α 's alone. Now consider the branch d_β . When we twist along the C_i 's, it is modified and covers the C -curve at its initial end, then itself, then the C -curve at its final end, but also it starts and ends on those two C -curves at its initial and final switches and in the same orientation (at the switch) as itself at the switches. Call these initial and final branches c_{init} and c_{fin} . Then when we twist along the D_j 's, by the same reasoning, the initial switch that c_{init} runs over is d_β , and the final switch that c_{fin} runs over is d_β , each with the appropriate orientation. Thus the ordered list σ_β begins and ends with d_β (with the correct orientations). Thus $s_\beta = 2/2 = 1$ and also there are no flips.

Thus we get a (negative) contribution of $\sum_i s_i = \sum_\beta 1$. This equals the number of branches in the D -curves. Because each switch has two d -branches and two c -branches incident to it, and each branch has two ends, we get that this equals the number of switches, which is $\sum_{i,j} |C_i \cap D_j|$.

It remains to examine the singularities. These correspond to regions in the complement of the C 's and D 's. We can determine the number of these using Euler characteristic: the number of switches minus the number of branches plus the number of these regions equals the Euler characteristic of the surface, which is $2 - 2g$. We know the number of switches is $\sum_{i,j} |C_i \cap D_j|$ and the number of branches is $2 \sum_{i,j} |C_i \cap D_j|$. Thus there are $2 - 2g + \sum_{i,j} |C_i \cap D_j|$ singularities.

We claim that each of these singularities is fixed and unrotated. To determine this requires understanding the transverse list ν_β for each branch d_β : the initial and final branches in this list (with orientation) determine which regions are mapped to the regions on either side of d_β and how rotated they are.²¹ Thus we must show that the initial and final branches in this list are both d_β with the correct orientation.

When we perform a Dehn twist along $D_{[\beta]}$, the branches that cover a given branch d_β are those c_α 's which meet $D_{[\beta]}$. These cover d_β on the side corresponding to the side of $D_{[\beta]}$ we've twisted on. The outermost branches covering d_β (ie the initial and final branches in ν_β) are those which meet d_β at a switch which are on the opposite side of the switch from d_β . Thus the outermost branches covering d_β are the outermost

²¹Note that because each region borders some d_β 's, we won't need to consider the c_α 's separately.

branches covering these c_α 's after twisting along $C_{[\alpha]}$ (for then these will be carried along with c_α and cover d_β when we twist along $D_{[\beta]}$). These are likewise the d -branches meeting the c_α at a switch on the opposite side of the switch. One of these is d_β , and we can see that in fact this branch is the outermost when c_α covers d_β , as desired.

Summing up, we've shown that

$$\begin{aligned} HF_*(\phi) &= \text{tr}(M) + f - \sum_i s_i + \sum_p (r_p - u_p) \\ &= \left(2 \sum_{i,j} |C_i \cap D_j| + 4 \sum_{i,j} |C_i \cap D_j|^2 \right) + 0 \\ &\quad - \sum_{i,j} |C_i \cap D_j| - \left(2 - 2g + \sum_{i,j} |C_i \cap D_j| \right) \\ &= 4 \sum_{i,j} |C_i \cap D_j|^2 + 2g - 2. \quad \square \end{aligned}$$

Remark 6.2 This easily extends to $\prod_j \tau_{D_j}^{-2\ell} \prod_i \tau_{C_i}^{2k}$. Only $\text{tr}(M)$ is changed (we replace A by kA and B by ℓB). We get $HF_*(\phi) = 4k\ell \sum_{i,j} |C_i \cap D_j|^2 + 2g - 2$. In fact, any composition of positive powers of $\prod_j \tau_{D_j}^{-2}$ and positive powers of $\prod_i \tau_{C_i}^2$ can be computed in this manner. The correction terms again remain the same, so we only need to know how to compute $\text{tr}(M)$.

References

- [1] **M Bestvina, M Handel**, *Train-tracks for surface homeomorphisms*, *Topology* 34 (1995) 109–140 [MR1308491](#)
- [2] **J S Birman, M E Kidwell**, *Fixed points of pseudo-Anosov diffeomorphisms of surfaces*, *Adv. in Math.* 46 (1982) 217–220 [MR679909](#)
- [3] **F Bourgeois, Y Eliashberg, H Hofer, K Wysocki, E Zehnder**, *Compactness results in symplectic field theory*, *Geom. Topol.* 7 (2003) 799–888 [MR2026549](#)
- [4] **A J Casson, S A Bleiler**, *Automorphisms of surfaces after Nielsen and Thurston*, *London Math. Soc. Student Texts* 9, Cambridge Univ. Press (1988) [MR964685](#)
- [5] **K Cieliebak, K Mohnke**, *Compactness for punctured holomorphic curves*, *J. Symplectic Geom.* 3 (2005) 589–654 [MR2235856](#)
- [6] **V Colin, K Honda**, *Reeb vector fields and open book decompositions* [arXiv:0809.5088](#)

