

A GLOBAL LE CALVEZ-YOCCOZ PROPERTY OF AREA-PRESERVING SURFACE DIFFEOMORPHISMS AND THREE-DIMENSIONAL REEB FLOWS

DAN CRISTOFARO-GARDINER AND ROHIL PRASAD

ABSTRACT. We prove that for any monotone area-preserving diffeomorphism of a closed surface, or Reeb flow on a closed contact three-manifold with torsion Chern class, the complement of a compact invariant set is never minimal. As a corollary, we obtain that, under the same assumptions, there are infinitely many distinct proper compact invariant sets whose union is dense in the manifold. No genericity assumptions are required. The former class of systems includes all Hamiltonian diffeomorphisms of closed surfaces and the latter includes all Finsler geodesic flows on closed surfaces. We can view our results as generalizations, in the smooth symplectic setting, to higher genus surfaces and three-manifolds of results of Le Calvez–Yoccoz, Franks, and Salazar for homeomorphisms of the two-sphere. Our results also give such a generalization for Finsler geodesic flows on closed surfaces.

Along the way, we prove a result, in any dimension, of potentially independent interest, detecting invariant sets via sequences of low-action pseudoholomorphic curves with controlled topology; as a corollary, this generalizes Ginzburg–Gürel’s “crossing energy bound” for Floer cylinders to punctured holomorphic curves, of arbitrary topology, in symplectizations, resolving an open question posed by them. Another feature of the argument which also might be of independent interest is a probabilistic result stating that almost all ECH/PFH U -map curves have Euler characteristic at least -2 under the above assumptions.

CONTENTS

1. Introduction	1
2. Proofs of main dynamical results	10
3. Low-action curves from embedded contact homology	16
4. Low-action curves from periodic Floer homology	24
5. Invariant sets from low-action holomorphic curves	28
Appendix A. Verifying assumptions	47
References	49

1. INTRODUCTION

1.1. Le Calvez–Yoccoz phenomena. The detection and classification of compact invariant sets is a fundamental question in dynamical systems. For example, a natural question asks how much of the parameter space is seen by a given trajectory; the closure of any such trajectory is a compact invariant set. For diffeomorphisms of the circle

and smooth flows on the plane, the theorems of Denjoy [20] and Poincaré–Bendixson [5] give a nearly complete picture. The situation in higher-dimensions, on the other hand, is much more mysterious, with a far greater diversity in the possible behaviors. This paper is about the conservative setting, in one dimension higher than Denjoy and Poincaré–Bendixson: area-preserving diffeomorphisms of 2-manifolds and volume-preserving flows on 3-manifolds.

A first observation is that without further assumptions, there might be no interesting invariant sets at all:

Example 1.1. Let f be an irrational translation of a two-torus; then, this is area-preserving, but the only compact invariant set is the entire manifold. One can similarly consider an irrational volume-preserving flow on the three-torus.

The main theme of our paper is that when one adds some further natural conditions of a symplectic nature, the situation changes completely. Our inspiration comes from the following result of Le Calvez–Yoccoz. Recall that an invariant set U of a map or flow, closed or not, is called *minimal* if the orbit of each initial condition $p \in U$ is dense in U . A groundbreaking 1997 paper by Le Calvez–Yoccoz [42], improving on an earlier result of Handel [31], showed that for any homeomorphism of S^2 the complement of an invariant finite set of points is never minimal. Their result resolved the 2-dimensional case of an old question of Ulam from the Scottish Book [43, p. 208].

Our first results give a generalization of this, in the smooth symplectic setting, to higher-genus surfaces and 3-manifolds.

Theorem 1. *Let Σ be a closed, oriented surface and let $\phi : \Sigma \rightarrow \Sigma$ be any monotone area-preserving diffeomorphism. Then for any proper compact invariant set $\Lambda \subset \Sigma$, the complement $\Sigma \setminus \Lambda$ is not minimal.*

Theorem 2. *Let Y be a closed, oriented 3-manifold equipped with a co-oriented contact structure ξ with torsion first Chern class. Let λ be any contact form defining ξ and let $\phi = \{\phi^t\}_{t \in \mathbb{R}}$ denote the Reeb flow of λ . Then for any proper compact invariant set $\Lambda \subset Y$, the complement $Y \setminus \Lambda$ is not minimal.*

What is novel about these theorems in the context of conservative dynamics is the level of generality. The simplest kinds of compact invariant sets are *periodic orbits*: this means that the invariant set is homeomorphic to a circle (in the case of flows) or a finite set of points with a transitive action (in the discrete setting). Previous results have established theorems like the above under strong dynamical assumptions such as the existence of only finitely many periodic orbits [29, 9]; we discuss this further in Remark 1.6. That these theorems hold much more generally, at least in low-dimensions, seems to us to be a quite new and perhaps unexpected phenomenon.

The assumptions of these theorems also apply to broad classes of systems in Hamiltonian dynamics. For example, any Hamiltonian diffeomorphism of a closed symplectic surface is monotone. As we will see, any geodesic flow on a closed Finsler surface corresponds to a Reeb flow with torsion first Chern class. Any Reeb flow on a rational homology 3-sphere of course also has torsion Chern class, as does any Reeb flow on a closed 3-manifold with contact structure supporting a contact Anosov flow [32, Theorem 4.1].

Theorem 1 and Theorem 2 guarantee the existence of an abundance of compact invariant sets. Moreover, the compact invariant sets are spread out in the manifold. For example, we obtain the following corollaries.

Corollary 1.1. *Under the assumptions of Theorem 1, the map ϕ has infinitely many distinct proper compact invariant sets whose union is dense in Σ .*

Corollary 1.2. *Under the assumptions of Theorem 2, the Reeb flow has infinitely many distinct proper compact invariant sets whose union is dense in Y .*

Corollary 1.1 and 1.2 are clearly false if one requires the compact invariant sets to be periodic orbits. For example, an irrational rotation of a two-sphere satisfies the assumptions of Theorem 1, but has just two periodic points.

Remark 1.3. The generality of the above theorems strongly precludes the space of possible improvements, without the imposition of additional restrictions. For example, one could hypothetically ask whether the invariant sets we detect support interesting invariant measures. However, Anosov–Katok famously constructed [2] an area-preserving diffeomorphism of S^2 whose invariant measures are as simple as possible: the only ergodic invariant measures are a pair of fixed points and the area measure. Corollary 1.1 applies to the Anosov–Katok example to produce many distinct proper compact invariant sets, but they all must therefore support essentially the same ergodic invariant measures.

Remark 1.4. The novelty of Theorems 1 and 2 is that they make no genericity assumptions at all. Indeed, prior results on the closing lemma (see [40, 3, 19, 22]) show that a C^∞ -generic system of the type we consider has a dense set of periodic points.

Remark 1.5. The arguments in Le Calvez - Yoccoz are local in nature, and produce for example invariant sets near any irrationally elliptic fixed point; see for example the exposition in [9]. As we will see, the invariant sets we detect arise from fundamentally different considerations: roughly speaking, they produce invariant sets near any proper compact invariant set containing all periodic orbits, hence have a global character.

1.2. Geodesic flows on surfaces. In the setting of Riemannian or Finsler manifolds, it is natural to ask how much of the manifold is visited by a given geodesic. Perhaps the simplest dichotomy in this direction is between the dense and non-dense geodesics. At one extreme, one could imagine that every geodesic is dense; this can not occur, because a closed geodesic can not be dense, but it is natural to wonder how far off it is from the actual behavior. In fact, for surfaces we have the following result contrasting this sharply:

Theorem 3. *Let F be any closed Finsler surface. Then there exists an infinite collection G of geodesics such that*

- (a) $\gamma(\mathbb{R})$ is not dense in F for any $\gamma \in G$;
- (b) The union $\bigcup_{\gamma \in G} \gamma(\mathbb{R})$ is dense in F ;
- (c) The closures $\overline{\gamma(\mathbb{R})}$ and $\overline{\gamma'(\mathbb{R})}$ are distinct for any pair of distinct elements $\gamma, \gamma' \in G$.

In other words, a dense set of points have a non-dense geodesic going through them; and, moreover, these geodesics all visit different sections of the surface.

In the negative curvature case, or the generic case, the above theorem is well-known, since in fact the stronger statement holds that the closed geodesics are dense. However, we make no curvature assumption at all, and in this generality this property seems to be quite new. Moreover, the theorem is false if one requires the geodesics to parameterize closed loops. For example, there exist Finsler metrics on S^2 with exactly two geometrically distinct closed geodesics [2].

In a different direction, our methods also give a new proof of an important recent result of Contreras–Mazzucchelli [12], giving conditions for a geodesic flow to be Anosov. We discuss this further in §1.6.1 below.

1.3. The work of Franks and Salazar. In fact, Theorem 1 and Theorem 2 follow from slightly more general (but slightly harder to state) results, which we now explain. This level of generality is also important for the applications to geodesic flows above.

Shortly after the work by Le Calvez–Yoccoz, Franks discovered the following refinement of their theorem. Recall that a compact invariant set Λ of a homeomorphism or flow on a compact manifold is called *locally maximal* if any sufficiently C^0 -close compact invariant set must be contained in Λ . If a compact invariant set Λ is *not* locally maximal, then any neighborhood U of Λ contains a point $z \notin U \setminus \Lambda$ with orbit closure contained in U . Franks [26] showed that for any homeomorphism of S^2 , the union of periodic points is either infinite or not locally maximal. A subsequent refinement in the conservative case by Salazar [49] showed that for any area-preserving homeomorphism of S^2 and any compact invariant set $\Lambda \subseteq S^2$ containing all periodic points, either $\Lambda = S^2$ or Λ is not locally maximal. We are able to generalize these results in the smooth symplectic case as well:

Theorem 4. *Let Σ be a closed, oriented surface and let $\phi : \Sigma \rightarrow \Sigma$ be any monotone area-preserving diffeomorphism. Then for any compact invariant set $\Lambda \subseteq \Sigma$ containing all periodic orbits of ϕ , either $\Lambda = \Sigma$ or Λ is not locally maximal.*

Theorem 5. *Let Y be a closed, oriented 3-manifold equipped with a co-oriented contact structure ξ with torsion first Chern class. Let λ be any contact form defining ξ and let $\{\phi^t\}_{t \in \mathbb{R}}$ denote the Reeb flow of λ . Then for any compact invariant set $\Lambda \subseteq Y$ containing all periodic orbits of $\{\phi^t\}_{t \in \mathbb{R}}$, either $\Lambda = Y$ or Λ is not locally maximal.*

As we will explain, Theorem 4 implies Theorem 1. The analogous chain of reasoning holds starting from Theorem 5.

Remark 1.6. Related results for Hamiltonian diffeomorphisms of $\mathbb{C}P^n$ with finitely many periodic points and dynamically convex Reeb flows on S^{2n+1} with finitely many closed orbits were respectively proved by Ginzburg–Gürel [29] and Cineli–Ginzburg–Gürel–Mazzucchelli [9]. There is no dimensional restriction in these results, and we say a bit more about this in connection to our results in §1.6.4.

Remark 1.7. Le Calvez–Yoccoz, Franks, and Salazar used very different methods from ours. Le Calvez–Yoccoz and Franks argue by contradiction with the key technical step that after possibly iterating the map, the Lefschetz index of the fixed point set is ≤ 0 ,

contradicting the condition that $\chi(S^2) > 0$; Salazar’s argument also uses this. Since any oriented surface of higher genus does not have positive Euler characteristic, we do not know how to approach Theorem 4 using their methods; the case of flows on three-manifolds via their methods is also unclear to us.

Remark 1.8. Our proof of Theorem 3, after making minor modifications, yields a geodesic flow analogue of Franks’ theorem: either the union of closed geodesics is dense in F or we can find non-dense geodesics which are arbitrarily nearby its closure.

1.4. Invariant sets from low-action holomorphic curves. We now explain a general theorem that we prove, applicable in any dimension and of independent interest, for extracting invariant sets; it is used to prove all of the above dynamical results. One can view this as a global detection method via “low-action” holomorphic curves.

1.4.1. The setup. We work in the general setting of punctured holomorphic curves in symplectizations $\mathbb{R} \times Y$ over framed Hamiltonian manifolds. A *framed Hamiltonian structure* on a smooth, oriented manifold Y of dimension $2n + 1 \geq 3$ is a pair $\eta = (\lambda, \omega)$ of a 1-form λ and a closed 2-form ω such that $\lambda \wedge \omega^n > 0$. The *Hamiltonian vector field* R_η is defined implicitly by

$$\lambda(R_\eta) \equiv 1, \quad \omega(R_\eta, -) \equiv 0.$$

The flow of R_η preserves ω and the volume form $\lambda \wedge \omega^n$. This setup is an abstraction of many important classes of systems in symplectic and conservative dynamics, including mapping torii of symplectic diffeomorphisms, Reeb and stable Hamiltonian flows, and volume-preserving flows on three-manifolds. For example, if $\omega = d\lambda$, then λ is a contact form and R_η is its Reeb vector field.

We follow the classical setup of holomorphic curve theory in symplectizations introduced by Hofer. Fix a Riemann surface (C, j) . A *J-holomorphic curve* is a proper smooth map $u : C \rightarrow \mathbb{R} \times Y$ satisfying the Cauchy–Riemann equation

$$J \circ Du = Du \circ j$$

where J is an η -adapted almost-complex structure on $\mathbb{R} \times Y$. This is a translation-invariant almost-complex structure restricting to a compatible almost-complex structure on the symplectic bundle $(\ker(\lambda), \omega)$ and sending $-R_\eta$ to the vector field ∂_a defined by the \mathbb{R} -coordinate on $\mathbb{R} \times Y$. We say u is *standard* if the domain C is homeomorphic to the complement of a finite subset of a closed Riemann surface. The geometry of a J -holomorphic curve in $\mathbb{R} \times Y$ is controlled by the *action* and *Hofer energy*¹, defined respectively as

$$\mathcal{A}(u) := \int_C u^* \omega, \quad \mathcal{E}(u) := \sup_{s \in \mathbb{R}} \int_{C \cap u^{-1}(\{s\} \times Y)} u^* \lambda.$$

The action controls how far on average the tangent planes of C , which are J -invariant, are from the vertical plane spanned by ∂_a and R_η . Therefore, a low-action holomorphic curve should approximate the vector field R_η very well. The Hofer energy is, informally, the maximum length of the level sets of C in $\mathbb{R} \times Y$.

¹This is not Hofer’s original definition. Finiteness of $\mathcal{E}(C)$, however, is equivalent to finiteness of the original Hofer energy.

1.4.2. *The limit set.* A key object in our method is the “limit set” of a sequence of holomorphic curves in a symplectization, which we now introduce. For any closed, smooth, odd-dimensional manifold Y , write $\mathcal{D}(Y)$ for the space of pairs (η, J) where η is a framed Hamiltonian structure and J is an η -adapted almost-complex structure. Equip it with the topology of C^∞ convergence in both η and J . Fix a pair $(\eta, J) \in \mathcal{D}(Y)$ and a sequence $\{(\eta_k, J_k)\}$ in $\mathcal{D}(Y)$ converging to (η, J) . Fix a sequence $\{u_k : C_k \rightarrow \mathbb{R} \times Y\}$ where u_k is J_k -holomorphic for each k .

It is convenient to define $X := (-1, 1) \times Y$, and to define for any $s \in \mathbb{R}$ the shift map $\tau_s : \mathbb{R} \times Y \rightarrow \mathbb{R} \times Y$, mapping $(t, y) \mapsto (t - s, y)$. Define the *limit set* \mathcal{X} of the sequence $\{u_k\}$ to be the collection of all closed subsets $K \subseteq (-1, 1) \times Y$ arising as subsequential Hausdorff limits as $k \rightarrow \infty$ of height-2 slices of u_k . That is, there exists a sequence $\{s_k\}$ of real numbers such that a subsequence of

$$\tau_{s_k} \cdot \left(u_k(C_k) \cap (s_k - 1, s_k + 1) \times Y \right) \subseteq X$$

converges in the Hausdorff topology to K . The limit set \mathcal{X} is a subset of $\mathcal{K}(X)$, the space of all closed subsets of X equipped with the topology of Hausdorff convergence. See § 2.2 for a definition of the Hausdorff topology.

The limit set has the following very important connectivity property:

Proposition 1.9. *Fix a closed, smooth, oriented, odd-dimensional manifold Y and fix a sequence $\{(\eta_k, J_k)\}$ converging in $\mathcal{D}(Y)$ to a pair (η, J) . Let $\{u_k : C_k \rightarrow Y\}$ denote a sequence where u_k is a standard J_k -holomorphic curve for each k . Then there exists a subsequence $\{u_{k_j}\}$ whose limit set \mathcal{X} is connected with respect to the Hausdorff topology.*

A harder theorem, which we prove, is the following:

Theorem 6. *Fix a closed, smooth, oriented, odd-dimensional manifold Y and a sequence $\{(\eta_k, J_k)\}_{k \geq 1}$ converging in $\mathcal{D}(Y)$ to a pair (η, J) . Let $\{u_k : C_k \rightarrow Y\}_{k \geq 1}$ denote a sequence where u_k is a standard J_k -holomorphic curve for each k and let $\mathcal{X} \subseteq \mathcal{K}(X)$ denote their limit set. Assume in addition that*

$$\lim_{k \rightarrow \infty} \mathcal{A}(u_k) = 0 \quad \text{and} \quad \inf_k \chi(C_k) > -\infty$$

. *Then every set $\bar{\Lambda} \in \mathcal{X}$ is equal to $(-1, 1) \times \Lambda$, where $\Lambda \in \mathcal{K}(Y)$ is non-empty and invariant under the flow of the Hamiltonian vector field R_η .*

The novelty² of Theorem 6 is that it extracts invariant sets without requiring that the Hofer energies $\{\mathcal{E}(u_k)\}$ admit a finite k -independent upper bound, or even that any of the Hofer energies $\mathcal{E}(u_k)$ are finite. Bounds on Hofer energy are a standard assumption in the vast majority of the symplectic field theory literature; the only exceptions known to us are [24, 47] about “feral” curves³. Our proof of Theorem 6 is inspired by, and builds on, ideas in these works, though one should emphasize that the setting here, of

²Indeed, in the case where $\sup_{k \geq 1} \mathcal{E}(u_k) \leq E$ for some finite E , Theorem 6 follows from the original work of Hofer, and moreover one obtains the stronger conclusion that any $\bar{\Lambda} \in \mathcal{X}$ is a cylinder over a finite union of periodic orbits.

³Loosely speaking, these are curves with unbounded Hofer energy; they do not play a role in our arguments.

a sequence of standard curves with action tending to zero, is quite different. As we will see in the proof, this leads to new topological challenges, with an essential point being that the curves we accommodate can have genus and many punctures. Indeed, it is essential for our arguments to be able to allow such curves.

Remark 1.10. The Euler characteristic $\chi(C_k)$ is finite for each k because we are assuming each curve is standard. The assumption of a finite k -independent lower bound on $\chi(C_k)$ is essential for our proof of Theorem 6. Removing this assumption would be of interest: it would shorten the proofs of our main dynamical Theorems 4 and 5 and extend them to any rational area-preserving diffeomorphism of a closed surface and any Reeb flow on a closed 3-manifold, respectively.

Remark 1.11. Fish-Hofer in [24, Definition 4.46] also define a kind of limit set (the “x-limit set”). Our notion of limit set differs from theirs in several ways which are crucial for our arguments. In particular, [24, Definition 4.46] has an asymptotic condition, which we do not want to require for our applications. And, their limit set is a single invariant set, while ours is a connected family of invariant sets. This distinction is actually a key innovation, since the connectedness is exploited to show that such families contain many distinct invariant sets.

1.5. The crossing energy theorem in symplectizations. To explain our final new result, which is more technical, we need to recall the “crossing energy theorem” for Hamiltonian diffeomorphisms. The crossing energy theorem is a powerful tool introduced by Ginzburg–Gürel [27]. Recall that a neighborhood U of a locally maximal compact invariant set Λ of a homeomorphism or flow is *isolating* if any compact invariant set $\Lambda' \subset U$ is a subset of Λ . The crossing energy theorem asserts that if Λ is a locally maximal invariant set of a Hamiltonian diffeomorphism (for example, a hyperbolic fixed point), U is an isolating neighborhood, and $V \subset U$ is such that $\bar{V} \subset U$, then any “Floer cylinder” crossing the shell $U \setminus V$ must have a uniform lower bound on its Floer energy. Analogues have been established for generating functions [1], gradient flow lines of the energy functional on loop space [30], and Floer cylinders in symplectic homology [9]. It is central to many results, such as Conley conjecture type results on the multiplicity of periodic points [4, 27, 28], dynamics of Hamiltonian and Reeb pseudorotations [29, 9], and the study of topological entropy via barcode invariants [8, 30].

Thus, one would like to generalize it for Reeb flows. This was first posed as a question in 2012 by Ginzburg–Gürel, but prior to our work it had not been clear how to prove it; see e.g. the discussion in [30, p. 4]. In fact, Theorem 6, which uses new tools that did not exist at the time of [27], provides this theorem as a corollary, for any framed Hamiltonian flow and for holomorphic curves with domain any closed Riemann surface with finitely many punctures removed. Here is the precise statement:

Theorem 7. *Fix a closed framed Hamiltonian manifold (Y, η) and an η -adapted almost-complex structure J . Let Λ be a locally maximal R_η -invariant set, U an isolating neighborhood and $V \subset U$ such that $\bar{V} \subset U$. Fix an integer $T > 0$ and let $u : C \rightarrow \mathbb{R} \times Y$ be any standard J -holomorphic curve with $\chi(C) \geq -T$. Then there is a constant $c = c(\eta, J, \Lambda, U, V, T) > 0$ such that*

$$\mathcal{A}(u) > c > 0$$

whenever there exists $s_-, s_+ \in \mathbb{R}$ with

$$(1) \quad u(C) \cap \{s_-\} \times Y \subset V, \quad u(C) \cap \{s_+\} \times Y \not\subset U.$$

In analogy with the progress for Hamiltonian diffeomorphisms summarized above, one hopes that Theorem 7 has many potential applications concerning the dynamics of Reeb flows.

Remark 1.12. Theorem 7 directly generalizes the crossing energy theorem for Hamiltonian diffeomorphisms from [27, 29]. Given a Hamiltonian diffeomorphism ϕ , the mapping torus Y_ϕ carries a natural framed Hamiltonian structure η such that R_η generates the suspension flow. There is an explicit correspondence between Floer cylinders for a choice of Hamiltonian H generating ϕ and holomorphic cylinders in $\mathbb{R} \times Y_\phi$ ⁴. The Floer energy of a Floer cylinder is equal to the action of its corresponding holomorphic cylinder. Thus, it suffices to apply Theorem 7 for $Y = Y_\phi$ and pass through this correspondence. On the other hand, there are also very interesting crossing energy theorems proved in the recent works [9, 10, 45]. The results in [9, 10] are for Floer cylinders in completed Liouville domains and the results in [45] are for Floer strips with Lagrangian boundary conditions; these do not similarly follow from Theorem 7.

