Contact homology lecture notes [working draft!]

Dan Cristofaro-Gardiner and Alvin Jin

Abstract

These are lecture notes to accompany a graduate level "Morse theory" course taught by Dan Cristofaro-Gardiner in the Fall of 2017. The final part of the class was an introduction to contact homology. These notes and the accompanying exercises are meant to serve as a reference for this material.

The pedagogical idea for this portion of the class was to look at cylindrical contact homology as an analogue of Morse homology for the symplectic action functional; then, discuss the difficulties that come up and some possible fixes, particularly an approach by Hutchings-Nelson in dimension 3; and finally, discuss related homology theories, particularly embedded contact homology. Our discussions involved some substantial input from the theory of J-holomorphic curve theory. In the notes, we do not assume that the reader is familiar with any J-holomorphic curve theory, although we certainly do not give complete proofs of many of the facts about J-holomorphic curves that we will need.

The notes contain some expository material that, to the authors' knowledge, does not appear elsewhere and so might be of broader interest. For example, there is a section on the proof that every Reeb vector field on a closed three-manifold has at least two orbits. While we have tried to proofread everything, please email one of us if you find any mistakes.

1 Introduction

These notes are about contact homology, which is a tool for better understanding contact manifolds. A *contact form* on an oriented $2n + 1$ dimensional manifold is a differential one-form λ satisfying

$$\lambda \wedge d\lambda \wedge \ldots \wedge d\lambda > 0. \tag{1}$$

We can think of this as an odd-dimensional cousin of a symplectic form.

There are two quantities associated to λ that we would like to better understand. The first is

$$\xi := \ker(\lambda), \tag{2}$$

called the *contact structure*. Recall that Frobenius' integrability theorem states that ξ is *integrable* if $d\lambda|_\xi = 0$. The condition (1) guarantees that $d\lambda|_\xi$ is nondegenerate. We therefore colloquially refer to ξ as a "maximally nonintegrable" hyperplane distribution[1]. The second quantity is the *Reeb vector field* R, defined by the equations

$$\lambda(R) = 1 \qquad d\lambda(R, \cdot) = 0.$$

Our two goals are now to define an invariant of contact structures ξ, and to better understand the dynamics of Reeb vector fields. It turns out that these two aims are related.

Our perspective will be heavily influenced by Morse homology. The key insight is that there is a functional

$$\mathcal{A} : C^\infty(S^1, Y) \to \mathbb{R} \tag{3}$$

given by

$$\gamma \to \int_\gamma \lambda$$

called the *symplectic action functional.*

It will turn out that formally:

- The critical points of $\mathcal{A}$ are closed orbits of R, namely smooth maps

$$\gamma : \mathbb{R}/T\mathbb{R} \to Y, \qquad \gamma' = R. \tag{4}$$

- The flow lines of $\mathcal{A}$ are cylinders asymptotic to Reeb orbits, satisfying the *J-holomorphic map equation* (5) defined below.

We will then attempt to define a homology theory from these objects, which we can call a "contact homology".

We now elaborate on the J-holomorphic map equation. We first find a smooth bundle map $J : TX \to TX$ satisfying $J^2 = -1$. This is called an *almost complex structure.* We can now study maps

$$u : (\Sigma, j) \to (X, J)$$

of Riemann surfaces (Σ, j) into X, that satisfy the equation

$$du \circ j = J \circ du. \tag{5}$$

We should think of the equation (5) as asserting that the map u intertwines the almost complex structure on the domain with the almost complex structure on the target. We call a map u satisfying (5) a *J-holomorphic map.* In our case, Σ will always be a (possibly disconnected) closed Riemann surface, minus a finite number of punctures. We will demand that u is asymptotic to a (possibly multiply covered) Reeb orbit γ at each puncture.

We now have to understand more about almost complex structures. Do they even always exist? The point is that $X = \mathbb{R} \times Y$ has a natural symplectic form $\omega = d(e^s\lambda)$, where s is the coordinate on $\mathbb{R}$. We call X the *symplectization* associated to Y.

[1]In fact, to define a contact structure, we do not need (2) to hold for a globally defined one-form satisfying (1); we could instead find a collection of locally defining 1-forms. For various reasons, however, we will not take this approach.

Exercise 1. *Check that ω is a symplectic form on X.*

It turns out that every symplectic manifold admits many such J. In fact, we can assume in addition that:

- J is $\mathbb{R}$-invariant,
- J preserves ξ,
- J takes ∂_s to R,
- The pairing $g(u, v) := \omega(u, Jv)$ is a Riemanninan metric.

We call such a J *compatible* with λ; a central fact is that the space of compatible J is contractible. The idea of studying J-holomorphic curves to construct powerful invariants goes back to Gromov, who used these concepts to prove his celebrated non-squeezing theorem and to prove other striking results. The condition in the fourth bullet point is especially neat: it relates Riemannin, complex, and symplectic geometry.

2 First attempt at a contact homology

We can summarize our hopes from the discussion in §1 as follows: we want to define a chain complex $CC_*(Y, \lambda, J)$. The generators of this chain complex should be closed orbits of the Reeb vector field. The differential d should count J-holomorphic maps for compatible J, with domain a cylinder. The differential d should satisfy $d^2 = 0$, so that the homology $CH_*(Y, \lambda, J)$ is defined.

What kind of invariance properties might we expect? In §1, we saw that the space of compatible J is contractible.

Exercise 2. *Fix a contact structure ξ. Show that the space of λ satisfying*

$$\ker(\lambda) = \xi$$

is contractible.

On the other hand, it turns out that if λ_1 and λ_2 define inequivalent contact structures, then one can not find a path of contact forms between them. Our basic expectation should then be that $CH_*(Y, \lambda, J)$ does not depend on the choice of compatible J, and is an invariant of the associated contact structure ξ. (Maybe by some miracle CH_* should only depend on Y, but we will set this aside for now and return to it later.)

We now attempt to make this precise, and see what kind of issues come up.

2.1 Potential problems

Our first issue is that solutions of (4) are never isolated objects — indeed, it follows from (4) that they come in S^1 families. In contrast, in our treatment of Morse homology we generally wanted a function whose critical points were nondegenerate and hence isolated.

There are several possible ways to deal with this, and we will pick one. Namely, the functional (3) is invariant under the action of $\mathrm{Diff}(S^1)$. We will pass to the quotient, and try to write down the formal Morse homology there instead. This amounts to the following: our chain complex will be generated by closed orbits of (4), modulo reparametrization of the domain; we call these objects *Reeb orbits*. Our differential will count J-holomorphic cylinders modulo reparametrization of the domain as well. To put this in a more general context, call the equivalence class of any J-holomorphic map under reparametrization a *J-holomorphic curve*. When the domain of a J-holomorphic curve is a cylinder, call it a *J-holomorphic cylinder.*

We can now try again to mimic the definition of Morse homology. However, we run into several more problems:

- Fix γ_+, γ_- Reeb orbits, and let $\mathcal{M}_J(\gamma_+, \gamma_-)$ be the space of J-holomorphic cylinders. Ideally, we would like this to be a manifold, of dimension "index(γ_+)-ind(γ_-)." But what is the index of $\gamma_\pm$? We could try to look at a formal Hessian of $\mathcal{A}$, and define the index in analogy with the Morse case. But it turns out this gives something infinite. In addition, the claim that the space $\mathcal{M}_J(\gamma_+, \gamma_-)$ is a manifold can be problematic.

- In Morse homology, flow lines only break along other flow lines. One could imagine a family of J-cylinders breaking into some other J-cylinders. Unfortunately, there are weirder things that can happen. For example, one could also imagine the cylinder breaking into a pair of pants, a plane, and a cylinder.

2.2 The Fredholm index

The first issue in the bullet point above turns out to be not so problematic — what is the analogue of the index of a critical point?

The idea is to look instead at the relative index between two Reeb orbits. To put this in a more general context, let C be a J-holomorphic curve in the symplectization X. Define the *index* of C

$$\mathrm{ind}(C) := (n-3)\chi(C) + 2c_\tau(C) + CZ_\tau^{\mathrm{ind}}(C). \tag{6}$$

Here, $2n$ is the dimension of X, $\chi(C) = 2 - 2g - \#(\text{punctures})$ is the Euler characteristic of the domain of C, $c_\tau(C)$ and $CZ_\tau^{\mathrm{ind}}(C)$ are the relative Chern class and Conley-Zehnder index, defined below, and τ is a symplectic trivialization of ξ over all Reeb orbits. We are assuming that λ is nondegenerate, which we will give a definition for below.

The motivation for (6) is that this is the "expected dimension" of the space of J-holomorphic curves near C. Namely, when C is cut out transversely in a suitable sense, then $\mathcal{M}$ is a manifold near C, of dimension equal to $\mathrm{ind}(C)$. To see that this number is given by the right hand side of (6), we recall the index theory from a previous lecture. We can linearize the J-holomorphic curve equation, and it turns out that the linearized operator is Fredholm. The index of this operator is what is expressed in (6). We can compute it in various ways in order to deduce the expression (6), see for example [2, 10].

A very important point, already alluded to above, is that the curve C might not be transverse. In this case, there is no guarantee that the space of curves is a manifold near C. Nevertheless, one can still compute the index of the linearized operator, and the identity (6) still hods.

