Symplectic Topology and Capacities[1]

talk given in Princeton to celebrate the 250th Anniversary of the founding of the university

Dusa McDuff

Revised June 26, 1997

I am going to talk about symplectic topology, a field that has seen a remarkable development in the past 15 years. Let's begin with what was classically known – which was rather little. We start with a symplectic form ω, that is, a closed 2-form which is nondegenerate. This last condition means that ω is defined on an even dimensional manifold M^{2n} and that its wedge with itself n times is a top dimensional form which never vanishes:

$$\omega \wedge \ldots \wedge \omega \neq 0 \qquad \text{everywhere.}$$

For example, in R^{2n} we can take

$$\omega_0 = \sum_{i=1}^{n} dx_i \wedge dy_i \tag{1}$$

where $x_1, y_1 \ldots, x_n, y_n$ are the coordinates of R^{2n}.

The structure given by this basic form ω_0 first arose in Hamilton's formulation of classical mechanics in the mid 19th century, where he displayed the symplectic form as the mediator between the energy function and the time evolution equations of the system. More precisely, any function $H : R^{2n} \to R$ defines a vector field X_H via the equation:

$$\omega_0(X_H, \cdot) = dH(\cdot).$$

The flow ϕ_t generated by X_H is a family of transformations from R^{2n} to itself that has the following interpretation: for each point p in the phase space R^{2n}, the image point $\phi_t(p)$ represents the state at time t of a system that was in state p at time 0. Moreover the number $H(p)$ is the energy of the system in state p. It is not hard to check that the transformations ϕ_t in this flow preserve the symplectic form, that is $\phi_t^*(\omega_0) = \omega_0$. Transformations (or diffeomorphisms) with this property are called *symplectomorphisms*. Obviously it is of great interest to understand their geometric properties, and, if possible, to find ways of distinguishing them from arbitrary smooth diffeomorphisms.

The first theorem in the subject is due to Darboux:

> *any symplectic form ω on any manifold M^{2n} is locally diffeomorphic to ω_0.*

[1] published in *Prospects in Mathematics* ed Rossi, Americal Mathematical Society (1999)

In other words there are coordinates in a neighborhood of any point so that the form looks like (1). This is an example of a *local* to *global* phenomenon that occurs everywhere in symplectic topology. By this I mean the following. If first you look linearly, at a tangent space for example, then up to a linear coordinate change there is only one linear nondegenerate 2-form, namely (1). Now, Darboux's theorem tells you that the uniqueness that you have at a point extends to uniqueness in a neighborhood. One can also interpret this in terms of a passage from linear to nonlinear: a linear phenomenon persists in the nonlinear situation. [2]

The second basic theorem is Moser's stability theorem concerning families of forms. Suppose we have a family of symplectic forms $\omega_t, 0 \leq t \leq 1$, on some closed manifold (compact without boundary.) Let's also suppose that the cohomology class $[\omega_t] \in H^2$ of these forms is fixed, so that the variation is just by exact forms: $\omega_t = \omega_0 + d\sigma_t$. Then the theorem is that the forms are all basically the same. That is, there exists a family of diffeomorphisms $\psi_t, 0 \leq t \leq 1$, of the manifold starting at the identity (that is $\psi_0 = \mathrm{id}$) such that $\psi_t^* \omega_t = \omega_0$ for all t.

What this says is there is no interesting deformation theory of symplectic forms: if you try to change the forms within a fixed cohomology class you can't do it. The proof involves finding a (time-dependent) vector field X_t which integrates to give the isotopy ψ_t. For this to work the manifold must be closed, since otherwise the flow ψ_t might run off it at infinity. However, there is an extension of this theorem to open manifolds if you put in some controls at infinity.

So that is what was known: not much at this point. The story that I really want to tell you begins with Gromov's work [10] in the mid 80s where he introduced elliptic methods into symplectic topology and proved a whole array of wonderful results.[3] First I shall state some of them and then give you an idea of how to prove them.

Gromov's results

(I) The first result is known as the **nonsqueezing theorem**. The basic question here is to understand the possible shapes of the image of a standard ball under a symplectic transformation. Gromov picked out as decisive the question of when a ball can be symplectically mapped inside a cylinder. Here the ball is the standard (compact) ball of radius r sitting in standard Euclidean space, $B^{2n}(r) \subset (R^{2n}, \omega_0)$, and the cylinder $B^2(\lambda) \times R^{2n-2}$ is the product of a 2-disc of radius λ with Euclidean $(2n-2)$-space. It is important here that the latter is a symplectic product; that is the symplectic form on the first factor is the

[2] For proofs of this and other results, and further references, the interested reader can consult McDuff and Salamon [21].

