WHAT IS SYMPLECTIC GEOMETRY?

DUSA MCDUFF

Abstract. In this talk we explain the elements of symplectic geometry, and sketch the proof of one of its foundational results — Gromov's nonsqueezing theorem — using J-holomorphic curves.

1. First notions

Symplectic geometry is an even dimensional geometry. It lives on even dimensional spaces, and measures the sizes of 2-dimensional objects rather than the 1-dimensional lengths and angles that are familiar from Euclidean and Riemannian geometry. It is naturally associated with the field of complex rather than real numbers. However, it is not as rigid as complex geometry: one of its most intriguing aspects is its curious mixture of rigidity (structure) and flabbiness (lack of structure). In this talk I will try to describe some of the new kinds of structure that emerge.

First of all, what is a symplectic structure? The concept arose in the study of classical mechanical systems, such as a planet orbiting the sun, an oscillating pendulum or a falling apple. The trajectory of such a system is determined if one knows its position and velocity (speed and direction of motion) at any one time. Thus for an object of unit mass moving in a given straight line one needs two pieces of information, the position q and velocity (or more correctly momentum) $p := \dot{q}$. This pair of real numbers $(x_1, x_2) := (p, q)$ gives a point in the plane $\mathbb{R}^2$. In this case the symplectic structure ω is an area form (written $dp \wedge dq$) in the plane. Thus it measures the area of each open region S in the plane, where we think of this region as oriented, i.e. we choose a direction in which to traverse its boundary ∂S. This means that the area is signed, i.e. as in Figure 1.1 it can be positive or negative depending on the orientation. By Stokes' theorem, this is equivalent to measuring the integral of the action $p\,dq$ round the boundary ∂S.

This might seem a rather arbitrary measurement. However, mathematicians in the nineteenth century proved that it is preserved under time evolution. In other words, if a set of particles have positions and velocities in the region S_1 at the time t_1 then at any later time t_2 their positions and velocities will form a region S_2 with the same area. Area also has an interpretation in modern particle (i.e. quantum) physics. Heisenberg's Uncertainty Principle says that we can no longer know both position and velocity to

Date: March 31, 2009, revised June 23 2009. Talk given at the European Women in Mathematics conference of Sept 2007 in Cambridge, UK..

Partially supported by NSF grant DMS 0604769.

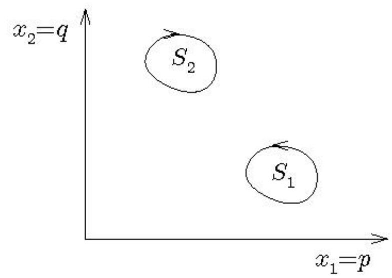

FIGURE 1.1. The area of the region S_1 is positive, while that of S_2 is negative.

an arbitrary degree of accuracy. Thus we should not think of a particle as occupying a single point of the plane, but rather lying in a region of the plane. The Bohr-Sommerfeld quantization principle says that the area of this region is quantized, i.e. it has to be an integral multiple of a number called Planck's constant. Thus one can think of the symplectic area as a measure of the entanglement of position and velocity.

An object moving in the plane has two position coordinates q_1, q_2 and correspondingly two velocity coordinates $p_1 = \dot{q}_1, p_2 = \dot{q}_2$ that measure its speed in each direction. So it is described by a point

$$(x_1, x_2, x_3, x_4) = (p_1, q_1, p_2, q_2) \in \mathbb{R}^4$$

in the 4-dimensional space $\mathbb{R}^4$. The symplectic form ω now measures the (signed) area of 2-dimensional surfaces S in $\mathbb{R}^4$ by adding the areas of the projections of S to the (x_1, x_2)-plane and the (x_3, x_4)-plane. Thus, as is illustrated in Figure 1.2,

$$\omega(S) = \text{area}(pr_{12}(S)) + \text{area}(pr_{34}(S)).$$

Notice that $\omega(S)$ can be zero: for example S might be a little rectangle in the x_1, x_3 directions which projects to a line under both pr_{12} and pr_{34}.

More technically, ω is a differential 2-form written as

$$\omega = dx_1 \wedge dx_2 + dx_3 \wedge dx_4,$$

and we evaluate the area $\omega(S) = \int_S \omega$ by integrating this form over the surface S. A similar definition is made for particles moving in n-dimensions. The symplectic area form ω is again the sum of contributions from each of the n pairs of directions:

$$\omega_0 = dx_1 \wedge dx_2 + dx_3 \wedge dx_4 + \cdots + dx_{2n-1} \wedge dx_{2n}. \tag{1.1}$$

We call this form ω_0 because it is the standard symplectic form on Euclidean space. The letter ω is used to designate any symplectic form.

To be even more technical, one can define a symplectic form ω on any even dimensional smooth (i.e. infinitely differentiable) manifold M as a closed, nondegenerate 2-form, where the nondegeneracy condition is that for each nonzero tangent direction v there is another direction w such that the area $\omega(v, w)$ of the little (infinitesimal) parallelogram spanned by these vectors is nonzero. (For a geometric interpretation of

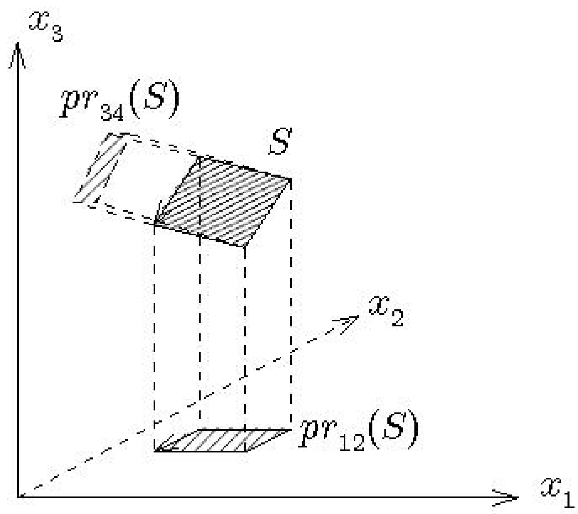

FIGURE 1.2. The symplectic area $\omega(S)$ is the sum of the area of its projection $pr_{12}(S)$ to the plane $(x_1, x_2) = (p_1, q_1)$ given by the velocity and position in the first direction together with the area of the corresponding projection $pr_{34}(S)$ for the two coordinates in the second direction. I have drawn the first 3 coordinates; the fourth is left to your imagination.