We now explain the Conley-Zehnder and Chern class terms in (6). We begin with the Conley-Zehnder term. This is a sum, over (possibly multiply covered) orbits at which C has ends, of Conley-Zehnder terms associated to each orbit. So, let γ be a Reeb orbit, and choose a parametrization as in (4). By linearizing the flow, we get a family of maps $\phi_t : \xi_{\gamma(0)} \to \xi_{\gamma(t)}$, for any positive number t; we can regard this as a family of symplectic matrices once a trivialization is chosen. Letting t range from 0 to T, we then get a family connecting the identity matrix to some other matrix corresponding to a map $P : \xi_{\gamma(0)} \to \xi_{\gamma(0)}$. We call the map P the *linearized return map*, and we say that γ is *nondegenerate* if 1 is not an eigenvalue of this map. When γ is nondegenerate, the Conley-Zehnder index associates an integer to this path of matrices.

We will only give the explicit formulas in the case where $\dim(X) = 4$, which is the main case of interest in our lectures. (For the other cases, see [5].) In this case, the eigenvalues of P are either both real, in which case we call γ *elliptic*, or they both lie on the unit circle, in which case we call γ *hyperbolic*. In the elliptic case, we have

$$CZ_\tau(\gamma) = 2\lfloor\theta\rfloor + 1,$$

where θ is the *monodromy angle*. The monodromy angle is defined as follows: in the elliptic case, one can choose the trivialization τ so that the path of matrices corresponding to ϕ_t is rotation, by numbers $2\pi\theta_t$ chosen so that $\theta_0 = 0$. We then define the monodromy angle to be θ_T.

In the hyperbolic case, we can choose an eigenvector v of ϕ_T, and consider the family of vectors $\phi_t(v)$; this rotates by angle πk as t ranges from 0 to T, and we define

$$CZ_\tau(\gamma) = k.$$

So, in the elliptic case, the local flow is close to a rotation, and the Conley-Zehnder index is determined by the angle of this rotation, and in the hyperbolic case the local flow expands in one direction and contracts in the other. We further call a hyperbolic orbit *positive hyperbolic* if the eigenvalues of the first return map are positive, and we call a hyperbolic orbit *negative hyperbolic* if the eigenvalues are negative.

Now let C be a J-holomorphic curve in X. Then we can divide up the punctures on the domain of C into positive and negative punctures based on whether u is asymptotic to $\mathbb{R} \times \gamma$ as $s \to \infty$ or as $s \to -\infty$. We now define

$$CZ_\tau(C) = \sum_{\gamma \text{a postitive puncture}} CZ_\tau(\gamma) - \sum_{\gamma' \text{a negative puncture}} CZ_\tau(\gamma').$$

What about the Chern class term $c_\tau(C)$?

Let's first do a brief crash course on characteristic classes. Let M be any CW complex, and Z a complex line bundle over M. The first Chern class $c_1(Z) \in H^2(M, \mathbb{Z})$ is the "first obstruction" to finding a nowhere vanishing section of Z. How might we find a nowhere zero section of Z? We first pick a nowhere section s_0 over the 0-cells; there is no obstruction to this. We now want to extend s_0 to a nowhere section s_1 over the 1-cells; there is again no obstruction, since $\mathbb{C} - \{0\}$ is connected. Now we try to extend s_1 over the 2-skeleton, keeping it nowhere zero. Here we do meet a potential obstruction, because $\mathbb{C} - \{0\}$ is not simply connected. We get an obstruction given by the winding number. This obstruction is a map from 2-cells to the integers, and so defines a chain in the second cohomology of M. One can check that in fact it gives a cohomology class, which does not depend on any of the choices required in the construction [4]. This is the first Chern class $c_1(Z)$. When M is a closed two-dimensional manifold, we can evaluate

$$\langle c_1(Z), [M] \rangle, \tag{7}$$

where $[M]$ is the fundamental class of M. This ends up being the same as a signed count of zeros of any section of Z that is transverse to the zero section.

This motivates the definition of the relative Chern class $c_\tau(C)$. We again only do the case where $\dim(X) = 4$. We want to look at the bundle ξ, restricted to C. This is a complex vector bundle, using the almost complex structure J. The curve C is topologically equivalent to a curve with boundary, so in contrast to the closed case below, we want to define a relative class. This is where the trivialization τ comes in. We think of u as a map from a manifold with boundary on some Reeb orbits. Now, we define $c_\tau(C)$ to be a signed count of zeros of a section of $u^*\xi$ extending the section τ and transverse to the zero section.

2.3 Transversality problems

Even though we have resolved our issues about finding an analogue of the Morse index, we still have to worry about whether $\mathcal{M}(\gamma_1, \gamma_2)$ is in fact a manifold.

This is, in fact, problematic.

To clarify the nature of the problem, the following definition is key.

Definition 3. *A holomorphic curve $u : \Sigma \to X$ is* somewhere injective *if there exists $z \in \Sigma$ such that $u^{-1}(u(z)) = \{z\}$ and du_z is injective.*

A nice fact [12] is that for a somewhere injective curve, the somewhere injective points are dense.

Theorem 4. *For generic J, somewhere injective curves are cut out transversely. That is, for generic J, the space of curves near any somewhere injective curve is a manifold, of dimension* $\mathrm{ind}(C)$.

Proof. This is somewhat like our earlier discussions of Sard-Smale. See McDuff-Salamon for the closed case, and Chris Wendl's notes for the general case [12]. □

So, somewhere injective curves work the way we would like. It is also worth noting that we can identify a somewhere injective curve with its image. This is useful, for trying to get a feel for the theory.

Our problems come from the existence of multiply covered curves. That is, once we have a curve $u : \sigma \to X$, we can precompose with a branched covering map of Riemann surfaces to get a map

$$\tilde{u} : \Sigma' \to \Sigma \to X.$$

Exercise 5. *Show that multiply covered curves cannot be always cut out transversely, even for generic J. For example, fill in the following. Let γ be a Reeb orbit (simple) and consider $\mathbb{R} \times \gamma$. This is the "trivial cylinder" over γ. It has a natural parameterization from a twice punctured sphere*

$$u : S^2 - \{0, \infty\} \to X$$

which is somewhere injective.

1. *Show that* $\operatorname{ind}(u) = 0$.
2. *Harder: find branched covers $\tilde{u} : \Sigma' \to S^2 - \{0, \infty\} \to X$ such that* $\operatorname{ind}(\tilde{u}) = 0$ *but $\tilde{u}$ has branched points.*
3. *Why is this a contradiction?*

Remark 6. *In fact, we can come up with examples where u is somewhere injective and $Ind(\tilde{u}) < 0$. If u exists, then $\tilde{u}$ exists and most definitely cannot be cut out transversely.*

The upshot of all this is that we definitely can't have transversality for all curves for generic J.

2.4 Possible fixes

So, we still have transversality problems, and we also have a problem that cylinders need not break into cylinders.

Here are three possible fixes that one could try:

- Find special conditions (usually on λ) to rule out bad multiple covers. For example, assume $Y = S^3$ and assume that $CZ_\tau(\gamma) \geq 3$ for all Reeb orbits γ, for a global trivialization τ. Then it turns out that the naive definition in terms of counting cylinders actually works. This leads to *CCH*, called *cylindrical contact homology.*
- The second option is to give up counting cylinders. Recall that we could have a sequence of cylinders degenerating to something nasty. We could instead think of having a genus zero curve with one positive puncture breaking into many $(g, p_+) = (0, 1)$ curves (but with negative punctures). The idea is that there is a maximum principal for the J-holomorphic map equation, but no minimum principle, so the singularities can only appear downward. We define the differential d as curves with one positive end and an arbitrary number of negative ends and $g = 0$. We get a different invariant $CHA(Y, \lambda)$, called the *contact homology algebra.* The generators are monomials in Reeb orbits. The differential d counts curves and extends to monomials via the Leibniz rule. Transversality is still a major problem, but can be fixed, eg with Pardon's work [9].

- We could try to do something completely different. We could attempt to count everything - have arbitrary numbers of punctures and genus. Notationally, (g, p_+, p_-) could be anything we want. This leads to *symplectic field theory*. This should work, but the foundations still need to be worked out.

- We could try to count similarly to symplectic field theory, but only count some special subset of curves. The key example for our purposes is *embedded contact homology*. For this, when $\dim(Y) = 3$, Hutchings defines a topological index I called the ECH index, and he shows that one can define a homology theory by counting $I = 1$ curves.

3 Cylindrical contact homology for dynamically convex contact forms

We will now explain some recent work by Hutchings-Nelson [7], working out the details of cylindrical contact homology for an important class of contact forms on three-manifolds.

To elaborate, let Y be a closed three-manifold. A contact form λ on Y is called *dynamically convex* if either λ has no contractible Reeb orbits, or $c_1(\xi)|_{\pi_2(Y)} = 0$ and $CZ(\gamma) \geq 3$ for all contractible Reeb orbits γ. Here, $CZ(\gamma) = CZ_\tau(\gamma)$ for any trivialization τ that extends to a trivialization over a disc bounding γ.

