Lectures on Symplectic Topology*
Park City 1997

Dusa McDuff

August, 1997

Instead of trying to give a comprehensive overview of the subject, I will concentrate on explaining a few key concepts and their implications, notably "Moser's argument" (or the homotopy method) in Lecture 2, capacity in Lecture 4 and Gromov's proof of the nonsqueezing theorem in Lecture 5. The first exhibits the flexibility of symplectic geometry while the latter two show its rigidity. Quite a lot of time is spent on the linear theory since this is the basis of everything else. The last lecture sketches the bare outlines of the theory of J-holomorphic spheres, to give an introduction to a fascinating and powerful technique.

Throughout the notation is consistent with that used in [MS1] and [MS2]. Readers may consult those books for more details on almost every topic mentioned here, as well as for a much fuller list of references.

I wish to thank Jenn Slimowitz for taking the notes and making useful comments on an earlier version of this manuscript.

1 Lecture 1: Basics

Symplectic geometry is the geometry of a skew-symmetric form. Let M be a manifold of dimension $2n$. A symplectic form (or symplectic structure) on M is a closed nondegenerate 2-form ω. Nondegeneracy means that $\omega(v, w) = 0$ for all $w \in TM$ only when $v = 0$. Therefore the map

$$I_\omega : T_pM \to T_p^*M : \quad v \mapsto \iota(v)\omega = \omega(v, \cdot)$$

is injective and hence an isomorphism. The basic example is

$$\omega_0 = dx_1 \wedge dy_1 + \ldots + dx_n \wedge dy_n$$

on $\mathbf{R}^{2n}$.

*this article is published in *IAS/Park City Math. Series* vol 7, ed Eliashberg and Traynor, AMS (1998)

Here are some fundamental questions.

• Can one get a geometric understanding of the structure defined by a symplectic form?

• Which manifolds admit symplectic forms?

• When are two symplectic manifolds (eg two open sets in $(\mathbf{R}^{2n}, \omega_0)$) symplectomorphic?

Definition 1.1 A diffeomorphism $\phi : (M, \omega) \to (M'\omega')$ is called a symplectomorphism if $\phi^*(\omega') = \omega$. The group of all symplectomorphisms is written $\mathrm{Symp}(M)$.

Existence of many symplectomorphisms

Given a function $H : M \to \mathbf{R}$ – often called the *energy* function or *Hamiltonian* – let X_H be the vector field defined by

$$\iota(X_H)\omega = dH.$$

(Observe that $X_H = (I_\omega)^{-1}(dH)$ is well defined because of the nondegeneracy of ω. Also, many authors put a minus sign in the above equation.) When M is compact, X_H integrates to a flow ϕ_t^H that preserves ω because

$$\mathcal{L}_{X_H}\omega = \iota(X_H)d\omega + d(\iota(X_H)\omega) = ddH = 0.$$

Here we have used both that ω is closed and that it is nondegenerate. The calculation

$$dH(X_H) = (\iota(X_H))\,\omega(X_H) = \omega(X_H, X_H) = 0$$

shows that X_H is tangent to the level sets of H. Thus the flow of ϕ_t^H preserves the function H.

Example 1.2 With $H : \mathbf{R}^2 \to \mathbf{R}$ given by $H(x, y) = y$ and with $\omega = dx \wedge dy$, we get

$$X_H = \frac{\partial}{\partial x} \quad \text{and} \quad \phi_t^H(x, y) = (x + t, y).$$

If $H : \mathbf{R}^{2n} \to R$ and $\omega = \omega_0$, we get

$$X_H = \sum_i \frac{\partial H}{\partial y_i}\frac{\partial}{\partial x_i} - \frac{\partial H}{\partial x_i}\frac{\partial}{\partial y_i}.$$

The solution curves $(x_i(t), y_i(t)) = \phi_t(x(0), y(0))$ satisfy Hamilton's equations

$$\dot{x}_i = \frac{\partial H}{\partial y_i}, \quad \dot{y}_i = -\frac{\partial H}{\partial x_i}.$$

With $H = \frac{1}{2}\sum x_j^2 + y_j^2$, the orbits of this action are circles. In complex coordinates $z_j = x_j + iy_j$ we have

$$\phi_t^H(z_1, \ldots, z_n) = (e^{-it}z_1, \ldots, e^{-it}z_n).$$

Thus we get a circle action on $\mathbf{R}^{2n} = \mathbf{C}^n$. The function H that generates it is called the *moment map* of this action.

Exercise 1.3 Often it is useful to consider Hamiltonian functions that depend on time: viz:

$$H : M \times [0,1] \to \mathbf{R}, \quad H(p,t) = H_t(p).$$

Then one defines X_{H_t} as before, and gets a smooth family ϕ_t^H of symplectomorphisms with $\phi_0^H = id$ which at time t are tangent to X_{H_t}:

$$\frac{d}{dt}(\phi_t^H(p)) = X_{H_t}(\phi_t^H(p)), \quad p \in M, t \in [0,1].$$

Such a family is called a *Hamiltonian isotopy*. Show that the set of all time-1 maps ϕ_1^H forms a subgroup of Symp(M). This is called the group of *Hamiltonian symplectomorphisms* Ham(M). Its elements are also often called *exact* symplectomorphisms.

Linear symplectic geometry

To get a better understanding of what is going on, let's now look at what happens at a point. As we shall see, linear symplectic geometry contains a surprising amount of structure. Moreover, most of this structure at a point corresponds very clearly to nonlinear phenomena. One example of this is Darboux's theorem. We shall see in a minute that there is only one symplectic structure on a given (finite-dimensional) vector space, up to isomorphism. Darboux's theorem says that, locally, there is only one symplectic form on a smooth manifold. In other words, every symplectic form ω on M is locally symplectomorphic to the standard form ω_0 on $\mathbf{R}^{2n}$. One might think that this implies there is no interesting local structure (just as if one were in the category of smooth manifolds.) But this is false, since, as we shall see, the standard structure ω_0 on $\mathbf{R}^{2n}$ is itself very interesting.

So let V be a vector space (over $\mathbf{R}$) with a nondegenerate skew bilinear form ω. Thus

$$\omega(v,w) = -\omega(w,v), \quad \omega(v,w) = 0 \text{ for all } v \in V \text{ implies } w = 0.$$

The basic example is $\mathbf{R}^{2n}$ with the form ω_0 considered as a bilinear form.

Given a subspace W define its *symplectic orthogonal* W^ω by:

$$W^\omega = \{v : \omega(v,w) = 0 \text{ for all } w \in W\}.$$

Lemma 1.4 $\dim W + \dim W^\omega = \dim V$.

Proof: Check that the map

$$I : V \to W^* : \quad v \mapsto \omega(v, \cdot)|_W$$

is surjective with kernel W^ω. □

A subspace W is said to be *symplectic* if $\omega|_W$ is nondegenerate. It is easy to check that:

Lemma 1.5 *W is symplectic* $\Longleftrightarrow W \cap W^\omega = \{0\} \Longleftrightarrow V \cong W \oplus W^\omega$.

Proof: Exercise. □

Further we say that W is *isotropic* iff $W \subset W^\omega$ and that W is *Lagrangian* iff $W = W^\omega$. In the latter case $\dim W = n$ by Lemma 1.4.

Proposition 1.6 *Every symplectic vector space is isomorphic to* $(\mathbf{R}^{2n}, \omega_0)$.

Proof: A basis $u_1, v_1, \ldots, u_n, v_n$ of (V, ω) is said to be standard if, for all i, j,

$$\omega(u_i, u_j) = \omega(v_i, v_j) = 0, \quad \omega(u_i, v_j) = \delta_{ij}.$$

Clearly $(\mathbf{R}^{2n}, \omega_0)$ posses such a basis. Further any linear map that takes one such basis into another preserves the symplectic form. Hence we just have to construct a standard basis for (V, ω).

