The smooth closing lemma for area-preserving maps of surfaces

Dan Cristofaro-Gardiner

University of Maryland

Brown University
November 3, 2021

Section 1

The Closing Lemma

Some questions

Question (Smale, Problem 10: "The Closing Lemma", 1998)

Let p be a non-wandering point of a diffeomorphism $S : M \longrightarrow M$ of a compact manifold. Can S be arbitrarily well approximated in C^r by $T : M \longrightarrow M$, so that p is a periodic point of T?

Non-wandering point p: $S^k U \cap U \neq \emptyset$ for each neighborhood U of p. Pugh: true in C^1 topology (1967).

Question (Franks-Le Calvez, '00; Xia: Poincaré '99)

For a generic C^r area-preserving diffeomorphism of a compact surface, is the union of periodic points dense?

Pugh-Robinson ('80s): true in the C^1 topology.

Today's theorem

Theorem (CG., Prasad, Zhang)

A generic element of Diff(Σ, ω) has a dense set of periodic points. More precisely, the set of elements of Diff(Σ, ω) without dense periodic points forms a meager subset in the C^∞-topology.

Definition of meager: countable union of nowhere dense subsets.

Remarks. Let Σ be a closed surface:

- Case $\Sigma = S^2$ previously shown by Asaoka-Irie (2015); more generally for any Hamiltonian diffeomorphism of any Σ.
- Case $\Sigma = T^2$ proved simultaneously to us by Edtmair-Hutchings using related, but different methods; more generally for any Σ when a certain Floer-homological condition holds. We later showed (with Pomerleano) this condition holds generically.

Section 2

Background: the Calabi invariant

More background: Hamiltonian flows

A pair (M^{2n}, ω) with ω a differential 2-form is called a **symplectic manifold** if $d\omega = 0, \omega \wedge \ldots \wedge \omega$ a volume form.

Example: any surface with area form.

Any $H : S^1 \times M^{2n} \longrightarrow \mathbb{R}$ induces a corresponding (possibly time varying) **Hamiltonian vector field** X_{H_t} by the rule

$$\omega(X_{H_t}, \cdot) = dH_t(\cdot).$$

Denote its flow by ψ_H^t.

Definition of the Calabi invariant

Let $Diffeo_c(D^2, dx \wedge dy)$ denote the set of diffeomorphisms

$$f : D^2 \longrightarrow D^2, f^*(dx \wedge dy) = dx \wedge dy, f = id \text{ near } \partial D^2.$$

There is a surjective homomorphism **Calabi**

$$Cal : Diffeo_c(D^2, dx \wedge dy) \longrightarrow \mathbb{R},$$

defined as follows:

- Given $\varphi \in Diffeo_c(D^2, dxdy)$, write $\varphi = \varphi^1_H$, $H = 0$ near ∂D^2.
- Define $Cal(\varphi) := \int_{D^2} \int_{S^1} Hdtdxdy$.
- Fact: $Cal(\varphi)$ doesn't depend on choice of H!

The Calabi invariant

Calabi measures the “average rotation” of the map φ:

$$Cal(\varphi) = \int\int Var_{t=0}^{t=1} Arg(\varphi_H^t(x) - \varphi_H^t(y))dxdy.$$

Section 3

A Weyl law and the idea of the proof

Warm-up case: compactly supported disc maps

We'll first explain the idea in the case of $G := Diffeo_c(D^2, dx \wedge dy)$. We'll define a sequence of maps

$$c_d : Diffeo_c(D^2, dx \wedge dy) \longrightarrow \mathbb{R}$$

with the following properties:

- (*Continuity.*) Each c_d is continuous (e.g. in C^0 topology).
- (*Spectrality.*) For any $\varphi \in G$, $c_d(\varphi)$ is the action of a set of periodic points of φ.
- (*Weyl Law.*) $lim_{d \longrightarrow \infty} \frac{c_d(\varphi)}{d} = Cal(\varphi)$

We can now sketch proof of the key fact: *given U open, nonzero $H \geq 0$ supported in U, $\varphi \circ \psi_H^t$ has a periodic point in U for some $0 \leq t \leq 1$.*

Background: the action

What is the action?

Background: On (S^2, ω), any $H \in C^\infty(S^1 \times S^2)$ has an associated **action functional**

$$\mathcal{A}_H(z,u) = \int_0^1 H(t, z(t))dt + \int_{D^2} u^*\omega$$

defined on **capped loops** (z, u).

