MATH 257A Symplectic Geometry

Shintaro Fushida-Hardy
381B Sloan Hall, Stanford University, CA

This document contains course notes from MATH 257A (taught at Stanford, fall 2019) transcribed by Shintaro Fushida-Hardy. This was largely live-TeXed during lectures, but some additional theorems, propositions, thoughts and mistakes have been added.

Contents

Chapter 1

Symplectic geometry

1.1 Lecture 1

1.1.1 Admin

Most probable course outline:

- First half of the course: First 8 lectures of Weinstein - "Lectures on Symplectic Manifolds"
- Next quarter of the course: Gromov non-squeezing, possibly following Gromov's paper.
- Final quarter of the course: "Symplectic rigidity: Lagrangian submanifolds" by Audin-Lalonde-Polterovich.

Prerequisites:

- Familiarity with vector fields and flows on smooth manifolds.
- Familiarity with differential forms and integration, and basic de Rham theory.
- Transversality theorem.

Other resources:

- McDuff-Salamon (big, good reference book)
- Da Silva (easy to read, similar to Weinstein)
- Arnol'd - "Methods of Classical Mechanics" (for the physically minded)

1.1.2 Motivation

Symplectic geometry is the study of symplectic manifolds, that is, the study of smooth manifolds equipped with a closed non-degenerate 2-form. More explicitly, a symplectic manifold is the data (M, ω), where ω satisfies the following properties:

1. $\omega \in \Omega^2(M)$, i.e. ω is an anti-symmetric bilinear form on T_pM for each p in M, which varies smoothly on M.

2. ω is non-degenerate. This means for each $p \in M$, $v \in T_pM$ non-zero, there exists $v' \in T_pM$ such that $\omega(v, v') \neq 0$.

3. ω is closed. This means $\mathrm{d}\omega = 0$, where d is the exterior derivative.

This last condition is difficult to motivate. However, all three conditions are necessary for *Darboux's theorem* to hold, which can be thought of as motivation.

Theorem 1.1.1 (Darboux)**.** *Let (M, ω) be a symplectic manifold, and $p \in M$. Then there are local coordinates $\{x_1, \ldots, x_n, y_1, \ldots, y_n\}$ such that $\omega = \mathrm{d}x_1 \wedge \mathrm{d}y_1 + \cdots + \mathrm{d}x_n \wedge \mathrm{d}y_n$.*

This result classifies all symplectic manifolds locally: the only local invariant up to symplectomorphism (defined later) is dimension. On the other hand, familiar structures such as Riemannian metrics have non-trivial local invariants (e.g. curvature).

Why study symplectic manifolds?

1. They arise naturally in other areas of maths:

   (a) Classical mechanics, or more specifically, Hamiltonian mechanics. Given a manifold X of "coordinates in space", the cotangent bundle T^*X defines the "phase space". The Hamiltonian is a smooth function

   $$H : T^*X \to \mathbb{R}.$$

   A symplectic structure allows the Hamiltonian to describe time evolution (dynamics) on X.

   (b) Complex geometry. Any affine variety which is also a complex manifold (more generally, a Stein manifold) has a natural symplectic structure which is unique up to symplectomorphism.

   (c) Lie groups/Lie algerbas. Let G be a Lie group and $\mathfrak{g}$ its Lie algebra. Recall the *adjoint representation* $G \to \mathrm{Aut}(\mathfrak{g})$, which sends each $g \in G$ to $\mathrm{Ad}(g) = \mathrm{d}\Phi_g$, where Φ_g is the inner automorphism of g. The *coadjoint representation* is then a map $G \to \mathrm{Aut}(\mathfrak{g}^*)$, defined by $g \mapsto \mathrm{Ad}(\mathrm{g}^{-1})^*$. The orbit of each $F \in \mathfrak{g}^*$ under the coadjoint action is a submanifold of $\mathfrak{g}^*$, and in fact carry natural symplectic structures.

(d) Representation varieties in low dimensional topology.

2. Worth studying for its own sake: On one hand, there are no local invariants, but on the other hand, there are beautiful and complicated global invariants. For closed symplectic manifolds, a simple global invariant is volume. Despite the lack of local invariants, symplectic manifolds are still "interesting", in the sense that the choice of ω is important. We will observe that symplectic structures usually have large symmetry groups, while metrics generally have small symmetry groups.

3. Finally, there is input from other fields.

1.2 Lecture 2

1.2.1 Different levels of symplectic structures

1. Symplectic vector space
2. Symplectic vector bundle
3. Symplectic manifold

Definition 1.2.1. Let V be a finite dimensional vector space, and $\Omega : V \times V \to \mathbb{R}$ a bilinear map. Note that the data of Ω is equivalent to that of a linear map $\widetilde{\Omega} : V \to V^*$, defined by $\widetilde{\Omega}(v) : w \mapsto \Omega(v, w)$. (V, Ω) is a *symplectic vector space* if Ω is antisymmetric and non-degenerate (in the sense that $\widetilde{\Omega}$ is an isomorphism). Note that there is not yet enough structure to make sense of Ω being closed. In this picture, a *symplectomorphism* is any bijective linear map which preserves the symplectic structure. Explicitly, a linear bijection $T : (V_1, \Omega_1) \to (V_2, \Omega_2)$ is a symplectomorphism if $\Omega_1(v, w) = \Omega_2(Tv, Tw)$ for all $v, w \in V_1$. More concisely, this means $\Omega_1 = T^*\Omega_2$.

Example. Let U be a real vector space. Then $U \oplus U^*$ has a natural symplectic structure defined by $\Omega(u+u^*, v+v^*) = v^*(u) - u^*(v)$. Suppose we choose a basis $\{e_1, \ldots, e_n, f_1, \ldots, f_n\}$. Then it is clear that Ω agrees with Darboux's theorem. Therefore, all symplectic vector spaces arise as $(U \oplus U^*, \Omega)$ for some U.

Definition 1.2.2. Let X be a manifold, and $\pi : E \to X$ a vector bundle. E is a *symplectic vector bundle* if it is equipped with a smoothly varying symplectic form Ω_x (as in definition 1.2.1) defined on the fibres E_x for each $x \in X$. More formally, smoothly varying means that Ω is a smooth section of the tensor bundle $E^* \otimes E^* \to X$. Here a *symplectomorphism* is a vector bundle isomorphism which pulls back the symplectic form.

Example. Recall that all symplectic vector spaces arise as direct sums $U \oplus U^*$. Do all symplectic vector bundles arise as Whitney sums $P \oplus P^* \to X$?

No. The sphere $\mathbb{S}^2$ admits a symplectic structure on its tangent bundle. However, any line bundle on $\mathbb{S}^2$ is trivial, so if the tangent bundle of $\mathbb{S}^2$ cannot be a sum bundle.

Definition 1.2.3. Let X be a manifold. A *symplectic manifold* is the data (X, ω) where ω is a symplectic structure on $TX \to X$ as in definition 1.2.2, which is also closed ($\mathrm{d}\omega = 0$). In this picture a *symplectomorphism* is a diffeomorphism which pulls back the symplectic form.

Example. Let V be a vector space, considered as a manifold. Then the manifold T^*V is diffeomorphic to $V \times V^*$, and so it has a natural symplectic structure as follows: At any point (v, v^*) in T^*V, $T_{(v,v^*)}(T^*V) = V \oplus V^*$ which carries a natural symplectic form $\Omega_{(v,v^*)}$. Define ω to be this symplectic form at each tangent space. Since it does not depend on the base point, it is exact.

Consider local coordinates $(q_1, \ldots, q_n, p_1, \ldots, p_n)$, with the q_i induced by e_i and p_i induced by e_i^*. The q_i represent position, while p_i represent momentum. Then

$$\omega = \mathrm{d}p_1 \wedge \mathrm{d}q_1 + \cdots + \mathrm{d}p_n \wedge \mathrm{d}q_n. \tag{1.1}$$

Example. Now suppose X is a manifold. Then T^*X has a natural symplectic structure as in the previous example. One method of construction is to consider an atlas of X. This induces an atlas of T^*X. Each of these charts is diffeomorphic to a subset of $\mathbb{R}^{2n}$, so they can be endowed with the natural symplectic structure from the above example. Moreover, they can be shown to glue correctly.

Another approach is to construct the so-called *tautological 1-form.* Consider the projection map $\pi : T^*X \to X$. Each point p_x in T^*X is of the form (x, ξ) where $\xi \in T_x^*X$. Thus, define $\lambda \in T^*(T^*X)$ pointwise by

$$\lambda_{(x,\xi)} = \xi \circ \mathrm{d}\pi = \mathrm{d}\pi^*\xi. \tag{1.2}$$

Given a vector $v \in T(T^*X)$, $\lambda_{(x,\xi)}$ first projects v onto T_xX, and then since $\xi \in T_x^*X$, the above formula defines a 1-form. The map λ is called the *tautological 1-form.* In fact, $\mathrm{d}\lambda$ defines a symplectic form, called the *canonical 2-form.*

Closedness of $\mathrm{d}\lambda$ is clear. To prove that $\mathrm{d}\lambda$ is non-degenerate and antisymmetric, we compute in coordinates. Let $\{U_\alpha\}$ be an atlas of X. Choose any point (x, ξ) in T^*U_α. In local coordinates, $(x, \xi) = \{q^1, \ldots, q^n, p_1, \ldots, p_n\}$ as in the previous example. Choose a vector $v \in T_{(x,\xi)}(T^*X)$, and choose α^i, β^i such that $v_{(x,\xi)} = \alpha^i \frac{\partial}{\partial q^i} + \beta^i \frac{\partial}{\partial p^i}$. Then

$$\begin{aligned}
\lambda_{(x,\xi)}(v) &= \xi\Big(\mathrm{d}\pi\Big(\alpha^i \frac{\partial}{\partial q^i} + \beta^i \frac{\partial}{\partial p^i}\Big)\Big) \\
&= p_j\Big(\alpha^i \frac{\partial}{\partial q^i}\Big) \\
&= p_i \alpha^i.
\end{aligned}$$

The last line follows from the fact that p_j are induced by the dual basis to $\frac{\partial}{\partial q^i}$. On the other hand,

$$\begin{aligned} p_j \mathrm{d}q^j(v) &= p_j \mathrm{d}q^j \left(\alpha^i \frac{\partial}{\partial q^i} + \beta^i \frac{\partial}{\partial p^i} \right) \\ &= p_i \alpha^i. \end{aligned}$$

Since these must hold for all (x, ξ) and v, it follows that $\lambda_{(x,\xi)} = p_i \mathrm{d}q^i$ in local coordinates. But now $\mathrm{d}\lambda = \mathrm{d}p_i \wedge \mathrm{d}q^i$, which is a non-degenerate antisymmetric form.

This last example is very significant because it is defined globally on T^*X, for any X.

1.3 Lecture 3

1.3.1 Naturality of the tautological form

Let $\varphi : M \to N$ be a diffeomorphism. Then there is a natural map $\mathrm{d}\varphi : TM \to TN$, and in turn, a natural map $\mathrm{d}\varphi^* : T^*N \to T^*M$. Explicitly, $\mathrm{d}\varphi$ and $\mathrm{d}\varphi^*$ are defined by

$$\mathrm{d}\varphi(v)(f) = v(\varphi^* f) = v(f \circ \varphi), \quad \mathrm{d}\varphi^*(\alpha) = \alpha \circ \mathrm{d}\varphi. \tag{1.3}$$

In fact, $\mathrm{d}\varphi^*$ is a diffeomorphism. This is immediate if $\mathrm{d}\varphi$ is a diffeomorphism. The latter is easily seen to be injective. For surjectivity, let $u \in TN$. Consider $v = u \circ (\varphi^{-1})^* \in TM$. Then $\mathrm{d}\varphi(v)(f) = u \circ (\varphi^{-1})^*(\varphi^* f) = u(f)$, since $(\varphi^{-1})^* = (\varphi^*)^{-1}$.

Hereafter we denote $\mathrm{d}\varphi^*$ by Φ. Recall that T^*M and T^*N are equipped with tautological 1-forms, λ_M and λ_N. We claim that $\mathrm{d}\Phi^* \lambda_M = \lambda_N$. (Here is a helpful? diagram for myself.)

$$\begin{tikzcd}
T(T^*N) \arrow[d, "\mathrm{d}\Phi"] \arrow[rr, bend left, "\mathrm{d}\pi_N"] & T^*N \arrow[d, "\Phi = \mathrm{d}\varphi^*"] \arrow[rr, bend left, "\pi_N"] & TN & N \\
T(T^*M) \arrow[rr, bend right, "\mathrm{d}\pi_M"'] & T^*M \arrow[rr, bend right, "\pi_M"'] & TM \arrow[u, "\mathrm{d}\varphi"] & M \arrow[u, "\varphi"]
\end{tikzcd}$$

Theorem 1.3.1. *Let $\varphi : M \to N$ be a diffeomorphism, and define $\Phi : T^*N \to T^*M$ as above. Then $\mathrm{d}\Phi^* \lambda_M = \lambda_N$, where $\lambda_{M,N}$ are the tautological forms on T^*M, T^*N.*

Proof. Choose any $v \in T_{(z,\zeta)}(T^*N)$. Choose the unique (x, ξ) such that $\varphi(x) = z$ and $\Phi_x \zeta = \xi$. Then

$$\begin{aligned} (\mathrm{d}\Phi^* \lambda_M)_{(z,\zeta)} v &= \lambda_{M\,\Phi(z,\zeta)}(\mathrm{d}\Phi(v)) \\ &= \lambda_{M\,(x,\xi)}(\mathrm{d}\Phi(v)) \\ &= \xi(\mathrm{d}\pi_M(\mathrm{d}\Phi(v))) \\ &= \zeta \circ \mathrm{d}\varphi_x(\mathrm{d}\pi_M(\mathrm{d}\Phi(v))). \end{aligned}$$

But by the commutative diagram above, $\mathrm{d}\varphi \circ \mathrm{d}\pi_M \circ \mathrm{d}\Phi = \mathrm{d}\pi_N$. Hence we have

$$(\mathrm{d}\Phi^*\lambda_M)_{(z,\zeta)}v = \zeta \circ \mathrm{d}\varphi_x(\mathrm{d}\pi_M(\mathrm{d}\Phi(v))) = \zeta(\mathrm{d}\pi_{N(z,\zeta)}(v)) = \lambda_N(v)$$

as required. $\square$

This can also be derived in coordinates. Let $p_i\mathrm{d}q^i$ be the tautological form of T^*N, and $\widetilde{p}_i\mathrm{d}\widetilde{q}_i$ the tautological form of T^*M. Since $\mathrm{d}q_i$ are contravariant while p_i are covariant, $\widetilde{p}_i\mathrm{d}\widetilde{q}_i$ is pulled back to $p_i\mathrm{d}q^i$.

Corollary 1.3.2. *Let $\omega_{M,N}$ denote the canonical 2-form on T^*M, T^*N respectively. Then with φ, Φ as above, $\mathrm{d}\Phi^*\omega_M = \omega_N$.*

In fact, a converse also holds. That is, if $\Phi : T^*N \to T^*M$ is a diffeomorphism such that $\mathrm{d}\Phi^*\lambda_M = \lambda_N$, then $\Phi = \mathrm{d}\varphi^*$ for some diffeomorphism $\varphi : M \to N$. To prove this, we introduce a certain canonical vector field.

1.3.2 Converse to the naturality result

Definition 1.3.3. Let X be a smooth manifold. The *Liouville vector field* V_X on T^*X is defined by $\iota_{V_X}\omega = \lambda$. Since ω is non-degenerate, this is well defined.

We now determine V in coordinates. Initially suppose $V = v^i\frac{\partial}{\partial p^i} + u^i\frac{\partial}{\partial q^i}$, and let $W = w^i\frac{\partial}{\partial p^i} + s^i\frac{\partial}{\partial q^i}$. Then we require $\omega(V, W) = \lambda(W)$, so in coordinates,

$$\begin{aligned}(\mathrm{d}p^j \wedge \mathrm{d}q^j)\Big(v^i\frac{\partial}{\partial p^i} + u^i\frac{\partial}{\partial q^i}, w^i\frac{\partial}{\partial p^i} + s^i\frac{\partial}{\partial q^i}\Big) &= v^j s^j - u^j w^j \\ &= p_j s^j = (p_j\mathrm{d}q^j)\Big(w^i\frac{\partial}{\partial p^i} + s^i\frac{\partial}{\partial q^i}\Big).\end{aligned}$$

Therefore, $V = p^i\frac{\partial}{\partial p^i}$. This is a "fibre-wise radial vector field". That is, given any $(x, \xi) \in T^*X$, $V_{(x,\xi)}$ is a vector in $T_{(x,\xi)}(T^*X)$ which points radially away from $(x, 0)$.

We now make some observations regarding V. Let M, N be manifolds, and $\Phi : T^*N \to T^*M$ a diffeomorphism. Since ω and λ are preserved by the pullback $\mathrm{d}\Phi^*$, and Φ is a diffeomorphism, the pushfoward of V_N to V_M is well defined, and

$$\mathrm{d}\Phi V_N = V_M. \tag{1.4}$$

Secondly, Let $\varphi^t_{N,M}$ denote the flows of $-V_{N,M}$ respectively. Then by equation 1.4, for all t,

$$\Phi \circ \varphi^t_N = \varphi^t_M \circ \Phi. \tag{1.5}$$

We are now ready to prove that Φ arises as $\mathrm{d}\varphi^*$ for some diffeomorphism $\varphi : M \to N$. This is done in three steps.

(1) Let Z_N be the zero section of T^*N. That is, the collection of all points $(z, 0)$ for $z \in N$. By the coordinate expression of the Liouville vector field, we see that it vanishes exactly on zero sections. But now by equation 1.4, V_M must vanish on $\Phi(Z_N)$. Therefore $\Phi(Z_N)$ is contained in Z_M (where the latter is the zero section of T^*M.) Since Φ is a diffeomorphism, arguing with the inverse shows that Z_N is mapped onto Z_M.

(2) We now show a stronger result, that Φ maps cotangent fibres to cotangent fibres. By our explicit form of the Liouville vector field, given any (z, ζ) in T^*N, we have $\varphi_N^t(z, \zeta) \to (z, 0)$ as $t \to \infty$. Let $\Psi : Z_N \to Z_M$ be the restriction of Φ to the zero sections as above. We wish to show that Φ sends (z, ζ) to the fibre above $\Psi(z, 0)$. But this is true by equation 1.5, since

$$\varphi_M^t(\Phi(z, \zeta)) = \Phi(\varphi_N^t(z, \zeta)) \to \Psi(z, 0) \quad \text{as} \quad t \to \infty,$$

and φ_M^t is purely radial, so $\Phi(z, \zeta)$ must lie above $\Psi(z, 0)$.

(3) The remainder of the proof is a coordinate calculation. Let p_i, q_i denote canonical coordinates on T^*M, and $\widetilde{p}_i, \widetilde{q}_i$ denote canonical coordinates on T^*N. Let $\varphi : M \to N$ be the diffeomorphism induced by Ψ. Since we showed above that cotangent fibres map to cotangent fibres, we must have $\widetilde{q}_i = \varphi(q_i)$. On the other hand, there is some map ρ such that $p_i = \rho(\widetilde{p}_i, \widetilde{q}_i)$. This gives

$$\rho(\widetilde{p}_i, \widetilde{q}_i)\mathrm{d}q^i = p_i \mathrm{d}q^i = \widetilde{p}_i \mathrm{d}\widetilde{q}^i = \widetilde{p}_i \mathrm{d}(\varphi(q_i)) = \widetilde{p}_i \mathrm{d}\varphi \circ \mathrm{d}q_i = (\mathrm{d}\varphi^* \widetilde{p}_i)\mathrm{d}q^i.$$

But this forces $\rho(\widetilde{p}_i, \widetilde{q}_i) = \mathrm{d}\varphi^* \widetilde{p}_i$. Since Φ sends $\widetilde{p}_i$ to p_i, this proves that $\Phi = \mathrm{d}\varphi^*$. In summary,

Theorem 1.3.4. *Let $\Phi : T^*N \to T^*M$ be a diffeomorphism such that $\mathrm{d}\Phi^* \lambda_M = \lambda_N$. Then $\Phi = \mathrm{d}\varphi^*$ for some diffeomorphism $\varphi : M \to N$.*

Example. In the above theorem, is not enough that Φ pulls back symplectic forms. (That is, it is essential that it pulls back the tautological 1-form.) To see this, consider the cotangent bundle of the circle, $T^*\mathbb{S}^1 = \mathbb{R}/\mathbb{Z} \times \mathbb{R}$, with p, q coordinates. Consider a Dehn twist $T : T^*\mathbb{S}^1 \to T^*\mathbb{S}^1$, with the latter space expressed in $\widetilde{p}, \widetilde{q}$ coordinates. The twist leaves p coordinates unchanged, while q gets mapped to $\widetilde{q} = q + f(p)$. Therefore

$$\mathrm{d}\widetilde{p} \wedge \mathrm{d}\widetilde{q} = \mathrm{d}p \wedge (\mathrm{d}q + \mathrm{d}f \mathrm{d}p) = \mathrm{d}p \wedge \mathrm{d}q.$$

That is to say, T is a diffeomorphism which preserves the canonical symplectic form. However, to see that T does not arise as $\mathrm{d}\varphi^*$ for some $\varphi : \mathbb{S} \to \mathbb{S}$, it suffices to show that T does not preserve the tautological form. We have

$$\widetilde{p}\mathrm{d}\widetilde{q} = p(\mathrm{d}q + \mathrm{d}f \mathrm{d}p) \neq p\mathrm{d}q,$$

since $\mathrm{d}f$ is non-zero for a Dehn twist. In fact, this example is of great importance, since Dehn twists can be embedded in most surfaces (as their support is an annulus which can be made arbitrarily small).

1.4 Lecture 4

1.4.1 Sub-objects of symplectic vector spaces

Definition 1.4.1. Let (V, Ω) be a symplectic vector space, and $W \subset V$ a subspace. Then

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

is the *symplectic complement* or *orthogonal* of W.

Proposition 1.4.2. $\dim W^\Omega + \dim W = \dim V$.

Proof. Define $f : V \to W^*$ by $v \mapsto \Omega(v, \cdot)$. Then $\ker f = W^\Omega$. Non-degeneracy of Ω ensures that $\operatorname{im} f = W^*$. Therefore by the rank nullity theorem,

$$\dim W + \dim W^\Omega = \dim W^* + \dim W^\Omega = \dim \operatorname{im} f + \dim \ker f = \dim V.$$

$\square$

Proposition 1.4.3. $W = (W^\Omega)^\Omega$.