To do this, start with any $u_1 \neq 0$. Choose v so that $\omega(u_1, v) = \lambda \neq 0$ and set $v_1 = v/\lambda$. Let W be the span of u_1, v_1. Then W is symplectic, so $V = W \oplus W^\omega$ by Lemma 1.5. By induction, we may assume that W^ω has a standard basis $u_2, v_2, \ldots, u_n, v_n$. It is easy to check that adding u_1, v_1 to this makes a standard basis for (V, ω). □

Exercise 1.7 (i) Show that if L is a Lagrangian subspace of the symplectic vector space (V, ω), any basis $u_1, \ldots, u_n$ for V can be extended to a standard basis $u_1, v_1, \ldots, u_n, v_n$ for (V, ω). (Hint: choose $v_1 \in W^\omega$ where W is the span of $u_2, \ldots, u_n$.)
(ii) Show that (V, ω) is symplectomorphic to the space $(L \oplus L^*, \tau)$ where

$$\tau((\ell, v^*), (\ell', v'^*)) = v'^*(\ell) - v^*(\ell').$$

The next exercise connects the linear theory with the Hamiltonian flows we were considering earlier.

Exercise 1.8 (i) Check that every codimension 1 subspace W is *coisotropic* in the sense that $W^\omega \subset W$. Note that W^ω is 1-dimensional. For obvious reasons its direction is called the *null* direction in W.
(ii) Given $H : M \to \mathbf{R}$ let $Q = H^{-1}(c)$ be a regular level set. Show that $X_H(p) \in (T_pQ)^\omega$ for all $p \in Q$. Thus the direction of X_H is determined by the level set Q. (Its size is determined by H.)

Exercise 1.9 Show that if ω is any symplectic form on a vector space of dimension $2n$ then the nth exterior power ω^n does not vanish. Deduce that the nth exterior power $\Omega = \omega^n$ of any symplectic form ω on a $2n$-dimensional manifold M is a volume form. Further every symplectomorphism of M preserves this volume form.

The cotangent bundle

This is another basic example of a symplectic manifold. The cotangent bundle TX carries a canonical 1-form λ_{can} defined by

$$(\lambda_{can})_{(x,v^*)}(w) = v^*(\pi_*(w)), \quad \text{for} \quad w \in T_{(x,v^*)}(T^*X),$$

where $\pi : TX \to X$ is the projection. (Here x is a point in X and $v^* \in T^*_xX$.) Then $\Omega_{can} = -d\lambda_{can}$ is a symplectic form. Clearly the fibers of $\pi : T^*X \to X$ are Lagrangian with respect to Ω_{can}, as is the zero section. Moreover, it is not hard to see that:

Lemma 1.10 *Let $\sigma_\alpha : X \to T^*X$ be the section determined by the 1-form α on X. Then $\sigma_\alpha^*(\lambda_{can}) = \alpha$. Hence the manifold $\sigma_\alpha(X)$ is Lagrangian iff α is closed.*

Exercise 1.11 (i) Take a function H on X and let $\widetilde{H} = H \circ \pi$. Describe the resulting flow on T^*X.
(ii) Every diffeomorphism ϕ of X lifts to a diffeomorphism $\widetilde{\phi}$ of T^*X by

$$\widetilde{\phi}(x, v^*) = (\phi(x), (\phi^{-1})^*v^*).$$

Show that $\widetilde{\phi}^*(\lambda_{can}) = \lambda_{can}$.
(iii) Let ϕ_t be the flow on X generated by a vector field Y. If $\widetilde{\phi}_t$ is the lift of this flow to T^*X show that the Hamiltonian $H : T^*X \to \mathbf{R}$ that generates this flow has the form

$$H(x, v^*) = v^*(Y(x)).$$

Hint: use (ii) and write down the defining equation for $\widetilde{Y} = X_H$ in terms of λ_{can}.

2 Lecture 2: Moser's argument

In this lecture I will show you a powerful argument due to Moser [M] which exhibits the "flabbiness" or lack of local structure in symplectic geometry. Here is the basic argument.

Lemma 2.1 *Suppose that ω_t is a family of symplectic forms on a closed manifold M whose time derivative is exact. Thus*

$$\dot{\omega}_t = d\sigma_t,$$

where σ_t is a smooth family of 1-forms. Then there is a smooth family of diffeomorphisms ϕ_t with $\phi_0 = id$ such that

$$\phi_t^*(\omega_t) = \omega_0.$$

Proof: We construct ϕ_t as the flow of a time-dependent vector field X_t. We know

$$\begin{aligned}
\phi_t^*(\omega_t) = \omega \quad &\Longleftrightarrow \quad \frac{d}{dt}(\phi_t^* \omega_t) = 0 \\
&\Longleftrightarrow \quad \phi_t^*\left(\dot{\omega}_t + \mathcal{L}_{X_t}\omega_t\right) = 0 \\
&\Longleftrightarrow \quad \dot{\omega}_t + \iota(X_t)d\omega_t + d(\iota(X_t)\omega_t) = 0 \\
&\Longleftrightarrow \quad d(\sigma_t + \iota(X_t)\omega_t) = 0.
\end{aligned}$$

This last equation will hold if $\sigma_t + \iota(X_t)\omega_t = 0$. Observe that for any choice of 1-forms σ_t the latter equation can always be solved for X_t because of the nondegeneracy of the ω_t. Therefore, reading this backwards, we see that we can always find an X_t and hence a family ϕ_t that will do what we want. □

Remark 2.2 (i) The condition $\dot{\omega}_t = d\sigma_t$ is equivalent to requiring that the cohomology class $[\omega_t]$ be constant. For if this class is constant the derivative $\dot{\omega}_t$ is exact for each t so that for each t there is a form σ_t with $\dot{\omega}_t = d\sigma_t$. Thus the problem is to construct these σ_t so that they depend smoothly on t. This can be accomplished in various ways (eg by using Hodge theory, or see Bott–Tu [BT].)

(ii) The previous lemma uses the fact that the forms ω_t are closed and the fact that the equation $\sigma_t + \iota(X_t)\omega_t$ can always be solved. This last is possible only for nondegenerate 2-forms and for nonvanishing top dimensional forms. In particular the argument does apply to volume forms. Note that this case is very different from the symplectic case because there is never any problem in constructing homotopies of volume forms. Indeed, the set of volume forms in a given cohomology class is convex: if ω_0, ω_1 are volume forms with the same orientation the forms $(1-t)\omega_0 + t\omega_1, 0 \le t \le 1$ are also volume forms. Thus all such forms are diffeomorphic. This is not true for symplectic forms. (Exercise: find an example.)

The previous remarks show that one cannot get interesting new symplectic structures by deforming a given structure within its cohomology class, ie:

Corollary 2.3 (Moser's stability theorem) *If $\omega_t, 0 \le t \le 1$, is a family of cohomologous symplectic forms on a closed manifold M then there is an isotopy ϕ_t with $\phi_0 = id$ such that $\phi_t^*(\omega_t) = \omega_0$ for all t.*

Other corollaries apply Moser's argument to noncompact manifolds M. In this case, to be able to define the flow of the vector field X_t one must be very careful to control its support. Since $X_t = 0 \iff \sigma_t = 0$ the problem becomes that of controlling the support of the forms σ_t. We illustrate what is involved by proving Darboux's theorem.