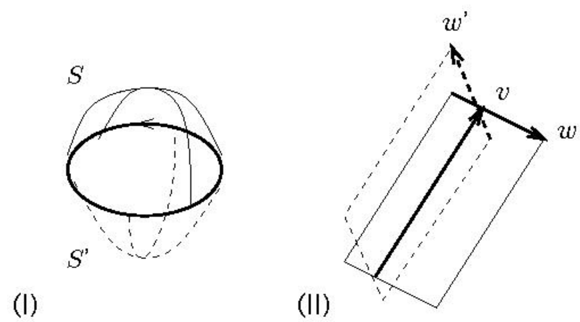

FIGURE 1.3. The fact that ω is *closed* means that the symplectic area of a surface S with boundary does not change as S moves, provided that the boundary is fixed. Thus in (I) the surfaces S and S' have the same area. Diagram (II) illustrates the *nondegeneracy* condition: for any direction v at least one of the family of 2-planes spanned by v and a varying other direction w has non zero area.

these conditions on ω see Figure 1.3.) A manifold is said to *be symplectic* or to *have a symplectic structure* if it is provided with a symplectic form.

The first important theorem in symplectic geometry is that locally[1] all symplectic forms are the same.

Darboux's Theorem: *Given a symplectic form ω on a manifold M and any point on M one can always find coordinates $(x_1, \dots, x_{2n})$ defined in an open neighborhood U of this point such that in this coordinate system ω is given on the whole open set U by formula (1.1).*

This is very different from the situation in the usual (Riemannian) geometry where one can make many local measurements (for example involving curvature) that distinguish among different structures. Darboux's theorem says that *all* symplectic structures are locally indistinguishable. Of course, as mathematicians have been discovering in the past 20 years, there are many very interesting global invariants that distinguish different symplectic structures. But most of these are quite difficult to define, often involving deep analytic concepts such as the Seiberg–Witten equations or J-holomorphic curves.

A symplectic form ω has an important invariant, called its cohomology class $[\omega]$. This class is determined by the areas $\omega(S)$ of *all* closed[2] surfaces S in M. In fact, for compact M the class $[\omega]$ is determined by a finite number of these areas $\omega(S_i)$ and so contains only a finite amount of information. Cf. Figure 1.3 where we pointed out that the area $\omega(S)$ does not change if we move S around.

There is a similar flabbiness in the symplectic structure itself. A fundamental theorem due to Moser says that one cannot change the symplectic form in any important way by deforming it, provided that the cohomology class is unchanged. More precisely, if $\omega_t, t \in [0,1]$, is a smooth path of symplectic forms such that $[\omega_0] = [\omega_t]$ for all t, then all these forms are "the same" in the sense that one can make them coincide by moving the points of M appropriately.[3] The important point here is that we cannot find new structures by deforming the old ones, provided that we fix the integrals of our forms over all closed surfaces. This result is known as **Moser's Stability Theorem**, and is an indication of robustness of structure.[4]

2. Symplectomorphisms

Another consequence of the lack of local features that distinguish between different symplectic structures is that there are many ways to move the points of the underlying space M without changing the symplectic structure ω. Such a movement is called a *symplectomorphism*. This means first that

---

[1] This means "on suitably small open sets".

[2] A *closed* surface is something like the surface of a sphere or donut; it has no edges and no holes.

[3] In technical language we say that these forms are all diffeomorphic.

[4] For precise statements, many proofs, and a list of references on all the topics mentioned here see [9]. There are also other more elementary books such as Cannas [2]. For simplicity, we shall only work here in dimensions 2 and 4. But all the results have higher dimensional analogs.

• ϕ is a *diffeomorphism*, that is, it is a bijective (one to one and onto) and smooth (infinitely differentiable) map $\phi : M \to M$, giving rise to the movement $x \mapsto \phi(x)$ of the points x of the space M;

and second that

• it *preserves symplectic area*, i.e. $\omega(S) = \omega(\phi(S))$ for *all* little pieces of surface S. The important point here is that this holds for all S, no matter how small or large. (Technically it is better to work on the infinitesimal level, looking at the properties of the derivative $d\phi$ of ϕ at each point.)

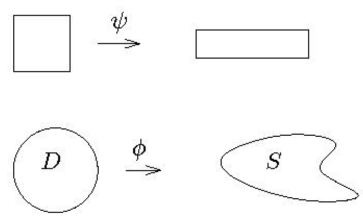

FIGURE 2.1. Symplectomorphisms in dimension 2.

In 2-dimensions, a symplectomorphism ϕ is simply an area preserving transformation. For example the map ψ in Figure 2.1 is given by the formula $\psi(x_1, x_2) = (2x_1, \frac{1}{2}x_2)$. Since it multiplies one coordinate by two and divides the other by two it does not change area. More generally, one form of Moser's theorem says the following:

Characterization of plane symplectomorphisms: *Suppose that S is a region in the plane $\mathbb{R}^2$ that is diffeomorphic to a disc D and has the same area as D. Then there is a symplectomorphism $\phi : D \to S$.*

The above statement means that we can choose the diffeomorphism $\phi : D \to S$ so that it preserves the area of *every* subset of D not just of D itself.

In 4-dimensions the situation is rather different. Gromov was the first to try to answer the question:

What are the possible shapes of a symplectic ball?

