Lecture XIV: Pseudoholomorphic curves II

Jonathan Evans

23rd December 2010

In today's lecture, being the last of this course, we'll pull out all the stops and demonstrate some of the enormous power of the pseudoholomorphic curve machinery introduced last time. Our aim will be to prove the following things:

- Gromov's theorem that the (infinite-dimensional) symplectomorphism group of $(S^2 \times S^2, \omega \oplus \omega)$ retracts onto the finite-dimensional subgroup of isometries for the Kähler metric (product of round metrics).
- McDuff's construction of two symplectic forms on $S^2 \times S^2 \times T^2$ which are deformation equivalent (through non-cohomologous forms) but not isotopic (through cohomologous forms).

The first is extraordinarily powerful, illustrates how useful curves can be in 4-d and uses many techniques we've now learned. The second is slightly off-the wall, doesn't rely on 4-d, illustrates how hard things get in higher dimensions and is an extremely beautiful proof.

First I will quickly recap the pseudoholomorphic curve existence proof from last time because I rushed it. The aim was to prove

Theorem

Let $(X, \omega) = (\mathbb{CP}^1 \times V, \omega_{FS} \oplus \omega_V)$ be a product symplectic manifold of dimension $2n$ where the minimal area of a homology class in V is at least $\int_{\mathbb{CP}^1} \omega_{FS} = 1$. Then for any ω-compatible almost complex structure J and any point p there is a J-holomorphic sphere in the homology class $A = [S^2 \times \{\star\}]$ passing through p.

The idea was to consider the moduli space of marked simple J-holomorphic spheres

$$(\mathcal{M}^*(A, J) \times \mathbb{CP}^1)/\mathbb{P}SL(2, \mathbb{C})$$

with its evaluation map $\mathrm{ev} : \mathcal{M}^*(A, J) \to X$. Since the moduli space (for generic J) is a compact manifold (by minimality of A) of dimension $2n + 2c_1(A) + 2 - 6 = 2n$ (plus 2 from a marked point, minus 6 from reparametrisations) we can define the degree of the evaluation map. Degree is invariant under bordism so the number $\deg \mathrm{ev}_J$ is independent of (generic choice of) J.

A product almost complex structure on $X = S^2 \times V$ is regular for curves in the class A (i.e. a regular value of the Fredholm projection map from the universal moduli space to the underlying space of almost complex structures) - this is because the curves are of the form $S^2 \times \{\star\}$ and the normal bundle to such a curve splits as a sum of complex line subbundles with Chern class zero. Therefore it's easy to see that the evaluation map has degree 1 for any regular J (so there's a J-curve in the class A through every point). To see that there's a J-curve in the class A through every point for a non-regular J, recall that the space of regular J is dense so we can find a sequence $J_i \to J$ of regular J_i approximating our irregular J. Take a J_i-curve u_i homologous to A through p. Consider the Gromov-limit of the u_i (which exists by Gromov compactness). Since the class A is minimal there is no bubbling and the limit curve is just a smooth J-curve homologous to A through p.

It turns out in dimension 4 we can say even more.

Theorem (Gromov)

If $(X, \omega) = (S^2 \times S^2, \omega_{FS} \oplus \omega_{FS})$ and J is an ω-compatible almost complex structure then there is a J-holomorphic sphere in each homology class $A = [S^2 \times \{\star\}]$, $B = [\{\star\} \times S^2]$ through every point. In fact there is exactly one J-holomorphic A-sphere α_p and one J-holomorphic B-sphere β_p through any given p and α_p intersects β_p exactly once, transversely.

The only part of this theorem we haven't proved is the uniqueness and intersection property. This phenomenon is very special to 4-dimensional symplectic geometry and I'll explain it in a moment. We will use this theorem to show that the symplectomorphism group of (X, ω) is homotopy equivalent to $SO(3) \times SO(3) \ltimes \mathbb{Z}/2$.

Positivity of intersections in 4d

Theorem (Positivity of intersections, Gromov-McDuff)

Suppose a pair of J-holomorphic curves in a 4-manifold intersect in some set of points p_i and that the intersection at p_i has multiplicity m_i (e.g. $m_i = 1$ means transverse, $m_i = 2$ means tangent but the second derivatives are linearly independent, etc.). Then $m_i > 0$. That is to say each intersection point contributes a positive amount to the homological intersection of the two curves and that contribution is 1 if and only if the intersection is transverse.

