

NOTES ON SYMPLECTIC TOPOLOGY

DANNY CALEGARI

ABSTRACT. These are notes on symplectic topology, based on a graduate course taught at the University of Chicago in Winter 2022

CONTENTS

1. Linear algebra	1
2. Symplectic Manifolds	10
3. Hamiltonian Flows	18
4. Holomorphic Curves	25
5. Floer Homology	43
References	46

1. LINEAR ALGEBRA

“You come to grad school and you think you’re going to prove the Poincaré Conjecture. And then six years later your thesis comes down to computing the entries of a three-by-three matrix.”

– Vinay Kathotia, motivational speech to incoming Berkeley math grads, 1995.

1.1. Symplectic vector spaces.

Definition 1.1 (Symplectic vector space). A *symplectic vector space* is a real, finite dimensional vector space V with a bilinear form $\omega : V \times V \rightarrow \mathbb{R}$ satisfying

- (1) antisymmetry: for $v, w \in V$ we have $\omega(v, w) = -\omega(w, v)$; and
- (2) nondegeneracy: for all nonzero $v \in V$ there is $w \in W$ with $\omega(v, w) \neq 0$.

An ω satisfying these conditions is said to be a *symplectic form* on V .

Example 1.2. Let $V = \mathbb{R}^{2n}$ with basis $x_1, \dots, x_n, y_1, \dots, y_n$. The *standard symplectic form* ω on \mathbb{R}^{2n} satisfies $\omega(x_i, x_j) = \omega(y_i, y_j) = 0$ for all i, j and $\omega(x_i, y_j) = -\omega(y_j, x_i) = \delta_{ij}$.

In terms of matrices,

$$\omega(v, w) = \langle Jv, w \rangle = w^T Jv = v^T J^T w = -v^T Jw$$

where J is the $2n \times 2n$ matrix which in block form is $J := \begin{pmatrix} 0 & -I_n \\ I_n & 0 \end{pmatrix}$. Note that $J^{-1} = -J = J^T$ and $J^2 = -I_{2n}$.

Date: March 11, 2022.

Antisymmetry means just that $\omega \in \Lambda^2 V$. Nondegeneracy is an open condition, and is dense if V is even dimensional (if V is odd dimensional any alternating form has a nontrivial kernel). If V has dimension $2n$ we can form the top exterior power $\omega^n := \omega \wedge \cdots \wedge \omega$ of any $\omega \in \Lambda^2 V$. Then ω is nondegenerate if and only if $\omega^n \neq 0$ in $\Lambda^{2n} V$. The group $\text{GL}(V)$ acts on V and therefore also on $\Lambda^2 V$, and it permutes the nondegenerate forms. In fact, the set of nondegenerate forms is a single $\text{GL}(V)$ orbit. In other words:

Lemma 1.3. *Any two symplectic forms on a real vector space are related by an automorphism.*

Proof. We shall show if ω is a symplectic form on V , then V is even dimensional, and there is a basis for V in which ω becomes the standard symplectic form.

First of all, for any bilinear form q and any subspace $W \subset V$ we define the *complement* W^\perp to be the set of $v \in V$ so that $q(w, v) = 0$ for all $w \in W$. There is a map $\phi : V \rightarrow W^*$ defined by $\phi(v) = q(\cdot, v)$ and if q is nondegenerate then ϕ is surjective with kernel W^\perp . Thus $\dim(W^\perp) + \dim(W) = \dim(V)$. Furthermore, (V, ω) splits as an orthogonal sum $(W, \omega|_W) + (W^\perp, \omega|_{W^\perp})$ if and only if $\omega|_W$ is nondegenerate.

Now for any nonzero $v \in V$ by nondegeneracy we can find w with $\omega(v, w) = 1$, and since ω is antisymmetric, v and w are not linearly dependent, so they span a subspace W of dimension 2 on which the restriction of ω is the standard symplectic form (in the basis v, w). But then $V = W + W^\perp$ and we are done by induction. \square

The stabilizer of $\omega \in \Lambda^2 V$ under $\text{GL}(V)$ is a subgroup called the *symplectic group* $\text{Sp}(V, \omega)$. If V has dimension $2n$ then $\text{GL}(V)$ has dimension $4n^2$ and $\Lambda^2 V$ has dimension $n(2n - 1)$. Thus $\text{Sp}(V, \omega)$ is a closed subgroup of $\text{GL}(V)$ (i.e. a Lie group) of dimension $n(2n + 1)$. For the standard symplectic form on \mathbb{R}^{2n} we write this group as $\text{Sp}(2n, \mathbb{R})$; it is the group of $2n \times 2n$ real matrices A for which $A^T J A = J$.

Note that A is symplectic if and only if A^T is. For, $A^T J A = J$ implies $A^T = J A^{-1} J^{-1}$ so

$$A J A^T = A J J A^{-1} J^{-1} = -J^{-1} = J$$

Lemma 1.4. *If we write M in $n \times n$ block form as $M = \begin{pmatrix} A & B \\ C & D \end{pmatrix}$ then M is symplectic (for the standard symplectic form on \mathbb{R}^{2n}) if and only if*

$$M^{-1} = \begin{pmatrix} D^T & -B^T \\ -C^T & A^T \end{pmatrix}$$

Proof. Write out $M^T J M = J$ as four equations in $n \times n$ matrices. \square

Lemma 1.5 (Lefschetz). *Let ω be a symplectic form on V of dimension $2n$. Then $\wedge \omega^j : \Lambda^{n-j} V \rightarrow \Lambda^{n+j} V$ is an isomorphism for all $0 \leq j \leq n$.*

Proof. Note that since $\Lambda^{n-j} V$ and $\Lambda^{n+j} V$ have the same dimension it suffices to prove that $\wedge \omega^j$ is injective. It is furthermore convenient to let $V = W^*$ for some vector space W and then observe for any $\xi \in W$ that contraction with ξ satisfies $\iota_\xi \wedge \omega^j = j \omega^{j-1} \wedge \iota_\xi(\omega)$.

By the definition of a symplectic form $\wedge \omega^n$ is an isomorphism so let's suppose by induction that $\wedge \omega^j$ is an isomorphism for some $j > 0$. Suppose $\alpha \in \Lambda^{n-j+1}$ satisfies $\alpha \wedge \omega^{j-1} = 0$. Then certainly $\alpha \wedge \omega^j = 0$ so for any $\xi \in W$

$$0 = \iota_\xi(\alpha \wedge \omega^j) = \iota_\xi(\alpha) \wedge \omega^j \pm j \alpha \wedge \omega^{j-1} \wedge \iota_\xi(\omega) = \iota_\xi(\alpha) \wedge \omega^j$$

Thus $\iota_\xi(\alpha) = 0$ for all $\xi \in W$ so that $\alpha = 0$ and the lemma is proved by induction.

This proof is apparently due to Calabi. \square

Lemma 1.6. *Let $A \in \text{Sp}(V, \omega)$. Then the characteristic polynomial of A is real, monic and palindromic (i.e. the coefficients of t^j and t^{2n-j} are equal for all j).*

Proof. The coefficients of the characteristic polynomial of a matrix A on a vector space V are (up to sign) the traces of A on the $\Lambda^j V$. But if A is symplectic, multiplication by ω^{n-j} defines an A -equivariant isomorphism from Λ^j to Λ^{2n-j} . \square

It follows that the eigenvalues of A (counted with multiplicity) are invariant under complex conjugation, and under inversion in the unit circle.

Let us now give some examples of symplectic vector spaces.

Example 1.7. If V is any finite dimensional real vector space, then there is a natural symplectic form ω on $V \oplus V^*$ that restricts to 0 on V and V^* , and for $v \in V$, $\alpha \in V^*$ it satisfies $\omega(v, \alpha) = \alpha(v) = -\omega(\alpha, v)$. The group $\text{GL}(V)$ sits naturally inside $\text{Sp}(V \oplus V^*, \omega)$ as block diagonal matrices.

Example 1.8. Let M^n be an n dimensional Riemannian manifold, and let γ be a geodesic. A *Jacobi field* along γ is a vector field V along γ satisfying $R(\gamma', V)\gamma' = V''$ where V'' means the second derivative of V along γ (with respect to the covariant derivative), and R the curvature operator. The space of Jacobi fields along V is $2n$ dimensional, since the Jacobi equation is linear of second order, so a Jacobi field is determined by its value and derivative at any point on γ .

There is a symplectic form on the space of Jacobi fields, defined by

$$\omega(U, V) := \langle U, V' \rangle_p - \langle U', V \rangle_p$$

for any $p \in \gamma$. This is evidently antisymmetric and nondegenerate (since $U(p), U'(p)$ can have any value in $T_p M$). To see it is independent of the point p , we parameterize γ by arclength t and define the quantity above to be $\omega_t(U, V)$ where $p = \gamma(t)$. Then differentiating,

$$\frac{d}{dt} \omega_t(U, V) = \langle U, V'' \rangle_{\gamma(t)} - \langle U'', V \rangle_{\gamma(t)} = \langle U, R(\gamma', V)\gamma' \rangle_{\gamma(t)} - \langle R(\gamma', U)\gamma', V \rangle_{\gamma(t)} = 0$$

by symmetries of the curvature tensor.

Example 1.9. Let M be a closed, oriented manifold of dimension $4n + 2$. The *intersection pairing* on $H^{2n+1}(M; \mathbb{R})$ is given by the formula

$$\alpha \cdot \beta := (\alpha \cup \beta) \cap [M]$$

where $[M]$ is the fundamental class in $H_{4n+2}(M; \mathbb{Z})$. This pairing is antisymmetric because $2n + 1$ is odd, and is nondegenerate by Poincaré duality. Thus $H^{2n+1}(M; \mathbb{R})$ becomes a symplectic vector space with respect to this pairing. This example is interesting even (especially?) for $n = 0$.

Example 1.10. Although it is not finite dimensional, let V be the space of smooth real-valued functions on the circle. There is a pairing ω on V defined by

$$\omega(f, g) = \int_{S^1} f dg$$

Integration by parts shows this is antisymmetric; it is also nondegenerate. The ‘meaning’ of this form is as follows. A pair of functions $f, g \in V$ define the coordinates of a smooth map from S^1 to \mathbb{R}^2 . The value of $\omega(f, g)$ is the (algebraic) area enclosed by the image of this map. In particular, the group of diffeomorphisms of S^1 , which acts on V by composition, simply reparameterizes such curves, and therefore preserves ω .

1.2. Lagrangian subspaces.

Definition 1.11. A subspace $W \subset V$ is *isotropic* if $W \subset W^\perp$, is *coisotropic* if $W^\perp \subset W$, and is *Lagrangian* if $W = W^\perp$.

Evidently Lagrangian subspaces have dimension half that of V . Note that if W is coisotropic, ω descends to a symplectic form on W/W^\perp .

Lemma 1.12. *Every isotropic subspace is contained in a Lagrangian subspace. If L_1, L_2 are any two Lagrangian subspaces then any linear isomorphism $L_1 \rightarrow L_2$ extends to a symplectic automorphism of V .*

Proof. Let $W \subset W^\perp$ be isotropic but not Lagrangian, and let $w_1 \in W^\perp - W$ be nonzero. Then $W + \langle w_1 \rangle$ is isotropic. Thus a maximal isotropic subspace is Lagrangian.

Without loss of generality we can take $V = \mathbb{R}^{2n}$ and L_2 to be the Lagrangian subspace spanned by x_1, \dots, x_n . Suppose L_1 is Lagrangian with basis v_1, \dots, v_n . Let W_1 be the isotropic subspace spanned by v_2, \dots, v_n . Since we have inclusions $W_1 \subset L_1 \subset W_1^\perp$, each inclusion of codimension 1, there is a $w_1 \in W_1^\perp$ satisfying $\omega(v_1, w_1) = 1$.

The linear map taking v_1, w_1 to x_1, y_1 respectively is symplectic on its domain. Furthermore, $\langle v_1, w_1 \rangle^\perp \cap L_1 = \langle v_2, \dots, v_n \rangle$ and $\langle x_1, y_1 \rangle^\perp \cap L_2 = \langle x_2, \dots, x_n \rangle$ so by induction on the dimension of the complements, this linear map extends to a symplectic automorphism of V taking v_i to x_i for all i . \square

Example 1.13. Let V, ω be a symplectic vector space. Then $V \oplus V$ is symplectic with respect to the form $\omega \oplus -\omega$. If $\phi : V \rightarrow V$ is linear, we can form the graph $\Gamma_\phi \subset V \oplus V$. Then Γ_ϕ is Lagrangian if and only if ϕ is symplectic.

Example 1.14. Consider \mathbb{R}^{2n} with the standard symplectic form. The subspaces X and Y spanned by x_1, \dots, x_n and by y_1, \dots, y_n are both Lagrangian, and $\mathbb{R}^{2n} = X \oplus Y$. Let ϕ be a linear map from X to Y and let Γ_ϕ be the graph of ϕ in \mathbb{R}^{2n} . We can express ϕ as an $n \times n$ matrix A_ϕ in terms of the bases x_j and y_j respectively. Then Γ_ϕ is Lagrangian if and only if A_ϕ is symmetric.

Example 1.15. Let M be a compact, oriented 3-manifold with boundary ∂M . The inclusion $\partial M \rightarrow M$ induces maps on homology. We claim the kernel L is Lagrangian.

First we show it is isotropic. Let $\alpha, \beta \in L$. Then by the definition of the kernel, there are oriented immersed surfaces F_α, F_β mapping to M with (oriented) boundaries mapping to α and β . If we make these surfaces transverse, their intersection is a 1-manifold which pairs the intersection points of $\alpha \cap \beta$ with opposite sign.

To show L is half-dimensional, let $\alpha_i \subset M$ be a basis for the image of $H_1(\partial M)$. Each α_i is homologous by a surface F_i to some $\beta_i \subset \partial M$. Let G_i be proper surfaces representing classes in $H_2(M, \partial M)$ dual to the α_i , so that $\alpha_i \cap G_j = \delta_{ij}$. Then by making F_i and G_j transverse, we see $\beta_i \cap \partial G_j = \delta_{ij}$. Since $\partial G_j \in L$ it follows that the dimension of L is at least equal to that of the image of $H_1(\partial M)$.

Let L be Lagrangian, and let G be the stabilizer of L in $\mathrm{Sp}(V, \omega)$. By Lemma 1.12, restriction to L defines a surjective homomorphism $G \rightarrow \mathrm{GL}(L)$. The kernel is isomorphic to $\mathbb{R}^{n(n+1)/2}$; to see this, we can let L be the Lagrangian subspace of \mathbb{R}^{2n} with the standard symplectic form with basis x_1, \dots, x_n . If $g \in G$ fixes each x_j , then to preserve ω it must take each y_i to $y_i + \sum S_{ij}x_j$ for some symmetric $n \times n$ matrix S ; i.e. g is a block matrix of the form $\begin{pmatrix} I_n & S \\ 0 & I_n \end{pmatrix}$ (this is just Example 1.14 again). So G has dimension $n(3n+1)/2$.

Let \mathcal{L}_n denote the space of Lagrangian subspaces of the standard symplectic \mathbb{R}^{2n} . This is a closed compact submanifold of the (compact) Stiefel manifold of all n dimensional subspaces. We have shown that $\mathrm{Sp}(2n, \mathbb{R})$ acts transitively on \mathcal{L}_n with point stabilizers of dimension $n(3n+1)/2$. Thus the dimension of \mathcal{L}_n is $n(2n+1) - n(3n+1)/2 = n(n+1)/2$.

Example 1.16. \mathcal{L}_1 is the space of lines in \mathbb{R}^2 ; thus it is naturally isomorphic to \mathbb{RP}^1 .