More precisely, let B be the round ball of radius one in $\mathbb{R}^4$. Thus

$$B = \left\{(x_1, \ldots, x_4) : x_1^2 + x_2^2 + x_3^2 + x_4^2 \le 1\right\}$$

consists of all points whose (Euclidean) distance from the origin $\{0\}$ is at most one. What can one say about the set $\phi(B)$ where ϕ is any symplectomorphism? Can $\phi(B)$ be long and thin? Can its shape be completely arbitrary? The analog of the 2-dimensional result would be that $\phi(B)$ could be any set that is diffeomorphic to B and also has the same volume.[5]

[5]Since $\omega \wedge \omega$ is a volume form, any symplectomorphism preserves volume. The fact that it is impossible to give a completely elementary proof of this (e.g. one that does not involve the concept

It is possible for $\phi(B)$ to be long and thin. For example one can stretch out the coordinates x_1, x_3 while shrinking the pair x_2, x_4 as in the map

$$\phi\big((x_1, x_2, x_3, x_4)\big) = (2x_1, \tfrac{1}{2}x_2, 2x_3, \tfrac{1}{2}x_4).$$

But the map

$$\psi\big((x_1, x_2, x_3, x_4)\big) = (\tfrac{1}{2}x_1, \tfrac{1}{2}x_2, 2x_3, 2x_4)$$

will *not* do since the area of rectangles in the x_1, x_2 plane are divided by 4. Note that ϕ preserves the pairs (x_1, x_2) and (x_3, x_4) and is made by combining area preserving transformations in each of these 2-planes. One might ask if there is a symplectomorphism that mixes these pairs, for example rotates in the x_1, x_3 direction. There certainly are such maps. For example

$$\phi\big((x_1, x_2, x_3, x_4)\big) = \tfrac{1}{2}(x_1 - x_3,\ x_2 - x_4,\ x_1 + x_3,\ x_2 + x_4)$$

is a symplectomorphism that rotates anticlockwise by 45 degrees in both the x_1, x_3 plane and the x_2, x_4 plane.

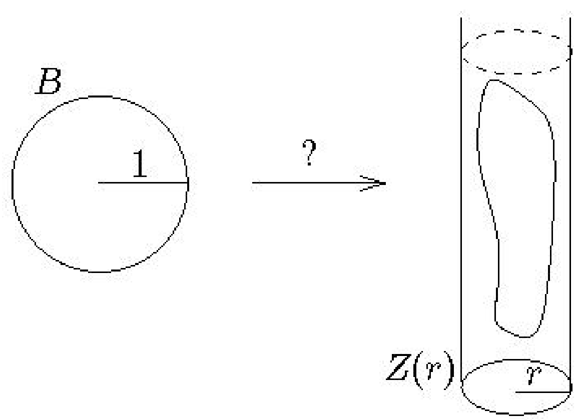

FIGURE 2.2. Can the unit ball B be squeezed into the cylinder $Z(r)$?

Nevertheless, Gromov in his nonsqueezing theorem showed that as far as the large geometric features of the space are concerned one can still see this splitting of $\mathbb{R}^4$ into the product of the (x_1, x_2) plane and the (x_3, x_4) plane. He described this in terms of maps of the unit ball B into the cylinder

$$Z(r) := D^2(r) \times \mathbb{R}^2 = \big\{(x_1, \ldots, x_4) : x_1^2 + x_2^2 \le r^2\big\} \ \subset \ \mathbb{R}^4 \tag{2.1}$$

of radius r, showing that one cannot squeeze a large ball into a thin cylinder of this form.

Gromov's Nonsqueezing Theorem: *If $r < 1$ there is no symplectomorphism ϕ such that $\phi(B) \subset Z(r)$.*

---

of a differential form) reflects the fact that to nineteenth century mathematicians this was a nontrivial result; cf. [1].

Although the nonsqueezing theorem might seem quite special and therefore unimportant (though perhaps cute), the property expressed here, that symplectomorphisms cannot squeeze a set in a pair of "symplectic directions" such as x_1, x_2, turns out to be absolutely fundamental: when properly formulated it gives a necessary and sufficient condition for a diffeomorphism to preserve the symplectic structure. Thus this theorem should be understood as a geometric manifestation of the very nature of a symplectic structure.

Another similar problem is that of the *Symplectic Camel*.[6] Here the camel is represented by a round 4-dimensional ball of radius 1 say, and the eye of the needle is represented by a "hole in a wall". That is to say, the wall with a hole removed is given by

$$W = \{(x_1, x_2, x_3, x_4) \in \mathbb{R}^4 \;:\; x_1 = 0, x_2^2 + x_3^2 + x_4^2 \geq 1\}$$

and we ask whether a (closed) round ball B of radius 1 can be moved from one side of the wall to the other in such a way as to preserve the symplectic form. (Note that because the ball is closed and the 2-sphere $\{x_2^2 + x_3^2 + x_4^2 = 1, x_1 = 0\}$ is contained in the wall, the ball will get stuck half way if one just tries moving it by a translation.)

It is possible to do this if one just wants to preserve volume. This is easy to see if one restricts to the three-dimensional case (by forgetting the last coordinate x_4); one can imagine squeezing a sufficiently flexible balloon through any small hole while preserving its volume. However, as Gromov showed, the symplectic case is more rigid.

The Symplectic Camel: *It is impossible to move a ball of radius ≥ 1 symplectically from one side of the wall to the other.*

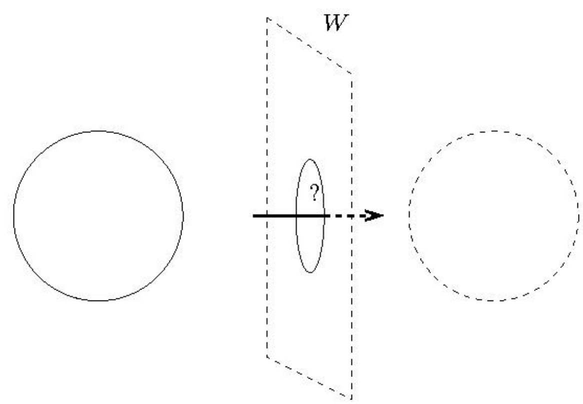

FIGURE 2.3. Can the ball go through the hole?

[6]The name of this problem is a somewhat "in" joke, of the kind appreciated by many mathematicians. The reference is to the saying that it would be easier for a camel to go through the eye of needle than for a rich man to get into heaven. (This saying is probably a mistranslation of a sentence in the bible.)