Proof. We first prove that $W \subset (W^\Omega)^\Omega$. Let $w \in W$. Choose $v \in W^\Omega$. Then $\Omega(w, v) = 0$, so $w \in (W^\Omega)^\Omega$. In fact, by proposition 1.4.2, we have $\dim W = \dim(W^\Omega)^\Omega$, so equality follows. $\square$

These two propositions raise the question of whether or not $W \oplus W^\Omega = V$ holds in general (as it does for inner products). We find that this is not true, giving rise to the following definitions:

Definition 1.4.4. Let W be a subspace of (V, Ω).

- W is *coisotropic* if $W^\Omega \subset W$.
- W is *isotropic* if $W \subset W^\Omega$.
- W is *Lagrangian* if it is isotropic and coisotropic.
- W is *symplectic* if $W \cap W^\Omega = \{0\}$.

Example. Suppose $W \subset V$ has dimension 1. Choose $w \in W$. Then for any other $w' \in W$, there exists a real number r such that $\Omega(w, w') = \Omega(w, rw) = r\Omega(w, w) = 0$. Since w is orthogonal to all $w' \in W$, we have $W \subset W^\Omega$, that is, W is isotropic. Similarly, every codimension 1 subspace is coisotropic.

Example. Consider $(V \oplus V^*, \Omega)$ where Ω is the canonical symplectic form. Then V and V^* are clearly Lagrangian subspaces. This raises the question of: what are all of the Lagrangian subspaces of $V \oplus V^*$?

Proposition 1.4.5. $W \subset V$ is Lagrangian if and only if W is (co)isotropic and half-dimensional. (i.e. $\dim W = \frac{1}{2}\dim V$.)

Proof. If W is Lagrangian, it is clearly half-dimensional and both isotropic and coisotropic. Conversely, suppose W is isotropic and half-dimensional. Then by 1.4.2, W is Lagrangian. $\square$

When is the graph of a linear map $f : V \to V^*$ Lagrangian? Since the dimension of the graph is the rank of f, by proposition 1.4.5, $G(f)$ must be isotropic and injective. That is, we must have $\Omega(v \oplus f(v), w \oplus f(w)) = 0$ for any $v \oplus f(v), w \oplus f(x)$ in $G(f)$, with f injective. But the former property is exactly the requirement that

$$f(w)v = f(v)w.$$

Therefore, the graph of a linear map $f : V \to V^*$ Lagrangian if $\widetilde{f} : V \times V \to \mathbb{R}$ defined by $\widetilde{f} : (v, w) \mapsto f(v)w$ is a symmetric bilinear form. A Lagrangian subspace arising as a graph is called *graphical*. The previous example shows that not every Lagrangian subspace is graphical (namely V^*).

1.4.2 Sub-objects of symplectic vector bundles and manifolds

Definition 1.4.6. Let $E \to M$ be a symplectic vector bundle, and $P \subset E$ a subbundle. P is *coisotropic (isotropic, Lagrangian, symplectic)*, if for every $x \in M$, the fibre $P_x \subset E_x$ is a coisotropic (isotropic, Lagrangian, symplectic) subspace.

Definition 1.4.7. Let M be a symplectic manifold, and $N \subset M$ a submanifold. N is *coisotropic (isotropic, Lagrangian, symplectic)*, if $TN \subset TM|_N$ is a coisotropic (isotropic, Lagrangian, symplectic) subbundle.

Example. By proposition 1.4.5, a submanifold N of a symplectic manifold (M, ω) is Lagrangian if $\omega|_N = 0$ and $\dim N = \frac{1}{2}\dim M$.

If a vector bundle is of the form $E \oplus E^* \to X$, then E and E^* are Lagrangian subbundles. When is a symplectic vector bundle of the form $E \oplus E^*$?

Theorem 1.4.8. *Let $P \to X$ be a symplectic vector bundle. Then $P \to X$ is symplectomorphic to $E \oplus E^* \to X$ for some E if and only if $P \to X$ admits a Lagrangian subbundle.*

The "only if" direction is clear; we prove this theorem in lecture 7.

Remark. There exist symplectic vector bundles which admit no Lagrangian subbundles, but do admit half-rank subbundles. This is explored at a later point.

What are some Lagrangian submanifolds in symplectic manifolds?

Example. Let Σ be a surface equipped with an area-form, and L any one dimensional submanifold. Then L is Lagrangian. Note that an area form on a surface is a symplectic form. Let $x \in L$. We have that $T_x L$ is a half dimensional subspace of $T_x \Sigma$. Moreover, it is a one dimensional subspace, so it is isotropic. Therefore it is a Lagrangian subspace. Thus L is a Lagrangian submanifold.

Example. Let $M = T^*X$ be the symplectic manifold equipped with the canonical form $d\lambda$. Let $Z \subset T^*X$ denote the *zero section*, $Z = \{(x, 0) \in T^*X\}$. Then Z is a Lagrangian submanifold. To see this, note that

$$\mathrm{d}\lambda = \mathrm{d}p^1 \wedge \mathrm{d}q^1 + \cdots + \mathrm{d}p^n \wedge \mathrm{d}q^n.$$

Restricting to Z forces each p to be identically zero. Thus $\mathrm{d}\lambda|_Z = 0$, so Z is isotropic. Since Z is half-dimensional, Z is Lagrangian. Similarly, any cotangent fibre is Lagrangian.

Proposition 1.4.9. Let $s : X \to T^*X$ be a section of the canonical projection. Then $s^*\lambda = s$, where λ denotes the tautological form on T^*X.

Proof. Choose a point $x \in X$. Then $s(x) = (x, \xi)$ for some $\xi \in T_x^*X$. The tautological form at $s(x)$ is given by $\xi \circ \mathrm{d}\pi$. But then for any $v \in T(T * X)$, we have

$$(s^*\lambda_{s(x)})(v) = \lambda_{(s(x))}(\mathrm{d}sv) = \xi(\mathrm{d}\pi(\mathrm{d}sv)) = \xi(v).$$

The third equality is because $\pi \circ s = \mathrm{id}_X$, so $\mathrm{id}_{TX} = \mathrm{d}(\pi \circ s) = \mathrm{d}\pi \circ \mathrm{d}s$. Since $\xi(v)$ is $s_x(v)$, we are done. □

In fact, the tautological one-form is the unique horizontal one-form that "cancels" a pullback. It follows that if ω is the canonical form, then

$$s^*\omega = s^*(\mathrm{d}\lambda) = \mathrm{d}(s^*\lambda) = \mathrm{d}s.$$

1.5 Lecture 5

1.5.1 More Lagrangian examples

In lecture 4, we defined what it means for a submanifold of a symplectic manifold to be Lagrangian. We can generalise this definition by stating it in terms of embeddings.

Definition 1.5.1. Let (M, ω) be a $2n$ dimensional symplectic manifold. Let $f : N \to M$ be an embedding of an n-manifold. f is a *Lagrangian embedding* if $f^*\omega = 0$.

This implies that $\operatorname{im} f$ is isotropic and half dimensional, hence a Lagrangian submanifold.

Example. As a corollary of 1.4.9, we obtain a new class of Lagrangian submanifolds. Let s be a section of T^*X, and suppose s is closed. Then s is a Lagrangian embedding of X, since $s^*\omega = \mathrm{d}s = 0$. Therefore $\operatorname{im} s$ is a Lagrangian submanifold. In fact, since s is a section of the cotangent bundle, $\operatorname{im} s$ can be interpreted as the graph of a function.

This example is surprisingly deep, as the calculation also shows that $\operatorname{im} s$ is a Lagrangian submanifold *if and only if s is closed.* There is a natural bijective correspondence between Lagrangian submanifolds of T^*X which project smoothly onto X and closed 1-forms on X.

We now create a family of graphical Lagrangian submanifolds. Recall that a symmetric bilinear form $V \times V \to \mathbb{R}$ gives rise to a graphical Lagrangian subspace of $V \oplus V^*$.

Let $(X_1, \omega_1), (X_2, \omega_2)$ be symplectic manifolds, and $\varphi : X_1 \to X_2$ a symplectomorphism. The canonical symplectic structure on the product space is $\pi_1^*\omega_1 + \pi_2^*\omega_2$. However, we can also define a *twisted product* by

$$X_1 \widetilde{\times} X_2 = (X_1 \times X_2, \pi_1^*\omega_1 - \pi_2^*\omega_2).$$

Moreover, the graph $G(\varphi)$ is a Lagrangian submanifold of $X_1 \widetilde{\times} X_2$. To see this, define the embedding $\Phi : X_1 \to X_1 \widetilde{\times} X_2$ by $\Phi : x \to (x, \varphi(x))$, so that $\operatorname{im} \Phi = G(\varphi)$. Then we need only show that Φ is a Lagrangian embedding, i.e. it pulls back the symplectic form to 0. We see that

$$\Phi^*(\pi_1^*\omega_1 - \pi_2^*\omega_2) = \omega_1 - \varphi^*\omega_2 = \omega_1 - \omega_1 = 0,$$

so Φ is a Lagrangian embedding as required.

In fact, the calculation shows that the converse also holds. If the graph of a diffeomorphism $\varphi : X_1 \to X_2$ is *not* Lagrangian in $X_1 \widetilde{\times} X_2$, then φ is not a symplectomorphism. This gives some insight into the philosophy of Weinstein that "symplectomorphisms are special cases of Lagrangian submanifolds".

1.5.2 Normal form theorem for symplectic vector spaces

At the beginning of lecture 1, Darboux's theorem was stated as motivation for the requirement that a symplectic form be closed. We now prove Darboux's theorem in the algebraic setting, that is, every symplectic vector space V arises as $W \oplus W^*$ equipped with the canonical form, for some half dimensional vector space W. The proof approach is as follows: first we show that V admits a Lagrangian subspace W. Then we show that W admits a Lagrangian complement, $W \oplus W' = V$. Finally we use this to construct the canonical basis of V.

First choose any 1-dimensional subspace W_1 of V. Then W_1 is isotropic. If W_1 is half-dimensional, we are done. Otherwise, choose any $w \in W_1^\Omega$ such that $w \notin W_1$, and define W_2 to be the span of W_1 and w. Since $w \in W_1^\Omega$, it is orthogonal to both W_1 and its own span, giving $W_2 \subset W_2^\Omega$. Inductively this process continues until $W = W_n$ is half-dimensional and isotropic, thus Lagrangian.

Next, suppose $W \subset V$ is Lagrangian. Let $w_1' \in V \setminus W$. Then W_1', (the span of w_1'), is isotropic and satisfies $W_1' \cap W = \{0\}$. If W_1' is half-dimensional, we are done. Otherwise choose $w_2' \in V \setminus \{W \cup W_1'\}$. As above, this is still isotropic and trivial intersects W. Note that 1.4.2 ensures the existence of w_2'. Inductively following this process, eventually $W' = W_n'$ is half dimensional, at which point it is also a Lagrangian subspace and has trivial intersection with W. Since they are both half-dimensional, it follows that $W \oplus W' = V$.

Now consider the map $\alpha : W' \to W^*$ defined by $\alpha(w') : w \mapsto \Omega(w, w')$. If w' is an element of $\ker \alpha$, then w' belongs to $W^\Omega = W$. Since W intersects W' trivially, $w' = 0$. But now α is an injection between two equal-dimensional vector spaces, so it is an isomorphism. This induces an isomorphism $\varphi : W \oplus W' \to W \oplus W^*$. It remains to prove that φ is in fact a symplectomorphism from $(W \oplus W', \Omega)$ to $(W \oplus W^*, \Omega_W)$.

For any $w \oplus w', v \oplus v' \in W \oplus W'$, we have

$$\begin{aligned}
(\varphi^*\Omega_W)(w \oplus w', v \oplus v') &= \Omega_W(\varphi(w \oplus w'), \varphi(v \oplus v')) \\
&= \Omega_W(w \oplus \alpha(w'), v \oplus \alpha(v')) \\
&= \alpha(v')w - \alpha(w')v \\
&= \Omega(w, v') - \Omega(v, w') = \Omega(w \oplus w', v \oplus v').
\end{aligned}$$

The last equality comes from the fact that $W \oplus W'$ is a Lagrangian splitting, so the cross terms vanish. Therefore φ is a symplectomorphism. In summary, we have the following theorem.

Theorem 1.5.2. *Let (V, Ω) be a symplectic vector space. Let $W \subset V$ be a Lagrangian subspace. (Such a subspace exists.) Then (V, Ω) is symplectomorphic to $(W \oplus W^*, \Omega_W)$ (where Ω_W denotes the canonical symplectic form on $W \oplus W^*$).*

The above theorem statement doesn't give any insight as to the amount of freedom we have in choosing the symplectomorphism between V and $W \oplus W^*$. We find that topologically there is only one choice of symplectomorphism (in the sense that the collection of choices is a contractible space).

To see this, first note that the above proof establishes the following bijective correspondence:

$$\left\{ \begin{array}{c} \text{Symplectomorphisms} \\ V \to W \oplus W^* \text{ which extend} \\ \text{the identity } W \to W. \end{array} \right\} \longleftrightarrow \left\{ \begin{array}{c} \text{Lagrangian complements to} \\ W \text{ inside } V. \end{array} \right\}$$

But a Lagrangian complement to W inside $V \cong W \oplus W^*$ projects onto W^*, so it is a graphical Lagrangian over W^*. Thus it corresponds to a symmetric bilinear form $W^* \times W^* \to \mathbb{R}$. In summary we have the bijective correspondence

$$\left\{ \begin{array}{c} \text{Symplectomorphisms} \\ V \to W \oplus W^* \text{ which extend} \\ \text{the identity } W \to W. \end{array} \right\} \longleftrightarrow \left\{ \begin{array}{c} \text{Symmetric bilinear forms} \\ W^* \times W^* \to \mathbb{R}. \end{array} \right\}$$

But the right hand side is a contractible space, since given any two symmetric bilinear forms g and g', $(1-t)g + tg'$ is a symmetric bilinear form.

Chapter 2

Complex geometry

2.1 Lecture 6

2.1.1 Corollaries of the normal form theorem

We can now classify all skew symmetric forms on vector spaces:

Corollary 2.1.1. *Let (V, Ω) be a vector space equipped with a skew-symmetric form. Then there exists a basis $e_1, \ldots, e_n, f_1, \ldots, f_n, w_1, \ldots, w_m$ of V such that for all i, j,*

$$\Omega(e_i, e_j) = \Omega(f_i, f_j) = \Omega(w_i, \cdot) = 0, \quad \Omega(e_i, f_j) = \delta_{ij}.$$

Proof. Let $W = \{w \in V : \Omega(w, \cdot) = 0\}$. Choose any $\widetilde{V} \subset V$ such that $W \oplus \widetilde{V} = V$. Then Ω restricted to $\widetilde{V}$ is a symplectic form. Let $e_1, \ldots, e_n, f_1, \ldots, f_n$ be the canonical basis for $\widetilde{V}$, and $w_1, \ldots, w_m$ any basis for W. $\square$

Corollary 2.1.2. *Suppose Ω is a skew-symmetric bilinear form on a vector space of dimension $2n$. Then Ω is non-degenerate if and only if $\Lambda^n \Omega \neq 0$.*

Proof. This follows immediately from expressing Ω in terms of the basis in the above corollary. $\square$

Corollary 2.1.3. *If ω is a symplectic form on a manifold of dimension $2n$, then $\Lambda^n \omega$ is a non-vanishing $2n$-form, i.e. a volume form.*

2.1.2 Complex structures on symplectic vector spaces

Definition 2.1.4. Let V be a real vector space. A *complex structure* on V is a linear map $J : V \to V$ such that $J^2 = \text{id}$. Let (V, Ω) be a symplectic vector space. A complex structure J on (V, Ω) is said to be

- *almost compatible* with Ω if $\Omega = J^*\Omega$ (i.e. $\Omega(v, w) = \Omega(Jv, Jw)$ for all $v, w \in V$),

- *tame* if $\Omega(v, Jv) > 0$ for all non-zero $v \in V$,
- *compatible* if it is almost compatible and tame.

As with the various sub-objects in the previous two lectures, the definition of a compatible complex structure extends to symplectic vector bundles and symplectic manifolds in the obvious way. For example, if (M, ω) is a symplectic manifold, then a *compatible almost complex structure* on M is a smooth section J of $\mathrm{End}(TM) \to M$ such that for each $x \in M$, $J_x^2 = -\mathrm{id}_{T_xM}$, and J_x is compatible with ω_x.

Remark. We referred to J above as an *almost complex structure* rather than a complex structure. This is because a smooth manifold is called *complex* if it has charts mapping into $\mathbb{C}^n$ with holomorphic transition maps. Every complex structure can be realised as an almost complex structure, but there exist 4-real-dimensional almost complex manifolds which are not complex. It is still open as to whether there are such manifolds in higher dimensions.

The definition of compatibility above gives rise to the famous *2-out-of-3 property* of symplectic structures, complex structures, and inner products.

Definition 2.1.5. An inner product g is said to be *compatible* with a complex structure J if and only if $g = J^*g$.

Proposition 2.1.6. Let (V, Ω, J) be a symplectic vector space equipped with a complex structure. Then $\Omega(\cdot, J\cdot)$ is an inner product if and only if J is compatible with Ω. Moreover, $\Omega(\cdot, J\cdot)$ is compatible with J.

Proof. Suppose J is compatible. Then for any $v, w \in V$,

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

Therefore $\Omega(\cdot, J\cdot)$ is symmetric. Moreover, since J is tame, $\Omega(\cdot, J\cdot)$ is positive definite, so it is an inner product. The converse is similar. Compatibility of $\Omega(\cdot, J\cdot)$ with J is immediate from the almost-compatibility of Ω with J. □

Proposition 2.1.7. Let (V, g, J) be an inner product space equipped with a compatible complex structure. Then $g(J\cdot, \cdot)$ is a symplectic form which is compatible with J.

Proof. Let $v, w \in V$. Then

$$g(Jv, w) = g(J^2v, Jw) = -g(v, Jw) = -g(Jw, v),$$

so $g(J\cdot, \cdot)$ is skew-symmetric. It is non-degenerate since the induced map $V \to V^*$ is defined by $-\widetilde{g} \circ J$, where $\widetilde{g}$ is the isomorphism from V to V^* given by $v \mapsto g(\cdot, v)$, and J is an automorphism of V. Compatibility is again immediate. □

The final direction is to prove that given a symplectic vector space equipped with an inner product, there is a canonical complex structure which is compatible with both the symplectic structure and inner product. This takes more work - we must introduce polar decompositions.

Proposition 2.1.8. Let (V, g) be a real inner product space. Let $T : V \to V$ be a linear bijection. Then there exist unique $P, \Theta : V \to V$ such that

- P is positive definite and self-adjoint (i.e. $g(Pv, v) > 0$ whenever $v \in V$ is non-zero, and $g(Pv, w) = g(v, Pw)$ for all $v, w \in V$).
- Θ is isometric (i.e. $g(\Theta v, \Theta w) = g(v, w)$ for all $v, w \in V$).
- $T = P\Theta$.

Proof. Consider the isomorphism $TT^* : V \to V$ (where V^* is the adjoint of V). TT^* is self adjoint and positive definite. By the spectral theorem, TT^* admits a positive definite square root - i.e. a map $P : V \to V$ such that $P^2 = TT^*$. Moreover, P is self-adjoint. Now define $\Theta = P^{-1}T$. Θ is an isometry if $\Theta^*\Theta = \mathrm{id}$, since we then have

$$g(\Theta v, \Theta w) = g(v, \Theta^*\Theta w) = g(v, w).$$

This is indeed the case, because

$$\Theta^*\Theta = (P^{-1}T)^*(P^{-1}T) = T^*P^{-1}P^{-1}T = T^*(P^2)^{-1}T = T^*(TT^*)^{-1}T = \mathrm{id}.$$

Clearly this construction gives $T = P\Theta$.

We next show uniqueness. Suppose $T = P'\Theta'$ for some isometry Θ' and positive definite self adjoint P'. Then $TT^* = (P'\Theta')(P'\Theta')^* = P'\Theta'\Theta'^*P' = P'^2$. Since positive definite self adjoint operators admit unique positive definite square roots, $P' = P$, and this forces $\Theta' = \Theta$. □

Proposition 2.1.9. Let (V, g) be a real inner product space, and $T : V \to V$ a linear bijection. Let $T = P\Theta$ be the polar decomposition of T. if T is normal ($TT^* = T^*T$) then $T = \Theta P$.

Proof. Since T is normal, we have

$$P^2 = TT^* = T^*T = (\Theta^*P)P\Theta = \Theta^*P^2\Theta.$$

Since positive definite self adjoint matrices admit unique positive definite square roots, it must be the case that P is the unique positive definite square root of $\Theta^*P^2\Theta$. However, we see that $\Theta^*P\Theta$ is positive definite since $g(\Theta^*P\Theta v, v) = g(P\Theta v, \Theta v) > 0$ for all $v \neq 0$, and moreover, it is a square root of $\Theta^*P^2\Theta$. It follows that $P = \Theta^*P\Theta$, so $\Theta P = P\Theta$. □

2.2 Lecture 7

2.2.1 2-out-of-3 property

In lecture 6 we showed that symplectic + complex ⇒ inner product, and inner product + complex ⇒ symplectic. We now complete the result by establishing that symplectic + inner product ⇒ complex.

More precisely, we establish that given a symplectic vector space (V, Ω) there is a canonical map Θ which admits a canonical section s, giving rise to the following correspondence:

$$\{\text{Inner products } g \text{ on } V.\} \underset{s}{\overset{\Theta}{\rightleftarrows}} \left\{ \begin{array}{c} \text{Complex structures on } V, \\ \text{compatible with } \Omega. \end{array} \right\}$$

Moreover, given any g, we show that $\Theta(g)$ is compatible with g.

Remark. While this is referred to as a correspondence, note that in general, many inner products may map to the same complex structure.