Theorem 2.1 (Darboux) *Every symplectic form on M is locally diffeomorphic to the standard form ω_0 on $\mathbf{R}^{2n}$.*

Proof: Given a point p on M let $\psi : nbhd(p) \to \mathbf{R}^{2n}$ be a local chart that takes p to the origin 0. We have to show that the form ω' obtained by pushing ω forward by ψ is diffeomorphic to the standard form ω_0 near 0. By Proposition 1.6 we can choose ψ so that $\omega' = \omega_0$ at the point 0. Now consider the family

$$\omega_t = (1-t)\omega_0 + t\omega'.$$

Since $\omega_t = \omega_0$ at 0 by construction and nondegeneracy is an open condition, there is some open ball U containing 0 on which all these forms are nondegenerate. Observe that $\dot{\omega}_t = \omega' - \omega_0$. Since U is contractible there is a 1-form σ such that $d\sigma = \omega' - \omega_0$. Moreover, by subtracting the constant form $\sigma(0)$ we can arrange that $\sigma = 0$ at the point 0. Thus the corresponding family of vector fields X_t vanishes at 0. Let ϕ_t be the partially defined flow of X_t. Since 0 is a fixed point, it is easy to see that there is a very small neighborhood V of 0 such that the orbits $\phi_t(p), 0 \le t \le 1$, of the points p in V remain inside U. Thus the ϕ_t are defined on V and $\phi_1^*(\omega') = \omega_0$. □

For another proof of Darboux's theorem (together with much else) see Arnold [A]. The next applications apply this idea to neighborhoods of submanifolds of M. The basic proposition is:

Proposition 2.4 *Let ω_0, ω_1 be two symplectic forms on M whose restrictions to the full tangent bundle of M along some submanifold Q of M agree: ie*

$$\omega_0|_{T_pM} = \omega_1|_{T_pM} \quad \textit{for} \quad p \in Q.$$

Then there is a diffeomorphism ϕ of M such that

$$\phi(p) = p, \quad \textit{for} \quad p \in Q, \qquad \phi^*(\omega_1) = \omega_0 \quad \textit{near } Q.$$

Proof: Again look at the forms $\omega_t = (1-t)\omega_0 + t\omega_1$. As before these are nondegenerate in some neighborhood of Q. Moreover

$$\dot{\omega}_t = \omega_1 - \omega_0$$

is exact near Q. If we find a 1-form σ that vanishes at all points of Q and is such that $d\sigma = \omega_1 - \omega_0$, then the corresponding vector fields X_t will also vanish along Q and will integrate to give the required diffeomorphisms ϕ_t near Q. Such a form σ can be constructed by suitably adapting the usual proof of Poincaré's lemma: see [BT], for example. □

We can get better results by considering special submanifolds Q. Consider for example a symplectic submanifold Q of (M, ω).[1] Then by Lemma 1.5 the normal bundle $\nu_Q = TM/TQ$ of Q may be identified with the symplectic orthogonal $(TQ)^\omega$ to TQ. Moreover ω restricts to give a symplectic structure on ν_Q: this means that each fiber has a natural symplectic structure that is preserved by the transition functions of the bundle. (See Lecture 3.)

Corollary 2.5 (Symplectic neighborhood theorem) *If ω_0, ω_1 are symplectic forms on M that restrict to the same symplectic form ω_Q on the submanifold Q, then there is a diffeomorphism ϕ of M that fixes the points of Q and is such that $\phi^*(\omega_1) = \omega_0$ near Q provided that ω_0 and ω_1 induce isomorphic symplectic structures on the normal bundle ν_Q.*

Proof: The hypothesis implies that there is a linear isomorphism

$$L : TM|_Q \to TM|_Q$$

that is the identity on the subbundle TQ and is such that $L^*(\omega_1) = \omega_0$. It is not hard to see that L may be realised by a diffeomorphism ψ of M that fixes the points of Q. In other words there is a diffeomorphism with $d\psi_p = L_p$ at each point of Q. Then

$$\omega_0|_{T_pM} = \psi^*\omega_1|_{T_pM}, \quad p \in Q,$$

and so the result follows from Proposition 2.4. □

We will see in the next lecture that giving an isomorphism class of symplectic structures on a bundle is equivalent to giving an isomorphism class of complex structures on it. Hence the normal data needed to make ω_0 and ω_1 agree near Q is quite weak. For example, if Q has codimension 2, all we need to check is that the two forms induce the same orientation on the normal bundle since the Euler class (or first Chern class) of ν_Q is determined up to sign by its topology.

Another important case is when Q is Lagrangian. In this case one can check that the normal bundle ν_Q is canonically isomorphic to the dual bundle TQ^*. Moreover this dual bundle is also Lagrangian. (Cf Exercise 1.7.) Thus when Q is Lagrangian with respect to both ω_0 and ω_1 there always is a linear isomorphism

$$L : TM|_Q \to TM|_Q$$

that is the identity on the subbundle TQ and is such that $L^*(\omega_1) = \omega_0$. Moreover, just as in the case of Darboux's theorem there is a standard model for Q, namely the zero section in the cotangent bundle (T^*Q, Ω_{can}). Thus we have:

[1] A submanifold Q of M is called symplectic if ω restricts to a symplectic form on Q, or, equivalently, if all its tangent spaces $T_pQ, p \in Q$, are symplectic subspaces. Similarly, Q is Lagrangian if $\omega|_Q \equiv 0$ and $\dim Q = n$.

Corollary 2.6 (Weinstein's Lagrangian neighborhood theorem) *If Q is a Lagrangian submanifold of (M, ω) there is a neighborhood of Q that is symplectomorphic to a neighborhood of the zero section in the cotangent bundle (T^*Q, Ω_{can}).*

Exercise 2.7 Given any two diffeomorphic closed smooth domains U, V in $\mathbf{R}^n$ that have the same total volume, show that there is a diffeomorphism $\phi : U \to V$ that preserves volume. Hint: first choose any diffeomorphism $\psi : U \to V$ and look at the forms $\omega_0, \omega_1 = \psi^*(\omega_0)$ on U. Adjust ψ by hand near the boundary ∂U so that $\omega_0 = \omega_1$ at all points on ∂U. Then use a Moser type argument to make the forms agree in the interior.

The last important result of this kind is the symplectic isotopy extension theorem due to Banyaga. The proof is left as an exercise.

Proposition 2.8 (Isotopy extension) *Let Q be a compact submanifold of (M, ω) and suppose that $\phi_t : M \to M$ is a family of diffeomorphisms of M starting at $\phi_0 = id$ such that $\phi_t^*(\omega) = \omega$ near Q. Then, if for every relative cycle $Z \in H_2(M, Q)$*

$$\int_Z \phi_t^*(\omega) = \int_Z \omega,$$

there is a family of symplectic diffeomorphisms ψ_t and a neighborhood U of Q such that $\psi_t(p) = \phi_t(p)$ for all $p \in U$.

3 Lecture 3: The linear theory

We will consider the vector space $\mathbf{R}^{2n}$ with its standard symplectic form ω_0. This may be written in vector notation as

$$\omega_0(v, w) = w^T J_0 v,$$

where w^T is the transpose of the column vector w and J_0 is the block diagonal matrix

$$J_0 = \operatorname{diag}\left(\begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix}, \ldots, \begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix} \right).$$

The symplectic linear group $\mathrm{Sp}(2n, \mathbf{R})$ (sometimes written $\mathrm{Sp}(2n)$) consists of all matrices A such that

$$\omega_0(Av, Aw) = \omega_0(v, w),$$

or equivalently of all A such that

$$A^T J_0 A = J_0.$$

Clearly $\mathrm{Sp}(2n, \mathbf{R})$ is a group. The identity $J_0^T = -J_0 = J_0^{-1}$ gives rise to interesting algebraic properties of this group. Firstly, it is closed under transpose, and secondly every symplectic matrix is conjugate to its inverse transpose $(A^{-1})^T$. The former statement is proved by inverting the identity

$$(A^{-1})^T J_0 A^{-1} = J_0,$$

and the second by multiplying the defining equation $A^T J_0 A = J_0$ on the left by $(J_0)^{-1}(A^T)^{-1}$.

Exercise 3.1 Show that if $\lambda \in \mathbf{C}$ is an the eigenvalue of a symplectic matrix A then so are $1/\lambda, \bar{\lambda}, 1/\bar{\lambda}$. What happens when $\lambda \in \mathbf{R}$, or $|\lambda| = 1$?

Recall that we are identifying $\mathbf{C}^n$ with $\mathbf{R}^{2n}$ by setting $z_j = x_j + iy_j$. Under this identification, J_0 corresponds to multiplication by i. Hence we may consider $\mathrm{GL}(n, \mathbf{C})$ to be the subgroup of $\mathrm{GL}(2n, \mathbf{R})$ consisting of all matrices A such that $AJ_0 = J_0 A$.

Exercise 3.2 Given an $n \times n$ matrix A with complex entries, find a formula for the corresponding real $2n \times 2n$ matrix.