Here is a particularly important class of dynamically convex contact forms:

Example 7. *Let $Y = S^3$, then $c_1(\xi) = 0$ since $H^2(S^3) = 0$. To get $CZ(\gamma) \geq 3$ for all γ, take λ to be the restriction of $\frac{1}{2}\sum x_i dy_i - y_i dx_i$ to ∂Z where Z is some convex subset of $\mathbb{R}^4$; the fact that this is a dynamically convex contact form is a famous lemma of Hofer-Wysocki-Zehnder.*

So, convexity implies dynamical convexity. We remark that it is a very interesting question whether or not these are in fact equivalent conditions.

3.1 Definition of CCH

Our goal is now to rigorously define $CCH(Y, \lambda)$, when λ is dynamically convex, following the work of Hutchings-Nelson. The homology $CCH(Y, \lambda)$ is the homology of a chain complex $CCC(Y, \lambda)$. The chain complex $CCC(Y, \lambda)$ is freely generated over $\mathbb{Q}$ by *good* Reeb orbits. A Reeb orbit is called *bad* if it is an even multiple cover of a negative hyperbolic orbit, and good Reeb orbits are precisely those Reeb orbits that are not bad.

Notice that we already have two seeming differences from our naive ideas, based on our experience with Morse homology. Namely, our complex is defined over $\mathbb{Q}$, in contrast to the case of Morse homology (which was defined over $\mathbb{Z}$); also, we have to throw out bad Reeb orbits. There are various ways to see why we have to do this. Later, we explain this from the point of view of gluing.

We now want to define the differential d by counting J-holomorphic cylinders. To make this precise, define

$$\delta\alpha = \sum_{\beta} \sum_{u \in \mathcal{M}_1^J \frac{(\alpha,\beta)}{\mathbb{R}}} \frac{\epsilon(u)}{m(u)} \beta$$

where α, β are good Reeb orbits, and $\mathcal{M}_1^J(\alpha, \beta)/\mathbb{R}$ is the space of J-holomorphic cylinders from α to β with Fredholm index 1, modulo translation in the $\mathbb{R}$ direction. The term $\epsilon(u) \in \{\pm 1\}$ is a sign determined by the orientation, and $m(u)$ is the covering multiplicity of the cylinder.

For a Reeb orbit α, we also want to define $\kappa(\alpha) = m(\alpha)\alpha$ where m is the multiplicity of α. We'll show $\delta\kappa\delta = 0$, by analyzing breakings of ind $= 2$ cylinders. This will imply that $d = \delta\kappa$ is a differential.

3.2 Analyzing the possible breakings

Recall that in Morse homology we analyzed the space of ind $= 2$ flow lines to show that $d^2_{Morse} = 0$. In analogy, we should try to do the same with the space of ind $= 2$ J-holomorphic cylinders.

So, we have to ask ourselves, how can an ind $= 2$ J-holomorphic cylinder break? A powerful theorem, called *SFT compactness* [1], says that we can compactify the space of ind $= 2$ J-holomorphic cylinders by adding cylindrical broken J-holomorphic buildings. A *J-holomorphic building* is a sequence of J-holomorphic curves, $\{u_1, \ldots, u_n\}$ in X, such that the negative asymptotics of each u_i agree with the positive asymptotics of the u_{i+1}. The u_i might not be connected, and we call each u_i a *level*. We call such a building *cylindrical* if the topological gluing given by the bijection between the ends at different levels gives a cylinder.

The upshot of the previous paragraph is that a sequence of cylinders has to break into a cylindrical building. Ideally, we would like to show that this building consists of two levels, each of which are cylinders. However, this is far from obvious. We could certainly imagine very complicated cylindrical buildings, with arbitrarily many levels.

What kind of tools do we have to analyze what sort of buildings are on the boundary of the ind $= 2$ moduli space?

Exercise 8. *The index is additive over gluing.*

Define the index of a building to be the sum of the indices of each level of the building. We now have the following proposition:

Proposition 9 (Hutchings-Nelson)**.** *Let J be generic, let (Y, λ) be dynamically convex, and let $B = (u_1, \ldots, u_n)$ be a cylindrical building. Then:*

- $\mathrm{Ind}(B) \geq 1$.
- *If* $\mathrm{Ind}(B) = 1$*, then B has one level.*
- *If* $\mathrm{Ind}(B) = 2$*, then either:*

– *B has one level.*

– *B has two levels, both cylinders.*

– *B has two levels, $B = (u_1, u_2)$ with u_1 is a branched cover of a trivial cylinder with index 0, and u_2 is a plane union a trivial cylinder.*

We will in a moment comment on the proof of the proposition, but we first discuss its significance. First note that the second bullet point tells us that the space of ind $= 1$ cylinders is compact, which is key for the differential d being defined.

We now comment on the significance of the third bullet point. The first two items of the third bullet point are what we want to hold, to show that $d^2 = 0$. The third possibility, however, is problematic: we could in principle have a sequence of cylinders degenerating into something which can not obviously be seen as contributing to d^2. We call the third possibility the *bad breaking*.

How might we rule out the third?

To simplify the exposition, assume in addition that $CZ(\gamma) = 3$ implies that γ is simple. This is a rather mild assumption, in view of the following:

Exercise 10. *Prove that if (Y, λ) is dynamically convex, and $\pi_1(Y)$ contains no torsion, then $CZ(\gamma) = 3$ implies that γ is simple.*

Now, if there are no contractible orbits, then the bad breaking clearly can not occur. If there are contractible orbits, but $CZ(\gamma) \geq 3$, and $CZ(\gamma) = 3$ implies that γ is embedded, then the only way the bad breaking could occur is if the plane in the lower level is asymptotic to a simple Reeb orbit. Thus, we only have to worry about a very specific kind of degeneration.

We now discuss the proof of the key Proposition 9. For time reasons, we leave this as a (hard) exercise; it is worth attempting to get a feel for the kind of arguments that are possible. Here are the tools needed for the proof. Stare at index formulas, and use the following facts:

- Somwhere injective curves are transverse for generic J.
- In particular, somewhere injective curves have index ≥ 0, with equality only for $\mathbb{R}$ invariant cylinders.
- Any curve that is not somewhere injective is a cover of a somewhere injective curve.
- The *Riemann-Hurwitz* formula: for a $d:1$ covering of Riemann surfaces $\tilde{\Sigma} \to \Sigma$, we have $\chi(\tilde{C}) = d\chi(C) - \#($ ramification points$)$. (Here, the count is a weighted count).

3.3 Adjunction formulas and asymptotic analysis

To completely rule out the breaking, in other words to deal with the final case of the bad breaking where the plane has simple asymptotics, we need some more sophisticated tools.

The first key point is the *adjunction formula.*

The warm-up case is where X is a closed symplectic 4-manifold, instead of the symplectization. Let C be a somewhere injective J-holomorphic curve in X. The "adjunction formula" gives us a way of relating the singularities of C to topological quantities. We specifically have:

$$\langle c_1(TX), C\rangle = \chi(C) + [C]\cdot[C] - 2\delta(C), \tag{8}$$

where $\delta(C)$ is a nonnegative count of singularities of C.

Exercise 11. *Prove* (8) *in the special case where C is in addition immersed, with only nodal singularities. (Note that nodes count as one in the count for δ.)*

In the symplectization case, a similar formula holds, but there are more terms because our situation is topoloically equivalent to a manifold with boundary, and we have to account for the boundary. We namely have:

$$c_\tau(C) = \chi(C) + Q_\tau(C) + w_\tau(C) - 2\delta(C), \tag{9}$$

when C is a somewhere injective curve, asymptotic to Reeb orbits. Here, $w_\tau(C)$ is the *asymptotic writhe* and $Q_\tau(C)$ is the *relative self-intersection*, both defined below.

We start with $Q_\tau(C)$. Recall $[C]\cdot[C]$ in the closed case counts intersections of C with a transversely intersecting curve $\tilde{C}$ in the same homology class. We define:

$$Q_\tau(C) = C\cdot\tilde{C} - \ell_\tau(C,\tilde{C}), \tag{10}$$

where $C, \tilde{C}$ have the same asymptotics, intersect transversely, and have the same relative homology class. The ℓ term is an *asymptotic linking number* term, which we now define.

To define the asymptotic linking term $\ell_\tau(C,\tilde{C})$, we need to spell out more about the asymptotics of C near a Reeb orbit. It turns out [6] that if s is sufficiently large, then $C\cap(Y\times\{s\})$ is a link around some Reeb orbits whose isotopy type does not depend on s. At an embedded Reeb orbit γ, we can use the trivialization τ to think of links near γ as a link in $S^1\times D^2\subset\mathbb{R}^3$, and we can then define the linking number at α, $\ell_\tau(C,\tilde{C},\alpha)$ as the ordinary linking number of links in $\mathbb{R}^3$. We can do an analogous thing for s sufficiently negative, and for $\tilde{C}$. We now define

$$\ell_\tau(C,\tilde{C}) := \sum_{\text{orbits } \alpha \text{ at which } C,\tilde{C} \text{ have positive ends}} \ell_\tau(C,\tilde{C},\alpha) - \sum_{\text{negative ends}} \ell_\tau(C,\tilde{C},\beta).$$

We need to be a little careful with this. It might not be the case that there exists another J-holomorphic curve $\tilde{C}$; however, we can just take $\tilde{C}$ to be some surface intersecting C transversely that is asymptotic to the same Reeb orbits as C, and has $\tilde{C}\cap(Y\times\{\pm s\})$ a link whose isotopy class does not depend on $\pm s$ for sufficiently large or sufficiently negative $\pm s$.