- [7] **A Cotton-Clay**, *A sharp bound on fixed points of area-preserving surface diffeomorphisms*, in preparation
- [8] **S Dostoglou, D A Salamon**, *Self-dual instantons and holomorphic curves*, Ann. of Math. (2) 139 (1994) 581–640 [MR1283871](#)
- [9] **E Eftekhary**, *Floer homology of certain pseudo-Anosov maps*, J. Symplectic Geom. 2 (2004) 357–375 [MR2131640](#)
- [10] **Y Eliashberg, A Givental, H Hofer**, *Introduction to symplectic field theory*, GAFA 2000 (Tel Aviv, 1999), Geom. Funct. Anal. (2000) 560–673 [MR1826267](#)
- [11] **O Fabert**, *Contact homology of Hamiltonian mapping tori* [arXiv:math/0609406v3](#)
- [12] **A Fathi, F Laudenbach, V Poenaru** (editors), *Travaux de Thurston sur les surfaces*, second edition, Astérisque 66, Soc. Math. France, Paris (1991) [MR568308](#) Séminaire Orsay, With an English summary
- [13] **A L Fel'shtyn**, *Floer homology, Nielsen theory, and symplectic zeta functions*, Tr. Mat. Inst. Steklova 246 (2004) 283–296 [MR2101299](#)
- [14] **A Floer**, *Morse theory for Lagrangian intersections*, J. Differential Geom. 28 (1988) 513–547 [MR965228](#)
- [15] **A Floer**, *Symplectic fixed points and holomorphic spheres*, Comm. Math. Phys. 120 (1989) 575–611 [MR987770](#)
- [16] **A Floer, H Hofer, D Salamon**, *Transversality in elliptic Morse theory for the symplectic action*, Duke Math. J. 80 (1995) 251–292 [MR1360618](#)
- [17] **R Gautschi**, *Floer homology of algebraically finite mapping classes*, J. Symplectic Geom. 1 (2003) 715–765 [MR2039162](#)
- [18] **M Hutchings**, *An index inequality for embedded pseudoholomorphic curves in symplectizations*, J. Eur. Math. Soc. (JEMS) 4 (2002) 313–361 [MR1941088](#)
- [19] **M Hutchings**, *Reidemeister torsion in generalized Morse theory*, Forum Math. 14 (2002) 209–244 [MR1880912](#)
- [20] **M Hutchings**, *Floer homology of families. I*, Algebr. Geom. Topol. 8 (2008) 435–492 [MR2443235](#)
- [21] **M Hutchings, M Sullivan**, *The periodic Floer homology of a Dehn twist*, Algebr. Geom. Topol. 5 (2005) 301–354 [MR2135555](#)
- [22] **M Hutchings, M Sullivan**, *Rounding corners of polygons and the embedded contact homology of T^3* , Geom. Topol. 10 (2006) 169–266 [MR2207793](#)
- [23] **S Jabuka, T Mark**, *Heegaard Floer homology of certain mapping tori*, Algebr. Geom. Topol. 4 (2004) 685–719 [MR2100677](#)
- [24] **B J Jiang, J H Guo**, *Fixed points of surface diffeomorphisms*, Pacific J. Math. 160 (1993) 67–89 [MR1227504](#)