[3] Another important theme in the development of symplectic topology came with Conley and Zehnder's use of variational methods in their proof in 1983 of Arnold's conjecture for the torus: see [3] and Hofer–Zehnder's book [11].

area form $dx_1 \wedge dy_1$, and that on the second is the basic form in the remaining variables.

Gromov's theorem tells us that there is such a symplectic embedding if and only if $r \leq \lambda$:

$$B^{2n}(r) \hookrightarrow B^2(\lambda) \times R^{2n-2} \qquad \Longleftrightarrow \qquad r \leq \lambda.$$

Obviously if $r \leq \lambda$, you can just include the ball inside the cylinder. The force of this theorem is that, if $r > \lambda$, it is impossible to take the round ball of radius r and squeeze it symplectically to fit it into the cylinder, no matter how long you make it: see Figure 2. This is a very clean statement of what a symplectic map cannot do, and shows there is some kind of fatness in the ball that you can't squeeze.[4]

Observe also that this property distinguishes symplectic from volume-preserving transformations. Every symplectic map preserves volume, since volume is given by the form $\omega_0^n = \omega_0 \wedge \ldots \wedge \omega_0$. Moreover, it is not hard to see that one can map *any* ball into *any* cylinder if all one requires is that the volume form is preserved, since maps of the form

$$(x_1, y_1, x_2, y_2, \ldots) \mapsto (\mu x_1, \mu y_1, \frac{1}{\mu} x_2, \frac{1}{\mu} y_2, \ldots)$$

preserve volume. Thus the nonsqueezing theorem shows that symplectic transformations are much more limited than volume-preserving ones. In particular, it is not possible to approximate an arbitrary volume-preserving transformation by a symplectic one in the uniform (or C^0) topology. Incidentally, one can also think of this result as another instance of the linear to nonlinear phenomenon: it is very easy to check that there is no *linear* (or affine) symplectic transformation of Euclidean space that squeezes a ball into a thinner cylinder, and now we see that this cannot be done by an arbitrary nonlinear symplectic transformation either.

(II) Another very nice result is the solution of the "camel" problem. This problem is the following: imagine you are in Euclidean space R^{2n} minus a wall W (Fig. 1).

INSERT FIGURE 1

This wall consists of the hyperplane $x_1 = 0$ with a hole of radius 1 in it and is defined by the equations

$$x_1 = 0\ , \quad \sum_{i>1} {x_i}^2 + \sum_{i\geq 1} {y_i}^2 \geq 1.$$

[4]Lalonde and McDuff [13] have recently shown that a similar result is true for embeddings $B^{2n}(r) \to B^2(\lambda) \times M^{2n-2}$ for any symplectic manifold M, closed or not.

Now imagine putting a big ball $B^{2n}(r), r > 1$ on the left side of Fig. 1, and ask: is there a way of deforming this ball symplectically so as to take it through the hole and over to the other side? One might first of all see whether it is possible to go through the hole preserving volume. But now you don't see the roundness in the ball: you can make the ball long and thin preserving volume, and then slide it through. The nonsqueezing theorem says that this manoever is impossible symplectically. However it still might be possible to get the ball through the hole in some other more complicated way — all we need now is to squeeze the ball in one hyperplane. Nevertheless, as you might guess, the answer to this problem is no. There is some sort of obstruction that prevents the ball from going through.

(III) Another result in which people are very interested these days is the uniqueness of the symplectic structure on R^4 and CP^2. The preceding results work in all dimensions but this one is definitely 4-dimensional. What Gromov proved is the following. Suppose we have a 4-dimensional symplectic manifold (X, ω) which, outside a compact set, is symplectomorphic to R^4 minus a ball. Then, provided that $\pi_2 X = 0$, this symplectomorphism extends to a symplectomorphism from the whole of X onto R^4. Note that the conclusion includes in it a statement of what the diffeomorphism type of X is, while all that is assumed is some homotopic theoretic knowledge and information about the symplectic structure on X at infinity.

The above result is actually equivalent to a uniqueness result for CP^2, that can be stated as follows. Suppose you have a compact symplectic manifold (Z, ω) whose second homotopy group $\pi_2(Z)$ is generated by a symplectically embedded 2-sphere S with self intersection $+1$. (This means that $S \cdot S = 1$ or, equivalently, that the normal bundle to S has Chern number 1.) Then this manifold (Z, ω) is symplectomorphic to CP^2 with its standard symplectic structure, appropriately scaled. To see the equivalence, observe that if you remove a neighborhood of the sphere S from Z you get a manifold whose boundary is symplectomorphic to the standard 3-sphere in R^4 and which can therefore be extended to a symplectic manifold X that looks like R^4 at infinity. Conversely, we can compactify X by removing a standard collar at infinity and then attaching a 2-sphere S to produce such a Z.