Both these results show that there is some rigidity in symplectic geometry. Exactly how this is expressed is still not fully understood, especially in dimensions > 4. However, recently progress has been made on another fundamental embedding problem in dimension 4. The question here is to understand the conditions under which one ellipsoid embeds symplectically[7] in another. Here, by the ellipsoid $E(a, b)$ we mean the set

$$E(a, b) := \Big\{(x_1, \dots, x_4) \ : \ \frac{x_1^2 + x_2^2}{a} + \frac{x_3^2 + x_4^2}{b} \le 1\Big\}.$$

In this language, the ball $B(r)$ of radius r is $E(r^2, r^2)$; in other words, the numbers a, b are proportional to areas, not lengths. Thus the question is: when does $E(a, b)$ embed symplectically in $E(a', b')$? Here we will fix notation by assuming that $a \le b$ and $a' \le b'$.

If the first ellipsoid is a ball $E(a, a)$ then the answer is given by the Nonsqueezing Theorem:

> *a necessary and sufficient condition for embedding $E(a, a)$ into $E(a', b')$ (where $a' \le b'$) is that $a \le a'$.*

(This condition is obviously sufficient since $E(a, b)$ is a subset of $E(a', b')$ when $a \le a'$ and $b \le b'$. On the other hand, it is necessary because if $E(a, a)$ embeds in $E(a', b')$ then, since $E(a', b') \subset Z(\sqrt{a'})$, it also embeds in $Z(\sqrt{a'})$, so that by the nonsqueezing theorem we must have $a \le a'$.)

But if the target is a ball $E(a', a')$ and the domain $E(a, b)$ is an arbitrary ellipsoid the answer is not so easy. It was proved in the 90s that when $a \le b \le 2a$ the situation is rigid: to embed $E(a, b)$ into $E(a', a')$ it is necessary and sufficient that $b \le a'$. In other words, the ellipsoid does *not* bend in this case. However, as soon as $b > 2a$ some flexibility appears and it is possible to embed $E(a, b)$ into $E(a', a')$ for some $a' < b$. Then of course one wants to know how much flexibility there is. What other obstructions are there to performing such an embedding besides the obvious one of volume? Notice that because a, b, a' are areas the volume obstruction to the existence of an embedding is that $ab \le (a')^2$.

This question was nicely formulated in a paper by Cieliebak, Hofer, Latschev and Schlenk [3] called *Quantitative Symplectic Geometry* in terms of the following function: define $c(a)$ for $a \ge 1$ by[8]

$$c(a) := \inf\{c' : E(a, b) \text{ embeds symplectically in } E(c', c')\}.$$

When [3] was written, this function was largely a mystery except that one knew that $c(a) = a$ for $a \le 2$ (rigidity). Now methods have been developed to understand it, and it should be fully known soon for all a: see McDuff and Schlenk [12] and also [11]. As

[7]We say that the set U embeds symplectically in V if there is a symplectomorphism ϕ such that $\phi(U) \subset V$.

[8]Note that $E(a, b)$ may not embed in $E\big(c(a), c(a)\big)$ itself – one usually needs a little extra room so that the boundary of $E(a, b)$ does not fold up on itself. However, one can show that the *interior* of $E(a, b)$ does embed in $E\big(c(a), c(a)\big)$.

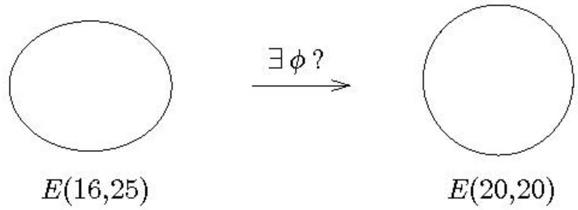

FIGURE 2.4. Does $E(a, b)$ embed symplectically in $E(a', b')$? In the case illustrated here, $ab = a'b'$, so there is no volume obstruction to the embedding, but the embedding does not exist because $25 > 20$. In fact, if we rescale by dividing all areas by 16, the problem is equivalent to embedding $E(1, \frac{25}{16})$ into $E(\frac{20}{16}, \frac{20}{16})$. But this is impossible by Equation (2.2) below.

a first step, work of Opshtein [14] can be used to evaluate $c(a)$ in the range $1 \le a \le 4$. Surprisingly, it turns out that

$$c(a) = a, \ \text{if } 1 \le a \le 2, \quad c(a) = 2 \ \text{if } 2 \le a \le 4. \tag{2.2}$$

In other words the graph is *constant* in the range $a \in [2, 4]$. To prove this one only needs to show that $c(4) = 2$. Because c is nondecreasing, if $c(2) = c(4) = 2$ then c must be constant on this interval. On the other hand, the statement $c(4) = 2$ implies that we can fill the volume of the ball $E(2, 2)$ by the interior of the ellipsoid $E(1, 4)$, which is a somewhat paradoxical state of affairs. Why can you fill all the volume of $E(2, 2)$ by the interior of $E(1, 4)$ when you cannot fill $E(\sqrt{2}, \sqrt{2})$ by the interior of $E(1, 2)$?

It turns out to be important that 4 is a perfect square. For any positive integer k, Opshtein discovered an explicit way to embed $E(1, k^2)$ into $E(k, k)$ that embeds the ellipsoid $E(1, k^2)$ into a neighborhood of a degree k curve such as $z_0^k + z_1^k + z_2^k = 0$ in the complex projective plane. Thus there is a clear geometric reason why the case $a = k^2$ is different from the general case. Many more things are now known about the function c: Figure 2.5 gives an idea of its graph. Here I would just like to point out that its behavior on the interval $[1, 4]$ is typical in symplectic geometry: either the situation is rigid (for $a \in [1, 2]$, the ellipse does not bend at all) or it is as flexible as it could possibly be (for $a \in [2, 4]$, the ellipsoid bends as much as is consistent with the obvious constraints coming from volume and the constraint at $a = 2$).

3. Almost complex structures and J-holomorphic curves

The remainder of this note tries to give a rough idea of how Gromov proved his results. Nowadays there are many possible approaches to the proof. But we shall explain Gromov's original idea that uses J-holomorphic curves. These provide a special way of cutting the cylinder into 2-dimensional slices of area πr^2 as in Figure 3.5, and we shall see that these provide an obstruction to embedding a ball of radius 1. Similarly, because one can fill the hole in the wall by these slices, the size of a ball that can be moved through the hole is constrained to be < 1.