Suppose g is an inner product on (V, Ω). Define $K_g : V \to V$ by $g(K_g v, w) = \Omega(v, w)$ for all $v, w \in V$. Since g and Ω are non-degenerate, this gives a well defined map. Since K_g is a linear bijection, it has a unique polar decomposition (by proposition 2.1.8) into $K_g = P_g J_g$, where J_g is isometric and P_g is positive definite and self adjoint. Define the map Θ by

$$\Theta : g \mapsto J_g.$$

We claim that θ is the desired map in the above correspondence. First we show that for any g, J_g is indeed a complex structure compatible with Ω. Observe that K_g is skew adjoint in the sense that

$$g(K_g v, w) = \Omega(v, w) = -\Omega(w, v) = -g(K_g w, v) = g(v, -K_g w)$$

for all $v, w \in K$. This implies that K_g is normal, so $P_g J_g = J_g P_g$. In particular, this implies that $J_g = P_g^{-1} K_g = K_g P_g^{-1}$. But then

$$J_g = K_g P_g^{-1} = -K_g^*(P_g^*)^{-1} = -(P_g^{-1} K_g)^* = -J_g^*,$$

so $J_g^2 = -J_g^* J_g = -\text{id}$. J_g is almost compatible, since for any $v, w \in V$,

$$\Omega(J_g v, J_g w) = g(K_g J_g v, J_g w) = g(P_g J_g v, w) = g(K_g v, w) = \Omega(v, w).$$

J_g is tame, since for any $v \in V$ non-zero, P_g positive definite gives

$$\Omega(v, J_g v) = g(K_g v, J_g v) = g(P_g J_g v, J_g v) > 0.$$

Therefore J_g is a compatible complex structure as required. In fact, by construction, J_g is isometric, so it is compatible with any metric in $\Theta^{-1}(J_g)$. Finally, we claim that the map

s defined by $J \mapsto \Omega(\cdot, J\cdot)$ is a section of Θ. Choose any J compatible with Ω. Let $\widetilde{\Theta}$ be defined by $g \mapsto K_g$ (so that $\Theta(g)$ is the isometric part of $\widetilde{\Theta}(g)$.) Then for any $v, w \in V$,

$$\Omega(-J\widetilde{\Theta}(s(J))v, w) = \Omega(\widetilde{\Theta}(s(J))v, Jw) = s(J)(\widetilde{\Theta}(s(J))v, w) = \Omega(v, w).$$

It follows that $J\widetilde{\Theta}(s(J)) = -\mathrm{id}$, so $J = \widetilde{\Theta}(s(J))$. In particular, it also follows that $\widetilde{\Theta}(s(J))$ is isometric, so $\widetilde{\Theta}(s(J)) = \Theta(s(J))$ by the uniqueness of the polar decomposition. Therefore s is a section of Θ. This result can be summarised by the following theorem:

Theorem 2.2.1 (2-out-of-3)**.** *Let (V, g, ω, J) be a symplectic inner product space equipped with a complex structure. Suppose $g(\cdot,\cdot) = \omega(\cdot, J\cdot)$. Then any one structure (g, ω, or J) is uniquely determined by the other two structures.*

Corollary 2.2.2. *The space C of compatible complex structures on a symplectic vector space (V, Ω) is contractible.*

Proof. The space of all inner products on V is contractible (since it is convex). Choose any inner product g on V. Using the correspondence above, define $\varphi : C \times [0, 1] \to C$ by

$$\varphi_t : J \mapsto \Theta \circ ((1-t)s(J) + tg).$$

This is a contraction, since $\varphi_0 = \mathrm{id}_S$, $\varphi_1 = \Theta(g)$. □

Corollary 2.2.3. *Every symplectic vector bundle (symplectic manifold) admits a compatible complex structure. Moreover, any two compatible complex structures are homotopic.*

There is also a canonical compatible complex structure that can be equipped on a symplectic vector space, given canonical coordinates. Suppose (V, Ω) is symplectic, with canonical basis $e_1, \dots, e_n, f_1, \dots, f_n$. Then define J component-wise by

$$Je_i = f_i, \quad Jf_i = -e_i.$$

It is easy to see that J is a compatible complex structure.

One can also ask: in what way does a choice of compatible complex structure add rigidity to a symplectic vector space? One answer is that it provides Lagrangian complements: Suppose $W \subset V$ is a Lagrangian subspace. Then $JW \cap W$ is trivial, and JW is Lagrangian. Therefore $JW \oplus W = V$. To see this, consider the inner product g defined by $g(\cdot,\cdot) = \Omega(\cdot, J\cdot)$. Since W is Lagrangian, we have

$$\begin{aligned} W &\subset \{v \in V : \Omega(v, w) = 0 \text{ for all } w \in W\} \\ &= \{v \in V : g(v, w) = 0 \text{ for all } w \in JW\}. \end{aligned}$$

Because g is an inner product, this shows that $W \cap JW$ is trivial. To see that JW is Lagrangian, we observe that

$$\begin{aligned} JW &= \{v \in V : g(v, w) = 0 \text{ for all } w \in W\} \\ &= \{v \in V : \Omega(v, w) = 0 \text{ for all } w \in JW\} = JW^{\Omega}. \end{aligned}$$

As a corollary, we can finally give a simple proof that a symplectic vector bundle admits a Lagrangian subbundle if and only if it splits into the canonical form.

Theorem 2.2.4. *Let $P \to X$ be a symplectic vector bundle. If P admits a Lagrangian subbundle L, then P arises as a sum bundle $L \oplus L^*$ equipped with the canonical symplectic form.*

Proof. Suppose P admits a Lagrangian subbundle L. Since every real vector bundle admits a metric, by the 2-out-of-3 property, P admits a complex structure. But now JL is a complementary Lagrangian subbundle to L. By the proof of the normal form theorem, $L \oplus JL$ is isomorphic to $L \oplus L^*$. $\square$

2.3 Lecture 8

2.3.1 Exploration of complex vector bundles following the 2-out-of-3 property

In the previous lecture it was shown that every symplectic vector bundle admits a compatible complex structure, and any two such structures are homotopic. It turns out that this still holds when we interchange the roles of the symplectic and complex structures. In fact, it follows that any two symplectic vector bundles are isomorphic if and only if their corresponding complex vector bundles are isomorphic.

Proposition 2.3.1. Every symplectic vector bundle admits a compatible complex structure. Moreover, any two compatible complex structures are homotopic. (Proved previously.)

Proposition 2.3.2. Every complex vector bundle admits a compatible symplectic structure. Moreover, any two compatible symplectic structures are homotopic.

Proof. This is a counterpart to correspondence between inner products and compatible complex structures established in lecture 7. Let (V, J) be a complex vector space. It is easy to verify that the following (not necessarily bijective) correspondence holds, where s, defined by $\Omega(\cdot, \cdot) \mapsto \Omega(\cdot, J\cdot)$, and Φ is constructed in a similar manner to the correspondence in lecture 7:

$$\{\text{Inner products } g \text{ on } V.\} \underset{s}{\overset{\Phi}{\rightleftharpoons}} \left\{ \begin{array}{c} \text{Symplectic structures on } V, \\ \text{compatible with } J. \end{array} \right\}$$

An analogous construction to the proof of corollary 2.2.2 shows that the space of compatible symplectic structures on V is contractible. The desired result follows. $\square$

Theorem 2.3.3. *Let (E_1, ω_1, J_1) and (E_2, ω_2, J_2) be symplectic vector bundles over X equipped with compatible complex structures. The vector bundles are isomorphic as symplectic vector bundles if and only if they are isomorphic as complex vector bundles.*

Proof. First suppose (E_1, ω_1, J_1) and (E_2, ω_2, J_2) are isomorphic as complex vector bundles. Then there is a vector bundle isomorphism $\Phi : E_1 \to E_2$ such that $\Phi^* J_2 = J_1$. Therefore $\Phi^*\omega_2$ and ω_1 are both compatible with J_1. It follows by the previous proposition that there is a one parameter family ω_t of compatible structures, with $\omega_0 = \Phi^*\omega_2$. But now we can find a smooth family of vector bundle isomorphisms $\Psi_t : E_1 \to E_1$ such that $\Psi_t^* J_t = J_1$ (McDuff-Salamon). In particular, it follows that $\Phi \circ \Psi_0$ is a vector bundle isomorphism preserving the complex structure. The converse is very similar. $\square$

Corollary 2.3.4. *Suppose (E, ω) is a symplectic vector bundle and J_1, J_2 are two compatible complex structures. Then (E, J_1) and (E, J_2) are isomorphic as complex vector bundles.*

Example. Suppose $E \oplus E^* \to X$ is a symplectic vector bundle. By the above, up to isomorphism, there is only one complex structure on the vector bundle. What is the associated complex structure? It is in fact the complexification of $E \to X$:

Given a real vector bundle $E \to X$, the *complexification* of E is the vector bundle $E \oplus E \to X$ equipped with the complex structure $J' : (e, \widetilde{e}) \mapsto (-\widetilde{e}, e)$, which we denote $E_{\mathbb{C}} \to X$. Equivalently, the complexification is the tensor bundle $E \otimes \mathbb{C} \to X$. Complex multiplication is then defined by $\alpha(v \otimes \beta) = v \otimes (\alpha\beta)$.

At the end of lecture 7 we showed that if L is a Lagrangian subbundle of P, then $L \oplus JL$ is a Lagrangian decomposition of P. Therefore by the previous theorem, every compatible complex structure on $E \oplus E^*$ is isomorphic to the complexification $E_{\mathbb{C}} \cong E \oplus JE$.

The above result has a "reverse analogue" for symplectic subbundles:

Proposition 2.3.5. Let (V, Ω, J) be a symplectic vector space equipped with a tame complex structure (not necessarily compatible). A subspace $W \subset V$ is symplectic if $JW \subset W$.

Proof. Let $v \in W \cap W^{\Omega}$. Suppose $JW \subset W$. Then $Jv \in W$, so $\Omega(v, Jv) = 0$. Thus $v = 0$. $\square$

Corollary 2.3.6. *Let N be an almost complex submanifold of a symplectic manifold (M, ω, J) equipped with a tame almost complex structure. Then N is symplectic. In particular, if N is an almost complex submanifold of an almost complex manifold which admits a compatible symplectic form, then N is symplectic.*

At the end of the previous lecture, we proved that a symplectic vector bundle splits as $E \oplus E^*$ if and only if it admits a Lagrangian subbundle. A natural follow up question is whether or not there are symplectic vector bundles which admit half dimensional subbundles but no Lagrangian subbundles. This can be answered in the affirmative using Chern classes. We now proceed with a number of results whose proofs will be omitted, since they are beyond the scope of the course (at least for now).

The *Chern classes* are invariants associated to complex vector bundles. These have corresponding *Chern numbers* which are numerical invariants. In the case of complex vector bundles over oriented closed 2-real-dimensional manifolds (which we call *surfaces*), the first Chern number completely captures the first Chern class. According to McDuff-Salamon, the Chern number can be defined axiomatically as follows:

Theorem 2.3.7. *There exists a unique map $c_1 : E \to \mathbb{Z}$ called the* first Chern number*, which assigns an integer to every complex vector bundle E over a surface Σ and satisfies the following axioms:*

1. *Two complex vector bundles over Σ are isomorphic if and only if they have the same rank and first Chern number.*

2. *Whenever $\varphi : \Sigma_2 \to \Sigma_1$ is a smooth map between surfaces and $E \to \Sigma_1$ is a complex vector bundle,*
$$c_1(\varphi^* E) = \deg(\varphi) c_1(E).$$

3. *If E_1, E_2 are complex vector bundles over Σ,*
$$c_1(E_1 \oplus E_2) = c_1(E_1) + c_1(E_2.)$$

4. *The first Chern number of the tangent bundle of Σ is the Euler characteristic of Σ.*

By theorem 2.3.3, the first Chern number can also be defined as an invariant of symplectic vector bundles over surfaces. Moreover, the axioms imply that the first Chern number vanishes if and only if the vector bundle is trivial. Thus the first Chern number is an obstruction to triviality. We now use the first Chern number to show that there exist symplectic vector bundles admitting half-dimensional subbundles, but no Lagrangian subbundles.

Proposition 2.3.8. The first Chern number of a complexification of a real vector bundle over a surface is zero.

Every surface Σ has an orientation reversing automorphism φ. Let $E \to \Sigma$ be a complex vector bundle. Then $\varphi^* E = \overline{E}$, the "conjugate bundle". Therefore by the second axiom above, $c_1(\overline{E}) = -c_1(E)$. Now suppose E is a complexification of a real vector bundle R, so that $J : R \oplus R \to R \oplus R$ is defined by $J : (e, \widetilde{e}) \mapsto (-\widetilde{e}, e)$. Then the conjugate complex structure is $\overline{J} : (e, \widetilde{e}) \mapsto (\widetilde{e}, -e)$. Clearly E and $\overline{E}$ are isomorphic complex vector bundles, by the bundle map $(e, \widetilde{e}) \mapsto (\widetilde{e}, e)$. Therefore by the first axiom above, $c_1(\overline{E}) = c_1(E)$. It follows that $c_1(E) = 0$.

Example. Let T be the tangent bundle of $\mathbb{S}^2$, and $\mathbb{R}_{\mathbb{C}}$ the complexification of the trivial real line bundle over $\mathbb{S}^2$. Then T admits a symplectic structure (e.g. the normalised area form of $\mathbb{S}^2$) and $\mathbb{R}_{\mathbb{C}}$ admits a compatible symplectic structure. Therefore $T \oplus \mathbb{R}_{\mathbb{C}} \to \mathbb{S}^2$

is a symplectic vector bundle, and it clearly admits a half-dimensional subbundle. On the other hand, if $T \oplus \mathbb{R}_{\mathbb{C}} \to \mathbb{S}^2$ admits a Lagrangian subbundle L, then it splits as $L \oplus L^*$ by theorem 2.2.4. But then any compatible complex structure is isomorphic to the complexification $L_{\mathbb{C}} \to \mathbb{S}^2$, and so its first Chern number must vanish. But we have $c_1(T \oplus \mathbb{R}_{\mathbb{C}}) = c_1(T) + c_1(\mathbb{R}_{\mathbb{C}}) = 2$, so $T \oplus \mathbb{R}_{\mathbb{C}}$ does not admit Lagrangian subbundles.

2.3.2 Integrability of almost complex manifolds

Definition 2.3.9. An *almost complex manifold* is a smooth manifold whose tangent bundle is equipped with a complex structure. A *complex manifold* is a smooth manifold covered by charts mapping into open subsets of $\mathbb{C}^n$, whose transition maps are biholomorphic.

Proposition 2.3.10. Every complex manifold can be realised as an almost complex manifold. The converse is not true in general.

To see this, suppose M is a complex manifold with local coordinates $z^j = x^j + iy^j$. Then one can define a complex structure by

$$J\frac{\partial}{\partial x^j} = \frac{\partial}{\partial y^j}, \quad J\frac{\partial}{\partial y^j} = -\frac{\partial}{\partial x^j}.$$

Conversely, there exist almost complex manifolds in real-dimension 2 which are not induced from holomorphic structures. We now investigate when an almost complex manifold does in fact arise from a holomorphic structure.

By the Frobenius integrability theorem, a subbundle of the tangent bundle of the manifold arises as the union of tangent spaces of a foliation if and only if the subbundle D satisfies $[X, Y] \in D$ for each $X, Y \in D$. This gives rise to the following definitions.

Definition 2.3.11. Let M be a manifold and D a subbundle of TM. Then D is called a *distribution*, and is said to be *integrable* if $[D, D] \subset D$.

Suppose (M, J) is an almost complex manifold. Consider the complexification $TM_{\mathbb{C}}$ of TM - note that this is $4n$-real-dimensional. Complexifying J gives a $\mathbb{C}$ linear map $J_{\mathbb{C}} : TM_{\mathbb{C}} \to TM_{\mathbb{C}}$.

Define $T_{1,0}M \coloneqq \{X - iJ_{\mathbb{C}}X : X \in TM\}, T_{0,1}M \coloneqq \{X + iJ_{\mathbb{C}}X : X \in TM\}$. Then for any $V \in T_{1,0}M$, we have

$$J_{\mathbb{C}}V = J_{\mathbb{C}}(X - iJX) = JX + iX = i(X - iJX).$$

Therefore $T_{1,0}M$ is the subbundle of $TM_{\mathbb{C}}$ consisting of eigenvectors of $J_{\mathbb{C}}$ which correspond to the eigenvalue i. Similarly, $T_{0,1}M$ is the subbundle of $TM_{\mathbb{C}}$ consisting of eigenvectors corresponding to eigenvalue $-i$. Moreover, it is easy to show that $TM_{\mathbb{C}} = T_{1,0}M \oplus T_{0,1}M$. These are called the *holomorphic* and *antiholomorphic* tangent bundles of (M, J). The celebrated Newlander–Nirenberg theorem is can be stated as:

Theorem 2.3.12. *Let (M, J) be an almost complex manifold. Then J is induced from a holomorphic structure if and only if the distribution $T^{0,1}M \subset TM_{\mathbb{C}}$ is integrable.*

This is highly non-trivial and not in the scope of this course. (In fact, the previous couple of pages in these notes are just me binge-typing notes about things that weren't actually covered in the lecture.)

Theorem 2.3.13. *An almost complex manifold is induced from a holomorphic structure if and only if the* Nijenhuis tensor $N^J(X, Y) = [X, Y] + J[JX, Y] + J[X, JY] - [JX, JY]$ *vanishes.*

Proof. We take for granted that N^J is truly a tensor, and prove that $T^{0,1}M$ is integrable if and only if N^J vanishes. Suppose $A + iJA, B + iJB$ are sections of $T^{0,1}M$. Then we have

$$[A + iJA, B + iJB] = [A, B] - [JA, JB] + i([JA, B] + [A, JB]).$$

Therefore $T^{0,1}M$ is integrable if and only if $J([A, B] - [JA, JB]) = [JA, B] + [A, JB]$, which is equivalent to the vanishing of the Nijenhuis tensor. $\square$

Remark. In symplectic geometry, almost complex manifolds (as opposed to complex manifolds) are sufficient 95% of the time.

2.4 Lecture 9

2.4.1 Stein Manifolds

The complex manifold $\mathbb{C}^n$ has a canonical symplectic structure, namely through the identification with $(\mathbb{R}^{2n}, \sum_i \mathrm{d}x^i \wedge \mathrm{d}y^i)$. In particular, complex submanifolds of $\mathbb{C}^n$ have induced symplectic structures by corollary 2.3.6.

Definition 2.4.1. A complex manifold X is *Stein* if there exists a proper holomorphic embedding of X into $\mathbb{C}^n$, for some n. Recall that a map is *proper* if the preimage of a compact map is compact.

Example. Subsets of $\mathbb{C}^n$ that arise as zero sets of polynomials (i.e. varieties) are Stein. However, not all Stein manifolds are affine.

Remark. In general there may be many proper embeddings of X into $\mathbb{C}^n$, each resulting in a different symplectic form. However, these are not so different as shown in the following result:

Theorem 2.4.2 (Gromov-Eliashberg)**.** *Any two symplectic structures on X arising from proper embeddings into $\mathbb{C}^n$ are symplectomorphic.*

What are some notable properties of Stein manifolds? Firstly they necessarily have infinite volume. A more interesting property is that they satisfy the *weak Lefschetz theorem*. That is, if X^{2n} is Stein, then $H_k(X, \mathbb{Z}) = 0$ for all $k > n$.

In fact, it was Shown by Eliashberg that Stein manifolds have a purely symplectic characterisation: all Stein manifolds can be obtained by Weinstein handle attachments. He called such manifolds *Weinstein manifolds* and constructed a functor

$$\text{Wein} : \text{Stein} \to \text{Weinstein}.$$

2.4.2 Kähler manifolds

We showed earlier that every symplectic manifold admits a compatible complex structure. A particularly well behaved case is when we can equip a symplectic manifold with an integrable compatible complex structure.

Definition 2.4.3. Let (M, ω, J) be a symplectic manifold equipped with a complex structure. M is a *Kähler manifold* if J is integrable.

Example. $\mathbb{C}^n$ is Kähler when equipped with the standard structures. Any almost complex submanifold of a Kähler manifold is Kähler. In particular, Stein manifolds are Kähler.

Proposition 2.4.4. Let (M, g, J) be a Riemannian manifold equipped with a compatible complex structure. Let $\omega(\cdot, \cdot) = g(J\cdot, \cdot)$. Then $\nabla J = 0$ if and only if (M, ω, J) is Kähler.

Proof. Suppose $\nabla J = 0$. Since ω is non-degenerate and skew-symmetric, it suffices to show that $\mathrm{d}\omega = 0$, and that the Nijenhuis tensor vanishes. First note that

$$\nabla(Jv) = (\nabla J)(v) + J\nabla v = J\nabla v.$$

In other words, J commutes with ∇. It follows that

$$\begin{aligned}
\mathrm{d}\omega(u, v) = \mathrm{d}g(Ju, v) &= g(\nabla(Ju), v) + g(Ju, \nabla v) \\
&= g(J\nabla u, v) + g(Ju, \nabla v) = \omega(\nabla u, v) + \omega(u, \nabla v).
\end{aligned}$$

Therefore, given a third vector field w, we have

$$w\omega(u, v) = \omega(\nabla_w u, v) + \omega(u, \nabla_w v).$$

But now by the formula for the exterior derivative on two forms,

$$\begin{aligned}
\mathrm{d}\omega(u, v, w) &= u\omega(v, w) + v\omega(w, v) + w\omega(u, v) - \omega([u, v], w) - \omega([v, w], u) - \omega([w, u], v) \\
&= u\omega(v, w) + v\omega(w, v) + w\omega(u, v) - \omega(\nabla_u v, w) + \omega(\nabla_v u, w) \\
&\qquad - \omega(\nabla_v w, u) + \omega(\nabla_w v, u) - \omega(\nabla_w u, v) + \omega(\nabla_u w, v) \\
&= 0.
\end{aligned}$$

Next we show that the Nijenhuis tensor N^J vanishes. Again by expressing each $[\varphi, \psi]$ as $\nabla_\varphi \psi - \nabla_\psi \varphi$, we immediately have

$$N^J = J(\nabla_X J)Y - J(\nabla_Y J)X + (\nabla_{JY} J)X - (\nabla_{JX} J)Y.$$

Therefore if $\nabla J = 0$, N^J vanishes.

Conversely, suppose (M, ω) is symplectic, and that J is integrable. To show that $\nabla J = 0$, one can work in local coordinates. It is intuitively clear since a holomorphic structure is locally unchanging. □

Another example of a Kähler manifold is $(\mathbb{CP}^n, \omega)$, where ω is the *Fubini Study metric*. The remainder of the lecture is dedicated to constructing the Fubini Study metric.