Lemma 3.3

$$\mathrm{Sp}(2n, \mathbf{R}) \cap \mathrm{O}(2n) = \mathrm{Sp}(2n, \mathbf{R}) \cap \mathrm{GL}(n, \mathbf{C}) = \mathrm{O}(2n) \cap \mathrm{GL}(n, \mathbf{C}) = \mathrm{U}(n).$$

Proof: Exercise. □

Our first main result is that $\mathrm{U}(n)$ is a maximal compact subgroup of $\mathrm{Sp}(2n)$ and hence, by the general theory of Lie groups, the quotient space $\mathrm{Sp}(2n)/\mathrm{U}(n)$ is contractible. The first statement above means that any compact subgroup G of $\mathrm{Sp}(2n)$ is conjugate to a subgroup of $\mathrm{U}(n)$. We won't prove this here since we will not use it. However, we will give an independent proof of the second.

To begin, recall the usual proof that $\mathrm{GL}(n, \mathbf{R})/\mathrm{O}(n)$ is contractible. One looks at the polar decomposition

$$A = (AA^T)^{\frac{1}{2}} O$$

of A. Here $P = AA^T$ is a symmetric, positive definite[2] matrix and hence diagonalises with real positive eigenvalues. In other words P may be written $X \Lambda X^{-1}$ where Λ is a diagonal matrix with positive entries. One can therefore define an arbitrary real power P^α of P by

$$P^\alpha = X \Lambda^\alpha X^{-1}.$$

[2]Usually a positive definite matrix is assumed to be symmetric (ie $P = P^T$). However, in symplectic goemetry one does come across matrices that satisfy the positivity condition $v^T P v > 0$ for all nonzero v but that are not symmetric. Hence it is better to mention the symmetry explicitly.

It is easy to check that

$$O = (AA^T)^{-\frac{1}{2}}A$$

is orthogonal. Hence one can define a deformation retraction of $\mathrm{GL}(2n, \mathbf{R})$ onto $\mathrm{O}(2n)$ by

$$A \mapsto (AA^T)^{\frac{1-t}{2}}O, \quad 0 \le t \le 1.$$

The claim is that this argument carries over to the symplectic context. To see this we need to show:

Lemma 3.4 *If $P \in \mathrm{Sp}(2n)$ is positive and symmetric then all its powers P^α, $\alpha \in \mathbf{R}$ are also symplectic.*

Proof: Let V_λ be the eigenspace of P corresponding to the eigenvalue λ. Then, if $v_\lambda \in V_\lambda, v_{\lambda'} \in V_{\lambda'}$,

$$\omega_0(v_\lambda, v_{\lambda'}) = \omega_0(Pv_\lambda, Pv_{\lambda'}) = \omega_0(\lambda v_\lambda, \lambda' v_{\lambda'}) = \lambda\lambda'\omega_0(v_\lambda, v_{\lambda'}).$$

Hence $\omega_0(v_\lambda, v_{\lambda'}) = 0$ unless $\lambda' = 1/\lambda$. In other words the eigenspaces $V_\lambda, V_{\lambda'}$ are symplectically orthogonal unless $\lambda' = 1/\lambda$. To check that P^α is symplectic we just need to know that

$$\omega_0(P^\alpha v_\lambda, P^\alpha v_{\lambda'}) = \omega_0(v_\lambda, v_{\lambda'}),$$

for all eigenvectors $v_\lambda, v_{\lambda'}$. But this holds since

$$\omega_0(P^\alpha v_\lambda, P^\alpha v_{\lambda'}) = \omega_0(\lambda^\alpha v_\lambda, (\lambda')^\alpha v_{\lambda'}) = (\lambda\lambda')^\alpha \omega_0(v_\lambda, v_{\lambda'}) = \omega_0(v_\lambda, v_{\lambda'}).$$

(Observe that everything vanishes when $\lambda\lambda' \neq 1$!) □

Thus the argument given above in the real context extends to the symplectic context, and we have:

Proposition 3.5 *The subgroup $\mathrm{U}(n)$ is a deformation retract of $\mathrm{Sp}(2n)$.*

ω-compatible almost complex structures

An almost complex structure on a vector space V is a linear automorphism $J : V \to V$ with J^2 equal to $-\mathbb{1}$. Thus one can define an action of $\mathbf{C}$ on V by

$$(a + ib)v = a + Jb,$$

so that (V, J) is a complex vector space. If V also has a symplectic form ω we say that ω and J are compatible if for all nonzero v, w,

$$\omega(Jv, Jw) = \omega(v, w), \quad \omega(v, Jv) > 0.$$

The basic example is the pair (ω_0, J_0) on $\mathbf{R}^{2n}$.

Any such pair (J, ω) defines a corresponding metric (inner product) g_J by

$$g_J(v, w) = \omega(v, Jw).$$

This is symmetric because

$$g_J(w, v) = \omega(w, Jv) = \omega(Jw, J^2 v) = \omega(Jw, -v) = \omega(v, Jw) = g_J(v, w).$$

Exercise 3.6 Show that J is ω-compatible iff there is a standard basis of the form

$$u_1, v_1 = Ju_1, \ldots, u_n, v_n = Ju_n.$$

Deduce that there is a linear symplectomorphism $\Phi : (\mathbf{R}^{2n}, \omega_0) \to (V, \omega)$ such that $J = \Phi J_0 \Phi^{-1}$.

Proposition 3.7 *The space of ω-compatible almost complex structures J on V is contractible.*

Proof: Without loss of generality we may suppose that (V, ω) is standard Euclidean space $(\mathbf{R}^{2n}, \omega_0)$. Clearly, $\mathrm{Sp}(2n)$ acts on the space $\mathcal{J}(\omega_0)$ of ω_0-compatible almost complex structures on $\mathbf{R}^{2n}$ by

$$A \cdot J = AJA^{-1}.$$

The preceding exercise shows that this action is transitive, since every J may be written as $J = AJ_0A^{-1}$ and so is in the orbit of J_0. The kernel of the action consists of elements that commute with J_0, in other words of unitary transformations. (See Lemma 3.3.) Thus $\mathcal{J}(\omega_0)$ is isomorphic to the homogeneous space $\mathrm{Sp}(2n)/\mathrm{U}(n)$ and so is contractible by Proposition 3.5. □

Exercise 3.8 Define the form ω_B by

$$\omega_B(v, w) = w^T B J_0 v.$$

Under what conditions on B is ω_B compatible with J_0? Deduce that for each fixed almost complex structure J on V the space of compatible ω is contractible.

Vector bundles

A (real) $2n$-dimensional vector bundle $\pi : E \to B$ is said to be *symplectic* if it has an atlas of local trivializations $\tau_\alpha : \pi^{-1}U_\alpha \to \mathbf{R}^{2n} \times U_\alpha$ such that for all $p \in U_\alpha \cap U_\beta$ the corresponding transition map

$$\phi_{\alpha,\beta}(p) = \tau_\alpha \circ (\tau_\beta)^{-1} : \mathbf{R}^{2n} \times p \to \mathbf{R}^{2n} \times p$$

is a linear symplectomorphism. Using a parametized version of Proposition 1.6, one can easily show that this is equivalent to requiring that there is a bilinear skew form σ on E that is nondegenerate on each fiber. For, given such σ, one can use Proposition 1.6 to choose the trivializations τ_α so that at each point $p \in B$ they pull back the standard form on $\mathbf{R}^{2n} \times p$ to the given form $\sigma(p)$. Then the transition maps have to be symplectic. Conversely, if the transition maps are symplectic the pull backs of the standard form by the τ_α agree on the overlaps to give a well-defined global form σ.

A σ-compatible almost complex structure J on E is an automorphism of E that at each point $p \in B$ is a $\sigma(p)$-compatible almost complex structure on the fiber.