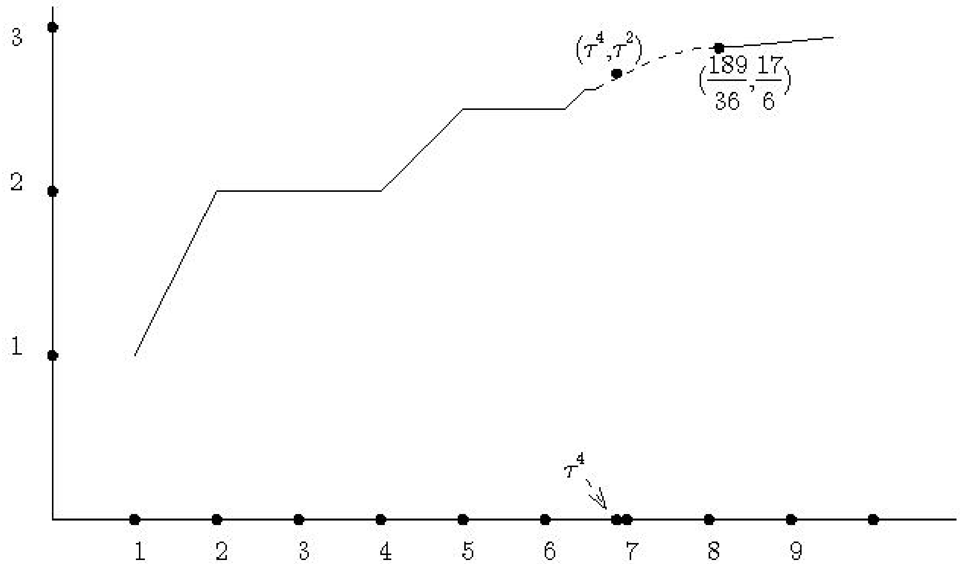

FIGURE 2.5. The graph of c: it appears to have an infinite staircase that converges to the point (τ^4, τ^2), where τ is the golden ratio, and it equals $\sqrt{a}$ for $a \geq \frac{189}{36}$. The graph between the points $\tau^4 < a < \frac{1}{36}$ is not yet completely known.

The concept of J-holomorphic curves has turned out to be enormously fruitful. Gromov's introduction of this idea in 1985 was one of the main events that initiated the modern study of symplectic geometry.

Gromov's key idea was to exploit the connection between symplectic geometry and the complex numbers.

A differentiable manifold M is a space in which one can do calculus: locally it looks like Euclidean space, but it can have interesting global structure.[9] As in calculus, one often approximates curves or surfaces near a given point $x \in M$ by the closest linear objects, tangent lines or planes as the case may be. The collection of all possible tangent directions at a point x is called the tangent space T_xM to M at x. It is a linear (or vector) space of the same dimension as M. As the point x varies over M the collection $\cup_{x\in M} T_xM$ of all these planes forms what is called the *tangent bundle* of M. If $M = \mathbb{R}^{2n}$ is Euclidean space itself, then one can identify each of its tangent spaces $T_x\mathbb{R}^{2n}$ with $\mathbb{R}^{2n}$, but most manifolds (such as the sphere) curve around and do not contain their tangent spaces.

An *almost complex structure at a point x* of a manifold M is a linear transformation J_x of the tangent space T_xM at x whose square is -1. Geometrically, J_x rotates by a quarter turn (with respect to a suitable coordinate system at x.) Thus the tangent space T_xM becomes a complex vector space (with the action of J_x playing the role of

[9]For a wonderful introduction to this subject see Milnor [13].

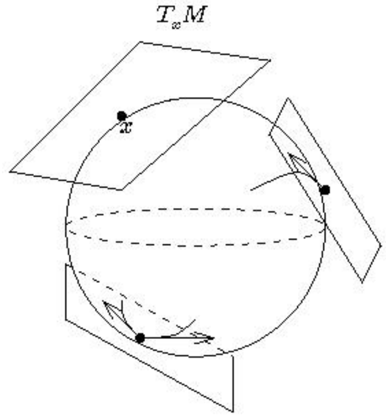

FIGURE 3.1. Some curves and tangent vectors on the two-sphere, together with some tangent spaces T_xM.

multiplication by $\sqrt{-1}$.) An *almost complex structure* J on M is a collection J_x of such transformations, one for each point of x, that varies smoothly as a function of x. If M has dimension 2 one can always choose local coordinates on M to make the function $x \to J_x$ constant. However in higher dimensions this is usually impossible. If such coordinates exist J is said to be *integrable*. What this means is explained more fully in Equation (3.2).

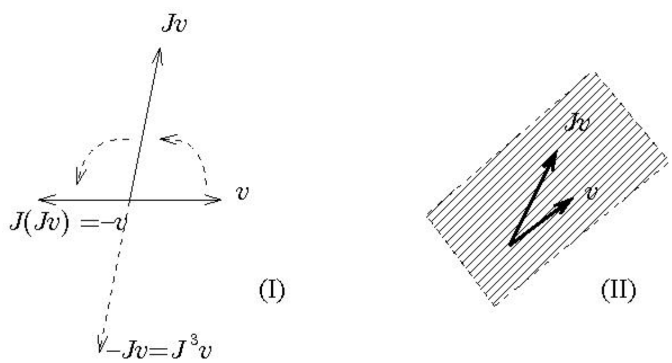

FIGURE 3.2. (I) pictures $J = J_x$ as a (skew) rotation; (II) shows the complex line spanned by v, Jv.