Note that Gromov can only establish uniqueness under the assumption that the manifold Z contains a symplectic 2-sphere S. One of the triumphs of the recent work of Taubes [25] on Seiberg–Witten theory is to remove this assumption. Thus one now knows that a symplectic 4-manifold that is diffeomorphic to CP^2 is actually symplectomorphic to CP^2.

(IV) As a counterpoint to the above uniqueness results, Gromov showed that when $n > 1$ there exists an exotic symplectic structure on R^{2n}. In other words, there is a structure ω that does not live in standard Euclidean space(R^{2n}, ω_0)

in the sense that there is no embedding

$$\psi : R^{2n} \to R^{2n}$$

such that $\psi^*(\omega_0) = \omega$. Here the difficulty is not so much in constructing something that you think should be exotic, but in proving that it is exotic.

Gromov's criterion hinges on properties of Lagrangian submanifolds of Euclidean space. An n-dimensional submanifold L of a symplectic manifold (M^{2n}, ω) is called Lagrangian if the restriction of ω to L is identically zero. This implies that if ω is exact (i.e. $\omega = d\lambda$), the 1-form λ restricts to a closed form λ_L on L since $d\lambda_L = \omega|_L = 0$. Hence λ defines a de Rham cohomology class $[\lambda_L] \in H^1(L)$. If in turn this class $[\lambda_L]$ is zero (or, equivalently, if $\lambda|_L$ itself is exact, $\lambda_L = dF$), the Lagrangian submanifold L is said to be *exact*. Gromov's main result is that *no* closed Lagrangian submanifold in the standard (R^{2n}, ω_0) is exact. He then constructed a symplectic structure ω on R^{2n} that contained an exact closed Lagrangian submanifold L, and concluded that ω had to be exotic since otherwise L would give rise to a forbidden exact Lagrangian submanifold in standard Euclidean space.

Nobody has yet managed to do much more with this problem, for example, showing that there is more than one exotic structure on Euclidean space, or finding an exotic structure that is standard at infinity (by (III) above this would have to live in dimension at least 6.)

J-holomorphic curves

I'll now try to give you an idea of how Gromov proved these results. Observe first that if you have a symplectic manifold (M, ω) then you can always find an almost complex structure J on the manifold. That's an automorphism of the tangent bundle TM which, like multiplication by i, satisfies the equation $J^2 = -\mathrm{Id}$. Moreover you can require this almost complex structure to be related to ω by the taming or positivity) condition:

$$\omega(v, Jv) > 0, \quad \text{whenever} \quad v \in T_p, v \neq 0, \tag{2}$$

so that there is an associated Riemannian metric on M, given by the symmetrization of $\omega(v, Jv')$. It is easy to check that the set of J satisfying these conditions (these are called ω-tame J) is nonempty and contractible. Therefore one can try to find invariants of (M, ω) by looking for invariants of the almost complex manifold (M, J) that do not depend on the choice of ω-tame J.

Gromov constructed his invariants by looking at maps of Riemann surfaces[5] into the almost complex manifold

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

[5] A Riemann surface is a compact 2-manifold Σ with complex structure j.

which satisfy the generalized Cauchy-Riemann equation:

$$du \circ j = J \circ du. \tag{3}$$

Such maps u are called J-holomorphic curves. Because equation (3) is elliptic, its solution spaces have very nice properties. For example, for generic ω-tame J, the space of solutions in a fixed homology class A is a finite dimensional manifold $\mathcal{M}(A, J)$. It's not true that these solution (or moduli) spaces are compact because curves can degenerate. But the positivity condition (2) allows one to control and understand the degenerations that occur and hence describe the compactified moduli spaces of curves. Moreover, because the space of ω-tame almost complex structures J is path-connected, any two generic J can be connected by a path J_t such that the set of all J_t-holomorphic curves for $t \in [0, 1]$ forms a cobordism between the moduli spaces at $t = 0, 1$. In many situations, this cobordism is compact, which means that any property of the moduli space $\mathcal{M}(A, J)$ which is a cobordism invariant depends only on the underlying symplectic form ω and not on J. For example, the Gromov invariants are defined by counting the number of J-holomorphic curves (with given genus and in a given homology class) that go through a fixed number of points or cycles.

I'll now show how these ideas can be used to prove the nonsqueezing theorem. Consider the ball $B^{2n}(r)$ with its standard complex structure J_0. First observe that if you have a J_0-holomorphic curve S that goes through the center of the ball and is properly embedded (so it goes all the way to the boundary) then the symplectic area of S is at least πr^2. In fact, J_0-holomorphic curves in R^{2n} are just complex curves in the usual sense, and so are minimal surfaces with respect to the usual metric g_0. Moreover, it is easy to see that the symplectic area $\int_S \omega_0$ of a complex curve S is just its area with respect to the usual metric. Hence the above result holds because of the well known fact that the surface of minimal area through the center of $B^{2n}(r)$ is the flat disc, which has an area πr^2.