Example 1.17. \mathcal{L}_2 is homeomorphic to the twisted (i.e. non-orientable) S^2 bundle over S^1 . One way to see this is as follows. As in Example 1.13, consider $\mathbb{R}^4 = \mathbb{R}^2 \oplus \mathbb{R}^2$ with the symplectic form ω that agrees with the standard form on the first \mathbb{R}^2 factor and with the negative of the standard form on the second \mathbb{R}^2 factor. There are two kinds of Lagrangian subspaces: type (i) are graphs Γ_ϕ of symplectic maps $\mathbb{R}^2 \rightarrow \mathbb{R}^2$, and type (ii) are subspaces of the form $\langle v, w \rangle$ where v is in the first \mathbb{R}^2 factor and w is in the second \mathbb{R}^2 factor.

The type (i) Lagrangians are parameterized by $\mathrm{Sp}(2, \mathbb{R}) = \mathrm{SL}(2, \mathbb{R})$ which is homeomorphic to an open solid torus N . The type (ii) Lagrangians are parameterized by a product of two \mathbb{RP}^1 s (i.e. a torus T). There is a natural compactification of N to a solid torus \bar{N} . Multiplication by $-\mathrm{Id} \in \mathrm{SL}(2, \mathbb{R})$ extends to an involution ι of $\partial\bar{N}$, and the quotient may be identified with T . In other words, \mathcal{L}_2 is homeomorphic to the quotient of \bar{N} by ι on $\partial\bar{N}$, which is the twisted S^2 bundle over S^1 .

Example 1.18. Here is another way to think about \mathcal{L}_2 . If π is a plane in \mathbb{R}^4 , any two basis elements U, V determine a 2-vector $\phi = U \wedge V \in \Lambda^2\mathbb{R}^4$, unique up to scale. Conversely a 2-vector $\phi \in \Lambda^2\mathbb{R}^4$ arises (uniquely) in this way if and only if $\phi \wedge \phi = 0$. Thus, planes in \mathbb{R}^4 are parameterized by a quadric Q in $\mathbb{RP}^5 = \mathbb{P}\Lambda^2\mathbb{R}^4$.

A choice of a volume form on \mathbb{R}^4 makes $\Lambda^2\mathbb{R}^4$ into its own dual with respect to wedge product. Thus a 2-form ω determines a hyperplane H_ω in $\Lambda^2\mathbb{R}^4$ consisting of all α with $\alpha \wedge \omega = 0$. The 2-form ω is nondegenerate iff $L_\omega := H_\omega \cap Q$ is nonsingular, in which case L_ω becomes the space of planes in \mathbb{R}^4 Lagrangian with respect to ω .

Example 1.19. And here is yet another way to think about \mathcal{L}_2 . Let $H_\omega \subset \Lambda^2\mathbb{R}^4$ be the space of 2-forms α with $\alpha \wedge \omega = 0$. The restriction of \wedge to H_ω is a nondegenerate symmetric form q and one may check that the signature is $(2, 3)$.

The action of $\mathrm{Sp}(4, \mathbb{R})$ on \mathbb{R}^4 preserves volume, and therefore preserves q on H_ω . This defines a map from $\mathrm{Sp}(4, \mathbb{R})$ to $\mathrm{SO}(2, 3)$ which factors through the quotient by $\pm I$ (which acts trivially on Λ^2) and reveals the exceptional isomorphism between $\mathrm{Sp}(4, \mathbb{R})$ and $\mathrm{Spin}(2, 3)$.

The group $\mathrm{SO}(2, 3)$ acts on the hyperboloid $X := \{v : q(v, v) = 1\}$ and the stabilizer of $v \in X$ is a copy of the Lorentz group $\mathrm{SO}(1, 3)$. Thus the restriction of q to X makes it into a 4-dimensional Lorentz manifold of constant curvature -1 known as *anti de-Sitter space*. The ideal (conformal) boundary of X is the projectivization of the ‘null cone’ $q(v, v) = 0$; but this is exactly L_ω .

The Lorentz structure on X defines a cone field in the tangent space at each point, consisting of the vectors timelike with respect to $q|TX$. This limits to a cone field on L_ω . At a Lagrangian $\pi \in L_\omega$ we may identify $T_\pi L_\omega$ with Lagrangian graphs Γ_A in $\pi \oplus \pi'$ of suitable linear $A : \pi \rightarrow \pi'$ for some Lagrangian complement π' in \mathbb{R}^4 . Identifying π with $\langle x_1, x_2 \rangle$ and π' with $\langle y_1, y_2 \rangle$ in standard symplectic \mathbb{R}^4 , the linear maps A for which Γ_A is Lagrangian are symmetric 2×2 matrices. The cone field of ‘timelike’ vectors in $T_\pi L_\omega$ consist of symmetric matrices A that are (positive) definite.

1.3. Maslov cycle. Let V, ω be a symplectic vector space, and let $\pi \subset V$ be a Lagrangian subspace. Arnold [1] defines the *train* of π , denoted $\mathcal{V}(\pi)$, to be the set of Lagrangian subspaces π' that are *not* transverse to π ; i.e. those for which $\pi \cap \pi'$ has positive dimension. Thus $\pi \in \mathcal{V}(\pi)$. A symplectic automorphism of V taking π to π' evidently takes $\mathcal{V}(\pi)$ to $\mathcal{V}(\pi')$. The train $\mathcal{V}(\pi)$ is a singular hypersurface in \mathcal{L}_n .

If we let X and Y denote the standard Lagrangian subspaces of standard symplectic \mathbb{R}^{2n} spanned by the x_j and y_j respectively, Lagrangian subspaces of the train $\mathcal{V}(Y)$ are called *vertical*, and $\mathcal{V}(Y)$ is called the *Maslov cycle*, and denoted $\Sigma_n \subset \mathcal{L}_n$. The complement of the Maslov cycle consists exactly of graphs Γ_ϕ of linear $\phi : X \rightarrow Y$ represented in terms of the standard bases by symmetric matrices A_ϕ . Thus the complement of any $\mathcal{V}(\pi)$ in \mathcal{L}_n is diffeomorphic to $\mathbb{R}^{n(n+1)/2}$.

Away from the codimension 2 subset $\mathcal{V}(X) \cap \mathcal{V}(Y)$, the space $\mathcal{V}(X) - \mathcal{V}(Y)$ therefore parameterizes symmetric $n \times n$ matrices with 0 as an eigenvalue. So $\mathcal{V}(X) - \mathcal{V}(Y)$ divides the space of $n \times n$ symmetric matrices into *chambers*, corresponding to symmetric matrices A_ϕ associated to *nondegenerate* symmetric forms of signature $(p, n - p)$. Note that these separate chambers become connected in $\mathcal{V}(Y) - \mathcal{V}(X)$. Crossing $\mathcal{V}(X)$ transversely away from its singular set changes one eigenvalue from positive to negative or conversely. Thus, although $\mathcal{V}(X)$ is not typically oriented (unless $n = 1$ in which case $\mathcal{V}(X)$ is a point), it is *co-oriented* by the direction in which eigenvalues go from negative to positive. Intersection with any $\mathcal{V}(\pi)$ therefore determines an element μ of $H^1(\mathcal{L}_n; \mathbb{Z})$ called the *Maslov cocycle*.

Example 1.20. We have $\mathcal{L}_1 = \mathbb{R}\mathbb{P}^1$ and the Maslov cocycle is the (oriented) generator of $H^1 = \mathbb{Z}$. Likewise, \mathcal{L}_2 is the twisted S^2 bundle over S^1 ; again $H^1 = \mathbb{Z}$ and the Maslov cocycle is the (oriented) generator.

Example 1.21. Let’s revisit Example 1.8. Let M^n be a Riemannian manifold, and let γ be a geodesic segment. Let J be the space of normal Jacobi fields along γ ; we have seen that this is naturally a symplectic vector space of dimension $2n - 2$. For every t the normal Jacobi fields that vanish at $\gamma(t)$ form a Lagrangian subspace $\pi(t)$ of J ; thus we obtain a path π in \mathcal{L}_{n-1} .

Let $p = \gamma(0)$. A point $q = \gamma(t)$ on γ is a *conjugate point* of p along γ if there is a (nonzero) normal Jacobi field along γ that vanishes at both p and q . The *Morse index* of γ is the number of conjugate points to p along γ , counted with multiplicity. This is equal to the dimension of the space of normal variations of γ on which the Hessian of energy is negative definite.

From the discussion above, it is evident that the conjugate points $q = \gamma(t)$ are precisely where $\pi(t)$ intersects the train of $\pi(0)$, and one can check that all intersections are positively oriented, so that the signed intersection at $\pi(t)$ is equal to the multiplicity of q as a

conjugate point (i.e. the dimension of the space of normal Jacobi fields vanishing at p and q). Thus the Morse index of γ is equal to its Maslov index — the signed intersection of $\pi(t)$ rel. endpoints with $\mathcal{V}(\pi(0))$.

1.4. Lie algebra. If $G(t)$ is a one parameter family of matrices in $\mathrm{Sp}(2n, \mathbb{R})$ with $G(0) = \mathrm{Id}$ then $H := G'(0)$ satisfies $H^T J + JH = 0$ and conversely given H the family $G(t) := e^{tH}$ is in $\mathrm{Sp}(2n, \mathbb{R})$. Matrices of the form H as above make up the Lie algebra $\mathfrak{sp}(2n, \mathbb{R})$.

Since $J^T = -J$ it follows that $H \in \mathfrak{sp}(2n, \mathbb{R})$ if and only if JH is *symmetric* (this gives another proof that the dimension of $\mathrm{Sp}(2n, \mathbb{R})$ is $n(2n + 1)$). For $H \in \mathfrak{sp}(V)$ let $h(v, w) := -\omega(Hv, w)/2$ be the associated symmetric quadratic form. This function is called the *Hamiltonian*. Consider the linear ODE $\dot{v} = Hv$ for $v \in V$ or in standard coordinates,

$$\dot{x} = -\frac{\partial h}{\partial y}, \quad \dot{y} = \frac{\partial h}{\partial x}$$

Then $\dot{h} = 2h(\dot{v}, v) = \omega(\dot{v}, Hv) = \omega(\dot{v}, \dot{v}) = 0$. In other words, h is a *conserved quantity* for this ODE.

1.5. Complex structures. Let \mathbb{C}^n have a standard basis z_1, \dots, z_n . If we write $z_j := x_j + iy_j$ then $x_1, \dots, x_n, y_1, \dots, y_n$ becomes a standard basis for the underlying real vector space \mathbb{R}^{2n} . The group $\mathrm{GL}(n, \mathbb{C})$ of complex automorphisms of \mathbb{C}^n thereby embeds in $\mathrm{GL}(2n, \mathbb{R})$; its image is the group of $2n \times 2n$ real matrices A which commute with J .

The following Lemma is known colloquially as “two out of three”:

Lemma 1.22 (Two out of three). *For the standard embedding of $\mathrm{GL}(n, \mathbb{C})$ in $\mathrm{GL}(2n, \mathbb{R})$ we have*

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

Proof. For a matrix $A \in \mathrm{GL}(2n, \mathbb{R})$ we have

- (1) $A \in \mathrm{O}(2n, \mathbb{R})$ iff $A^T A = \mathrm{Id}$;
- (2) $A \in \mathrm{Sp}(2n, \mathbb{R})$ iff $A^T J A = J$;
- (3) $A \in \mathrm{GL}(n, \mathbb{C})$ iff $AJ = JA$.

Thus A is in any two of the groups iff it is in the third. We may take $\mathrm{U}(n) := \mathrm{O}(2n, \mathbb{R}) \cap \mathrm{GL}(n, \mathbb{C})$ as a definition. \square

Note that a matrix $\begin{pmatrix} A & -B \\ B & A \end{pmatrix}$ is in $\mathrm{O}(2n, \mathbb{R}) \cap \mathrm{Sp}(2n, \mathbb{R})$ if and only if $A + iB$ is unitary (in the usual sense).

Another way to express this lemma is in terms of Hermitian forms. Let V be a finite dimensional complex vector space, and let $V_{\mathbb{R}}$ denote the same vector space thought of as a real vector space (of twice the dimension). A *Hermitian form* on V is a map $h : V \times V \rightarrow \mathbb{C}$ satisfying

- (1) symmetry: $h(v, w) = \overline{h(w, v)}$ for all $v, w \in V$;
- (2) sesquilinearity: $h(\alpha v, \beta w) = \alpha \bar{\beta} h(v, w)$ for any $v, w \in V$ and $\alpha, \beta \in \mathbb{C}$; and
- (3) positivity: $h(v, v)$ is real and positive for all nonzero v .

If we choose a basis, we can represent h by an $n \times n$ complex matrix Q with $Q^* = Q$ so that $h(v, w) = v^T Q \bar{w}$ where $*$ means conjugate transpose.

Warning 1.23. Conventions for Hermitian forms differ as to which argument is complex linear and which is complex anti-linear. The opposite convention is more common with physicists.

Let h be a Hermitian form. We can separate real and imaginary parts to write $h = g - i\omega$. Then g is symmetric positive definite, and ω is antisymmetric nondegenerate. Note that g determines ω and conversely, since $g(iv, w) = \omega(v, w)$ (this also shows that ω is nondegenerate). The group $U(h)$ is the subgroup of $GL(V, \mathbb{C})$ preserving h ; by the discussion above it is isomorphic to $GL(V, \mathbb{C}) \cap O(V_{\mathbb{R}}, g) = GL(V, \mathbb{C}) \cap Sp(V_{\mathbb{R}}, \omega)$.

Let h be a Hermitian form on an n -dimensional complex vector space V , and let $\omega = \Im(h)$ be the associated symplectic form on $V_{\mathbb{R}}$. Then an n -dimensional real subspace $L \subset V_{\mathbb{R}}$ is Lagrangian if and only if $h|_L$ is totally real; i.e. $h(v, w) \in \mathbb{R}$ for all $v, w \in L$.

Lemma 1.24. *The group $U(n)$ acts transitively on \mathcal{L}_n , with stabilizers isomorphic to $O(n)$. Thus $\mathcal{L}_n = U(n)/O(n)$.*

Proof. Let L_1, L_2 be any two real n -dimensional subspaces of V . Then any (real) linear isometry $\phi : L_1 \rightarrow L_2$ extends to a complex linear isometry $\phi_{\mathbb{C}} : L_1 \oplus iL_1 \rightarrow L_2 \oplus iL_2$. But if $h|_{L_j}$ is totally real, then $L_j \oplus iL_j$ is (complex-linearly) isomorphic and isometric to V for $j = 1, 2$. Thus $\phi_{\mathbb{C}} \in U(V)$. \square

The homotopy groups of $U(n)/O(n)$ stabilize for large enough n , and Bott periodicity says that in the limit they are periodic with period 8:

$$\pi_k(U(n)/O(n)) = 0, \mathbb{Z}, \mathbb{Z}/2\mathbb{Z}, \mathbb{Z}/2\mathbb{Z}, 0, \mathbb{Z}, 0, 0 \text{ depending on } k \bmod 8 \text{ if } k \ll n$$

By examining the homotopy long exact sequence in low dimensions one sees in fact that $H^1(U(n)/O(n)) = \pi_1(U(n)/O(n)) = \mathbb{Z}$ for all $n \geq 1$.