Proposition 3.9 *Every symplectic vector bundle (E, σ) admits a contractible family of compatible almost complex structures, and hence gives rise to a complex structure on E that is unique up to isomorphism. Conversely, any complex vector bundle admits a contractible family of compatible symplectic forms, and hence has a symplectic structure that is unique up to isomorphism. Thus classifying isomorphism classes of symplectic bundles is the same as classifying isomorphism classes of complex bundles.*

Proof: (Sketch) One way of proving the first statement is to note that the space of compatible almost complex structures on E forms a fiber bundle over B that, by Proposition 3.7, has contractible fibers. Another way is to start from the contractible space of inner products on E and to show that each such inner product gives rise to a unique almost complex structure. (The details of this second argument can be found in 2.5,6 of [MS2].) The second statement follows by similar arguments, using Exercise 3.8. □

Clearly the tangent bundle TM of every symplectic manifold (M, ω) is a symplectic vector bundle with symplectic structure given by ω. The previous proposition shows that TM has a well-defined complex structure, and so, in particular, has Chern classes $c_i(TM)$. The first Chern class $c_1(TM) \in H^2(M, \mathbf{Z})$ is a particularly useful class as it enters into the dimension formula for moduli spaces of J-holomorphic curves. (See Lecture 5.)

The Lagrangian Grassmannian

Another interesting piece of linear structure concerns the space $\mathcal{L}(n)$ of all Lagrangian subspaces of a $2n$-dimensional symplectic vector space (V, ω). This is also known as the Lagrangian Grassmannian. Here we will consider the space of unoriented Lagrangian subspaces, but it is easy to adapt our remarks to the oriented case.

Lemma 3.10 *Let J be any ω-compatible almost complex structure on the symplectic vector space (V, ω). Then the subspace $L \subset V$ is Lagrangian if and only if there is a standard basis $u_1, v_1, \ldots, u_n, v_n$, for (V, ω) such that $u_1, \ldots, u_n$ span L and $v_j = Ju_j$ for all j.*

Proof: Let g_J be the associated metric and choose a g_J-orthonormal basis for L. Then

$$\omega(u_i, Ju_j) = -g_J(u_i, u_j) = \delta_{ij}.$$

Hence we get a standard basis by setting $v_j = Ju_j$ for all j. The converse is clear. □

Corollary 3.11 $\mathcal{L}(n) \cong \mathrm{U}(n)/\mathrm{O}(n)$.

Proof: We may take $(V, \omega, J) = (\mathbf{R}^{2n}, \omega_0, J_0)$. Let L_0 be the Lagrangian spanned by the vectors $\frac{\partial}{\partial x_1}, \dots, \frac{\partial}{\partial x_n}$. The previous lemma shows that every Lagrangian subspace L is the image $A(L_0)$ of L_0 under the unitary transformation A that takes $\frac{\partial}{\partial x_j}$ to u_j and $\frac{\partial}{\partial y_j}$ to v_j. Moreover, $A(L_0) = L_0$ exactly when A belongs to the orthogonal subgroup $\mathrm{O}(n) \subset \mathrm{U}(n)$. □

Exercise 3.12 Show that the group $\mathrm{Sp}(2n)$ acts transitively on pairs of transversally intersecting Lagrangians.

Lemma 3.13 $\pi_1(\mathcal{L}(n)) = \mathbf{Z}$.

Proof: The long exact homotopy sequence of the fibration $\mathrm{O}(n) \to \mathrm{U}(n) \to \mathcal{L}(n)$ contains the terms

$$\pi_1(\mathrm{O}(n)) \to \pi_1(\mathrm{U}(n)) \to \pi_1(\mathcal{L}(n)) \to \pi_0(\mathrm{O}(n)) = \mathbf{Z}/2\mathbf{Z}.$$

It is easy to check that the map

$$\mathrm{U}(n) \to S^1 : \quad A \mapsto \det(A)$$

induces an isomorphism on π_1 (where det denotes the determinant over $\mathbf{C}$.) On the other hand $\pi_1(\mathrm{O}(n))$ is generated by a loop in $\mathrm{O}(2)$ and the inclusion $\mathrm{O}(2) \hookrightarrow \mathrm{U}(2)$ takes its image in $\mathrm{SU}(2)$. Hence the map $\pi_1(\mathrm{O}(n)) \to \pi_1(\mathrm{U}(n))$ is trivial. □

It is not hard to check that a generating loop of $\pi_1(\mathcal{L}(n)$ is

$$t \mapsto (e^{\pi i t}\mathbf{R}) \oplus \mathbf{R} \oplus \cdot \oplus \mathbf{R} \subset \mathbf{C}^n, \quad 0 \le t \le 1.$$

The Maslov index

There are several ways to use the structure of the Lagrangian Grassmannian to get invariants. Typically the resulting invariants are called the "Maslov index". Here is one way that is relevant when considering the Lagrangian intersection problem. Suppose we are given two Lagrangian submanifolds Q_0, Q_1 in (M, ω) that intersect transversally. For example Q_0 might be the zero section of the cotangent bundle T^*Q and Q_1 might be the graph of an exact 1-form df

with nondegenerate zeros. In this case one can assign an index to each transversal intersection point $x \in Q_0 \cap Q_1$ by using the usual Morse index for critical points of the function f. Although this is not possible in the general situation, we will now explain how it is possible to define a relative index of pairs x_+, x_- of intersection points. If one has chosen a homotopy class of connecting trajectories u, this index takes values in $\mathbf{Z}$. Here, a connecting trajectory means a map $u : D^2 \to M$ such that

$$\begin{aligned} u(1) &= x_+, & u(-1) &= x_-, \\ u(e^{\pi i t}) &\in Q_0, & 0 &\le t \le 1, \\ u(e^{\pi i t}) &\in Q_1, & 1 &\le t \le 2. \end{aligned}$$

Let us first see how to use this data to define a closed loop $L(t), 0 \le t \le 4$, in $\mathcal{L}(n)$. Note that $u^*(TM)$ is a symplectic bundle over the disc and so is symplectically trivial. Choose a trivialization $\phi : u^*(TM) \to D^2 \times \mathbf{R}^{2n}$. Then, define

$$L(t) = \begin{cases} \phi(u^*(T_{u(e^{\pi i t})} Q_0)), & 0 \le t \le 1, \\ \phi(u^*(T_{u(e^{\pi i t})} Q_1)), & 2 \le t \le 3. \end{cases}$$

For $t \in [1, 2]$ choose any path in $\mathcal{L}(n)$ from $L(1)$ to $L(2)$. To complete the loop, observe that by Exercise 3.12 there is $A \in \mathrm{Sp}(n)$ such that

$$A(L(0)) = L(1), \quad A(L(3)) = L(2).$$

Therefore, we may set

$$L(3 + s) = A(L(2 - s)), \quad 0 \le s \le 1.$$

The Maslov index $\mu_u(x_-, x_+)$ is now defined to be the element in $\pi_1(\mathcal{L}(n) \cong \mathbf{Z}$ represented by this path.

Exercise 3.14 Check that this index is independent of choices.

4 Lecture 4: The Nonsqueezing theorem

In Lecture 2, I explained various results that showed how flexible symplectomorphisms are and how little local structure a symplectic manifold has. Now I want to show you the phenomenon of symplectic rigidity that is encapsulated in Gromov's nonsqueezing theorem [G]. We will consider the cylinder

$$Z(r) = B^2(r) \times \mathbf{R}^{2n-2} = \{(x, y) \in \mathbf{R}^{2n} : x_1^2 + y_1^2 \le r\}$$

with the restriction of the usual symplectic form ω_0.

Theorem 4.1 (Gromov) *If there is a symplectomorphism that maps the the unit ball $B^{2n}(1)$ in $(\mathbf{R}^{2n}, \omega_0)$ into the cylinder $Z(r)$ then $r \ge 1$.*

This deceptively simple result is, as we shall see, enough to characterise symplectomorphisms among all diffeomorphisms. It clearly shows that symplectomorphisms are different from volume-preserving diffeomorphisms since it is easy to construct a volume-preserving diffeomorphism that squeezes the unit ball into an arbitrarily thin cylinder. We will begin discussing the proof at the end of this lecture. For now, let's look at its implications.