1.6. Further remarks.

1.6.1. *The C^2 -stability conjecture.* A recent breakthrough result by Contreras–Mazzucchelli [12] established the C^2 -stability conjecture for geodesic flows on Riemannian surfaces, namely that the C^2 -structurally stable flows are exactly the Anosov ones. Our methods give an alternative proof of the following key ingredient in their proof, giving a sufficient criterion for a geodesic flow to be Anosov:

Corollary 1.13 ([12, Theorem D]). *Let F be a closed Finsler surface, and let $\mathcal{P} \subseteq SF$ denote the closure of the union of closed orbits of the geodesic flow. Then if i) \mathcal{P} is uniformly hyperbolic and ii) the geodesic flow is Kupka–Smale, the geodesic flow is Anosov.*

It is important to note that Corollary 1.13 is only a special case of [12, Theorem D], which holds for all Reeb flows on closed 3-manifolds. However, the special case suffices for the application to the stability conjecture.

1.6.2. *Non-rational maps.* We note that, in the case of the torus, our results are close to being sharp. Indeed, any non-rational ϕ is either i) Hamiltonian isotopic to a translation $(x, y) \mapsto (x+a, y+b)$ where $(a, b) \notin \mathbb{Q}^2$ or ii) Hamiltonian isotopic to a smooth conjugate of an affine map

$$(x, y) \mapsto (x + ny, y + b),$$

where n is a nonzero integer and $b \notin \mathbb{Q}$. The examples in case i) have no proper compact invariant sets when both a and b are irrational and the examples in case ii) never have any proper compact invariant subsets.

It would be interesting to see whether our results hold for rational maps in higher genus.

⁴An explicit derivation for the 2-disk, which generalizes to arbitrary symplectic manifolds, can be found in [6, Lemma 20].

1.6.3. *Three-dimensional energy surfaces.* Density results such as Corollary 1.1 and Corollary 1.2 also hold for compact regular Hamiltonian hypersurfaces in some symplectic 4-manifolds, including \mathbb{R}^4 , without any contact type condition. This will be discussed in a future work [46].

1.6.4. *Higher dimensions.* We close with some speculations on the extension of our results to higher dimensions. As we mentioned in Remark 1.6, higher-dimensional versions of the theorems in this paper were proved in [29, 9], but with the dynamical assumption that the systems must have finitely many periodic orbits.

It would be very interesting to find the weakest possible dynamical assumptions for which our theorems extend to higher-dimensional Hamiltonian diffeomorphisms and Reeb flows. Our Theorem 6 on the extraction of invariant sets from low-action holomorphic curves with bounded topology works in all dimensions, and as mentioned in Remark 1.12, implies a “crossing energy bound” which is a key technical ingredient (among many) in [29]. Proposition 1.9 also clearly works in any dimension. However, our existence results for low-action holomorphic curves rely on deep properties of ECH/PFH, which are invariants defined for area-preserving surface diffeomorphisms and three-dimensional Reeb flows, respectively.

Finding low-action holomorphic curves with bounded topology in higher dimensions in high generality seems like it will require substantial new ideas. This is consistent with a more general theme in a range of problems in current symplectic research — ranging from the kind of questions considered in this paper to problems about the algebraic structure of certain homeomorphism groups to questions like symplectic packing stability [14, 13] — where one would like analogues of various properties related to ECH/PFH in higher dimensions.

On a more optimistic note, the fact that Theorem 6 works for *non-cylindrical* curves is an asset. It opens up for the first time the possibility of using powerful theories such as contact homology or SFT, which count non-cylindrical curves, to explore invariant sets of higher-dimensional Reeb flows. Indeed, as we have seen here, one needs to consider non-cylindrical curves in our arguments to get our results.

1.7. **Outline of article.** § 2 proves Theorems 4, 5, and 7. The proofs of Theorems 4 and 5 require Theorem 6 and two propositions (Propositions 2.2 and 2.3). These respectively assert that monotone area-preserving maps and Reeb flows of torsion contact forms have many low-action holomorphic curves with controlled topology. § 3 proves Proposition 2.3 using embedded contact homology. § 4 proves Proposition 2.2 using periodic Floer homology. § 5 proves Theorem 6.

1.8. **Acknowledgements.** Crucial discussions around this project occurred at the July 2023 workshop *Dynamical Systems* at the Oberwolfach Research Institute. We thank the institute for their hospitality. We also thank Erman Cineli, Viktor Ginzburg, Basak Gürel, and Marco Mazzucchelli for extremely helpful comments on an earlier version of our paper.

DCG: The origins of my interest in the topology of ECH curves are in my joint work [16] with Michael Hutchings and Dan Pomerleano. I thank them for many extremely helpful discussions about this topic. I also thank Umberto Hryniewicz and Hui Liu

for further discussions about this around our joint work [15]. I also thank Giovanni Forni for helpful discussions and I think Rich Schwartz for helpful comments on an earlier draft of this work. In addition, I thank the National Science Foundation for their support under agreements DMS-2227372 and DMS-2238091.

RP: I am grateful for useful conversations with Joel Fish, Viktor Ginzburg, Helmut Hofer, and Patrice Le Calvez. I would also like to acknowledge support from the National Science Foundation under agreement DGE-1656466 and from the Miller Institute at the University of California Berkeley.

2. PROOFS OF MAIN DYNAMICAL RESULTS

This section is primarily concerned with the proofs of Theorems 4 and 5. The proofs we give here rely on Theorem 6 and two propositions (Propositions 2.2 and 2.3), whose proofs are all deferred to subsequent sections. After these results are proved, we explain how they imply the other dynamical theorems stated in the introduction. We conclude the section with a proof of Theorem 7, which is a straightforward corollary of Theorem 6.

Before proceeding, let us explain the basic ideas behind the proofs. Proposition 2.3 asserts that any nondegenerate torsion contact form on a closed 3-manifold Y admits sequences of holomorphic curves in $\mathbb{R} \times Y$ with uniformly bounded topology, finite Hofer energy, and arbitrarily low action. Moreover, the curves in the sequence can be taken to pass through any point $(0, z)$ in $\mathbb{R} \times Y$ that is not on any closed Reeb orbit. The low action and bounded topology produce, via Theorem 6 and Proposition 1.9, a connected family \mathcal{X} of compact invariant sets. The finite Hofer energy and point constraints are then exploited, via an elementary topological argument, to conclude Theorem 5. The theorem is proved for degenerate contact forms by approximating them by nondegenerate contact forms and using the resulting holomorphic curves; to make this work, we require a certain amount of quantitative control over the relevant curves, which is why Proposition 2.3 is quantitative in nature. The path from Proposition 2.2 to Theorem 4 is completely analogous.

Remark 2.1. Our style of argument is robust enough to generalize to other settings where low action holomorphic curves with bounded topology are present. For example, it follows from [29] that the mapping torus of any Hamiltonian pseudorotation of $\mathbb{C}\mathbb{P}^n$ admits holomorphic cylinders of arbitrarily low action, with finite Hofer energy, passing through any point in the symplectization. Applying our argument proves the analogue of Theorems 4 and 5 for these maps.

2.1. Existence of low-action holomorphic curves with bounded topology. The following two key propositions show that mapping torii of monotone area-preserving surface diffeomorphisms and Reeb flows of torsion 3-dimensional contact forms have, after possibly making small perturbations, many low-action holomorphic curves with bounded topology. We start with the statement for area-preserving surface diffeomorphisms. Relevant notations and definitions are found in Section 4.

Proposition 2.2. *Let ϕ be a monotone area-preserving diffeomorphism of a closed symplectic surface (Σ, ω) . Let Y_ϕ denote the mapping torus of ϕ . There exists a positive integer $d_0 \geq 1$, depending only on the Hamiltonian isotopy class of ϕ , such that the*

following holds for all $d \geq d_0$ and any nondegenerate Hamiltonian perturbation ϕ' of ϕ : For any fixed $z' \in Y'_\phi$, not on any closed Reeb orbit, and generic choice of ϕ' -adapted J' , there exists a standard J' -holomorphic curve $u : C \rightarrow \mathbb{R} \times Y_{\phi'}$ such that:

- (a) $(0, z') \in u(C)$;
- (b) $\mathcal{E}(u) \leq d$;
- (c) $\mathcal{A}(u) \leq d^{-1/2}$;
- (d) $\chi(C) \geq -2$.

Next, the statement for Reeb flows. Relevant notations and definitions are found in Section 3.

Proposition 2.3. *Let λ be a torsion contact form on a closed 3-manifold Y . Then there exists a positive integer $k_0 \geq 0$ such that the following holds for any $k \geq k_0$ and any C^∞ -small nondegenerate perturbation λ' of λ : For any fixed $z' \in Y$, not on any closed Reeb orbit, and generic choice of λ' -adapted J' , there exists a standard J' -holomorphic curve $u : C \rightarrow \mathbb{R} \times Y$ such that:*

- (a) $(0, z') \in u(C)$;
- (b) $\mathcal{E}(u) \leq k^{3/4}$;
- (c) $\mathcal{A}(u) \leq k^{-1/16}$;
- (d) $\chi(C) \geq -2$.

Remark 2.4. In fact, we will see in the proofs of Proposition 2.2 and Proposition 2.3 that not only does there exist curves with the above properties, but that the ECH/PFH curves satisfy these properties under the assumptions of the proposition with probability 1. More precisely, as we will explain later, for k (resp. d) as in the statement of the propositions, the non-triviality of the “ U ”-map implies the existence of approximately k (resp. d) curves, and we can look at the proportion that satisfy the conclusions of the propositions. Our arguments imply that this number limits to 1.

2.2. The Hausdorff topology. Let Z denote any separable and locally compact metric space (e.g. any second countable topological manifold). Recall that $\mathcal{K}(Z)$ denotes the space of all closed subsets of Z , equipped with the topology of Hausdorff convergence. We collect some basic facts about the Hausdorff topology here.

2.2.1. Hausdorff convergence. Fix any sequence $\{\Lambda_k\}$ in $\mathcal{K}(Z)$. We define $\liminf \Lambda_k \in \mathcal{K}(Z)$ to be the set of $z \in Z$ such that each neighborhood of z intersects all but finitely many of the Λ_k . We define $\limsup \Lambda_k \in \mathcal{K}(Z)$ to be the set of $z \in Z$ such that each neighborhood of z intersects infinitely many of the Λ_k . It is clear that $\liminf \Lambda_k \subseteq \limsup \Lambda_k$. Then $\Lambda_k \rightarrow \Lambda$ in the Hausdorff topology if and only if $\liminf \Lambda_k = \Lambda = \limsup \Lambda_k$. We recall [44, Corollary 2.2] that, under the imposed conditions on Z , the space $\mathcal{K}(Z)$ is compact and metrizable.

2.2.2. Continuity lemmas. The following lemmas are about continuity properties of maps between $\mathcal{K}(Z)$'s. They are elementary and we omit the proofs. The first lemma asserts that taking a union with a closed set is continuous.

Lemma 2.5. *Let Z be a separable and locally compact metric space. For any $\Lambda' \in \mathcal{K}(Z)$, the map $\Lambda \mapsto \Lambda \cup \Lambda'$ is a continuous map from $\mathcal{K}(Z)$ to itself.*

On the other hand, taking an intersection with a closed set is usually not continuous. The next lemma asserts that it is continuous in a rather specific situation.

Lemma 2.6. *Let Y be any separable and locally compact metric space and set $X := (-1, 1) \times Y$, equipped with the product metric. Then for any sequence $\{\Lambda_k\}$ in $\mathcal{K}(Y)$ such that*

$$(-1, 1) \times \Lambda_k \rightarrow (-1, 1) \times \Lambda \in \mathcal{K}(X),$$

we have $\Lambda_k \rightarrow \Lambda$.

The last lemma asserts that, when Z_1 is a compact metric space, pushing forward a closed set by a continuous map $f : Z_1 \rightarrow Z_2$ is Hausdorff continuous.

Lemma 2.7. *Let $f : Z_1 \rightarrow Z_2$ be a continuous map where Z_1 is a compact metric space and Z_2 is a separable and locally compact metric space. Then the map $\Lambda \mapsto f(\Lambda)$ defines a Hausdorff continuous map $\mathcal{K}(Z_1) \rightarrow \mathcal{K}(Z_2)$.*

2.2.3. Invariant sets of flows. Suppose that Z is a smooth and compact manifold and let R denote a vector field on Z . Let $\mathcal{K}(Z, R) \subseteq \mathcal{K}(Z)$ denote the subspace of all closed subsets invariant under the flow of R . This subspace is closed and therefore compact, but it may not be connected. The following lemma states and proves an important property of locally maximal $\Lambda \in \mathcal{K}(Z, R)$.

Lemma 2.8. *If an element $\Lambda \in \mathcal{K}(Z, R)$ is locally maximal, then it is maximal with respect to inclusion in any connected subspace $\mathcal{Z} \subseteq \mathcal{K}(Z, R)$ containing Λ .*

Proof. Conley proved [11, Theorem 3.5] that Λ is locally maximal if and only if it is maximal with respect to inclusion in some open and closed subspace $\mathcal{Z}' \subseteq \mathcal{K}(Z, R)$. Any connected subspace \mathcal{Z} such that $\Lambda \in \mathcal{Z}$ must be contained in \mathcal{Z}' and the lemma follows. \square

2.3. Proofs of Theorem 4 and Theorem 5. In this section, we prove Theorems 4 and Theorem 5. The proofs are virtually identical, so for brevity we will only give the proof of Theorem 4.

The main ingredients in the proof of Theorem 4 are Theorem 6, Proposition 2.2, and Proposition 1.9. Choose any monotone area-preserving diffeomorphism ϕ of a closed symplectic surface (Σ, ω) . Let Y_ϕ denote its mapping torus and $\eta = (dt, \omega_\phi)$ denote its associated framed Hamiltonian structure. We observe that Theorem 4 follows from proving its analogue in the mapping torus:

Proposition 2.9. *For any closed R_η -invariant set $\Lambda \subseteq Y_\phi$ containing all periodic orbits of R_η , either $\Lambda = Y_\phi$ or Λ is not locally maximal.*

Proof of Proposition 2.9. Let $\Lambda \subseteq Y_\phi$ be any compact invariant set. We assume without loss of generality that it is proper, since if $\Lambda = Y_\phi$ we are already done. Our argument will go via approximation to the nondegenerate case, so choose a sequence $\{H_k\}$ of smooth functions $H_k : \mathbb{R}/\mathbb{Z} \times \Sigma \rightarrow \mathbb{R}$ converging in C^∞ to 0 such that for each k , the map $\phi_k := \phi \circ \psi_{H_k}^1$ is nondegenerate. For each k , the pair $\eta_k := (dt, \omega_\phi + dH_k \wedge dt)$ is a framed Hamiltonian structure, and the map

$$(t, p) \mapsto (t, (\psi_{H_k}^t)^{-1}(p))$$

descends to an isomorphism

$$(Y_\phi, \eta_k) \rightarrow (Y_{\phi_k}, (dt, \omega_{\phi_k}))$$

of framed Hamiltonian manifolds. Since H_k converges to 0 in the C^∞ topology as $k \rightarrow \infty$, we have that η_k converges in C^∞ to η as $k \rightarrow \infty$. For each k we choose an η_k -adapted J_k such that the sequence $\{(\eta_k, J_k)\}$ converges to a pair (η, J) in $\mathcal{D}(Y)$.

Now fix any point $z \in Y_\phi$ not in Λ . Then there is a sequence of points $\{z_k\}$ converging to z , such that z_k is not on any closed Reeb orbit for ϕ_k ; this follows from the fact that the ϕ_k are nondegenerate, hence the union of their closed Reeb orbits has measure zero. By Proposition 2.2, after possibly making an arbitrarily small perturbation to each J_k , for each sufficiently large d and each sufficiently large k there exists a J_k -holomorphic curve

$$u_{d,k} : C_{d,k} \rightarrow \mathbb{R} \times Y_\phi$$

such that:

- (i) $(0, z_k) \in u_{d,k}(C_{d,k})$;
- (ii) $\mathcal{E}(u_{d,k}) \leq d$;
- (iii) $\mathcal{A}(u_{d,k}) \leq d^{-1/2}$;
- (iv) $\chi(C_{d,k}) \geq -2$.

Write $\mathcal{P} \subseteq Y_\phi$ for the union of closed orbits of R_η and write $\overline{\mathcal{P}}$ for its closure. For each d and k , write $\mathcal{P}_d(k)$ for the union of closed orbits of R_{η_k} of period at most d . By the Hofer energy bound in (ii), we have that the level sets concentrate around closed orbits of period $\leq d$ as $s \rightarrow \infty$. We also note that since $R_{\eta_k} \rightarrow R_\eta$, we have for each fixed d that

$$\limsup_{k \rightarrow \infty} \mathcal{P}_d(k) \subseteq \mathcal{P},$$

that is periodic orbits of R_{η_k} with bounded period converge to periodic orbits of R_η . It follows that for each d , we can choose some large $k_d \gg 1$ and some $s_d \in \mathbb{R}$ such that

$$(2) \quad \limsup_{d \rightarrow \infty} \tau_{s_d} \cdot \left(u_{d,k_d}(C_{d,k_d}) \cap (s_d - 1, s_d + 1) \times Y_\phi \right) \subseteq (-1, 1) \times \overline{\mathcal{P}}.$$

Write $X := (-1, 1) \times Y_\phi$ and write $\mathcal{X} \subseteq \mathcal{K}(X)$ for the limit set of the sequence $\{u_{d,k_d}\}_{d \geq 1}$; by passing to a subsequence, we can assume by Proposition 1.9 that this is connected. Let $f : \mathcal{K}(X) \rightarrow \mathcal{K}(Y_\phi)$ denote the map $\overline{\Lambda} \mapsto \overline{\Lambda} \cap \{0\} \times Y_\phi$ and write $\mathcal{Y} := f(\mathcal{X})$. By Lemma 2.6, the map f is continuous at each point of \mathcal{X} , so it follows that \mathcal{Y} is a connected subset of $\mathcal{K}(Y)$.

By (iii), (iv), and Theorem 6, we have that $\mathcal{Y} \subseteq \mathcal{K}(Y_\phi, R_\eta)$. By (2), some $\Lambda' \in \mathcal{Y}$ is contained entirely in $\overline{\mathcal{P}}$, and is therefore contained in Λ . By (i), some $\Lambda'' \in \mathcal{Y}$ contains the point z . Now let $\mathcal{Z} \subset \mathcal{K}(Y_\phi, R_\eta)$ denote the collection of closed invariant sets equal to a union $K \cup \Lambda$ for some $K \in \mathcal{X}$. By Lemma 2.5, \mathcal{Z} is the image of a connected set by a continuous map, so it is connected. Moreover, $\Lambda \in \mathcal{Z}$, since $\Lambda = \Lambda' \cup \Lambda$, and $\Lambda'' \cup \Lambda \in \mathcal{Z}$ by definition. However, $\Lambda'' \cup \Lambda$ is not a subset of Λ , since it contains z . Hence, by Lemma 2.8, Λ is not locally maximal. □

2.4. Proofs of other dynamical results.

2.4.1. *Proofs of Theorems 1 and 2.* We prove Theorem 1 using Theorem 4. Theorem 2, the version for three-dimensional Reeb flows, follows from the same formal argument using Theorem 5, and so its proof will be skipped for brevity.

Proof of Theorem 1. Let $\phi : \Sigma \rightarrow \Sigma$ be a monotone area-preserving diffeomorphism of a closed surface and let $\Lambda \subset \Sigma$ be a proper closed invariant set. Set $U := \Sigma \setminus \Lambda$. Let $\mathcal{P} \subseteq \Sigma$ denote the union of all periodic orbits. If $\mathcal{P} \not\subseteq \Lambda$, then there exists a periodic orbit contained in U . Therefore, U is not minimal. If $\mathcal{P} \subseteq \Lambda$, then by Theorem 4 there exists some proper closed invariant subset Λ' arbitrarily C^0 -close to Λ but not equal to Λ . By taking Λ' sufficiently close to Λ , we can ensure that $\Lambda' \cap U$ is not equal to U . Then any point $z \in \Lambda' \cap U$ will not have dense orbit in U , so U is not minimal. \square

2.4.2. *Proofs of Corollaries 1.1 and 1.2.* We prove Corollary 1.1; the same formal argument proves Corollary 1.2.

Proof of Corollary 1.1. Let $\phi : \Sigma \rightarrow \Sigma$ be a diffeomorphism satisfying the assumptions of Theorem 1. The corollary is equivalent to the statement that there exists infinitely many proper compact R_ϕ -invariant subsets whose union is dense in the mapping torus Y_ϕ . Let $\mathcal{K}' := \mathcal{K}(Y_\phi, R_\phi) \setminus \{Y_\phi\}$ denote the set of all proper compact R_ϕ -invariant sets. The set \mathcal{K}' admits a natural partial order defined by inclusion. By Theorem 1, \mathcal{K}' has no maximal element. The set \mathcal{K}' is non-empty since ϕ has at least one periodic orbit (see Proposition 4.1). The set \mathcal{K}' is infinite, since otherwise it would have a maximal element. Now consider the set

$$Z := \overline{\bigcup_{\Lambda \in \mathcal{K}'} \Lambda} \subseteq Y_\phi,$$

the closure of the union of all proper compact invariant sets. The set Z is compact, invariant, and contains all proper compact invariant sets. Since \mathcal{K}' has no maximal element, Z must be equal to Y_ϕ . \square

2.4.3. *Proof of Theorem 3.*

Proof. Fix a closed Finsler surface F as in the statement of the theorem. Now we introduce the notion of a “projected invariant set” of the geodesic flow, which will be useful for our proof. Recall that there is a Reeb vector field whose orbits project to geodesics. We write SF for the unit tangent bundle and R for this vector field. Let $\pi : SF \rightarrow F$ denote the bundle projection. We call a compact subset $\Xi \subseteq F$ a *projected invariant set* if there exists some compact R -invariant subset $\Lambda \subseteq SF$ such that $\pi(\Lambda) = \Xi$. A projected invariant set Ξ is called *proper* if $\Xi \neq F$. The standard example of a projected invariant set is the closure $\overline{\gamma(\mathbb{R})}$ of some geodesic γ . We use Theorem 5 to prove the following analogue of Corollaries 1.1 and 1.2.

Lemma 2.10. *There exists infinitely many distinct proper projected invariant sets whose union is dense in F .*

Proof. Let $\Xi \subseteq F$ be any proper projected invariant set. We claim that there exists a proper projected invariant set Ξ' such that Ξ is a strict subset of Ξ' . Once this claim is established, an analogous argument to the proof of Corollary 1.1 proves the lemma. We omit this part and focus on proving the claim.

Fix a proper projected invariant set Ξ as above. We construct the set Ξ' in two cases, depending on whether Ξ contains all of the closed geodesics or not. First, suppose that there exists a closed geodesic $\gamma : \mathbb{R} \rightarrow F$ such that $\gamma(\mathbb{R})$ is not a subset of Ξ . Then, we set $\Xi' = \Xi \cup \gamma(\mathbb{R})$.