So now take this ball and suppose it is embedded in a cylinder. On the image you have the pushforward of J_0, which you can always extend to an ω_0-tame J that is standard near the boundary of the cylinder. Then J extends to the compact manifold $S^2 \times R^{2n-2}$ obtained by closing up the cylinder. Now, if J is the product almost complex structure on $S^2 \times R^{2n-2}$ there is a flat J-holomorphic 2-sphere through every point, that is unique modulo reparametrization. Moreover one can show that this product J is generic in the requisite sense (that is, it is a regular point for an appropriate operator). Hence the Gromov invariant that counts the number of spheres in the class $[S^2 \times pt]$ through a fixed point takes the value $+1$. The theory outlined above implies that this number is independent of the choice of J. Therefore, for any ω-tame J on the product and any point p, there is precisely one J-holomorphic 2-sphere through p in the homology class of $S^2 \times pt$, when these are counted with appropriate signs and modulo parametrization. In particular, for our special J which equals the pushforward of the standard structure on the image of the ball, there is at least one J- holomorphic sphere Σ going through the image of the center of the ball. Note that the symplectic area of Σ is determined by its homology class and so is $\pi(\lambda+\varepsilon)^2$ for some arbitrarily small $\varepsilon > 0$. (This ε appears because we slightly increase the size of the cylinder when we close it up.) Now look at the inverse image S of Σ in the ball. (see Figure 2).

INSERT FIGURE 2 HERE

By our previous result, the area of S is at least πr^2. Moreover, since symplectomorphisms preserve area, this has to be strictly less than the area of Σ, which is $\pi(\lambda+\varepsilon)^2$. Since this is true for all $\varepsilon > 0$ we find that $r \leq \lambda$.

A similar but more elaborate version of this idea proves the camel theorem: one just shows that there are a lot of J-holomorphic disks with boundaries on the hole and area at most π that give obstructions to putting the ball through. Finally, for the uniqueness of structure theorem one shows that there are so many spheres that one can somehow use them as coordinates to construct a diffeomorphism from one manifold to another. The argument about Lagrangian submanifolds L is somewhat more complicated but it also uses the behavior of J-holomorphic curves (in this case, discs with boundary on L) as its technical base.

Since their introduction in the mid-80s these elliptic methods have been very influential, especially when combined with Floer's approach to Morse theory: see for example Floer [8]. A detailed description of the construction of genus 0 Gromov invariants, together with a discussion of their relation to quantum cohomology and problems in enumerative geometry, may be found in Ruan–Tian [23] and McDuff–Salamon [20]. The higher genus case is treated in Ruan–Tian [24].

Capacities

One set of applications of Gromov's ideas revolves around the notion of symplectic capacity that was formalized by Ekeland and Hofer [7]. A symplectic capacity c is a function of subsets in Euclidean space (or more generally, in arbitrary symplectic manifolds of a fixed dimension) taking values in $[0,\infty]$, and is a measure of how big this subset is. It satisfies these conditions:

i) If $(U,\omega) \subset (V,\omega)$ then $c(U,\omega) \leq c(V,\omega)$, that is, c is monotonic.

ii) If $\phi : (U,\omega) \to (V,\omega')$ is a diffeomorphism which is conformally symplectic, that is $\phi^*(\omega') = \lambda\omega$ for some constant λ, then $\lambda c(U,\omega) = c(V,\omega')$.

iii) The third condition is a normalization condition which specifies that the capacity of a ball with its standard structure is positive: $c(B^{2n}(r)) > 0$.

iv) Finally you need something which tells you that you are not measuring a volume; you want something that is definitely two dimensional. This is formulated as: $c(B^2(r)) \times R^{2n-2} < \infty$.

An example of capacity is this:

$$c_G(M) = sup\{\pi r^2 : B^{2n}(r) \hookrightarrow M \text{ symplectically}\}$$

Here you are just measuring the size of the largest ball which embeds symplectically in M. This obviously satisfies i) and ii). Condition iii) follows from Darboux's theorem: there are small symplectic balls in M. Finally, the non-squeezing theorem gives (iv):

$$c_G(B^2(\lambda) \times R^{2n-2}) = \pi\lambda^2$$

I shall call this capacity defined by embedding balls the *Gromov capacity* c_G. There are other symplectic capacity functions developed by Ekeland–Hofer and Hofer–Zehnder (for precise references see [21, 11]) that arise from variational problems, but I won't go into those.