The clearest way to understand the force of Theorem 4.1 is to use the Ekeland–Hofer idea of capacity. A *symplectic capacity* is a function c that assigns an element in $[0, \infty]$ to each symplectic manifold of dimension $2n$ and satisfies the following axioms:

(i) (*monotonicity*) if there is a symplectic embedding $\phi : (U, \omega) \to (U', \omega')$ then $c(U, \omega) \le c(U', \omega')$.
(ii) (*conformal invariance*) $c(U, \lambda\omega) = \lambda^2 c(U, \omega)$.
(iii) (*nontriviality*)

$$0 < c(B^{2n}(1), \omega_0) = c(Z(1), \omega_0) < \infty.$$

It is the last property $c(Z(1), \omega_0) < \infty$ that implies that capacity is an essentially 2-dimensional invariant, for example that it cannot be a power of the total volume. Sometimes one considers capacities that satisfy a less stringent version of (iii): namely
(iii′)

$$0 < c(B^{2n}(1), \omega_0), \quad c(Z(1), \omega_0) < \infty.$$

However, below we will use the strong form (iii).

The interesting question is: do symplectic capacities exist? A moment's reflection shows that the fact that they do is essentially equivalent to the nonsqueezing theorem. Let us define the Gromov capacity c_G by

$$c_G(U, \omega) = \sup\{\pi r^2 : B^{2n}(r) \text{ embeds symplectically in } U\}.$$

Then c_G clearly satisfies the conditions (i), (ii), and also $c_G(B^{2n}(r)) = \pi r^2$. The only difficult thing to check is that $c_G(Z(1)) < \infty$, but in fact

$$c_G(Z(r)) = \pi r^2$$

by the nonsqueezing theorem. Thus c_G is a capacity. There are now several other known capacities, (cf work by Ekeland–Hofer [EH], Hofer–Zehnder [HZ], Viterbo [V]) mostly defined by looking at properties of the periodic flows of certain Hamiltonian functions H that are associated to U.

The main result is

Theorem 4.2 (Ekeland–Hofer) *A (local) orientation-preserving diffeomorphism ϕ of $(\mathbf{R}^{2n}, \omega_0)$ is symplectic iff it preserves the capacity of all open subsets of $\mathbf{R}^{2n}$, ie iff there is a capacity c such that $c(\phi(U)) = c(U)$ for all open U.*

The proof is based on the corresponding result at the linear level.

Proposition 4.1 *A linear map L that preserves the capacity of ellipsoids is either symplectic or antisymplectic, ie $L^*(\omega_0) = \pm\omega_0$.*

Proof: If L is neither symplectic nor antisymplectic the same can be said of its transpose L^T. Therefore there are vectors v, w so that

$$\omega_0(v, w) \neq \pm\,\omega_0(L^T v, L^T w).$$

By perturbing v, w and using the openness of the above condition we can suppose that both $\omega_0(v, w)$ and $\omega_0(L^T v, L^T w)$ are nonzero. Then, replacing L^T by its inverse if necessary, we can arrange that

$$0 < \lambda^2 = |\omega_0(L^T v, L^T w)| < \omega_0(v, w) = 1.$$

Now construct two standard bases of $\mathbf{R}^{2n}$, the first starting as

$$u_1 = v, \quad v_1 = w, \quad u_2, \ldots,$$

and the second starting as

$$u_1' = \frac{L^T v}{\lambda}, \quad v_1' = \pm\frac{L^T w}{\lambda}, \quad u_2', \ldots.$$

Let A, resp A', be the symplectic linear map that takes the standard basis $e_1, e_2, e_3, \ldots$ of $\mathbf{R}^{2n}$ to $u_1, v_1, u_2, \ldots$, resp $u_1', v_1', u_2', \ldots$. Then, setting $C = (A')^{-1} L^T A$, we have

$$C(e_1) = \lambda e_1, \quad C(e_2) = \lambda e_2.$$

Thus the matrices for C and C^T have the form

$$C = \left(\begin{array}{cc|ccc} \lambda & 0 & * & \ldots & * \\ 0 & \lambda & * & \ldots & * \\ \hline 0 & 0 & * & \ldots & * \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ 0 & 0 & * & \ldots & * \end{array}\right), \qquad C^T = \left(\begin{array}{cc|ccc} \lambda & 0 & 0 & \ldots & 0 \\ 0 & \lambda & 0 & \ldots & 0 \\ \hline * & * & * & \ldots & * \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ & * & * & \ldots & * \end{array}\right).$$

It is now easy to check that C^T maps the unit ball into the cylinder $Z(\lambda)$. But because A, L and A' preserve capacity, so does $C^T = A^T L (A'^T)^{-1}$. This contradiction proves the result. Note that we have only needed the fact that C^T preserves the capacity of the unit ball. Hence we only need to know that L preserves the capacity of all sets that are images of the ball by symplectic linear maps, ie the ellipsoids. $\square$

Proof of Theorem 4.2 We want to show that the derivative $d\phi_p$ of ϕ at every point p in its domain is symplectic. By pre- and post-composing with suitable translations, it is easy to see that it suffices to consider the case when $p = 0$ and $\phi(0) = 0$. Then the derivative $d\phi_0$ is the limit in the compact open topology of the diffeomorphisms ϕ_t given by

$$\phi_t(v) = \frac{\phi(tv)}{t}.$$

Because ϕ preserves capacity, and capacity behaves well under rescaling (see condition (ii)), the diffeomorphisms ϕ_t also preserve capacity. Moreover, by the exercise below, the capacity of convex sets is continuous with respect to the Hausdorff topology on sets. Thus the uniform limit $d\phi_0$ of the ϕ_t preserves capacity and so must be either symplectic or anti-symplectic.

To complete the proof, we must show that $d\phi_0$ is symplectic rather than anti-symplectic. If n is odd this follows immediately from the fact that $d\phi_0$ preserves orientation. If n is even, repeat the previous argument replacing ϕ by $id_{R^2} \times \phi$. □

Exercise 4.2 Recall that the Hausdorff distance $d(U, V)$ between two subsets U, V of $\mathbf{R}^{2n}$ is defined to be

$$d(U,V) = \max_{x \in U}\left(\min_{y \in V} \|x - y\|\right) + \max_{y \in V}\left(\min_{x \in U} \|x - y\|\right).$$

Suppose that U is a convex set containing the origin. Show that for all $\varepsilon > 0$ there is $\delta > 0$ such that

$$(1-\varepsilon)U \subset V \subset (1+\varepsilon)U, \quad \text{whenever} \quad d(U,V) < \delta.$$

Using this, prove the claim in the previous proof that $d\phi_0$ preserves capacity.

Corollary 4.3 (Eliashberg, Ekeland–Hofer) *The group* $\mathrm{Symp}(M, \omega)$ *is* C^0*-closed in the group of all diffeomorphisms.*

Proof: We must show that if ϕ_n is a sequence of symplectomorphisms that converge uniformly to a diffeomorphism ϕ_0 then ϕ_0 is itself symplectic. But ϕ_0 preserves the capacity of ellipsoids because capacity is continuous with respect to the Hausdorff topology on convex sets. Hence result. □

Note that these results give us a way of defining symplectic homeomorphisms. In fact, there are two possibly different definitions. One says that a (local) homeomorphism of $\mathbf{R}^{2n}$ is symplectic if it preserves the capacity of all open sets, the other that it is symplectic if it preserves the capacity of all sufficiently small ellipsoids. Very little is known about the properties of such homeomorphisms. In particular, it is unknown whether these two definitions agree and the extent to which they depend on the particular choice of capacity.

Theorem 4.2 makes clear that symplectic capacity is the basic symplectic invariant from which all others are derived. The fact that capacity is C^0-continuous shows the robustness of the property of being symplectic, and is really the reason why there is an interesting theory of symplectic topology. There is much recent work that develops the ideas presented here. Here is a short list of key references: Floer–Hofer [FH] on the theory of symplectic homology, Cieliebak–Floer–Hofer–Wysocki [CFHW] on its appplications, Hofer [H] and Lalonde–McDuff [LM] on the Hofer norm on the group $\mathrm{Ham}(M,\omega)$, and Polterovich [P] on its applications.