- [25] **P Kronheimer, T Mrowka**, *Monopoles and three-manifolds*, New Math. Monogr. 10, Cambridge Univ. Press (2007) [MR2388043](#)
- [26] **E Lebow**, *Embedded contact homology of 2-torus bundles over the circle*, PhD thesis, University of California, Berkeley (2007)
- [27] **Y-J Lee**, *Heegaard splittings and Seiberg–Witten monopoles*, from: “Geometry and topology of manifolds”, (U Boden, Hans, I Hambleton, AJ Nicas, BD Park, editors), Fields Inst. Commun. 47, Amer. Math. Soc. (2005) 173–202 [MR2189932](#)
- [28] **Y-J Lee**, *Reidemeister torsion in Floer–Novikov theory and counting pseudo-holomorphic tori. I*, J. Symplectic Geom. 3 (2005) 221–311 [MR2199540](#)
- [29] **Y-J Lee**, *Reidemeister torsion in Floer–Novikov theory and counting pseudo-holomorphic tori. II*, J. Symplectic Geom. 3 (2005) 385–480 [MR2198782](#)
- [30] **Y-J Lee, CH Taubes**, *Periodic Floer homology and Seiberg–Witten Floer cohomology* [arXiv:0906.0383](#)
- [31] **G Liu, G Tian**, *On the equivalence of multiplicative structures in Floer homology and quantum homology*, Acta Math. Sin. (Engl. Ser.) 15 (1999) 53–80 [MR1701133](#)
- [32] **D McDuff, D Salamon**, *Introduction to symplectic topology*, second edition, Oxford Math. Monogr., The Clarendon Press, Oxford Univ. Press, New York (1998) [MR1698616](#)
- [33] **L Mosher**, *Train track expansions of measured foliations*, Preprint (2003)
- [34] **P Ozsváth, Z Szabó**, *Holomorphic disks and topological invariants for closed three-manifolds*, Ann. of Math. (2) 159 (2004) 1027–1158 [MR2113019](#)
- [35] **R C Penner**, *A construction of pseudo-Anosov homeomorphisms*, Trans. Amer. Math. Soc. 310 (1988) 179–197 [MR930079](#)
- [36] **R C Penner, J L Harer**, *Combinatorics of train tracks*, Annals of Math. Studies 125, Princeton Univ. Press (1992) [MR1144770](#)
- [37] **T Perutz**, *A hypercube for fixed-point Floer homology*, Lecture, UCLA/USC (2008)
- [38] **T Perutz**, *Lagrangian matching invariants for fibred four-manifolds. II*, Geom. Topol. 12 (2008) 1461–1542 [MR2421133](#)
- [39] **M Pozniak**, *Floer homology, Novikov rings and clean intersections*, PhD thesis, University of Warwick (1994)
- [40] **J Robbin, D Salamon**, *The spectral flow and the Maslov index*, Bull. London Math. Soc. 27 (1995) 1–33 [MR1331677](#)
- [41] **D Salamon, E Zehnder**, *Morse theory for periodic solutions of Hamiltonian systems and the Maslov index*, Comm. Pure Appl. Math. 45 (1992) 1303–1360 [MR1181727](#)
- [42] **M Schwarz**, *Cohomology operations from S^1 -cobordisms in Floer homology*, PhD thesis, ETH Zürich (1995)

- [43] **P Seidel**, *The symplectic Floer homology of a Dehn twist*, Math. Res. Lett. 3 (1996) 829–834 [MR1426539](#)
- [44] **P Seidel**, *More about vanishing cycles and mutation*, from: “Symplectic geometry and mirror symmetry (Seoul, 2000)”, (K Fukaya, Y-G Oh, K Ono, G Tian, editors), World Sci. Publ., River Edge, NJ (2001) 429–465 [MR1882336](#)
- [45] **P Seidel**, *Symplectic Floer homology and the mapping class group*, Pacific J. Math. 206 (2002) 219–229 [MR1924827](#)
- [46] **CH Taubes**, *Seiberg Witten and Gromov invariants for symplectic 4-manifolds*, (R Wentworth, editor), First Intl. Press Lecture Ser. 2, Intl. Press, Somerville, MA (2000) [MR1798809](#)
- [47] **CH Taubes**, *The Seiberg–Witten equations and the Weinstein conjecture*, Geom. Topol. 11 (2007) 2117–2202 [MR2350473](#)
- [48] **CH Taubes**, *Embedded contact homology and Seiberg–Witten Floer cohomology I–IV*, Preprint (2008)
- [49] **WP Thurston**, *On the geometry and dynamics of diffeomorphisms of surfaces*, Bull. Amer. Math. Soc. (N.S.) 19 (1988) 417–431 [MR956596](#)
- [50] **M Usher**, *Vortices and a TQFT for Lefschetz fibrations on 4-manifolds*, Algebr. Geom. Topol. 6 (2006) 1677–1743 [MR2263047](#)
- [51] **M-L Yau**, *A holomorphic 0-surgery model for open books*, Int. Math. Res. Not. (2007) Art. ID rnm041, 23pp [MR2346631](#)

Department of Mathematics, Harvard University
One Oxford St, Cambridge, MA 01238, USA

acotton@math.harvard.edu

<http://math.harvard.edu/~acotton>

Proposed: Peter Ozsváth

Seconded: Yasha Eliashberg, Leonid Polterovich

Received: 18 July 2008

Revised: 29 April 2009