One important use of symplectic capacities, noted by Eliashberg [6] and Ekeland–Hofer, is to give a topological criterion for deciding whether a diffeomorphism is symplectic or not.[6] Let's suppose that you have a capacity (such as c_G) with the specific normalization condition that the capacity of the ball of radius R is the same as that of the cylinder of radius R:

$$c(B^{2n}(R)) = c(B^2(R) \times R^{2n-2}).$$

Then the statement is:

> *a diffeomorphism ϕ of R^{2n} is symplectic if and only if ϕ preserves capacity: that is, $c(\phi(W)) = c(W)$ for all open subsets W of R^{2n}.*

The proof of this statement is based on the fact that capacity is invariant under rescaling. In order to establish that a map ϕ is symplectic one just has to show that its derivative at each point is a linear symplectic map. But when you take a derivative at a point you look at a little piece of space and magnify it until in the limit you get the derivative. Thus taking the derivative is just the limit of a rescaling process, and so it is not hard to see that if a map ϕ preserves capacity its derivative also does. Now it is just a matter of linear algebra to show that the resulting linear map is symplectic: in fact, any linear map L that preserves capacity is symplectic or antisymplectic,[7] and it is not hard to rule out the antisymplectic case.

[6] This criterion is called topological because it makes no use of the derivative, and hence implies that the group of symplectomorphisms is closed in the group of all diffeomorphisms with respect to the uniform topology.

[7] that is $L^*(\omega_0) = \pm\omega_0$

The symplectic topology of Euclidean space

Because Darboux's theorem shows that any symplectic form is locally like the standard structure in Euclidean space, there is a temptation to to think that there is nothing to be said about local symplectic geometry, all forms are locally the same and that is that. But it turns out that the standard symplectic structure has a whole lot of interesting properties.

Here I'll mention some things that can be proved using the idea of capacity. First, consider an open polydisc $P(r_1, r_2, ..., r_n)$ — the product of n open 2-dimensional discs of radii $r_1, \dots, r_n$, with normalization condition $r_1 \le r_2 \le \dots \le r_n$. Gromov asked: are these symplectic manifolds (with the standard structure inherited from Euclidean space) all different? The point of his question is this. A closed (smooth) domain in Euclidean space has a characteristic Hamiltonian flow on its boundary and you can easily distinguish among closed polydiscs by looking at this boundary flow.[8] If you look at open polydiscs you don't have that flow, but you might suspect that the flow on the boundary would leave enough trace inside to force the open domains to be different.

In fact, when $n = 2$ it is easy to show that this suspicion is true by using the Gromov capacity. Since

$$c_G(P(r_1, r_2)) = \pi r_1^2, \qquad \mathrm{vol}\,(P(r_1, r_2)) = \pi^2 r_1^2 r_2^2,$$

one can distinguish two 4-dimensional polydiscs using these two invariants. In other words, $P(r_1, r_2)$ is symplectomorphic to $P(r_1', r_2')$ if and only if $r_1 = r_1', r_2 = r_2'$. But when $n \ge 3$ these two invariants are no longer enough. However Floer and Hofer developed a theory of symplectic homology (a kind of Floer homology built out of action functionals coming from Hamiltonians) that extends the notion of capacity and, using this, Floer–Hofer–Wysocki showed in [9] that the polydiscs are all distinct. A similar result holds for ellipsoids.

Another kind of question is this, consider the set $Emb(B^{2n}(r), U)$ of all symplectic embeddings of a ball into the set U – here U could be a domain in R^{2n} or any symplectic manifold. Capacity tells you the maximum size of a ball that you can embed, but now you can ask whether this space of embeddings is connected. This can be thought of as an extension of the camel problem, because the camel problem tells you that if the ball embedded on one side of a wall with a hole has radius larger than that of the hole, this space is disconnected. One could ask:

> *for which domains in Euclidean space is this embedding space disconnected?*

[8] This flow is simply the Hamiltonian flow – generated by the vector field X_H defined earlier – of a smooth function H that is constant on the boundary. The fact that a closed polydisc has nonsmooth boundary causes no real problem here. Any symplectomorphism ϕ that takes one closed polydisc P to another P' maps a sequence of smooth approximations to the boundary ∂P into a similar approximating sequence for $\partial P'$. But a symplectomorphism defined only on an open polydisc need not extend to the boundary and so need not take an approximating sequence to ∂P into an approximating sequence for $\partial P'$.