Rather few manifolds have integrable almost complex structures. (To be technical for a minute, this happens if and only if M has a complex structure, i.e. if and only if one can glue M together from its locally Euclidean pieces $U \subset \mathbb{C}^n$ by using holomorphic functions.) However, many manifolds have almost complex structures.[10] In particular,

[10]For example the 6-dimensional sphere S^6 has an almost complex structure. It is a famous unsolved problem to decide whether it has a complex structure.

symplectic manifolds always do. In fact, in this case one can choose J to be *compatible* with the symplectic form ω, i.e. so that at all points $x \in M$

$$\omega(J_x v, J_x w) = \omega(v, w), \quad \text{and} \quad \omega(v, J_x v) > 0, \tag{3.1}$$

for all nonzero tangent vectors $v, w \in T_x M$. [11] The first equation here says that rotation by J_x preserves symplectic area, while the inequality (called the *taming condition*) says that every complex line has positive symplectic area. Note that complex lines have 2 real dimensions; they are spanned over $\mathbb{R}$ by two vectors of the form v and $J_x v =$ "iv".

For any given ω there are many compatible almost complex structures; in fact there is a contractible set of such structures. Associated to each such J there is a Riemannian metric, i.e. a symmetric inner product g_J on the tangent space $T_x M$. It is given by the formula

$$g_J(v, w) := \omega(v, Jw), \qquad v, w \in T_x M.$$

As with any metric, this gives a way of measuring lengths and angles. However, it depends on the choice of J and so is not determined by ω alone. Nevertheless, because via J it has a very geometric relationship to ω it is often useful to consider it. [12]

As an example that will be useful later, observe that the usual (integrable!) complex structure J_0 on $\mathbb{C}^2 = \mathbb{R}^4$ is compatible with the standard symplectic form ω_0 and the associated metric g_0 is the usual Euclidean distance function.

J-holomorphic curves: A (real) curve in a manifold M is a path in M; that is, it is the image of a map $f : U \to M$ where U is a subinterval of the real line $\mathbb{R}$. A J-holomorphic curve in an almost complex manifold (M, J) is the complex analog of this. It has one complex dimension (but 2 real dimensions) and is the image of a "complex" map $f : \Sigma \to M$ from some complex curve Σ into (M, J). Here we shall take the domain Σ to be either a 2-dimensional disc D (consisting of a circle in the plane together with its interior) or the 2-sphere $S^2 = \mathbb{C} \cup \{\infty\}$, which we shall think of as the complex plane $\mathbb{C}$ completed by adding a point at ∞; see Figure 3.3.

If J is integrable, we can choose local complex coordinates on the target space M of the form $z_1 = x_1 + ix_2,\ z_2 = x_3 + ix_4$ so that at each point x the linear transformation $J = J_x$ acts on the tangent vectors $\frac{\partial}{\partial x_j}$ by "multiplication by i": namely,

$$J\Big(\frac{\partial}{\partial x_1}\Big) = \frac{\partial}{\partial x_2},\ \ J\Big(\frac{\partial}{\partial x_2}\Big) = -\frac{\partial}{\partial x_1},\ \ J\Big(\frac{\partial}{\partial x_3}\Big) = \frac{\partial}{\partial x_4},\ \ J\Big(\frac{\partial}{\partial x_4}\Big) = -\frac{\partial}{\partial x_3}. \tag{3.2}$$

Then there is an obvious notion of "complex" map: in terms of a local coordinate $z = x + iy$ on the domain and this coordinate system on the target, f is given by two power series $f_1(z),\ f_2(z)$ with complex coefficients a_k, b_k:

$$f(z) = \big(f_1(z), f_2(z)\big) = \Big(\sum_{k\ge 0} a_k z^k, \sum_{k\ge 0} b_k z^k\Big),$$

[11] Here I have used the language of differential 2-forms; but readers can think of $\omega(v, w)$ as the symplectic area of a small (infinitesimal) parallelogram spanned by the vectors v, w.

[12] For example, the associated metric on the loop space of M leads to a very natural interpretation for gradient flows on this loop space. This is the basis of Floer theory; see [8].

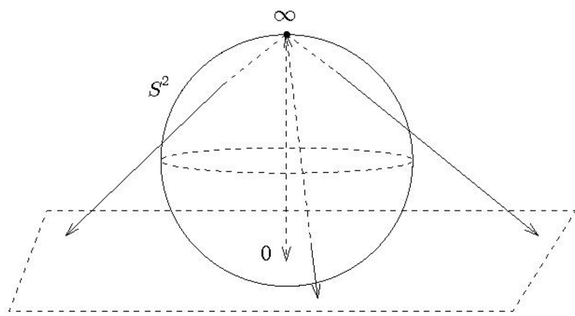

FIGURE 3.3. The 2-sphere S^2 as the completion of the complex plane $\mathbb{C}$. Often one puts ∞ at the north pole of the 2-sphere and identifies $S^2 \setminus \{\infty\}$ with the plane via stereographic projection.

i.e. f is holomorphic. Such functions can be characterized by the behavior of their derivatives: these must satisfy the Cauchy–Riemann equation

$$\frac{\partial f}{\partial x} + J\frac{\partial f}{\partial y} = 0.$$

This equation still makes sense even if J is not integrable, and so given such J we say that $f : \Sigma \to (M, J)$ is J-holomorphic if it satisfies the above equations.

J-holomorphic curves as minimal surfaces: The images $f(\Sigma)$ of such maps have very nice properties. In particular, their area with respect to the associated metric g_J equals their symplectic area. We saw earlier that the symplectic area of a surface is invariant under deformations of the surface that fix its boundary. (Cf. Figure 1.3 (I).) It follows easily that their metric area can only *increase* under such deformations, i.e. J-holomorphic curves are so-called g_J-minimal surfaces. Thus we can think of them as the complex analog of a real geodesic.[13]

Minimal surfaces have the following very important property that we will use later. Let g_0 be the usual Euclidean metric on $\mathbb{R}^4$ (or, in fact, on any Euclidean space $\mathbb{R}^d$). Suppose that S is a g_0-minimal surface in the ball B of radius 1 that goes through the center of the ball and has the property that its boundary lies on the surface of the ball. (Technically, we say that S is *properly embedded* in B.) Then

(3.3) $$\text{the } g_0\text{-area of } S \text{ is } \geq \pi.$$

In fact the g_0-minimal surface of *least* area that goes through the center of a unit ball is a flat disc of area π. All others have nonpositive curvature, which means that at each point they bend in opposite directions like a saddle and so, unless they are completely flat, have area greater than that of the flat disc.