Lemma 1.25. *With respect to this identification, the map $\det^2 : U(n)/O(n) \rightarrow S^1$ pulls back the oriented generator of $H^1(S^1)$ to the Maslov class $\mu \in H^1(\mathcal{L}_n)$.*

Proof. Identify the standard symplectic space with \mathbb{C}^n . Let $\varphi : [0, 1] \rightarrow U(n)$ take t to the matrix with $e^{i\pi t}$ on the diagonal. Let γ be the loop of Lagrangians which is the orbit of $\mathbb{R}^n \subset \mathbb{C}^n$ under $\varphi([0, 1])$. For small t the Lagrangian $\varphi(t)\mathbb{R}^n$ is (in standard real coordinates) the graph of the diagonal linear matrix $A_t : \mathbb{R}^n \rightarrow \mathbb{R}^n$ with diagonal entries $\tan(t)$. The loop γ intersects the train $\mathcal{V}(\mathbb{R}^n)$ only at \mathbb{R}^n though not transversely. If we perturb γ to be transverse to the train, then we see that n eigenvalues go from negative to positive near $t = 0$ so that $\mu(\gamma) = n$. Since $\det^2 \varphi(t) = e^{2\pi i n t}$ we see that the two classes agree and are nonzero on γ . Since $H^1(U(n)/O(n)) = \mathbb{Z}$ the two classes are equal. \square

Every $A \in GL(n, \mathbb{R})$ has a unique *polar decomposition* $A = PU$ where $U \in O(n, \mathbb{R})$ and P is symmetric and positive definite. To see this, observe that AA^T is symmetric and positive definite; therefore it makes sense to write

$$A = (AA^T)^{1/2}(AA^T)^{-1/2}A$$

and observe that $P := (AA^T)^{1/2}$ is symmetric positive definite, and $U := (AA^T)^{-1/2}A$ is orthogonal.

Lemma 1.26. *Let $A \in \mathrm{Sp}(2n, \mathbb{R})$ and let $A = PU$ be its polar decomposition. Then P and U separately are in $\mathrm{Sp}(2n, \mathbb{R})$. Furthermore, for every positive real $t > 0$ the matrix P^t is in $\mathrm{Sp}(2n, \mathbb{R})$. Thus there is a deformation retraction from $\mathrm{Sp}(2n, \mathbb{R})$ to $\mathrm{U}(n)$ sending each PU to the 1-parameter family P^tU as t goes from 1 to 0.*

Proof. Since A is symplectic, so is A^T and thus also AA^T . Since AA^T is symmetric and positive definite it makes sense to write $(AA^T)^t$ for any real positive t and by considering the action of a symmetric positive definite matrix on its eigenspaces, one sees that $(AA^T)^t$ is symplectic for all t . Thus $P := (AA^T)^t$ is symplectic, and so (therefore) is U . Since U is symplectic and orthogonal, it is unitary, and P^tU is a path in $\mathrm{Sp}(2n, \mathbb{R})$ from A to U . Since every unitary matrix is already symplectic, the deformation retraction surjects onto $\mathrm{U}(n)$. \square

It is convenient to spell out the meaning of the two out of three lemma in the abstract. This is captured by the following two definitions.

Definition 1.27. Suppose V, ω is a symplectic vector space. A *compatible* inner product is a positive definite symmetric inner product $g : V \times V \rightarrow \mathbb{R}$ so that if $J \in \mathrm{GL}(V)$ is the unique automorphism with $\omega(X, Y) = g(JX, Y)$, then $J^2 = -\mathrm{Id}$.

Note that J as above is always g -skew-adjoint.

Definition 1.28. Suppose V, ω is a symplectic vector space. A *compatible* complex structure is a linear map $J : V \rightarrow V$ with $J^2 = -\mathrm{Id}$ such that $\omega(Jv, Jw) = \omega(v, w)$ for all v, w , and $\omega(v, Jv) > 0$ for all $v \neq 0$.

A compatible complex structure defines a Hermitian form h with real part $\omega(\cdot, J\cdot)$ and imaginary part $-\omega(\cdot, \cdot)$. Conversely, for any complex vector space, the imaginary part of any Hermitian form is a symplectic form compatible with multiplication by i . Thus: if we choose an isomorphism of V, ω with the standard symplectic \mathbb{R}^{2n} , the space of compatible complex structures is isomorphic to the coset space $\mathrm{Sp}(2n, \mathbb{R})/\mathrm{U}(n)$.

1.6. The Siegel Upper Half-Space. It follows from Lemma 1.26 that $\mathrm{U}(n)$ is a maximal compact subgroup of $\mathrm{Sp}(2n, \mathbb{R})$, and that the coset space $\mathrm{Sp}(2n, \mathbb{R})/\mathrm{U}(n)$ admits a homogeneous complete $\mathrm{Sp}(2n, \mathbb{R})$ -invariant metric of non-positive curvature, unique up to scale.

The *Siegel upper half-space* \mathfrak{H}^n is a model for this homogeneous space. It is the space of $n \times n$ complex symmetric matrices $Z = X + iY$ for which Y is positive definite. This is an open subset of $\mathbb{C}^{n(n+1)/2}$.

Theorem 1.29 (Siegel). *The symplectic group $\mathrm{Sp}(2n, \mathbb{R})$ acts on \mathfrak{H}^n as follows. If we write a symplectic matrix M in $n \times n$ block form as $M = \begin{pmatrix} A & B \\ C & D \end{pmatrix}$ then*

$$M \cdot Z = (AZ + B)(CZ + D)^{-1}$$

This action takes \mathfrak{H}^n to \mathfrak{H}^n , is transitive, and the stabilizer of the matrix $i \cdot I_n$ is precisely $\mathrm{U}(n)$.

Proof. This proof is due to Siegel [16]. A complex $n \times n$ matrix Z is symmetric and has positive definite imaginary part if and only if

$$(Z^T \ I_n) J \begin{pmatrix} Z \\ I_n \end{pmatrix} = 0 \text{ and } \frac{1}{2i} (Z^* \ I_n) J \begin{pmatrix} Z \\ I_n \end{pmatrix} > 0$$

If $M \cdot \begin{pmatrix} Z \\ I \end{pmatrix} = \begin{pmatrix} E \\ F \end{pmatrix}$ then since $M^T J M = J$ and $M^* = M^T$, we have

$$F^T E - E^T F = (E^T \ F^T) J \begin{pmatrix} E \\ F \end{pmatrix} = 0 \text{ and } \frac{1}{2i} (E^* \ F^*) J \begin{pmatrix} E \\ F \end{pmatrix} = -\frac{1}{2i} (E^* F - F^* E) > 0$$

The matrix $F := CZ + D$ is invertible, since if $Fv = 0$ then $v^* F^* = 0$ and therefore $v^* (E^* F - F^* E) v = 0$. But $-(1/2i)(E^* F - F^* E) > 0$ so $v = 0$.

Since $E^T F = F^T E$ we have

$$(EF^{-1})^T = F^{-T} E^T = F^{-T} E^T F F^{-1} = F^{-T} F^T E F^{-1} = E F^{-1}$$

so EF^{-1} is symmetric. Likewise, from $-(1/2i)(E^* F - F^* E) > 0$ we have

$$-\frac{1}{2i} F^* ((EF^{-1})^* - EF^{-1}) F > 0 \text{ and therefore } -\frac{1}{2i} ((EF^{-1})^* - EF^{-1}) > 0$$

so that EF^{-1} has positive imaginary part. This concludes the proof that the action is well-defined, and takes \mathfrak{H}^n to itself.

We now show the action is transitive. If $Z \in \mathfrak{H}^n$ is arbitrary, we can write $Z = X + iY$ where X and Y are both real and symmetric, and Y is positive. The matrix

$$M := \begin{pmatrix} Y^{1/2} & XY^{-1/2} \\ 0 & Y^{-1/2} \end{pmatrix}$$

is symplectic and takes $i \cdot I_n$ to Z .

Finally we show that the stabilizer of $i \cdot I_n$ is $U(n)$. A symplectic matrix $M = \begin{pmatrix} A & B \\ C & D \end{pmatrix}$ stabilizes $i \cdot I_n$ if and only if $D = A$ and $C = -B$. But this is equivalent to $MJ = JM$. \square

We shall discuss the geometry of \mathfrak{H}^n further in § 2.2.

2. SYMPLECTIC MANIFOLDS

2.1. Symplectic Manifolds.

Definition 2.1 (Symplectic manifolds). A *symplectic manifold* is a smooth $2n$ -dimensional manifold M with a smooth 2-form ω satisfying

- (1) nondegeneracy: ω^n is nowhere vanishing; and
- (2) integrability: $d\omega = 0$.

A diffeomorphism between symplectic manifolds is a *symplectomorphism* if it pulls back the symplectic form on the range to the symplectic form on the domain.

A symplectic manifold inherits a natural orientation, for which ω^n is positive. Since $d\omega = 0$ there is a well-defined class $[\omega] \in H^2(M)$, and if M is closed, then $\int_M \omega^n = [\omega]^n([M]) > 0$. Thus for M a closed symplectic manifold, the cohomology classes $[\omega]^k$ are nontrivial for $k = 0, \dots, n$.

Example 2.2. \mathbb{R}^{2n} with the standard basis is a symplectic manifold with respect to the linear form $\omega := \sum dx_j \wedge dy_j$. Any open subset of \mathbb{R}^{2n} becomes a symplectic manifold by restriction. Thus any finite CW complex is homotopy equivalent to a (noncompact) symplectic manifold.

Example 2.3. Open submanifolds of symplectic manifolds are symplectic. Products of symplectic manifolds are symplectic.

Example 2.4 (Immersion). Smale's *immersion theorem* says (amongst other things) that if M^{2n} is a *noncompact* smooth $2n$ -manifold with trivial tangent bundle TM , then for any trivialization of TM we may find an immersion $\phi : M \rightarrow \mathbb{R}^{2n}$ pulling back the standard trivialization of $T\mathbb{R}^{2n}$ to a trivialization of TM homotopic to the given one. If ω denotes the standard symplectic form on \mathbb{R}^{2n} then $\phi^*\omega$ becomes a symplectic form on M .

More generally, if M is noncompact, N is symplectic of the same dimension as M , and there is some homotopy class of map $f : M \rightarrow N$ for which f^*TN is isomorphic to TM , then there is an immersion $M \rightarrow N$ homotopic to f , which pulls back the symplectic form on N to a symplectic form on M .

Example 2.5 (Surfaces). On a surface a symplectic form is the same as a (signed) area form. Thus every orientable surface admits a symplectic structure.

Example 2.6 ($M \times S^1$). Let M be an oriented 3-manifold fibering over S^1 . In other words, there is a smooth submersion $M \rightarrow S^1$ whose fibers F are smooth, closed, oriented surfaces. The angle form on S^1 pulls back to a closed nondegenerate 1-form α on M that vanishes on each TF and as is well-known we may also find a closed 2-form ω_M that restricts to an area form on each TF . Evidently $\omega_M \wedge \alpha$ is a volume form on M .

On $M \times S^1$ define $\omega := \omega_M + \alpha \wedge d\theta$ where $d\theta$ is the angle form on the S^1 factor. Evidently ω is closed, and ω^2 is a volume form on $M \times S^1$. Thus $M \times S^1$ is symplectic. Note that one can also think of $M \times S^1$ as an F bundle over T^2 . A similar construction shows that for any smooth F bundle $F \rightarrow E \rightarrow B$ where B is a symplectic manifold and $[F]$ is nontrivial in $H_2(E; \mathbb{R})$, the manifold E may also be given a symplectic structure. This construction is due to Thurston.

It is a deep theorem [6] depending on the work of many authors that if a closed 4-manifold of the form $M \times S^1$ is symplectic, then M fibers over S^1 .

Example 2.7 (Cotangent Bundles). Let M be any smooth manifold and let T^*M denote its cotangent bundle with projection map $\pi : T^*M \rightarrow M$. There is a *tautological* 1-form λ on T^*M defined as follows. A point in T^*M is a pair (p, q) where $p \in M$ and $q \in T_p^*M$ is a linear form on T_pM . Let $T_{p,q}$ denote the tangent space to T^*M at (p, q) . The map $d\pi$ takes $T_{p,q}$ surjectively to T_pM , and the kernel is $T_qT_p^*M$ which is canonically isomorphic to T_p^*M ; i.e. there is a short exact sequence of vector spaces $T_p^*M \rightarrow T_{p,q} \rightarrow T_pM$. The linear form q on T_pM pulls back under $d\pi$ to a linear form $d\pi^*q$ on $T_{p,q}$. We define λ to be the 1-form whose value at (p, q) is $d\pi^*q$.

If x_j are local coordinates on M , we define y_j to be local coordinates on T^* dual to the vector fields ∂x_j (i.e. y_j is the coefficient of dx_j in some 1-form locally). Thus in local coordinates, $\lambda = \sum y_j dx_j$. The form $\omega := -d\lambda = \sum dx_j \wedge dy_j$ is evidently closed, and locally isomorphic to the standard symplectic form on \mathbb{R}^{2n} (and therefore nondegenerate). We call this the *tautological* symplectic structure on T^*M .

Note that since the definition of ω does not depend on the choice of coordinates, any diffeomorphism between smooth manifolds induces a symplectomorphism between the symplectic structures on their cotangent spaces.

Example 2.8 (Kähler Manifolds). An *almost Hermitian manifold* is a smooth manifold M of dimension $2n$ together with a Riemannian metric g and an endomorphism $J : TM \rightarrow TM$ squaring to $-\text{Id}$ (i.e. an almost complex structure) and preserving g . Associated to g we can define a smoothly varying field of Hermitian forms h on TM by $h(v, w) = g(v, w) + ig(v, Jw)$ pointwise, and observe that $h = g - i\omega$ for a non-degenerate 2-form ω .

An almost Hermitian manifold is

- (1) *almost Kähler* if $d\omega = 0$;
- (2) *Hermitian* if J is integrable (i.e. if there is a holomorphic atlas on M for which J is multiplication by i);
- (3) *Kähler* if J or equivalently ω is parallel (i.e. if $\nabla J = 0$ equivalently $\nabla\omega = 0$ for ∇ the Levi-Civita connection associated to g).

The two conditions $\nabla J = 0$ and $\nabla\omega = 0$ are equivalent by Lemma 1.22. Furthermore, Kähler is equivalent to Hermitian + almost Kähler.

Any symplectic manifold M, ω admits an almost Kähler structure. By Lemma 1.22 this is equivalent to the choice of a *tame* almost complex structure J preserving ω , which always exists because the coset space Sp/U is contractible (more on this in § 4.1).

Example 2.9 (Complex Projective Varieties). Complex space \mathbb{C}^{n+1} admits a Hermitian metric $h = \sum dz_j d\bar{z}_j$ with respect to the usual coordinates $z := z_0, z_1, \dots, z_n$. The restriction of this form to the unit sphere $S^{2n+1} \subset \mathbb{C}^{n+1}$ is invariant under the diagonal action of $\text{U}(1)$ so the restriction to the normal bundle of the $\text{U}(1)$ orbits descends to a Hermitian form on $\mathbb{C}\mathbb{P}^n := \text{U}(1) \backslash \mathbb{C}^{n+1}$ called the *Fubini–Study form*. Thus $\mathbb{C}\mathbb{P}^n$ is Kähler.

The Kähler property is inherited by complex submanifolds, so every smooth complex projective variety is Kähler with respect to the restriction of the Fubini–Study form.

Example 2.10 (Lefschetz Manifold). A symplectic manifold of dimension $2n$ is a *Lefschetz manifold* if it satisfies the strong Lefschetz property: i.e. that $\cup[\omega^j] : H^{n-j}(M; \mathbb{R}) \rightarrow H^{n+j}(M; \mathbb{R})$ is an isomorphism for all $0 \leq j \leq n$. Note that for any symplectic manifold $\wedge\omega^j : \Omega^{n-j}(M) \rightarrow \Omega^{n+j}(M)$ is always an isomorphism at the level of *forms*. This is actually true for any 2-form ω with ω^n nowhere zero, closed or not, by applying Lemma 1.5 fiberwise.