Next, suppose that Ξ contains the images of all closed geodesics. Let $\Lambda \subset SF$ denote the closure of the union of all compact invariant sets projecting into Ξ ; note that $\pi(\Lambda) = \Xi$. Since Ξ contains the images of all closed geodesics, it follows from the construction that Λ must contain all closed orbits of the geodesic flow. By Theorem 5 and Lemma A.4, Λ is not locally maximal. Therefore, there exists a Hausdorff convergent sequence $\Lambda_k \rightarrow \Lambda$ such that Λ_k is not a subset of Λ for each k . We can assume without loss of generality that Λ is a strict subset of Λ_k for each k . This is done by replacing Λ_k with $\Lambda_k \cup \Lambda$ for each k ; these sets still Hausdorff converge to Λ by Lemma 2.5. By the definition of Λ , it follows that Ξ is a strict subset of $\pi(\Lambda_k)$ for each k . Moreover, it follows from Lemma 2.7 that $\pi(\Lambda_k) \rightarrow \Xi$, so $\pi(\Lambda_k)$ is a proper projected invariant set for sufficiently large k . We set $\Xi' = \pi(\Lambda_k)$ for any choice of sufficiently large k . \square

We now discuss how to prove Theorem 3 using the lemma. By Lemma 2.10, there exists an infinite collection $\mathcal{F} \subseteq \mathcal{K}(F)$ of projected proper invariant sets whose union is dense in F . For any proper projected invariant set $\Xi \subset F$ and any point $z \in F$, there exists a geodesic $\gamma : \mathbb{R} \rightarrow F$ such that $\gamma(0) = z$ and $\overline{\gamma(\mathbb{R})} \subseteq \Xi$. Hence, for each $\Xi \in \mathcal{F}$, we obtain a countable sequence of geodesics γ_k^Ξ , such that i) $\gamma_k^\Xi(\mathbb{R}) \subset \Xi$ for each k and ii) $\cup_k \gamma_k^\Xi(\mathbb{R})$ is dense in Ξ . After passing to a subset, we may assume that these geodesics have distinct closures. Let G be the collection of all geodesics γ_k^Ξ over all $\Xi \in \mathcal{F}$ and all k . Then G is infinite, because \mathcal{F} is, and satisfies Theorem 3(a–c). \square

2.4.4. Proof of Corollary 1.13.

Proof. Contreras–Mazzucchelli proved [12, Section 3.4, p. 18] that, given conditions i) and ii), the invariant set \mathcal{P} is locally maximal. Their argument is short and we paraphrase it here for the convenience of the reader. A uniformly hyperbolic invariant is locally maximal if and only if it has “local product structure” [25, Theorem 6.2.7]. By the definition of local product structure [12, Section 3.2], the set \mathcal{P} has local product structure if and only if for any pair of sufficiently close periodic orbits γ_1 and γ_2 , the local stable/unstable manifolds intersect in \mathcal{P} . The existence of a non-empty transverse intersection follows from Condition ii) and the fact that γ_1 and γ_2 are close. To prove that the intersection is contained in \mathcal{P} , Contreras–Mazzucchelli observe that if γ_1 and γ_2 are close then they are in the same homoclinic class. By the Birkhoff–Smale horseshoe theorem [25, Theorem 6.5.2], it follows that any intersection of stable/unstable manifolds is approximated by periodic orbits.

Now, the corollary follows immediately from our main Theorem 5. Since \mathcal{P} is locally maximal, Theorem 5 and Lemma A.4 imply that $\mathcal{P} = SF$. By i), we have that SF is uniformly hyperbolic, and therefore that the geodesic flow is Anosov. \square

2.5. Proof of crossing energy theorem. We prove Theorem 7 using Theorem 6, Proposition 1.9, and the lemmas from this section.

Proof of Theorem 7. Assume for the sake of contradiction that the corollary is false. Then there exists a sequence of J -holomorphic curves $\{u_k : C_k \rightarrow \mathbb{R} \times Y\}$ satisfying (1), $\chi(C_k) \geq -T$ and $\mathcal{A}(u_k) \leq 1/k$. By Theorem 6, every element of the limit set \mathcal{X} is a cylinder over a closed invariant set. By Proposition 1.9, after replacing $\{u_k\}$ with a subsequence if necessary, the limit set is connected. Let $f : \mathcal{K}(X) \rightarrow \mathcal{K}(Y)$ denote the map $\bar{\Lambda} \mapsto \bar{\Lambda} \cap \{0\} \times Y$ and set $\mathcal{Y} := f(\mathcal{X})$. By Lemma 2.6 and Theorem 6, \mathcal{Y} is a connected subspace of $\mathcal{K}(Y, R_\eta)$.

Let $\mathcal{Z} \subseteq \mathcal{K}(Y)$ denote the collection of closed invariant sets which are a union $K \cup \Lambda$ for some $K \in \mathcal{Y}$; this is connected by Lemma 2.5. The first item in (1) implies \mathcal{Y} contains an invariant set Λ' contained in U , and since U is an isolating neighborhood of Λ , this implies $\Lambda' \subseteq \Lambda$. It follows that $\Lambda \in \mathcal{Z}$. Since Λ is locally maximal, every element of \mathcal{Z} is a subset of Λ by Lemma 2.8. Hence, every element of \mathcal{Y} is a subset of Λ . This in turn implies that for sufficiently large k , every level set of $u_k(C_k)$ lies inside $\mathbb{R} \times \bar{V}$. We now arrive at a contradiction by appealing to the second condition of (1). \square

3. LOW-ACTION CURVES FROM EMBEDDED CONTACT HOMOLOGY

This section proves Proposition 2.3. Previously statements like this have been proved under the assumption of two periodic Reeb orbits [39], or later under the assumption of finitely many periodic Reeb orbits [16]; this section shows that this phenomenon in fact holds much more generally.

The basic strategy of proof follows the strategy developed in [39, 16], whereby one considers a tower of U -curves and compares the ECH index I to the J_0 index, which controls the topology of the curves — we will review these terms below. What is new here is the argument to bound the difference between I and J_0 : this is controlled by the Chern class and the Conley-Zehnder index, and some new ideas are needed to bound these terms without assuming finitely many orbits.

3.1. Embedded contact homology. We review the basic features of embedded contact homology [33, 35] here. Fix a closed, smooth, connected, oriented three-manifold Y and a contact structure ξ .

3.1.1. Reeb flow basics. Fix any contact form λ defining ξ , i.e. satisfying the identity $\ker(\lambda) = \xi$. Recall that the *Reeb vector field* R is the unique vector field solving the equations

$$\lambda(R) \equiv 1, \quad d\lambda(R, -) \equiv 0.$$

A *closed Reeb orbit* is a smooth map $\gamma : \mathbb{R}/T\mathbb{Z} \rightarrow Y$ for some $T > 0$ such that $\dot{\gamma}(t) = R(\gamma(t))$ for all t ; as is standard we will make no distinction between two closed Reeb orbits that agree up to a reparameterization of the domain. A closed Reeb orbit is *simple* if γ is injective. For any closed Reeb orbit γ , we write $\gamma^k : \mathbb{R}/kT\mathbb{Z} \rightarrow Y$ for its k -th iteration. The number T is the *action* of γ , denoted by

$$\mathcal{A}(\gamma) := \int_\gamma \lambda = T.$$

The time T linearized flow of R determines a symplectic isomorphism $\xi_{\gamma(0)} \rightarrow \xi_{\gamma(0)}$ called the *Poincaré return map*. The orbit γ is *nondegenerate* if the return map does

not have an eigenvalue equal to 1. A nondegenerate orbit γ is *hyperbolic* if P_γ has real eigenvalues and *elliptic* if P_γ has complex eigenvalues of unit length. We say λ is *nondegenerate* if all closed Reeb orbits are nondegenerate. For any fixed contact structure ξ , a generic defining contact form λ is nondegenerate.

3.1.2. *ECH generators.* A *Reeb orbit set* is a (possibly empty) finite set $\alpha = \{(\alpha_i, m_i)\}$ of pairs (α_i, m_i) , where α_i is a simple closed Reeb orbit and $m_i \in \mathbb{N}$ is a positive integer multiplicity. An *ECH generator* is a Reeb orbit set $\alpha = \{(\alpha_i, m_i)\}$ such that i) each of the α_i are pairwise distinct and ii) $m_i = 1$ if α_i is hyperbolic. Denote by $\text{ECC}(Y, \lambda)$ the $\mathbb{Z}/2$ -vector space generated by the set of ECH generators. Any Reeb orbit set α has a homology class $[\alpha] := \sum_i m_i [\alpha_i] \in H_1(Y; \mathbb{Z})$. For each $\Gamma \in H_1(Y; \mathbb{Z})$, let $\text{ECC}(Y, \lambda, \Gamma)$ denote the sub-module generated by ECH generators homologous to Γ .

3.1.3. *ECH differential.* Assume that λ is nondegenerate. Choose a generic λ -adapted almost-complex structure J on $\mathbb{R} \times Y$. The ECH differential

$$\partial_J : \text{ECC}(Y, \lambda) \rightarrow \text{ECC}(Y, \lambda)$$

is defined by counting certain “ J -holomorphic currents” which we now define. We say that a J -holomorphic curve $u : C \rightarrow \mathbb{R} \times Y$ is *somewhere injective* if there exists $\zeta \in C$ such that $u^{-1}(u(\zeta)) = \{\zeta\}$ and Du is injective at ζ . A *J -holomorphic current* is a finite set $\mathcal{C} = \{(C_k, d_k)\}$ of pairs where the C_k denote distinct standard, somewhere injective J -holomorphic curves with finite Hofer energy, and the d_k are positive integer multiplicities. We say that \mathcal{C} is *somewhere injective* if $d_k = 1$ for each k and *embedded* if the C_k are pairwise disjoint and embedded. For any J -holomorphic current $\mathcal{C} = \{(C_k, d_k)\}$, the slices

$$\mathcal{C} \cap \{s\} \times Y = \{(C_k \cap \{s\} \times Y, d_k)\}$$

form for $|s| \gg 1$ a weighted collection of embedded loops in Y . The slices converge as 1-dimensional currents to Reeb orbit sets α and β as $s \rightarrow \infty$ and $s \rightarrow -\infty$, respectively. For any pair of Reeb orbit sets α and β with $[\alpha] = [\beta]$, we let $\mathcal{M}(\alpha, \beta)$ denote the moduli space of J -holomorphic curves with positive asymptotic limit at α and negative asymptotic limit at β .

Any $\mathcal{C} \in \mathcal{M}(\alpha, \beta)$ has an associated *ECH index* $I(\mathcal{C}) \in \mathbb{Z}$, defined below, and for each $k \in \mathbb{Z}$ we let $\mathcal{M}_k(\alpha, \beta)$ denote the subspace of curves of ECH index k . When J is sufficiently generic, the space $\mathcal{M}_1(\alpha, \beta)$ is a smooth 1-dimensional manifold. Moreover, it has a free \mathbb{R} -action given by the translation action on $\mathbb{R} \times Y$, and the quotient $\mathcal{M}_1(\alpha, \beta)/\mathbb{R}$ is a finite set of points. The matrix coefficient of the ECH differential with respect to a pair of ECH generators α and β is defined by the identity

$$\langle \partial_J \alpha, \beta \rangle := \#_2 \mathcal{M}(\alpha, \beta)/\mathbb{R}$$

where $\#_2$ denotes the modulo 2 count of points. By [37, 38], $\partial_J^2 = 0$, and therefore $(\text{ECC}(Y, \lambda), \partial_J)$ is a chain complex. The *embedded contact homology* $\text{ECH}(Y, \xi)$ is its homology group. A consequence of Taubes’ isomorphism of ECH with monopole Floer homology [51] is that $\text{ECH}(Y, \xi)$ does not depend on the choice of contact form λ defining ξ or the choice of J used to define the ECH differential. The ECH differential preserves the homology class $[\alpha] \in H_1(Y; \mathbb{Z})$ of an orbit set α ; for any $\Gamma \in H_1(Y; \mathbb{Z})$ we write $\text{ECH}(Y, \xi, \Gamma)$ for the homology of $(\text{ECC}(Y, \lambda, \Gamma), \partial_J)$.

3.1.4. *The U -map.* For a generic choice of λ -adapted almost-complex structure J and any point $z \in Y$ not on any closed Reeb orbit, there exists a chain map

$$U_{J,z} : \text{ECC}(Y, \lambda, \Gamma) \rightarrow \text{ECC}(Y, \lambda, \Gamma)$$

defined as follows. Write $\mathcal{M}_2(\alpha, \beta; z)$ for the space of all J -holomorphic currents with ECH index 2 whose support contains $(0, z) \in \mathbb{R} \times Y$. For a generic choice of J and any z not on any closed Reeb orbit, the space $\mathcal{M}_2(\alpha, \beta; z)$ is a finite set of points. The matrix coefficient of $U_{J,z}$ with respect to α and β is defined by the identity

$$\langle U_{J,z}\alpha, \beta \rangle = \#_2 \mathcal{M}(\alpha, \beta; z).$$

The map $U_{J,z}$ descends to a map on homology that we call the U -map:

$$U : \text{ECH}(Y, \xi, \Gamma) \rightarrow \text{ECH}(Y, \xi, \Gamma).$$

The map $U_{J,z}$ may vary with different choices of J and z . However, the chain homotopy class of $U_{J,z}$ does not depend on the choice of J and z , so the induced map on homology does not depend on the choice of J and z .

3.1.5. *U -towers and the volume property.* Fix any $\Gamma \in H_1(Y; \mathbb{Z})$. A U -tower is a sequence of nonzero classes

$$\{\sigma_k\}_{k \geq 0} \subset \text{ECH}(Y, \xi, \Gamma)$$

such that i) $U\sigma_k = \sigma_{k-1}$ for each $k > 0$ and ii) $U\sigma_0 = 0$. Taubes' isomorphism [51] and a computation by Kronheimer–Mrowka [41, Chapter 35] prove that $\text{ECH}(Y, \xi, \Gamma)$ contains a U -tower whenever the class $c_1(\xi) + 2\text{PD}(\Gamma) \in H^2(Y; \mathbb{Z})$ is torsion. Here $c_1(\xi)$ denotes the first Chern class of ξ with respect to any complex structure which rotates positively with respect to $d\lambda$.

For any nonzero class $\sigma \in \text{ECH}(Y, \xi)$, define its *spectral invariant* $c_\sigma(\lambda) \in \mathbb{R}$ to be the infimum of all L such that σ is represented by a cycle in $\text{ECC}(Y, \lambda)$ with all constituent generators having action $\leq L$. Here the action of an ECH generator $\alpha = \{(\alpha_i, m_i)\}$ is defined to be

$$\mathcal{A}(\alpha) := \sum_i m_i \mathcal{A}(\alpha_i).$$

A quantitative version of the proof that ECH is independent of the choice of J used to define the ECH differential shows that the spectral invariants do not depend on J either. However, the spectral invariants can and usually do vary with λ . Each spectral invariant $c_\sigma(\lambda)$ is C^0 -continuous with respect to λ ; this allows us to extend the definition of $c_\sigma(\lambda)$ to degenerate λ . The following lemma records some relevant chain-level information that we can extract from a U -tower.

Lemma 3.1. *Assume that there exists a U -tower $\{\sigma_k\}_{k \geq 0} \subset \text{ECH}(Y, \xi, \Gamma)$ for some $\Gamma \in H_1(Y; \mathbb{Z})$. Assume that λ is nondegenerate and choose generic J , and $(0, z)$ not on any closed Reeb orbit, so that the ECH differential ∂_J and the chain map $U_{J,z}$ are well-defined. Then for each $\epsilon > 0$ and each $k \geq 1$, there exists an ECH generator α_k such that*

- (a) $\mathcal{A}(\alpha_k) \leq c_{\sigma_k}(\lambda)$;
- (b) $U_{J,z}^k(\alpha_k) \neq 0$.

Proof. For any fixed $\epsilon > 0$ and k , there exists a cycle $x \in \text{ECC}(Y, \lambda, \Gamma)$ representing σ_k which splits as a sum of ECH generators $x = \sum_{i=1}^N x_i$ with action less than $c_{\sigma_k}(\lambda) + \epsilon$. Since there are only finitely many Reeb orbit sets of action $\leq c_{\sigma_k}(\lambda) + 1$, as λ is nondegenerate, we can therefore find a cycle x as above with action exactly $c_{\sigma_k}(\lambda)$. Since $U_{j,z}^k(x) \neq 0$, it follows that $U_{j,z}^k(x_i) \neq 0$ for some i . Set $\alpha_k := x_i$. \square

One of the most powerful properties of the ECH spectral invariants is the *volume property* proved by Cristofaro-Gardiner–Hutchings–Ramos [17]. Their result, stated in the following proposition, shows that the spectral invariants of a U -tower asymptotically recover the contact volume.

Theorem 8 (ECH volume property, [17]). *Assume that there exists a U -tower $\{\sigma_k\}_{k \geq 0} \subset \text{ECH}(Y, \xi, \Gamma)$ for some $\Gamma \in H_1(Y; \mathbb{Z})$. Then for any contact form λ we have*

$$(3) \quad \lim_{k \rightarrow \infty} c_{\sigma_k}(\lambda)^2 / 2k = \int_Y \lambda \wedge d\lambda.$$

3.1.6. *The ECH index.* We now give the previously deferred definition of the ECH index. Fix a nondegenerate contact form λ and a pair of ECH generators $\alpha = \{(\alpha_i, m_i)\}$ and $\beta = \{(\beta_j, n_j)\}$. Let $H_2(Y, \alpha, \beta)$ denote the space of equivalence classes of integral 2-chains with boundary $\alpha - \beta$, where two such chains are equivalent if and only if they differ by a 2-boundary. The *ECH index* of a class $Z \in H_2(Y, \alpha, \beta)$ is an integer defined by the formula

$$(4) \quad I(Z) := c_\tau(Z) + Q_\tau(Z) + \sum_i \sum_{k=1}^{m_i} \text{CZ}_\tau(\alpha_i^k) - \sum_j \sum_{l=1}^{n_j} \text{CZ}_\tau(\beta_j^l).$$

Definitions of these terms can be found in [35]. We will narrow our discussion of the ECH index to exactly those terms which are useful for the proof of Proposition 2.3, namely the relative Chern class and the Conley-Zehnder index. The *relative Chern class* $c_\tau(Z)$ is defined as follows. We choose an oriented smooth surface $S \subset Y$ representing Z and choose a section $\psi : S \rightarrow \xi$, transverse to the zero section, such that ψ is a nonzero constant on each component of ∂S and we set

$$c_\tau(Z) := \#\psi^{-1}(0)$$

where $\#$ denotes the oriented count of points. As for the Conley-Zehnder index, we define it by

$$\text{CZ}_\tau(\gamma) = [\theta_\tau(\gamma)] + [\theta_\tau(\gamma)]$$

where $\theta_\tau(\gamma)$ denotes the “monodromy number” of the linearized flow along γ in the trivialization τ . We refer the reader to [35] for a definition of the monodromy number. We define the ECH index of a curve to be the ECH index of its homology class, emphasizing that this is independent of the choice of τ .

3.1.7. *Topological complexity of U -map curves.* A variant of the ECH index called the J_0 index plays a key role in the proof of Propositions 2.2 and 2.3. In the notation of

(4) we write

$$\begin{aligned}
(5) \quad J_0(Z) &:= -c_\tau(Z) + Q_\tau(Z) + \sum_i \sum_{k=1}^{m_i-1} \text{CZ}_\tau(\alpha_i^k) - \sum_j \sum_{l=1}^{n_j-1} \text{CZ}_\tau(\beta_j^l) \\
&= I(Z) - 2c_\tau(Z) - \left(\sum_i \text{CZ}_\tau(\alpha_i^{m_i}) - \sum_j \text{CZ}_\tau(\beta_j^{n_j}) \right).
\end{aligned}$$

The J_0 index controls the topological complexity of a holomorphic current counted by the U -map. To state the bound precisely, we need to recall the following structural property of U -map currents, a proof of which is found in [35, Proposition 3.7]. Any current $\mathcal{C} \in \mathcal{M}_2(\alpha, \beta; z)$ counted by the U -map splits as a disjoint union $\mathcal{C}_0 \sqcup \mathcal{C}_2$ where \mathcal{C}_0 is a union of trivial cylinders with multiplicities and \mathcal{C}_2 is an embedded J -holomorphic curve with $I(\mathcal{C}_2) = 2$.

Proposition 3.2 ([34, Section 6]). *Fix a generic J and point z so that the chain map $U_{J,z}$ is defined, and let $\mathcal{C} = \mathcal{C}_0 \sqcup \mathcal{C}_2 \in \mathcal{M}_2(\alpha, \beta; z)$ be a J -holomorphic current counted by the U -map. Then*

$$(6) \quad J_0(\mathcal{C}) \geq -\chi(\mathcal{C}_2).$$

3.2. The based rotation number. To prove what we need to know about the ECH curves, we will also need to recall some information about the “rotation number” of flows. Our treatment here is inspired by, and closely follows [7], though we handle a few points in a different way that is better suited for our purposes. For any closed Reeb orbit $\gamma : \mathbb{R}/T\mathbb{Z} \rightarrow Y$ and any choice of (positive) symplectic trivialization τ there is a number $\rho_\tau(\gamma, \xi) \in \mathbb{R}$ called the *based rotation number* of γ , which measures how ξ rotates under the linearization of the Reeb flow. It is defined as follows. The linearized flow is symplectic; apply polar de-composition and take the unitary part. The unitary part descends to a flow on the oriented real projectivization $P(\xi)$. Pull back by γ and conjugate with the trivialization τ to define a flow

$$\bar{\Phi} : \mathbb{R} \times (\mathbb{R}/T\mathbb{Z} \times \mathbb{R}/\mathbb{Z}) \rightarrow \mathbb{R}/T\mathbb{Z} \times \mathbb{R}/\mathbb{Z}$$

generated by a vector field \bar{R} . We write $\theta : \mathbb{R}/T\mathbb{Z} \times \mathbb{R}/\mathbb{Z} \rightarrow \mathbb{R}/\mathbb{Z}$ for the angular coordinate on the target. Let $\bar{\theta} : [0, T] \rightarrow \mathbb{R}$ be the unique real-valued lift of the circle-valued map $t \mapsto \theta(\bar{\Phi}_t(0, 0))$ satisfying the initial condition $\bar{\theta}(0) = 0$. Then the based rotation number is

$$\rho_\tau(\gamma, \xi) := \bar{\theta}(T).$$

This depends only on the homotopy class, rel endpoints, of the path of symplectic matrices arising from the linearized flow.

The next lemma relates the based rotation number for the contact structure to the Conley–Zehnder index.