Another interesting question to which I don't have an answer is this:

> *is there a closed manifold M such that this embedding space is disconnected?*

It is likely that the answer depends on r: the embedding space might always be connected for small enough r but then become disconnected for large r, much as happens in the camel problem. So far, the only results on this question are in dimension 4. For example I showed using J-holomorphic curves that this space is always connected when the target U is a 4-ball or more generally a star-shaped region.[9] There is also one relevant result of Floer–Hofer–Wysocki that considers what happens when the domain and target spaces both have corners: they prove that $Emb(P(r_1, r_2), P(1, 1))$ is disconnected if $r_1, r_2 < 1$ but $r_1^2 + r_2^2 > 1$. The two embeddings which are not isotopic are the obvious inclusions (see Fig. 3). INSERT FIGURE 3 HERE

You see that to get from one image to the other you have to rotate a quarter turn. However, using symplectic homology one can show that if $r_1^2 + r_2^2 > 1$ the corners of the domain are too large for it to be able to be turned in this way. (The corners are detected using products in the homology theory.)

Symplectic deformations and isotopies

I have now told you something of what is known about the symplectic topology of Euclidean space. Now, I would like to talk about symplectic structures on general manifolds. Remember that Moser's stability theorem says that if you have a family of symplectic forms $\omega_t, 0 \le t \le 1$, on a closed manifold M with constant cohomology class then this family is just gotten by a isotopy ψ_t of the underlying manifold. Thus $\psi_t^*(\omega_t) = \omega_0$ for all t. Such a family ω_t is called an *isotopy*. There is another related notion where you have a family of symplectic forms as above, but you don't make any assumption on the cohomology class, instead allowing it to vary. That is called a *deformation*. And you might ask:

> *Suppose you have a deformation $\{\omega_t\}$ such that $[\omega_0] = [\omega_1]$, is the family homotopic through deformations with fixed endpoints to an isotopy?*

That is, you ask if you can get anything new in a given cohomology class if you allow yourself to make deformations.

Examples in 6-dimensions

[9] Some recent work in progress of Damiano and Laudenbach [4] may shed light on these questions.

There are certain examples which tell you that in dimension six and above deformations are different from isotopies. I can just briefly describe them. Take $M = S^2 \times S^2 \times T^2$ and let the symplectic form

$$\omega = \sigma_0 \oplus \sigma_1 \oplus \sigma_2, \tag{4}$$

be the sum of the area forms on each factor. What is very important is that you insist that the area of the first factor is the same as the area of the second factor. Suppose you take a diffeomorphism ϕ which twists the first factor as follows. Let (z, w, s, t) be the coordinates for M, where $z, w \in S^2, s, t \in T^2 = R^2/Z^2$. Then define

$$\phi(z, w, s, t) = (z, R_{z,t}w, s, t)$$

where $R_{z,t}$ is rotation of the sphere by $2\pi t$ about the axis through z. (Here we think of the 2-sphere S^2 embedded in 3-space R^3 in the usual way so that the axis through z is just the radius $0z$.) Then the claim is that the pull back $\phi^*\omega_0$ of ω_0 by this twisting ϕ is deformation equivalent to ω_0, but not isotopic to it. Interestingly enough, if you put a constant $\lambda > 1$ in (4)

$$\omega^\lambda = \lambda\sigma_0 \oplus \sigma_1 \oplus \sigma_2,$$

so that the first sphere is bigger than the second one, then the forms $\phi^*(\omega^\lambda)$ and ω^λ are actually isotopic. (This immediately implies that $\phi^*(\omega)$ and ω are deformation equivalent.) But when $\lambda = 1$ you can't undo the effect of the twisting. In fact, you can see this twisting symplectically by looking at the family of spheres $z \mapsto (z, w_0, s_0, t)$ for $t \in S^1$. These spheres are J-holomorphic when J is the product almost complex structure J_{split}, and persist as J varies among all ω-tame almost complex structures. In other words for each such J one has a map $S^2 \times S^1 \to M$ with image equal to a family of J-holomorphic spheres. Moreover, for each J, the composite of this map with the projection pr_2 onto the second sphere is bordant to a constant map.[10] However, if one looks at J that are tamed by the twisted form $\phi^*(\omega)$ the story is different. For example if one takes $J = \phi^*(J_{split})$ the family is

$$z \mapsto \phi^{-1}(z, w_0, s_0, t) = (z, R_{z,-t}w_0, s_0, t),$$

and so, composition with projection onto the second sphere gives the map

$$e_J : S^2 \times S^1 \to S^2, \qquad (z, t) \mapsto R_{z,-t}(w_0).$$

Using a version of the Hopf invariant, one easily sees that this map is not bordant to a constant map. A similar map e_J can be defined for any J that is tamed by a

[10] Two maps $f_i : X_i \to Y$ are said to be bordant if there is a manifold W with boundary equal to $X_0 \cup X_1^{opp}$ and a map $g : W \to Y$ such that g restricts to f_i on X_i for $i = 0, 1$. Here I am assuming that all the manifolds X_i, W are oriented and compact and am writing X^{opp} for X with the opposite orientation. For example, homotopic maps are bordant, though not conversely.

form isotopic to $\phi^*(\omega)$, and it follows from the general theory that all such maps are bordant. Hence ω cannot be isotopic to $\phi^*(\omega)$. The argument breaks down when $\lambda > 1$ since in this case the moduli space of curves we are considering fails to be compact. This implies that our method only gives a noncompact bordism $g : W \to S^2$ between the maps e_J as J varies. Since all maps are noncompactly bordant, the invariant disappears.