Every closed Kähler manifold is a Lefschetz manifold (this is the so-called *Hard Lefschetz theorem*). This follows two facts, both miraculous:

- (1) the symplectic form ω is harmonic for the Kähler metric; and
- (2) for any harmonic form α the wedge product $\alpha \wedge \omega$ is also harmonic.

Since a harmonic form is nontrivial in cohomology unless it is identically zero, and since $\wedge\omega^j$ is an isomorphism at the level of forms, it follows that $\cup[\omega^j]$ is injective on $H^{n-j}(M; \mathbb{R})$ and therefore an isomorphism. Note that the wedge product of two harmonic forms is rarely harmonic (even on a Kähler manifold).

If M is a Lefschetz manifold then for j odd the dimension of $H^j(M; \mathbb{R})$ is necessarily even. Let's see why. By Poincaré duality it suffices to prove this for $j \leq n$. There is a

pairing $H^j \times H^j \rightarrow \mathbb{R}$ defined at the level of forms by

$$([\alpha], [\beta]) \rightarrow \int_M \alpha \wedge \omega^{n-j} \wedge \beta$$

This pairing is just the composition of $[\alpha] \rightarrow [\alpha] \cup [\omega^{n-j}]$ with the ordinary nondegenerate (by Poincaré duality) cup product pairing $H^{2n-j} \times H^j \rightarrow H^{2n} \cong \mathbb{R}$. Thus on a Lefschetz manifold the pairing $H^j \times H^j \rightarrow \mathbb{R}$ must also be nondegenerate. But for j odd this pairing is antisymmetric, so $\dim H^j$ must be even.

We have already seen examples of symplectic manifolds that are not Lefschetz manifolds. Let $W = M \times S^1$ where M is a 3-manifold fibering over the circle. This is symplectic by Example 2.6. Then $\dim H^1(W) = \dim H^1(M) + 1$ by the Künneth formula. But there are many 3-manifolds that fiber over the circle with $\dim H^1$ even; for example, the T^2 bundle whose monodromy acts on $H_1(T^2)$ by the matrix $\begin{pmatrix} 1 & \\ 0 & 1 \end{pmatrix}$ has $\dim H^1 = 2$.

Example 2.11 (Coadjoint Orbits). Let G be a connected Lie group, and let \mathfrak{g} be its Lie algebra. The group G acts on itself by inner automorphisms, and the induced action on $\mathfrak{g} = T_e G$ is the adjoint action $\text{Ad} : G \rightarrow \text{GL}(\mathfrak{g})$ given by

$$\text{Ad}_g X = \left. \frac{d}{dt} \right|_{t=0} gh(t)g^{-1}$$

where $h(t)$ is a smooth path in G with $h(0) = e$ and $h'(0) = X$. Dualizing gives the *coadjoint* action $\text{Ad}^* : G \rightarrow \text{GL}(\mathfrak{g}^*)$ defined by $\text{Ad}_g^* \xi(X) = \xi(\text{Ad}_{g^{-1}} X)$. Note that $\text{Ad}_g^* = (\text{Ad}_{g^{-1}})^*$.

The derivatives of Ad and Ad^* at e are $\text{ad} : \mathfrak{g} \rightarrow \text{End}(\mathfrak{g})$ and $\text{ad}^* : \mathfrak{g} \rightarrow \text{End}(\mathfrak{g}^*)$ given by formulae

$$\text{ad}_X Y = [X, Y] \quad (\text{ad}_X^* \xi)(Y) = -\xi([X, Y])$$

for $X, Y \in \mathfrak{g}$ and $\xi \in \mathfrak{g}^*$.

Let O be an orbit of $\text{Ad}^*(G)$ in \mathfrak{g}^* , and let $\xi \in O \subset \mathfrak{g}^*$. The map $X \rightarrow \text{ad}_X^* \xi$ takes \mathfrak{g} surjectively onto the tangent space $T_\xi O$ with kernel \mathfrak{g}_ξ equal to the $X \in \mathfrak{g}$ with $\text{ad}_X^* \xi = 0$. Thus $T_\xi O$ inherits a canonical nondegenerate 2-form defined by the pairing $X, Y \rightarrow \xi([X, Y])$.

For $X \in \mathfrak{g}/\mathfrak{g}_\xi$ let $X^\# := \text{ad}_X^* \xi \in T_\xi O$. Let ω be the 2-form on O defined by $\omega(X^\#, Y^\#) = \xi([X, Y])$ for $X, Y \in \mathfrak{g}/\mathfrak{g}_\xi$. Then for any X, Y, Z ,

$$d\omega(X^\#, Y^\#, Z^\#) = X^\#(\omega(Y^\#, Z^\#)) - \omega([X^\#, Y^\#], Z^\#) + \text{cyclic permutations}$$

For any X, Y, Z one can compute

$$X^\# \omega(Y^\#, Z^\#) = \left. \frac{d}{dt} \right|_{t=0} (\text{Ad}_{e^{tX}}^* \xi)([Y, Z]) = \left. \frac{d}{dt} \right|_{t=0} \xi(\text{Ad}_{e^{-tX}} [Y, Z]) = -\xi([X, [Y, Z]])$$

Likewise

$$\omega([X^\#, Y^\#], Z^\#) = \omega([X, Y]^\#, Z^\#) = \xi([[X, Y], Z])$$

So $d\omega = 0$ by the Jacobi identity, and one obtains:

Theorem 2.12 (Kirillov–Kostant). *The coadjoint orbits of a connected Lie group are symplectic manifolds with respect to ω .*

Example 2.13. For the unitary group U the adjoint representation is isomorphic to the coadjoint one; thus the adjoint orbits have a symplectic structure. The Lie algebra $\mathfrak{u}(n+1)$ consists of skew-Hermitian matrices (i.e. $A^* = -A$) acting on \mathbb{C}^{n+1} . The adjoint orbit of a rank 1 skew-Hermitian matrix is isomorphic to $\mathbb{C}\mathbb{P}^n$.

Example 2.14 (Character Varieties; Goldman [7]). Let S be a closed oriented surface of genus $g > 1$ and let $\pi = \pi_1(S)$. Let G be a semisimple Lie group (i.e. \mathfrak{g} is a direct sum of nonabelian Lie algebras without any non-zero proper ideals). Let $\text{Hom}(\pi, G)$ denote the space of representations of π to G , and let $\text{Hom}^*(\pi, G) \subset \text{Hom}(\pi, G)$ denote the open subset of irreducible representations; i.e. representations whose image has centralizer of minimum possible dimension (equal to the dimension of $Z(G)$). This is a smooth manifold of dimension $(2g - 1) \dim(G)$.

The group G mod its center acts locally freely on $\text{Hom}^*(\pi, G)$ by conjugation, and we denote the quotient by $X^*(\pi, G)$. Since G is semisimple, the center is discrete, so this is a smooth manifold of dimension $(2g - 2) \dim(G)$. An (irreducible) representation $\rho : \pi \rightarrow G$ makes \mathfrak{g} into a module over π with respect to $\text{ad} \circ \rho$, and the tangent space to $X^*(\pi, G)$ at the conjugacy class of ρ is naturally isomorphic to $H_\rho^1(\pi, \mathfrak{g})$ where the subscript ρ denotes the module structure.

For any Lie group there is a symmetric ad-invariant bilinear form β on \mathfrak{g} called the *Killing form*, defined by $\beta(U, V) := \text{tr}(\text{ad}_U \text{ad}_V)$. This is nondegenerate precisely when G is semisimple. Using β one defines a pairing

$$H_\rho^1(\pi, \mathfrak{g}) \times H_\rho^1(\pi, \mathfrak{g}) \xrightarrow{\cup} H_\rho^2(\pi, \mathfrak{g} \otimes \mathfrak{g}) \xrightarrow{\beta} H^2(\pi, \mathbb{R}) \xrightarrow{\int} \mathbb{R}$$

where the last map identifies $H^2(\pi, \mathbb{R}) = H^2(S, \mathbb{R})$ and integrates over a fundamental class $[S]$. It turns out that this pairing is nondegenerate, and the associated 2-form it defines on $X^*(\pi, G)$ is closed.

We shall see some more examples of symplectic manifolds in the sequel.

2.2. The geometry of \mathfrak{H}^n . The Siegel upper-half space is Kähler (and therefore symplectic). We briefly indicate how \mathfrak{H}^n and its geometry fits into the general framework of symmetric spaces, by recalling several facts without proof; see e.g. Helgason [9] or Cheeger-Ebin [4] Ch. 3 for details.

2.2.1. The geometry of G/K . Any noncompact semisimple Lie group G has a so-called Cartan decomposition $G = KAK$ where K is a maximal compact subgroup, and A is the identity component of a maximal split torus (i.e. a maximal abelian subgroup diagonalizable over \mathbb{R}). The decomposition of an element $g = k_1 a k_2$ is unique up to conjugating a by an element of the Weyl group $W(G, A)$, which is the normalizer of A in K . A coset representative of a under W may be found in A^+ , the (closed) positive Weyl chamber in A .

Let \mathfrak{g} be the Lie algebra of G and let \mathfrak{k} be the Lie subalgebra of K . Recall that the Killing form β on \mathfrak{g} is the symmetric bilinear form $\beta(U, V) = \text{trace}(\text{ad}_U \text{ad}_V)$. Cartan showed there is an involution $d\theta$ on G for which there is a splitting $\mathfrak{g} = \mathfrak{p} \oplus \mathfrak{k}$ where \mathfrak{k} is a -1 eigenspace for $d\theta$ and \mathfrak{p} is a $+1$ eigenspace. The Killing form β is negative definite on \mathfrak{k} and positive definite on \mathfrak{p} , and furthermore $[\mathfrak{k}, \mathfrak{k}] \subset \mathfrak{k}$, $[\mathfrak{k}, \mathfrak{p}] \subset \mathfrak{p}$, $[\mathfrak{p}, \mathfrak{p}] \subset \mathfrak{k}$. The inner product $\beta|_{\mathfrak{p}}$ defines,

by quotient and left translation, a G -invariant metric on G/K making it into a symmetric space (i.e. for every $p \in G/K$ there is an isometry of G/K acting as -1 on T_p).

With respect to this metric the formula for the sectional curvature is as follows. For vectors u, v at a point p in G/K there are unique left-invariant vector fields U, V on G in \mathfrak{p} which project to u, v at the coset p respectively (here we implicitly use the identification of \mathfrak{g} with the left-invariant vector fields on G). Then

$$\langle R(u, v)v, u \rangle = \beta([U, V], [U, V])$$

Note that since $U, V \in \mathfrak{p}$ we have $[U, V] \in \mathfrak{k}$ where β is negative definite. Thus $K(u, v) \leq 0$ with equality if and only if U, V commute. It follows that for any $k \in K$ the orbit $A \cdot k$ in G/K is a *maximal flat*; i.e. a maximal totally geodesic subspace isometric to Euclidean space.

2.2.2. The Riemannian metric on \mathfrak{H}^n . For $G = \mathrm{Sp}(2n, \mathbb{R})$ the maximal compact subgroup K is $\mathrm{U}(n)$. The Cartan decomposition gives $\mathrm{Sp}(2n, \mathbb{R}) = KAK$ where $K = \mathrm{U}(n)$ and A is the group of matrices of the form $\begin{pmatrix} D & 0 \\ 0 & D^{-1} \end{pmatrix}$ where D is positive-definite and diagonal.

The Cartan involution $\theta : g \rightarrow (g^T)^{-1}$ fixes $\mathrm{U}(n)$ and its Lie algebra \mathfrak{k} , and $d\theta$ acts as -1 on \mathfrak{p} , the space of real $2n \times 2n$ matrices of the form $P := \begin{pmatrix} Y & X \\ X & -Y \end{pmatrix}$ where X and Y are themselves both symmetric. There is an isomorphism $\phi : \mathfrak{p} \rightarrow T$ where T is the tangent space to \mathfrak{H}^n at $i \cdot I_n$ (equivalently, T is the space of symmetric complex $n \times n$ matrices). This isomorphism takes P as above to $2(X + iY)$.

The Killing form is $\beta(U, V) = 2 \mathrm{tr}(UV)$ (up to a multiplicative factor depending on n); thus the norm squared of a vector is twice the sum of the squares of its eigenvalues (this normalization agrees with Siegel and simplifies formulae in what follows).

Let's compute the Riemannian metric on the tangent space T to \mathfrak{H}^n at $i \cdot I_n$. Let $Z_j := X_j + iY_j$ for $j = 1, 2$. Then identifying T with \mathfrak{p} as above, and using the formula for β , we obtain

$$\langle Z_1, Z_2 \rangle = \frac{1}{2} \mathrm{tr} \left(\begin{pmatrix} Y_1 & X_1 \\ X_1 & -Y_1 \end{pmatrix} \begin{pmatrix} Y_2 & X_2 \\ X_2 & -Y_2 \end{pmatrix} \right) = \mathrm{tr}(X_1 X_2 + Y_1 Y_2) = \mathrm{Re}(\mathrm{tr}(Z_1 \bar{Z}_2))$$

To compute the metric on the tangent space at an arbitrary $Z \in \mathfrak{H}^n$, define

$$M := \begin{pmatrix} Y^{1/2} & XY^{-1/2} \\ 0 & Y^{-1/2} \end{pmatrix}$$

then

$$M \cdot (iI + tY^{-1}U) = X + iY + tY^{-1/2}UY^{1/2} + O(t^2)$$

so $dM(Y^{-1}U) = \mathrm{Ad}(Y^{-1/2})U$. Thus we obtain Siegel's formula for the Riemannian metric: at an arbitrary $Z \in \mathfrak{H}^n$ the metric g is equal to the real part of $\mathrm{tr}(Y^{-1}dZY^{-1}d\bar{Z})$.

2.2.3. Kähler structure. The tangent space T at iI has a natural complex structure, and $\mathrm{U}(n)$ acts complex linearly. In fact, \mathfrak{H}^n is an open subset of the space of symmetric $n \times n$ complex matrices, and inherits the structure of a complex manifold. Evidently the action of $\mathrm{Sp}(2n, \mathbb{R})$ preserves this structure.

For $t \in \mathbb{R}$ define

$$R(t) := \begin{pmatrix} \cos(t)I & -\sin(t)I \\ \sin(t)I & \cos(t)I \end{pmatrix}$$

Then $R(t) \in U(n)$ and acts on T as multiplication by e^{2ti} (thus the matrix $j := 2^{-1/2} \begin{pmatrix} I & -I \\ I & I \end{pmatrix}$ corresponds to multiplication by i).

For $M = \begin{pmatrix} E & 0 \\ 0 & -E \end{pmatrix}$ with E symmetric, $\beta(M, jMj^{-1}) = 0$. It follows that the totally real and imaginary subspaces of T are orthogonal, and the Riemannian metric g on \mathfrak{H}^n is equal to the real part of an invariant Hermitian form $h := \text{tr}(Y^{-1}dZY^{-1}d\bar{Z})$.

The imaginary part of h is a non-degenerate 2-form $-\omega$. Since h is parallel, so is ω ; thus \mathfrak{H}^n is Kähler.