The links arising from the asymptotics of C are actually braid closures, so we sometimes call them braids.

We can define the asymptotic writhe $w_\tau(C)$ analogously: we have

$$w_\tau(C) := \sum_{\text{orbits } \alpha \text{ at which } C \text{ have positive ends}} w_\tau(C,\alpha) - \sum_{\beta \text{ negative ends}} w_\tau(C,\beta),$$

where we are intersecting C with large $|s|$ slices to get links, using the trivialization τ to view these as links in $\mathbb{R}^3$, and computing the writhe of this link in $\mathbb{R}^3$. (Namely, we identify a neighborhood of this link with $A \times I$, where A is an annulus. We project to an annulus and count crossing with signs.)

3.4 Finishing off the bad breaking

We now have all the tools we need to rule out the bad breaking.

Assume that we have a cylinder C in the symplectization X that is close to breaking into the bad breaking (u_1, u_2).

We now chop C into two curves C_1, C_2, by chopping right below the positive asymptotics of u_2. (This means that we fix some large positive value of s, and we define C_1 to be the portion of C close to breaking into the part of u_2 with $\mathbb{R}$-coordinate no more than s, while we define C_2 to be the closure of the complement of C_1 in C.)

Then C_1 and C_2 are both somewhere injective curves in X (because for example u_2 has a somewhere injective component), with boundary. We can then apply the relative adjunction formula (9).

For C_1 we have

$$-1 + w_\tau(\xi_+) - w_\tau(\xi_1 \cup \xi_2) = 2\triangle_+ \geq 0,$$

where ξ_+ is the braid corresponding to the positive asymptotics of C, while $\xi_1 \cup \xi_2$ corresponds to the negative asymptotics of C_1. We can in addition look at the cylindrical component of C_2 to conclude that

$$w_\tau(\xi_1) - w_\tau(\xi_-) = 2\triangle_- \geq 0.$$

We need one more tool from the asymptotic analysis to deal with the writhe terms. This is the *writhe bound.* Assume that we have a braid corresponding to a positive end. Then we have

$$w_\tau(\xi) \geq (d-1)\operatorname{wind}_\tau(\xi),$$

where $\operatorname{wind}_\tau(\xi)$ is the linking number of ξ with the Reeb orbit. We also have

$$\operatorname{wind}_\tau(\xi) \leq \lfloor CZ_\tau(\gamma)^d \rfloor / 2.$$

We can now finish the proof. What is $w_\tau(\xi_1 \cup \xi_2)$?

Exercise 12. $w_\tau(\xi_1 \cup \xi_2) = w_\tau(\xi_1) + w_\tau(\xi_2) + 2d \operatorname{wind}_\tau(\xi_2)$.

Note that in general, instead of the winding number above, one would have the asymptotic linking number. The point of the exercise is that in our situation, the braid ξ_1 is inside the braid ξ_2.

Moreover, we have

Exercise 13. $w_\tau(\xi_2) = 0$.

Putting everything we have said together [this could use a little more detail], we get

$$-1 + w_\tau(\xi_+) - 2d \operatorname{wind}_\tau(\xi_2) - w_\tau(\xi_-) \geq 0.$$

Exercise 14. *Show this is a contradiction using the writhe bounds.*

Thus, the bad breaking is ruled out.

3.5 Automatic transversality

We have so far showed that a sequence of index 2 cylinders can only break into two index 1 cylinders. This is what we were hoping for.

Before moving on, we need to discuss transversality a little more. For generic J, we can guarantee that somewhere injective cylinders are transverse. But what about multiply covered cylinders?

By using ideas going back to Gromov, Hofer, Wendl, and others, we can find criteria guaranteeing that certain curves are "automatically" transverse: in other words, they are transverse, whether or not J is chosen generically.

The specific criteria is:

Proposition 15. *Let C be an immersed J-holomorphic curve in X. Assume that*

$$2g(C) - 2 + h_+(C) < \mathrm{ind}(C), \tag{11}$$

where $h_+(C)$ is the number of ends of C at positive hyperbolic orbits. Then C is cut out transversely.

Exercise 16. *Using this, show that all of the cylinders counted by d are transverse for generic J.*

3.6 Gluing

We are now almost done! We have to this point shown that, assuming dynamical convexity, index 2 cylinders can only break into two index 1 cylinders. To finish the proof, we need to discuss how gluing of J-holomorphic curves works.

So, imagine we have a broken ind $= 2$ cylinder $B = (u_+, u_-)$ with 2 levels. How many ends of the ind $= 2$ moduli space of cylinders are close to breaking along B? In other words, how many index 2 cylinders are close to breaking to our given picture? Define $k = \gcd(m(u_+), m(u_-))$.

Lemma 17. *If γ_0 is good, then there are $\frac{km(\gamma_0)}{m(u_+)m(u_-)}$ ways to glue. If γ_0 is bad, then there are 0 ways to glue.*

Exercise 18. *Assuming Lemma 17, prove that $d^2 = 0$.*

We now discuss why Lemm 17 is true. To glue, we first define a "pregluing" map, which gives a cylinder; we then try to perturb to get an honest cylinder. Assume that the building B breaks along an orbit γ_0. The appearance of the combinatorial factors is partially due to the fact that the ends of $u_\pm$ near γ_0 define $m(\gamma_0)$-fold covers of γ_0. The pregluing requires choosing an isomorphism of these covering spaces, and there are $d(\gamma_0)$ such choices.

More explicitly, to preglue, we first pick a point $p \in \gamma_0$ and parameterizations $\phi_\pm : \mathbb{R} \times S^1 \to \mathbb{R} \times Y$ of $u_\pm$ such that

$$\lim_{t \to \pm\infty} \phi_\pm(t, 1) = p.$$

There are $\frac{d(\gamma_0)}{d(u_\pm)}$ ways to do this up to $\mathbb{R}$-actions on the domain and the target. Since everything is transverse, the gluing theory says that there is a unique way to perturb each pregluing to an honest curve. Thus, there are

$$\frac{d(\gamma_0)^2}{d(u_+)d(u_-)}$$

ways to glue. How many of these are the same? We have an action by the covering group $\mathbb{Z}/d(\gamma_0)$. Two gluings are the same if and only if they are in the same orbit.

What is the action on solutions? We can compute the following:

$$j \in \mathbb{Z}/d(\gamma_0)$$

fixes (ϕ_+, ϕ_-) means that $\frac{d(\gamma_0)}{d(u_+)} | j$ and $\frac{d(\gamma_0)}{d(u_-)} | j$. By definition of k, this means that $\frac{d(\gamma_0)}{k} | j$, which implies each orbit has size $d(\gamma_0)/k$. This implies that, after equivalence, there are $\frac{d(\gamma_0)^2}{d(u_+)d(u_-)} \frac{k}{d(\gamma_0)}$ ways of gluing, which is our number.

In the "bad" case, a similar analysis gives an even number of ways to glue always, and one can check that the signs cancel. (Of course, we haven't had time to explain how to define the signs.)

3.7 Upshot

We can now put all of this together. By Proposition 9 and Exercise 16, the space of index 1 cylinders that we need to consider is a compact manifold, that is 0-dimensional after $\mathbb{R}$ translation. Moreover, for any orbit α, the sum

$$\sum_\beta \#\mathcal{M}_J^{\mathrm{ind}=1}(\alpha, \beta). \tag{12}$$

is a finite sum. This is because of the following. Recall the action functional (3).

Exercise 19.

- *Show that if there exists a J-holomorphic curve C from α to β, then*

$$\mathcal{A}(\alpha) \geq \mathcal{A}(\beta),$$

with equality only if the image of C is an $\mathbb{R}$-invariant cylinder.

- *Conclude that if λ is nondegenerate, then the sum in* (12) *is a finite sum.*

We therefore have a well-defined map d on the CCH chain complex. By Exercise 18, the differential d satisfies $d^2 = 0$. So, the homology $CCH(Y, \lambda, J)$ is indeed well-defined! Why is it independent of J, and λ within the same contact structure? This is beyond the scope of this class: Hutchings-Nelson give a clever argument using a related "nonequivariant" theory.

4 Contact structures on T^3

We now present an application of cylindrical contact homology, We'll prove:

Theorem 20. *There are infinitely many distinct contact structures on T^3.*

Sketch. To get started with the proof, consider $\lambda_n = \cos(nz)\,dx + \sin(nz)\,dy$ for $n \geq 1$.

Exercise 21. *Check that λ_n is a contact form on T^3.*

The question is, are all the $\xi_n = \ker(\lambda_n)$ the same? We call two contact structures ξ_1, ξ_2 over Y contactomorphic if there is a diffeomorphism $f : Y \to Y$ such that $f_*\xi_1 = \xi_2$. So, more specifically, we want to know whether or not the ξ_n are contactomorphic.

Exercise 22. *(A little silly) If ξ_1, ξ_2 are contactomorphic contact structures, $\mathrm{Ker}(\lambda_i) = \xi_i$, and the J_i are λ_i-admissible, then $CCH(Y, \lambda_1, J_1) \cong CCH(Y, \lambda_2, J_2)$. (This is assuming $CCH(Y, \lambda_i, J_i)$ is defined, and only depends on $\mathrm{Ker}(\lambda_i)$.)*

So, we just need to compute $CCH(T^3, \lambda_n)$. What are the Reeb orbits?