[13]Remember that a geodesic in a Riemannian manifold (M, g_J) is a path that minimizes the length between any two of its points (provided these are sufficiently close.) The metric area of a surface is a measure of its energy. Thus a minimal surface has minimal energy and is, for example, the shape taken up by a soap film in 3-space that spans a wire frame.

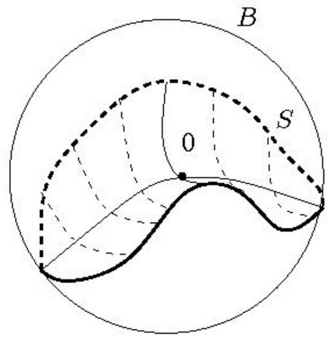

FIGURE 3.4. Any g_0-minimal surface S through 0 has area $\geq \pi$.

Families of J-holomorphic curves: Holomorphic (or complex) objects are much more rigid than real ones. For example there are huge numbers of differentiable real valued functions on the 2-sphere S^2, but the only complex (or holomorphic) functions on S^2 are *constant*.[14] There are infinitely many holomorphic functions if one just asks that they be defined in some small open subset of S^2, but the condition that they be globally defined is very strong. Something very similar happens with complex curves.

If one fixes a point $x \in M$ there are infinitely many real curves through x. In fact there is an infinite dimensional family of such curves, i.e. the set of all such curves can be given the topology of an infinite dimensional space. For real curves it does not matter if we look at little pieces of curves or the whole of a closed curve (e.g. the image of a circle). In the complex case, there still are infinite dimensional families of curves through x if we just look at little pieces of curves. But if we look at closed curves, e.g. maps whose domain is the whole of the 2-sphere, then there is at most a finite dimensional family of such curves. Moreover, Gromov discovered that under many circumstances the most important features of the behavior of these curves does not depend on the precise almost complex structure we are looking at.

For example, if (M, ω) is the complex projective plane with its usual complex structure J_0, then a complex line can be parametrized by a (linear) holomorphic map $f : S^2 \to (M, J_0)$ and so can be thought of as a J_0-holomorphic curve. If we perturb J_0 to some other ω-compatible almost complex structure J, then each complex line perturbs to a J-holomorphic curve. Gromov showed that, just as there is exactly one complex line through each pair of distinct points x, y, there is exactly one of these J-holomorphic curves through each x, y.[15]

[14]A well known result in elementary complex analysis is that every bounded holomorphic function that is defined on the whole of the complex plane $\mathbb{C}$ is constant. (These are known as entire functions.) Since the Riemann sphere $S^2 = \mathbb{C} \cup \{\infty\}$ contains $\mathbb{C}$, the same is true for S^2.

[15]This very sharp result uses the fact that the complex projective plane has 4 real dimensions. In higher dimensional complex projective spaces, Gromov showed that one can count these curves with appropriate signs and that the resulting sum is one. But now there may be more than one actual curve through two points. Thus the theory is no longer so geometric. The effect is that we know much more

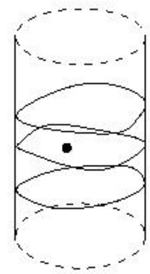

FIGURE 3.5. Slicing the cylinder $Z(r)$ with J-holomorphic discs of symplectic area πr^2. In dimension 4, provided that we put on suitable boundary conditions, or better still compactify as explained below, there is precisely one such disc through each point. In higher dimensions, with $Z(r) := D^2(r) \times \mathbb{R}^{2n-2}$, there is at least one.

What we need to prove the nonsqueezing theorem is a related result about cylinders $Z(r)$ as defined in Equation (2.1).

Slicing cylinders: *Let $\big(Z(r), \omega_0\big)$ be the cylinder in $(\mathbb{R}^4, \omega_0)$, and let J be any ω_0-tame almost complex structure on $Z(r)$ that equals the usual structure outside a compact subset of the interior of $Z(r)$. Then there is a J-holomorphic disc $f : (D^2, \partial D^2) \to (Z(r), \partial Z)$ of symplectic area πr^2 through every point of $Z(r)$.*

Note that here we are interested in discs whose boundary circle ∂D^2 is taken by f to the boundary ∂Z of the cylinder. The above statement is true for the usual complex structure J_0. In fact if $w_0 = (z_0, y_0) \in D^2(r) \times \mathbb{R}^2$ then the map

$$f(z) = \left(\frac{z}{r}, y_0\right) \in D^2(r) \times \mathbb{R}^2$$

goes through the point w_0. Because there is essentially one map of this kind (modulo reparametrizations of f), a deformation argument implies that there always is at least one such map no matter what J we choose.[16]

3.1. **Sketch proof of the nonsqueezing theorem.** Suppose that there is a symplectic embedding

$$\phi : B^4(1) \to Z(r) = D^2(r) \times \mathbb{R}^2.$$

We need to show that $r \geq 1$. Equivalently, by slightly increasing r, we may suppose that the image of the ball lies inside the cylinder, and then we need to show that $r > 1$. We shall do this by using J-holomorphic slices as described above, but where J is

about symplectic geometry in 4-dimensions than we do in higher dimensions. However, results like the nonsqueezing theorem are known in all dimensions.

[16] To make this argument precise we should partially compactify the target by identifying each boundary circle $\partial D^2(r) \times \{y\}, y \in \mathbb{R}^2$, to a point. The target then becomes $S^2 \times \mathbb{R}^2$. Correspondingly we should look at maps with domain $S^2 = D^2/\partial D^2$. Then the count of J-holomorphic curves through w_0 can be rephrased in terms of the degree of a certain map.

chosen very carefully. (Really the whole point of this argument is to choose a J that is related to the geometry of the problem.)

This is how we manage it. In order to make the slicing arguments work we need our J to equal the standard Euclidean structure J_0 near the boundary of $Z(r)$ and also outside a compact subset of $Z(r)$.[17] But because the image $\phi(B)$ of the ball is strictly inside the cylinder, we can also make J equal to any specified ω_0-tame almost complex structure on $\phi(B)$. In particular, we may assume that J equals the pushforward of the standard structure $\phi_*(J_0)$ on $\phi(B)$. In other words, *inside* the embedded ball J is "standard".