- Critical points of H: capped 1-periodic orbits of φ^t_H.
- Critical values of H: called the **action spectrum** $Spec(H)$:, has Lebesgue measure 0.
- Fact: Each $c_d(\varphi^1_H) \in Spec_d(H)$ the **degree** d **action spectrum**, also has measure 0.

More general surfaces

A similar argument works over an arbitrary closed surface Σ. Main challenge: in finding a Weyl law, Calabi homomorphism not in general defined. For example, $Diff(S^2, \omega_{std})$ is a simple group!

Solution: We prove a "relative" Weyl law recovering a "relative" Calabi invariant.

Statement of relative Weyl law: take $\varphi \in Diff(\Sigma, \omega)$, fix $U \subset \Sigma$ open, H compactly supported in U. Then we define c_d analogously to above and show the relative Weyl law:

$$lim_{d \longrightarrow \infty} \frac{c_d(\varphi \circ \psi_H^1) - c_d(\varphi)}{d} = \int_0^1 \int_U H \omega dt.$$

Section 4

The Periodic Floer homology spectral invariants

PFH: the setup

Our proof builds on a great story due to Hutchings, Lee, Taubes.

Let $\varphi \in \mathit{Diffeo}(\Sigma, \omega)$. Recall the **mapping torus**

$$Y_\varphi = \Sigma_x \times [0,1]_t / \sim, \quad (x,1) \sim (\varphi(x), 0).$$

Has a canonical vector field

$$R := \partial_t,$$

a canonical two-form ω_φ induced by ω, and a canonical plane field $\xi = Ker(dt)$.

The definition of PFH

Useful for us to assume **monotonicity equation**:

$$c_1(\xi) + 2PD(\Gamma) = \lambda[\omega_\varphi]$$

for some $\Gamma \in H_1(Y_\varphi), \lambda \in \mathbb{R}$. There's a **degree map** $d : H_1(Y_\varphi) \longrightarrow H_1(S_1) = \mathbb{Z}$, and we also assume $d(\Gamma)$ sufficiently large.

The $\mathbb{Z}_2$ vector space $PFH(\varphi, \Gamma)$ is homology of a chain complex $PFC(\varphi, \Gamma)$, (for nondegenerate φ). Details of $PFC(\varphi, \Gamma)$:

- Freely generated by sets $\{(\alpha_i, m_i)\}$, where
- α_i distinct, embedded closed periodic orbits of R
- m_i positive integer; ($m_i = 1$ if α_i is hyperbolic)
- $\sum m_i[\alpha_i] = \Gamma$.

The differential

- Differential ∂ counts $I = 1$ J-holomorphic curves in $X := \mathbb{R} \times Y_\varphi$, for generic J, where I is the "ECH index". That is:
$$\langle \partial \alpha, \beta \rangle = \#\mathcal{M}_J^{I=1}(\alpha, \beta)$$
- $J : TX \longrightarrow TX, J^2 = -1, \mathbb{R}$-invariant (and admissible)
- ECH index beyond scope of talk; basic idea: $I = 1$ forces curves to be mostly embedded,
- Definition of J-holomorphic curve:
$u : (C, j) \longrightarrow (X, J), \quad du \circ j = J \circ du.$