There are now many known proofs of the nonsqueezing theorem that are based on the different notions of capacity that have been developed: see for example Ekeland–Hofer [EH] and Viterbo [V]. We shall follow the original proof of Gromov [G] that uses J-holomorphic curves.

Preliminaries on J-holomorphic curves

A J-holomorphic curve (of genus 0) in an almost complex manifold (M,J) is a map

$$u : (S^2, j) \to (M, J): \quad J \circ du = du \circ j,$$

where j is the usual almost complex structure on S^2. This equation may be rewritten as

$$\overline{\partial}_J u = \frac{1}{2}(du + J \circ du \circ j) = 0.$$

In local holomorphic coordinates $z = s + it$ on S^2, j acts by $j(\frac{\partial}{\partial s}) = \frac{\partial}{\partial t}$ and so this translates to the pair of equations:

$$\frac{\partial u}{\partial s} + J(u)\frac{\partial u}{\partial t} = 0, \quad \frac{\partial u}{\partial t} - J(u)\frac{\partial u}{\partial s} = 0.$$

Note that the second of these follows from the first by multiplying by J. Further, if J were constant in local coordinates on M (which is equivalent to requiring that J be integrable) these would reduce to the usual Cauchy–Riemann equations. As it is, these are quasi-linear equations that agree with the Cauchy–Riemann equations up to terms of order zero. Hence they are elliptic.

There is one very important point about J-holomorphic curves in the case when J is compatible with a symplectic form ω. We then have an associated metric g_J and we find (in obvious but rather inexact notation)

$$\begin{aligned}
\int_{S^2} u^*(\omega) &= \int_{S^2} \omega(\frac{\partial u}{\partial s}, \frac{\partial u}{\partial t}) \\
&= \int_{S^2} \omega(\frac{\partial u}{\partial s}, J\frac{\partial u}{\partial s}) \\
&= \int_{S^2} g_J(\frac{\partial u}{\partial s}, \frac{\partial u}{\partial s})\, ds\, dt
\end{aligned}$$

$$
\begin{aligned}
&= \frac{1}{2}\int_{S^2}\left(\left|\frac{\partial u}{\partial s}\right|^2+\left|\frac{\partial u}{\partial t}\right|^2\right)\,ds\,dt\\
&= g_J\text{-area of } \operatorname{Im} u.
\end{aligned}
$$

Thus the g_J-area of a J-holomorphic curve is determined entirely by the homology class A that it represents. Note that $\omega(A)$ is always strictly positive unless $A = 0$: indeed the restriction of ω to a J-holomorphic curve is nondegenerate at all nonsingular points. Further the next exercise implies that such curves are g_J-minimal surfaces. It is possible to develop much of the theory of J-holomorphic curves using this fact. (This is Gromov's original approach. More details can be found in some of the articles in Audin–Lafontaine [AL].) In the next lecture we will sketch the outlines of a rather different approach using standard elliptic analysis.

Exercise 4.4 Let $(V, \omega J)$ be a symplectic vector space with compatible almost complex structure J and associated inner product g_J. Given two vectors v, w denote by $P(v, w)$ the parallelogram they span. Show that

$$\omega(v, w) \le g_J\text{-area of } P(v, w)$$

with equality if and only if $w = Jv$. Deduce that J-hol curves in (M, ω, J) are g_J-minimal surfaces.

5 Lecture 5

Sketch of the proof of the nonsqueezing theorem.

Suppose that $\phi : B^{2n}(1) \to Z(r)$ is a symplectic embedding. Its image lies in some compact subset $B^2(r) \times K$ of $Z(r)$ that can be considered as a subset of the compact manifold $(S^2 \times T^{2n-2}, \Omega)$, where Ω is the sum $\sigma \oplus \kappa\omega_0$ of a symplectic form σ on S^2 with total area $\pi r^2 + \varepsilon$ and a suitable multiple of the standard form ω_0 on T^{2n-2}. Let J_0 be the usual almost complex structure on $\mathbf{R}^{2n}$ and let J be an Ω-compatible almost complex structure on $S^2 \times T^{2n-2}$ that restricts to $\phi_*(J_0)$ on the image of the ball. (It is easy to construct such J using the methods of proof of Proposition 3.9.) As we will see below, the theory of J-holomorphic curves ensures that there is at least one J-holomorphic curve through each point of $S^2 \times T^{2n-2}$ in the class $A = [S^2 \times pt]$. Let C be such a curve through the image $\phi(0)$ of the origin, and let S be the component of the inverse image $\phi^{-1}(C)$ that goes through the origin. Then S is a proper[3] J_0-holomorphic curve in the ball $B^{2n}(1)$ through 0. Since J_0 is the usual complex structure, this means that S is a g_0-minimal surface (where g_0 is the usual metric

[3]This means that the intersection of S with any compact subset of the ball is compact. Thus it goes all the way out to the boundary.

on $\mathbf{R}^{2n}$.) But it is well-known that the proper surface of smallest area through the center of a ball of radius 1 is a flat disc with area π. Hence

$$\pi \le g_0\text{-area of } S = \int_S \omega_0 = \int_{\phi(S)} \Omega < \int_C \Omega = \int_{S^2 \times pt} \Omega = \pi r^2 + \varepsilon.$$

Since this is true for all $\varepsilon > 0$ we must have $r \ge 1$. □

What we have used here from the theory of J-holomorphic curves is the existence of a curve in class $A = [S^2 \times pt]$ through an arbitrary point in $S^2 \times T^{2n-2}$. It is easy to check that when J equals the product almost complex structure J_{split} there is exactly one such curve through every point. For in this case the two projections are holomorphic so that every J_{split}-holomorphic curve in $S^2 \times \mathbf{R}^{2n-2}$ is the product of curves in each factor. But the curve in T^{2n-2} represents the zero homology class and so must be constant. Now, the basic theory of J-holomorphic curves is really a deformation theory: if you know that curves exist for one J you can often prove they exist for all other J.[4] That is exactly what we need here. Here is an outline of how this works. For more details see [MS1] as well as the Park City lectures by Salamon.

Fredholm theory

Let $\mathcal{M}(A, \mathcal{J})$ be the space of all pairs (u, J), where $u : (S^2, j) \to (M, J)$ is J-holomorphic , $u_*([S^2]) = A \in H_2(M)$, and $J \in \mathcal{J}(\omega)$. One shows that a suitable completion of $\mathcal{M}(A, \mathcal{J})$ is a Banach manifold and that the projection

$$\pi : \mathcal{M}(A, \mathcal{J}) \to \mathcal{J}$$

is Fredholm of index

$$2(c_1(A) + n),$$

where $c_1 = c_1(TM, J)$. In this situation one can apply an infinite dimensional version of Sard's theorem (due to Smale) that states that there is a set $\mathcal{J}_{reg}$ of second category in $\mathcal{J}$ consisting of regular values of π. Moreover by the implicit function theorem for Banach manifolds the inverse image of a regular value is a smooth manifold of dimension equal to the index of the Fredholm operator. Thus one finds that for almost every J

$$\pi^{-1}(J) = \mathcal{M}(A, J)$$

is a smooth manifold of dimension $2(c_1(A) + n)$. Moreover, by a transversality theorem for paths, given any two elements $J_0, J_1 \in \mathcal{J}_{reg}$ there is a path $J_t, 0 \le t \le 1$, such that the union

$$W = \cup_t \mathcal{M}(A, J_t) = \pi^{-1}(\cup_t J_t)$$

[4] An *existence* theory for J-holomorphic curves had to wait until the recent work of Donaldson and Taubes.

is a smooth (and also oriented) manifold with boundary

$$\partial W = \mathcal{M}(A, J_1) \cup -\mathcal{M}(A, J_0).$$

It follows that the evaluation map

$$ev_J : \mathcal{M}(A, J) \times_G S^2 \to M, \quad (u, z) \mapsto u(z),$$

is independent of the choice of (regular) J up to oriented bordism.[5] (Here $G = PSL(2, \mathbf{C})$ is the reparametrization group and has dimension 6.) In particular, *if* we could ensure that everything is compact and *if* we arrange that ev maps between manifolds of the same dimension then the degree of this map would be independent of J.