2.3. Moser and Darboux Theorems. Let $X(t)$ be a (possibly) time dependent vector field on a compact manifold M , generating a 1-parameter isotopy $\phi(t)$. If M is symplectic with symplectic form ω , then $\omega(t) := \phi(t)^*\omega$ are also symplectic. Furthermore, since each $\phi(t)$ is isotopic to the identity, the cohomology classes $[\omega(t)]$ are constant and equal to $[\omega]$. Moser's theorem is a sort of converse to this fact:

Theorem 2.15 (Moser). *Let $\omega(t)$ be a smooth 1-parameter family of symplectic forms on a closed manifold M , and suppose that the cohomology classes $[\omega(t)]$ are constant. Then the $\omega(t)$ are all pulled back from some fixed symplectic form by a 1-parameter isotopy.*

Proof. Fix a smooth metric on M for convenience. By hypothesis the forms $\dot{\omega}(t)$ are exact, and by Hodge for each t there is a unique 1-form $\alpha(t)$ that minimizes the L^2 norm amongst all forms satisfying $d\alpha(t) = \dot{\omega}(t)$. Again by Hodge, $\alpha(t)$ is smooth and depends smoothly on t . Since $\omega(t)$ is nondegenerate for each t , there is a unique vector field $X(t)$ so that $\alpha(t)(\cdot) = \omega(t)(X(t), \cdot)$; in other words $\iota_{X(t)}\omega(t) = \alpha(t)$. But by Cartan's magic formula,

$$\mathcal{L}_{X(t)}\omega(t) = \iota_{X(t)}d\omega(t) + d\iota_{X(t)}\omega(t) = d\alpha(t) = \dot{\omega}(t)$$

In other words, the $\omega(t)$ are all pulled back from $\omega(0)$ by the flow $\phi(t)$ generated by $X(t)$, as claimed. \square

As a Corollary of Moser's theorem, we may prove a theorem of Darboux to the effect that symplectic manifolds are all locally standard:

Theorem 2.16 (Darboux). *Let M^{2n}, ω be a symplectic manifold. Then every point has a neighborhood symplectomorphic to an open subset of \mathbb{R}^{2n} with the standard symplectic form.*

Proof. Let $p \in M$ be arbitrary, and let U be an open neighborhood of p for which there is some local chart $\phi : U \rightarrow \mathbb{R}^{2n}$ taking p to 0 (say) and for which $d\phi$ pulls back the standard form on $T_0\mathbb{R}^{2n}$ to $\omega|_{T_pM}$. Let ω' be the pullback of the standard symplectic form on \mathbb{R}^{2n} under $d\phi$. Thus ω and ω' are both symplectic forms on U which agree at p . Since nondegeneracy is an open condition, there is a smaller open neighborhood V of U for which $\omega(t) := t\omega + (1-t)\omega'$ is nondegenerate (and hence symplectic) for $t \in [0, 1]$.

Using ϕ we may identify V with its image in \mathbb{R}^{2n} and think of $\omega(t)$ as a 1-parameter family of symplectic forms on a neighborhood of 0. Since \mathbb{R}^{2n} is contractible, we can write $\dot{\omega}(t) = d\alpha(t)$ for some smooth family of 1-forms $\alpha(t)$, and by subtracting the constant value at the origin from $\alpha(t)$ we may assume $\alpha(t)$ vanishes at the origin. Let $X(t)$ be the vector fields defined on a neighborhood of 0 by $\iota_{X(t)}\omega(t) = \alpha(t)$. They generate a partially defined flow that vanishes at 0, and therefore $\omega(0)$ and $\omega(1)$ are symplectomorphic on a sufficiently small neighborhood of 0. \square

Actually the same proof shows more:

Theorem 2.17 (Symplectic Neighborhood Theorem). *Suppose $Q \subset M$ is a submanifold, and $\omega(0), \omega(1)$ are symplectic forms on M whose restrictions to $TM|_Q$ agree. Then there is a diffeomorphism ϕ of M to itself, fixed pointwise on Q , and with $\phi^*(\omega(1)) = \omega(0)$ on some open neighborhood of Q .*

2.4. Lagrangian submanifolds.

Definition 2.18. Let M^{2n}, ω be symplectic. A submanifold $L \subset M$ is *Lagrangian* if $T_p L$ is a Lagrangian subspace of $T_p M$ for every $p \in L$.

Theorem 2.19 (Weinstein). *Let $L \subset M$ be Lagrangian. Then there is an open neighborhood of L which is symplectomorphic to a neighborhood of the zero section of T^*L .*

Proof. Choose a compatible Riemannian metric on M near L , and consider the normal bundle νL . Then for each $p \in L$ the normal ν_p is also Lagrangian, and $\omega|_{T_p M}$ makes $T_p L \oplus \nu_p$ canonically isomorphic to $T_p L \oplus T_p^* L$. Thus the restriction of ω to $TM|_L$ is isomorphic to the restriction of the tautological symplectic form to $TT^*L|_L$. Now apply the Symplectic Neighborhood Theorem 2.17. \square

Example 2.20 (Closed 1-forms). Let M be a smooth n -manifold and let T^*M have its tautological symplectic structure. It turns out that a 1-form on M is a Lagrangian submanifold of T^*M if and only if it is closed.

Lemma 2.21. *Let M be a smooth n -manifold, and let T^*M have the tautological symplectic structure. A 1-form β on M determines an n -dimensional submanifold (i.e. a section) L_β of T^*M . Then L_β is Lagrangian if and only if β is closed. In particular, every exact 1-form $d\mathbf{f}$ defines a Lagrangian section of T^*M .*

Proof. Since the property of being Lagrangian is local, and since all the notions pertaining to the tautological symplectic structure are diffeomorphism invariant as objects on M , we may assume $M = \mathbb{R}^n$ and the tautological symplectic structure agrees with the standard symplectic structure on $T^*\mathbb{R}^n = \mathbb{R}^{2n}$.

At a point $(p, \beta(p)) \in T^*\mathbb{R}^n$ the submanifold L_β is tangent to the graph of $\nabla\beta(p)$ where ∇ is the covariant derivative on \mathbb{R}^n . This is Lagrangian if and only if $\nabla\beta(p)$ is symmetric, as in Example 1.14. But on any Riemannian manifold, exterior d on forms is the antisymmetrization of covariant derivative, so $\nabla\beta$ is symmetric pointwise if and only if $d\beta = 0$. \square

Example 2.22. Let $L \subset \mathbb{R}^{2n}$ be a compact oriented Lagrangian submanifold. By Weinstein's theorem the normal bundle of L is isomorphic to T^*L , so the obstruction to pushing L off itself is (the negative of) $\chi(L)$. But L can evidently be pushed off itself in \mathbb{R}^{2n} (just by translating it further than its diameter). So necessarily $\chi(L) = 0$.

Example 2.23. Let $L \subset \mathbb{R}^{2n}$ be a Lagrangian submanifold of standard symplectic \mathbb{R}^{2n} . There is a *Gauss map* $g : L \rightarrow \mathcal{L}_n$ that takes $p \in L$ to $T_p L \in \mathcal{L}_n$. The pullback of the Maslov class $\mu \in H^1(\mathcal{L}_n; \mathbb{Z})$ defines a class $g^*\mu \in H^1(L; \mathbb{Z})$ (sometimes denoted μ_L) also called the Maslov class.

In fact there is a canonical representative of the de Rham class corresponding to μ associated to the geometry of L . This observation is due to Morvan [15]. Recall, for any submanifold S of any Riemannian manifold the *second fundamental form* is a symmetric quadratic form \mathbb{I} on TS taking values in the normal bundle $\nu(S)$ defined as follows. For vectors $v, w \in T_p S$ extend these to smooth vector fields V, W on a neighborhood of p , and define $\mathbb{I}(v, w)$ to be the normal component of $\nabla_V(W)$. The *mean curvature* vector H is the trace of \mathbb{I} ; this is a section of the normal bundle. Now let L be a Lagrangian submanifold of \mathbb{R}^{2n} with mean curvature H , and let $\sigma_L \in \Omega^1(L)$ denote the 1-form $\sigma_L := \omega(H, \cdot)$.

Lemma 2.24. *The restriction of σ_L to L is closed, and the de Rham cohomology class of $-\sigma_L/\pi$ is μ_L .*

Proof. It suffices to show that the integral of σ_L along a loop is equal to (an appropriate multiple of) μ_L .

Let $p \in L$. After a unitary change of coordinates, we may assume that $T_p L$ is the totally real subspace $\mathbb{R}^n \subset \mathbb{C}^n$. Let $e_j = \partial/\partial x_j$ so that the e_j are an orthonormal basis for $T_p L$. Let $\gamma : (-\epsilon, \epsilon) \rightarrow L$ be a smooth path with $\gamma(0) = p$, and let $X := \gamma'(0)$. We denote ∇_{e_j} by ∇_j (this is both the Levi-Civita connection on $\mathbb{R}^{2n} = \mathbb{C}^n$ and ordinary partial differentiation with respect to x_j).

Since L is Lagrangian,

$$\sigma_L(X) = \sum_j \omega(\nabla_j e_j, X) = \sum_j \langle J \nabla_j e_j, X \rangle = - \sum_j \langle J e_j, \nabla_j X \rangle = - \sum_j \langle J e_j, \nabla_X e_j \rangle$$

because J commutes with ∇ , and $\langle J e_j, X \rangle = 0$ since L is Lagrangian.

Now let $A(t)$ be the unique 1-parameter family of unitary automorphisms of \mathbb{C}^n so that $A(t)$ takes $T_p L$ to $T_{\gamma(t)} L$ by parallel transport along γ . The j th diagonal entry of $A'(0)$ is i times the component of $\nabla_X e_j$ in the direction $\partial/\partial y_j$. Thus trace $A'(0) = -i\sigma_L(\gamma'(0))$. The claim follows. \square

We shall say more about the Maslov class for Lagrangian submanifolds of arbitrary symplectic manifolds in § 4.4.

3. HAMILTONIAN FLOWS

Let M, ω be symplectic. By Cartan's formula, a vector field X preserves ω if and only if $\omega(X, \cdot)$ is closed, and conversely every closed 1-form α satisfies $\alpha = \omega(X, \cdot)$ for some unique X preserving ω .

A vector field X is *Hamiltonian* if $\omega(X, \cdot)$ is exact, in which case there is a smooth function H (unique up to a constant) so that $dH = -\omega(X, \cdot)$. Conversely, given H we define X_H by $dH = -\omega(X_H, \cdot)$. The function H is also (rather confusingly) called the Hamiltonian (or, more usefully, the *energy*). Since $X_H(H) = dH(X_H) = -\omega(X_H, X_H) = 0$ it follows that H is constant on the orbits of X_H ; one may interpret this as ‘‘conservation of energy’’, hence the terminology.

If we choose Darboux coordinates $x_1, \dots, x_n, y_1, \dots, y_n$ near a point, the flow associated to the vector field X_H may be expressed by Hamilton's equations

$$\dot{x} = -\frac{\partial H}{\partial y}, \quad \dot{y} = \frac{\partial H}{\partial x}$$

Compare with the linear case § 1.4.

Example 3.1 (Geodesic flow). Let M be a Riemannian manifold. If $\beta \in T_p^*M$ then $\beta^\sharp \in T_pM$ is the unique vector with $\langle \beta^\sharp, v \rangle = \beta(v)$ for all $v \in T_pM$. Define $H(\beta) := -\|\beta^\sharp\|^2/2$. Then the associated Hamiltonian flow X_H is (after identifying T^*M with TM) the geodesic flow on TM associated to the metric.

To see this, choose standard local coordinates x_i, y_i on T^*M and write the metric $g_{ij}(x)dx_i dx_j$. Let g^{ij} denote the inverse (i.e. $g^{ij}g_{jk} = \delta_k^i$ with Einstein summation convention for repeated indices). Then $H(x, y) = -(1/2)g^{ij}(x)y_i y_j$ and Hamilton's equations become

$$\dot{x}_i = -\frac{\partial H}{\partial y_i} = g^{ij}(x)y_j, \quad \dot{y}_i = \frac{\partial H}{\partial x_i} = -\frac{1}{2} \frac{\partial g^{jk}(x)}{\partial x_i} y_j y_k$$

The first equation gives $y_j = g_{ij}\dot{x}_i$. Substituting in the second equation gives the geodesic equation $\ddot{x}_i + \Gamma_{jk}^i \dot{x}_j \dot{x}_k = 0$ for the Christoffel symbols $\Gamma_{jk}^i := (1/2)g^{il}(g_{lj,k} + g_{lk,j} - g_{jkl})$.

Example 3.2 (Magnetic flow). Let M be a Riemannian manifold and let B be a closed 2-form on M . If ω is the canonical symplectic form on T^*M then $\omega_B := \omega + \pi^*B$ is a symplectic form on T^*M with $\omega_B^n := \omega^n$. Let $H(\beta) := -\|\beta^\sharp\|^2/2$ be the Hamiltonian function associated to the metric as in Example 3.1. The flow X_H with respect to the form ω_B is called the *magnetic flow* on T^*M (here one interprets B as the magnetic field).

Raising one index by using the Riemannian metric, the 2-form B determines a field of infinitesimal antisymmetric endomorphisms of M (i.e. a field of infinitesimal rotations); the flow X_H is (infinitesimally) the composition of the geodesic flow with this rotation field. Thus (for example) if M is a hyperbolic surface and B is the area form the orbits of X_H project to curves on M with constant geodesic curvature. If $|v|$ is small, the orbit of v projects to a closed round circle in M ; if $|v|$ is large the orbit is almost a geodesic. The orbit of v with $|v| = 1$ projects to a horocycle; thus the restriction of X_H to the unit (co)tangent bundle is the horocycle flow on M .

Example 3.3. Let $G = \text{SO}(3)$ and let $X^*(\pi, \text{SO}(3))$ denote the (irreducible part of the) character variety of $\pi = \pi_1(S)$ for S a closed oriented surface of genus $g > 1$. Let $g \in \pi$ be the conjugacy class of an essential separating simple closed curve, so that we may write π as the amalgam of two subgroups A and B along the cyclic group generated by g . For $\rho \in X^*$ let $H(\rho) = \text{tr } \rho(g)$.

For a matrix $M = \rho(g)$ in $\text{SO}(3)$, the trace is equal to $1 + 2\cos(\theta)$ where M is rotation through angle θ around some axis ℓ . The Hamiltonian flow associated to H acts on (conjugacy classes of) representations by keeping them fixed on A , and conjugating B by rotation about the axis ℓ (which commutes with $\rho(g)$).

3.1. Characteristics. Let Q be a smooth codimension one submanifold of a symplectic manifold M . Every codimension 1 subspace of a symplectic vector space is coisotropic, so $T_q Q^\perp \subset T_q Q$ for all $q \in Q$ and these tangents integrate to a nonsingular 1-dimensional foliation on Q called the *characteristic foliation*.

Now suppose there is some smooth function H , and a regular value t for H so that Q is a component of $H^{-1}(t)$. Then X_H is nowhere zero on Q , and for any $V \in T_q Q$ we have $-\omega(X_H, V) = dH(V) = 0$. It follows that X_H is tangent to the characteristic foliation, so that the (unparameterized) flowlines of X_H on Q actually do not depend on H , but only

on Q itself. In particular, the question of whether the flow X_H has periodic orbits on Q or not depends only on Q and not on H .

Example 3.4 (Principle of least action). On T^*M with its canonical symplectic form $\omega = -d\lambda$ choose a Hamiltonian H . For points $p, q \in M$ the cotangent spaces T_p^* and T_q^* are Lagrangian. Let $Q \subset T^*M$ be a level set of H , and for $\gamma : [0, 1] \rightarrow Q$ with $\gamma(0) \in T_p^*$ and $\gamma(1) \in T_q^*$ define the *action*

$$A(\gamma) = \int_{\gamma} \lambda = \int_0^1 \lambda(\gamma'(t)) dt$$

Let $\Gamma(s, t)$ be a 1-parameter variation of $\gamma(t)$ so that $\Gamma(0, t) = \gamma(t)$, and denote $S := \partial_s \Gamma$ and $T := \partial_t \Gamma$. Then

$$\delta A(\gamma)(S(0, \cdot)) = \frac{d}{ds} \Big|_{s=0} \int_0^1 \lambda(T(s, t)) dt = \int_0^1 S(0, t) \lambda(T(0, t)) dt$$

Since S and T commute, $S\lambda(T) = -\omega(S, T) + T\lambda(S)$. The term $T\lambda(S)$ integrates to $\lambda(S(0, 1)) - \lambda(S(0, 0)) = 0$ because $S(0, \cdot)$ is tangent to T_p^* and T_q^* at the endpoints, and therefore in the kernel of λ . It follows that γ is critical for A if and only if γ' is in TQ^\perp .