Since the curves α_p and β_p in the previous theorem had homological intersection 1 and $[\alpha_p] \cdot [\alpha_p] = [\beta_p] \cdot [\beta_p] = 0$, the unproven parts of that theorem follow from this positivity of intersections property. Positivity of intersections is not easy to prove, but it *is* easy to see that it fails in higher dimensions, even for integrable complex structures: a pair of complex lines in $\mathbb{CP}^3$ can happen to intersect, but one can easily disjoin them because they each have real codimension 4.

Gromov's calculations for $\mathrm{Symp}(S^2 \times S^2, \omega \oplus \omega)$

Theorem

The symplectomorphism group of $(S^2 \times S^2, \omega \oplus \omega)$ *is homotopy equivalent to* $SO(3) \times SO(3) \ltimes \mathbb{Z}/2$*, which can be identified as the subgroup of Kähler isometries (acting by rotations of each factor and the involution* $(x, y) \mapsto (y, x)$*. The symplectomorphism group of* $(S^2 \times S^2, \lambda\omega \oplus \omega)$ *for* $\lambda > 1$ *has an element of infinite order in its fundamental group.*

Exercise

$\clubsuit$: *I will not prove the second statement, I'll just leave it as an (unfairly hard!) exercise with the hint that now the (formerly Lagrangian) antidiagonal can sometimes be represented by a pseudoholomorphic sphere with* $c_1(\overline{\Delta}) = -2$ *(if you get stuck, look at Gromov's original paper (Section 2.4.*C_2*) for a stronger hint and then go and look at later papers of McDuff, Abreu and Anjos for a full proof and indeed a complete working-out of the homotopy type of the symplectomorphism group).*

Proof

We'll actually just show that

$$\mathrm{Symp}_0 = \{\phi \in \mathrm{Symp}(S^2 \times S^2, \omega \oplus \omega) : \phi_* : H_*(S^2 \times S^2; \mathbb{Z}) \circlearrowleft = \mathrm{id}\}$$
$$\simeq SO(3) \times SO(3)$$

Fix a point $(p, q) \in S^2 \times S^2$ and define

$$\mathcal{M} = \{(u, v, J) : \overline{\partial}_J u = \overline{\partial}_J v = 0,\ u_*[S^2] = [S^2 \times \{\star\}],$$
$$v_*[S^2] = [\{\star\} \times S^2],\ (p, q) \in \mathrm{im}(u) \cap \mathrm{im}(v)\}$$

Note that $G \times G$ ($G = \mathbb{P}SL(2, \mathbb{C})$) acts by reparametrisations on $\mathcal{M}$ and by our (uniqueness and) existence theorem the quotient is precisely the space $\mathcal{J}$ of compatible almost complex structures. Note that $G \times G$ retracts onto the subgroup $SO(3) \times SO(3)$.

Consider the diagram

$$\begin{array}{ccccc} SO(3)\times SO(3) & \xrightarrow{\iota} & \mathrm{Symp}_0 & \longrightarrow & \mathrm{Symp}_0/(SO(3)\times SO(3)) \\ \downarrow & & \downarrow \tau & & \downarrow \\ G\times G & \longrightarrow & \mathcal{M} & \longrightarrow & \mathcal{M}/(G\times G) \end{array}$$

where $\tau(\psi)=(\psi\circ\alpha,\psi\circ\beta,\psi_*J_0)$, $\alpha(x)=(x,q')$, $\beta(y)=(p',y)$, $(p',q')=\psi^{-1}(p,q)$ and J_0 is the standard product complex structure. Define

$$\tau\circ\iota(SO(3)\times SO(3))=\mathcal{M}_0\subset\mathcal{M}$$

Since the horizontal maps in the diagram form fibre sequences, since $\mathcal{J}$ is contractible and since $SO(3)\times SO(3)\subset G\times G$ is a deformation retract, it's not hard to see that $\mathcal{M}_0$ is a deformation retract of $\mathcal{M}$.