Exercise 23. *The Reeb vector field associated to λ_n is $R_n = \cos(nz)\partial_x + \sin(nz)\partial_y$.*

Exercise 24. *There are no contractible orbits for λ_n, and for each $(a, b, 0) \in H_*(T^3, \mathbb{Z})$ such that $(a, b, 0) \neq 0$, there are exactly n S^1 families of orbits in the homology class $(a, b, 0)$.*

Because our orbits come in families, this is analogous to a Morse-Bott situation. We want to perturb λ_n to a dynamically convex contact form that is nondegenerate. There is a standard way to handle this. We perturb λ_n by taking a Morse function on S^1. Specifically, we replace λ_n with $f \cdot \lambda_n$ where $f : Y \to \mathbb{R}_{>0}$ interpolates between the identity function away from neighborhoods of the Reeb orbit families and the pull back of the Morse function on S^1 along the Reeb orbit families. After perturbation, this procedure gives an orbit in class $(a, b, 0)$ for each critical point of our Morse function. If we take the standard function on S^1, we then get two orbits $e_{a,b}$ and $h_{a,b}$, with one elliptic and one hyperbolic.

Dynamical convexity now follows from the fact that there are no contractible orbits.

What are the J-holomorphic curves? Any J-holomorphic cylinder preserves the class of orbits in $H_1(T^3)$. So, if C is a J-holomorphic cylinder from α to β, then $[\alpha], [\beta] \in H_1(T^3)$ are the same. What about cylinders between the $2n$ orbits in class $(a, b, 0)$? We'll only sketch it. Before perturbation, we can analyze all J-cylinders. To do this, recall the action (3), and recall Exercise 37.

Exercise 25. *Each orbit in class $(a, b, 0)$ of any λ_n has action*

$$\mathcal{A} = 2\pi\sqrt{a^2 + b^2}.$$

So, before perturbation, there are no cylinders between distinct orbits in the same homology class either. What the Morse-Bott machinery gives you [needs to be added] is that after perturbation, the only new cylinders correspond to flow lines for the Morse function that we perturbed with. Thus, there are two cylinders for each new orbit, but their signs cancel, as in the case of the standard Morse function on S^1, so that the differential is identically 0.

Thus, we have

$$CCH(T^3, \lambda_n, (a, b, 0)) = \mathbb{Q}^{2n},$$

where $CCH(T^3, \lambda_n, (a, b, 0))$ denotes the homology of the subcomplex generated by Reeb orbits in class $(a, b, 0)$. As expected from the fact that the continuation map should also count (broken) J-holomorphic cylinders, the invariance property [say a little more about this] of CCH implies that $CCH(T^3, \lambda_n, (a, b, 0))$ only depends on λ_n. This proves the theorem.

□

5 Embedded contact homology

5.1 Introduction

In previous sections, we defined cylindrical contact homology, and showed an application. But our definition required (Y, λ) to be dynamically convex. What if we wanted to be able to define a homology theory given any nondegenerate contact form on a 3-manifold?

One could proceed by virtual techniques, as in the work of Pardon [9], to get the Contact Homology Algebra introduced in §2; presumably one could also define Symplectic Field Theory this way, although the details have not been fully worked out. There are other possibilities too that are beyond the scope of this course.

We will describe a different invariant, called *embedded contact homology*. We can define the differential and prove that $d^2 = 0$ without using any virtual techniques. (We can also prove invariance without virtual techniques, but this currently requires Seiberg-Witten theory.)

The idea behind embedded contact homology is that we introduce a new index. Specifically, given a curve C, we define the ECH index:

$$I(C) = c_\tau(C) + Q_\tau(C) + CZ^I_\tau(C),$$

where $c_\tau(C)$ is the relative first Chern class and $Q_\tau(C)$ is the relative self-intersection, both defined previously. Meanwhile, the "total Conley-Zehnder index" $CZ^I_\tau(C)$ is defined by:

$$\sum_i \sum_{j=1}^{m_i} CZ_\tau(\alpha^j) - \sum_i \sum_{j=1}^{n_i} CZ_\tau(\beta^j)$$

where the outer sum is indexed over orbits at which C has positive (respectively negative) ends and the m_i (respectively n_i) is the total multiplicity of all the ends, namely the sum of the multiplicities of all of the ends.

A simple example illustrates the difference between CZ^I and CZ^{ind}.

Example 26. *Let C be a curve with two positive ends, one at an orbit α^2 and the other at an orbit α^3. Let C have negative ends at an orbit β_1^2 and an orbit β_2. Then:*

$$CZ_\tau^I(C) = \sum_{j=1}^{5} CZ_\tau(\alpha^j) - \sum_{j=1}^{2} CZ_\tau(\beta_1^j) - CZ_\tau(\beta_2),$$

while

$$CZ_\tau^{\text{ind}}(C) = CZ_\tau(\alpha^2) + CZ_\tau(\alpha^3) - CZ_\tau(\beta_1^2) - CZ_\tau(\beta_2).$$

Our hope is that we can define a homology theory by counting ECH index 1 curves. This actually works! Note that in contrast to cylindrical contact homology, we are allowing many positive and negative ends, and terms like our CZ^I term only look at the total multiplicities of ends. So, our complex should be generated by "orbit sets", instead of individual orbits.

We now provide the details. Define an *orbit set* to be a finite set $\{(\alpha_i, m_i)\}$, where the α_i are distinct embedded Reeb orbits, and the m_i are positive integers, which should be thought of as covering multiplicities. We sometimes write an orbit set in the multiplicative notation $\alpha = \prod_i \alpha_i^{m_i}$.

Now let $ECC(Y, \lambda, J)$ be the chain complex freely generated over $\mathbb{Z}$ by orbit sets $\alpha = \{(\alpha_i, m_i)\}$, such that in addition each $m_i = 1$ whenever α_i is hyperbolic. (We call such an orbit set *admissible*.) The differential d counts index 1 curves:

$$\langle d\alpha, \beta \rangle = \#\mathcal{M}_J^{I=1}(\alpha, \beta)/\mathbb{R}, \tag{13}$$

where $\mathcal{M}_J^{I=1}(\alpha, \beta)$ is the moduli space of curves from the orbit set α to the orbit set β with $I = 1$. Here, we mod out curves by equivalence *as currents*. We will explain this in a little bit, but to get intuition, note that a connected[2] somewhere injective current is determined by its image, while a connected multiply covered current is determined by its image, together with its covering multiplicity. So, we should think of a J-holomorphic current as a set $\{(C_i, m_i)\}$, where the C_i are distinct somewhere injective curves and the m_i are positive integers.

Remark 27. *The only difference between the "usual" equivalence relation on curves, namely automorphisms of the domain, and equivalence of currents, comes from how we treat multiple covers. For currents, only multiplicity is relevant. Otherwise, we need to keep track of branch points. One reason it is appropriate to look at currents is because the ECH index only depends on the relative homology[3] class $[C] \in H_2(Y, \alpha, \beta)$.*

Hutchings and Taubes showed that $d^2 = 0$ so the homology is well-defined. What goes into the proof? Why is d even well-defined?

5.2 The index inequality

To have any hope that any of this makes sense, we definitely want $\mathcal{M}_J^{I=1}$ to be a 1-manifold. So, we need a relationship between ind and I.

[2]By this, we mean a curve with connected domain, rather than connected image.

[3]This is homology induced by 2-chains with Z with $\partial Z = \alpha - \beta$. [probably say a little more about this in next draft]

Proposition 28 (Index-inequality)**.** *Let C be somewhere injective. Then*

$$\operatorname{ind}(C) \le I(C) - 2\delta(C), \tag{14}$$

where $\delta \ge 0$ is a count of singularities.

Exercise 29. *Prove* (14) *when C is a somewhere injective curve in a closed* 4*-manifold X. In fact, show that equality holds here: we have*

$$I(C) = c(C) + Q(C),$$

where $c(C) = \langle c_1(TX), C\rangle + [C] \cdot [C]$ and $\operatorname{ind}(C) = -\chi(C) + 2\langle c_1(TX), C\rangle$. To do this, use the adjunction formula.

Now assume J is generic, C is somewhere injective and connected, and $I(C) = 1$. Then the left handside of (14) nonnegative, while $I(C) = 1$ and $\delta(C) \ge 0$. This means $\delta(C) = 0$ and that C is embedded. (This is the basic idea behind why the theory is called embedded contact homology). This argument also implies that $\operatorname{ind}(C) = 1$, after using:

Exercise 30. *If the image of C is an $\mathbb{R}$-invariant cylinder, then $I(C) = 0$.*

This all looks quite promising. A connected somewhere injective $I = 1$ curve has to be embedded, and have Fredholm index 1. But what about the multiple covers?

Proposition 31. *Assume that $I(C) = 1$, but not necessarily somewhere injective. Then*

$$C = C_0 \sqcup C_1,$$

where C_1 is embedded with $\operatorname{ind}(C_1) = 1 = I(C_1)$ and C_0 is a union of covers of trivial cylinders.

We leave the proof as a hard but beautiful exercise.