Definition 2.4.5. The *complex projective space* $\mathbb{CP}^n$ is the quotient of $\mathbb{S}^{2n+1} \subset \mathbb{C}^{n+1}$ obtained by identifying all points of $\mathbb{S}^{2n+1}$ which lie on the same complex line. Equivalently, $\mathbb{CP}^n$ is the quotient of $\mathbb{S}^{2n+1}$ by the diagonal group action of the multiplicative group $U(1)$. (Here $U(1) \cong \mathbb{S}^1$ is the *circle group*, $\{e^{i\theta} : \theta \in [0, 2\pi)\}$.)

If π denotes the quotient map, then $\pi : \mathbb{S}^{2n+1} \to \mathbb{CP}^n$ is a circle bundle over $\mathbb{CP}^n$. In fact, this is exactly the Hopf fibration! Also note that $\mathbb{S}^{2n+1}$ inherits a metric from $\mathbb{C}^{n+1}$ (namely the *round metric* g°). Let the distribution D on $\mathbb{S}^{2n+1}$ be defined by

$$D_x = \{v \in T_x\mathbb{S}^{2n+1} : v \text{ is orthogonal to the fibres of } \pi\}.$$

For each $x \in \mathbb{S}^{2n+1}$,

$$\mathrm{d}\pi_x|_{D_x} : D_x \to T_{[x]}\mathbb{CP}^n$$

is now a surjection. We can define a metric h on $T\mathbb{CP}^n$ by $h(u, v) = g^\circ(\widetilde{u}, \widetilde{v})$ for any lifts $\widetilde{u}$ and $\widetilde{v}$ by $\mathrm{d}\pi$. To see that this is well defined, we note that the circle metric is invariant under the $U(1)$ action. h is called the *Fubini-Study metric* on $\mathbb{CP}^n$. Moreover, let J denote the ambient complex structure on $\mathbb{C}^{n+1}$. Then each D_x is closed under J, and $J|_{D_x}$ is invariant under the $U(1)$ action. Therefore $J|_{D_x}$ descends to a complex structure on $\mathbb{CP}^n$. One can show that J is covariantly constant on $\mathbb{CP}^n$ to conclude that $\mathbb{CP}^n$ is a Kähler manifold with the Fubini-Study metric and induced complex structure.

Corollary 2.4.6. *Let A be a complex projective variety. Then A is a Kähler manifold. More generally, all complex submanifolds of $\mathbb{CP}^n$ are Kähler.*

2.5 Lecture 10

2.5.1 K3 surfaces

In the previous lecture we introduced Kähler manifolds, and remarked observed that all complex submanifolds of $\mathbb{C}^n$ and $\mathbb{CP}^n$ are Kähler. One might then ask - are there Kähler manifolds that are not projective? It turns out there are many such examples, and some arise as *K3 surfaces*.

Definition 2.5.1. A *K3 surface* X is a compact connected complex surface with trivial *canonical bundle* and *irregularity* zero. Here the canonical bundle refers to the 2nd exterior power (in general it is the nth exterior power of a complex n-manifold). The irregularity is the hodge number $h^{0,1}$, or equivalently, the dimension of the sheaf cohomology group $H^1(X, \mathcal{O}_X)$.

Unfortunately understanding this definition is beyond the scope of the course. However, I attempt to write some exposition to understand the hodge number. Some notable properties are as follows:

Proposition 2.5.2. Some properties of K3 surfaces:

- K3 surfaces exist.
- Every K3 surface is a Kähler manifold.
- Any two K3 surfaces are diffeomorphic.
- By Shing-Tung Yau's solution to the Calabi conjecture, every K3 surface is a Calabi-Yau manifold. (i.e. the Kähler metric is Ricci flat.)
- A K3 surface is projective if and only if it is algebraic. That is, it can be embedded in $\mathbb{CP}^n$ if and only if it arises as a variety.

Since every K3 surface is diffeomorphic, we might ask - how much choice do we have in the complex structure? It turns out the space of K3 surfaces is 20 dimensional. Moreover, the space of algebraic K3 surfaces can be shown to be 19 dimensional. It follows that there are K3 surfaces which cannot be holomorphically embedded in $\mathbb{CP}^n$.

Example. An example of a K3 surface is the *Fermat quartic*, defined to be the zero set of $x^4 + y^4 + z^4 + w^4$. This can be embedded in $\mathbb{CP}^3$.

2.5.2 Dolbeault cohomology

For a long time it was believed that every symplectic manifold was Kähler. In this section we prove that there exist non-Kähler symplectic manifolds, by constructing the *Kodaira-Thurston manifold*. However, to appreciate the construction, we must first develop some Dolbeault theory.

Recall from lecture 8 the holomorphic and anti-holomorphic tangent bundles, $T_{1,0}M$ and $T_{0,1}M$, of an almost complex manifold (M, J), defined by

$$T_{1,0}M = \{X - iJX : X \in TM\}$$
$$T_{0,1}M = \{X + iJX : X \in TM\}.$$

Moreover, we had $TM_{\mathbb{C}} = T_{1,0}M \oplus T_{0,1}M$. An easy way to see this is to consider the maps

$$\pi_{1,0} : TM \to T_{1,0}M, \quad \pi_{1,0} : v \mapsto \frac{1}{2}(v - iJv)$$
$$\pi_{0,1} : TM \to T_{0,1}M, \quad \pi_{0,1} : v \mapsto \frac{1}{2}(v + iJv).$$

Then $\pi_{1,0}$ and $\pi_{0,1}$ are (real) vector bundle isomorphisms, which satisfy $\pi_{1,0} \circ J = i\pi_{1,0}$, and $\pi_{0,1} \circ J = -i\pi_{0,1}$. This extends to a complex vector bundle isomorphism

$$(\pi_{1,0}, \pi_{0,1}) : TM_{\mathbb{C}} \to T_{1,0}M \oplus T_{0,1}M$$

by extending $\pi_{1,0}$ and $\pi_{0,1}$ to projections on $TM_{\mathbb{C}}$. Similarly, we decompose the cotangent bundle, and higher exterior powers: define $\Lambda^{1,0}M$ and $\Lambda^{0,1}M$ by

$$\Lambda^{1,0}M := \{\xi \in \Lambda^1 M_{\mathbb{C}} : \xi(v) = 0 \text{ for all } v \in T_{0,1}M\}$$
$$\Lambda^{0,1}M := \{\xi \in \Lambda^1 M_{\mathbb{C}} : \xi(v) = 0 \text{ for all } v \in T_{1,0}M\}.$$

It can be shown that $\Lambda^{1,0}M = \{\xi - i\xi \circ J : \xi \in \Lambda^1 M\}$, and $\Lambda^{0,1}M = \{\xi + i\xi \circ J : \xi \in \Lambda^1 M\}$. This gives rise to maps $\pi^{1,0}$ and $\pi^{0,1}$ as above, such that

$$(\pi^{1,0}, \pi^{0,1}) : \Lambda^1 M_{\mathbb{C}} \to \Lambda^{1,0}M \oplus \Lambda^{0,1}M$$

is an isomorphism.

We then define $\Lambda^{k,0}M$ and $\Lambda^{0,k}M$ to the kth exterior powers of $\Lambda^{1,0}M$ and $\Lambda^{0,1}M$ respectively. Finally we define $\Lambda^{k,l}M$ to be $\Lambda^{k,0}M \otimes \Lambda^{0,l}M$.

For vector spaces, we have $\Lambda^k(U \oplus V) = \bigoplus_{i=0}^k \Lambda^i U \otimes \Lambda^{k-i} V$. Therefore

$$\Lambda^k M_{\mathbb{C}} = \Lambda^k(\Lambda^{1,0}M \oplus \Lambda^{0,1}M) = \bigoplus_{i=0}^{k} \Lambda^{i,0}M \oplus \Lambda^{0,k-i}M = \bigoplus_{i=0}^{k} \Lambda^{i,k-i}M.$$

Proposition 2.5.3. Let (M, J) be an almost complex manifold. Then J is induced from a complex structure if and only if

$$\mathrm{d}(\Gamma(\Lambda^{p,q}M)) \subset \Gamma(\Lambda^{p+1,q}M \oplus \Lambda^{p,q+1}M).$$

Proof. We use the characterisation that J is induced from a complex structure if and only if $T_{0,1}M$ is integrable. We first prove that $T_{0,1}M$ is integrable if and only if $\mathrm{d}(\Gamma(\Lambda^{1,0}M)) \subset \Gamma(\Lambda^{2,0}M \oplus \Lambda^{1,1}M)$. In other words, we must show that $T_{0,1}M$ is integrable if and only if the $\Lambda^{0,2}M$ component of $\mathrm{d}\xi$ vanishes, for any section ξ on $\Lambda^{1,0}M$. By definition, the $\Lambda^{0,2}M$ component of $\mathrm{d}\xi$ vanishes if and only if $\mathrm{d}\xi(u, v) = 0$ for all $u, v \in \Gamma(T_{0,1}M)$. But since

$$\mathrm{d}\xi(u, v) = u\xi(v) - v\xi(u) - \xi([u, v]) = -\xi([u, v])$$

for u, v sections of $T_{0,1}M$, the $\Lambda^{0,2}M$ component of $\mathrm{d}\xi$ vanishes if and only if $T_{0,1}M$ is integrable, as required.

This now extends to the fact that $T_{0,1}M$ is integrable if and only if $\mathrm{d}(\Gamma(\Lambda^{p,q}M)) \subset \Gamma(\Lambda^{p+1,q}M \oplus \Lambda^{p,q+1}M)$, by applying the product rule. $\square$

Definition 2.5.4. On a complex manifold, the holomorphic and antiholomorphic differential operators $\partial : \Lambda^{p,q}M \to \Lambda^{p+1,q}M$ and $\overline{\partial} : \Lambda^{p,q}M \to \Lambda^{p,q+1}M$ are defined by projecting the image of the exterior derivative d.

By the above proposition, it is clear that $\mathrm{d} = \partial + \overline{\partial}$ on complex manifolds. This also gives that ∂^2, $\overline{\partial}^2$, and $\partial\overline{\partial} + \overline{\partial}\partial$ all vanish. Hence for every k, we have a chain complex

$$\Lambda^{k,0}M \xrightarrow{\overline{\partial}} \Lambda^{k,1}M \xrightarrow{\overline{\partial}} \Lambda^{k,2}M \xrightarrow{\overline{\partial}} \Lambda^{k,3}M \xrightarrow{\overline{\partial}} \cdots$$

which we call the *Dolbeault cocomplex*. Along with it, we have *Dolbeault cohomology groups*

$$H^{k,l}_{Dol}(M) = \frac{\ker \overline{\partial} : \Lambda^{k,l} \to \Lambda^{k,l+1}}{\operatorname{im} \overline{\partial} : \Lambda^{k,l-1} \to \Lambda^{k,l}}.$$

The Dolbeault cohomology is a useful tool in Kähler geometry.

Theorem 2.5.5. *On a compact Kähler manifold (M, ω, J), the Dolbeault cohomology groups satisfy*

$$H^m_{dR}(M; \mathbb{C}) \cong \bigoplus_{k+l=m} H^{k,l}_{Dol}(M).$$

In particular, $H^{k,l} \cong \overline{H^{l,k}}$, and each $H^{k,l}$ is finite dimensional. This is known as the Hodge decomposition.

Let b^m denote the Betti numbers of the de Rham cohomology groups. We write $h^{k,l}$ to denote the dimensions of the Dolbeault cohomology groups $H^{k,l}_{Dol}(M)$. These are called the *Hodge numbers*.

Remark. In the definition of a K3 surface, we required that $h^{0,1} = 0$. In other words, we require that every closed antiholomorphic 1-form arises as $\overline{\partial} f$ for some smooth function f.

Corollary 2.5.6. *All odd Betti numbers of a Kähler manifold are even.*

This is an immediate consequence of the Hodge decomposition, since

$$b^{2m+1} = \sum_{k+l=2m+1} h^{k,l} = 2\sum_{k=0}^{m} h^{k,2m+1-k}.$$

A useful diagram is the so-called *Hodge diamond*:

$$\begin{array}{ccccc} & & \vdots & & \\ h^{2,0} & & h^{1,1} & & h^{0,2} \\ & h^{1,0} & & h^{0,1} & \\ & & h^{0,0} & & \end{array}$$

Example. Earlier in the lecture we introduced *K3 surfaces*. One can show that all K3 surfaces have the following Hodge diamond:

$$\begin{array}{ccccc} & & 1 & & \\ & 0 & & 0 & \\ 1 & & 20 & & 1 \\ & 0 & & 0 & \\ & & 1 & & \end{array}$$

Example. The *Kodaira-Thurston manifold* is an example of a closed symplectic manifold which is not Kähler. It is constructed to be the quotient of $(\mathbb{R}^4, \omega)$ (where ω is the standard symplectic form) by the relations

$$\begin{aligned}(x_1, \ldots, x_4) &\sim (x_1 + 1, x_2, x_3 + x_4, x_4)\\ (x_1, \ldots, x_4) &\sim (x_1, x_2 + 1, x_3, x_4)\\ (x_1, \ldots, x_4) &\sim (x_1, x_2, x_3 + 1, x_4)\\ (x_1, \ldots, x_4) &\sim (x_1, x_2, x_3, x_4 + 1).\end{aligned}$$

Therefore it is something like a 4-torus with a Dehn twist. One can visualise the Kodaira-Thurston manifold as a torus-bundle over the torus, where the base space is restricted to the x_1 and x_2 coordinates. Then a loop in the x_1 direction corresponds to a Dehn twist upstairs, while the x_2 direction has no monodromy.

One can show that $b^1 = 3$ for the Kodaira-Thurston manifold. Since b^1 is odd, it cannot be a Kähler manifold.

2.5.3 Construction of symplectic manifolds

The fact that the above "counter example" was required to show that not all symplectic manifolds are Kähler highlights an underlying problem in symplectic geometry: the construction problem. There is currently no method of constructing a "generic symplectic manifold".

Two operations that exist are the symplectic blow up and symplectic sum. The symplectic blow up topologically corresponds to the usual blow up, in which, essentially, a sphere of a neighbourhood is replaced with a projective plane. One can show that a symplectic structure is induced.

The symplectic sum is a connected sum which inherits the symplectic structure. Suppose and M_1, M_2 are symplectic $2n + 2$ manifolds, and V a symplectic $2n$-manifold, embedded in M_1 and M_2 via ι_1, ι_2. Moreover, suppose the Euler classes of the normal bundles are opposite: $e(N_{M_1}V) = -e(N_{M_2}V)$. Then M_1 and M_2 are "topologically gluable", and it was proven by Gompf that there is a canonical isotopy class of symplectic structures on the connected sum $(M_1, V)\#(M_2, V)$. By using symplectic sums, Gompf proved the following theorem:

Theorem 2.5.7 (Gompf)**.** *Arbitrary finitely presented groups arise as fundamental groups of closed symplectic 4-manifolds.*

2.6 Lecture 11

2.6.1 Symplectic reduction

A couple of lectures ago, we constructed the Fubini-Study metric on $\mathbb{CP}^n$. This is in fact a special case of a more general procedure, which we call *symplectic reduction*.

Example (Symplectic reduction of symplectic vector spaces)**.** Symplectic reduction in the linear setting is easy. Let (V, Ω) be a symplectic vector space, and $W \subset V$ any subspace. Then $W/W \cap W^\Omega$ is naturally symplectic. To see this, note that the kernel of the quotient map $\pi : W \to W \cap W^\Omega$ is exactly $W \cap W^\Omega$, so one can define Ω_W by $\Omega_W(u, v) = \Omega(u', v')$ for any lifts u' and v' by π. This is non-degenerate since all non-degeneracy is precisely removed by the quotient.

Usually symplectic reduction is applied to coisotropic subspaces, in which the reduction is W/W^Ω (since $W \cap W^\Omega = W^\Omega$).

Example (Symplectic reduction of symplectic vector bundles)**.** This case is still easy, but requires an additional subtle assumption. If $X \to M$ is a symplectic vector bundle, and $Y \subset X$ a submanifold, it is not generally true that $Y \cap Y^\Omega$ has constant rank. Therefore it is not generally a vector bundle, so neither is the quotient $Y/Y \cap Y^\Omega$. However, by adding this as an additional premise, symplectic reduction extends to symplectic vector bundles.

The case of symplectic manifolds is the most involved, since specifying a subbundle of the tangent bundle (i.e. a distribution) does not immediately guarantee the existence of a submanifold whose tangent bundle is the prescribed subbundle.

Proposition 2.6.1. Suppose (M, ω) is a symplectic manifold, and $X \subset M$ a submanifold. Suppose $TX \cap TX^\omega$ has constant rank (and hence defines a distribution in TX). Then $TX \cap TX^\omega$ is integrable.

Proof. Let u, v be vector fields in the distribution $D = TX \cap TX^\omega$. Let $\omega_X = \omega|_X$. If $\omega_X([u, v], \cdot) = 0$, then by definition $[u, v] \in D$, which will prove that D is integrable. Hence it will be shown that $\omega_X([u, v], \cdot) = 0$.

Suppose η is a vector field on X. The explicit formula for the exterior derivative of a 2-form gives

$$\begin{aligned}\mathrm{d}\omega_X(u,v,\eta) = u\omega_X(v,\eta) + v\omega(\eta,u) + \eta\omega(u,v) \\ - \omega_X([u,v],\eta) - \omega_X([v,\eta],u) - \omega_X([\eta,u],v).\end{aligned}$$

The first three terms and last two terms immediately vanish since $u, v \in TX^\omega$, which leaves

$$0 = \mathrm{d}\omega_X(u,v,\eta) = \omega_X([u,v],\eta)$$

where the left side vanishes since symplectic forms are closed. Therefore D is integrable. $\square$

By the Frobenius integrability theorem, it follows that $TX \cap TX^\omega$ is the tangent space of some foliation $\mathcal{F}$ in X.

Example. If X is coisotropic, this process gives the *characteristic foliation.* One can also observe that the leaves are isotropic. In particular, if X has codimension 1, it is automatically coisotropic, and $TX \cap TX^\omega$ is one dimensional. This is automatically integrable.

The condition that $TX \cap TX^\omega$ has constant rank is not yet enough to define symplectic reduction for manifolds. This requires one more premise:

Proposition 2.6.2 (Symplectic reduction for symplectic manifolds)**.** Let (M,ω), X, and $\mathcal{F}$ be as above. Suppose there is a smooth manifold B (the *leaf space*) such that each leaf of $\mathcal{F}$ is given by the fibre of a surjective submersion $\pi : X \to B$. Suppose moreover that the fibres are connected. Then B has a canonical symplectic structure ω_r such that $\pi^*\omega_r = \omega_X$. (B,ω_r) is the *symplectic reduction* of X.

Proof. By the "linear case", for each $x \in X$, $T_xX/T_xX \cap T_xX^\omega$ has a linear symplectic structure. Let $x, x' \in X$ be lifts of $b \in B$. By construction the solid arrows in the diagram below are isomorphisms (by recalling that the tangent space at x of a leaf in the foliation is $T_xX \cap T_xX^\omega$). Therefore, if the dashed line preserves the symplectic structure, a canonical non-degenerate skew symmetric form is induced on T_bB.

$$\begin{tikzcd} T_xX/T_xX \cap T_xX^\omega \arrow[rr, dash, dotted] \arrow[dr] & & T_{x'}X/T_{x'}X \cap T_{x'}X^\omega \arrow[dl] \\ & T_bB & \end{tikzcd}$$

To prove that the dashed line preserves the symplectic structures on the quotients, we begin with a lemma: suppose u is a vector field on X, tangent to $TX \cap TX^\omega$. Then $\mathcal{L}_u\omega_X = 0$. (That is, the flow of u preserves ω_X.) This follows from the Cartan magic formula:

$$\mathcal{L}_u\omega_X = \iota_u\mathrm{d}\omega_X + \mathrm{d}(\iota_u\omega_X).$$

The first term on the right vanishes by closedness of ω, and the second term vanishes by the choice of u.

We now return to the original claim. Suppose x, x' are lifts of b which lie in the same submersion chart U. One can construct a vector field supported in U which flows x to x' in time 1. By the previous lemma, ω_X is preserved. More generally, given any lifts x and x', there is a path between them (by connectivity). Cover this path in submersion charts and choose a finite subcover. The result now follows inductively from the single-cover case.

This shows that the dashed line does in fact preserve ω_X, so a non-degenerate skew form ω_r is induced on B. Finally to see that it is closed (and hence a symplectic form), choose any section s of π. Then

$$\mathrm{d}\omega_r = \mathrm{d}(s^*\omega_X) = s^*(\mathrm{d}\omega_X) = 0.$$

$\square$

Example. Recall the construction of the Fubini-Study metric on $\mathbb{CP}^n$ from two lectures ago. Riemannian and holomorphic structures were used to determine that it was Kähler, but symplectic reduction provides an easy method for working in the symplectic setting.

Chapter 3

Physics

3.1 Lecture 12

3.1.1 Solutions to PDEs

A heuristic from linear algebra is as follows: given k equations and n unknowns, if $k < n$, a solution is guaranteed to exist, while if $k > n$, if the equations are genuinely distinct, we cannot find solutions. Applying this to PDEs, we have the following intuition: given n functions (in some number of variables) and k partial differential equations in these functions, if $k \leq n$, we can solve the system and vice versa.

Unfortunately the heuristic is wrong. consider the differential equation $\mathrm{d}\alpha = \beta$, where $\alpha \in \Omega^2(\mathbb{R}^4), \beta \in \Omega^3(\mathbb{R}^4)$. Expressing this in index notation, there are 6 functions and 4 equations. However, existence of solutions is classified into the following cases:

- If $\mathrm{d}\beta \neq 0$, there are no solutions.
- If $\mathrm{d}\beta = 0$, there are infinitely many solutions.

Another example is to consider a symplectic manifold (M, ω), and let $\varphi : M \to M$ be a smooth map. Then $\varphi^*\omega = \omega$ can be expressed as a system of differential equations, with $O(n)$ functions and $O(n^2)$ equations. The heuristic suggests there are no solutions, but usually there do exist solutions.

3.1.2 Hamiltonian vector fields

In the previous lecture, the construction of symplectic reduction for manifolds relied on a certain vector flow preserving the symplectic structure. Such vector fields are also of importance outside of this construction.

Definition 3.1.1. A vector field $v \in \Gamma(TM)$ is *symplectic* if $\mathcal{L}_v\omega = 0$.

Again by the Cartan magic formula, this is equivalent to the condition that $\iota_v\omega$ is a closed 1-form. In fact, by non-degeneracy of ω, every 1-form arises as $\iota_u\omega$ for some u. Hence the Cartan magic formula establishes the following one-to-one correspondence.

$$\{ \text{ Symplectic vector fields. } \} \longleftrightarrow \{ \text{ Closed 1-forms. } \}$$

Definition 3.1.2. A vector field v is *Hamiltonian* if $\iota_v\omega$ is exact. (Every Hamiltonian vector field is automatically Symplectic.)