I would also like to mention that there are some examples of Ruan [22] which give non deformation equivalent symplectic structures on 6 dimensional manifolds of the form $M^6 = X^4 \times S^2$. The idea is that you take two non diffeomorphic but homotopy equivalent symplectic four manifolds X, Y. When you stabilize them (cross with S^2) they become diffeomorphic. However, the pattern of J-holomorphic curves they contain persists under stabilization and therefore they are often symplectically distinct. For example, one can take X equal to the complex projective plane blown up at 8 points and Y equal to the Barlow surface. Then $X \times S^2$ contains 8 families of J-holomorphic spheres corresponding to the blown up points, while $Y \times S^2$ contains no J-holomorphic spheres except for those in class $[pt \times S^2]$.

The 4-dimensional case

Finally, something about 4-dimensional manifolds. As already mentioned, Taubes–Seiberg–Witten theory allows a significant strengthening of Gromov's uniqueness results. You see, Gromov had to assume in his statement that some structure was known – either there was a symplectically embedded 2-sphere or in the case of Euclidean space the structure was known at infinity. But now this is unnecessary, at least in the compact case, since Taubes has shown that the Seiberg-Witten invariants of a closed symplectic 4-manifold do not vanish, and that nonvanishing invariants produce symplectically embedded curves. Together with some work of Li–Liu [16], Liu [17] and Lalonde-McDuff [14], this completely settles questions on the symplectic structure of projective space and ruled manifolds: for a survey and detailed references see Lalonde-McDuff [15].

Ruan's 6-dimensional examples mentioned above have the property that the two symplectic forms are not deformation equivalent even after pull back by an arbitrary diffeomorphism. It is not known in 4 dimensions whether such examples can exist. In other words one can ask:

> *is there a 4-manifold X that supports two symplectic forms ω_0 and ω_1 with the property that $\psi^*\omega_1$ is not deformation equivalent to ω_0 for any diffeomorphism ψ?*

Another question is:

> *are two deformation equivalent and cohomologous forms ω_0, ω_1 on a 4-manifold necessarily isotopic?*

The point here is that isotopy is the basic relation if you are interested in the symplectic geometry of a manifold, but the Gromov invariants defined by counting J-holomorphic curves are invariant under deformation. In fact, in 4-dimensions Taubes has shown that the Gromov invariants coincide with the Seiberg–Witten invariants and so depend only on the smooth structure. (For more information on different ways of counting curves see Ionel–Parker [12].) The 6-dimensional examples of deformation equivalent but nonisotopic forms used an invariant of the moduli space that was more complicated than simply counting the numbers of curves through a set of points (or, more generally, through some k-cycle.) Therefore, one might hope that there would be symplectic 4-manifolds X with positive dimensional moduli spaces of curves that one could use to create new invariants fine enough to distinguish some nonisotopic forms. In order to get an invariant, these moduli spaces would have to be nonempty for generic J, and the easiest way to ensure this is to ask that the corresponding Gromov invariant is nonzero.[11] In the simplest case, these moduli spaces would appear as solution spaces to the Seiberg–Witten equations for some Spin^c structure, and so X would have to have $b_2^+ = 1$ and be of so-called nonsimple SW-type. Such manifolds do exist: for example blow-ups of ruled surfaces or of the complex projective plane. However, there is an inflation technique due to Lalonde–McDuff (see [14]) that allows one to use an embedded J-holomorphic curve to change the cohomology class of a family of symplectic forms, and if there are enough such curves one can alter this cohomology class almost at will. Hence one can show:

> *if X is a symplectic manifold of nonsimple SW-type, any deformation ω_t between two cohomologous symplectic forms ω_0 and ω_1 is homotopic through deformations with fixed endpoints to an isotopy.*