Lemma 3.3. *For any closed Reeb orbit γ and any trivialization τ , we have the bound*

$$(7) \quad |\text{CZ}_\tau(\gamma) - 2\rho_\tau(\gamma, \xi)| \leq 6.$$

Proof. The Conley–Zehnder index is defined by

$$\text{CZ}_\tau(\gamma) = [\theta_\tau(\gamma)] + [\theta_\tau(\gamma)]$$

where $\theta_\tau(\gamma)$ denotes the monodromy number of γ in the trivialization τ . Define $\rho'_\tau(\gamma, \xi)$ analogously to $\rho_\tau(\gamma, \xi)$, but using the full linearized flow rather than just the unitary part. It is proved in [7, Lemma 2.6] that

$$|\theta_\tau(\gamma) - \rho'_\tau(\gamma, \xi)| \leq 1.$$

It remains to understand the relationship between ρ and ρ' . We claim that

$$(8) \quad |\rho_\tau(\gamma, \xi) - \rho'_\tau(\gamma, \xi)| \leq 1,$$

which implies the lemma in view of the above. To see why (8) holds, we first note that if we choose as our basepoint (i.e. our trivialization τ) an eigenvector of the positive-definite symmetric part of the polar decomposition of the time T linearized flow, then in fact $\rho'_\tau = \rho_\tau$. Indeed, the space of positive-definite symmetric and symplectic matrices is contractible, and the rotation number only depends on the homotopy class, rel endpoints, so we can replace the positive-definite part of the path of matrices arising from the linearized flow by symmetric and symplectic positive-definite matrices which all have τ as an eigenvector with positive eigenvalue. The claimed inequality (8) now follows from [7, Lemma 2.6], which bounds the difference between the based rotation number measured with respect to two different basepoints. \square

3.3. Proof of Proposition 2.3. We can now give the proof of Proposition 2.3.

Proof. Step 1: To deal with the fact that we are considering contact structures that are torsion, but possibly non-trivial, we will need to work with an “ n^{th} -power” construction. This step collects the results we will need about this.

Fix an integer $n \geq 1$ such that $n \cdot c_1(\xi) = 0$. Write $\xi_n = \xi \otimes \dots \otimes \xi$ for the n -fold (complex) tensor product of ξ . This is a (trivial) complex line bundle. Choose a unitary trivialization τ of ξ over the simple closed Reeb orbits. This induces a trivialization τ_n of ξ_n . The line bundle ξ_n has a relative Chern class, defined analogously to the contact case. We first note that the relative Chern class of ξ_n with respect to τ_n is computed as follows:

$$(9) \quad c_{\tau_n}(Z, \xi_n) = n \cdot c_\tau(Z, \xi).$$

Next, it is useful to understand how the Chern class depends on the choice of trivialization. Fix any pair of unitary trivializations τ, τ' and any simple closed Reeb orbit $\gamma : \mathbb{R}/T\mathbb{Z} \rightarrow Y$. The trivializations define unitary bundle isomorphisms

$$\tau, \tau' : \gamma^* \xi_n \rightarrow \mathbb{R}/T\mathbb{Z} \times \mathbb{C};$$

the composition $\tau \circ (\tau')^{-1}$ defines a smooth map $\mathbb{R}/T\mathbb{Z} \rightarrow U(1)$. Denote the degree of this map by $\text{wind}_\gamma(\tau, \tau'; \xi_n)$. Then we have the identity

$$(10) \quad c_\tau(Z, \xi_n) - c_{\tau'}(Z, \xi_n) = - \sum_i m_i \text{wind}_{\alpha_i}(\tau, \tau'; \xi) + \sum_j n_j \text{wind}_{\beta_j}(\tau, \tau'; \xi_n).$$

This is proved by the same argument as in the case of contact structures [33].

There is an analogous story for the based rotation number. The unitary part of the linearized flow, being complex linear, defines a map on the complex tensor product ξ_n and we can define the based rotation number analogously, which we call the *induced based rotation number* (in the trivialization τ) on ξ_n , denoted $\rho_\tau(\gamma, \xi_n)$. We now prove

some basic properties of the induced based rotation number analogous to the observed properties of the relative Chern class.

Lemma 3.4. *The unitary component of the based rotation number and the induced based rotation number satisfy the following basic properties:*

- (Change of trivialization) For any pair τ, τ' of unitary trivializations we have

$$(11) \quad \rho_\tau(\gamma, \xi_n) - \rho_{\tau'}(\gamma, \xi_n) = \text{wind}_\gamma(\tau, \tau'; \xi_n).$$

- (Additive under tensor product) For any $n \geq 1$ we have

$$(12) \quad \rho_{\tau_n}(\gamma, \xi_n) = n \cdot \rho_\tau(\gamma, \xi).$$

Proof. Let $\bar{\Phi}$ and $\bar{\Phi}'$ be the respective flows on $\mathbb{R}/T\mathbb{Z} \times \mathbb{R}/\mathbb{Z}$ defined by τ and τ' . After applying a constant rotation to one of the trivializations, we may assume without loss of generality that $\tau^{-1}(0, 0) = (\tau')^{-1}(0, 0)$. It follows that

$$\bar{\Phi}_t(0, 0) = (\tau \circ (\tau')^{-1}) \circ \bar{\Phi}'_t(0, 0)$$

for each $t \in \mathbb{R}$. The lifts $\bar{\theta}$ and $\bar{\theta}'$ corresponding to τ and τ' differ by the lift of the map $\mathbb{R}/T\mathbb{Z} \rightarrow U(1)$ defined by $\tau \circ (\tau')^{-1}$. The map $\mathbb{R}/T\mathbb{Z} \rightarrow U(1)$ has degree $\text{wind}_\gamma(\tau, \tau'; \Xi)$, so it follows that

$$\theta(T) - \theta'(T) = \text{wind}_\gamma(\tau, \tau'; \Xi)$$

which proves (11).

Now fix any $n \geq 1$ and let $\bar{\Phi}$ and $\bar{\Phi}_n$ be the respective flows on $\mathbb{R}/T\mathbb{Z} \times \mathbb{R}/\mathbb{Z}$ defined by τ and τ_n . We observe that

$$(\bar{\Phi}_n)_t = n \cdot \bar{\Phi}_t$$

for each $t \in \mathbb{R}$. This implies that $\bar{\theta}_n = n \cdot \bar{\theta}$ where $\bar{\theta}$ and $\bar{\theta}_n$ are the lifts corresponding to τ and τ_n . Evaluating both sides at T yields (12). \square

Step 2: Recall that the tensor power ξ_n is a trivial complex line bundle. We fix a *global* unitary trivialization \mathfrak{t} of ξ_n over the entire manifold Y . Let λ' be any nondegenerate contact form. Let $\alpha = \{(\alpha_i, m_i)\}$ and $\beta = \{(\beta_j, n_j)\}$ be any pair of homologous ECH generators such that $\mathcal{A}(\beta) \leq \mathcal{A}(\alpha)$. This step proves that there exists a constant $\delta > 0$ depending only on the background metric, the C^2 norm of λ' , and the trivialization \mathfrak{t} such that for any $Z \in H_2(Y, \alpha, \beta)$ we have

$$(13) \quad |I(Z) - J_0(Z)| \leq \delta \mathcal{A}(\alpha).$$

Write $K(Z) := I(Z) - J_0(Z)$. Choose any symplectic trivialization τ of ξ over the simple closed Reeb orbits. Then $K(Z)$ expands as

$$K(Z) = 2c_\tau(Z) + \sum_i CZ_\tau(\alpha_i^{m_i}) - \sum_j CZ_\tau(\beta_j^{n_j}).$$

We define

$$K_{\text{approx}}(Z, \xi) := 2c_\tau(Z, \xi) + 2 \sum_i \rho_\tau(\alpha_i^{m_i}, \xi) - 2 \sum_j \rho_\tau(\beta_j^{n_j}, \xi)$$

and a corresponding version

$$(14) \quad K_{\text{approx}}(Z, \xi_n) := 2c_{\tau'}(Z, \xi_n) + 2 \sum_i \rho_{\tau'}(\alpha_i^{m_i}, \xi_n) - 2 \sum_j \rho_{\tau'}(\beta_j^{n_j}, \xi_n)$$

for ξ_n , where τ' denotes a unitary trivialization of ξ_n over the simple closed Reeb orbits. It follows from (10) and (11) that the definition of $K_{\text{approx}}(Z, \xi_n)$ does not depend on the choice of τ' . It follows from (7) that

$$(15) \quad |K(Z) - K_{\text{approx}}(Z, \xi)| \leq 6 \sum_i m_i + 6 \sum_j m_j \leq 12T_{\min}(\lambda')^{-1} \mathcal{A}(\alpha)$$

where $T_{\min}(\lambda')$ denotes the minimal period of a closed Reeb orbit of λ' . Note that $T_{\min}(\lambda')$ admits a positive lower bound depending only on the C^2 norm of λ' .

We now bound $K_{\text{approx}}(Z, \xi)$. Since the left-hand side in (14) does not depend on the choice of τ' on the right-hand side, we set $\tau' = \tau_n$ and use (9) and (12) to show that

$$(16) \quad K_{\text{approx}}(Z, \xi) = n^{-1} \cdot K_{\text{approx}}(Z, \xi_n).$$

Now we set $\tau' = \mathfrak{t}$ and expand

$$K_{\text{approx}}(Z, \xi_n) := 2c_{\mathfrak{t}}(Z, \xi_n) + 2 \sum_i \rho_{\mathfrak{t}}(\alpha_i^{m_i}, \xi_n) - 2 \sum_j \rho_{\mathfrak{t}}(\beta_j^{m_j}, \xi_n).$$

It follows that immediately that $c_{\mathfrak{t}}(Z, \xi_n) = 0$. It remains to bound $\rho_{\mathfrak{t}}(\gamma, \xi_n)$ for any closed Reeb orbit $\gamma : \mathbb{R}/T\mathbb{Z} \rightarrow Y$. To do so, it is convenient to observe that the based rotation number along γ can be computed by integrating the “rotation density” of the flow with respect to \mathfrak{t} . To be precise, the global trivialization \mathfrak{t} and the action of the unitary part of the linearized flow on ξ_n define a flow

$$\bar{\Phi}_{\mathfrak{t}} : \mathbb{R} \times Y \times \mathbb{R}/\mathbb{Z} \rightarrow Y \times \mathbb{R}/\mathbb{Z}$$

generated by a vector field \bar{R} . The Lie derivative of the \mathbb{R}/\mathbb{Z} -coordinate on the target is a smooth function $r_{\mathfrak{t}} : Y \rightarrow \mathbb{R}$, which restricts to the rotation density on any simple closed Reeb orbit. The function $r_{\mathfrak{t}}$ depends only on \mathfrak{t} and the linearized Reeb flow, so $|r_{\mathfrak{t}}|$ admits a finite upper bound $c > 0$ depending only on \mathfrak{t} and the C^2 norm of λ' . We conclude that

$$|\rho_{\mathfrak{t}}(\gamma, \xi_n)| \leq \sup |r_{\mathfrak{t}}| \cdot \mathcal{A}(\gamma) \leq \delta_1 \cdot \mathcal{A}(\gamma).$$

It follows from the above bound that

$$(17) \quad |K_{\text{approx}}(Z, \xi_n)| \leq 2\delta_1 \left(\sum_i m_i \mathcal{A}(\alpha_i) + \sum_j n_j \mathcal{A}(\beta_j) \right) \leq 4\delta_1 \mathcal{A}(\alpha).$$

Combine (15), (16), and (17) to show

$$|K(Z)| \leq (12T_{\min}(\lambda')^{-1} + 4\delta_1) \mathcal{A}(\alpha)$$

which proves (13) with $\delta := 12T_{\min}(\lambda')^{-1} + 4\delta_1$.

Step 3: To simplify the notation, write $c'_k = c_{\sigma_k}(\lambda')$. Choose generic J and z not on a closed Reeb orbit such that the chain map $U_{J,z}$ is well-defined. By Lemma 3.1, for any $k \geq 1$, there exists an ECH generator α_k such that

- (i) $\mathcal{A}(\alpha_k) \leq c'_k$;
- (ii) $U_{J,z}^k(\alpha_k) \neq 0$.

It follows that there exists a sequence of ECH generators $\{\beta_j\}_{j=0}^k$, each with $\mathcal{A}(\beta_j) \leq c'_k$, and J -holomorphic currents $\{\mathcal{C}_j\}_{j=1}^k$ such that $\mathcal{C}_j \in \mathcal{M}(\beta_j, \beta_{j-1})$, $I(\mathcal{C}_j) = 2$, and $(0, z) \in \text{supp}(\mathcal{C}_j)$ for each j . Now set $Z := \sum_{j=1}^k [\mathcal{C}_j]$. Using the fact that $I(\mathcal{C}_j) = 2$ for each j and the bound (13) we derive the bound

$$(18) \quad \sum_{j=1}^k J_0(\mathcal{C}_j) = J_0(Z) \leq 2k + \delta \mathcal{A}(\alpha_k) \leq 2k + 2\delta c'_k.$$

It is an immediate consequence of (6) that $J_0(\mathcal{C}_j) \geq -1$ for each j . Write S_1 for the set of indices j such that

$$J_0(\mathcal{C}_j) \geq 3.$$

It follows that $3S_1 - (k - S_1) \leq 2k + \delta c'_k$, hence

$$(19) \quad \#S_1 \leq \frac{3}{4}k + \frac{\delta}{4}c'_k.$$

Write S_2 for the set of all indices j such that $\mathcal{A}(\mathcal{C}_j) \geq k^{-1/16}$. Since $\sum_{j=1}^k \mathcal{A}(\mathcal{C}_j) = \mathcal{A}(\alpha_k) - \mathcal{A}(\beta_0) \leq c'_k$, and the action is nonnegative, it follows that

$$(20) \quad \#S_2 \leq c'_k k^{1/8}.$$

The quantity $c_k(\lambda)$ is $O(k^{1/2})$, in view of (3), and for λ' sufficiently close to λ $c'_k \leq 2c_k(\lambda)$.

Thus, by (19) and (20) there exists an index $0 \leq j \leq k$ in neither S_1 nor S_2 . Take C_k to be any component of \mathcal{C}_j passing through $(0, z)$. By (6) it follows that $\chi(C_k) \geq -J_0(\mathcal{C}_j)$. Thus, C_k satisfies the requirements of Proposition 2.3. \square

4. LOW-ACTION CURVES FROM PERIODIC FLOER HOMOLOGY

This section proves Proposition 2.2 using the theory of periodic Floer homology (PFH), an analogue of ECH for area-preserving surface maps defined in [33, 36]. The proof is relatively simple compared to the proof of Proposition 2.3 above, which uses deep quantitative properties of ECH. We exploit an *algebraic* aspect of PFH that is not present in ECH, namely that PFH has many “ U -cycles” [22, 18].

4.1. Periodic Floer homology. We review the theory of PFH. We will discuss both the basics and some key new developments in the theory. Fix a closed, oriented surface Σ of genus g , an area form ω , and a diffeomorphism $\phi : \Sigma \rightarrow \Sigma$ preserving the area form.

4.1.1. Basics. The *mapping torus* of ϕ is the 3-manifold

$$Y_\phi := [0, 1] \times \Sigma / (1, p) \sim (0, \phi(p)).$$

Write t for the coordinate on the interval $[0, 1]$. The one-form dt on $[0, 1] \times \Sigma$ descends to a closed 1-form, also denoted by dt , on Y_ϕ . The area form ω defines a closed two-form ω_ϕ on Y_ϕ . The pair $\eta = (dt, \omega_\phi)$ is a framed Hamiltonian structure and the Reeb vector

field $R_\phi := R_\eta$ generates the suspension flow of ϕ . The mapping torus Y_ϕ fibers over the circle; write $V_\phi \rightarrow Y_\phi$ for the vertical tangent bundle.

Several key definitions carry over to this setting from ECH. In analogy with ECH, we will call periodic orbits of R_ϕ *closed Reeb orbits*. The definitions of elliptic/hyperbolic orbits from ECH have analogues here, replacing the bundle ξ with the bundle V_ϕ . Moreover, the ECH and J_0 indices are also defined in this setting, again replacing ξ with V_ϕ .

4.1.2. Rationality and monotonicity. We say ϕ is *rational* if the cohomology class $[\omega_\phi]$ is a real multiple of a rational class. We say that ϕ is *monotone* if it is rational and we have $c_1(V_\phi) = c[\omega_\phi]$ for some constant $c \in \mathbb{R}$. When $g \neq 1$, our monotonicity condition coincides with the monotonicity condition introduced by Seidel [50]. When $g = 1$, we show in Lemma A.3 that $c_1(V_\phi)$ always vanishes, so any rational area-preserving diffeomorphism is monotone in this case.

4.1.3. Definition of PFH. The definition of the version of PFH that we will use requires that ϕ is rational and also *nondegenerate*, meaning every closed Reeb orbit is either elliptic or hyperbolic. Choose a generic η -adapted almost-complex structure J on $\mathbb{R} \times Y_\phi$. Choose a union γ of embedded loops, transverse to Y_ϕ , called a *reference cycle*. Let Σ denote the homology class of a fiber of the map $Y_\phi \rightarrow S^1$. The *degree* $d(\gamma)$ of γ is the oriented intersection number of γ with $[\Sigma] \in H_2(Y_\phi; \mathbb{Z})$. We assume that $d(\gamma) > \max(0, g - 1)$ and that γ is *monotone*. This means that the homology class $\Gamma := [\gamma] \in H_1(Y_\phi; \mathbb{Z})$ satisfies the identity

$$(21) \quad c_1(V_\phi) + 2 \text{PD}(\Gamma) = c \cdot [\omega_\phi]$$

for some constant $c \neq 0$. The constant c is explicitly computable: pairing both sides of (21) with $[\Sigma]$ shows that the constant $c = 2A^{-1}(d - g + 1)$, where $A := \int_\Sigma \omega$. We note that (21) has a solution if and only if ϕ is rational, and that if (21) has a solution, it has solutions of arbitrarily high degree. Finally, let $K_\phi := \ker(\omega_\phi)$ denote the subgroup of all integral homology classes on which ω_ϕ integrates to 0.

The PFH chain complex $\text{PFC}_*(\phi, \gamma)$ is defined to be the vector space over $\mathbb{Z}/2$ freely generated by pairs $\Theta = (\alpha, Z)$ that we call *anchored ECH generators*. Here $\alpha = \{(\alpha_i, m_i)\}$ is an ECH generator such that $[\alpha] = [\gamma]$ and Z is an element of $H_2(Y_\phi, \alpha, \gamma)/K_\phi$ (recall that $H_2(Y_\phi, \alpha, \gamma)$ is an affine space over $H_2(Y_\phi; \mathbb{Z})$).

The differential ∂_J is defined similarly to the ECH differential, although now we take the relative homology classes of the holomorphic curves into account. Write $\mathcal{M}(\alpha, \beta, W)$ for the moduli space of holomorphic currents from α to β that represent the class $W \in H_2(Y, \alpha, \beta)$; let $\mathcal{M}_k(\alpha, \beta, W)$ denote the subspace of currents with ECH index k . Fix a pair of anchored ECH generators $\Theta = (\alpha, Z)$, $\Theta' = (\beta, Z')$. The matrix coefficient of ∂_J with respect to Θ and Θ' is defined by the formula

$$\langle \partial\Theta, \Theta' \rangle := \#_2 \mathcal{M}_1(\alpha, \beta, Z - Z') / \mathbb{R}.$$

Write $\text{PFH}_*(\phi, \gamma)$ for the homology of the complex $(\text{PFC}_*(\phi, \gamma), \partial_J)$. The PFH chain complex and homology group carry some additional basic features that we now review. There is a natural action of $H_2(Y_\phi; \mathbb{Z})$ on $\text{PFC}_*(\phi, \gamma, J)$; a class $W \in H_2(Y_\phi; \mathbb{Z})$ acts on a generator (α, Z) by sending it to $(\alpha, Z + W)$. This action commutes with the

differential and so descends to an action on $\text{PFH}_*(\phi, \gamma)$ as well. The U -map on PFH is also defined analogously to the U -map for ECH.

After choosing a framing of V_ϕ over γ , the PFH complex also comes equipped with a \mathbb{Z} -grading, which is defined for each anchored ECH generator by the formula

$$(22) \quad I(\Theta) := c_\tau(Z) + Q_\tau(Z) + \sum_i \sum_{k=1}^{m_i} \text{CZ}_\tau(\alpha_i^k).$$

The differential and U -map have degree -1 and -2 with respect to this grading. The H_2 -action shifts the grading as follows:

$$(23) \quad I(W \cdot \Theta) = I(\Theta) + \langle c_1(V_\phi) + 2\text{PD}(\Gamma), W \rangle = I(\Theta) + 2A^{-1}(d - g + 1) \int_W \omega_\phi$$

for any anchored ECH generator Θ and any $W \in H_2(Y_\phi; \mathbb{Z})$. The last line uses (21) and our computation of the monotonicity constant above. The identity (23) also shows that the \mathbb{Z} -grading is well-defined.

4.1.4. The U -cycle property. The analogue of a U -tower in ECH is a U -cycle. Assume that ϕ is nondegenerate and rational and choose a monotone reference cycle γ so that PFH is well-defined. A nonzero element $\sigma \in \text{PFH}_*(\phi, \gamma, G)$ is U -cyclic of order m for some integer $m \geq 1$ if

$$U^{m(d(\gamma)-g+1)}\sigma = (-m[\Sigma]) \cdot \sigma.$$

It is known that every nonzero element of PFH is U -cyclic as long as γ has sufficiently high degree.

Proposition 4.1 (Existence of U -cyclic elements, [18]). *Assume that ϕ is nondegenerate and rational and fix a monotone reference cycle γ . There exists an integer $d_0 > \max(0, g - 1)$, depending only on the Hamiltonian isotopy class of ϕ , such that if $d(\gamma) \geq d_0$, then $\text{PFH}_*(\phi, \gamma) \neq 0$ and every nonzero class is U -cyclic.*

The following lemma is a ‘‘chain-level’’ version of Proposition 4.1.

Lemma 4.2. *Assume that ϕ is nondegenerate and rational and fix a monotone reference cycle γ . There exists an integer $d_0 > \max(0, g - 1)$, depending only on the Hamiltonian isotopy class of ϕ , such that the following holds. Choose any monotone reference cycle γ such that $d(\gamma) \geq d_0$. Choose generic J and $z \in Y_\phi$ so that the chain-level map $U_{J,z}$ is well-defined. Then there exist positive integers m_0 and m_1 and a sequence $\{\Theta_j\}_{j=1}^{m_1}$ of nonzero generators of $\text{PFC}_*(\phi, \gamma)$ such that*

$$\langle U_{J,z}^{m_0(d(\gamma)-g+1)}\Theta_j, \Theta_{j+1} \rangle \neq 0$$

for each $j \in \{1, \dots, m_1 - 1\}$ and

$$\langle U_{J,z}^{m_0(d(\gamma)-g+1)}\Theta_{m_1}, m_0 m_1 [\Sigma] \cdot \Theta_1 \rangle \neq 0.$$

Proof. Suppose γ has sufficiently high degree so that Proposition 4.1 holds. Choose a trivialization of the restriction of V_ϕ to γ and use this to define a \mathbb{Z} -grading on PFH. Fix a grading k for which $\text{PFH}_k(\phi, \gamma) \neq 0$. Write $Z_k \subset \text{PFC}_k(\phi, \gamma)$ for the space of cycles of degree k , and B_k for the space of boundaries of degree k . The proof will take 3 steps.

Step 1: This step shows that Z_k and B_k have finite dimension over $\mathbb{Z}/2$. By the change of grading formula (23), it follows that for each ECH generator α there exists at most one anchored ECH generator $\Theta = (\alpha, Z)$ such that $I(\Theta) = k$. Since ϕ is nondegenerate, it has finitely many ECH generators representing any given homology class in $H_1(Y_\phi; \mathbb{Z})$. This implies that $\text{PFC}_k(\phi, \gamma)$ contains finitely many anchored ECH generators, so it is a finite-dimensional $\mathbb{Z}/2$ -vector space. This implies that Z_k and B_k have finite dimension as well.