Hamiltonian vector fields can be realised as energy preserving vector fields by the following construction:

Proposition 3.1.3. Let $H : M \to \mathbb{R}$ be a smooth map (the "Hamiltonian"). Then the vector field X_H defined by

$$\omega(X_H, \cdot) = \mathrm{d}H$$

is a Hamiltonian vector field.

Since ω is non-degenerate this truly defines a vector field. Moreover, X_f is tautologically Hamiltonian. By this construction, two functions give rise to the same Hamiltonian vector field if and only if they differ by a locally constant function.

Proposition 3.1.4. Let $H : M \to \mathbb{R}$. The flow of X_H preserves H. In particular, the level sets of H (i.e. energy levels) are preserved.

Proof.

$$X_H(H) = \mathrm{d}H(X_H) = \omega(X_H, X_H) = 0.$$

□

Suppose c is a regular value of H, i.e. a value such that $\mathrm{d}H_x$ is surjective (equivalently, non-zero) for each x in the preimage of c. By a dimensionality argument, $H^{-1}(c)$ has codimension 1, so it is a coisotropic submanifold, and $TH^{-1}(c) \cap TH^{-1}(c)^{\omega}$ is integrable. Therefore $H^{-1}(c)$ has a characteristic foliation as discussed in the previous lecture. This is called the *characteristic line field* and depends on the geometry of $H^{-1}(c)$ but not on H itself.

To see this, choose v to be orthogonal to $H^{-1}(c)$. Then

$$0 = \mathrm{d}H(v) = \omega(X_H, v),$$

so v is ω-orthogonal to X_H, and the integral curves of X_H follow the leaves of the characteristic foliation.

It is well known that symplectic geometry supposedly models Hamiltonian mechanics. Now that Hamiltonian vector fields have been introduced, it is only natural that Hamilton's equations are derived. Consider $(\mathbb{R}^{2n}, \mathrm{d}p^i \wedge \mathrm{d}q^i)$, $H : \mathbb{R}^{2n} \to \mathbb{R}$. Then

$$\iota_{X_H}\omega = \frac{\partial H}{\partial p^i}\mathrm{d}p^i + \frac{\partial H}{\partial q^i}\mathrm{d}q^i.$$

Let $X_H = f^i \frac{\partial}{\partial p^i} + g^i \frac{\partial}{\partial q^i}$ with a view to solving for f^i and g^i. This gives $f_i = \frac{\partial H}{\partial q^i}, g_i = -\frac{\partial H}{\partial p^i}$, so any integral curve $(p(t), q(t))$ of X_H satisfies

$$\dot{q}_i(t) = -\frac{\partial H}{\partial p^i}(p(t), q(t))$$
$$\dot{p}_i(t) = \frac{\partial H}{\partial q^i}(p(t), q(t)).$$

These are the famous Hamilton equations! (up to a difference in sign convention.)

In particular, suppose we choose $H = p^2/2m + V(q)$. Then the above equations become

$$\dot{q}_i(t) = -\frac{p_i}{m}$$
$$\dot{p}_i(t) = \frac{\partial V}{\partial q^i}(q(t)).$$

Thus $\frac{\partial V}{\partial q^i}(q(t)) = -m\ddot{q}_i(t)$. Or, more concisely, $F = ma$.

Proposition 3.1.5. Let $V \subset \mathbb{R}^{2n}$ be compact, and $H : \mathbb{R}^{2n} \to \mathbb{R}$ the Hamiltonian. Consider the time evolution of V by the flow ϕ_t of X_H. Then

$$\mathrm{vol}(\phi_t(V)) = \mathrm{vol}(V)$$

for all t.

Proof. Since X_H preserves ω, in particular it preserves ω^n. □

3.1.3 Poisson bracket

Definition 3.1.6. Let $f, g \in C^\infty(M)$. The *Poisson bracket* $\{f, g\}$ of f and g is the smooth function defined by $\{f, g\} = \omega(X_g, X_f)$.

Proposition 3.1.7. Let u, v be symplectic vector fields. Then $\omega([u, v], \cdot) = \mathrm{d}(\omega(v, u))$. I.e. $[u, v]$ is the Hamiltonian vector field of $\omega(v, u)$.

Proof. On one hand, the Leibniz rule gives

$$\mathcal{L}_u(\omega(v, \cdot)) = (\mathcal{L}_u\omega)(v, \cdot) + \omega(\mathcal{L}_u v, \cdot) = \omega([u, v], \cdot).$$

On the other hand, the Cartan magic formula gives

$$\mathcal{L}_u(\omega(v, \cdot)) = \iota_u(\mathrm{d}(\omega(v, \cdot))) + \mathrm{d}(\omega(v, u)) = \mathrm{d}(\omega(v, u)).$$

□

Proposition 3.1.8. Given any two smooth functions f, g, $X_{\{f,g\}} = [X_f, X_g]$.

Proof. Set $u = X_f, v = X_g$. Then by the previous proposition,

$$\omega([X_f, X_g], \cdot) = \mathrm{d}(\omega(X_g, X_f)) = \mathrm{d}\{f, g\}.$$

By definition of Hamiltonian vector fields, it follows that $X_{\{f,g\}} = [X_f, X_g]$. □

Proposition 3.1.9. The Poisson bracket satisfies the Jacobi identity.

Proof. By the previous proposition,

$$\{f, \{g, h\}\} = -\omega(X_f, X_{\{g,h\}}) = -\omega(X_f, [X_g, X_h]).$$

Another way to write this is

$$\{f, \{g, h\}\} = \omega(X_{\{g,h\}}, X_f) = \mathrm{d}\{g, h\}(X_f) = X_f\{g, h\} = X_f\omega(X_h, X_g)$$

Combing these two expressions,

$$2(\{f, \{g, h\}\} + \{g, \{h, f\}\} + \{h, \{f, g\}\}) = \mathrm{d}\omega(X_f, X_h, X_g) = 0.$$

□

In fact, bilinearity and skew symmetry are immediate, so the Poisson bracket is a Lie bracket.

3.2 Lecture 13

The fruits of the previous lecture can be summarised in a single exact sequence of Lie algebras:

$$\begin{tikzcd}
0 \arrow[r] & H^0(M) \arrow[r, hook, "\iota"] & C^\infty(M) \arrow[r] \arrow[dr, "\mathrm{d}"'] & \mathrm{SympVect}(M) \arrow[r] \arrow[d, dash, dotted] & \dfrac{\mathrm{SympVect}(M)}{\mathrm{HamVect}(M)} \arrow[r] \arrow[d, dash, dotted] & 0 \\
 & & & \{\text{closed 1-forms}\} \arrow[r, two heads, "\pi"] & H^1(M) \arrow[ur] &
\end{tikzcd}$$

Exactness is immediate: $H^0(M)$ corresponds to locally constant functions, so ι is simply an inclusion map, implying exactness at $H^0(M)$. Moreover, since $H^0(M)$ corresponds to locally constant functions, these are exactly the kernel of the exterior derivative, giving exactness at $C^\infty(M)$. Exactness at $\mathrm{SympVect}(M)$ is immediate since the image of d is the space of exact 1-forms, which is precisely quotiented out by π. Finally exactness at $H^1(M)$ is the most obvious of all, since quotient maps are surjective.

The less immediate statement is that this is an exact sequence of Lie algebras. $H^0(M)$ and $C^\infty(M)$ are equipped with the Poisson bracket, $\mathrm{SympVect}(M)$ with the Lie bracket of vector fields, and $H^1(M)$ the trivial Lie bracket. Since ι is an inclusion, it is immediately a Lie algebra homomorphism. The map $C^\infty(M) \to \mathrm{SympVect}(M)$ sends $f \mapsto X_f$. By a proposition at the end of the previous lecture, $\{f, g\} \mapsto X_{\{f,g\}} = [X_f, X_g]$, hence $C^\infty(M) \to \mathrm{SympVect}(M)$ is also a Lie algebra homomorphism. Finally the map $\mathrm{SympVect}(M) \to \mathrm{SympVect}(M)/\mathrm{HamVect}(M)$ is a Lie algebra homomorphism, since it was shown that $[u, v]$ is Hamiltonian whenever u and v are symplectic. This means $[u, v]$ is mapped to zero, which is exactly the result of the trivial bracket applied to the images of u and v.

3.2.1 Various isotopies

Given two maps, a *homotopy*, loosely speaking, is a continuous deformation of one map to the other. This continuous deformation does not need to preserve any structure. On the other hand, an *isotopy* between two maps must preserve additional structure assigned to the two maps. For example, a homotopy H of simple closed curves is just a continuous deformation from one curve to the other, so at an arbitrary time t, H_t is generally just a closed curve. However, if H is an isotopy of simple closed curves, then H_t must also be simple for all t.

Since it can be ambiguous to call all isotopies an isotopies irrespective of context, more precise terms are introduced during this lecture.

Definition 3.2.1. Let $\Phi : M \times [0, \varepsilon] \to M$ be smooth. Denote $\Phi|_{M\times\{t\}}$ by Φ_t. Φ is a *diffeotopy* if $\Phi_0 = \mathrm{id}$, and Φ_t is a diffeomorphism for all t.

There is a correspondence between diffeotopies and time dependent vector fields: suppose $\Phi : M \times I \to M$ is a diffeotopy, where I is a closed interval. This induces a diffeomorphism $\widetilde{\Phi} : M \times I \to M \times I$, defined by $\widetilde{\Phi}(x, t) = (\Phi_t(x), t)$. (Thus it is a diffeomorphism preserving constant time slices.) It follows that each $(y, t) \in M \times I$ belongs to a unique "curve" $\widetilde{\Phi}(\{x\} \times I)$. This can be formalised: for each $x \in M$, let $\gamma_x : I \to M \times I$ be defined by $\gamma_x(t) = \widetilde{\Phi}(\{x\} \times I)$. The tangent vectors of these curves are given by the pushforward of the unit vector field $\partial/\partial t$ in I. Hence the collection of all γ_x gives a non-vanishing vector field

$$X_\Phi : M \times I \to T(M \times I)$$

where each (x, t) is sent to the pushforward of the unit vector field $\partial/\partial t$ by γ_x. But then projecting back down by $M \times I \to I$ recovers the unit vector field, so there is a map $V : M \times I \to TM$ such that

$$X_\Phi(y, t) = (V(y, t), \partial/\partial t) \in T_y M \times \mathbb{R} = T_{y,t}(M \times I).$$

V is the *time dependent vector field* corresponding to Φ. More explicitly, one can write

$$V(y, t) = \frac{\partial \Phi}{\partial t}(\Phi_t^{-1}(y), t).$$

Conversely, given certain restrictions on time dependent vector fields, there exist diffeotopies giving rise to them.

Theorem 3.2.2. *Let V be a time dependent vector field $V : M \times I \to TM$, with bounded velocity. (i.e. M has a complete metric, and there is a constant K such that $|V|$ is bounded by K.) Then V generates a unique diffeotopy Φ of M satisfying*

$$\frac{\partial \Phi}{\partial t}(x,t) = V(\Phi(x,t),t).$$

Note that the "initial condition" of this ODE is embedded in the definition of a diffeotopy: $\Phi_0 = \mathrm{id}$.

Definition 3.2.3. Let $\Phi : M \times I \to M$ be a diffeotopy. Then

- Φ is a *symplectic isotopy* if Φ preserves the symplectic structure at each t. (I like to call these symplectopies.)
- If $f : M \to M$ is a diffeomorphism which arises as the time-1 map of a $[0,1]$-dependent Hamiltonian vector field, f is a *Hamiltonian diffeomorphism.*
- If Φ is Hamiltonian at each t, Φ is a *Hamiltonian isotopy.*

The above definitions were concise, but will be made more explicit for my own benefit: Suppose V is a time dependent vector field. If there exists a smooth family of smooth functions $H_t : M \to \mathbb{R}$ such that

$$\iota_{V_t}\omega = \mathrm{d}H_t,$$

then the diffeotopy corresponding to V is called a *Hamiltonian isotopy.* A symplectomorphism φ is called a *Hamiltonian diffeomorphism* if there exists a Hamiltonian isotopy Φ such that $\varphi = \Phi_1$.

Remark. Evidently the two definitions given above for Hamiltonian diffeomorphisms differ from one another. It is non-trivial to show that they agree, but they are in fact equivalent definitions. Similarly the two definitions of Hamiltonian isotopies are equivalent.

Definition 3.2.4. $\mathrm{Symp}(M)$ and $\mathrm{Ham}(M)$ denote the groups of symplectomorphisms and Hamiltonian diffeomorphisms respectively. $\mathrm{Symp}_0(M)$ denotes the identity component of $\mathrm{Symp}(M)$.

It is immediate that $\mathrm{Ham}(M) \subset \mathrm{Symp}(M)$. However, stronger results can be obtained: we have

$$\mathrm{Ham}(M) \lhd \mathrm{Symp}_0(M) \lhd \mathrm{Symp}(M)$$

as (infinite dimensional) Lie groups. On the Lie algebra level, this corresponds to

$$\mathrm{HamVect}(M) \cong \frac{C^\infty(M)}{\text{locally constant functions}} \hookrightarrow \mathrm{SympVect}(M).$$

Definition 3.2.5. Let Σ be two dimensional, and $u : \Sigma \to (M, \omega)$. The (symplectic) area of u is defined by

$$\mathrm{Area}(u) := \int_\Sigma u^*\omega.$$

Symplectic (and related) isotopies can be characterised by their effects on symplectic areas, as seen in the following chain of propositions.

Proposition 3.2.6. $\varphi : M \to M$ is a symplectomorphism if and only if φ preserves symplectic areas.

Proof. First suppose φ is a symplectomorphism. Then

$$\int_\Sigma (\varphi \circ u)^*\omega = \int_\Sigma u^*\varphi^*\omega = \int_\Sigma u^*\omega.$$

Conversely, if φ is not a symplectomorphism, choose some x and vectors v, v' such that $\omega_x(v, v') > \omega_{\varphi(x)}(\mathrm{d}\varphi_x v, \mathrm{d}\varphi_x v')$. Then for a sufficiently small disk D with an embedding $u : D \to M$, $u(0) = x$, we have

$$\int_D u^*\omega > \int_D (\varphi \circ u)^*\omega.$$

$\square$

Proposition 3.2.7. Let $\Phi : M \times I \to M$ be a diffeotopy. This is a symplectopy if and only if symplectic areas are preserved at each t, if and only if symplectic areas traced by contractible loops are zero for all t.

Proof. First some terminology is clarified. The "symplectic area traced by a contractible loop at t" refers to the symplectic area of $\mathbb{S}^1 \times [0, t] \hookrightarrow M \times I \to M$.

The equivalence of the first two statements is clear. it remains to prove that symplectic areas are preserved at each t if and only if symplectic areas traced by contractible loops are zero for all t. Choose $\gamma : \mathbb{S}^1 \to M$ contractible, i.e. such that there exists $u : D \to M$ with $\partial u = \gamma$. Define $f : D \times [0, \varepsilon] \to M$ such that $(z, t) \mapsto \Phi_t(u(z))$. Then by Stokes' theorem,

$$\int_{\partial(D\times[0,\varepsilon])} (\partial f)^*\omega = \int_{D\times[0,\varepsilon]} f^*\mathrm{d}\omega = 0.$$

But the left side can be decomposed into three pieces:

$$\int_{\partial(D\times[0,\varepsilon])} (\partial f)^*\omega = \int_{\{\varepsilon\}\times D} f^*\omega - \int_{\{0\}\times D} f^*\omega + \int_{[0,\varepsilon]\times\partial D} f^*\omega.$$

The first two terms cancel for any ε and contractible loop if and only if Φ preserves all areas. But the first two terms cancel if and only if the last term vanishes, as required. $\square$

Proposition 3.2.8. Let $\Phi : M \times I \to M$ be a diffeotopy. This is a Hamiltonian isotopy if and only if symplectic areas traced by loops are zero for all t. (Not necessarily contractible loops.)

3.3 Lecture 14

3.3.1 Symplectic isotopies

In this lecture we continue from the previous lecture. We proved equivalence between properties (1), (3), and (4).

Proposition 3.3.1. Let $\Phi : M \times [0, \varepsilon] \to M$ be a diffeotopy. The following are equivalent:

1. Φ is a symplectopy.
2. Φ corresponds to a $[0, \varepsilon]$-dependent vector field which is symplectic for every $t \in [0, \varepsilon]$.
3. Φ preserves symplectic areas. More precisely, if Σ is an oriented compact manifold with boundary, and $u : \Sigma \to M$ is smooth, then $\mathrm{Area}(u) = \mathrm{Area}(\Phi_t \circ u)$ for any t.
4. Symplectic areas traced by contractible loops vanish. More precisely, if $\gamma : \mathbb{S}^1 \to M$ is a loop such that there exists $u : D \to M$ with $\partial u = \gamma$, then $\mathbb{S}^1 \times [0, t] \to M$ defined by $(\theta, k) \mapsto \Phi_k(\gamma(\theta))$ has zero area.

It remains to prove the equivalence between (1) and (2). Let Φ be a diffeotopy, and V the corresponding time dependent vector field. Then $\frac{\partial}{\partial t}\Phi(x, t) = \Phi_t^* V_t$. Observe that Φ is symplectic if and only if

$$\frac{\partial}{\partial t}\Phi_t^*\omega = 0.$$

But one can show that

$$\frac{\partial}{\partial t}\Phi_t^*\omega = \Phi_t^*(\mathcal{L}_{V_t}\omega),$$

which shows that Φ is symplectic if and only if V_t is symplectic for each t. Due to technical difficulties, this lecture didn't cover any more content.

3.4 Lecture 15

3.4.1 Hamiltonian isotopies

Let $H, G : M \times [0, \varepsilon] \to \mathbb{R}$ be smooth. They give rise to time dependent Hamiltonian vector fields, which in turn give rise to Hamiltonian flows $\varphi_H^t, \varphi_G^t : M \times [0, \varepsilon] \to M$.

Lemma 3.4.1. At any t, $\varphi_H^t \circ \varphi_G^t : M \times \{t\} \to M$ is the Hamiltonian flow of the smooth map $H_t + G_t \circ (\varphi_H^t)^{-1} : M \times \{t\} \to \mathbb{R}$.

Proof. Let X_H, X_G be the corresponding time dependent vector fields. By the chain rule,

$$\frac{\mathrm{d}}{\mathrm{d}t}(\varphi_H^t \circ \varphi_G^t)(x)\Big|_{t=t_0} = X_{H t_0}(\varphi_G^{t_0}(x)) + (\varphi_H^{t_0})_* X_{G t_0}(\varphi_G^{t_0}(x)).$$

This shows that $X_{Ht} + (\varphi_H^t)_* X_{Gt}$ generates the flow $\varphi_H^t \circ \varphi_G^t$. But for any t, we see that $(\varphi_H^t)_* X_{Gt} = X_{G_t \circ (\varphi_H^t)^{-1}}$ from the following calculation: (The ts have been dropped for clarity.)

$$\begin{aligned}
\omega((\varphi_H)_* X_G, V) &= (\varphi_H)^* \omega(X_G, (\varphi_H^{-1})_* V) \\
&= \omega(X_G, (\varphi_H^{-1})_* V) \\
&= \mathrm{d}G((\varphi_H^{-1})_* V) \\
&= (\varphi_H^{-1})^* \mathrm{d}G(V) = \mathrm{d}(G \circ \varphi_H^{-1})(V).
\end{aligned}$$

Therefore $X_{Ht} + (\varphi_H^t)_* X_{Gt} = X_{Ht} + X_{G_t \circ (\varphi_H^t)^{-1}}$, so the corresponding Hamiltonian is $H_t + G_t \circ (\varphi_H^t)^{-1}$. $\square$

Recall the series of equivalent conditions for an isotopy to be symplectic. Similar results exists for Hamiltonian isotopies:

Proposition 3.4.2. Let $\Phi : M \times [0, \varepsilon] \to M$ be a diffeotopy. The following are equivalent:

1. Φ is a Hamiltonian isotopy.
2. The corresponding time dependent vector field is Hamiltonian.
3. For every loop $\gamma : \mathbb{S}^1 \to M$, the area traced by $u : \mathbb{S}^1 \times [0, t] \to M$ is zero at every time.

The third condition holds for all loops, rather than just contractible loops, distinguishing it from symplectopies. The following diagram illustrates the difference.

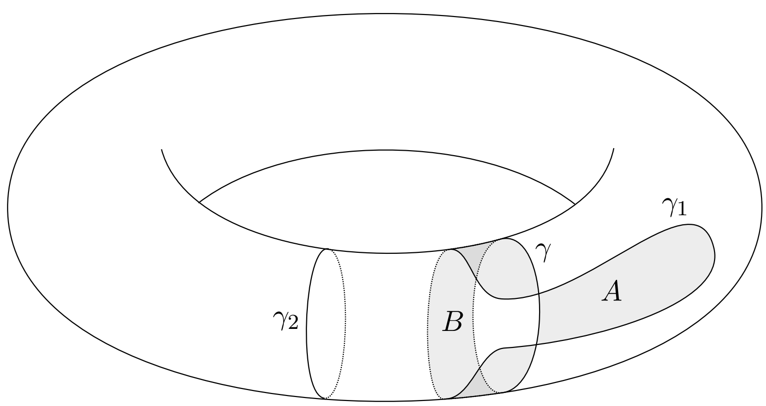

- The isotopy taking γ to γ_2 is tracing out a non-zero area. However, γ is not contractible. Hence the isotopy might be a symplectopy, although it cannot be a Hamiltonian isotopy.

- The isotopy taking γ to γ_2 might be a Hamiltonian isotopy if $A = B$. Otherwise it cannot be Hamiltonian.

The equivalence between (1) and (2) was not proven in the lecture, although it was remarked that $2 \Rightarrow 1$ is immediate while the converse is non-trivial. Condition (3) is now analysed.

Recall the gradient lemma:

Lemma 3.4.3. A 1-form is exact if and only if the line integral of the form depends only on the endpoints of the curve, or equivalently, if the integral around any smooth closed curve is zero.