This rules out any easy way of finding nonisotopic but cohomologous symplectic forms on a 4-manifold. It also allows one to answer a question about the uniqueness of symplectic blowing up that I have been thinking about for quite some time. When you blow up a symplectic manifold you embed a lot of disjoint balls, cut out their interiors and then glue up their boundaries: for details see [21]. The cohomology class of the blow up is determined by the size of the balls. Moreover, in dimension 4, isotopy classes of forms on the blow up are in bijective correspondence with path components of the relevant space of embeddings of balls. Since these embedding spaces are clearly path connected if the balls can be reduced in size, it is easy to see that any two blow up forms are deformation equivalent. Therefore,

> *if X satisfies the hypotheses above, in particular if X is a rational or ruled surface, cohomologous blow ups are unique up to isotopy.*

Details of this argument may be found in McDuff [19]. One can also use similar ideas to construct certain blow ups thereby solving the symplectic packing

[11] It is no good taking X to be something like the 4-torus T^4 that has families of J-holomorphic tori of real dimension 2 for some special J – a product, for example – but no tori for other J.

problem for these manifolds: see Biran [2].

In conclusion I would like to mention Donaldson's work [5] (see also Auroux [1]) that constructs geometrically interesting "almost J-holomorphic" sections of "sufficiently positive" line bundles on a symplectic manifold (of any dimension). In particular, he has shown that if the cohomology class of the form ω is integral, it is possible to represent a sufficiently large multiple of the Poincaré dual of this class by a codimension 2 symplectic submanifold that is well-defined up to isotopy. This promises both to give interesting obstructions to the existence of symplectic forms in dimensions > 4 and to yield many more invariants of symplectic manifolds. In particular, one may be able to use it to prove Donaldson's stabilization conjecture which says that if X, Y are symplectic 4-manifolds with the same homotopy type then the obvious symplectic forms on $X \times S^2$ and $Y \times S^2$ are deformation equivalent if and only if X and Y are diffeomorphic. If something like this were true, then proceeding inductively one might find that deformation classes of symplectic manifolds of any dimension larger than 4 are as complicated as diffeomorphism classes of symplectic 4-manifolds.

References

[1] D. Auroux, Asymptotically holomorphic families of symplectic submanifolds, preprint (1997).

[2] P. Biran, Symplectic packing in dimension 4, preprint (1996), to appear in *Geom. and Functional Analysis.*

[3] C. Conley and E. Zehnder, The Birkhoff–Lewis fixed point theorem and a conjecture of V.I. Arnold, *Inventiones Math.* **73** (1983), 33–49.

[4] Damian and F. Laudenbach, in preparation.

[5] S. K. Donaldson, Symplectic submanifolds and almost-complex geometry, *Journ. of Differential Geometry*, **44** (1996), 666 – 705.

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

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

[8] A. Floer, Symplectic fixed points and holomorphic spheres. *Communications in Mathematical Physics*, **120** (1989), 575–611.

[9] A. Floer, H. Hofer and K. Wysocki, Applications of symplectic homology I. *Mathematische Zeitschrift* **217** (1994), 577–606.

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

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

[12] E. Ionel and T. Parker, The Gromov invariants of Ruan–Tian and Taubes, preprint (1997).

[13] F. Lalonde and D. McDuff, The geometry of symplectic energy, *Annals of Mathematics*, **141** (1995), 349–371.

[14] F. Lalonde and D McDuff, The classification of ruled symplectic 4-manifolds, *Math. Research Letters* **3** (1996), 769–778.

[15] F. Lalonde and D McDuff, J-curves and the classification of rational and ruled symplectic 4-manifolds, in *Symplectic and Contact Geometry*, ed. C. Thomas, Camb Univ Press (1996).

[16] T.J. Li and A. Liu, General wall crossing formula, *Math. Research Letters* **2** (1995), 797–810.

[17] A. Liu, Some new applications of general wall crossing formula, *Math. Research Letters*, **3** (1996), 569–586.

[18] D. McDuff, Examples of symplectic structures, *Inventiones Mathematicae*, **89**, (1987) 13–36.

[19] D. McDuff, From symplectic deformation to isotopy, to appear in the Proceedings of the conference at Irvine 1996, ed Stern, International Press.

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

[21] D. McDuff and D. Salamon, *Introduction to Symplectic Topology*, OUP, Oxford, UK (1995)

[22] Y. Ruan, Symplectic topology on algebraic 3-folds, *Journal of Differential Geometry* **39** (1994), 215–227.

[23] Y. Ruan and G. Tian, A mathematical theory of quantum cohomology, *Journal of Differential Geometry*, **42**, (1995), 259–367.

[24] Y. Ruan and G. Tian, Higher genus symplectic invariants and sigma model coupled with gravity, *Inventiones Mathematicae*, **130** (1997), 455–516.

[25] C. H. Taubes, The Seiberg–Witten and the Gromov invariants, *Math. Research Letters* **2** (1995), 221-238.