Step 2: This step uses Proposition 4.1 to show that there is a nonzero cycle $x \in Z_k$ fixed up to a shift by an iterate of the U -map. Fix generic J and z so that the chain-level map $U_{J,z}$ is well defined. Proposition 4.1 implies that there exists an integer $m_0 \geq 1$ such that for any nonzero cycle $x \in Z_k$, there exists a chain z such that

$$(24) \quad m_0[\Sigma] \cdot U_{J,z}^{m_0(d(\gamma)-g+1)} x = x + \partial z.$$

Let T be the restriction of $m_0[\Sigma] \cdot U_{J,z}^{m_0(d(\gamma)-g+1)}$ to Z_k . Then (24) implies that $\text{Im}(T - 1) \subseteq B_k$. Since $\text{PFH}_k(\phi, \gamma) \neq 0$, it follows that $\dim(Z_k) > \dim(B_k)$. This implies that the operator $T - 1$ has nonzero kernel and therefore there exists some nonzero $x \in Z_k$ such that $Tx = x$.

Step 3: This step completes the proof. Expand the element x from the previous step into a sum $\sum_{i=1}^N x_i$ where each x_i is an anchored ECH generator. The desired cyclic sequence $\{\Theta_j\}_{j=1}^{m_1}$ of anchored ECH generators will be picked out from the x_i using a short combinatorial argument. Define a directed graph G as follows. The vertex set of G is $\{1, \dots, N\}$ and there is an edge from i to j if and only if $\langle Tx_i, x_j \rangle \neq 0$. We allow edges to start and end at the same vertex. It is well-known that any directed graph with no *sources*, i.e. vertices which have no incoming edges, has a directed cycle. Now, G has no sources: this follows because $Tx = x$ implies that for each j , the identity $\langle Tx, x_j \rangle \neq 0$ holds, which in turn implies that there exists some i such that G has an edge from i to j . Thus, G has a cycle. Thus, there exists a set $\{x'_j\}_{j=1}^{m_1}$ of anchored ECH generators such that

$$(25) \quad \langle Tx'_j, x'_{j+1} \rangle \neq 0, \quad \langle Tx'_{m_1}, x'_1 \rangle \neq 0,$$

for each $j \in \{1, \dots, m_1 - 1\}$. For each j , set $\Theta_j := -(j - 1)m_0[\Sigma] \cdot x'_j$. Then, the Θ_j satisfy the conditions of the lemma by (25), since $T = m_0[\Sigma] \cdot U_{J,z}^{m_0(d(\gamma)-g+1)}$ by definition. \square

4.2. Proof of Proposition 2.2. We now suppose that ϕ is *monotone*, which we recall means $c_1(V_\phi) = c[\omega_\phi]$ for some constant $c \in \mathbb{R}$. The proof of Proposition 2.2 is an immediate consequence of the following result, since the monotonicity condition is preserved under Hamiltonian isotopy.

Proposition 4.3. *Assume that ϕ is nondegenerate and monotone. There exists an integer $d_0 \geq \max(0, g - 1)$ depending only on g and the Hamiltonian isotopy class such that for any $z \in Y_\phi$ not on any closed Reeb orbit, and generic J , there exists a standard J -holomorphic curve $u_d : C_d \rightarrow \mathbb{R} \times Y_\phi$ satisfying the following properties:*

- (a) $(0, z) \in u_d(C_d)$.
- (b) $\mathcal{E}(u_d) \leq d$.

- (c) $\mathcal{A}(u_d) \leq d^{-1/2}$.
(d) $\chi(C_d) \geq -2$.

Proof. Fix $d_0 > 0$ so that Lemma 4.2 holds and fix any monotone reference cycle γ with degree $d := d(\gamma) \geq d_0$. The proof will take 2 steps.

Step 1: Fix generic J and z so that the map $U_{J,z}$ is well-defined on $\text{PFC}_*(\phi, \gamma)$. Let $\{\Theta_j\}_{j=1}^{m_1}$ denote the sequence of generators provided by Lemma 4.2. Then, by Lemma 4.2, we obtain a sequence $\mathcal{C}_1, \dots, \mathcal{C}_{m_0 m_1 (d-g+1)}$ of J -holomorphic currents counted by the U -map such that

$$\sum_{i=1}^{m_1 m_0 (d(\gamma) - g + 1)} [\mathcal{C}_i] = m_0 m_1 [\Sigma] \in H_2(Y_\phi, \alpha_1, \alpha_1) = H_2(Y_\phi; \mathbb{Z}) / K_\phi.$$

By additivity of the action and of J_0 we therefore obtain

$$(26) \quad \sum_{i=1}^{m_1 m_0 (d(\gamma) - g + 1)} \mathcal{A}(\mathcal{C}_i) = m_0 m_1 A, \quad \sum_{i=1}^{m_1 m_0 (d(\gamma) - g + 1)} J_0(\mathcal{C}_i) = 2m_0 m_1 (d(\gamma) + g - 1).$$

In the equality for J_0 , we have used the fact that ϕ is monotone, which implies that $c_1(V_\phi)$ has zero pairing with K_ϕ , together with the fact that $J_0([\Sigma]) = 2(d(\gamma) + g - 1)$.

Step 2: This step finishes the proof of the proposition. Write S_1 for the set of i such that $\mathcal{A}(\mathcal{C}_i) > d(\gamma)^{-1/2}$ and write S_2 for the set of i such that $J_0(\mathcal{C}_i) \geq 3$. Then, by nonnegativity of the action of pseudoholomorphic curves, and (26), we have

$$(27) \quad \#S_1 \leq A m_0 m_1 (d(\gamma))^{1/2}.$$

Since the J_0 index is bounded below by -1 , the bound (26) implies

$$(28) \quad \#S_2 \leq m_0 m_1 (3d(\gamma) + g - 1) / 4.$$

Thus, after possibly increasing d_0 , we have the strict inequality

$$\#(S_1 \cup S_2) < m_0 m_1 (d(\gamma) - g + 1).$$

This implies that there exists some i such that $\mathcal{A}(\mathcal{C}_i) \leq (d(\gamma))^{-1/2}$ and $J_0(\mathcal{C}_i) \leq 2$. Thus, the component $u_d : C_d \rightarrow \mathbb{R} \times Y_\phi$ of \mathcal{C}_i containing $(0, z)$ has $\mathcal{A}(u_d) \leq (d(\gamma))^{1/2}$ and $\chi(C_d) \geq -2$, by (6). It remains to show that $\mathcal{E}(u_d) \leq d(\gamma)$: this follows since, as dt is closed, the integral of dt over any level set of \mathcal{C}_i is equal to the pairing $\langle dt, [\gamma] \rangle = d(\gamma)$. \square

5. INVARIANT SETS FROM LOW-ACTION HOLOMORPHIC CURVES

The purpose of this section is to prove Theorem 6. For the remainder of the section, we fix a closed, smooth, connected, oriented manifold Y of odd dimension $2n + 1 \geq 3$.

5.1. Notational preliminaries. Let us begin by reviewing the setup.

5.1.1. Stable constants. The statements and proofs below will involve several constants which depend on Y , η , and J , where η is a framed Hamiltonian structure on Y and J is an η -adapted almost-complex structure. We say that such a constant is *stable* if it can be taken to be invariant under C^∞ -small perturbations of η and J .

5.1.2. *Geometry of symplectizations.* Let $\mathcal{D}(Y)$ be the space of pairs (η, J) where η is a framed Hamiltonian structure and J is an η -adapted almost-complex structure of $\mathbb{R} \times Y$. We equip $\mathcal{D}(Y)$ with the topology of C^∞ -convergence. That is, a sequence $\{(\eta_k = (\lambda_k, \omega_k), J_k)\}$ in $\mathcal{D}(Y)$ converges to $(\eta = (\lambda, \omega), J)$ if and only if the sequences $\{\lambda_k\}$, $\{\omega_k\}$, and $\{J_k\}$ converge in the C^∞ -topology to λ , ω , and J , respectively. Choose a pair $(\eta = (\lambda, \omega), J) \in \mathcal{D}(Y)$. To this pair we associate the following translation-invariant and J -invariant Riemannian metric on $\mathbb{R} \times Y$:

$$g := da \otimes da + \lambda \otimes \lambda + \omega(-, J-).$$

We fix notation for norms of tensors with respect to g . For any smooth tensor \mathcal{T} on $\mathbb{R} \times Y$, write $|\mathcal{T}|_g$ for its pointwise g -norm, which is a smooth function on $\mathbb{R} \times Y$. Write $\|\mathcal{T}\|_g := \sup_{z \in \mathbb{R} \times Y} |\mathcal{T}|_g(z)$ for the C^0 norm of \mathcal{T} with respect to g . We fix notation for the metric balls of g . Let dist_g denote the distance function of g . Omitting the dependence on g for brevity, we let

$$\overline{B}_r(z) := \{w \in \mathbb{R} \times Y \mid \text{dist}_g(z, w) \leq r\}$$

denote the closed metric ball of radius $r > 0$ centered at $z \in \mathbb{R} \times Y$.

5.1.3. *Geometry of J -holomorphic curves.* Fix a J -holomorphic curve $u : C \rightarrow \mathbb{R} \times Y$. We say u is *compact* and *connected* if the domain C is respectively compact and connected. We say u is *generally immersed* if the critical point set $\text{Crit}(u)$ is discrete. This is always true if C is connected and u is not a constant map. We say u is *boundary immersed* if the restriction of u to ∂C is an immersion. We assume for the sake of convenience that any J -holomorphic curves is generally immersed and boundary immersed unless stated otherwise.

We let $\gamma := u^*g$ denote the pullback metric on C , which is defined at any point $z \in C$ such that $du(z) \neq 0$. In particular, for a generally immersed curve, the metric is defined outside of a discrete subset of points. Let $\alpha := u^*\lambda$ denote the pullback of λ . Let $|\mathcal{T}|_\gamma$ and $\|\mathcal{T}\|_\gamma$ denote the pointwise and C^0 norms of a tensor \mathcal{T} with respect to γ . For any domain $U \subseteq C$, we define its *area* to be the integral of the volume form of γ over the set $U \setminus \text{Crit}(u)$:

$$\text{Area}_\gamma(U) := \int_{U \setminus \text{Crit}(u)} \text{dvol}_\gamma.$$

5.2. **The connected-local area bound and its significance.** The main estimate required for the proof of Theorem 6 is a so-called ‘‘connected-local area bound’’. In this section we state this estimate, deferring the proof to Section 5.5, and then use it to prove Theorem 6.

Given a J -holomorphic curve $u : C \rightarrow \mathbb{R} \times Y$, any point $\zeta \in C$, and any $r > 0$, let $S_r(\zeta)$ denote the connected component of $u^{-1}(\overline{B}_r(u(\zeta)))$ containing ζ . Our estimate gives an a priori bound on the area of $S_r(\zeta)$ assuming that $\mathcal{A}(u)$ is small and r is small. The bound depends on the Euler characteristic of C , which is the primary reason why Euler characteristic bounds are assumed in Theorem 6.

Proposition 5.1 (Connected-local area bound for low-action curves). *Fix $(\eta, J) \in \mathcal{D}(Y)$. There exists stable constants $\epsilon_7 = \epsilon_7(\eta, J) > 0$ and $\epsilon_8 = \epsilon_8(\eta, J) > 0$ such that*

the following holds. Let $u : C \rightarrow \mathbb{R} \times Y$ be a standard J -holomorphic curve such that $\mathcal{A}(u) \leq \epsilon_7$. Then for any point $\zeta \in C$, we have the bound

$$(29) \quad \text{Area}_\gamma(S_{\epsilon_8}(\zeta)) \leq \epsilon_8^{-1}(\chi(C)^2 + 1).$$

Proposition 5.1 is inspired by [24, Thm. 5], the ‘‘asymptotic connected-local area bound’’ for ‘‘feral’’ curves, a fundamental result in the new Fish-Hofer theory. There are several new aspects here. The main point is that [24, Thm. 5] is an asymptotic result for a fixed curve (of possibly unbounded Hofer energy, but finite action.) This allows for the reduction to annular curves, whereas in our setting we need to consider much more complicated topologies; new arguments are required for this. Another novelty is the replacement of the asymptotic condition with a bound on the action instead. Both here and in [24] one also has to be very careful with the constants to ensure stability.

Let us now explain why Theorem 6 follows from Proposition 5.1. The argument for this appeared in the feral context in [24, Proposition 4.47].

Proof of Theorem 6. As the argument from here is essentially the same as [24, Proposition 4.47] except for minor modifications, we will only provide a sketch highlighting the key points. Fix $(\eta, J) \in \mathcal{D}(Y)$ and a sequence $\{(\eta_k, J_k)\}$ in $\mathcal{D}(Y)$ converging to it. Fix a sequence $\{u_k : C_k \rightarrow Y\}$ of standard J_k -holomorphic curves such that $\lim_{k \rightarrow \infty} \mathcal{A}(u_k) = 0$ and $\inf_k \chi(C_k) > -\infty$. Let $\mathcal{X} \subset \mathcal{K}(X)$ denote the limit set and choose any $\bar{\Lambda} \in \mathcal{X}$. When k is sufficiently large, and therefore $\mathcal{A}(u_k)$ is sufficiently small, it follows that $(a \circ u_k)(C_k) = \mathbb{R}$, hence $\bar{\Lambda}$ is non-empty. To see that $\bar{\Lambda} = (-1, 1) \times \Lambda$, where $\Lambda \in \mathcal{K}(Y)$ is R_η -invariant, it suffices to show that for any $z := (t, y) \in \bar{\Lambda}$ there exist some $\epsilon > 0$ such that

$$(t + \tau, y) \in \bar{\Lambda}, \quad (t, \phi^\tau(y)) \in \bar{\Lambda}$$

for any $\tau \in (-\epsilon, \epsilon)$, where $\{\phi^t\}_{t \in \mathbb{R}}$ denotes the flow of R_η . This is proved by a standard application of Fish’s target-local Gromov compactness theorem (the version stated in [24, Theorem 2.36]). Fix points $\zeta_k \in C_k$ such that $u_k(\zeta_k) \rightarrow z$ and define local patches $S_k := S_{\epsilon_8}(\zeta_k) \subset C_k$. Proposition 5.1 gives a k -independent upper bound on $\text{Area}_\gamma(S_k)$ in view of the fact that $\mathcal{A}(u_k) \rightarrow 0$, $\{\chi(C_k)\}$ is uniformly bounded, and all constants are stable. The surfaces S_k have uniformly bounded genus as well, because of the uniform bound on $\chi(C_k)$. Thus, we are justified in applying target-local Gromov compactness, and after passing to the limit obtain a non-constant holomorphic curve passing through z and contained in $\bar{\Lambda}$. The curve has action 0, due to the fact that $\mathcal{A}(u_k) \rightarrow 0$, and hence has tangent plane at any immersed point equal to $\text{Span}(\partial_a, R_\eta)$, from which the desired property follows. \square

5.3. Properties of the limit set. Before diving into the proof of the connected-local area bound, we provide a proof of Proposition 1.9 from the introduction. We start with a pair of lemmas. To state the first lemma, it is useful to fix the following notation. For any $(\eta', J') \in \mathcal{D}(Y)$, any standard J' -holomorphic curve $u : C \rightarrow \mathbb{R} \times Y$ defines a map

$$\begin{aligned} \mathcal{S}_u : \mathbb{R} &\rightarrow \mathcal{K}(X), \\ s &\mapsto \tau_s \cdot \left(u(C) \cap (s - 1, s + 1) \times Y \right). \end{aligned}$$

We make the following assertion:

Lemma 5.2. *For any $(\eta', J') \in \mathcal{D}(Y)$ and any standard J' -holomorphic curve $u : C \rightarrow \mathbb{R} \times Y$, the map \mathcal{S}_u is continuous.*

The next lemma is a purely topological statement.

Lemma 5.3. *Let K be any compact and metrizable topological space. Let $Z_k \subseteq K$ be a sequence of connected subsets. Then there exists a subsequence $W_j := Z_{k_j}$ such that their set of subsequential limit points is connected.*

Lemma 5.2 has a short and elementary proof.

Proof of Lemma 5.2. It suffices to prove that for any sequence $s_k \rightarrow s$ in \mathbb{R} , we have $\mathcal{S}_u(s_k) \rightarrow \mathcal{S}_u(s)$ in $\mathcal{K}(X)$. The proof will take 2 steps.

Step 1: This step proves that $\limsup \mathcal{S}_u(s_k) \subseteq \mathcal{S}_u(s)$. Choose any point $z = (t, y) \in \limsup \mathcal{S}_u(s_k) \subseteq X$. Then any neighborhood of z meets infinitely many of the sets

$$\mathcal{S}_u(s_k) = \tau_{s_k} \cdot \left(u(C) \cap (s_k - 1, s_k + 1) \times Y \right).$$

After passing to a subsequence, there exists a sequence of points $z_k = (t_k, y_k) \in \mathcal{S}_u(s_k)$ such that $t_k \rightarrow t$ and $y_k \rightarrow y$. Choose some small $\epsilon > 0$ such that $t \pm \epsilon \in (-1, 1)$. Note that $\tau_{-s_k}(z_k) \in u(C)$ for each k . Since $s_k \rightarrow s$ and $t_k \rightarrow t$, it follows that

$$\tau_{-s_k}(z_k) \in u(C) \cap [s + t - \epsilon, s + t + \epsilon] \times Y$$

for sufficiently large k . Since $u(C)$ is proper, the set on the right-hand side is closed. The points $\tau_{-s_k}(z_k)$ converge to $\tau_{-s}(z)$, so it follows that

$$\tau_{-s}(z) \in u(C) \cap [s + t - \epsilon, s + t + \epsilon] \times Y \subset u(C) \cap (s - 1, s + 1) \times Y$$

and therefore that $z \in \mathcal{S}_u(s)$.

Step 2: This step proves that $\mathcal{S}_u(s) \subseteq \liminf \mathcal{S}_u(s_k)$. Fix any $z = (t, y) \in \mathcal{S}_u(s)$. We must show that any neighborhood of z has non-empty intersection with all but finitely many of the sets $\mathcal{S}_u(s_k)$. Write $z_k := \tau_{s_k - s}(z) = (-s_k + s + t, y)$. For all sufficiently large k , we have that $-s_k + s + t \in (-1, 1)$, so it follows that $\tau_{-s_k}(z_k) \in u(C) \cap (s_k - 1, s_k + 1) \times Y$. This implies $z_k \in \mathcal{S}_u(s_k)$ for sufficiently large k . Since $z_k \rightarrow z$, we conclude that any neighborhood of z meets $\mathcal{S}_u(s_k)$ for sufficiently large k . □

Next, we prove Lemma 5.3.

Proof of Lemma 5.3. Let $W_j := Z_{k_j}$ be any subsequence for which there exist points $w_j \in W_j$ converging to some $w \in K$. Let W denote the set of subsequential limit points of $\{W_j\}$. The proof that W is connected will take 4 steps.

Step 1: This step makes some simplifying observations. It suffices to show that, for any pair of disjoint open subsets U and V such that $W \subseteq U \sqcup V$, we have either $W \subseteq U$ or $W \subseteq V$. Without loss of generality, we will assume that the point $w \in W$ is contained in U , and then show that $W \subseteq U$.

Step 2: This step proves that $W_j \subseteq U \sqcup V$ for all sufficiently large j . We prove it by contradiction. Assume that the claim is false. Then, after passing to a subsequence, there exists for each j some $w'_j \in W_j$ such that $w'_j \notin U \sqcup V$ for any j . Since K is

compact and metrizable, it is sequentially compact, and therefore a subsequence of $\{w'_j\}$ converges to some $w' \in W$. However, the complement of $U \sqcup V$ is closed, so w' does not lie in $U \sqcup V$, contradicting the initial assumption that $W \subseteq U \sqcup V$.

Step 3: This step proves that $W_j \subseteq U$ for all sufficiently large j . By Step 2 and the fact that W_j is connected for each j , we either have $W_j \subseteq U$ or $W_j \subseteq V$ for large j . Since $w_j \rightarrow w \in U$, it follows that W_j has non-empty intersection with U for sufficiently large j , which implies $W_j \subseteq U$.

Step 4: This step completes the proof. By Step 3, it follows that W lies inside the closure \overline{U} . Any open set intersecting \overline{U} must intersect U , so V is disjoint from \overline{U} . We conclude that $W \subseteq U$. □

We now give the promised proof of Proposition 1.9.

Proof of Proposition 1.9. Fix $(\eta, J) \in \mathcal{D}(Y)$ and a sequence $\{(\eta_k, J_k)\}$ in $\mathcal{D}(Y)$ converging to it. Fix a sequence $\{u_k : C_k \rightarrow \mathbb{R} \times Y\}$ of standard J_k -holomorphic curves. Recall that each curve u_k defines a continuous map $\mathcal{S}_k := \mathcal{S}_{u_k}$ from \mathbb{R} to $\mathcal{K}(X)$. Write $Z_k := \mathcal{S}_k(\mathbb{R})$ for each k . By Lemma 5.2, Z_k is the image by a continuous map of a connected space, so it is connected. Thus, $\{Z_k\}$ is a sequence of connected subsets of the compact and metrizable space $\mathcal{K}(X)$. The limit set \mathcal{X} is equal to the set of subsequential limit points of the sequence $\{Z_k\}$. The proposition now follows from Lemma 5.3. □

5.4. Preliminaries from feral curve theory. It remains to prove Proposition 5.1, which will take up the remainder of the paper. To do this, we need to first collect and review some more preliminaries from the work of Fish–Hofer [24], which is the purpose of this section.

5.4.1. Perturbed holomorphic curves. Many of the estimates in [24] are easiest when the height function $a \circ u$ of the pseudoholomorphic curve is Morse. Unfortunately, this is not in general the case, so the notion of a “perturbed curve” was introduced there to remedy this. More precisely, a *perturbed J -holomorphic curve* is a pair (u, f) where $u : C \rightarrow \mathbb{R} \times Y$ is a J -holomorphic curve and $f : C \rightarrow \mathbb{R}$ is a smooth function which is compactly supported in the open subset $C \setminus (\partial C \cup \text{Crit}(u))$. Perturbing u in the vertical direction by f defines a new map

$$\tilde{u} : \zeta \mapsto \exp_{u(\zeta)}^g(f(\zeta)\partial_a).$$

Write $\tilde{\gamma} := \tilde{u}^*g$ for the induced pullback metric. Define an almost-complex structure \tilde{j} on C as the unique almost-complex structure which is a $\tilde{\gamma}$ -isometry and coincides with j on the complement of $\text{supp}(f)$. We then define a one-form

$$\tilde{\alpha} := -(\tilde{u}^*da \circ \tilde{j})$$

on C . This should be thought of as a perturbation of $\alpha = u^*\lambda = -(u^*da \circ j)$.