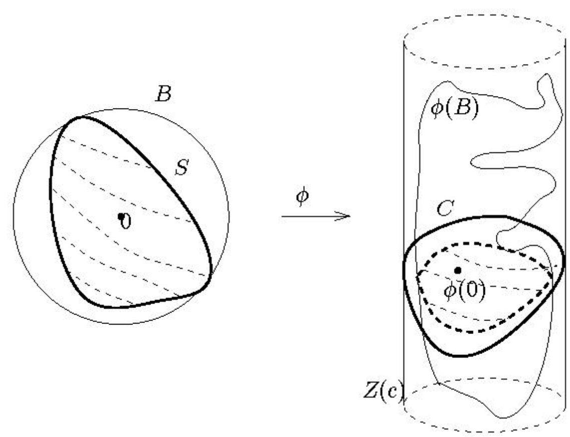

FIGURE 3.6. The g_0-minimal surface S is taken by the embedding ϕ into the intersection (outlined in heavy dots) of the slice C with the image $\phi(B)$ of the ball.

Then, the statement above about slicing cylinders says that there is a J-holomorphic disc

$$f : D^2 \to Z(r), \quad f(0) = \phi(0),$$

that goes through the image $\phi(0)$ of the center of the ball and also has boundary on the boundary of the cylinder. Further, the symplectic area of the slice $C = f(D^2)$ is πr^2.

Now consider the intersection $C_B := C \cap \phi(B)$ of the slice with the embedded ball $\phi(B)$. By construction, this goes though the image $\phi(0)$ of the center 0 of the ball B. We now look at this situation from the vantage point of the original ball B. In other words, we look at the inverse image $S := \phi^{-1}(C_B)$ of the curve C_B under ϕ as in Figure 3.6. This consists of all points in the ball that are taken by ϕ into C_B and forms a curve in B that goes through its center. The rest of our argument involves understanding the properties of this curve S in B.

[17]This technical condition is needed so that we can compactify the domain and target as explained earlier.

One very important fact is that S is holomorphic in the usual sense of this word, i.e. it is holomorphic with respect to the usual complex structure J_0 on $\mathbb{R}^4 = \mathbb{C}^2$. This follows from our choice of J: by construction, J equals the pushforward of J_0 on the image $\phi(B)$ of the ball, and so, because C_B lies in $\phi(B)$ and is J-holomorphic, it pulls back to a curve S that is holomorphic with respect to the pullback structure J_0.

As we remarked above, this means that S is a minimal surface with respect to the standard metric g_0 on $\mathbb{R}^4$ associated to ω_0 and J_0. So by Equation (3.3) the area of S with respect to g_0 is at least π. But because S is holomorphic, this metric area is the same as its symplectic area $\omega_0(S)$. This, in turn, equals the ω_0-area of the image curve $\phi(S) = C_B$, because ϕ preserves ω_0.

Finally note that, by construction, C_B is just part of the J-holomorphic slice C through $\phi(0)$. It follows that C_B has strictly smaller ω_0-area than C. (This follows from the taming condition $\omega_0(v, Jv) > 0$ of Equation (3.1), which, because it is a pointwise inequality, implies that every little piece of a J-holomorphic curve — such as $C \smallsetminus C_B$ — has strictly positive symplectic area.) But our basic theorem about slices says that $\omega_0(C) = \pi r^2$. Putting this all together, we have the following string of inequalities and equalities:

$$\pi \ \le\ g_0\text{-area } S \ =\ \omega_0\text{-area } S \ =\ \omega_0\text{-area } \phi(S) \ <\ \omega_0\text{-area } C \ =\ \pi r^2.$$

Thus $\pi < \pi r^2$. This means that $r > 1$, which is precisely what we wanted to prove.

Acknowledgements: I wish to thank Yael Degany and James McIvor for their useful comments on an earlier draft of this note.

References

[1] V. I. Arnold, *Mathematical methods in Classical Mechanics* Springer (1978).

[2] A. Cannas, *Lectures on Symplectic Geometry*, Lecture Notes in Mathematics vol 1764, 2nd ed. Springer (2008).

[3] K. Cieliebak, H. Hofer, J. Latschev and F. Schlenk, Quantitative symplectic geometry, arXiv:math/0506191, *Dynamics, Ergodic Theory, Geometry MSRI*, **54** (2007), 1–44.

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

[5] C. Marle, The inception of Symplectic Geometry: the works of Lagrange and Poisson during the years 1808–1810, arxiv:math/0902.0685.
An article about the early history of the idea of a symplectic form. A more standard description of the history can be found in [1].

[6] D. McDuff, Symplectic structures – a new approach to geometry, *Notices of the Amer. Math. Soc.* **45** (1998), 952–960.
This is slightly more technical than the current article, but is still quite accessible.

[7] D. McDuff, A glimpse into symplectic geometry, in *Mathematics: Frontiers and Perspectives 2000*, AMS, Providence
another survey article that emphasizes the two-fold nature of symplectic geometry.

[8] D. McDuff, Floer Theory and Low dimensional Topology, *Bull Amer. Math. Soc.* (2006)
a survey article describing some recent applications of symplectic geometry to low dimensional topology

[9] D. McDuff and D. Salamon, *Introduction to Symplectic Topology*, 2nd edition (1998) OUP, Oxford, UK

[10] D. McDuff and D.A. Salamon, *J-holomorphic curves and symplectic topology*. Colloquium Publications **52**, American Mathematical Society, Providence, RI, (2004).
the two above books are comprehensive (and not very elementary) introductions to the field
[11] D. McDuff, Symplectic embeddings of 4-dimensional ellipsoids, to appear in *Journ. of Top.* (2009).
[12] D. McDuff and F. Schlenk, The embedding capacity of 4-dimensional symplectic ellipsoids, in preparation.
[13] J. Milnor, *Topology from a differentiable viewpoint*
an elementary and very clear introduction to this fascinating subject
[14] E. Opshtein, Maximal symplectic packings of $\mathbb{P}^2$, arxiv:0610677.

Department of Mathematics, Barnard College, Columbia University
E-mail address: `dusa@math.columbia.edu`