Let V_t be the time dependent vector field corresponding to the diffeotopy $\Phi : M \times [0, \varepsilon] \to M$. Let $u : \mathbb{S}^1 \times [0, \varepsilon] \to M$ be defined by $(\theta, t) \mapsto \Phi_t(\gamma(\theta))$. Then $\frac{\mathrm{d}u}{\mathrm{d}t}(\theta, t) = V_t(\Phi_t(\gamma(\theta)))$. Define α to be the 1-form $\alpha_t(\cdot) = \omega(V_t, \cdot)$. Then

$$\int_{\mathbb{S}^1 \times [0,t]} u^*\omega = \int_0^t \int_{\mathbb{S}^1} \omega\Big(\frac{\mathrm{d}u}{\mathrm{d}t}, \frac{\mathrm{d}u}{\mathrm{d}\theta}\Big)\mathrm{d}\theta\mathrm{d}t = \int_0^t \int_{\mathbb{S}^1} \alpha_t\Big(\frac{\mathrm{d}u}{\mathrm{d}\theta}\Big)\mathrm{d}\theta\mathrm{d}t = \int_0^t \Big(\int_{\mathbb{S}^1} (u|_{\{t\}\times\mathbb{S}^1})^*\alpha_t\Big)\mathrm{d}t.$$

By definition, V_t is Hamiltonian if and only if α_t is exact. But by the previous lemma, α_t is exact if and only if the term on the right vanishes for all u and t. Hence V_t is Hamiltonian for each t if and only if the area traced by u vanishes, as required.

Remark. Let $\Phi : [0, \varepsilon] \times M \to M$ be a symplectopy. This determines a canonical element of $H^1(M, \mathbb{R})$ which represents the "area". If the flux is zero at all times, Φ is Hamiltonian.

3.4.2 Infinitesimal group actions

Definition 3.4.4. Let $\mathfrak{g}$ be a Lie algebra. An *infinitesimal action* on M is a Lie algebra homomorphism $\mathfrak{g} \to \Gamma(TM)$. More specifically, an *infinitesimal symplectic action* is a Lie algebra homomorphism $\mathfrak{g} \to \mathrm{SympVect}(M)$. Finally, a *Hamiltonian action* is a Lie algebra homomorphism $\mathfrak{g} \to C^\infty(M)$, where the latter is equipped with the Poisson bracket.

Assume the image of $\mathfrak{g} \to \mathrm{SympVect}(M)$ lies in $\mathrm{HamVect}(M)$. Does this imply that the action arises from a Hamiltonian action? Not necessarily! There is an obstruction the existence of the lift below:

$$\begin{tikzcd} & \mathfrak{g} \arrow[d] \arrow[dl, dotted, "\exists ?"'] \\ C^\infty(M) \arrow[r, "f \mapsto X_f"'] & \mathrm{HamVect}(M) \end{tikzcd}$$

Precisely, the obstruction is the second Lie algebra cohomology with coefficients in $H^0(M, \mathbb{R})$. More concretely, choose a basis $g_1, \ldots, g_n$ of $\mathfrak{g}$, and lift their images in $\mathrm{HamVect}(M)$ to $C^\infty(M)$. It is not immediate that $[g_i, g_j]$ agrees with $\{f_i, f_j\}$ where $g_i \mapsto f_i$.

3.5 Lecture 16

3.5.1 Lie theory

Definition 3.5.1. A *Lie group* is a smooth manifold G equipped with a compatible group structure. That is, multiplication and inversion must be smooth.

Proposition 3.5.2. Each Lie group G has a corresponding Lie algebra

$$T_{\mathrm{id}}G \cong \{v \in \Gamma(TG) : (m_g)_* v = v\} = \{\text{left invariant vector fields}\}.$$

Here m_g is the map $m_g : h \mapsto gh$. Given any $v \in T_{\mathrm{id}}G$, the corresponding left-invariant vector field is defined by $\xi_v(g) = (m_g)_* v$. The Lie bracket on left invariant vector fields is exactly the Lie bracket of vector fields.

Proposition 3.5.3. There is a one-to-one correspondence

$$\{\text{connected simply connected Lie groups}\} \longleftrightarrow \{\text{finite dimensional Lie algebras}\}.$$

Lemma 3.5.4. Let $f : G \to H$ be a Lie group homomorphism. Then $\mathrm{d}f$ is a Lie algebra homomorphism.

Proof. Let $v \in T_{\mathrm{id}}G$, $f_* v \in T_{\mathrm{id}}H$. These correspond to left invariant vector fields ξ_v on G and $\xi_{f_* v}$ on H. These vector fields are f-related:

$$(\xi_{f_* v} \circ f)(g) = (m_{f(g)})_* f_* v = f_* (m_g)_* v = (f_* \xi_v)(g).$$

Since f-relatedness is preserved by Lie brackets, $\mathrm{d}f$ preserves Lie brackets. Thus $\mathrm{d}f$ is a Lie algebra homomorphism. $\square$

3.5.2 Adjoint and coadjoint representations

A Lie group action is a smooth group action $G \times M \to M$. There is a correspondence between Lie group actions and infinitesimal group actions:

$$\{G \text{ action on } M\} \longleftrightarrow \{\text{infinitesimal action } \mathfrak{g} \to \Gamma(TM)\}.$$

There is always a map from the left to right, namely

$$a : G \times M \to M \longmapsto \frac{\mathrm{d}}{\mathrm{d}t} a(\exp(tv), x)\Big|_{t=0}.$$

If G is connected and simply connected, and there is a basis of $\mathfrak{g}$ consisting of complete vector fields, given any infinitesimal action, there is a Lie group action inducing it.

Definition 3.5.5. Let Φ_g denote the inner automorphism of g for each $g \in G$. By an earlier lemma, $\mathrm{d}\Phi_g$ is a Lie algebra homomorphism. In fact, it is an automorphism. The *adjoint representation* is

$$\mathrm{Ad} : G \to \mathrm{Aut}(\mathfrak{g})$$

defined by $g \mapsto \mathrm{d}\Phi_g$.

The adjoint representation induces infinitesimal actions: consider $T_{\mathfrak{g}}$ to be a manifold. Then there is a canonical action $\mathrm{ad} : \mathfrak{g} \to \Gamma(T_{\mathfrak{g}})$ defined by $v \mapsto [v, \cdot]$. This is induced from Ad in the sense that $\mathrm{ad} = \mathrm{d}\,\mathrm{Ad}$, where $\Gamma(T_{\mathfrak{g}})$ is identified with the Lie algebra corresponding to the Lie group $\mathrm{Aut}(\mathfrak{g})$. Thus given elements $x, y \in \mathfrak{g}$,

$$\mathrm{ad}_x\, y = [x, y].$$

There is a very evident parallel with the Lie derivative, in which

$$\mathcal{L}_X Y = [X, Y]$$

for all vector fields X and Y on a manifold.

Definition 3.5.6. The *coadjoint representation* is $K : G \to \mathrm{Aut}(\mathfrak{g}^*)$ defined by

$$g \mapsto \mathrm{Ad}(g^{-1})^*.$$

The coadjoint representation is of interest since its orbits admit symplectic structures as we now explore.

Definition 3.5.7. Let M be a closed manifold. A Lie algebra structure on $C^\infty(M)$ is a *Poisson structure* if it satisfies the following Leibniz rule:

$$[f, gh] = g[f, h] + h[f, g].$$

For example, $C^\infty(M)$ equipped with the Poisson bracket is a Poisson structure.

Proposition 3.5.8. A Poisson structure is equivalent to a bivector field $\pi \in \Gamma(\Lambda^2 M)$ satisfying $[\pi, \pi] = 0$ (where the brackets denote the Schouten bracket, a generalisation of the Lie bracket).

Let π be an alternating bi-vector field. This corresponds to a map

$$\widetilde{\pi} : T^*M \to TM,$$

and its image defines a "distribution" corresponding to a singular foliation by the Frobenius integrability theorem. The leaves of this foliation are symplectic.

Example. Consider the manifold $\mathfrak{g}^*$. Each $g \in G$ defines an orbit in $\mathfrak{g}^*$ via the coadjoint representation. These leaves are naturally symplectic.

More explicitly, at each $a \in \mathfrak{g}^*$, $T_a\mathfrak{g}^*$ is isomorphic to $\mathfrak{g}$. Thus one can define $\pi \in \Gamma(\Lambda^2\mathfrak{g})$ by $\pi_a(g, g') = a([g, g'])$ for $g, g' \in \mathfrak{g}$. This defines a symplectic structure on each coadjoint orbit.

3.6 Lecture 17

3.6.1 Coadjoint action examples

In the previous lecture, we observed that orbits of the coadjoint action of G on $\mathfrak{g}^*$ have natural symplectic structures. Are the orbits manifolds?

Proposition 3.6.1. Suppose G acts on a smooth manifold M. Its orbits are images of injective immersions.

Proof. Let $p \in M$. Then $\operatorname{Stab}(p) := \{g \in G : gp = p\} \subset G$ is a closed subset. The map

$$\iota : G/\operatorname{Stab}(p) \hookrightarrow M, \quad [g] \mapsto gp$$

is an injective smooth map whose image is the orbit of p. By the quotient manifold theorem, the domain is a manifold. ι has constant rank so it is an immersion. If G is compact, the orbits are submanifolds. $\square$

Theorem 3.6.2 (Kirilov-Kostant-Souriau)**.** *Every homogeneous symplectic manifold of a lie group G is, up to a possible covering, a coadjoint orbit of a central extension of G. (Homogeneous symplectic means it admits a transitive G action by symplectomorphisms.)*

Example. Let $G = \mathrm{SU}(2)$. Then $\mathfrak{g} = \mathfrak{su}(2)$. Identify $\mathfrak{g}$ with

$$\mathfrak{g}^* = \left\{ \begin{pmatrix} ai & z \\ -\overline{z} & -ai \end{pmatrix} : a \in \mathbb{R}, z \in \mathbb{C} \right\}.$$

Consider the G action on $\mathfrak{g}$ by conjugation. This corresponds to a rotation in $\mathbb{R}^3$, so the orbits are spheres.

Example. Let $G = \mathrm{SL}(2, \mathbb{R})$, so that $\mathfrak{g} = \mathfrak{sl}(2, \mathbb{R})$. $\mathfrak{sl}(2, \mathbb{R})$ can be identified with real matrices with trace zero. There are six types of conjugacy classes (i.e. orbits of the coadjoint action), as follows:

$$\begin{pmatrix} 0 & 0 \\ 0 & 0 \end{pmatrix}, \begin{pmatrix} 0 & 0 \\ 1 & 0 \end{pmatrix}, \begin{pmatrix} 0 & 1 \\ 0 & 0 \end{pmatrix}, \begin{pmatrix} \lambda & 0 \\ 0 & -\lambda \end{pmatrix}, \begin{pmatrix} 0 & -\lambda \\ \lambda & 0 \end{pmatrix}, \begin{pmatrix} 0 & \lambda \\ -\lambda & 0 \end{pmatrix}$$

More precisely, the above is a complete list of representatives of distinct conjugacy classes, where $\lambda > 0$ is arbitrary. Identifying $(2, \mathbb{R})$ with $\mathbb{R}^3$, the orbits are the submanifolds in the following figure:

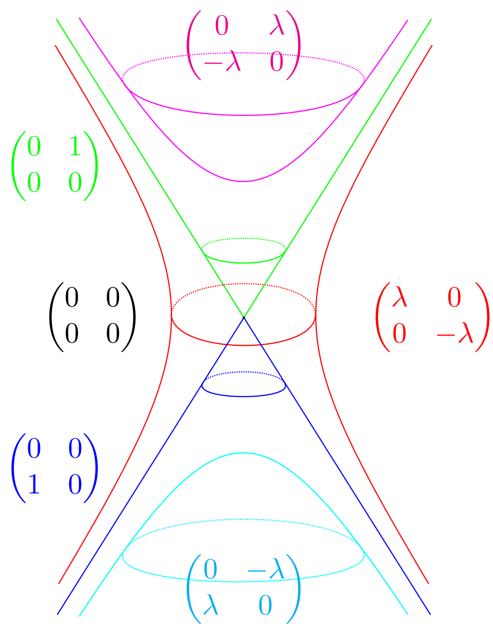

3.6.2 Moment and comoment maps

Recall that Hamiltonian actions, introduced at the end of lecture 15, were Lie algebra homomorphisms $\mathfrak{g} \to C^\infty(M)$, where the latter is equipped with the Poisson bracket.

Definition 3.6.3. A symplectic Lie group action $G \times M \to M$ is a *Hamiltonian Lie group action* if the corresponding infinitesimal action $\mathfrak{g} \to \mathrm{SympVect}(M)$ lifts to a Hamiltonian action $\mathfrak{g} \to C^\infty(M)$. The lift $\mathfrak{g} \to C^\infty(M)$ is called the *comoment map*.

Proposition 3.6.4. The comoment map α is G-equivariant.

Proof. First this statement is unpacked. Recall that G acts on $\mathfrak{g}$ by the adjoint action. G also has a natural action on $C^\infty(M)$ by precomposition by the left multiplication map. We must show that for each $g \in G$ and $v \in \mathfrak{g}$,

$$\alpha(\mathrm{Ad}(g)v) = \alpha(v) \circ m_g = m_g^* \alpha(v) = \alpha(m_{g*} v).$$

Observe that if $\alpha(v)$ generates the flow φ^t, then $\alpha(\mathrm{Ad}(g)v)$ generates the the flow $m_g \circ \varphi^t \circ m_g^{-1}$. On the other hand, in lecture 15 it was shown that the Hamiltonian vector field of $H \circ \varphi^{-1}$ is $\varphi_* X_H$. Therefore $\alpha(v) \circ m_g$ has the Hamiltonian vector field $m_{g*} X_{\alpha(v)}$. The result follows. $\square$

There is also a *moment map*.

Definition 3.6.5. The map $\mu : M \to \mathfrak{g}^*$ defined by

$$\mu : x \mapsto (g \mapsto \alpha(g)(x))$$

is the *moment map.* This is G-equivariant by the equivariance of the comoment map.

Remark. Hamiltonian actions can be defined via the moment map.

Example. Suppose $\mathcal{O} \subset \mathfrak{g}^*$ is a coadjoint orbit. Then it is a manifold, and the moment map $\mathcal{O} \to \mathfrak{g}^*$ is simply the inclusion map.

3.7 Lecture 18

3.7.1 Moment map examples

Example. Consider the sphere $\mathbb{S}^2 \subset \mathbb{R}^3$, with the standard area form as its symplectic structure. Consider rotation along the z-axis. This is Hamiltonian since all areas are preserved (so it is symplectic), and all loops are contractible. The standard area form, in coordinates, is given by

$$\omega = \mathrm{d}z\mathrm{d}\theta.$$

Hence the generator of the rotation ($\mathbb{S}^1$ action) is $\partial/\partial\theta$. The corresponding Lie algebra is $\mathbb{R}$. The moment map $\mu : \mathbb{S}^2 \to \mathbb{R}$ is exactly projection onto the z coordinate.

This example gives rise to Archimedes' theorem: the horizontal projection of a sphere onto a cylinder is area-preserving. Of course, this is obvious if one knows that an area form is given by $\mathrm{d}z\mathrm{d}\theta$. This is sometimes used as a map projection.

Example. SO(3) acts on $\mathbb{S}^2$ with a Hamiltonian action. This is in fact a coadjoint orbit of SO(3), so by the general theory, the moment map is the inclusion $\mu : \mathbb{S}^2 \hookrightarrow \mathfrak{g}^* \cong \mathbb{R}^3$.

Remark. Moment and comoment maps are not unique. For example, in the z-axis rotation of $\mathbb{S}^2$, the map $\mathbb{R} \to C^\infty(M)$ given by $1 \mapsto z + 1$ is a comoment map. This gives a corresponding distinct moment map.

Remark. $0 \in \mathfrak{g}^*$ is always a coadjoint orbit.

Proposition 3.7.1. The moment map is unique up to a constant of integration. More precisely, if μ_1, μ_2 are moment maps of a G-action, then $\mu_1 - \mu_2 \in [\mathfrak{g}, \mathfrak{g}]^0 \cong H^1(\mathfrak{g}, \mathbb{R})$. Therefore if $\mathfrak{g}$ is semisimple, the moment map is unique. On the other and, if $\mathfrak{g}$ is abelian, for any $c \in \mathfrak{g}^*$, $\mu + c$ is a moment map.

Example. $\mathbb{R}^3$ acts on $\mathbb{R}^3$ by translation. This incudes a Hamiltonian group action on $T^*\mathbb{R}^3$. The corresponding moment map is $T^*\mathbb{R}^3 \to \mathbb{R}^3$, which is given by projection onto the p-coordinate. (Exercise: flesh out details.)

Example. SO(3) acts on $\mathbb{R}^3$ by rotations. This incudes a Hamiltonian group action on $T^*\mathbb{R}^3$. The corresponding moment map is $T^*\mathbb{R}^3 \to \mathbb{R}^3$, which is given by $(p, q) \mapsto p \times q$. The former corresponds to linear momentum, while this example is angular momentum. (Exercise: flesh out details.) See recent book by Peter Woit.

3.7.2 Symplectic reduction revisited

Proposition 3.7.2. Let $G \times M \to M$ be a Hamiltonian action with moment map $\mu : M \to \mathfrak{g}^*$. Consider

$$\mathrm{d}\mu_p : T_pM \to T_{\mu(p)}\mathfrak{g}^* \cong \mathfrak{g}^*.$$

Then

$$\begin{aligned} \ker \mathrm{d}\mu_p &= (T_p \operatorname{Orb}(p))^\omega \\ \operatorname{im} \mathrm{d}\mu_p &= \operatorname{Ann}(\operatorname{LieAlg}(\operatorname{Stab}(p))) \subset \mathfrak{g}^*. \end{aligned}$$

Proof. Let $\alpha : \mathfrak{g} \to \Gamma(TM)$ be the corresponding infinitesimal action. Let $p \in M$. Then

$$\omega_p(\alpha(X)_p, v) = \mathrm{d}\mu_p(v)(X)$$

for any $X \in \mathfrak{g}, v \in T_pM$ by definition. Thus $w \in \ker \mathrm{d}\mu_p$ if and only if it lies in the symplectic orthogonal of all $\alpha(X)_p$. This is exactly the symplectic orthogonal of $T_p \operatorname{Orb}(p)$. The image is similar. $\square$

This looks rather simple but it has an important consequence. A number of lectures ago, symplectic reduction was introduced in a general setting. The usual setting of symplectic reduction is as follows:

Theorem 3.7.3 (Marsden-Weinstein-Meyer)**.** *Let $G \curvearrowright M$ be a Hamiltonian group action, and $\mu : M \to \mathfrak{g}^*$ a corresponding moment map. Assume G is compact and acts freely on $\mu^{-1}(0)$. Then $\mathrm{d}\mu_p$ is surjective for all $p \in \mu^{-1}(0)$, so 0 is a regular value, thus $\mu^{-1}(0)$ is a submanifold. In fact, $\mu^{-1}(0)$ is coisotropic. It has a characteristic foliation given by G-orbits. Therefore $\mu^{-1}(0)/G$ admits a symplectic structure. This is the* symplectic reduction *of M by G, denoted $M//G$.*

Chapter 4

Local and global invariants

4.1 Lecture 19

4.1.1 Darboux's theorem

Theorem 4.1.1. *Let (M, ω) be a connected symplectic manifold. For any $p, q \in M$, there is a Hamiltonian isotopy $\Phi : M \times I \to M$ such that $\Phi_1(p) = q$.*

Intuitively this theorem says that any two points in a symplectic manifold locally look the same. This theorem is used to prove Darboux's theorem.

Proof. Choose a smoothly embedded path $\gamma : [0, 1] \to M$ such that $\gamma(0) = p, \gamma(1) = q$, and let $v_t = \gamma'(t) \in T_{\gamma(t)}M$. Next we wish to find a $[0, 1]$ dependent Hamiltonian vector field which is supported near γ, which restricts to v_t. How do we do this?

Consider a neighbourhood N of $\gamma(t)$. Choose a diffeotopy $\Psi : M \times I \to M$ such that Ψ is the identity map outside of N, and $\Psi_t(\gamma(0)) = \gamma(t)$. Fix an embedding $E : B_\varepsilon(0) \hookrightarrow M$ which sends 0 to $\gamma(0)$, and consider $\Psi_t \circ E$. At each time t, solve the following in $B_\varepsilon(0)$ for a linear Hamiltonian:

$$X_{H_t}(\gamma(t)) = v_t.$$

Cut them off appropriately (i.e. multiply by smooth cutoff functions). This gives a time dependent Hamiltonian vector field. This corresponds to the desired Hamiltonian isotopy. $\square$

Theorem 4.1.2 (Darboux's theorem · Normal form theorem)**.** *Let (M, ω) be symplectic, and $x \in M$. Then there exists a coordinate system $p_1, \dots, p_n, q_1, \dots, q_n$ such that $\omega = \mathrm{d}p^i \wedge \mathrm{d}q^i$ in the domain of the coordinates.*

Proof. Let $y \neq x$. Suppose there exists a symplectic structure ω' on M such that $\omega' = \omega$ near x, and $\omega' = \mathrm{d}p \wedge \mathrm{d}q$ near y in some coordinate system. Then by the previous theorem, ω' is of the form $\mathrm{d}p \wedge \mathrm{d}q$ near x, but this is exactly ω near x.

Therefore we consider the following picture:

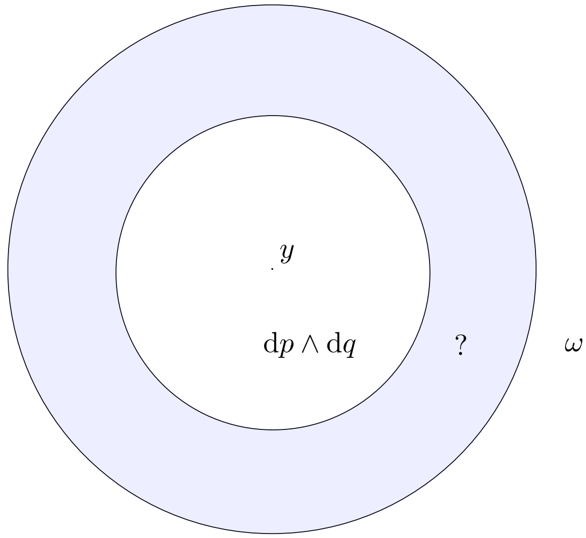

We wish to somehow interpolate and obtain a symplectic form on M.

Remark. In general, adding symplectic forms does not give a symplectic form. If $\rho_1+\rho_2 = 1$ is a partition of unity, $\rho_1\omega_1 + \rho_2\omega_2$ is not closed in general.