As in [24], the perturbation is subject to some quantitative controls. This is encoded in the notion of a (δ, ϵ) -tame perturbation, which is found in [24, Definition 4.24]. We will not repeat the precise definition of a (δ, ϵ) -tame perturbation, since we never make explicit use of it, but we will recall its key attributes. The constant $\delta > 0$ controls the support of f and the domain where $a \circ \tilde{u}$ is Morse. That is, f is supported outside

a $\delta/2$ -neighborhood of $\text{Crit}(u)$, and $a \circ \tilde{u}$ is Morse outside of a δ -neighborhood. The constant ϵ controls the C^2 norm of f . We require δ to be smaller than a constant depending on u and ambient geometry, and ϵ to be smaller than a constant depending on δ , u , and ambient geometry.

Given a tame perturbation (u, f) , several explicit estimates are proved in [24, Section 4.1] relating the perturbed map \tilde{u} to the unperturbed map u . We do not make direct use of the majority of them, although they are key ingredients in the proofs of other estimates from [24] that we cite. One estimate which we do use relates the perturbed and unperturbed pullback metrics:

Lemma 5.4 ([24, Equation 4.3]). *Fix $(\eta, J) \in \mathcal{D}(Y)$. Fix a J -holomorphic curve $u : C \rightarrow \mathbb{R} \times Y$ and any (δ, ϵ) -tame perturbation (u, f) . Then we have*

$$\gamma/2 \leq \tilde{\gamma} \leq 2\gamma,$$

where we recall that $\gamma = u^*g$ and $\tilde{\gamma} = \tilde{u}^*g$.

We also remark that if the curve u is immersed, then we can use much simpler perturbations, because there are no singularities for the perturbation to avoid. Call a perturbation ϵ -Morse if $a \circ \tilde{u}$ is Morse over the entire domain C and if the C^2 norm of f is smaller than ϵ . Our main dynamical results (Theorems 1–3) only require working with embedded holomorphic curves. A reader who is solely interested in these results can replace all instances of (δ, ϵ) -tame perturbations with ϵ -Morse perturbations, where ϵ is very small. Lemma 5.4 is also straightforward if u is immersed. In this case $\tilde{\gamma}$ is ϵ -close to γ in the C^1 topology. In the non-immersed setting, some care is required because γ degenerates as it approaches $\text{Crit}(u)$.

5.4.2. Tracts. Fix a perturbed J -holomorphic curve (u, f) . Tracts and strips, introduced in [24], are highly structured compact portions of the domain C . They are important because they satisfy several geometric estimates. A *tract* in (u, f) is a connected, compact embedded surface $\tilde{C} \subset C$, possibly with boundary and corners, such that:

- (i) The boundary $\partial\tilde{C}$ is disjoint from $\text{Crit}(a \circ \tilde{u})$;
- (ii) The boundary of \tilde{C} decomposes as $\partial_h\tilde{C} \cup \partial_v\tilde{C}$ where
 - (a) $\partial_h\tilde{C} \cap \partial_v\tilde{C}$ is a finite set of corners;
 - (b) If it is non-empty, each component of $\partial_v\tilde{C}$ is tangent to the gradient vector field $\text{grad}_{\tilde{\gamma}}(a \circ \tilde{u})$;
 - (c) The function $a \circ \tilde{u}$ is constant on each component of $\partial_h\tilde{C}$.

The sets $\partial_h\tilde{C}$ and $\partial_v\tilde{C}$ are respectively called the *horizontal* and *vertical* boundaries of the tract \tilde{C} .

5.4.3. Strips. For the proof of Proposition 5.12 at the end of this section, we need to introduce the notion of a strip. Strips, which are closely related to tracts, were originally defined in [24, Definition 4.17]; we write down a simpler version of their definition which is sufficient for our purposes. A *strip* (in our sense) in (u, f) is an embedded rectangle $\tilde{C} \subset C$, i.e. a compact domain of genus zero with four smooth




FIGURE 1. A tract (blue) and a strip (red).

boundary curves meeting in four corners, satisfying the following two properties. First, \tilde{C} contains no critical points of $a \circ \tilde{u}$. Second, the boundary $\partial\tilde{C}$ decomposes as a union $\partial_h\tilde{C} \cup \partial_v\tilde{C}$ with certain properties. The set $\partial_v\tilde{C}$ is called the *vertical boundary*, and consists of a pair of disjoint smooth curves tangent to $\text{grad}(a \circ \tilde{u})$. The set $\partial_h\tilde{C}$ is called the *horizontal boundary*, and consists of a pair of disjoint smooth curves, which each meet the segments of $\partial_v\tilde{C}$ at the corners. An illustration is provided in Figure 1. We label the two components of $\partial_h\tilde{C}$ with the notation $\partial_h^\pm\tilde{C}$. The labelling is uniquely determined by the condition $\sup_{\zeta \in \partial_h^+\tilde{C}}(a \circ \tilde{u})(\zeta) > \sup_{\zeta \in \partial_h^-\tilde{C}}(a \circ \tilde{u})(\zeta)$. We call $\partial_h^+\tilde{C}$ the *top horizontal boundary* and $\partial_h^-\tilde{C}$ the *bottom horizontal boundary*.

The strips we use below are always defined by fixing a segment \mathcal{I} mapping into a level set of $a \circ \tilde{u}$ and taking a union of gradient flow segments starting at points $\zeta \in \mathcal{I}$. The resulting strip \tilde{C} will have $\partial_h^-\tilde{C} = \mathcal{I}$ and $\partial_h^+\tilde{C}$ consisting of the endpoints of these gradient flow segments.

5.4.4. *Exponential area bound for tracts.* In the classical theory of J -holomorphic curves in symplectizations, area bounds are deduced from Hofer energy bounds. The following result provides an alternative in the absence of Hofer energy bounds. It is arguably the most crucial estimate in [24]. For example, it shows up in the proofs of Proposition 5.8, Lemma 5.6, and Lemma 5.7 below.

Proposition 5.5 ([24, Theorem 9]). *Fix $(\eta, J) \in \mathcal{D}(Y)$. Let (u, f) be a (δ, ϵ) -tame perturbed J -holomorphic curve with domain C and let $\tilde{C} \subset C$ be a tract of (u, f) , for which there exist constants $a_+ > a_-$ such that*

- (i) $(a \circ \tilde{u})(\tilde{C}) \subset [a_-, a_+]$;
- (ii) $(a \circ \tilde{u})(\partial_h\tilde{C}) \cap (a_-, a_+) = \emptyset$;
- (iii) a_+ and a_- are regular values of $a \circ \tilde{u}$.

Then there exists a stable constant $c_1 = c_1(\eta, J) \geq 1$ such that

$$(30) \quad \text{Area}_{\tilde{\gamma}}(\tilde{C}) \leq c_1 e^{c_1(a_+ - a_-)} \left(\int_{(a \circ u)^{-1}(a_-) \cap \tilde{C}} \tilde{\alpha} + \mathcal{A}(u) \right).$$

Given its importance, we briefly summarize the proof of Proposition 5.5 for the reader. One can compute

$$\text{Area}_{\tilde{\gamma}}(\tilde{C}) = \int_{\tilde{C}} \tilde{u}^* da \wedge \tilde{\alpha} + \tilde{u}^* \omega = \int_{a_-}^{a_+} \left(\int_{(a \circ \tilde{u})^{-1}(t) \cap \tilde{C}} \tilde{\alpha} \right) dt + \mathcal{A}(u).$$

Write $e(t)$ for the $\tilde{\alpha}$ -measure of $(a \circ \tilde{u})^{-1}(t) \cap \tilde{C}$. Whenever the interval $[t_-, t_+]$ contains no critical values, one proves the bound $e(t_+) - e(t_-) \lesssim \int_{t_-}^{t_+} e(s) ds + \delta$, where δ is some controllable error term. Using Gronwall's inequality, it follows that $e(t)$ grows exponentially in t . This is plugged into the computation above to deduce (30).

5.4.5. Strip estimates. We now state two technical estimates for strips of perturbed J -holomorphic curves, which are used only in the proof of Proposition 5.12 at the end of this section. Lemma 5.6, which is a simpler version of a lemma from [24], shows that the length of the bottom horizontal boundary of a strip is controlled by the length of the top horizontal boundary, provided that the strip is not too tall.

Lemma 5.6 ([24, Lemma 4.21]). *Fix $(\eta, J) \in \mathcal{D}(Y)$. Then there exists a stable constant $c_2 = c_2(\eta, J) \geq 1$ such that the following holds. Let \tilde{C} be a strip of a (δ, ϵ) -tame perturbed J -holomorphic curve (u, f) such that $a \circ \tilde{u}$ is constant on $\partial_h^- \tilde{C}$ and*

$$\sup_{\zeta \in \tilde{C}} (a \circ \tilde{u})(\zeta) - \inf_{\zeta \in \tilde{C}} (a \circ \tilde{u})(\zeta) \leq c_2^{-1}$$

Then we have the bound

$$\int_{\partial_h^- \tilde{C}} \tilde{\alpha} \leq 2 \int_{\partial_h^+ \tilde{C}} \tilde{\alpha} + 2c_2 \int_{\tilde{C}} u^* \omega.$$

Lemma 5.6 is proved using Proposition 5.5. To state the next lemma, we need another definition. A strip \tilde{C} is *rectangular* if both $\partial_h^\pm \tilde{C}$ are contained inside level sets of $a \circ \tilde{u}$. The lemma asserts that for any finite collection of strips which are not too tall or too short, one of them contains a gradient flow line of $a \circ \tilde{u}$ of controlled length running from the bottom horizontal boundary component to the top horizontal boundary component.

Lemma 5.7 ([24, Lemma 4.23]). *Fix $(\eta, J) \in \mathcal{D}(Y)$. There exists a stable constant $c_3 = c_3(\eta, J) \geq 1$ such that the following holds. Let (u, f) be a (δ, ϵ) -tame perturbed J -holomorphic curve. Let $\{\tilde{C}_k\}_{k=1}^n$ be a finite set of rectangular strips of (u, f) satisfying the following properties:*

- (i) $a_0 = \inf_{\zeta \in \tilde{C}_k} (a \circ \tilde{u})(\zeta)$ is independent of k ;
- (ii) $a_1 = \sup_{\zeta \in \tilde{C}_k} (a \circ \tilde{u})(\zeta)$ is independent of k ;
- (iii) $a_1 - a_0 \leq c_3^{-1}$;
- (iv) $\sum_{k=1}^n \int_{\tilde{C}_k} u^* \omega \leq (a_1 - a_0) \sum_{k=1}^n \int_{\partial_h^- \tilde{C}_k} \tilde{\alpha}$;
- (v) Each of the strips $\{\tilde{C}_k\}_{k=1}^n$ are pairwise disjoint.

Then there exists $k \in \{1, \dots, n\}$ and a smooth map $q : [0, s] \rightarrow \tilde{C}_k$ such that

$$\dot{q}(s) = \text{grad}_{\tilde{\gamma}}(a \circ \tilde{u})(q(s)), \quad (a \circ \tilde{u} \circ q)(0) = a_0, \quad (a \circ \tilde{u} \circ q)(s) = a_1$$

and

$$\text{length}_{\tilde{\gamma}}(q([0, s])) \leq c_3(a_1 - a_0).$$

The proof of Lemma 5.7 is rather intricate, since the lengths of gradient flow lines could be heavily distorted in the areas where the norm of $\text{grad}(a \circ \tilde{u})$ is small. We refer the reader to [24, Section 4.3.1] for details.

5.4.6. *Action quantization.* Fish–Hofer [24, Theorem 4] proved that a holomorphic curve $u : C \rightarrow \mathbb{R} \times Y$ has a positive lower bound on its action near any interior global maximum/minimum of the function $a \circ u$. Their lower bound depends on the genus of C , which suffices for our intended applications. For the sake of a cleaner statement, we note that the bound can be made genus-independent using the compactness theory of J -holomorphic currents [48, Remark 5.20]. We now state the precise quantization result.

Proposition 5.8 ([24, Theorem 4]). *Fix $(\eta, J) \in \mathcal{D}(Y)$. Fix any real number $s > 0$. There exists a stable constant $\hbar = \hbar(\eta, J, s) > 0$ such that, for any compact, connected J -holomorphic curve $u : C \rightarrow \mathbb{R} \times Y$, we have*

$$\mathcal{A}(u) \geq \hbar > 0$$

provided that the following properties are satisfied for some $a_0 \in \mathbb{R}$:

- (i) Either $\min_{\zeta \in C}(a \circ u)(\zeta)$ or $\max_{\zeta \in C}(a \circ u)(\zeta)$ is equal to a_0 ;
- (ii) $(a \circ u)(\partial C) \cap [a_0 - s, a_0 + s] = \emptyset$.

Proposition 5.8 is proved via a contradiction argument using Proposition 5.5 and Fish’s target-local Gromov compactness theorem [23].

5.4.7. *Geodesic distance lemma.* The following elementary lemma is used in the proof of Lemma 5.17 below.

Lemma 5.9 ([24, Lemma 4.29]). *Fix $(\eta, J) \in \mathcal{D}(Y)$. There exists a stable constant $\epsilon_4 = \epsilon_4(\eta, J) \in (0, 1/100)$ such that the following holds for any $\epsilon \leq \epsilon_4$. Each smooth unit-speed immersion $q : [0, T] \rightarrow \mathbb{R} \times Y$ such that*

- (i) $\lambda(\dot{q}(t)) > 0$ for each $t \in [0, T]$;
- (ii) $\epsilon \leq \int_q \lambda \leq 10\epsilon$;
- (iii) the set of $t \in [0, T]$ such that $\lambda(\dot{q}(t)) < 1/2$ has Lebesgue measure at most ϵ satisfies the bound

$$\text{dist}_q(q(0), q(T)) > \epsilon/2.$$

The bound in Lemma 5.9 is proved by direct computation in geodesic local coordinates y^1, \dots, y^{2n+1} , with the additional condition that λ is close to dy^1 . We need to state Lemma 5.9 here, instead of alongside Lemma 5.17, because the constant ϵ_4 appears in the statement of Proposition 5.10 below.

5.5. Proof of the connected local area bound. We now prove Proposition 5.1. In fact, this will be deduced from a more technical bound, Proposition 5.10, which we now state. The statement requires defining a stable constant

$$(31) \quad \epsilon_5 := 2^{-24} \min(c_2^{-1}, c_3^{-4}, \epsilon_4)$$

where c_2, c_3, ϵ_4 are the stable constants from Lemmas 5.6, 5.7, and 5.9, respectively.

Proposition 5.10. *Fix $(\eta, J) \in \mathcal{D}(Y)$. Let ϵ_4 and ϵ_5 denote the stable constants from Lemma 5.9 and (31), respectively. There exists a stable constant $c_6 = c_6(\eta, J) \geq 1$ such that the following holds. Let $u : C \rightarrow \mathbb{R} \times Y$ be a compact, connected J -holomorphic curve satisfying the following properties:*

- (L1) $(a \circ u)(\partial C) = \{a_0, a_1\}$ where $a_1 > a_0$;
- (L2) a_0 and a_1 are regular values of the projection $a \circ u : C \rightarrow \mathbb{R}$;
- (L3) $a_1 - a_0 \geq \epsilon_5/8$;
- (L4) $\sup_{\zeta \in C} (a \circ u)(\zeta) - \inf_{\zeta \in C} (a \circ u)(\zeta) \leq \epsilon_5$;
- (L5) $\mathcal{A}(u) \leq 2^{-48} \epsilon_4 \epsilon_5$;
- (L6) The set of all $\zeta \in (a \circ u)^{-1}(a_0)$ such that $|u^* \lambda|_\gamma(\zeta) < 1/2$ has Lebesgue measure at most ϵ_4 in $(a \circ u)^{-1}(a_0)$, where Lebesgue measure is defined using the pullback metric γ .

Then for each $\zeta \in C$, we have the bound

$$\text{Area}_\gamma(S_{\epsilon_5}(\zeta)) \leq c_6(\chi(C)^2 + 1).$$

Proposition 5.10 generalizes an estimate proved by Fish–Hofer (stated in [24, Proposition 4.30]). They made the additional assumption that the domain C is homeomorphic to a compact annulus. The main novelty in the proof of Proposition 5.10 is the introduction of combinatorial and topological arguments to deal with non-annular holomorphic curves. Let us defer the proof for the moment and first prove Proposition 5.1 assuming that Proposition 5.10 holds. The idea is that, assuming the action of u is sufficiently small, any point $\zeta \in C$ is contained inside a surface satisfying the assumptions of Proposition 5.10. In particular, to ensure that (L6) holds, we need the following technical lemma, which asserts that most tangent planes of a low-action holomorphic curve are nearly vertical.

Lemma 5.11. *Write $\mathcal{Q} \subset \mathbb{R}$ for the set of t such that i) t is a regular value of $a \circ u$ and ii) the set of all $\zeta \in (a \circ u)^{-1}(t)$ such that $|\alpha|_\gamma(\zeta) < 1/2$ has Lebesgue measure greater than ϵ_4 . Then the Lebesgue measure of \mathcal{Q} is at most $2\epsilon_4^{-1} \mathcal{A}(u)$.*

Proof. This follows immediately from [24, Lemma 4.27] with parameters $\delta = \epsilon_4$ and $\theta = 1/2$. \square

We now give the proof of Proposition 5.1.

Proof of Proposition 5.1. Let $u : C \rightarrow \mathbb{R} \times Y$ be a standard J -holomorphic curve. Pick any point $\zeta \in C$. We will show that, when $\mathcal{A}(u)$ is sufficiently small, there exists a compact surface $\tilde{C} \subset C$ containing ζ in the middle which satisfies (L1)–(L6) and has Euler characteristic bounded below by $\chi(C)$. Proposition 5.10 then implies the area bound. The proof will take 6 steps.

Step 1: This step defines the surface \widehat{C} . To define the surface, we need to assume that $\mathcal{A}(u) \leq 2^{-48}\epsilon_4\epsilon_5$. Now we choose a positive parameter $r \in (\epsilon_5/8, \epsilon_5/4)$ and set $a_0 := (a \circ u)(\zeta) - r$ and $a_1 := (a \circ u)(\zeta) + r$. We fix r so that i) both a_0 and a_1 are regular values of $a \circ u$ and ii) the set of all $\zeta' \in (a \circ u)^{-1}(a_0)$ such that $|u^*\lambda|_\gamma(\zeta') < 1/2$ has Lebesgue measure at most ϵ_4 . There exists a full measure set of r satisfying i) by Sard's theorem, and a positive measure set of r satisfying ii) by Lemma 5.11 and our assumed bound on $\mathcal{A}(u)$, so an r satisfying both conditions exists.

Set $C_0 := (a \circ u)^{-1}([a_0, a_1]) \subset C$. Write Δ for the disjoint union of compact components of $C \setminus \text{Int}(C_0)$ and set $C_1 := C_0 \cup \Delta$. Let \widehat{C} be the connected component of C_1 containing ζ .

Step 2: This step proves the two-sided bound

$$(32) \quad 2 \geq \chi(\widehat{C}) \geq \chi(C).$$

The upper bound follows from the fact that \widehat{C} is connected. To prove the lower bound, we must show $\chi(\Sigma) \leq 0$ where $\Sigma := C \setminus \text{Int}(\widehat{C})$. Since C is connected and has empty boundary, it follows from the inverse function theorem that any proper embedding of a surface with empty boundary into C is surjective. It follows that each connected component of Σ must have at least one boundary component. Since each connected component of Σ intersects \widehat{C} , it cannot entirely consist of components of C_0 and Δ . It follows that each connected component of Σ is non-compact. We have shown that each connected component of Σ is non-compact and has non-empty boundary, so Σ has non-positive Euler characteristic.

Step 3: Write \widehat{u} for the restriction of u to \widehat{C} . This step verifies that \widehat{u} satisfies (L2) and (L3). Condition (L2) is satisfied because a_0 and a_1 are regular values of $a \circ u$. Condition (L3) is satisfied because $r \in (\epsilon_5/8, \epsilon_5/4)$.

Step 4: Assume that $\mathcal{A}(u) < \hbar(\eta, J, \epsilon_5/8)$ where \hbar denotes the constant from Proposition 5.8. This step verifies that \widehat{u} satisfies (L1) and (L4) given this assumption.

Write

$$L_+ := \sup_{z \in \widehat{C}} (a \circ u)(z) - \sup_{z \in \partial \widehat{C}} (a \circ u)(z), \quad L_- := \inf_{z \in \partial \widehat{C}} (a \circ u)(z) - \inf_{z \in \widehat{C}} (a \circ u)(z).$$

Note that both L_+ and L_- are positive by definition. By the contrapositive of Proposition 5.8, it follows that both L_+ and L_- are at most $\epsilon/8$. Since $(a \circ \widehat{u})(\zeta) = (a_1 + a_0)/2$, we can rearrange the bounds for L_+ and L_- to conclude that

$$\sup_{z \in \partial \widehat{C}} (a \circ \widehat{u})(z) \geq (a_1 + a_0)/2 - \epsilon_5/8 > a_0, \quad \inf_{z \in \partial \widehat{C}} (a \circ \widehat{u})(z) \leq (a_1 + a_0)/2 + \epsilon_5/8 < a_1.$$

Since $(a \circ \widehat{u})(\widehat{C}) \subseteq \{a_0, a_1\}$ by construction, the bounds above are sufficient to conclude (L1). To prove (L4), we use (L1) and the upper bounds for L_+ and L_- :

$$\sup_{z \in \widehat{C}} (a \circ u)(z) - \inf_{z \in \widehat{C}} (a \circ u)(z) = L_+ + L_- + (a_1 - a_0) \leq \epsilon_5/8 + \epsilon_5/8 + \epsilon_5/2 = \epsilon_5.$$

Step 5: This step verifies that \widehat{u} satisfies (L5) and (L6). Both follow from our assumptions and the construction in Step 1.

Step 6: By Proposition 5.10, we have the bound

$$(33) \quad \text{Area}_\gamma(S_{\epsilon_5/16}(\zeta) \cap \widehat{C}) = \text{Area}_\gamma(S_{\epsilon_5/16}(\zeta)) \leq c_0(\chi(\widehat{C})^2 + 1)$$

where $c_0 = c_0(\eta, J) > 0$ is a stable constant. The area bound then follows from (32). \square

It remains to prove Proposition 5.10 and this will take up the remainder of the paper. As the proof gets rather technical, let us start by providing a sketch of the argument.

Outline of the proof of Proposition 5.10.

Let $u : C \rightarrow \mathbb{R} \times Y$ denote a compact, connected J -holomorphic curve satisfying (L1)–(L6). We begin by constructing a *tract decomposition* of u (Proposition 5.12, see also Figure 2). Recall that a *tract* is a compact embedded surface with corners in C , with *horizontal* boundary components mapping into level sets and *vertical* boundary components mapping into gradient flow lines of the function $a \circ u$. We show that, after perturbing u slightly, the domain C can be cut up into tracts with each boundary component having controlled length. The existence of such a decomposition, for annular domains, is itself implicit in the first five steps of the proof of [24, Proposition 4.30], though the explicit statement we formulate here is novel and care is required to get the right statement. Our proof mostly follows these steps, with a new argument to take care of the fact our domain might not be annular; an expository emphasis is also to isolate the key constants to clarify that they are stable. We defer it to the end of the body of the paper.