The differential

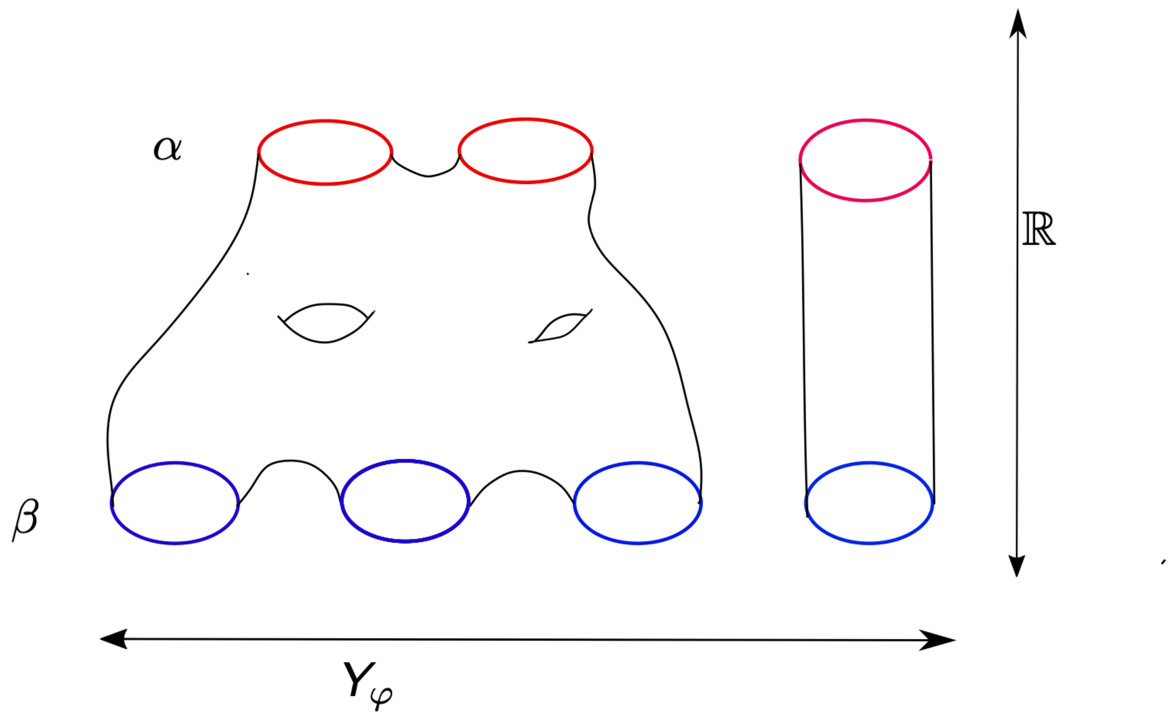

Figure: A J-hol curve contributing to $\langle \partial\alpha, \beta \rangle$.

Example 1: an irrational shift of T^2

Write $T^2 = [0,1]^2/\sim$.

Let $S : T^2 \longrightarrow T^2$ be an irrational shift. This has no periodic points at all! So *PFH* vanishes (other than the empty set).

Example 2: an irrational rotation of S^2

Let φ be an irrational rotation of S^2. This has two fixed points p_+, p_-. One can check $I(C) \in 2\mathbb{Z}$ for any curve C. Conclusion: differential vanishes.

So, degree 1 part generated by p_+, p_-; degree 2 part generated by p_+^2, p_+p_-, p_-^2 etc. $\implies$ Rank $PFH(S^2, d) = d + 1$.

The Lee-Taubes isomorphism

Lee-Taubes showed that there is a canonical isomorphism

$$PFH(\varphi, \Gamma) \cong \widehat{HM}_{c_-}(Y_\varphi, s_\Gamma),$$

where $\widehat{HM}_{c_-}$ is the (negative monotone) Seiberg-Witten Floer cohomology of Y_φ in the spin-c structure s_Γ corresponding to Γ.

This gives a bridge between low-dimensional topology and surface dynamics that is central to our proofs.

Application 1: generic non-vanishing of *PFH*

Theorem (CG., Prasad, Zhang)

Fix a closed surface Σ. Then for C^∞-generic φ, there exists classes $\Gamma_d \in H_1(Y_\varphi)$ with degrees tending to $+\infty$ such that

$$PFH(\Sigma, \varphi, \Gamma_d) \neq 0.$$

Compare with our earlier T^2 example. Upshot: there is a lot of nonzero homology for defining invariants.

Twisted PFH

To get quantitative information, Hutchings' observed one can work with a "twisted" version of PFH; homology of a complex $\widetilde{PFC}(\varphi, \Theta)$.

Details of $\widetilde{PFC}(\varphi, \Theta)$:

- Choose a (trivialized) reference cycle Θ with $[\Theta] = \Gamma$ in H_1.
- Generator of $\widetilde{PFC}(\varphi, d)$ a pair (α, Z), $Z \in H_2(\alpha, \Theta)$
- ∂ counts $I = 1$ curves C from (α, Z) to (β, Z'):
  - this means: C a curve from α to β, with $Z = [C] + [Z']$.

Then $\widetilde{PFH}$ has an action defined by $\mathcal{A}(\alpha, Z) = \int_Z \omega_\varphi$ and for any nonzero $\sigma \in \widetilde{PFH}(\varphi, \Theta)$ we can define $c_\sigma(\varphi)$ to be the minimum action required to represent it. We call this the **spectral invariant** associated to σ.

Section 5

Impressionistic sketch of the Weyl law

The Weyl law and Seiberg-Witten theory

The proof of the Weyl law is beyond the scope of the talk. Very rough idea: the Seiberg-Witten equations are equations for a pair (A, Ψ), where Ψ is a section of s_Γ and A is a spin-c connection.

The configurations with $\Psi = 0$ are called **reducible** and can be described explicitly. In fact, there is a Floer homology for reducibles computable by classical topology.

We define a "Seiberg-Witten" spectral invariant, compute it for the reducibles, and show that it does not change much when compared with PFH via the Lee-Taubes isomorphism.