3.2. Invariant volume and recurrence. Let Q be a smooth level set for H . Then dH is nonsingular on Q so we may write $\omega^n = \nu \wedge dH$ for ν some $(2n - 1)$ -form that restricts to a volume form on Q . Since Lie derivative acts as a derivation on forms, and since dH and ω^n are invariant under X_H , it follows that $(\mathcal{L}_{X_H} \nu) \wedge dH = 0$; i.e. the restriction $\nu|_{TQ}$ is invariant under X_H . In particular, by Poincaré recurrence, if Q is compact, and $U \subset Q$ is any (relatively) open set, any characteristic that intersects U must intersect it infinitely often.

3.3. Poisson bracket.

Definition 3.5 (Poisson bracket). For smooth functions F, G define the *Poisson bracket* $\{F, G\}$ by the formula

$$\{F, G\} := \omega(X_F, X_G) = -dF(X_G) = X_F(G)$$

In other words, the Poisson bracket measures the derivative of one function under the Hamiltonian flow generated by the other. In particular, two functions Poisson commute if and only if the values of one are constant on the orbits of the flow generated by the other.

Lemma 3.6. For Hamiltonian vector fields X_G, X_H we have $X_{\{G, H\}} = [X_G, X_H]$. Consequently, Poisson bracket makes smooth functions on M into a Lie algebra. Furthermore, Poisson bracket is a derivation; i.e. $\{F, GH\} = G\{F, H\} + \{F, G\}H$.

Proof. For three functions F, G, H

$$X_F(\omega(X_G, X_H)) = d\{G, H\}(X_F) = -\omega(X_{\{G, H\}}, X_F) = \{F, \{G, H\}\}$$

On the other hand, since X_F preserves ω ,

$$X_F(\omega(X_G, X_H)) = \omega([X_F, X_G], X_H) + \omega(X_G, [X_F, X_H])$$

Finally, since ω is closed,

$$0 = (d\omega)(X_F, X_G, X_H) = X_F(\omega(X_G, X_H)) - \omega([X_F, X_G], X_H) + \text{cyclic permutations}$$

Note by the way this already proves the Jacobi identity for Poisson bracket and therefore (since antisymmetry is obvious) proves the second statement of the Lemma. In any case we obtain

$$0 = \omega([X_F, X_G], X_H) + \omega([X_G, X_H], X_F) + \omega([X_H, X_F], X_G)$$

and therefore

$$\omega([X_G, X_H], X_F) = \omega(X_{\{G, H\}}, X_F)$$

Since both sides depend only on the value of X_F at each point, since this value can be arbitrary, and since ω is nondegenerate, $X_{\{G, H\}} = [X_G, X_H]$.

The derivation property follows from the derivation property for vector fields acting on functions, since $\{F, G\} = dG(X_F) = X_F(G)$. \square

Here is another way to summarize the discussion so far. Let $\mathfrak{symp}(M)$ denote the space of symplectic vector fields on M (this is a Lie algebra with respect to usual Lie bracket) and $\text{Ham}(M)$ the subspace of Hamiltonian vector fields. There are short exact sequences of Lie algebras

$$0 \rightarrow \text{Ham}(M) \rightarrow \mathfrak{symp}(M) \rightarrow H^1(M; \mathbb{R}) \rightarrow 0$$

where $H^1(M; \mathbb{R})$ is an abelian Lie algebra (i.e. Lie bracket is trivial) and

$$0 \rightarrow \mathbb{R} \rightarrow C^\infty(M) \rightarrow \text{Ham}(M) \rightarrow 0$$

For M compact, the second exact sequence splits:

Lemma 3.7. *Suppose M is compact. Then the sequence $\mathbb{R} \rightarrow C^\infty(M) \rightarrow \text{Ham}(M)$ splits, by assigning to every Hamiltonian vector field X the unique H with $X_H = X$ and $\int_M H\omega^n = 0$.*

Proof. We need to check that the splitting map is a Lie algebra homomorphism; equivalently, if G and H have average 0 then so does $\{G, H\}$. But

$$\int_M \{G, H\}\omega^n = \int_M (\mathcal{L}_{X_G}H)\omega^n = \int_M \mathcal{L}_{X_G}(H\omega^n)$$

which equals zero because the average of H is zero, and so is the average of any function obtained from H by pulling back by a volume-preserving diffeomorphism. \square

3.4. Hamiltonian circle actions. Suppose H is a Hamiltonian, and let Q be a (component of a) smooth level set of H for which the flow ϕ_t generated by X_H is periodic on Q (i.e. it defines a smooth, locally free S^1 action on Q). Because X_H is characteristic, X_H preserves the symplectic structure on TQ/TQ^\perp . It follows that the quotient Q/S^1 is symplectic.

Example 3.8. On \mathbb{C}^{n+1} the function $H := -\pi|z|^2$ generates a flow X_H which is periodic with period 1 on each nonsingular level set $|z|^2 = t > 0$. The quotient by the S^1 action is $\mathbb{C}\mathbb{P}^n$ with its standard (Fubini–Study) symplectic form.

3.5. Moment maps. Let (M, ω) be a symplectic manifold, and let G be a Lie group acting on M by symplectomorphisms. Taking derivatives gives a Lie algebra homomorphism $\mathfrak{g} \rightarrow \mathfrak{symplect}(M)$. For $\xi \in \mathfrak{g}$ denote the image by X_ξ . The action is said to be *weakly Hamiltonian* if every X_ξ is Hamiltonian. An action is weakly Hamiltonian if and only if the composition $\mathfrak{g} \rightarrow H^1(M; \mathbb{R})$ is zero; one can interpret this composition homologically as an element of $H^1(\mathfrak{g}) \otimes H^1(M; \mathbb{R})$; thus if $H^1(\mathfrak{g}) = 0$ every action of G is weakly Hamiltonian.

If the action is weakly Hamiltonian, then for every $\xi \in \mathfrak{g}$ we can choose a Hamiltonian function H_ξ with $-dH_\xi = \omega(X_\xi, \cdot)$. In fact, by choosing such an H_ξ on a basis, we can make a *linear* choice $c : \xi \rightarrow H_\xi$. The action is said to be *Hamiltonian* if we can choose c to be a Lie algebra homomorphism.

Warning 3.9. Some authors use the terminology ‘Hamiltonian’ for what we are calling ‘weakly Hamiltonian’, and ‘Poisson’ for what we are calling ‘Hamiltonian’.

Define a function τ by $\tau(\xi, \eta) = \{H_\xi, H_\eta\} - H_{[\xi, \eta]}$ so that τ vanishes if and only if c is a Lie algebra homomorphism. Since $\xi \rightarrow X_\xi$ is a Lie algebra homomorphism, it follows that τ takes values in \mathbb{R} ; i.e. $\tau \in \Lambda^2 \mathfrak{g}^*$. This implies that $\{H_{[\xi, \eta]}, H_\zeta\} = \{\{H_\xi, H_\eta\}, H_\zeta\}$ for all $\xi, \eta, \zeta \in \mathfrak{g}$.

Thinking of τ as a Lie algebra 2-cochain on \mathfrak{g} ,

$$\begin{aligned} (d\tau)(\xi, \eta, \zeta) &= -\tau([\xi, \eta], \zeta) + \text{cyclic permutations} \\ &= H_{[[\xi, \eta], \zeta]} - \{H_{[\xi, \eta]}, H_\zeta\} + \text{cyclic permutations} = 0 \end{aligned}$$

In other words, τ is a 2-cocycle. Furthermore, changing c by an element of \mathfrak{g}^* changes τ by an (arbitrary) 2-coboundary. Thus:

Lemma 3.10. *A weakly Hamiltonian action is Hamiltonian if and only if $[\tau] = 0$ in $H^2(\mathfrak{g})$.*

If an action is Hamiltonian, the map $\xi \rightarrow H_\xi$ is unique up to addition of a Lie algebra homomorphism $\mathfrak{g} \rightarrow \mathbb{R}$; i.e. an element of $H^1(\mathfrak{g})$.

A weakly Hamiltonian action is always Hamiltonian in the following important special cases:

- (1) For a semisimple Lie algebra we have $H^1(\mathfrak{g}) = H^2(\mathfrak{g}) = 0$ ([5], Thm. 21.1). Thus for G semisimple, any action by symplectomorphisms is Hamiltonian, and the map $\xi \rightarrow H_\xi$ is unique.
- (2) When M is compact, the splitting in Lemma 3.7 shows that any weakly Hamiltonian action is Hamiltonian.
- (3) If $M = T^*N$ for some smooth manifold N , and if the action of G preserves the tautological 1-form λ then $0 = \mathcal{L}_{X_\xi} \lambda = \iota_{X_\xi} d\lambda + d\iota_{X_\xi} \lambda$ so $\xi \rightarrow -\lambda(X_\xi)$ is a Hamiltonian function for X_ξ and one can check it defines a Lie algebra homomorphism.

Definition 3.11. Suppose G acts on M by symplectomorphisms, and the action is Hamiltonian. A *moment map* is a map $\mu : M \rightarrow \mathfrak{g}^*$ so that the map $\xi \rightarrow H_\xi(p)$ defined by

$$H_\xi(p) := \langle \mu(p), \xi \rangle$$

is a Lie algebra homomorphism generating the action.

Warning 3.12. If G is a Lie group acting on a manifold M , any $\xi \in \mathfrak{g}$ determines a vector field X_ξ on M by differentiating the orbits of $e^{t\xi} \cdot p$. This is an *anti-homomorphism* from the Lie algebra structure on \mathfrak{g} to the Lie algebra structure on vector fields on M .

Theorem 3.13. *Suppose G is connected, then for any Hamiltonian action of G on M , the moment map is G -equivariant. That is, $\mu(g \cdot p) = \text{Ad}_g^*(\mu(p))$ for every $g \in G$ and $p \in M$.*

Proof. We are trying to show, for every $\xi \in \mathfrak{g}$, that

$$H_\xi(g \cdot p) = \langle \mu(g \cdot p), \xi \rangle = \langle \text{Ad}_g^* \mu(p), \xi \rangle = \langle \mu(p), \text{Ad}_{g^{-1}} \xi \rangle = H_{\text{Ad}_{g^{-1}} \xi}(p)$$

Since G is connected, it suffices to prove this for g in the open subset of the form $e^{t\zeta}$ for some $\zeta \in \mathfrak{g}$. Furthermore, since the desired identity holds at $t = 0$, it suffices to show that the derivatives of both sides with respect to t agree.

The derivative of the right hand side is

$$\frac{d}{dt} H_{\text{Ad}_{e^{-t\zeta}} \xi}(p) = H_{-[\zeta, \text{Ad}_{e^{-t\zeta}} \xi]}(p) = \{H_\zeta, H_{\text{Ad}_{e^{-t\zeta}} \xi}\}(p) = \omega(X_\zeta, X_{\text{Ad}_{e^{-t\zeta}} \xi})(p)$$

where in the second term, $[\cdot, \cdot]$ means Lie bracket in \mathfrak{g} , and the action of G induces an anti-homomorphism from \mathfrak{g} to vector fields on M (hence the sign of the third term).

Whereas the derivative of the left hand side is

$$\frac{d}{dt} H_\xi(e^{t\zeta} \cdot p) = X_\zeta H_\xi(e^{t\zeta} \cdot p) = -\omega(X_\xi, X_\zeta)(e^{t\zeta} \cdot p) = -\omega(X_{\text{Ad}_{e^{-t\zeta}} \xi}, X_{\text{Ad}_{e^{-t\zeta}} \zeta})(p)$$

so the two sides are equal, since $\text{Ad}_{e^{-t\zeta}} \zeta = \zeta$. \square

It follows that the moment map takes G -orbits in M to coadjoint orbits in \mathfrak{g}^* . Kostant showed that if $H^1(\mathfrak{g}) = H^2(\mathfrak{g}) = 0$ (for instance, if G is semisimple), the map from each G -orbit to its corresponding coadjoint orbit is a covering map. We shall not use this fact in the sequel.

Example 3.14. Let G act on \mathfrak{g}^* by the coadjoint action (Example 2.11), and let $M \subset \mathfrak{g}^*$ be a coadjoint orbit. Then $\mu : M \rightarrow \mathfrak{g}^*$ is simply the inclusion map.

Example 3.15. If $H \subset G$ is a Lie subgroup, and $\mu_G : M \rightarrow \mathfrak{g}^*$ is a moment map associated to a Hamiltonian action of G on M , the composition of μ_G with the projection $\mathfrak{g}^* \rightarrow \mathfrak{h}^*$ dual to inclusion $\mathfrak{h} \rightarrow \mathfrak{g}$ is a moment map associated to the action of H .

Example 3.16. Let $U(n)$ act on \mathbb{C}^n by the standard matrix action. This is Hamiltonian, and if $\xi \in \mathfrak{u}(n)$ is a skew-Hermitian matrix then $H_\xi(z) = (1/2)iz^*\xi z$ where z^* is the conjugate transpose of the (column) vector z . The inner product $\langle \xi, \zeta \rangle := \text{tr}(\xi^*\zeta)$ identifies $\mathfrak{u}(n)$ with its dual; with respect to this identification $\mu(z) = -(1/2)izz^*$.

Let $T^n \subset U(n)$ denote the maximal torus consisting of diagonal matrices with entries of norm 1. The Lie algebra \mathbb{R}^n may be identified with imaginary diagonal complex matrices (note under this identification, each S^1 factor in T^n has length 2π). Thus the moment map $\mu : \mathbb{C}^n \rightarrow \mathbb{R}^n$ associated to the T^n action takes z to $-(1/2)(|z_1|^2, |z_2|^2, \dots, |z_n|^2)$. In particular, the image of each $U(n)$ orbit is a simplex.

Example 3.17 (Atiyah–Bott [3]). Let S be a compact oriented Riemann surface, and let G be a compact Lie group with Lie algebra \mathfrak{g} . Exterior d defines a connection on the trivial

bundle $S \times G$; any other connection d_A differs from this by a 1-form A with coefficients in \mathfrak{g} . Let $\mathcal{A} := \Omega^1(S, \mathfrak{g})$ denote the space of connections; therefore also $T_A\mathcal{A} = \Omega^1(S, \mathfrak{g})$.

An Ad-invariant inner product $\langle \cdot, \cdot \rangle$ on \mathfrak{g} (e.g. the Killing form if G is semisimple) defines a pairing $\mathfrak{g}^2 \rightarrow \mathbb{R}$ and lets us define a symplectic form on $T_A\mathcal{A}$ by

$$\omega(\alpha, \beta) := \int_S \langle \alpha \wedge \beta \rangle$$

Since S is a Riemann surface, Hodge star is defined on $\Omega^1(S)$ and therefore also on $\Omega^1(S, \mathfrak{g})$. It is compatible with the symplectic form so together they define a Hermitian form on $T_A\mathcal{A}$ which (under the natural identification of \mathcal{A} with $T_A\mathcal{A}$) makes \mathcal{A} into a(n infinite dimensional) Kähler manifold.