The next step after construction of the tract decomposition is to bound the area of each tract. We show that the number of vertical and horizontal boundary components are each controlled by $\chi(C)$ (Lemma 5.13) as is the total Euler characteristic of all the tracts (Lemma 5.14). Recalling that each horizontal boundary component has controlled length, we conclude that each tract has a uniform bound on the length of its entire bottom boundary, depending only on $\chi(C)$ and ambient geometry. Applying Proposition 5.5 bounds the area of each tract by a constant depending only on $\chi(C)$ and ambient geometry.

The final step is to cover $S_{\epsilon_5}(\zeta)$ by a controlled number of tracts. This gives a bound on its area since we have already bounded the area of each tract. A geodesic distance argument that we learned from [24] implies that $S_{\epsilon_5}(\zeta)$ cannot intersect both vertical boundary components of a “rectangular” tract, defined (analogously to rectangular strips) below. The topological lemmas mentioned above imply that most tracts are rectangular. These results are combined with a graph-theoretic argument to prove the desired covering bound.

5.5.1. *Statement of tract decomposition.* We now begin the process of making the above outline rigorous. The first part is the statement of the tract decomposition.

Proposition 5.12. *Fix $(\eta, J) \in \mathcal{D}(Y)$. Let $u : C \rightarrow \mathbb{R} \times Y$ be a compact, connected J -holomorphic curve satisfying (L1)–(L6). There exists a fixed constant $\delta_0 = \delta_0(u, \eta, J) > 0$ such that the following holds. For any sufficiently small $\epsilon > 0$, there exists a (δ_0, ϵ) -tame perturbation (u, f) and a finite set of tracts $\{\widetilde{C}_k\}_{k=1}^N$ satisfying the following properties:*




FIGURE 2. A schematic of the tract decomposition. One of the displayed tracts has positive genus and several horizontal/vertical boundary components.

- (a) $C = \bigcup_{k=1}^N \tilde{C}_k$.
- (b) For each $k \neq k'$, the intersection $\tilde{C}_k \cap \tilde{C}_{k'}$ is either empty or equal to a disjoint union of components of $\partial_v \tilde{C}_k$.
- (c) For each k , we have $(a \circ \tilde{u})(\partial_h \tilde{C}_k) = \{a_0, a_1\}$.
- (d) For each k and each component $L \in \pi_0(\partial_h \tilde{C}_k)$ such that $(a \circ \tilde{u})(L) = a_0$, we have $\int_L \tilde{\alpha} < 10\epsilon_4$. Moreover, if L is not a circle, then $\int_L \tilde{\alpha} > \epsilon_4$.
- (e) For each k and each component $L \in \pi_0(\partial_v \tilde{C}_k)$, we have

$$\text{length}_\gamma(L) \leq c_3(a_1 - a_0).$$

The simplest kinds of tracts in the decomposition are *rectangular* tracts, i.e. tracts with zero genus, two horizontal boundary components, and two vertical components. The tract decomposition could, however, contain tracts with positive genus and many horizontal/vertical boundary components. See Figure 2 for a schematic of what the tract decomposition might look like. We defer the proof of Proposition 5.12 for the moment, collecting some useful lemmas about the asserted tract decomposition first.

5.5.2. *Tract topology bounds.* Fix $(\eta, J) \in \mathcal{D}(Y)$ and a compact, connected J -holomorphic curve satisfying (L1)–(L6). Use Proposition 5.12 to construct a perturbation (u, f) and tract decomposition $\{\tilde{C}_k\}_{k=1}^N$. As mentioned above, the tracts \tilde{C}_k could have complicated topology. The next two lemmas provide some a priori topological control. Lemma 5.13 bounds the number of horizontal and vertical boundary components of each tract in terms of $\chi(C)$.

Lemma 5.13. *For each k , we have the bounds*

$$(34) \quad \#\pi_0(\partial_v \tilde{C}_k) \leq 2 - 2\chi(C), \quad \#\pi_0(\partial_h \tilde{C}_k) \leq 4 - 3\chi(C).$$

Proof. Fix any k . Write $M := \#\pi_0(\partial_v \tilde{C}_k)$ for the number of vertical boundary components of \tilde{C}_k . Let \hat{C} denote the closure of $C \setminus \tilde{C}_k$. The proof will take 2 steps.

Step 1: This step proves the first bound in (34). We assume without loss of generality that $M > 2$. If $M \leq 2$, then the desired bound is immediate because by (L1) we have $\chi(C) \leq 0$. It follows from the Mayer–Vietoris sequence that

$$M = \chi(\tilde{C}_k) + \chi(\hat{C}) - \chi(C).$$

We show $\chi(\tilde{C}_k) \leq 1$ and $\chi(\hat{C}) \leq M/2$. The upper bound on $\chi(\tilde{C}_k)$ follows from the fact that $\chi(\tilde{C}_k)$ is connected and has non-empty boundary. The upper bound on $\chi(\hat{C})$ is deduced as follows. Each connected component of \hat{C} is a tract which shares at least at least two vertical boundary components with \tilde{C}_k . Therefore, \hat{C} has at most $M/2$ connected components. Each connected component of \hat{C} has non-empty boundary and therefore has Euler characteristic ≤ 1 . Combine both of these upper bounds with the identity for M and re-arrange to get the first bound in (34).

Step 2: This step proves the second bound in (34). Each connected component of $\partial_h \tilde{C}_k$ is either i) a compact interval which intersects exactly two components of $\partial_v \tilde{C}_k$ or ii) a circle which is a connected component of ∂C . Any component of $\partial_v \tilde{C}_k$ intersects exactly two components of $\partial_h \tilde{C}_k$, so there are M components of the former type, which we showed in Step 1 is bounded above by $2 - 2\chi(C)$. There are at most $\#\pi_0(\partial C)$ components of the latter type. This gives the bound

$$\#\pi_0(\partial_h \tilde{C}_k) \leq 2 + \#\pi_0(\partial C) - 2\chi(C).$$

The second bound in (34) now follows from plugging in the inequality $\#\pi_0(\partial C) \leq 2 - \chi(C)$. \square

The next lemma collects an elementary identity, and some useful related observations, used in the proof of Lemma 5.16 below.

Lemma 5.14. *The Euler characteristic of the domain C satisfies the following identity:*

$$(35) \quad 2\chi(C) = \sum_{k=1}^N (2\chi(\tilde{C}_k) - \#\pi_0(\partial_v \tilde{C}_k)).$$

Moreover, for each k we have $2\chi(\tilde{C}_k) \leq \#\pi_0(\partial_v \tilde{C}_k)$, with equality if and only if either

- (a) $\chi(\tilde{C}_k) = 1$ and $\#\pi_0(\partial_v \tilde{C}_k) = 2$;
- (b) $C = \tilde{C}_k$ and $\chi(\tilde{C}_k) = \#\pi_0(\partial_v \tilde{C}_k) = 0$.

Proof. The proof of the lemma will take 3 steps.

Step 1: This step proves (35). For each k , write $M_k := \#\pi_0(\partial_v \tilde{C}_k)$ for the number of vertical boundary components of \tilde{C}_k . Write M for the number of gradient trajectories of $a \circ \tilde{u}$ which are vertical boundary components of some \tilde{C}_k . It follows that $2M = \sum_k M_k$ because each gradient trajectory is a vertical boundary component of exactly two tracts. The Mayer–Vietoris sequence then implies

$$2\chi(C) = \sum_k 2\chi(C_k) - 2M = \sum_k (2\chi(C_k) - M_k).$$

Step 2: Fix any k . The next two steps prove that $2\chi(\tilde{C}_k) \leq \#\pi_0(\partial_v \tilde{C}_k)$ and characterizes the equality cases. This step provides a proof assuming that $\#\pi_0(\partial_v \tilde{C}_k) = 0$, i.e.

the tract has no vertical boundary components. By Proposition 5.12(b) it follows that \tilde{C}_k is both closed and open in C . Since C is connected, it follows that $C = \tilde{C}_k$. From Proposition 5.12(c), we conclude that $\#\pi_0(\partial\tilde{C}_k) \geq 2$. This implies the bound

$$2\chi(\tilde{C}_k) \leq 0 = \#\pi_0(\partial_v\tilde{C}_k),$$

with equality if and only if $\chi(\tilde{C}_k) = \#\pi_0(\partial_v\tilde{C}_k) = 0$.

Step 3: This step has the same aim as the previous step, but in the case where $\#\pi_0(\partial_v\tilde{C}_k) > 0$. In this case, we must have $\#\pi_0(\partial_v\tilde{C}_k) \geq 2$, and that \tilde{C}_k has non-empty boundary, so we conclude that

$$2\chi(\tilde{C}_k) \leq 2 \leq \#\pi_0(\partial_v\tilde{C}_k).$$

Equality holds if and only if $\chi(\tilde{C}_k) = 1$ and $\#\pi_0(\partial_v\tilde{C}_k) = 2$. □

Remark 5.15. A tract \tilde{C} is *rectangular* if it has zero genus, two horizontal boundary components, and two vertical boundary components. A tract is rectangular if and only if $\chi(\tilde{C}) = 1$ and $\#\pi_0(\partial_v\tilde{C}) = 2$. The primary implication of the identity (35) is that, barring a degenerate case, the tract \tilde{C}_k is rectangular for all but at most $-\chi(C)$ indices k .

5.5.3. *Tract coverings of controlled size.* As above, fix a perturbed J -holomorphic curve (u, f) and a tract decomposition $\{\tilde{C}_k\}_{k=1}^N$. The following lemma asserts that local connected components can be covered by a controlled number of tracts.

Lemma 5.16. *For any point $\zeta \in C$, there exists a covering of the surface $\tilde{S}_{4\epsilon_5}(\zeta) := \tilde{u}^{-1}(B_{4\epsilon_5}(\tilde{u}(\zeta)))$ by $2 - 3\chi(C)$ tracts.*

Note that $\chi(C) \leq 0$ by (L1), so the number $2 - 3\chi(C)$ from the lemma is indeed positive. Recall that a tract \tilde{C} is *rectangular* if it has zero genus, two horizontal boundary components, and two vertical boundary components. The proof of Lemma 5.16 requires the following technical lemma, which asserts that $\tilde{S}_{4\epsilon_5}(\zeta)$ cannot intersect both vertical boundary components of a rectangular tract. The proof of this lemma is similar to the proof of [24, Lemma 4.35].

Lemma 5.17. *For any point $\zeta \in C$ and any k such that \tilde{C}_k is rectangular, the surface $\tilde{S}_{4\epsilon_5}(\zeta)$ does not intersect both vertical boundary components of \tilde{C}_k .*

Proof. Let $\tilde{C} := \tilde{C}_k$ denote any rectangular tract. Then \tilde{C} has two horizontal and two vertical boundary components. Let L denote the bottom horizontal boundary component, defined rigorously as the unique horizontal boundary component contained in $(a \circ \tilde{u})^{-1}(a_0)$. Then L is a compact interval connecting the two vertical boundary components. Write $\partial L = \zeta_+ - \zeta_-$ where ζ_{\pm} are distinct points in \tilde{C} . Write γ_{\pm} for the vertical boundary components intersecting L at ζ_{\pm} respectively. Write

$$d := \inf_{\substack{z_+ \in \gamma_+ \\ z_- \in \gamma_-}} \text{dist}_g(\tilde{u}(z_+), \tilde{u}(z_-))$$

for the extrinsic distance between γ_+ and γ_- . Since both γ_+ and γ_- have length at most $c_3(a_1 - a_0)$ by Proposition 5.12(e), it follows from the triangle inequality that

$$d \geq \text{dist}_g(\tilde{u}(\zeta_+), \tilde{u}(\zeta_-)) - 2c_3(a_1 - a_0).$$

Let $q : [0, T] \rightarrow C$ denote the unique unit-speed parameterization of L such that $q(0) = \zeta_-$ and $q_+ = \zeta_+$. Apply (L6), Proposition 5.12(d), and Lemma 5.9 to the curve $\tilde{q} := \tilde{u} \circ q$ to bound $\text{dist}_g(\tilde{u}(\zeta_+), \tilde{u}(\zeta_-))$ from below. We deduce the bound

$$d \geq \epsilon_4/2 - 2c_3(a_1 - a_0).$$

The right-hand side is seen to be strictly greater than $8\epsilon_5$ using (L4) and the bound $\epsilon_5 \leq 2^{-24} \min(c_3^{-3}, \epsilon_4)$. We conclude that γ_+ and γ_- have extrinsic distance greater than $8\epsilon_5$ from each other. By the triangle inequality, they cannot both intersect $\tilde{S}_{4\epsilon_5}(\zeta)$ for any choice of $\zeta \in C$. \square

Lemma 5.16 is proved by combining Lemmas 5.14 and 5.17 with a combinatorial argument.

Proof of Lemma 5.16. Let $\hat{C}_1, \dots, \hat{C}_D$ denote a minimal-size cover of $\tilde{S}_{4\epsilon_5}(\zeta)$ by tracts. Our goal is to prove the bound $D \leq 2 - 3\chi(C)$. We assume without loss of generality that $D \geq 2$. The proof will take 4 steps.

Step 1: Write Z for the number of indices i such that \hat{C}_i is rectangular. This step observes that $Z \geq D + \chi(C)$. This is a direct consequence of Lemma 5.14 and Remark 5.15. It is important that we assume $D \geq 2$ here, to avoid the degenerate case stated in Lemma 5.14(b).

Step 2: This step constructs a connected graph G as follows. The vertices are $\{1, \dots, D\}$. For any $i \neq j$, we add an edge between them if the tracts \hat{C}_i and \hat{C}_j share a vertical boundary component. The connected surface $\tilde{S}_{4\epsilon_5}(\zeta)$ intersects each of the tracts \hat{C}_i since they form a cover of minimal size. This implies that $\bigcup_{i=1}^D \hat{C}_i$ is connected, which in turn implies that G is connected.

Step 3: Let i be any vertex such that \hat{C}_i is rectangular. This step shows that i is a vertex of degree 1, which is equivalent to the assertion that $\tilde{S}_{4\epsilon_5}(\zeta)$ intersects exactly one vertical boundary component of \hat{C}_i . This assertion follows from applying Lemma 5.17 to the rectangular tract \hat{C}_i and the fact that G is connected.

Step 4: This step uses Lemma 5.14 and some basic graph theory to prove that $D \leq 2 - 3\chi(C)$. For each $i \in \{1, \dots, D\}$, let N_i denote the degree of the vertex i and M_i denote the number of vertical boundary components of \hat{C}_i . Note that $N_i \leq M_i$ for any i . By Step 3, we have $N_i = 1$ and $M_i = 2$ when \hat{C}_i is rectangular, so we get the improved bound $N_i \leq M_i - 1$ in this case. We deduce the inequality

$$(36) \quad 2\chi(C) \leq \sum_{i=1}^D (2\chi(\hat{C}_i) - M_i) \leq 2D - Z - \sum_{i=1}^D N_i.$$

The first inequality uses (35) and the assertion, proved in Lemma 5.14, that every term on its right-hand side is ≤ 0 . The second inequality uses the observed bounds for each M_i above and the bound $\chi(\hat{C}_i) \leq 1$. The last two terms on the right are controlled by D and $\chi(C)$. We proved that $Z \geq D + \chi(C)$ in Step 1. Since G is connected, it

has at least $D - 1$ edges. The sum $\sum_{i=1}^D N_i$ is twice the number of edges of G , so we conclude that $\sum_{i=1}^D N_i \geq 2D - 2$. Plug these bounds into (36) to get an upper bound on D :

$$2\chi(C) \leq 2D - (D + \chi(C)) - (2D - 2) \quad \Rightarrow \quad D \leq 2 - 3\chi(C).$$

□

5.5.4. *Proof of Proposition 5.10.* We can now explain the proof of the technical bound Proposition 5.10, contingent on Proposition 5.12.

Proof of Proposition 5.10. Use Proposition 5.12 to construct a (δ, ϵ) -tame perturbation (u, f) and a tract decomposition $\{\tilde{C}_k\}_{k=1}^N$. Recall that \tilde{u} denotes the perturbed map and that $\tilde{\gamma} = \tilde{u}^*g$ denotes the perturbed metric. The proposition states that $\epsilon > 0$ can be chosen to be arbitrarily small; we will choose ϵ to be smaller than ϵ_5 . This implies that $\text{dist}_g(\tilde{u}(\zeta'), u(\zeta')) < \epsilon_5$ for any $\zeta' \in C$. It follows from this that $S_{\epsilon_5}(\zeta) \subseteq \tilde{S}_{4\epsilon_5}(\zeta)$. Using this observation and Lemma 5.4, it follows that

$$\text{Area}_\gamma(S_{\epsilon_5}(\zeta)) \leq \text{Area}_\gamma(\tilde{S}_{4\epsilon_5}(\zeta)) \leq 2 \text{Area}_{\tilde{\gamma}}(\tilde{S}_{4\epsilon_5}(\zeta)).$$

So, to complete the proof it suffices to bound the area of $\tilde{S}_{4\epsilon_5}(\zeta)$ with respect to the metric $\tilde{\gamma}$. By Lemma 5.13, each tract has at most $4 - 3\chi(C)$ horizontal boundary components, and by Proposition 5.12(d) it follows that

$$\int_{\partial_h \tilde{C}_k \cap (a \circ \tilde{u})^{-1}(a_0)} \tilde{\alpha} \leq 10\epsilon_4(4 - 3\chi(C))$$

for each k . The area bound in Proposition 5.5 shows for each k the bound

$$\text{Area}_\gamma(\tilde{C}_k) \leq 2 \text{Area}_{\tilde{\gamma}}(\tilde{C}_k) \leq c_0(1 - \chi(C))$$

where $c_0(\eta, J) > 0$ is stable and k -independent. By Lemma 5.16, $\tilde{S}_{4\epsilon_5}(\zeta)$ is covered by $2 - 3\chi(C)$ tracts, and the desired area bound follows. □

5.5.5. *Proof of tract decomposition.* To conclude, we need to provide the promised proof of Proposition 5.12, which will take up the remainder of the paper. As we explained in our earlier outline of our arguments, a large part of the proof repeats arguments found in [24], so we only provide sketches for much of this part. On the other hand, many estimates from [24] and many of the assumptions (L1)–(L6) are used in the proof, and it is crucial to keep careful account of the relevant constants, so even in the sketched parts of the proof we are very precise about the estimates and assumptions used.

Proof of Proposition 5.12. The proof of Proposition 5.12 is simple when $\mathcal{A}(u) = 0$, and we begin by explaining this: in this case, the map u is a branched covering map from C onto $[a_0, a_1] \times \gamma$, where γ is a closed orbit of R_η . Cut up γ into intervals $\{\mathcal{I}_\ell\}_{\ell=1}^M$ with length in $(3\epsilon_4, 4\epsilon_4)$ (or leave it be if it is shorter than that) such that the segments $[a_0, a_1] \times \{z\}$ do not intersect a critical value of u for any ℓ and any endpoint $z \in \partial \mathcal{I}_\ell$. For each ℓ , define $\tilde{C}_\ell := u^{-1}([a_0, a_1] \times \mathcal{I}_\ell)$. The set $\{\tilde{C}_\ell\}_{\ell=1}^M$ is the desired tract decomposition.

Thus, we can assume $\mathcal{A}(u) > 0$, which will be a standing assumption for the rest of the proof. Fix $(\eta, J) \in \mathcal{D}(Y)$ and a compact, connected J -holomorphic curve $u : C \rightarrow \mathbb{R} \times Y$

satisfying (L1)–(L6). The proof takes 5 steps. The first four steps closely follow the first five steps in the proof of [24, Proposition 4.30]. The stable constants c_2, c_3, c_4 from Lemmas 5.6, 5.7, 5.9, respectively, and the stable constant ϵ_5 from (31) will appear frequently.

Step 1: This step fixes a (δ, ϵ) -tame perturbation (u, f) where δ and ϵ satisfy suitable bounds. We require δ to be smaller than a stable constant depending on the map u . We also require $\delta \ll c_2 \mathcal{A}(u)$, which is only possible because we are assuming $\mathcal{A}(u) > 0$. We then require ϵ to be smaller than a constant depending on u and δ and smaller than the constant ϵ_5 . The proof that such a perturbation exists is given in Step 1 of the proof of [24, Proposition 4.30].

Step 2: Steps 2–4 will show that for a large measure set of initial conditions $\zeta \in \partial_h^- C$, there exists a solution of the gradient flow equation

$$(37) \quad q : [0, T] \rightarrow C, \quad q'(s) = \text{grad}_{\tilde{z}}(a \circ \tilde{u})(q(s)), \quad q(0) = \zeta$$

terminating on $\partial_h^+ C$. As in [24], we define some relevant sets:

$$(38) \quad \begin{aligned} \mathcal{C} &:= \{\zeta' \in \text{Crit}(u) \mid r(\zeta') = \delta/2\}, \\ \mathcal{D} &:= \{\zeta' \in \text{Crit}(u) \mid r(\zeta') < \delta/2\}. \end{aligned}$$

The set \mathcal{D} is a union of small disks of radius $\delta/2$, centered at the critical points of u , and the set \mathcal{C} is the union of the boundaries of these disks.

Step 2 and its proof in particular follows Step 2 of the proof of [24, Proposition 4.30]. Its goal is to show that solutions to (37) only pass through \mathcal{D} for a small $\tilde{\alpha}$ -measure set of initial conditions. Calling this set of initial conditions $\mathcal{D}^- \subseteq \partial_h^- C$, we can prove the bound

$$(39) \quad \int_{\mathcal{D}^-} \tilde{\alpha} \leq 4c_2 \mathcal{A}(u).$$

Here is an outline of the proof of (39). We note that any gradient flow line starting from a point in \mathcal{D}^- must intersect the set \mathcal{C}' . Then, using the gradient flow, we construct a disjoint union of strips such that i) their top boundaries lie in \mathcal{C} , ii) their bottom boundaries lie in \mathcal{D}^- and iii) the total $\tilde{\alpha}$ -measure of the bottom boundaries is close to that of \mathcal{D}^- . The $\tilde{\alpha}$ -measure of \mathcal{C} is $\lesssim \delta$, which is by Step 1 much less than $c_2 \mathcal{A}(u)$. Then (39) follows from the height bound (L4) and Lemma 5.6.

Step 3: This step follows Step 3 of [24, Proposition 4.30] and its proof. The perturbed height function $a \circ \tilde{u}$ is Morse on $C \setminus \mathcal{D}$. It follows that $a \circ \tilde{u}$ has finitely many critical points in $C \setminus \mathcal{D}$ and each one is non-degenerate. For each $k \in \{0, 1, 2\}$, write \mathcal{M}_k for the set of index- k critical points of $a \circ \tilde{u}$ in $C \setminus \mathcal{D}$. The goal of this step is to show, for each k , that the set of initial conditions in $\partial_h^- C$ whose gradient flow lines limit to a point in $\mathcal{M}_0, \mathcal{M}_1$, or \mathcal{M}_2 is small. It is clear that only a finite set $\mathcal{F}^- \subset \partial_h^- C$ of initial conditions have gradient flow lines limiting to a point in \mathcal{M}_0 or \mathcal{M}_1 , so it remains to control the $\tilde{\alpha}$ -measure of the set of initial conditions whose gradient flow lines limit to a point in \mathcal{M}_2 .