To combat this issue, we wish to patch together exact two forms (rather than closed two forms). Given $\mathrm{d}\lambda_1, \mathrm{d}\lambda_2$, the two-form $\mathrm{d}(\rho_1\lambda_1 + \rho_2\lambda_2)$ is clearly closed. However, it is difficult to ensure non-degeneracy: we must have bounds on $\rho_1, \rho_2, \mathrm{d}\rho_1$, and $\mathrm{d}\rho_2$. Thus we must find a one form θ compactly supported near y such that

- $\omega + \mathrm{d}\theta = \mathrm{d}p \wedge \mathrm{d}q$ in a neighbourhood of y,
- $\mathrm{d}\theta$ is small.

This is possible since $\omega - \mathrm{d}p \wedge \mathrm{d}q|_y = 0$. □

A more local (and arguably better) argument uses the *Moser trick* which we explore in the following lecture.

4.2 Lecture 20

Here we present another proof of Darboux's theorem using the Moser trick.

Proof. Set up: (M, ω) is symplectic, $x \in M$, with a neighbourhood U_1 around x with p, q coordinates such that $\omega_x = \mathrm{d}p \wedge \mathrm{d}q$. Consider $t\omega + (1-t)\mathrm{d}p \wedge \mathrm{d}q$. This is a symplectic form for all $t \in [0,1]$ for some neighbourhood $U_2 \subset U_1$ of x. We wish to find a time dependent vector field $V_t : [0,1] \to \Gamma(TU_2)$ such that:

1. $V_t(x) = 0$ for all $t \in [0,1]$. By general ODE theory, there then exists a neighbourhood $U_3 \subset U_2$ such that the flow of V_t leaves U_3 contained in U_2.

2. If the flow is denoted by φ_t, then $\varphi_t^*\omega_t = \omega$ for all t.

The second condition is equivalent to $\frac{\mathrm{d}}{\mathrm{d}t}\varphi_t^*\omega_t = 0$. But

$$\frac{\mathrm{d}}{\mathrm{d}t}\varphi_t^*\omega_t = \varphi_t^*\mathcal{L}_{V_t}\omega_t + \varphi_t^*(\frac{\mathrm{d}}{\mathrm{d}t}\omega_t).$$

(This follows from fixing each term and differentiating the other term.) Moreover, $\frac{\mathrm{d}}{\mathrm{d}t}\omega_t = -\mathcal{L}_{V_t}\omega_t = \omega - \mathrm{d}p \wedge \mathrm{d}q$. On the other hand, $\mathcal{L}_{V_t}\omega_t = \mathrm{d}(\iota_{V_t}\omega_t)$ by the Cartan magic formula. Now suppose $\omega - \mathrm{d}p \wedge \mathrm{d}q$ is closed (as we are still trying to construct a vector field). Then $\mathrm{d}\theta = \mathrm{d}(\iota_{V_t}\omega_t)$. What if we define $\theta = \iota_{V_t}(\omega_t)$? This will uniquely define V_t.

Thus it remains to find a θ such that $\mathrm{d}\theta = \mathrm{d}p \wedge \mathrm{d}q - \omega$. We simply need $\theta(x) = 0$. This can be achieved by choosing any θ and modifying by a constant, or by using integration as in the proof of the Poincaré lemma. □

Loosely speaking, the Moser trick is the process of reducing a problem to the "apparently more difficult" problem of finding a vector field, or one form etc, in the above fashion.

4.2.1 More theorems proved with Moser's trick

Proposition 4.2.1. Let (M, ω) be closed. Suppose ω_t is a $[0,1]$ dependent family of symplectic forms, with $\omega_0 = \omega$, and $[\omega] = [\omega_t] \in H^2_{dR}(M)$ for every t. Then there exists a diffeotopy φ_t such that $\varphi_t^*\omega_t = \omega$. In particular, (M, ω) is symplectomorphic to (M, ω_t) for each t.

Proof. Note that every vector field is complete. We use the Moser argument: i.e. attempt to construct φ_t as the flow of V_t as in the previous proof:

$$0 = \frac{\mathrm{d}}{\mathrm{d}t}\varphi_t^*\omega_t \Longleftarrow \mathrm{d}(\iota_{V_t}\omega_t) = -\frac{\mathrm{d}}{\mathrm{d}t}\omega_t.$$

Note that unlike the previous proof, the term on the right hand side is not constant. However, by the premise, $[\frac{\mathrm{d}}{\mathrm{d}t}\omega_t] = \frac{\mathrm{d}}{\mathrm{d}t}[\omega_t] = 0$. Thus we wish to choose $\zeta_t \in \Omega^1(M)$ such that

1. $\frac{\mathrm{d}}{\mathrm{d}t}\omega_t = \mathrm{d}\zeta_t$

2. ζ_t depends smoothly on t.

How do we find ζ_t? Not entirely trivial: cleanest method is to choose a metric, and use the Hodge theorem. (Using the Hodge star, define the adjoint d^*, and use this to find ζ_t.) Another method is to choose a covering of M and patch together local primitives. A third method is to consider the bundle of all possible ζ_t and find a smooth section. Intuitively, all of the methods work since there is a contractible space of choices. $\square$

Remark. The above theorem does *not* say that if ω_0 and ω_1 belong to the same cohomology class, then there is a smooth family ω_t that connects them. The straight line homotopy ω_t might be degenerate for some t.

Theorem 4.2.2. *Let Σ be a closed surface, with ω_0, ω_1 two area forms. Then (Σ, ω_0) and (Σ, ω_1) are symplectomorphic if and only if the two areas are equal.*

Proof. Key idea for the proof: For every $t \in [0,1]$, observe that ω_0 and ω_1 have the same sign (as the top cohomology is 1-dimensional: area forms essentially live in $(0, \infty)$. Then $t\omega_0 + (1-t)\omega_1$ is a symplectic form for each t. Result now follows from Moser. $\square$

More results can be proven using Moser:

Theorem 4.2.3 (Extension theorem)**.** *Let P be a manifold, and $N \subset P$ a closed submanifold.*

1. *Let Ω be a skew-symmetric bilinear form on $T_N P := TP|_N$, whose restriction to TN is a closed 2-form. Then Ω extends to a closed skew-symmetric 2-form in a neighbourhood of N. If Ω is symplectic, the extension is symplectic.*

2. *Let Ω_0, Ω_1 be symplectic structures on P whose restriction to $T_N P$ are equal. Then there are neighbourhoods $U, V \subset N$ and a symplectomorphism $f : (U, \Omega_0) \to (V, \Omega_1)$ such that $f|_N = \mathrm{id}$ and $\mathrm{d}f|_{T_N P} = \mathrm{id}$.*

Remark. The second result is similar to Darboux's theorem, in which we prescribed that $\omega = \mathrm{d}p \wedge \mathrm{d}q$ at a point x. Here we give this prescription to an entire submanifold.

4.3 Lecture 21

4.3.1 Extension theorem

Recall that at the end of the previous lecture, we stated the *extension theorem*. We give the proof here. The key is the *homotopy formula* from de Rham theory.

Proposition 4.3.1 (Homotopy formula)**.** Let $f, g : X \to Y$, and $F : X \times I \to Y$ a homotopy between them. Then $f^*, g^* : \Omega^*(X) \to \Omega^*(Y)$ are *chain homotopic.* More explicitly, define $h : \Omega^*(Y) \to \Omega^*(X)$ by

$$\omega \mapsto \int_0^1 (i_t^* \iota_{\partial/\partial t} F^*)\omega \mathrm{d}t$$

where i_t is the inclusion $X \to X \times I, x \mapsto (x,t)$. Then h is the chain homotopy.

Proof. This follows from the Cartan magic formula. Namely, one can prove that $f^* - g^* = \mathrm{d}h + h\mathrm{d}$ by computing $\int_0^1 \frac{\mathrm{d}}{\mathrm{d}t} F_t^* \omega$. □

Corollary 4.3.2. *Let $E \to M$ be a a vector bundle, with a disk-subbundle $D \subset E$, equipped with the zero-section z. Then there is a homotopy between $z \circ \pi$ and* id *from D to D, explicitly given by*

$$((v,m),t) \mapsto (tv,m).$$

If ω is a form on D, by the homotopy formula above,

$$\omega - (z \circ \pi)^* \omega = \mathrm{d}h\omega + h\mathrm{d}\omega.$$

Now for any $\alpha \in \Omega^*(D)$, $h\alpha\big|_{T_{z(M)}D} = 0$, since $\iota_{\partial/\partial t} F^* \alpha\big|_{z(M)\times I} = 0$. This is because $\iota_{\partial/\partial t} F^* \alpha(v_1, \ldots, v_k) = \alpha(F_* \partial/\partial t, F_* v_1, \ldots)$, and $F_* \partial/\partial t = 0$ on $z(M) \times I$. Moreover, $\mathrm{d}h\alpha\big|_{T_{z(M)}D} = 0$ if $\alpha\big|_{T_{z(M)}D} = 0$, so by the homotopy formula,

$$\alpha = \alpha - \pi^* z^* \alpha = h\mathrm{d}\alpha + \mathrm{d}h\alpha = h\mathrm{d}\alpha.$$

We are now ready to prove the extension theorem:

Proof. (Part one of the extension theorem.) Given Ω on $T_N P$, we need to extend this to a form in a neighbourhood P of N. We can assume that P is tubular by general differential topology.

Since P is a tubular neighbourhood of N, it is a disk subbundle of the normal bundle of N. Choose any extension of Ω to P, and choose any connection. Observe that at each point (x,v) in P, $T_{(x,v)}P \cong (T_N P)_x$. Call this extension of Ω α. By the homotopy formula,

$$\alpha - \pi^* z^* \alpha = \mathrm{d}h\alpha - h\mathrm{d}\alpha.$$

Then $\alpha - h\mathrm{d}\alpha$ is closed by observing that the two other terms above are closed. Moreover, $\mathrm{d}\alpha$ vanishes on N, so $\alpha - h\mathrm{d}\alpha$ also extends Ω. Thus we have obtained a closed extension of Ω to P, as required.

(Part two of the extension theorem.) Consider Ω_0 and Ω_1 on P, which have the same restriction on N. By the homotopy formula, we have a primitive of their difference: $\Omega_0 - \Omega_1 = \mathrm{d}\eta$. Note that $t\Omega_0 + (1-t)\Omega_1$ is symplectic in a neighbourhood of N. Proceed exactly as in the Darboux theorem and apply the Moser trick. That is, we must find a vector field V_t with $\iota_{V_t}\omega_t = \eta$.

More explicitly, we require f (representing the time-1 flow) to satisfy $\mathrm{d}f\big|_{T_N P} = \mathrm{id}$. We need $\mathrm{d}\varphi_{V_t}^t : T_n P \to T_n P, n \in N$ to be the identity. Check that this follows from $\mathrm{d}\eta = 0$. □

Remark. "This lecture will get more confusing as we keep going, with normal bundles of the subbundle of the normal bundle, there will be like two more layers..."

In fact, the lecture finished prematurely so we didn't observe the confusion remarked above.

4.4 Lecture 22

4.4.1 Applications of the extension theorem

Turned up to the lecture late! Oh no!

Last week we looked at the *extension theorem*, which states that if $N \subset P$ is a submanifold, then

1. if Ω on $T_N P$ is closed on TN, then it has a closed extension to a neighbourhood of N.

2. if $\Omega_0 = \Omega_1$ on $T_N P$, then there is a neighbourhood of N in which they are symplectomorphic.

Some applications of this theorem are as follows:

Example. Darboux's theorem: Let E be the normal bundle of $N \subset P$. By the tubular neighbourhood theorem, there is an isomorphism $T_N E \to T_N P$ which fixes the subbundle TN. Therefore there exist open neighbourhoods $U \supset N \subset E, V \supset N \subset P$, and a diffeomorphism $U \to V$, which induces the above isomorphism.

Example. Let $L \subset M$ be a Lagrangian submanifold. Then there is an isomorphism between its normal bundle and cotangent bundle. To see this, choose a Lagrangian complement subbundle to $TL \subset T_L M$. There is an isomorphism $T_L M \cong T^*L$. This is formalised by the Weinstein neighbourhood theorem:

Theorem 4.4.1 (Weinstein neighbourhood theorem)**.** *Let $L \subset M$. There exist neighbourhoods $U \supset L \subset T^*L$, $V \supset L \subset M$, such that there is a symplectomorphism between U and V which fixes L.*

Example. Let $C \subset M$ be isotropic. Here the normal bundle and $(T_C M, \Omega)$ depends on more than just C: by linear algebra, if V is a symplectic vector space, and $W \subset V$ is isotropic, then W^Ω is coistropic. The normal direction is as follows:

1. Choose any complement to W in W^Ω, say S.

2. S^Ω is also symplectic.

3. $W \subset S^\Omega$ is Lagrangian.

4. Choose Lagrangian complement to W in S^Ω.

Now V is symplectomorphic to $W \oplus W^* \oplus S$, where each W, W^*, S have canonical symplectic structures. In each case, the freedom in the choice of sympletic structure is contractible.

Now $T_C M \cong S(C) \oplus TC \oplus T^*C$, where $S(C) \cong TC^\Omega / TC$. Thus we have a symplectomorphism $T_C M/TC \cong S(C) \oplus T^*C$. By the extension theorem, this is a symplectomorphism between neighbourhoods of C. The neighbourhoods of C correspond to symplectic vector bundles over C with rank $\dim M - 2\dim C$. i.e. *all neighbourhoods of isotropic submanifolds look like $S(C) \oplus T^*C$.*

Remark. "Many years after Weinstein's work, someone published a two page paper on the above result which has thousands of citations but it's basically a homework problem!"

Example. What do neighbourhoods of symplectic submanifolds look like? How about neighbourhoods of coisotropic submanifolds?

4.4.2 Gromov non-squeezing

We established in the previous few lectures that there are no local invariants of symplectic manifolds. How about global invariants? Gromov discovered a global invariant (other than volume):

Theorem 4.4.2 (Gromov non-squeezing)**.** *Suppose $B(r) \subset \mathbb{R}^{2n}$ is a ball of radius r. Let $Z(R) = D(R) \times \mathbb{R}^{2n-2}$ be a cylinder, where $D(R) \subset \mathbb{R}^2$ is the disk of radius R. Then if $B(r)$ embeds symplectically into $Z(R)$, then $R \geq r$.*

Conversely, a slight modification can remove the obstruction all together.

Theorem 4.4.3 (Gutn, Polterovich)**.** *Let Σ be T^2 with an open disk removed. Suppose Σ has area 1. For all $r > 0$, $B(r)$ symplectically embeds in $\Sigma \times \mathbb{R}^2$.*

Observe that if $T^2 \setminus D$ is replaced with $\mathbb{S}^2 \setminus D$, we are in the realm of Gromov non-squeezing. This shows a dichotomy between full rigidity and zero rigidity.

Remark. Differential topology is very soft, while Riemannian geometry is very rigid. Symplectic geometry is in between, sometimes it is soft and sometimes it's rigid.

Remark. For any R, there are symplectic embeddings $B^n(R) \subset \mathbb{R}^n$ with arbitrarily large percentage (volumewise) of $B^n(\mathbb{R})$ in $Z(1)$. These are called *Katak embeddings.*

In the following lectures, a sketch of Gromov non-squeezing will be given. First some definitions are required:

Definition 4.4.4. Let (Σ, j) be a Riemann surface, and (M, J) an almost-complex manifold. A map $u : \Sigma \to M$ such that

$$\mathrm{d}u \circ j = J \circ \mathrm{d}u$$

is called a *(j, J)-holomorphic curve.* (Or just *J-holomorphic.*)

If one chooses holomorphic coordinates $z = s + it$ on Σ, then the above condition can be re-written as

$$\frac{\partial u}{\partial t} = J \circ \frac{\partial u}{\partial s}.$$

4.5 Lecture 23

4.5.1 Gromov non-squeezing: proof idea (in only one lecture?!)

Theorem 4.5.1 (Gromov, 85)**.** *Let $B(r) \subset \mathbb{R}^{2n}$ be an open ball of radius $r > 0$. Let $Z(R) \coloneqq D^2(R) \times \mathbb{R}^{2n-2} \subset \mathbb{R}^{2n}$ be the cylinder of radius $R > 0$. Then there is a symplectic embedding $\varphi : B(r) \hookrightarrow Z(R)$ if and only if $r \leq R$.*

Remark. Intuitively, any time symplectic structures appear to have rigidity that contradicts Riemannian rigidity, it can be thought of as a J-holomorphic curve adding an obstruction.

Proof. Assume there is an embedding $\varphi : B(r) \hookrightarrow Z(R)$. Let $\varepsilon > 0$. Then $\varphi(B(r - \varepsilon)) \subset D^2 \times [-A, A]^{2n-2}$. There is a symplectic embedding $D^2(R) \hookrightarrow S^2(\pi R^2 + \varepsilon)$. This gives a symplectic embedding

$$\widetilde{\varphi} : B(r - \varepsilon) \hookrightarrow S^2(\pi R^2 + \varepsilon) \times T^{2n-2}.$$

The latter space is equipped with two form $\sigma \oplus \tau$. The actual theorem we prove is the following:

Theorem 4.5.2. *Let (M, τ) be symplectic with dimension $2n - 2$. If there is exists a symplectic embedding $\varphi : B(r) \hookrightarrow S^2(a) \times M$, with $S^2(a) \times M$ equipped with $\sigma \otimes \tau$, then $\pi r^2 \leq a$.*

non-squeezing is now a corollary: If $\pi(r - \varepsilon)^2 \leq \pi R^2 + \varepsilon$ for all $\varepsilon > 0$, then $r \leq R$.

Recall the definition of a pseudo-holomorphic, or J-holomorphic curve: $u : (\Sigma, j) \to (X, J)$ is J-holomorphic, if $\mathrm{d}u \circ j = J \circ \mathrm{d}u$.

The above theorem is proven using the following two lemmas:

Lemma 4.5.3. (Main lemma) Suppose $\pi_2(M) = 0$. There exists an ω-compatible almost complex structure J on $S^2 \times M$ such that $\varphi^* J = i$ on $B(r)$, and a J-holomorphic curve $u : \mathbb{CP}^1 \to S^2 \times M$ such that

1. $[u] = [S^2 \times \{\mathrm{pt}\}] \in H_2(S^2 \times M)$.
2. $\varphi(0) \in \mathrm{im}(u)$.

Lemma 4.5.4. (Monotonicity lemma) (Minimal surface theory) Suppose $\widetilde{u} : (\Sigma, j) \to (B(r), i)$ is non-constant, proper, pseudo-holomorphic, with $0 \in \mathrm{im}(\widetilde{u})$. Then the area of $\widetilde{u}$ is greater than or equal to πr^2. Visually this corresponds to the interesting result that the way to slice a sphere with a surface that passes through the origin in a way that minimise area is to use a plane. Any perturbation will have slightly more area coming from additional wobbliness.

It follows that, if we assume the above two lemmas, then

$$\pi r^2 \leq \mathrm{Area}(\widetilde{u}) \leq \mathrm{Area}(u) \leq a,$$

where the first inequality comes from the monotonicity lemma and the second comes from the symplectic embedding. The last equality is because

$$\int u^*\omega = \omega([S^2 \times \mathrm{pt}]) = \sigma([u]) = a.$$

Remark. Note that the previous two non-main lemmas are slightly false, but they contain the right ideas.

The idea of the proof of the main lemma is as follows: Let

$$\mathcal{J} = \{\omega\text{-compatible almost complex structure on } S^2 \times M\},$$

suppose it is $W^{k,2}$. For $J \in \mathcal{J}$, let $\widetilde{M}_J = \{u : \mathbb{CP}^1 \to S^2 \times M : uJ\text{-holomorphic}, u(\infty) = \varphi(0), [u] = [S^2 \times \mathrm{pt}]\}$. We establish a correspondence

$$\widetilde{M}_J = \overline{\partial}_J^{-1}(0) \subset B_J \overset{\overline{\partial}_J}{\longleftrightarrow} \mathcal{E}_J = \bigcup_u W^{k-1}\Gamma\overline{\hom}_{\mathbb{C}}(T\mathbb{CP}^1, u^*T(S^2 \times M)).$$

Here B_J is the Banach manifold consisting of maps in $W^{k,p}$ rather than smooth maps, allowing us to use methods from analysis. $\overline{\partial}_J = J \circ \mathrm{d}u \circ j + \mathrm{d}u$ is the antiholomorphic differential. Thus it vanishes if and only if u is J-holomorphic. We wish to study the linearisation

$$D_u\overline{\partial}_J : T_uB_J \to \Gamma u^*T(S^2 \times M) \to \mathcal{E}_{Ju}.$$

This is a Cauchy Riemann operator, allowing us to use the Riemann Roch theorem!

Implicit function theorem: If $D_u\overline{\partial}_J$ is surjective for all $u \in \overline{\partial}_J^{-1}(0)$, then $\widetilde{M}_J$ is smooth with dimension 4. (The dimension comes from the Riemann Roch theorem: $n\chi(\mathbb{CP}^1) + c_1([u]) - 2n = 2n + 4 - 2n = 4$.

Observe that $\mathrm{Aut}(\mathbb{CP}^1, \infty)$ has dimension 4, since the automorphisms of $\mathbb{CP}^1$ are Mobius transformations which have dimension 6 (mapping any three points to any three points), but now we are forced to fix a point (any two points go to any two points). Thus

$$M_J = \widetilde{M}_J / \mathrm{Aut}(\mathbb{CP}^1, \infty)$$

is smooth with dimension zero. Since $[u] = [S^2 \times \mathrm{pt}]$, M_J is compact.

A proof outline for the remainder of the proof is as follows:

1. $J_0 = i_{\mathbb{CP}^1} \times i_{T^{2n-2}}$ is Fredholm regular, so the implicit function theorem applies. Thus $|M_{J_0}| = 1$.

2. There exists J_1 Fredholm regular such that $\varphi^* J_1 = i$. (Existence of J_1 is very lucky, due to the choice of homology class $[S^2 \times \mathrm{pt}]$.)

3. For a generic path $J_t, t \in [0,1]$ from J_0 to J_1, $M = \cup_{t\in[0,1]}\{t\} \times M_{J_t}$ has the structure of a compact 1-dimensional manifold, with ∂M the disjoint union of M_{J_0} and M_{J_1}. This gives a cobordism from M_{J_0} to M_{J_1}. But if M_{J_0} consists of a single point, the existence of the cobordism ensures that M_{J_1} is non-empty (since their cardinalities must agree mod 2).