Exercise 32. *Prove Proposition 31. Hint: use $\mathbb{R}$-translation, plus the index inequality in the somewhere injective case.*

Remark 33. *The result of the proposition explains why we are forced to consider curves as currents. Otherwise, multiply covered components would not be rigid modulo translation.*

The upshot of all this is $\mathcal{M}_J^{I=1}(\alpha, \beta)$ is a 1-manifold. We want to count points in it. The natural question is: is $\mathcal{M}_J^{I=1}(\alpha, \beta)$ compact?

We run SFT compactness as before; we have a building B, and we would like to show that it consists of just one nontrivial level.

We have:

Exercise 34. *(easy) The ECH index is additive over levels.*

We also have:

Exercise 35. *With J generic, $I(C) \geq 0$, with equality iff C is a union of $\mathbb{R}$-invariant cylinders. (The proof is similar to the proof of Proposition 31).*

This improves the situation a lot! The only possible breaking has 2 levels. [maybe say a word about multiple levels with branched covers of trivial cylinders in future draft]. One is a branched cover of the trivial cylinder. Further work shows:

Exercise 36 (Hard Exercise)**.** *Show that any multiply covered cylinder on the top or bottom of B has to be trivial.*

Thus, the count in the definition of the ECH differential is indeed finite. The sum (13) is well-defined, similarly to in the contact homology case. The key point is that the action functional extends to orbit sets linearly, and we still have:

Exercise 37. *Let α and β be orbit sets. If $C \in \mathcal{M}(\alpha, \beta)$, then $\mathcal{A}(\alpha) \geq \mathcal{A}(\beta)$ with equality only if $\alpha = \beta$ and C is a cover of an $\mathbb{R}$-invariant cylinder.*

5.3 Sketch of the proof that $d^2 = 0$

We now give a very rough sketch of the proof that $d^2 = 0$. This is a hard result – it is contained in a 200 page paper by Hutchings-Taubes, but we'll try our best to illustrate the key ideas.

Our plan, as always, is to look at the boundary of the ECH index 2 moduli space $\partial \mathcal{M}^{I=2}$, which will give us information about d^2. How can an $I = 2$ curve break? Assume that it broke into three levels, C_1, C_2 and C_3. (If there are more than 3 levels, then it turns out that we can reduce to the three-level case.) Similarly to the proof that the differential d is well-defined, we would need

$$I(C_1) + I(C_2) + I(C_3) = 2.$$

There are three possibilities:

- $I(C_1) = 1 = I(C_3)$ and $I(C_2) = 0$. Thus, C_2 is a branched cover of a trivial cylinder. This is the meat of the 200 page paper.
- $I(C_2) = I(C_3) = 1$ and $I(C_1) = 0$. Thus, C_1 is a branched cover of a trivial cylinder.
- $I(C_1) = I(C_2) = 1$ and $I(C_3) = 0$. Thus, C_3 is a branched cover of a trivial cylinder.

We can rule out the second and third bullet points above similarly to the proof that d is well-defined. (See Exercise 36.) For the first bullet point, we want to do some gluing analysis, similarly as in the case of cylindrical contact homology. We want to preglue C_1, C_2 and C_3 to form an "almost" J-holomorphic curve; in other words, the preglued curve will give a map u into X satisfying

$$du \circ j - J \circ du = \text{error}(p),$$

where error(p) is small. Now, we want to try to perturb, normal to the pregluied object, to get something J-holomorphic. There is an obstruction to always being able to do this, coming from the branch points. For example, consider an example (C_1, C_2, C_3) with one branch point. It cannot possibly be the case that we can always glue, because in this example we have 4 parameters for the gluing, namely 2-parameters for the branch point, and the $\mathbb{R}$-action on C_1 and C_3.

It is reasonable, then, to try fixing the branch point and see if we can glue or not. To make this more precise, define B to be the space of branched covers of a fixed trivial cylinder. Hutchings and Taubes define a bundle $E \to B$ and a section $\psi : B \to E$, called an *obstruction section* of the bundle E, with the property that $\psi^{-1}(0)$ is in bijection with the configurations that can be glued.

In addition, the rank of E is the dimension of B. So,

$$\#\psi^{-1}(0) = e(E),$$

the Euler number of E. Roughly speaking, half of the Hutchings-Taubes paper is dedicated to the gluing analysis sketched above. The remaining hundred pages shows that $e(E) = 1$ when the top and bottom levels have ECH index 1. Warning: there is some subtlety about taking a relative Euler class here, because B is not closed. We will completely gloss over this, but you should take a look at the paper!

5.4 Invariance and Seiberg-Witten

Now, we talk about invariance. In analogy with cylindrical contact homology, we hope that $ECH(Y, \lambda, J)$ only depends on the contact structure $\mathrm{Ker}(\lambda) = \xi$. How can we prove this?

The "usual" approach would be to build continuation maps by counting $I = 0$ curves. This might work, but the technical challenges are substantial. Luckily, there is a kind of shortcut. Namely, Taubes has shown that

$$ECH(Y, \lambda, J) \simeq \widehat{HM}(Y), \tag{15}$$

where $\widehat{HM}(Y)$ is the *Seiberg-Witten Floer cohomology* of Y, defined and studied by Kronheimer-Mrowka [3]. Very importantly, $\widehat{HM}(Y)$ only depends on Y! So, by the isomorphism (15), ECH is in fact not only an invariant of $\mathrm{Ker}(\lambda)$, but an invariant of Y.

A very brief explanation of $\widehat{HM}$ is in order. This is the homology of a chain complex. Roughly speaking, the generators are solutions to the three-dimensional Seiberg-Witten equations; this is a certain system of partial differential equations, originating in physics. The "gauge group" $C^\infty(Y, S^1)$ acts on solutions and we quotient out this equivalence. The differential counts gauge equivalence classes of four-dimensional solutions on $\mathbb{R} \times Y$. The three-dimensional Seiberg-Witten equations are variational — they are formally critical points of the "Cherns-Simons-Dirac" Functional, and one can think of $\widehat{HM}$ as a kind of Morse homology for this functional.

To sketch Taubes' proof, note that the $3d$-Seiberg-Witten equations are equations for a pair (A, ψ) with ψ a section of a bundle S and A a spin^c connection on S. Taubes' proof goes by deforming the equations by adding $r\lambda$ for large r to one side, and rescaling by $\sqrt{r}$. He shows that as $r \to \infty$, a sequence of solutions with an r-independent bound on "energy" must have curvature concentrating along some orbit set α. By using a similar argument concerning the four-dimensional equations, this produces a map from the Seiberg-Witten chain complex to the ECH chain complex. Taubes constructs the inverse to this map by using a related system of equations, called the "vortex" equations. For the (substantial) details of all of this, see [11].

5.5 The Weinstein conjecture

We can now give a quick application of these ideas.

Theorem 38 (Three-dimensional Weinstein conjecture)**.** *Any Reeb vector field on a closed three-manifold has at least one closed orbit.*

Proof. If λ is degenerate, then by definition there must be at least one Reeb orbit, hence the theorem. So, we can assume λ nondegenerate. By work of Kronheimer-Mrowka [3], $\widehat{HM}$ must have infinite rank. The theorem now follows from (15).

□

We remark that in fact Taubes' original proof of the $3d$ Weinstein conjecture did not require the full force of *ECH*. Rather, all that is needed is the procedure for producing an orbit set from a sequence of solutions with uniformly bounded energy that was sketched in the previous section, together with an argument showing that such sequences of solutions exist.

We also remark that it would be very interesting to prove the three-dimensional Weinstein conjecture in full generality without using Seiberg-Witten theory. Since the definition of ECH, and the proof that $d^2 = 0$, do not require Seiberg-Witten theory, ECH might be very useful for this.

6 Further applications

6.1 Two Reeb orbits

In three-dimensions, one can improve on the Weinstein conjecture as follows.

Theorem 39 (CG.-Hutchings)**.** *Any Reeb vector field on a closed 3-manifold has at least two closed orbits.*

We remark that Theorem 39 is in some sense optimal, in that we have seen examples with exactly two Reeb orbits, namely irrational ellipsoids.

6.1.1 The nondegenerate case

We start with a warm-up case. Namely, let λ be nondegenerate. Then $ECH(Y, \lambda)$ is defined. We know that there is at least one orbit. Call it γ. So, assume γ is the only orbit; we will find a contradiction. The idea is that we know that *ECH* must have infinite rank. Thus, γ can not be hyperbolic. In principle, it could be elliptic, but we will rule this out by using grading considerations.

To elaborate, $ECH(Y,\lambda)$ has a grading induced by the ECH index I, such that the differential d decreases the grading by 1. More explicitly, given orbit sets α and β, we can try to define a relative grading gr by declaring

$$\mathrm{gr}(\alpha,\beta) := I(Z),$$

where Z is any element in $H_2(Y,\alpha,\beta)$. However, this will depend on the choice of Z unless we take values in $\mathbb{Z}$, modulo a suitably chosen integer. The following exercise clarifies this:

Exercise 40. *If $Z, Z' \in H_2(Y,\alpha,\beta)$ then*

$$I(Z) - I(Z') = \langle Z - Z', c_1(\xi) + 2\,\mathrm{PD}(\Gamma), \tag{16}$$

whenever $[\alpha] = [\beta] = \Gamma \in H_1(Y)$.