The *gauge group* $\mathcal{G} := \text{Map}(S, G)$ acts on the space of trivializations of the bundle, and therefore acts on \mathcal{A} . As a formula, the action is given by $g^*A = g^{-1}dg + g^{-1}Ag$. The Lie algebra of \mathcal{G} is $\Omega^0(S, \mathfrak{g})$ and under this action, $\xi \in \Omega^0(S, \mathfrak{g})$ maps to the vector field whose value at $A \in \mathcal{A}$ is $d_A\xi := d\xi + [A, \xi]$. This action preserves the Hermitian form.

The *curvature* of a connection A is $F_A \in \Omega^2(S, \mathfrak{g})$ defined by $F_A := dA + (1/2)[A \wedge A]$. Note that curvature transforms tensorially under the gauge group: $F_{g^*A} = g^{-1}F_Ag$. The space $\Omega^2(S, \mathfrak{g})$ may be interpreted as the dual of the Lie algebra $\Omega^0(S, \mathfrak{g})$ under the pairing $\langle F, \xi \rangle = \int_S \langle F \wedge \xi \rangle$. Thus for any ξ , there is a function $H_\xi : \mathcal{A} \rightarrow \mathbb{R}$ defined by $A \rightarrow \int_S \langle F_A \wedge \xi \rangle$. For any $B \in \Omega^1(S, \mathfrak{g})$ we compute

$$\left. \frac{d}{dt} \right|_{t=0} F_{A+tB} = dB + [A \wedge B] = d_A B$$

and therefore

$$dH_\xi(B) = \int_S \langle d_A B, \xi \rangle = \int_S \langle B, d_A \xi \rangle = -\omega(d_A \xi, B)$$

where we use the formula $d\langle X, Y \rangle = \langle d_A X, Y \rangle + (-1)^{\deg(X)} \langle X, d_A Y \rangle$ and integration by parts.

It follows that H_ξ is a Hamiltonian for the vector field $d_A \xi$ and one can check that $\xi \rightarrow H_\xi$ defines a Lie algebra homomorphism (this follows from the transformation rule for curvature under the gauge group). Thus $\mu : \mathcal{A} \rightarrow \Omega^2(S, \mathfrak{g})$ defined by $\mu(A) = F_A$ is the moment map.

3.6. Marsden–Weinstein quotient. A coisotropic submanifold carries a canonical foliation with isotropic leaves:

Lemma 3.18. *If $Q \subset M$ is a coisotropic submanifold, the distribution $TQ^\perp \subset TQ$ is integrable.*

Proof. If X, Y are sections of TQ^\perp and Z is any section of TQ then

$$0 = d\omega(X, Y, Z) = X(\omega(Y, Z)) - \omega([X, Y], Z) + \text{cyclic permutations} = -\omega([X, Y], Z)$$

so $[X, Y]$ is a section of TQ^\perp . \square

Let $U \subset Q$ be an open submanifold on which the isotropic foliation is a product, and let $Q \rightarrow Q/\sim$ denote projection to the (local) leaf space. If X is a vector field tangent to the leaves, then $\mathcal{L}_X \omega = 0$. This shows that ω descends to a well-defined symplectic form on Q/\sim .

A foliation of a manifold M is *regular* if every point has a neighborhood foliated as a product that intersects each leaf at most once. If a foliation is regular, then its leaf space is a smooth manifold, and the projection map to the leaf space is a submersion. A coisotropic submanifold is *regular* if its isotropic foliation is regular. Thus we have shown:

Lemma 3.19. *Let Q be coisotropic, and suppose its isotropic foliation is regular. Then the leaf space Q/\sim is a symplectic manifold.*

Warning 3.20. If M is compact, then a regular foliation must have compact leaves; however the converse is false: Sullivan gave an example of a foliation of a compact 5-manifold by circles whose leaf space is not even Hausdorff. If there is a uniform bound on the volume of the leaves then M is ‘nearly’ regular: the leaf space has the natural structure of a smooth orbifold and the projection map is a submersion in the category of orbifolds.

Regular coisotropic submanifolds arise naturally as fibers of moment maps:

Theorem 3.21 (Marsden–Weinstein). *Let G be a connected Lie group, and suppose we are given a Hamiltonian action of G on M with moment map $\mu : M \rightarrow \mathfrak{g}^*$ such that*

- (1) 0 is a regular value of μ (so that $\mu^{-1}(0)$ is a smooth submanifold); and
- (2) G acts freely and properly on $\mu^{-1}(0)$.

Then $\mu^{-1}(0)$ is a coisotropic manifold, and the orbits of G are the isotropic leaves. Therefore the quotient $\mu^{-1}(0)/G$ is a symplectic manifold called the Marsden–Weinstein quotient and denoted $M//G$.

Proof. Let $p \in \mu^{-1}(0)$ and let $O := Gp$ denote the orbit of p under G . Since G acts freely, and 0 is a regular value, the dimension of O is equal to the codimension of $\mu^{-1}(0)$. Therefore the theorem will be proved if we can show that $T_p O \subset T_p \mu^{-1}(0)^\perp$.

By the definition of the moment map, for each $\xi \in \mathfrak{g}$ and $p \in \mu^{-1}(0)$ we have $H_\xi(p) = \langle \mu(p), \xi \rangle = 0$. Thus H_ξ is constant (in fact 0) on $\mu^{-1}(0)$ so for any $v \in T_p \mu^{-1}(0)$ we have

$$0 = vH_\xi = dH_\xi(v) = -\omega(X_\xi, v)$$

so that $X_\xi \in T_p \mu^{-1}(0)^\perp$. Since $T_p O$ is spanned by X_ξ we are done. □

4. HOLOMORPHIC CURVES

4.1. Compatible almost complex structures. Let M be a smooth manifold. An *almost complex* structure on M is a smooth bundle endomorphism $J : TM \rightarrow TM$ with $J^2 = -\text{Id}$. This is the same thing as an isomorphism from TM to the underlying real bundle $E_{\mathbb{R}}$ of some complex vector bundle E over M .

Definition 4.1. Let (M, ω) be symplectic. An almost complex structure J on TM is *tame* if $\omega(X, JX) > 0$ for all nonzero vectors $X \in T_p M$ and is *compatible* if it is tame, and if $\omega(JX, JY) = \omega(X, Y)$ for any two $X, Y \in T_p M$.

If J is tame, then $g(X, Y) := \omega(X, JY)$ defines a Riemannian metric on M , and if J is compatible, then $h(X, Y) := g(X, Y) - i\omega(X, Y)$ is a Hermitian metric on M giving it the structure of an almost Kähler manifold (see Example 2.8).

Let $\mathcal{J}(M, \omega)$ denote the space of compatible almost complex structures and $\mathcal{J}_\tau(M, \omega)$ the space of tame almost complex structures.

Lemma 4.2. *Both spaces $\mathcal{J}(M, \omega)$ and $\mathcal{J}_\tau(M, \omega)$ are nonempty and contractible, and the inclusion $\mathcal{J} \rightarrow \mathcal{J}_\tau$ is a homotopy equivalence.*

Proof. Let V, ω be a symplectic vector space. Let g be a positive definite symmetric inner product on V . Then we obtain an automorphism A of V by $\omega(X, Y) = g(AX, Y)$. Evidently A is g -skew-adjoint, so that $-A^2$ is g -self-adjoint and positive definite, and we can write $Q^2 = -A^2$ for some unique g -self-adjoint positive definite Q commuting with A . Then $Q^{-1}A$ satisfies $Q^{-1}AQ^{-1}A = Q^{-2}A^2 = -\text{Id}$. Furthermore,

$$\omega(X, Q^{-1}AX) = g(AX, Q^{-1}AX) > 0$$

and

$$\omega(Q^{-1}AX, Q^{-1}AY) = g(AQ^{-1}AX, Q^{-1}AY) = g(AX, Y) = \omega(X, Y)$$

Thus $Q^{-1}A$ is a compatible complex structure on V . If we define $g'(X, Y) := \omega(X, Q^{-1}AY)$ then $g' = g$ if and only if $Q = \text{Id}$, i.e. if and only if A is already a compatible complex structure.

Thus there is a natural retraction from the space of positive definite symmetric inner products on V to the space of compatible complex structures. This retraction depends continuously on the symplectic structure, so applying the retraction fiberwise to TM , we obtain a retraction from the space of Riemannian metrics on M to $\mathcal{J}(M, \omega)$. This proves that $\mathcal{J}(M, \omega)$ is nonempty and contractible.

We now show that $\mathcal{J} \rightarrow \mathcal{J}_\tau$ is a homotopy equivalence. This follows from the fiberwise statement, so we just need to show for V, ω a symplectic vector space, that the ω -tame complex structures deformation retract to the ω -compatible complex structures. Let J be ω -tame; i.e. $\omega(X, JX) > 0$ for all X . Then $g_J(X, Y) := (1/2)(\omega(X, JY) + \omega(Y, JX))$ is a positive definite symmetric inner product. As above, g_J defines a compatible almost complex structure. This defines a retraction $\mathcal{J}_\tau \rightarrow \mathcal{J}$ and one can check that the fibers are contractible. \square

In fact, the space of compatible complex structures on a symplectic vector space V^{2n} is evidently isomorphic to the Siegel upper-half space \mathfrak{H}^n , which we have already seen is contractible.

4.2. Calibrations.

Definition 4.3. If (M, g) is a Riemannian manifold, a p -form α is a *calibration* if

- (1) α is closed; and
- (2) for any p -vector ξ , we have $|\alpha(\xi)| \leq \text{vol}_g(\xi)$.

Let $G(\alpha) \subset \Lambda^p M$ be the space where $\alpha(\xi) = \text{vol}_g(\xi)$. A p -dimensional submanifold N is *calibrated* by α if $TN \subset G(\alpha)$.

A calibrated manifold N is volume minimizing in its homology class. For, if N' is homologous to N , then

$$\text{vol}(N) = \int_N d\text{vol} = \int_N \alpha = \int_{N'} \alpha \leq \int_{N'} d\text{vol} = \text{vol}(N')$$

Lemma 4.4. *Let J be a compatible almost complex structure on (M, ω) with associated Riemannian metric $g(X, Y) := \omega(X, JY)$. Then ω is a calibration with respect to g , and $G(\omega) \subset \Lambda^2 M$ consists precisely of the oriented J -invariant 2-planes.*

Proof. Let $X, Y \in T_p M$ be orthonormal with respect to g . Then $\omega(X, Y) = -g(X, JY)$ so $|\omega(X, Y)| \leq 1 = \|X \wedge Y\|$ and $\omega(X, Y) = 1$ if and only if $Y = JX$. \square

An immersed surface $u : S \rightarrow M$ is calibrated by ω if and only if TS is J -invariant and J acts as a positive rotation of TS by $\pi/2$ (note that this implies S is oriented). The metric gives S a conformal structure and therefore the structure of a Riemann surface; thus $u : S \rightarrow M$ is calibrated if and only if it is *holomorphic* in the sense that $du(iX) = Jdu(X)$ for any vector $X \in TS$ (the terms “ J -holomorphic” and “pseudoholomorphic” are also standard). We shall discuss holomorphic curves in § 4.5.

4.3. Characteristic classes. Since $\mathcal{J}(M^{2n}, \omega)$ is nonempty, every symplectic manifold admits an almost complex structure.

An n -dimensional complex vector bundle E over a base B has Chern classes $c_j(E) \in H^{2j}(B; \mathbb{Z})$. These are related to the characteristic classes of the underlying real bundle $E_{\mathbb{R}}$ as follows. Let $c_j := c_j(E)$ and denote the Stiefel–Whitney and Pontrjagin classes of $E_{\mathbb{R}}$ by $w_j := w_j(E_{\mathbb{R}}) \in H^j(B; \mathbb{Z}/2\mathbb{Z})$ and $p_j := p_j(E_{\mathbb{R}}) \in H^{4j}(B; \mathbb{Z})$. Then

- (1) $w_{2j} = c_j \bmod 2$ and $w_{2j+1} = 0$ for all j ;
- (2) $1 - p_1 + p_2 - \cdots \pm p_n = (1 - c_1 + c_2 - \cdots \pm c_n)(1 + c_1 + c_2 + \cdots + c_n)$;
- (3) $c_n = e(E_{\mathbb{R}})$, the Euler class.

Example 4.5. Let W be a closed simply-connected 4-manifold. Then $w_1 = w_3 = 0$ and w_2 admits an integral lift. If W admits an almost complex structure then $c_2([W]) = \chi(W)$. Furthermore, the signature σ satisfies

$$3\sigma = p_1([W]) = (c_1^2 - 2c_2)([W])$$

In other words, w_2 has an integral lift c_1 satisfying $c_1^2([W]) = 2\chi + 3\sigma$. Since w_2 is characteristic, $c_1^2([W]) = \sigma \bmod 8$ so $\chi + \sigma = 0 \bmod 4$. Thus (for example) S^4 does not admit an almost-complex structure.

A symplectic manifold is said to be *monotone* if $[\omega]$ is a positive multiple of c_1 .

4.4. Maslov class. Let M^{2n}, ω be a symplectic manifold and let J be a compatible almost complex structure with associated metric g making M almost-Kähler. For every $p \in M$ let \mathcal{L}_p denote the space of Lagrangian subspaces of $T_p M$. An identification of $T_p M$ with \mathbb{C}^n identifies \mathcal{L}_p with $U(n)/O(n)$ so that we may write $\det^2 : \mathcal{L}_p \rightarrow S^1$; likewise if \mathcal{L}_p^+ denotes the space of oriented Lagrangian subspaces of $T_p M$ (which is a connected orientable double cover of \mathcal{L}_p) we may write $\det : \mathcal{L}_p \rightarrow S^1$. However this map depends on the choice of complex coordinate.

What is well-defined independent of coordinates is the map $\Lambda^n : \mathcal{L}_p^+ \rightarrow \Lambda_{\mathbb{C}}^n T_p M$ that takes an oriented orthonormal basis x_1, \dots, x_n for $\pi \in \mathcal{L}_p^+$ to $x_1 \wedge \cdots \wedge x_n \in \Lambda_{\mathbb{C}}^n T_p M$ where $\Lambda_{\mathbb{C}}^n T_p M$ denotes the top exterior power of $T_p M$ thought of as a complex vector space. One may also write $\Lambda_{\mathbb{C}}^n T_p M$ suggestively as L_{-K} , the dual of the tautological line bundle $L_K := \Lambda_{\mathbb{C}}^n T_p^* M$. By abuse of notation we denote this map by $\det : \mathcal{L}_p^+ \rightarrow S^1(L_{-K})$ and likewise $\det^2 : \mathcal{L}_p \rightarrow S^1(L_{-2K})$.

For $L \subset M$ a Lagrangian submanifold there is an associated Gauss map pointwise from $p \in L$ to $T_p L \in \mathcal{L}_p$, and if L is orientable there is a lift of this map to \mathcal{L}_p^+ . We thereby

obtain a section $\sigma : L \rightarrow E$ (resp. $\sigma^+ : L \rightarrow E^+$ for L orientable) where E is the circle bundle over L obtained by restricting $S^1(L_{-2K})$ (resp. $S^1(L_{-K})$).

Recall that M is said to be *monotone* if $[\omega]$ is proportional to c_1 (by a positive constant of proportionality κ). This means that there is a real positive constant κ and a Hermitian connection $2A$ on L_{-2K} (unique up to gauge transformations) with curvature $2\kappa\omega$. In particular, the curvature of this connection vanishes identically on any Lagrangian L . Because the connection is flat, parallel transport defines a projection $E|_U \rightarrow S^1$ over sufficiently small open sets $U \subset L$ unique up to composition with a (locally constant!) rotation. Thus the angle form $d\theta$ on S^1 pulls back to a canonical *closed* 1-form on E , and the further pullback of this form under σ defines a de Rham cohomology class $\mu \in H^1(L; \mathbb{R})$ which is called the *universal Maslov class of L* (an analogous construction defines $\mu^+ \in H^1(L; \mathbb{R})$ with $2\mu^+ = \mu$ when L is orientable).