Therefore let $\Psi_t : \mathcal{M} \to \mathcal{M}$ be a homotopy from $\Psi_1 = \mathrm{id}$ to Ψ_0 where $\Psi_0(\mathcal{M}) \subset \mathcal{M}_0$ and $\Psi_t|_{\mathcal{M}_0} = \mathrm{id}_{\mathcal{M}_0}$. We will construct a map $\mathcal{F} : \mathcal{M} \to \mathrm{Symp}_0$ which is a left-inverse of τ, i.e. $\mathcal{F} \circ \tau = \mathrm{id}$. That is we will find a smooth choice of symplectomorphisms over $\mathcal{M}$ such that if $(u, v, J) = (\psi \circ \alpha, \psi \circ \beta, \psi_* J_0)$ then $\mathcal{F}(u, v, J) = \psi$. Then the composition $\mathcal{F} \circ \Psi_t \circ \tau$ will give a homotopy from id (at time 1) to a deformation retract of Symp_0 onto $SO(3) \times SO(3)$ at time 0.

The map $\mathcal{F}$ is constructed as follows. Starting with a triple (u, v, J) we want to construct a map $S^2 \times S^2 \circlearrowleft$. Gromov's theorem tells us there are foliations of $S^2 \times S^2$ by J-holomorphic spheres in the homology classes $[\alpha]$ and $[\beta]$. The curves u and v are each leaves in these respective foliations. For $(z_1, z_2) \in S^2 \times S^2$ define α_{z_2} to be the (unique) leaf of the α-foliation through $v(z_2)$ and β_{z_1} to be the leaf of the β-foliation through $u(z_1)$. These leaves intersect in a single point, which we denote $\phi(z_1, z_2)$ and the assignment ϕ turns out to be a diffeomorphism (which relies on the elliptic nature of the Cauchy-Riemann equations).
Note that ϕ preserves the curves u and v and fixes their intersection point (p, q).

Now we want to obtain a symplectomorphism ψ from ϕ in a canonical way and we will set $\mathcal{F}(u, v, J) = \psi$. This will be achieved by a Moser isotopy.

Lemma

The forms $\omega_t = (1-t)\omega + t\phi^*\omega$ are symplectic ($t \in [0,1]$).

Once we have the lemma, since the forms ω_t are cohomologous there exists a Moser isotopy θ_t such that $\theta_0 = \mathrm{id}$, $\theta_1^*\omega_1 = \omega$ and so $\psi = \phi \circ \theta_1$ is a symplectomorphism. The Moser isotopy is canonically determined once we have chosen an antiderivative for $\dot{\omega}_t = \phi^*\omega - \omega$ but we can simply chose the unique 1-form in the image of d^* (for a fixed metric) by Hodge theory.

Proof of lemma.

The ϕ-preimages of the rulings $\alpha = S^2 \times \{\star\}$ and $\beta = \{\star\} \times S^2$ are J-holomorphic curves by construction and hence ω-symplectic (by ω-tameness of J). Moreover $\phi^*\omega$ is positive on these preimages because ω is positive on the rulings (the symplectic form is just the product form). Now $\omega_t = (1-t)\omega + t\phi^*\omega$ is obviously closed, so it's symplectic if

$$0 < \omega_t^2 = (1-t)^2\omega^2 + t^2(\phi^*\omega)^2 + 2t(1-t)\omega \wedge \phi^*\omega$$

since both $\omega^2 > 0$ and $(\phi^*\omega)^2 > 0$ it suffices to check that the cross-term is positive. But if $v_1, \ldots, v_4$ is a local frame for $T(S^2 \times S^2)$ such that v_1, v_2 span $T(S^2 \times \{\star\})$ and v_3, v_4 span $T(\{\star\} \times S^2)$ then the cross-term evaluates on the frame to give terms like $\omega(v_1, v_2)\phi^*\omega(v_3, v_4)$ (which are manifestly positive) and $\omega(v_1, v_3)\phi^*\omega(v_2, v_4)$ (which vanish because the standard foliations are ω-orthogonal so $\phi^*\omega(v_2, v_4) = 0$) This completes the proof of the lemma. □

The proof of this theorem also applies to construct a symplectomorphism proving the following theorem which we used in our proof of Luttinger unknottedness.