Thus, we should fix the class γ of our orbit sets, and declare the relative grading as taking values in $\mathbb{Z}/\ell$ where ℓ is the divisibility of $c_1(\xi) + 2\,\mathrm{PD}(\Gamma)$ in $H^2/\text{torsion}$. To package this, we introduce $ECH(Y,\lambda,\Gamma)$ to be the homology of the subcomplex generated by orbit sets α with homology class $[\alpha] = \Gamma$, so that this admits a $\mathbb{Z}/\ell$.

The grading is particularly simple given:

Exercise 41. *Given any ξ, there always exists $\Gamma \in H_1(Y)$ such that $c_1(\xi) + 2\,\mathrm{PD}(\Gamma) = 0$.*

To proceed with the proof in the nondegenerate case, choose Γ such that $c_1(\xi) + 2\,\mathrm{PD}(\Gamma) = 0$. Then the relative grading gr takes values in $\mathbb{Z}$. Also, by work of Kronheimer-Mrowka, the group $ECH(Y,\lambda,\Gamma)$ has infinite rank. Now arbitrarily normalize the relative grading gr to an absolute grading $|\cdot|$ by declaring some orbit set to have grading 0.

We will now do some thinking as the grading goes to infinity.

Exercise 42. *Show that whenever $|\gamma^d|$ is defined, then $|\gamma^d| >> d$. That is*

$$\frac{|\gamma^d|}{d} \to \infty.$$

To get some intution for Exercise 42, you might start with CZ^I. This looks approximately like $\sum_{i=1}^d i\theta$, where θ is the monodromy angle for λ; this grows quadratically in d.

On the other hand, $\widehat{HM}(Y,\Gamma)$ is in fact 2-periodic in sufficiently high grading, with nonzero rank, by another calculation of Kronheimer-Mrowka. This completes the proof in the nondegenerate case.

Exercise 43. *Show using similar arguments that if λ is nondegenerate and there are exactly two orbits, then both are elliptic.*

Remark 44. *We have introduced the splitting*

$$ECH(Y,\lambda) = \bigoplus_{\Gamma \in H_1(Y)} ECH(Y,\lambda,\Gamma).$$

This is useful, and corresponds to the splitting $\widehat{HM}(Y) = \bigoplus_{\Gamma\in H_1(Y)} \widehat{HM}(Y,\Gamma)$, which is the Seiberg-Witten Floer homology in the spin^c structure $\mathfrak{s}_\xi + PD(\Gamma)$, where $\mathfrak{s}_\xi$ is a certain spin^c structure determined by ξ.

6.1.2 The degenerate case

Now, assume λ is degenerate. We obviously need to argue somewhat differently from the previous section, because $ECH(Y, \lambda, \Gamma)$ is not defined. Why can't we have exactly one orbit which is degenerate? Nondegeneracy is a generic condition, so we could try to perturb to the nondegenerate case, but we need the right tools. For example, why couldn't you have exactly two nondegenerate orbits colliding into a single degenerate one as we turn off the perturbation from the degenerate case?

To handle all of this, recall the symplectic action functional (3). We will define numbers $c_\sigma(\lambda)$ for any (possibly degenerate) contact form λ, given a nonzero class $\sigma \in ECH(Y)$. The numbers are the "symplectic action required to represent the class σ."

They are rigorously defined as follows. First assume λ is nondegenerate. Define $ECH^L(Y, \lambda, \Gamma)$ to be the homology of the subcomplex generated by orbit sets α with $\mathcal{A}(\alpha) \leq L$, where $L \geq 0$ is a real number. There is an inclusion induced map:

$$ECH^L(Y, \lambda, \Gamma) \to ECH(Y, \xi, \Gamma),$$

where $ECH(Y, \xi, \Gamma)$ is computed by any nondegenerate contact form giving ξ. We can now define

$$c_\sigma(\lambda) := \min\{L : \sigma \text{ in image of } ECH^L(Y, \lambda, \Gamma) \to ECH(Y, \xi, \Gamma)\}.$$

When λ is degenerate, define

$$c_\sigma(\lambda) := \lim_{n\to\infty} c_\sigma(\lambda_n)$$

where $\lambda_n \to \lambda$ in C^0.

Remark 45. *The numbers $c_\sigma(\lambda)$ are called "spectral invariants." This idea is used often in Floer homology. It is surprising that only C^0 convergence is needed.*

To use these $c_\sigma(\lambda)$, we need another asymptotic formula. As before, we choose $\Gamma \in H_1(Y)$ such that $c_1(\xi) + 2PD(\Gamma) = 0$, to get a grading $|\cdot|$.

Theorem 46 (CG, Hutchings, Ramos)**.** *Let $\{\sigma_n\}$ be a sequence of nonzero classes in $ECH(Y, \xi, \Gamma)$, with definite gradings tending to $+\infty$. As $n \to +\infty$,*

$$\frac{c^2_{\sigma_n}(\lambda)}{|\sigma_n|} \to \int_Y \lambda \wedge d\lambda. \tag{17}$$

Remark 47. *We should think of $\int_Y \lambda \wedge d\lambda$ is the volume of Y. The asymptotic formula* (17) *is called the* volume identity.

We can now finish the proof modulo some facts. Recall $\widehat{HM}$ is 2-periodic in sufficiently high grading. This implies by (15) that the same holds for ECH.

Now let $\{\sigma_n\}$ be a sequence of nonzero classes in ECH with $|\sigma_n| = |\sigma_{n-1}| + 2$. We have some facts:

- Any $c_\sigma(\lambda)$ is the action of some orbit set.

- We can choose the σ_n so that $c_{\sigma_n}(\lambda) < c_{\sigma_{n+1}}(\lambda)$.

Assuming the two bullet points above, we now finish the proof. Assume there is exactly one orbit γ, of action T. Then, each $c_{\sigma_n}(\lambda)$ is a multiple of T, with c_{σ_n} increasing with n. Thus,

$$\lim_{n\to\infty} \frac{c_{\sigma_n}^2}{|\sigma_n|} = \infty,$$

which contradicts the volume axiom (17).

Why are the bullet points true? The first essentially holds by definition. The second is harder, and requires explaining one additional structure on *ECH*. Namely, the 2-periodicity result about $\widehat{HM}$ cited above can be improved. There is a degree -2 map U on $\widehat{HM}$, and Kronheimer-Mrowka show that the U map is an isomorphism in sufficiently high degree. Under the isomorphism (15), the U map agrees with a degree -2 map on *ECH*, which we also call U. The ECH U-map is simple to explain: it counts $I = 2$ curves, through a marked point away from all Reeb orbits.

In the nondegenerate case, then, the second bullet follows by taking a sequence of nonzero classes σ_n with

$$U\sigma_n = \sigma_{n-1}.$$

One can now argue as in Exercise 37. Most of this also holds in the degenerate case, but now we have to worry about a sequence of J-holomorphic curves, all going through a marked point, with positive energies, but with energies going to 0. One can rule this out by a compactness argument; essentially, such a collection of curves would have to be converging to an $\mathbb{R}$-invariant cylinder, but this cylinder would have to pass through a marked point away from all Reeb orbits.

6.2 Symplectic embeddings

We now illustrate a different sort of application.

Given symplectic manifolds (X_1, ω_1), (X_2, ω_2), it turns out to often be quite interesting to study whether or not there exists a symplectic embedding

$$(X_1, \omega_1) \to (X_1, \omega_2).$$

The problem is particularly subtle when the dimensions of X_1 and X_2 are close together; when the dimension of X_2 is somewhat larger than the dimension of X_1, versions of Gromov's h-principle often apply. When

$$\dim(X_1) = \dim(X_2) = 4$$

we can sometimes use embedded contact homology to study this question by using ECH to define obstructions to symplectic embeddings. We will now explain how this works in an illustrative situation, namely where each of the X_i are irrational ellipsoids $E(a_i, b_i)$. Recall that the boundary of these ellipsoids have a standard nondegenerate contact form, given by restricting the standard contact form λ_{std} on $\mathbb{R}^4$.

6.2.1 Cobordism maps and symplectic embedding obstructions

The basic idea is as follows. Assume that there exists a symplectic embedding

$$\Psi : \overline{E(a,b)} \to E(c,d).$$

We can remove the interior of $\Psi(\overline{E(a,b)})$ to get a "symplectic cobordism" from $\partial E(c,d)$ to $\partial E(a,b)$. This will induce a map

$$\tilde{\Psi} : ECH(\partial E(c,d)) \to ECH(\partial E(a,b)). \tag{18}$$

The reason that we expect the map (18) is because we expect ECH to be a "field theory", similarly to the Symplectic Field Theory introduced in a previous lecture. We can now use the map (18) to find obstructions, because of the following two key points:

Proposition 48.