In the case of the nonsqueezing theorem we are interested in looking at curves in the class $A = [S^2 \times pt]$ in the cylinder $S^2 \times T^{2n-2}$. Thus $c_1(A) = 2$ since the normal bundle to $S^2 \times pt$ is trivial (as a complex vector bundle with the induced structure from J_{split}.) Thus

$$\begin{aligned} \dim(\mathcal{M}(A, J) \times_G S^2) &= 2(c_1(A) + n) - 6 + 2 \\ &= 4 + 2n - 6 + 2 = 2n = \dim(M). \end{aligned}$$

Further when $J = J_{split}$ the unparametrized moduli space $\mathcal{M}(A, J)/G)$ is compact (it is diffeomorphic to T^{2n-2}) and ev_J has degree 1. It is also possible to check that J_{split} is regular. So the problem is to check that compactness holds for all J. If so, we would know that ev_J has degree 1 for all regular J, ie there is at least one J-holomorphic curve through every point.

Compactness

This is the most interesting part of the theory and leads to all sorts of new developments such as the connection with stable maps and Deligne–Mumford compactifications.

We proved the following lemma at the end of Lecture 4. It is the basic reason why spaces of J-holomorphic curves can be compactified.

Lemma 5.1 *If u is J-holomorphic for some ω-compatible J then*

$$\|u\|_{1,2} = \int_{S^2} u^*(\omega) = g_J\text{-}area\ of\ (\operatorname{Im} u).$$

[5] Two maps, $e_i : M_i \to X$ for $i = 1, 2$, are said to be *oriented bordant* if there is an oriented manifold W with boundary $\partial W = M_1 \cup (-M_2)$ and a map $e : W \to X$ that restricts to e_i on the boundary component M_i. Often the compactness that is needed to get any results from this notion is built into the definition. For example, if all manifolds M_1, M_2, W are compact and if M_1, M_2 have no boundary then bordant maps e_i represent the same homology class.

Here $\|u\|_{1,2}$ denotes the L^2-norm of the first derivative of u. If we just knew a little more we would have compactness by the following basic regularity theorem for solutions of elliptic differential equations.

Lemma 5.2 *If $u_n : S^2 \to M$ are J-holomorphic curves such that for some $p > 2$ and $K < \infty$*

$$\|u_n\|_{1,p} \leq K,$$

then a subsequence of the u_n converges uniformly with all derivatives to a J-holomorphic map u_∞.

It follows from the above two lemmas that if $u_n \in \mathcal{M}(A, J)$ is a sequence with no convergent subsequence then the size of the derivatives du_n must tend to infinity. In other words

$$c_n = \max_{z \in S^2} |du_n(z)| \to \infty.$$

By reparametrizing by suitable rotations we can assume that this maximum is always assumed at the point $0 \in \mathbf{C} \subset \mathbf{C} \cup \infty = S^2$. The claim is that as $n \to \infty$ a "bubble" is forming at 0, i.e. the image curve is breaking up into two or more spheres. To see this analytically consider the reparametrized maps $v_n : \mathbf{C} \to M$ defined by

$$v_n(z) = u_n(z/c_n).$$

Then

$$|dv_n(0)| = 1 \quad \text{and} \quad |dv_n(z)| \leq 1, z \in \mathbf{C}.$$

Therefore by Lemma 5.2 a suitable subsequence of the v_n converge to a map $v_\infty : \mathbf{C} \to M$. Moreover because the energy (or symplectic area) of the image of the limit v_∞ is bounded (by $\omega(A)$), the image points $v_\infty(z)$ converge as $z \to \infty$. In other words v_∞can be extended to a map $v_\infty : S^2 \to M$. (Here we are applying a removable singularity theorem for J-holomorphic maps $v : D^2 - \{0\} \to M$ that have finite area.)

Usually the image curve $C_\infty = v_\infty(S^2)$ will be just a part of the limit of the set-theoretic limit of the curves $C_n = u_n(S^2)$. What we have done in constructing C_∞ is focus on the part of C_n that is the image of a very small neighborhood of 0, and there usually are other parts of C_n (separated by a "neck".) Thus typically the the curves C_n converge (as point sets) to a union of several spheres, and the bubble C_∞ is just one of them. (Such a union of spheres is often called a "cusp-curve" or reducible curve.) It can happen that the bubble C_∞ is the whole limit of the C_n. But in this case one can show that it is possible to reparametrize the original maps u_n so that they converge. In other words, the u_n converge in the space $\mathcal{M}(A, J)/\mathrm{G}$ of *unparametrized* curves. For example, if we started with a sequence of the form $u \circ \gamma_n$ where γ_n is a nonconvergent sequence in the reparametrization group G, then the effect of the reparametrizations v_n of u_n is essentially to undo the γ_n. More precisely, v_n would have the form $u \circ \gamma_n'$ where the γ_n' do converge in G.

This argument (when made somewhat more precise) shows that the only way the unparametrized moduli space $\mathcal{M}(A, J)/G$ can be noncompact is if there is a reducible J-holomorphic curve in class A consisting of several nontrivial spheres that represent classes $A_1, A_2, \ldots$. Since each $\omega(A_i) > 0$, this is possible only when $\omega(A)$ is not minimal (among all positive values of ω on spheres). In particular, when $M = S^2 \times T^{2n-2}$, $\pi_2(M)$ is generated by A. Therefore there is no reducible J-holomorphic curve in class A, and $\mathcal{M}(A, J)/G$ is always compact. Hence the space $\mathcal{M}(A, J) \times_G S^2$ is also always compact. This completes our sketch of the proof of the nonsqueezing theorem.

References

[A] V.I. Arnold, *Mathematical methods in classical mechanics.* Springer-Verlag, Berlin (1978).

[AL] M. Audin and F. Lafontaine, (ed.) *Holomorphic curves in symplectic geometry*. Progress in Mathematics **117**, Birkhäuser, Basel (1994).

[BT] R. Bott and G. Tu, *Differential Forms in Algebraic Topology.* Graduate Texts in Mathematics, **82**. Springer-Verlag, Berlin (1982).

[CFHW] K. Cieliebak, A. Floer, H. Hofer and K. Wysocki, Applications of symplectic homology II: stability of the action spectrum, *Math Zeitschrift*

[EH] I. Ekeland and H. Hofer, Symplectic topology and Hamiltonian dynamics. *Mathematische Zeitschrift*, **200** (1989), 355–78.

[E] Y. Eliashberg, A theorem on the structure of wave fronts and its applications in symplectic topology. *Functional Analysis and Applications*, **21**, (1987) 65–72.

[FH] A. Floer and H. Hofer, Symplectic homology I: Open sets in $\mathbf{C}^n$. *Mathematische Zeitschrift*, **215** (1994), 37–88.

[G] M. Gromov, Pseudo holomorphic curves in symplectic manifolds. *Inventiones Mathematicae*, **82** (1985), 307–47.

[H] H. Hofer, Estimates for the energy of a symplectic map. *Commentarii Mathematici Helvetici*, **68** (1993), 48–72.

[HZ] H. Hofer and E. Zehnder, *Symplectic capacities*. Birkhauser, Basel (1994).

[LM] F. Lalonde and D. McDuff, Hofer's L^∞ geometry: geodesics and stability, I, II. *Invent. Math.* **122** (1995), 1–33, 35–69.

[MS1] D. McDuff and D.A. Salamon, *J-holomorphic curves and quantum cohomology.* University Lecture Series, (1994), American Mathematical Society, Providence, RI.

[MS2] D. McDuff and D.A.Salamon, *Introduction to Symplectic Topology*, OUP (1995).

[M] J.K. Moser, On the volume elements on manifolds. *Transactions of the American Mathematical Society*, **120** (1965), 280–96.

[P] L. Polterovich, Symplectic aspects of the first eigenvalue, preprint (1997)

[V] C. Viterbo, Symplectic topology as the geometry of generating functions. *Mathematische Annalen*, **292** (1992), 685–710.