Theorem

A symplectic 4-manifold which is symplectomorphic to $\mathbb{C}^2$ outside a compact set (and which has $H_2 = 0$) is globally symplectomorphic to $\mathbb{C}^2$.

Standardness at infinity allows you to compactify to $S^2 \times S^2$. The hypothesis on H_2 is used to rule out bubbling of pseudoholomorphic spheres. If H_2 is allowed to be nonvanishing then you can also get blow-ups of $\mathbb{C}^2$.

The remainder of the lecture was not delivered during the course due to time constraints.

Non-isotopic symplectic forms after McDuff

On $X = S^2 \times S^2 \times T^2$, take the symplectic form $\Omega = \lambda\omega_{S^2} \oplus \omega_{S^2} \oplus \omega_{T^2}$ where each ω has area 1 on the corresponding factor and $\lambda \geq 1$. Let $\phi_{z,s}$ be the $2\pi s$-rotation of S^2 around the axis through the point $z \in S^2$. Define the diffeomorphism

$$\psi : X \circlearrowleft, \ \psi(x, y, (s, t)) = (x, \phi_{x,s}(y), (s, t))$$

where x and y are points on S^2 and $(s, t) \in T^2$. Let $\Omega' = (\psi^{-1})^*\Omega$.

Theorem (McDuff)

The forms Ω and Ω' are not isotopic (through cohomologous symplectic forms) when $\lambda = 1$ (though they are for any $\lambda > 1$).

Henceforth we'll assume $\lambda = 1$.

Proof

Let A be the homology class $S^2 \times \{\star\} \times \{\star\}$ and let J be the Ω-compatible complex structure (direct sum of standard complex structures on each factor). The moduli space of simple curves in the class A, $\mathcal{M}^*(A, J)$ is compact since A has minimal area. The expected dimension is $2n + 2c_1(A) - 6 = 2c_1(A) = 4$. Take two marked points on the domain

$$\mathcal{M}_{0,2}(A, J) = \left(\widetilde{\mathcal{M}}^*(A, J) \times S^2 \times S^2\right) / \mathbb{P}SL(2, \mathbb{C})_{\mathrm{diag}}$$

and let ev_i denote the evaluation map at the ith marked point. $\mathcal{M}_{0,2}(A, J)$ has expected dimension 8 (4 from the moduli space plus 2 from each marked point). It is not hard to check that the standard J is actually regular (this involves the computation of a Dolbeaut cohomology group[1]) and so it's unsurprising that the moduli space has the right dimension.

[1] $H^{0,1}_{\bar\partial}(\mathbb{CP}^1, \mathcal{O}(2) \oplus \mathcal{O}^2) = H^{1,0}_{\bar\partial}(\mathbb{CP}^1, \mathcal{O}(-2) \oplus \mathcal{O}^2)^* =$ $H^0(\mathbb{CP}^1, \mathcal{O}(-4) \oplus \mathcal{O}(-2)^2)^* = 0$ (by Serre duality) since the sphere has normal bundle $\mathcal{O}^2$.

We can generically choose J to make ev_2 transverse to the s-circle $\gamma = \{(x_0, y_0, (s, t_0)) : s \in S^1\}$ (indeed it is already transverse for our standard J). The space

$$\mathcal{M}(\gamma, J) = \mathrm{ev}_2^{-1}(\gamma)$$

is therefore a manifold of dimension $3 = 8 - 5$, 5 being the codimension of γ in X. Varying J gives a cobordism of moduli spaces. In particular a generic path J_t of compatible almost complex structures gives a cobordism $\mathcal{M}(\gamma, J_t)$ of $\mathcal{M}(\gamma, J_0)$ with $\mathcal{M}(\gamma, J_1)$ and the evaluation map $\mathrm{ev}_1 : \mathcal{M}(\gamma, J_t) \to X$ gives a bordism from $\mathrm{ev}_1 : \mathcal{M}(\gamma, J_0) \to X$ to $\mathrm{ev}_1 : \mathcal{M}(\gamma, J_1) \to X$.