We denote this set by $\mathcal{E}^- \subset \partial_h^- C$ and assert the bound

$$(40) \quad \left| \int_{\mathcal{E}^-} \tilde{\alpha} \right| \leq 4c_2 \mathcal{A}(u).$$

The proof of (40) is similar in style to the proof of (39). By the Morse lemma, there exists for each point $z \in \mathcal{M}_2$ a circle \mathcal{C}_z of arbitrarily small length, such that any gradient flow line limiting to z intersects \mathcal{C}_z exactly once transversely. We choose such circles \mathcal{C}_z such that their total length is at most $c_2 \mathcal{A}(u)$. Using the gradient flow, we construct a disjoint union of strips such that i) their top boundaries lie in \mathcal{C}_z for some $z \in \mathcal{M}_2$, ii) their bottom boundaries lie in \mathcal{E}^- and iii) the total $\tilde{\alpha}$ -measure of the bottom boundaries is close to that of \mathcal{E}^- . Then (40) follows from the height bound (L4) and Lemma 5.6.

Step 4: This step follows Steps 4 and 5 of [24, Proposition 4.30] and their proofs. It shows that for each closed interval $\mathcal{I} \subset \partial_h^- C$ satisfying

$$\int_{\mathcal{I}} \tilde{\alpha} \geq ((a_1 - a_0)^{-1} + 10c_2) \mathcal{A}(u)$$

there exists a solution $q : [0, T] \rightarrow C$ to the equation

$$q'(s) = \text{grad}_{\tilde{\gamma}}(a \circ \tilde{u})(q(s))$$

such that $q(0) \in \mathcal{I}$, $q(T) \in \partial_h^+ C$, and

$$\text{length}_{\tilde{\gamma}}(q([0, T])) \leq c_3(a_1 - a_0).$$

To prove this claim, we define $\mathcal{T} \subset \partial_h^- C$ to be the complement of the set $\mathcal{D}^- \cup \mathcal{E}^- \cup \mathcal{F}^-$. For any point $\zeta \in \mathcal{T}$, there exists a solution q to (37) such that $q(0) = \zeta$ and $q(T) \in \partial_h^+ C$. It follows from (39) and (40) that

$$\int_{\partial_h^- C \setminus \mathcal{T}} \tilde{\alpha} = \int_{\mathcal{D}^- \cup \mathcal{E}^- \cup \mathcal{F}^-} \tilde{\alpha} \leq 8c_2 \mathcal{A}(u).$$

It follows that there exists a finite set $\{\mathcal{T}_k\}_{k=1}^N$ of pairwise disjoint intervals, each contained in \mathcal{T} , such that

$$(41) \quad \sum_{k=1}^N \int_{\mathcal{T}_k} \tilde{\alpha} \geq \int_{\partial_h^- C} \tilde{\alpha} - 10c_2 \mathcal{A}(u).$$

Now write $\mathcal{I}' := \mathcal{I} \cap \cup_{k=1}^N \mathcal{T}_k$. It follows from (41) that

$$\int_{\mathcal{I}'} \tilde{\alpha} \geq \int_{\mathcal{I}} \tilde{\alpha} - 10c_2 \mathcal{A}(u) \geq (a_1 - a_0)^{-1} \mathcal{A}(u).$$

The strips associated to the intervals $\mathcal{T}_k \cap \mathcal{I}'$ satisfy the assumptions of Lemma 5.7. The only nontrivial assumptions to check are (iii) and (iv). Assumption (iii) follows from (L4) and Assumption (iv) follows from the inequality above. Applying Lemma 5.7 produces the desired gradient flow line.

Step 5: This step uses the result of the previous step to complete the proof of the proposition. Unlike the other steps, this step does not have any close counterpart in the proof of [24, Proposition 4.30]. Choose a finite cover $\{\mathcal{L}_j\}_{j=1}^M$ of $\partial_h^- C$ satisfying the following properties:

- (1) Each \mathcal{L}_j is homeomorphic to a closed interval or a circle.
- (2) $\int_{\mathcal{L}_j} \tilde{\alpha} < 5\epsilon_4$.
- (3) If \mathcal{L}_j is homeomorphic to a closed interval, then $\int_{\mathcal{L}_j} \tilde{\alpha} > 3\epsilon_4$.

(4) For any $j \neq j'$, the interiors of \mathcal{L}_j and $\mathcal{L}_{j'}$ (relative to $\partial_h^- C$) are disjoint.

For any j such that \mathcal{L}_j is homeomorphic to an interval, we define a sub-interval $\widehat{\mathcal{L}}_j \subset \mathcal{L}_j$ as follows. Write ζ_j^- and ζ_j^+ for the left and right endpoints of \mathcal{L}_j with respect to the orientation defined by $\tilde{\alpha}$. Fix interior points $\zeta_j^0, \zeta_j^1 \in \mathcal{L}_j$ such that ζ_j^0 is to the left of ζ_j^1 and the following holds. Write $\mathcal{L}_j^0, \mathcal{L}_j^1$, and \mathcal{L}_j^2 for the sub-intervals with oriented boundaries $\zeta_j^0 - \zeta_j^-$, $\zeta_j^1 - \zeta_j^0$, and $\zeta_j^+ - \zeta_j^1$, respectively. Then we require

$$\int_{\mathcal{L}_j^i} \tilde{\alpha} \in (\epsilon_4, 2\epsilon_4)$$

for each $i \in \{0, 1, 2\}$, and set $\widehat{\mathcal{L}}_j := \mathcal{L}_j^1$.

Note that by (L3), (L5), and the bound $c_2 \leq 2^{24}\epsilon_5^{-1}$ (see (31)), we have the bound $((a_1 - a_0)^{-1} + 10c_2) \mathcal{A}(u) \leq \epsilon_4$. Therefore, the interval $\widehat{\mathcal{L}}_j$ satisfies the required length lower bound in Step 4. It follows from Step 4 that for each $j \in \{1, \dots, M\}$ such that \mathcal{L}_j is homeomorphic to an interval, there exists a point $\zeta_j \in \widehat{\mathcal{L}}_j$ and a gradient flow trajectory $q_j : [0, T_j] \rightarrow C$ such that $q_j(0) = \zeta_j$, $q_j(T_j) \in \partial_h^+ C$, and the length of $q_j([0, T_j])$ is at most $c_3(a_1 - a_0)$.

Set

$$\dot{C} := C \setminus \cup_j q_j([0, T_j])$$

and write $\{\dot{C}_k\}_{k=1}^N$ for the connected components of \dot{C} . For each $k \in \{1, \dots, N\}$, the closure \widetilde{C}_k of \dot{C}_k relative to C is a tract. We verify that $\{\widetilde{C}_k\}_{k=1}^N$ satisfy the properties of Proposition 5.12. Proposition 5.12(a-c) are evident from the construction. The upper bound in Proposition 5.12(d) follows from the fact that any $L \in \pi_0(\partial_h^- \widetilde{C}_k)$ is contained in the union of at most two of the sets \mathcal{L}_j . The lower bound in Proposition 5.12(d) follows from the fact that if L is not a circle, then L must contain either \mathcal{L}_j^0 or \mathcal{L}_j^2 for some j . Proposition 5.12(e) follows from the fact that for each k and each component $L' \in \pi_0(\partial_v \widetilde{C}_k)$, there exists some j such that $L' = q_j([0, T_j])$, and therefore L' has length at most $c_3(a_1 - a_0)$. □

APPENDIX A. VERIFYING ASSUMPTIONS

This short appendix collects some elementary arguments verifying that important classes of maps and flows satisfy the assumptions of our main results. In the case of surface maps, we show that Hamiltonian surface diffeomorphisms and rational area-preserving 2-torus diffeomorphisms are monotone. We start with Hamiltonian diffeomorphisms.

Lemma A.1. *Any Hamiltonian diffeomorphism of a closed, oriented surface Σ is monotone.*

Lemma A.1 follows immediately from the next lemma and the easily verified fact that the identity map is monotone.

Lemma A.2. *Let ϕ and ϕ' be a pair of Hamiltonian isotopic area-preserving diffeomorphisms of a closed, oriented surface Σ equipped with an area form ω . Then ϕ' is monotone if and only if ϕ is.*

Proof. Choose a Hamiltonian function $H : \mathbb{R}/\mathbb{Z} \times \Sigma \rightarrow \mathbb{R}$ whose time-one Hamiltonian flow is $\phi^{-1}\phi'$. Write $\{\psi^t\}_{t \in \mathbb{R}}$ for the Hamiltonian flow of H . We use this choice to identify the mapping torii of ϕ and ϕ' :

$$f_H : j_\phi \simeq Y_{\phi'}, \quad (t, p) \mapsto (t, (\psi^t)^{-1}(p)).$$

We compute

$$f_H^* c_1(V_{\phi'}) = c_1(V_\phi), \quad f_H^*[\omega_{\phi'}] = [f_H^* \omega_\phi] = [\omega_\phi + dH \wedge dt] = [\omega_\phi].$$

It follows from this computation that ϕ' is monotone if and only if ϕ is. \square

Next, we show that any rational area-preserving torus diffeomorphism is monotone.

Lemma A.3. *Write $\mathbb{T}^2 := (\mathbb{R}/\mathbb{Z})^2$ and let $\omega := dx \wedge dy$ denote the standard area form. Any area-preserving diffeomorphism $\phi : \mathbb{T}^2 \rightarrow \mathbb{T}^2$ has $c_1(V_\phi) = 0$. Therefore, if ϕ is rational, then it is monotone.*

Proof. Fix any area-preserving diffeomorphism $\phi : \mathbb{T}^2 \rightarrow \mathbb{T}^2$. For any matrix $A \in \mathrm{SL}(2, \mathbb{Z})$, write ϕ_A for the torus map $w \mapsto Aw$. Any area-preserving diffeomorphism is isotopic through area-preserving diffeomorphisms to some ϕ_A . Such an isotopy identifies mapping torii and Chern classes. Therefore, it is sufficient to prove the lemma under the assumption that $\phi = \phi_A$ for some $A \in \mathrm{SL}(2, \mathbb{Z})$.

Write $r(A) := \mathrm{rank}(\ker(A - \mathrm{Id}))$. This is an integer between 0 and 2, inclusive. We give separate proofs that $c_1(V_\phi) = 0$ depending on the value of $r(A)$. If $r(A) = 0$, then it follows from the Mayer–Vietoris sequence that $b_2(Y_\phi) = 1 + r(A) = 1$ and that the second homology group of Y_ϕ is generated by a torus fiber. The class $c_1(V_\phi)$ has zero pairing with a torus fiber, so $c_1(V_\phi) = 0$.

Now, assume $r(A) \geq 1$. Then the matrix A fixes some nonzero vector $v \in \mathbb{R}^2$. Therefore, the differential of $\phi = \phi_A$ fixes the constant vector field v on \mathbb{T}^2 . This vector field defines a non-vanishing section of the vertical tangent bundle V_ϕ . We conclude that $c_1(V_\phi) = 0$. \square

Next, we consider flows on 3-manifolds. We show that the geodesic flow of a Finsler surface is the Reeb flow of a torsion contact form.

Lemma A.4. *Let F be a closed Finsler surface. Then there exists a torsion contact form α on the unit tangent bundle SF whose Reeb vector field generates the geodesic flow.*

Proof. We recall, following [21], how to realize a Finsler geodesic flow as a Reeb flow of a contact form. Let $v : TF \rightarrow \mathbb{R}$ denote the Finsler norm and set $H := v^2/2$. Define a 1-form α on TF by the local coordinate expression $\sum_{i=1}^2 \partial_{q^i} H dp^i$, where p^i and $q^i = \partial_{p^i}$ denote local coordinates on the base and fiber. It is proved in [21, Section 2] that α restricts to a contact form on SF and its Reeb vector field generates the geodesic flow.

Now, we claim that, when we endow it with a complex structure J , the 2-plane bundle $\xi := \ker(\alpha)$ has torsion Chern class. Our proof strategy is to construct a flat J -linear

connection on ξ . This suffices to prove the claim, because the cohomology class of the curvature is equal to the image of $c_1(\xi)$ in $H^2(SF; \mathbb{R})$. Write ℓ for the real line bundle defined as the tangent bundle along the fibers of SF . It follows from the definition of α that ℓ is a sub-bundle of ξ , so ξ splits as an internal direct sum $\ell \oplus J\ell$ of real line sub-bundles.

We observe that the real line bundle ℓ admits a flat connection ∇ ; this connection induces a candidate connection $\bar{\nabla}$ on ξ as follows. Any section s of ξ splits uniquely as a sum $s_0 + Js_1$ where s_0 and s_1 are sections of ℓ . Then, define $\bar{\nabla}s := (\nabla s_0) + J(\nabla s_1)$. The connection $\bar{\nabla}$ is J -linear by definition. It remains to prove that it is flat. Observe that $\bar{\nabla}$ is flat if and only if the identity of operators

$$\bar{\nabla}_V \bar{\nabla}_W - \bar{\nabla}_W \bar{\nabla}_V - \bar{\nabla}_{[V,W]} = 0$$

holds for any pair of vector fields V and W on SF . Testing the left-hand side against any section $s = s_0 + Js_1$, we verify this identity directly:

$$\begin{aligned} \left(\bar{\nabla}_V \bar{\nabla}_W - \bar{\nabla}_W \bar{\nabla}_V - \bar{\nabla}_{[V,W]} \right) \cdot s &= \left(\nabla_V \nabla_W - \nabla_W \nabla_V - \nabla_{[V,W]} \right) \cdot s_0 \\ &\quad + J \cdot \left(\nabla_V \nabla_W - \nabla_W \nabla_V - \nabla_{[V,W]} \right) \cdot s_1 \\ &= 0. \end{aligned}$$

The first equality follows from expanding $s = s_0 + Js_1$ and simplifying. The second equality follows from the fact that ∇ is a flat connection on ℓ . \square

REFERENCES

- [1] S. Allais. On the Hofer-Zehnder conjecture on $\mathbb{C}P^d$ via generating functions. *Internat. J. Math.*, 33(10-11):Paper No. 2250072, 58, 2022.
- [2] D. V. Anosov and A. B. Katok. New examples in smooth ergodic theory. *Ergodic diffeomorphisms. Trudy Moskov. Mat. Obšč.*, 23:3–36, 1970.
- [3] M. Asaoka and K. Irie. A C^∞ closing lemma for Hamiltonian diffeomorphisms of closed surfaces. *Geom. Funct. Anal.*, 26(5):1245–1254, 2016.
- [4] M. Batoréo. On hyperbolic points and periodic orbits of symplectomorphisms. *J. Lond. Math. Soc. (2)*, 91(1):249–265, 2015.
- [5] I. Bendixson. Sur les courbes définies par des équations différentielles. *Acta Math.*, 24(1):1–88, 1901.
- [6] B. Bramham. Periodic approximations of irrational pseudo-rotations using pseudoholomorphic curves. *Ann. of Math. (2)*, 181(3):1033–1086, 2015.
- [7] J. Chaidez and O. Edtmair. 3D convex contact forms and the Ruelle invariant. *Invent. Math.*, 229(1):243–301, 2022.
- [8] E. Cineli, V. L. Ginzburg, and B. Z. Gürel. Topological entropy of Hamiltonian diffeomorphisms: a persistence homology and Floer theory perspective. *arXiv preprint arXiv:2111.03983*, 2021.
- [9] E. Cineli, V. L. Ginzburg, B. Z. Gürel, and M. Mazzucchelli. Invariant sets and hyperbolic closed Reeb orbits. *arXiv preprint arXiv:2309.04576*, 2023.
- [10] E. Cineli, V. L. Ginzburg, B. Z. Gürel, and M. Mazzucchelli. On the barcode entropy of Reeb flows. *arXiv preprint arXiv:2401.01421*, 2024.
- [11] C. C. Conley. Some abstract properties of the set of invariant sets of a flow. *Illinois J. Math.*, 16:663–668, 1972.
- [12] G. Contreras and M. Mazzucchelli. Proof of the C^2 -stability conjecture for geodesic flows of closed surfaces. *To appear in Duke Math. J.*, *arXiv preprint arXiv:2109.10704*.

- [13] D. Cristofaro-Gardiner and R. Hind. Boundaries of open symplectic manifolds and the failure of packing stability. *arXiv preprint arXiv:2307.01140v2*, 2023.
- [14] D. Cristofaro-Gardiner, V. Humilière, and S. Seyfaddini. Proof of the simplicity conjecture. *Ann. of Math. (2)*, 199:181–257, 2024.
- [15] D. Cristofaro-Gardiner, M. Hutchings, U. Hryniewicz, and H. Liu. Proof of Hofer–Wysocki–Zehnder’s two or infinity conjecture. *arXiv preprint arXiv:2310.07636*, 2023.
- [16] D. Cristofaro-Gardiner, M. Hutchings, and D. Pomerleano. Torsion contact forms in three dimensions have two or infinitely many Reeb orbits. *Geom. Topol.*, 23(7):3601–3645, 2019.
- [17] D. Cristofaro-Gardiner, M. Hutchings, and V. G. B. Ramos. The asymptotics of ECH capacities. *Invent. Math.*, 199(1):187–214, 2015.
- [18] D. Cristofaro-Gardiner, D. Pomerleano, R. Prasad, and B. Zhang. A note on the existence of U-cyclic elements in periodic Floer homology. *arXiv preprint arXiv:2110.13844*, 2021.
- [19] D. Cristofaro-Gardiner, R. Prasad, and B. Zhang. Periodic Floer homology and the smooth closing lemma for area-preserving surface diffeomorphisms. *arXiv preprint arXiv:2110.02925*, 2021.
- [20] A. Denjoy. Sur les courbes définies par les équations différentielles à la surface du tore. *J. Math. Pures Appl. (9)*, 11:333–375, 1932.
- [21] M. Dörner, H. Geiges, and K. Zehmisch. Finsler geodesics, periodic Reeb orbits, and open books. *Eur. J. Math.*, 3(4):1058–1075, 2017.
- [22] O. Edtmair and M. Hutchings. PFH spectral invariants and C^∞ closing lemmas. *arXiv preprint arXiv:2110.02463*, 2021.
- [23] J. W. Fish. Target-local Gromov compactness. *Geom. Topol.*, 15(2):765–826, 2011.
- [24] J. W. Fish and H. Hofer. Feral curves and minimal sets. *Ann. of Math. (2)*, 197(2):533–738, 2023.
- [25] T. Fisher and B. Hasselblatt. *Hyperbolic flows*. Zurich Lectures in Advanced Mathematics. EMS Publishing House, Berlin, [2019] ©2019.
- [26] J. Franks. The Conley index and non-existence of minimal homeomorphisms. In *Proceedings of the Conference on Probability, Ergodic Theory, and Analysis (Evanston, IL, 1997)*, volume 43, pages 457–464, 1999.
- [27] V. L. Ginzburg and B. Z. Gürel. Hyperbolic fixed points and periodic orbits of Hamiltonian diffeomorphisms. *Duke Math. J.*, 163(3):565–590, 2014.
- [28] V. L. Ginzburg and B. Z. Gürel. Non-contractible periodic orbits in Hamiltonian dynamics on closed symplectic manifolds. *Compos. Math.*, 152(9):1777–1799, 2016.
- [29] V. L. Ginzburg and B. Z. Gürel. Hamiltonian pseudo-rotations of projective spaces. *Invent. Math.*, 214(3):1081–1130, 2018.
- [30] V. L. Ginzburg, B. Z. Gürel, and M. Mazzucchelli. Barcode entropy of geodesic flows. *To appear in J. Eur. Math. Soc.*, *arXiv preprint arXiv:2212.00943*.
- [31] M. Handel. There are no minimal homeomorphisms of the multipunctured plane. *Ergodic Theory Dynam. Systems*, 12(1):75–83, 1992.
- [32] S. Hozoori. Dynamics and topology of conformally Anosov contact 3-manifolds. *Differential Geometry and its Applications*, 73:101679, 2020.
- [33] M. Hutchings. An index inequality for embedded pseudoholomorphic curves in symplectizations. *J. Eur. Math. Soc. (JEMS)*, 4(4):313–361, 2002.
- [34] M. Hutchings. The embedded contact homology index revisited. In *New perspectives and challenges in symplectic field theory*, volume 49 of *CRM Proc. Lecture Notes*, pages 263–297. Amer. Math. Soc., Providence, RI, 2009.
- [35] M. Hutchings. Lecture notes on embedded contact homology. In *Contact and symplectic topology*, volume 26 of *Bolyai Soc. Math. Stud.*, pages 389–484. János Bolyai Math. Soc., Budapest, 2014.
- [36] M. Hutchings and M. Sullivan. The periodic Floer homology of a Dehn twist. *Algebr. Geom. Topol.*, 5:301–354, 2005.
- [37] M. Hutchings and C. H. Taubes. Gluing pseudoholomorphic curves along branched covered cylinders. I. *J. Symplectic Geom.*, 5(1):43–137, 2007.
- [38] M. Hutchings and C. H. Taubes. Gluing pseudoholomorphic curves along branched covered cylinders. II. *J. Symplectic Geom.*, 7(1):29–133, 2009.

- [39] M. Hutchings and C. H. Taubes. The Weinstein conjecture for stable Hamiltonian structures. *Geom. Topol.*, 13(2):901–941, 2009.
- [40] K. Irie. Dense existence of periodic Reeb orbits and ECH spectral invariants. *J. Mod. Dyn.*, 9:357–363, 2015.
- [41] P. Kronheimer and T. Mrowka. *Monopoles and three-manifolds*, volume 10 of *New Mathematical Monographs*. Cambridge University Press, Cambridge, 2007.
- [42] P. Le Calvez and J.-C. Yoccoz. Un théorème d’indice pour les homéomorphismes du plan au voisinage d’un point fixe. *Ann. of Math. (2)*, 146(2):241–293, 1997.
- [43] R. D. Mauldin, editor. *The Scottish Book*. Birkhäuser/Springer, Cham, second edition, 2015. Mathematics from the Scottish Café with selected problems from the new Scottish Book, including selected papers presented at the Scottish Book Conference held at North Texas University, Denton, TX, May 1979.
- [44] C. T. McMullen. *Renormalization and 3-manifolds which fiber over the circle*, volume 142 of *Annals of Mathematics Studies*. Princeton University Press, Princeton, NJ, 1996.
- [45] M. Meiwes. On the barcode entropy of Lagrangian submanifolds. *arXiv preprint arXiv:2401.07034*, 2024.
- [46] R. Prasad. High-dimensional families of holomorphic curves and the dynamics on three-dimensional energy surfaces. *In preparation*.
- [47] R. Prasad. Invariant probability measures from pseudoholomorphic curves I. *J. Mod. Dyn.*, 19:31–74, 2023.
- [48] R. Prasad. Invariant probability measures from pseudoholomorphic curves II: Pseudoholomorphic curve constructions. *J. Mod. Dyn.*, 19:75–160, 2023.
- [49] J. M. Salazar. Instability property of homeomorphisms on surfaces. *Ergodic Theory Dynam. Systems*, 26(2):539–549, 2006.
- [50] P. Seidel. Symplectic Floer homology and the mapping class group. *Pacific J. Math.*, 206(1):219–229, 2002.
- [51] C. H. Taubes. Embedded contact homology and Seiberg-Witten Floer cohomology I. *Geom. Topol.*, 14(5):2497–2581, 2010.