- *Normalize the relative grading* gr *to an absolute grading* $|\cdot|$ *by declaring the empty set of Reeb orbits to have grading* 0*. Then the map* $\tilde{\Psi}$ *preserves the grading* $|\cdot|$.
- *If* σ_1 *and* σ_2 *are nonzero classes with* $\tilde{\Psi}(\sigma_1) = \sigma_2$*, then*

$$c_{\sigma_2}(\lambda_{std}|_{\partial E(c,d)}) < c_{\sigma_1}(\lambda_{std}|_{\partial E(a,b)}).$$

To show that the map (18) exists, and to prove Proposition 48, we again use Seiberg-Witten theory. Kronheimer-Mrowka show that Seiberg-Witten Floer homology admits cobordism maps, so we can define (18) by using Taubes' isomorphism (15). By using ideas related to the ideas needed to prove (15), Hutchings-Taubes show that the map $\tilde{\Psi}$ satisfies a *holomorphic curve axiom*: namely, the map $\tilde{\Psi}$ is induced by a chain map Φ, with the property that if $\langle \Phi(\alpha), \beta \rangle$ is nonzero for orbit sets α and β, then there is a possibly broken J-holomorphic building B from α to β. The proof of the second bullet point now follows as in Exercise 37. The proof of the first bullet point also uses the holomorphic curve axiom. By comparing the ECH index formula with the formula for the expected dimension of the Seiberg-Witten moduli space, one can show that the J-holomorphic building B has ECH index 0, because the Seiberg-Witten cobordism map counts Fredholm index 0 monopoles.

6.2.2 The ECH of an irrational ellipsoid

We now compute the ECH of the boundary of an irrational ellipsoid $E(a,b)$.

Exercise 49.

- *Show that the Reeb vector field for* $\lambda_{std}|_{\partial E(a,b)}$ *is given in polar coordinates by*

$$R = \frac{2\pi}{a}\partial_{\theta_1} + \frac{2\pi}{b}\partial_{\theta_2}.$$

Conclude that if b/a *is irrational, then there are exactly two Reeb orbits,*

$$\gamma_1 = \{z_2 = 0\}, \qquad \gamma_2 = \{z_1 = 0\}.$$

- *Show that* γ_1 *and* γ_2 *are both elliptic.*

By the index ambiguity formula (16), the ECH index of any relative homology class $Z \in H_2(\partial E(a,b), \alpha, \beta)$ only depend on α and β. The quantities

$$c_\tau(\gamma_1^m \gamma_2^n) := c_\tau(Z), \quad Q_\tau(\gamma_1^m \gamma_2^n) := Q_\tau(Z),$$

where Z is any element in $H_2(\partial E(a,b), \alpha, \emptyset)$ thus do not depend on Z. We have

$$|\gamma_1^m \gamma_2^n| = c_\tau(\gamma_1^m \gamma_2^n) + Q_\tau(\gamma_1^m \gamma_2^n) + CZ_\tau^I(\gamma_1^m \gamma_2^n). \tag{19}$$

Exercise 50.

- *Show that $\mathcal{A}(\gamma_1) = a$ and $\mathcal{A}(\gamma_2) = b$.*
- *Show that under the identification*

$$T\mathbb{R}^4 = \mathbb{C} \oplus \mathbb{C},$$

the restriction of ξ to γ_1 agrees with the second summand, and the restriction of ξ to γ_2 agrees with the first.

- *Use the previous bullet point to define a trivialization τ, with the property that*

$$c_\tau(\gamma_1^m \gamma_2^n) = m + n, \quad Q_\tau(\gamma_1^m \gamma_2^n) = 2mn, \quad \theta_1 = a/b, \quad \theta_2 = b/a,$$

where the θ_i are the monodromy angles for γ_i.

We can now plug in the calculations from this exercise into (19). We get

$$|\gamma_1^m \gamma_2^n| = m + n + 2mn + \sum_{i=1}^{m} (2\lfloor i(a/b) \rfloor + 1) + \sum_{i=1}^{n} (2\lfloor i(b/a) \rfloor + 1). \tag{20}$$

It follows that $|\gamma_1^m \gamma_2^n|$ is always even, hence $d = 0$ identically. In fact, there is a beautiful interpretation of the right hand side of (20).

Exercise 51. *Show that $|\gamma_1^m \gamma_2^n|$ is twice the number of lattice points in the triangle in the first quadrant bounded by the axes and the line through the point (m,n) with slope $-b/a$.*

We can now complete our computation. We have:

$$ECH_*(\partial E(a,b)) = \mathbb{Z}, \text{ if }^* \text{ is an even nonnegative integer}; \quad 0 \text{ otherwise} \tag{21}$$

To see this, we just imagine moving the line with slope $-b/a$ away from the origin. Since b/a is irrational, this line hits lattice points one at a time, eventually hitting any lattice point in the first quadrant exactly once. Hence, (21) follows from (20).

6.2.3 Putting it all together

We now put all this together to show how to find obstructions to ellipsoid examples. To illustrate the idea, we will first show:

Proposition 52. *There does not exist a symplectic embedding*

$$E(1,2) \to B^4(2-c), \tag{22}$$

whenever $c > 0$.

Proof. Assume that such an embedding exists. Then there exists $\epsilon_1 > 0$ and $\epsilon_2 > 0$ such that there is an embedding

$$E(1, 2+\epsilon_1) \to E(2-c, 2-c+\epsilon_2),$$

with $(2+\epsilon_1)$ and $(2-c+\epsilon_2)/(2-c)$ irrational, and $2-c+\epsilon_2 < 2$. Then, by the procedure in §6.2.1, there exists a map

$$\phi : ECH(\partial E(2-c, 2-c+\epsilon_2)) \to ECH(\partial E(1, 2+\epsilon_1))$$

satisfying the conclusions of Proposition 48. Let $\gamma_1^{\pm}$ and $\gamma_2^{\pm}$ be the Reeb orbits for $\partial E(2-c, 2-c+\epsilon_2)$ and $\partial E(1, 2+\epsilon_1)$; the notation will be such that the plus orbits correspond to $E(2-c, 2-c+\epsilon_2)$, the negative orbits correspond to $\partial E(1, 2+\epsilon_1)$, and $\gamma_1^{\pm}$ is always shorter than $\gamma_2^{\pm}$.

Now consider the orbit set for each irrational ellipsoid with grading given by 4. It follows from (20) that for $\partial E(2-c, 2-c+\epsilon_2)$ this is given by γ_2^+ and for $\partial E(1, 2+\epsilon_1)$ this is given by $(\gamma_1^-)^2$. Then by Proposition 48, we would have to have

$$\mathcal{A}((\gamma_1^-)^2) < \mathcal{A}(\gamma_2^+).$$

On the other hand, we compute that

$$\mathcal{A}((\gamma_1^-)^2) = 2 \quad \mathcal{A}(\gamma_2^+) = 2-c+\epsilon_2,$$

which is a contradiction. □

Thus, we get a very strong obstruction to embeddings of $E(1,2)$ into a ball by looking at the portion of ECH in grading 4. (Notice that if $c \le 0$, then the embedding (22) always exists by inclusion.) If we want to get embedding obstructions for other ellipsoids, we can of course try to look at all the different graded pieces of ECH. To make this precise, define the sequence $N(a,b)$ whose k^{th} element, indexed starting at $k=0$, is the $(k+1)^{st}$ smallest element in the matrix

$$(ma+nb)_{(m,n)\in\mathbb{Z}_{\ge 0}\times\mathbb{Z}_{\ge 0}},$$

and let $N(a,b)_k$ denote the k^{th} element.

Exercise 53. *Assume b/a and c/d be irrational. Show that if there exists a symplectic embedding*

$$E(a,b) \to E(c,d),$$

then we must have

$$(N(a,b))_k \le N(c,d)_k$$

for all k.

In fact, it is shown by McDuff [8] that the sequences $N(a,b)$ and $N(c,d)$ completely characterize whether or not a symplectic embedding exists: if $N(a,b)_k \leq N(c,d)_k$ for all k, then McDuff shows that there is an embedding $E(a,b) \to E(c,d)$.

The numbers $N(a,b)$ intersect with a very interesting branch of symplectic geometry called *symplectic capacity* theory. One can use similar ideas to define a sequence of numerical invariants $c_k((X,\omega))$ associated to any symplectic 4-manifold, called *ECH capacities*. The numbers $N(a,b)$ are precisely the ECH capacities of the ellipsoid $E(a,b)$.

References

[1] F. Bourgeois, Y. Eliashberg, H. Hofer, K. Wysocki, and E. Zehnder, *Compactness results in symplectic field theory*, Geom. Top. 7 (2003) 799-888.

[2] D. Dragnev, *Fredholm theory and transversality for noncompact pseudoholomorphic maps in symplectizations*, CPAM 57 (2004), 726 - 763.

[3] P. Kronheimer and T. Mrowka, *Monopoles and three-manifolds*, Cambridge University Press, 2008.

[4] J, Milnor and J. Stasheff, *Characteristic classes*, Princeton University Press, 1974.

[5] J. Gutt, *The Conley-Zehnder index for a path of symplectic matrices*, arXiv:1201.3728.

[6] M. Hutchings, *Lecture notes on embedded contact homology*, arXiv:1303.5789.

[7] M. Hutchings and J. Nelson, *Cylindrical contact homology for dynamically convex contact forms in three dimensions,*, to appear in JSG.

[8] D. McDuff, *The Hofer conjecture on embedding symplectic ellipsoids*, JDG 88 (2011), 519-532.

[9] J. Pardon, *Contact homology and virtual fundamental cycles*, arXiv:1508.03873.

[10] M. Schwarz, *Cohomology operations from S^1-cobordisms in Floer homology*, PhD thesis, ETH Zurich 1995.

[11] C. H. Taubes, *Embedded contact homology and Seiberg-Witten Floer cohomology I-5*, Geom. Top. (2010).

[12] C. Wendl, *Lectures on symplectic field theory*, arXiv:1612.01009.