Suppose that Ω and Ω' were isotopic through a family Ω_t of symplectic forms and suppose J_t is a path of generic Ω_t-compatible almost complex structures joining the standard J with $\psi_* J$, its pushforward under ψ (which is compatible with $\Omega' = (\psi^{-1})^* \Omega$). Varying symplectic forms don't matter for any of the foregoing, because the transversality setup at no point uses the existence of a taming symplectic form. That is only used to prove compactness, and since we have taming symplectic forms all the way along our path we know that we get the relevant energy bounds to prove Gromov compactness. So the general theory of pseudoholomorphic curves gives us

Lemma

If Ω and Ω' were isotopic then the evaluation maps

$$\mathrm{ev}_1 : \mathcal{M}(\gamma, J) \to X \text{ and } \mathrm{ev}_1 : \mathcal{M}(\gamma, \psi_* J) \to X$$

would be bordant.

We will show that they are not bordant maps.

Caveat

Notice that we can deform Ω to Ω' through non-cohomologous symplectic forms (by letting λ get slightly larger than 1). Why does this not give us a bordism of evaluation maps? The point is that the homology class $[S^2 \times \{\star\} \times \{\star\}]$ is no longer minimal when $\lambda > 1$, so our moduli spaces are no longer compact manifolds, we need to take bubbling into account and this messes up our bordism. In other words, McDuff's beautiful argument really uses the symmetry of the symplectic form, just like Gromov's computation of the homotopy type of the symplectomorphism group of $S^2 \times S^2$ breaks down when you use a nonmonotone symplectic form.

Bordism

What is the space $\mathcal{M}(\gamma, J)$? Pretty clearly it's just $S^1 \times S^2$: through each $(x_0, y_0, (s, t_0)) \in \gamma$ there is a single J-sphere $(\{(x, y_0, (s, t_0)) : x \in S^2\}$ (remember there is still a marked point on the domain of a map in $\mathcal{M}(\gamma, J)$, so we can think of S^1 as γ and S^2 as the domain, giving $S^1 \times S^2$ overall). A well-known bordism invariant of maps is the Hopf invariant (of maps from S^3 to S^2)[2]. So to imitate the Hopf map we turn ev_1 into a map from $S^1 \times S^2$ to a 2-manifold, by projecting $X = S^2 \times S^2 \times T^2$ onto its second factor. This gives us maps

$$f_0, f_1 : S^1 \times S^2 \to S^2$$

associated to J and $\psi_* J$ respectively.

[2] As Michael Weiss kindly pointed out to me, the Hopf invariant isn't a bordism invariant. However it *is* invariant if the bordism in question satisfies the condition that the map on H_2 from the cobordism to S^2 vanishes, which is indeed what McDuff uses.

It's clear from the construction that $f_0(S^1 \times S^2) = \{y_0\}$, which is nullbordant. Since $S^1 \times S^2$ is the zero-surgery on the unknot in S^3 we can take the surgery cobordism with the constant map to y_0 as a bordism from the constant map $S^3 \to S^2$ to f_0. We'll construct a bordism from f_1 to the Hopf map $S^3 \to S^2$ and then existence of a bordism from f_0 to f_1 would give rise to a bordism from the constant map to the Hopf map, which contradicts bordism invariance of the Hopf invariant.
The map f_1 is given by

$$(s, x) \mapsto \phi_{x,s}(y_0)$$

This sends the sphere $\{(0, x) : x \in S^2\}$ to the point y_0 and the circle $K = \{(s, y_0) : s \in S^1\}$ to the point y_0. Hence it factors through

$$S^1 \times S^2 \to S^3 \to SO(3) \to S^2$$

The composite $S^3 \to SO(3) \to S^2$ is the Hopf map. It's easy to perturb f_1 to make it constant on a tubular neighbourhood of the circle K, then surger the circle out to get a cobordism to S^3 over which the map extends as a bordism to a (perturbed) Hopf map.

We have now opened up a can of pseudoholomorphic worms. There are two sensible directions to pursue for those who are interested. First, chapter 9 of McDuff-Salamon's big book which has lots more interesting applications of this theory (including obstructions to Lagrangian embeddings, existence of periodic orbits and much, much more). Second, the rest of McDuff-Salamon's big book proves all the analytical details which I have been too lazy/time-constrained to mention. If anyone's interested, we can set up some kind of reading group next term to cover this stuff, but someone else will have to write the lectures!