4. Since M_{J_1} is non-empty, and J_1 satisfies the conditions of the main lemma, the result follows.

□

4.6 Lecture 24

We began with a summary of the main points in the proof of Gromov non-squeezing, namely the two lemmas mentioned in the proof from the previous lecture.

Remark. Symplectic geometry has two eras... one is pre-Gromov, one is post-Gromov.

A point that needs justifying from the above proof: why is there a J-holomorphic curve passing through every point in $S^2(a) \times T^2(A)$ for generic J?

Remark. Riemannian geometry is a huge field... 50% of it is only known by Gromov.

4.6.1 Lagrangians in linear symplectic reductions

Recall symplectic reduction: for symplectic vector spaces V, given $Q \subset V$ a subspace, the symplectic reduction of Q is $Q/Q \cap Q^\Omega$. In a more specific setting, given a coisotropic $Q \subset V$, the symplectic reduction is simply Q/Q^Ω.

Lemma 4.6.1. Suppose $L \subset V$ is Lagrangian, and $W \subset V$ coisotropic. Then the image of $L \cap W \to W/W^\Omega$ is a Lagrangian subspace of W/W^Ω.

Proof. It must be shown that L_W (the image of $L \cap W$) is isotropic and coisotropic. It is immediately isotropic. For the reverse direction, choose any $[v] \in {L_W}^\Omega$, i.e. any v with $\omega(v,l) = 0$ for all $l \in L \cap W$, so $v \in (L \cap W)^\Omega$.

Observe that $(L \cap W)^\Omega = L^\Omega + W^\Omega = L + W^\Omega$. Choose any $v \in (L \cap W)^\Omega$. This decomposes as $v_1 + v_2$ in $L + W^\Omega$. Since $W^\Omega \subset W$, it must be the case that $v_2 \in W$. Thus $v_1 = v - v_2$ lies in W. This gives $v \in (L \cap W) + W^\Omega$, so $v \in L_W$ as required. □

Suppose L and W are transverse; $L+W=V$. Suppose V has dimension $2n$, so that L has dimension n. Then W has dimension $n+k$, and $L\cap W$ has dimension k, with W/W^Ω having dimension $(n+k)-(n-k)=2k$. The image of $L\cap W\to W/W^\Omega$ is isomorphic to $L\cap W/L\cap W^\Omega$ and has dimension k, so $L\cap W^\Omega=0$. Taking these into account, Umut suspects that there should be a simpler proof of the above lemma in this special case.

In summary, the following has been established: suppose V is a symplectic vector space, $W\subset V$ is coisotropic, and $L\subset V$ is Lagrangian.

- Let L_W denote the image of $L\cap W\to W/W^\Omega$. Then $L_W\subset W/W^\Omega$ is a Lagrangian subspace.
- If in addition L is transverse to W, then $L_W\cong L\cap W$.

4.7 Lecture 25

4.7.1 Lagrangians in general symplectic reductions

We continue the discussion from the previous lecture, in the vector bundle setting manifold settings.

Symplectic vector bundle case: Let $E\to M$ be a symplectic vector bundle, and $Q\subset E$ a coisotropic subbundle. This gives a symplectic vector bundle $Q/Q^\Omega\to M$. Let $L\subset E$ be a Lagrangian subbundle. We make the assumption that $L\cap Q$ has constant rank, so that it defines a subbundle. Then the image $L\cap Q\to Q/Q^\Omega$ is also a subbundle. To see this, observe that at any x in the base space, the image of the quotient is isomorphic to $L_x\cap Q_x/(L_x\cap Q_x^\Omega)$, but

$$L_x\cap Q_x^\Omega=L_x^\Omega\cap Q_x^\Omega=(L_x+Q_x)^\Omega,$$

and the fibres of the last object forms a bundle if $L_x\cap Q_x$ does. Therefore the image of $L\cap Q$ is a vector bundle; the *reduced Lagrangian subbundle.*

Symplectic manifold case: Suppose M is a symplectic manifold, $C\subset M$ coisotropic, and $L\subset M$ Lagrangian. We make the following assumptions:

- C is foliated by isotropic leaves.
- The leaf space X is itself a manifold, e.g. the leaves are the fibres of a submersion $C\to X$.

Then X has a natural symplectic structure. (See earlier lecture on integrability for details.) To obtain stronger results, we make more assumptions:

- L intersects C "cleanly", i.e. $L\cap C$ is a submanifold, and for all $a\in L\cap C$, $T_a(L\cap C)=T_aL\cap T_aC$.

Note that this assumption is weaker than transversality.

We wish to obtain a Lagrangian inside X, given a smooth map $g : L \cap C \hookrightarrow C \to X$. This corresponds to the differential

$$T_a L \cap T_a C \to T_a C/(T_a C)^{\Omega},$$

where the codomain is established by the vector bundle setting. This has a Lagrangian image: Since g has constant rank, by the constant rank theorem, the image is an immersed Lagrangian. If L and C are transverse, then $L \cap C \to X$ is an immersion.

Remark. The above map is not an embedding, since L can intersect a leaf more than once.

4.7.2 Symplectic category

Definition 4.7.1. A *Lagrangian correspondence*, or *canonical relation* from $(M, \omega_M) \to (N, \omega_N)$ is a Lagrangian submanifold in the twisted product $M \widetilde{\times} N$.

Recall that the twisted product is $M \times N$ equipped with the symplectic form $\pi_M^* \omega_M - \pi_N^* \omega_N$.

Example. Suppose $\varphi : M \to N$ is a symplectomorphism. Then the graph $G(\varphi) \subset M\widetilde{N}$ is a Lagrangian correspondence.

Example. Suppose G is a Hamiltonian action on M, with moment map $\mu : M \to \mathfrak{g}^*$. Assume the action of G on $\mu^{-1}(0)$ is free. Then $\mu^{-1}(0) \hookrightarrow M \times M /\!\!/ G$ is a Lagrangian correspondence. (The above map is injective, and its image lifted to $M \widetilde{\times} M$ is Lagrangian.) More explicitly, $\mu^{-1}(0)$ is a coisotropic submanifold with leaves given by the orbits of G.

Example. An example of the above example: $\mathbb{S}^{2n-1} \to \mathbb{C}^n \times \mathbb{CP}^{n-1}$.

Proposition 4.7.2. Lagrangian correspondences can be composed: If $L_1 \subset P_1 \widetilde{\times} P_2$, $L_2 \subset P_2 \widetilde{\times} P_3$, then there is

$$L_1 \circ L_2 \subset P_1 \widetilde{\times} P_3.$$

Proof. $L_1 \times L_2 \subset (P_1 \widetilde{\times} P_2) \times (P_2 \widetilde{\times} P_3)$ is Lagrangian. Consider the diagonal $\Delta_{P_2} = -P_2 \widetilde{\times} P_2$. Then $P_1 \times \Delta_{P_2} \times P_3$ is a coisotropic submanifold of $(P_1 \widetilde{\times} P_2) \times (P_2 \widetilde{\times} P_3)$. The reduction of this coistropic submanifold is exactly $P_1 \widetilde{\times} P_3$. By the earlier discussion of symplectic reduction for manifolds, if the intersection of $L_1 \times L_2$ with $P_1 \times \Delta_{P_2} \times P_3$ is clean, this gives an immersed Lagrangian in $P_1 \widetilde{\times} P_3$. $\square$

4.8 Lecture 26

4.8.1 Exact Lagrangians

Definition 4.8.1. Let $(T^*M, \mathrm{d}\lambda_{\mathrm{taut}})$ be the canonical symplectic manifold. A Lagrangian $L \subset T^*M$ is called *exact* if $\lambda|_L$ is exact.

Example.

- The zero section of T^*M is exact. In fact, this is very much the case as $\lambda|_{Z_M} \equiv 0$.
- Suppose $\varphi : T^*M \to T^*M$ is a Hamiltonian diffeomorphism. Then $\varphi(Z_M)$ is exact. This follows from a flux argument: Choose a Hamiltonian isotopy and compute integral of λ over 1-cycles on $\varphi(Z_M)$ via Stokes' theorem.

Nearby Lagrangian conjecture Any compact exact Lagrangian in T^*M is Hamiltonian-isotopic to the zero section.

Remark. The fact that this is an open problem is an example of the general construction problem in symplecitc geometry.

Definition 4.8.2. Let $\pi : E \to M$ be a submersion. (E.g. the projection map of a vector bundle.) Let $S : E \to \mathbb{R}$ be a function. Inside T^*E, there is a coisotropic submanifold

$$N \coloneqq \{(e, \xi) : \xi|_{\ker \mathrm{d}\pi_e} = 0\}.$$

Then there is an inclusion

$$N \hookrightarrow T^*E \to T^*M,$$

where $N \subset T^*E$ is coisotropic with isotropic leaves given by the fibres of $N \to T^*M$.

Now if the graph $G(\mathrm{d}S)$ in T^*E is transverse to N, then there is a reduced Lagrangian submanifold $\Sigma_S \coloneqq \iota_S : N \cap G(\mathrm{d}S) \to T^*M$. This is the *Lagrangian generated by* S.

There is an injection $\iota_S : \Sigma_S \hookrightarrow T^*E \to E$, allowing Σ_S to be considered as a submanifold of E. Σ_S corresponds to "vertical critical points" of S. ι_S records the horizontal part of $\mathrm{d}S$ at a point of $\Sigma_S \subset E$.

In coordinates, let $E = \mathbb{R}^{n+k}$, $M = \mathbb{R}^n$, $(\xi, \eta) \mapsto \eta$. Then

$$\Sigma_S = \{(\xi, \eta) : \partial S/\partial \xi(\xi, \eta) = 0\},$$

and the inclusion map is given by $\Sigma_S \hookrightarrow T^*\mathbb{R}^n \cong \mathbb{R}^n \times \mathbb{R}^n$, $(\xi, \eta) \mapsto (\partial S/\partial \xi(\xi, \eta), \eta)$.

Remark. This is always an immersion since Lagrangian reduction is an immersion, but it need not be an embedding.

Example. $S : \mathbb{R}_a \times R_x \to \mathbb{R}$, $(a, x) \mapsto a^3/3 + (x^2 - 1)a$. Then

$$\Sigma_S = \{a^2 + (x^2 - 1) = 0\} \subset \mathbb{R}^2, \iota_S(a, x) \mapsto (2xa, x).$$

Therefore the immersion sends circles at each a to a figure-8 with a double point at the origin.

In this example, "N" (which was seen in the definition of generating functions) is $\{(a, x, p_a, p_x) : p_a = 0\}$.

Lemma 4.8.3. $\iota_S^* \lambda_{\text{taut}} = \mathrm{d}(S|_{\Sigma_S})$, so ι_S is an exact Lagrangian immersion. This is trivial when computing in coordinates.

Theorem 4.8.4. *Every Lagrangian $L \subset T^*M$ (for L, M compact) that is Hamiltonian-isotopic to the zero section Z_M admits a generating function. In fact, this generating function is special; it can be chosen such that $S : \mathbb{R}^N \times M \to \mathbb{R}$ and S is quadratic at infinity (in the $\mathbb{R}^N$ direction).*

Proof. The following will be a construction that doesn't work immediately, but contains inspirational ideas. Let $\gamma : [0, 1] \to T^*M, \gamma(t) = (p(t), q(t))$. The *action functional*, where $H_t : T^*M \times [0, 1] \to \mathbb{R}$ is fixed, is given by

$$A(\gamma) = \int_{[0,1]} \gamma^* \lambda_{\text{taut}} + \int_{[0,1]} H_t(\gamma(t))\mathrm{d}t = \int_0^1 p(t)\dot{q}(t) + H_t(p(t), q(t))\mathrm{d}t.$$

Claim: consider $E \coloneqq \{\gamma : [0, 1] \to TM : \gamma(0) \in Z_M\}$. Then $E \to M, \gamma \mapsto \pi(\gamma(1))$ is a bundle, where $A : E \to \mathbb{R}$ is the generating function. Then A "generates" the time 1 image of Z_M under the flow of H_t.

Exercise for the reader: If $\gamma(t) = (p(t), q(t)) \mapsto (p(t) + \delta p(t), q(t) + \delta q(t))$, then $A(\gamma) \mapsto A(\gamma) + \delta A$, where δA is given by

$$\delta A = \int \delta p(t)(\dot{q}(t) + \partial H_t/\partial p(p(t), q(t))) - \int \delta q(t)(\dot{p}(t) - \partial H_t/\partial q(p(t), q(t))) + p(1)\delta q(1).$$

In this formula, $\delta q(1)$ is zero for vertical directions. That is, $\gamma \in \Sigma_S$ if and only if γ satisfies Hamilton's equations. Moreover, if $\Sigma_S \to T^*M$ is given by $\gamma \mapsto (p(1), q(1))$, the image is precisely the image of Z_M under the Hamiltonian flow of H_t. This is explored further in McDuff-Salamon. $\square$

4.9 Lecture 27

In subsequent lectures, each student (as well as the lecturer!) will give a 10 minute presentation on a topic from *Symplectic rigidity: Lagrangian submanifolds* they found interesting.

- Monday: Umut, Jordi
- Wednesday: Eric, Alec, Shintaro
- Friday: Jared, Darius, Juan

4.9.1 Poincaré's last geeometric theorem (Umut's talk)

Theorem 4.9.1. *Let $A = \mathbb{S}^1_\theta \times [0,1]_x$ be an annulus (with area form $\mathrm{d}\theta\mathrm{d}x$). If $\varphi : A \to A$ is an area-preserving diffeomorphism which "twists boundaries in opposite directions", then φ has at least two fixed points.*

This theorem was motivated by Poincaré's study of the 3-body problem. The idea was to consider a Hamiltonian motion in 4-dimensional phase space, and to then restrict to an energy level. Each energy level is 3-dimensional. We wish to understand the energy levels near a given energy level.

Claim: Nearby periodic orbits are in correspondence with fixed points of the *Poincaré return map.* This is an area preserving map which can, without loss of generality, be considered on the annulus. More precisely, the linearisation of the problem lives in the annulus.

This is captured by the following diagram:

$$\begin{tikzcd}
\mathbb{R} \times I \arrow[r] \arrow[d, "\widetilde{\varphi}"] & \mathbb{S}^1 \times I \arrow[d, "\varphi"] \\
\mathbb{R} \times I \arrow[r] & \mathbb{S}^1 \times I
\end{tikzcd}$$

We have a "boundary twisting property" which can be formalised as:

$$\widetilde{\varphi}(\widetilde{\theta} + 1, x) = \widetilde{\varphi}(\widetilde{\theta}, x) + (1, 0).$$

We wish to translate Poincaré's problem into one of fixed points of Hamiltonian diffeomorphisms. Assume for simplicity that near the boundary of the annulus $\widetilde{\varphi}$ is a translation in $\widetilde{\theta}$.

Yasha's trick: consider two strips placed side by side. Consider any arc on the annulus from one boundary component to another. On one of these strips, apply $\widetilde{\varphi}$. Glue the strips together where two arcs meet the boundary, consider area change. This is in a video by Yasha.

Claim: Every symplectomorphim of the torus which fixes a parallel and moves a meridian by zero flux has at least three fixed points. This is solved by a Moser type argument.

Monday: Arnol'd conjecture.

4.9.2 Wavefronts (Jordi's talk)

Consider Lagrangian immersions (rather than assuming Lagrangian submanifolds).

Definition 4.9.2. Let M be an n-manifold, and $f : M \to \mathbb{C}^n$. If $\mathrm{d}f$ is injective and its image is a Lagrangian submanifold, then f is a Lagrangian immersion.

Recall: canonical form $\mathrm{d}p \wedge \mathrm{d}q$ is exact; $\omega = \mathrm{d}\lambda$. Whenever $f^*\omega = 0$, $f^*\lambda$ is itself closed.

Definition 4.9.3. A Lagrangian immersion $f : M \to \mathbb{C}^n$ is called exact if $f^*\lambda$ is exact, i.e. there exists $F : M \to \mathbb{R}, f^*\lambda = \mathrm{d}F$.

Idea: $f \times F : M \to \mathbb{C}^n \times \mathbb{R}$, Legendrian immersion into a contact manifold.
We are ready to consider wavefronts: These will help understand n-dimensional submanifolds in $\mathbb{R}^n \times \mathbb{R}^n$ as n-dimensional submanifolds of $\mathbb{R}^n \times \mathbb{R}$.

$$\pi = \mathrm{Re}(f \times F) : M \to \mathbb{R}^n \times \mathbb{R}.$$

Definition 4.9.4. A *wavefront* is the image of the Lagrangian immersion

$$\mathrm{d}z - \sum_j y_j \mathrm{d}x_j, (f \times F)^*(\mathrm{d}z - \sum_j y_j \mathrm{d}x_j) = 0, z = \sum_j y_j x_j.$$

Example: Whitney immersion. $f : \mathbb{S}^n \to \mathbb{C}^n, (x, a) \mapsto (1 + 2ia)x$. Sends a "double tear drop" to a "figure 8". Somehow reversible!

4.10 Lecture 28

4.10.1 Surgery theory (Eric's talk)

See Jared Marx-Kuo for lecture notes (was late to class).

4.10.2 Soft Lagrangian obstructions (My talk)

A big general question in symplectic geometry is the following:

Which n-manifolds L admit Lagrangian embeddings into $\mathbb{C}^n$?

By a Lagrangian embedding, we mean the image of L is an embedded Lagrangian submanifold of $\mathbb{C}^n$, where $\mathbb{C}^n$ has the standard compatible symplectic structure.

The standard symplectic structure ω of $\mathbb{C}^n$ is defined by

$$\omega(z_j, w_j) = \Re(z_j)\Im(w_j) - \Im(z_j)\Re(w_j).$$

Using the fact that $g(-,-) = \omega(-, i-)$ is an inner product on $\mathbb{C}^n$, one can prove that a submanifold $L \subset \mathbb{C}^n$ is Lagrangian if and only if $TL \oplus iTL = T\mathbb{C}^n$.

Definition 4.10.1. A *totally real embedding* of a manifold L into $\mathbb{C}^n$ is an embedding such that T_xL and iT_xL are transverse. (Thus every Lagrangian embedding is totally real.) Totally real embeddings are referred to as *soft Lagrangian embeddings*.

We can now discuss soft and hard obstructions to the existence of Lagrangian embeddings in $\mathbb{C}^n$. "Soft obstructions" refer to obstructions to the existence of soft Lagrangian embeddings. These tend to be topological without requiring symplectic rigidity results. On the other hand, "hard obstructions" are specifically obstructions to Lagrangian embeddings. These truly require a consideration of the symplectic structure, e.g. Maslov classes and symplectic area classes.

A cool observation is that $\mathbb{S}^3$ has no soft obstructions, even though it admits no Lagrangian embedding in $\mathbb{C}^3$. In that sense, soft Lagrangian embeddings are truly softer than Lagrangian embeddings.

Definition 4.10.2. Let $f : L \to \mathbb{C}^n$ be an immersion of an n-real dimensional manifold L. There is a regular homotopy invariant $d(f)$ defined to be the signed count of double points of the immersion. (For n even and L oriented, there is a natural signed count. Otherwise, blackbox.)

Theorem 4.10.3 (Whitney, classical theorem)**.** *Suppose $n = 2k$ is even, and let $f : L \to \mathbb{C}^n$ be a Lagrangian immersion with normal crossings. Then*

$$d(f) = (-1)^{k+1}\chi(L)/2.$$

(equality is mod 2 if L is non-orientable.)

Theorem 4.10.4. *A closed connected n-manifold L admits totally real embeddings if and only if it is U-parallelisable and*

1. *if n is even, $\chi(L) = 0$. (Equality is mod 4 if L is not orientable.)*
2. *If $n \equiv 1 \mod 4$ and L is orientable, $\widehat{\chi}(L) = 0$.*

Here $\widehat{\chi}(L)$ is the *Kervaire semi-characteristic* which essentially plays the same role as the Euler characteristic.

I've introduced a definition without actually defining it. What is U-parallelisability? This is a complex version of the requirement that a manifold has a trivial tangent bundle. An n dimensional manifold M is *parallelisable* if

$$TM \cong M \times \mathbb{R}^n.$$

M is *U-parallelisable* if

$$TM \otimes_{\mathbb{R}} \mathbb{C} \cong L \times \mathbb{C}^n.$$

As we know, most spheres have tangent bundles which are non-trivial. The only counterexamples are $\mathbb{S}^1, \mathbb{S}^3$, and $\mathbb{S}^7$. (E.g. $\mathbb{S}^2$ has a non-trivial tangent bundle by the hairy ball theorem.) Using this fact we can attempt to determine which spheres have Lagrangian embeddings.

1. All spheres of dimension $n \neq 1, 3, 7$ have no soft Lagrangian embeddings (and hence no Lagrangian embeddings).

2. $\mathbb{S}^1$ has a Lagrangian embedding, since the inclusion map $\mathbb{S}^1 \hookrightarrow \mathbb{C}$ is a Lagrangian embedding.

3. By considering the forgetful functor from Lagrangian immersions to topological immersions, one can conclude that $\mathbb{S}^7$ has no totally real immersions, but $\mathbb{S}^3$ does have totally real immersions.

This is noteworthy, since $\mathbb{S}^3$ has a totally real immersion in $\mathbb{C}^3$, but no Lagrangian embedding into $\mathbb{C}^3$. One can see this by using a *hard Lagrangian obstruction*. I.e. noting that every closed Lagrangian submanifold of $\mathbb{C}^n$ has a non-vanishing symplectic area class, hence a non-trivial first cohomology group. This automatically bars $\mathbb{S}^n$ from having a Lagrangian embedding for all $n > 1$. In particular, this holds for $n = 3$.

4.10.3 h-principle (Alec's talk)

Recall: If $f : L \to \mathbb{C}^n$ is Lagrangian, then $TL \oplus iTL \cong T\mathbb{C}^n$. This gives an isomorphism $TL \otimes \mathbb{C} \to \mathbb{C}^n$, so any Lagrangian immersion gives rise to a U-parallelisation.

Proposition 4.10.5. All Lagrangian immersions have U-parallelisations.

h-principle is the converse.

Theorem 4.10.6. *The set of Lagrangian regular homotopy classes of Lagrangian immersions $L \to \mathbb{C}^n$ is a* 1-1 *correspondence with homotopy classes of U-parallelisations.*

What are some nice examples?
Any n-dimensional $\mathbb{C}$ vector bundle with n-dimensional base....
Out of time! Next lecture angrily stormed into the lecture theatre.