4.5. Holomorphic curves.

Definition 4.6 (Holomorphic Curve). Let M, ω be symplectic, and let J be a compatible almost-complex structure. Let Σ be a compact Riemann surface. A *J -holomorphic curve* (or just *holomorphic curve* if J is understood) is a smooth map $u : \Sigma \rightarrow M$ for which $du(iX) = Jdu(X)$ for all $X \in T\Sigma$.

If we write $\bar{\partial}_J u := (1/2)(du + Jdu i)$ then u is J -holomorphic if and only if $\bar{\partial}_J u = 0$. Let $z := x + iy$ be a local holomorphic coordinate on Σ . In a symplectic trivialization with local Darboux coordinates on M , the almost complex structure is given by a field of symplectic matrices $J(p)$ satisfying $J^2 = -\text{Id}$ pointwise. A map u is holomorphic if and only if it satisfies the system of first-order equations $\partial_y u = J(u(z))\partial_x u$. The symbol of this PDE is the same as that of the usual Cauchy–Riemann equations, so this PDE is elliptic.

One may consider holomorphic curves of any genus but for now we restrict ourselves to the best behaved examples, which are genus zero. We give S^2 its usual structure as a Riemann surface, by identifying it with the Riemann sphere.

Definition 4.7. Fix an almost compatible almost-complex structure J on M, ω . Fix a homology class $A \in H_2(M; \mathbb{Z})$. Let $\mathcal{M}(A, J)$ denote the set of J -holomorphic maps $u : S^2 \rightarrow M$ where $u_*([S^2]) = A$.

Let $\bar{\mathcal{M}}(A, J)$ denote the quotient of $\mathcal{M}(A, J)$ by the equivalence relation that identifies two maps if they differ by precomposition by a holomorphic automorphism of S^2 .

If A, J are understood we write simply \mathcal{M} and $\bar{\mathcal{M}}$. The group $\text{Aut}(S^2)$ of holomorphic automorphisms of S^2 is the Möbius group. It is 3 complex dimensional and noncompact, and $\bar{\mathcal{M}} = \mathcal{M}/\text{Aut}(S^2)$. Thus \mathcal{M} will never be compact unless $A = 0$ and \mathcal{M} consists only of constant maps. It is a miracle that $\bar{\mathcal{M}}$ is very often compact, or may be compactified by adding suitable ‘degenerate’ curves in a comprehensible way.

For $u \in \mathcal{M}$ let $\text{area}(u)$ denote the ‘area’ of $u(S^2)$; i.e. the integral over S^2 of the pullback of the area form in the Riemannian metric associated to the almost complex structure J . For a holomorphic curve, $\text{area}(u) := \int_{S^2} u^*\omega = [\omega](A)$. Thus $\text{area}(u)$ is constant over \mathcal{M} , and is strictly positive unless $A = 0$. In particular, if $A \neq 0$ and $[\omega](A) \leq 0$ then \mathcal{M} is empty. Furthermore, since the domain of u is a 2-sphere, \mathcal{M} is empty unless A is in the image of the Hurewicz map $\pi_2(M) \rightarrow H_2(M)$.

4.6. Local structure. Because the operator $\bar{\partial}_J$ agrees with the Cauchy–Riemann operator to leading order, we expect that the local structure of a J -holomorphic curve should resemble the local structure of an ‘honest’ holomorphic curve in an algebraic variety. This intuition is realized by the following theorem of Micallef and White [14], Thms. 6.1 and 6.2:

Theorem 4.8 (Micallef–White). *Let $u : \Sigma \rightarrow M^{2n}$ be a non-constant J -holomorphic curve for some Riemann surface Σ (not necessarily connected). Then for any point $x \in u(\Sigma)$ and for all $p_j \in u^{-1}(x)$ there are neighborhoods $p_j \in U_j \subset \Sigma$ and $x \in V \subset M$ and coordinate charts $\psi : V \rightarrow \mathbb{C}^n$, $\phi_j : U_j \rightarrow \mathbb{C}$ with $\psi(x) = 0$ and $\phi_j(p_j) = 0$, and so that*

$$\psi u \phi_j^{-1}(z) = (z^{Q_j}, f_j(z))$$

where $f_j(z) \in \mathbb{C}^{n-1}$ vanishes to order $\geq Q_j \geq 1$ at $z = 0$.

Furthermore, if J is C^2 , then ψ is C^1 and each ϕ_j is $C^{2,\alpha}$ for some positive α .

This implies (for instance) that multiple points x in the image of u are isolated unless u factors through a (branched) cover. It also lets us control the geometry of a singularity.

Theorem 4.8 is actually a corollary of a more general theorem about the local structure of singularities of minimal surfaces in Riemannian manifolds, and has nothing fundamentally to do with symplectic or almost-complex geometry per se. We refer the reader to [14] for the proof.

4.7. Fredholm theory. We would like to topologize \mathcal{M} and $\overline{\mathcal{M}}$. Our goal is to show that for ‘generic’ J and suitable A , the set \mathcal{M} can be given the structure of a smooth oriented finite-dimensional manifold, and $\overline{\mathcal{M}}$ a compact smooth oriented finite-dimensional manifold; and furthermore that the oriented cobordism class of $\overline{\mathcal{M}}$ is well-defined, independent of the choice of (generic) J . Throughout this section and the next we only give a brief overview; for details see [12], Chapter 3.

A bounded linear operator between Banach spaces is *Fredholm* if it has finite dimensional kernel and cokernel. The *index* is the difference of these dimensions. A smooth map between Banach manifolds is Fredholm if its differential is Fredholm as a map between tangent spaces at each point; the index is (locally) constant.

Our first step is to show that \mathcal{M} is the fiber of a Fredholm map between suitable Banach manifolds. Fix some unspecified sufficiently large positive integers k, l with $k \leq l + 1$. We will consider the following Banach manifolds:

Definition 4.9 (Banach manifolds). Let $\mathcal{J}^l := \mathcal{J}^l(M, \omega)$ denote the space of C^l almost-complex structures compatible with ω . This is a Banach manifold in the C^l topology (i.e. the uniform norm on the coefficients of J and its partial derivatives of order $\leq l$ in a local coordinate).

Let \mathcal{B} denote the space of smooth maps $u : S^2 \rightarrow M$ in the homology class of A , and let $\mathcal{B}^{k,p}$ denote the completion of \mathcal{B} with respect to the L^p norm on the partial derivatives of order $\leq k$. This is a Banach manifold locally modeled on a Sobolev space of the form $W^{k,p}$.

For a fixed J let $\mathcal{E} \rightarrow \mathcal{B}$ denote the vector bundle whose fiber over u is the space $\mathcal{E}_u := \Omega_j^{0,1}(u^*TM)$ of smooth J -antilinear 1-forms on S^2 with values in u^*TM , and let

$\mathcal{E}^{k-1,p} \rightarrow \mathcal{B}^{k,p}$ be the Sobolev completion. This is a (Sobolev) Banach space bundle over $\mathcal{B}^{k,p}$ for each J whose fibers are modeled on a Sobolev space of the form $W^{k-1,p}$.

The operator $\bar{\partial}_J$ defines a smooth section $\mathcal{B}^{k,p} \rightarrow \mathcal{E}^{k-1,p}$, and the moduli space \mathcal{M}_J is the preimage of the zero section. Thus we would like to understand when $\bar{\partial}_J$ is transverse to the zero section. Consider the linearization of the operator $\bar{\partial}_J$ at u :

$$(d\bar{\partial}_J)_u : T_u \mathcal{B}^{k,p} \rightarrow T_{\bar{\partial}_J u} \mathcal{E}^{k-1,p}$$

The tangent space to $\mathcal{E}^{k-1,p}$ at $\bar{\partial}_J u$ splits — although not canonically — as $T_u \mathcal{B}^{k,p} \oplus \mathcal{E}_u^{k-1,p}$. The reason the splitting is non-canonical is that there is no natural way to identify the fibers $\mathcal{E}_u^{k-1,p}$ of the bundle $\mathcal{E}^{k-1,p}$ at different u , since they are 1-forms on S^2 with values in different vector bundles. However, by using the Levi-Civita connection ∇ on M with respect to the metric associated to J we may parallel transport maps u' sufficiently close to u along short geodesics to u , and thereby identify $\mathcal{E}_{u'}$ with \mathcal{E}_u . Denote by D_u the projection of $(d\bar{\partial}_J)_u$ to $\mathcal{E}_u^{k-1,p}$ with respect to this splitting. In other words:

$$D_u : W^{k,p}(u^*TM) \rightarrow W^{k-1,p}(\Lambda^{0,1}T^*S^2 \otimes_J u^*TM)$$

For $\bar{\partial}_J$ to be transverse to the zero section at a holomorphic curve u is for D_u to be *surjective*.

A formula for D_u is

$$D_u \xi = (1/2) (\nabla \xi + J(u) \nabla \xi i + (\nabla_\xi J)(u) \partial_J(u) i)$$

Although J might not be integrable, the term $\nabla_\xi J$ is 0th order in ξ , so in local coordinates, $D_u \xi = \bar{\partial}_J \xi + 0$ th order terms in ξ . It follows that D_u is elliptic, so that there is an estimate of the form

$$\|\xi\|_{W^{k,p}} \leq C(\|D_u \xi\|_{W^{k-1,p}} + \|\xi\|_{W^{k-1,p}})$$

which implies by the usual argument of elliptic regularity that D_u has a closed range and a finite dimensional kernel. Integration by parts proves the same for the adjoint

$$D_u^* : W^{k,q}(\Lambda^{0,1}T^*S^2 \otimes_J u^*TM) \rightarrow W^{k-1,q}(u^*TM)$$

(where $1/p + 1/q = 1$) so D_u is Fredholm.

An almost-complex structure J is *regular* if D_u as above is surjective for every $u \in \mathcal{M}(A, J)$. A suitable version of the implicit function theorem says that if J is regular, then $\mathcal{M}(A, J)$ is a smooth manifold of dimension equal to the index of D_u . This is not a purely formal statement: for smooth maps between infinite dimensional manifolds one needs uniform control on a right inverse to construct a diffeomorphism from a neighborhood of $0 \in \ker(D_u)$ onto a neighborhood of $u \in \mathcal{M}$. The required estimates can be found in [12], § 3.3.

In summary:

Proposition 4.10. *For any J the operator D_u is Fredholm. Suppose J is regular so that D_u is surjective for every $u \in \mathcal{M}(A, J)$. Then $\mathcal{M}(A, J)$ is a smooth manifold of dimension equal to the index of D_u .*

In fact $\mathcal{M}(A, J)$ is an *oriented* manifold. The kernel of D_u is not quite J -invariant unless J is integrable; however one may homotop D_u in the obvious way to an operator whose kernel is J -invariant. This gives an almost complex structure and hence an orientation on the tangent space to $\mathcal{M}(A, J)$.

4.8. Regular almost complex structures. As we vary J in \mathcal{J}^l the moduli spaces $\mathcal{M}(A, J)$ vary.

Definition 4.11. The *universal* moduli space $\mathcal{M}^l(A)$ is

$$\mathcal{M}^l(A) := \{(u, J) \in \mathcal{B}^{k,p} \times \mathcal{J}^l \text{ such that } \bar{\partial}_J(u) = 0\}$$

We would like to show that $\mathcal{M}^l(A)$ is a smooth Banach submanifold of the Banach manifold $\mathcal{B}^{k,p} \times \mathcal{J}^l$.

The tangent space to \mathcal{J}^l at J is the space of C^l endomorphisms Y of TM that preserve ω (i.e. $\omega(Yv, w) + \omega(v, Yw) = 0$ pointwise) and anti-commute with J ; i.e.

$$0 = \left. \frac{d}{dt} \right|_{t=0} (J + tY)^2 = JY + YJ$$

Denote this tangent space by $C^l(\text{End}(TM, J, \omega))$.

By abuse of notation let $\mathcal{E}^{k-1,p}$ be the Banach bundle over $\mathcal{B}^{k,p} \times \mathcal{J}^l$ with fiber over (u, J) equal to

$$\mathcal{E}_{(u,J)}^{k-1,p} = W^{k-1,p}(\Lambda^{0,1}T^*S^2 \otimes_J u^*TM)$$

Let $\mathcal{F} : \mathcal{B}^{k,p} \times \mathcal{J}^l \rightarrow \mathcal{E}^{k-1,p}$ be the section $\mathcal{F}(u, J) = \bar{\partial}_J(u)$, so that $\mathcal{M}(A)$ is the preimage of the zero section of \mathcal{F} . Again, by using the Levi-Civita connection to identify nearby fibers of $\mathcal{E}^{k-1,p}$, we can compose $(d\mathcal{F})_{(u,J)}$ with projection to $\mathcal{E}_{(u,J)}^{k-1,p}$ to get

$$D\mathcal{F}_{(u,J)} : W^{k,p}(u^*TM) \times C^l(\text{End}(TM, J, \omega)) \rightarrow W^{k-1,p}(\Lambda^{0,1}T^*S^2 \otimes_J u^*TM)$$

A formula for $D\mathcal{F}$ is

$$D\mathcal{F}_{(u,J)}(\xi, Y) = D_u\xi + (1/2)Y(u)du i$$

The operator D_u is Fredholm, so its cokernel is finite dimensional. To show that $\mathcal{M}(A)$ is a smooth Banach manifold it suffices to show that $D\mathcal{F}$ is surjective, and since D_u is Fredholm, it is enough to show that its range is dense.

If not, there is some nonzero $\eta \in W^{k-1,q}(\Lambda^{0,1}T^*S^2 \otimes_J u^*TM)$ that vanishes on the image of $D\mathcal{F}$; i.e.

$$\int_{S^2} \langle \eta, D_u\xi \rangle = 0 \text{ and } \int_{S^2} \langle \eta, Y(u)du i \rangle = 0$$

for every ξ and every Y . The first equation says that $D_u^*\eta = 0$; elliptic regularity implies that if η vanishes on an open set then it is identically zero.

A holomorphic curve u is *injective* at $z \in S^2$ if $du(z) \neq 0$ and if $u^{-1}(u(z)) = z$. If z is injective, and η is nonzero at $z \in S^2$, then we can find a deformation Y of J supported near $u(z)$ that pairs nontrivially with η . Thus any η that vanishes on the image of $D\mathcal{F}$ must vanish near z and therefore is identically zero, so that $D\mathcal{F}$ is surjective.

This prompts the following definition:

Definition 4.12. A holomorphic curve $u : S^2 \rightarrow M$ is *simple* if it is non-constant, and does not factor through a branched cover $u = vf$ for some holomorphic v and some $f : S^2 \rightarrow S^2$.

A nonconstant holomorphic curve has finitely many critical points. The jet of a holomorphic curve at a point determines the germ there, so if u is simple, at any point where two local branches intersect their jets differ. It follows that any simple holomorphic curve is injective on an open dense set (in fact, the complement of a finite set).

Let $\mathcal{M}_s(A, J)$, $\mathcal{M}_s(A)$, $\mathcal{M}_s^l(A)$ and so on denote the space of *simple* holomorphic curves (of whatever kind being considered). The space $\mathcal{M}_s^l(A)$ is evidently open and by the argument above and the implicit function theorem,

Proposition 4.13. *The space $\mathcal{M}_s^l(A)$ is a smooth Banach manifold.*

For details see [12], Prop. 3.4.1.

Now, the projection $\pi : \mathcal{M}_s^l(A) \rightarrow \mathcal{J}^l$ is a smooth map between smooth Banach manifolds. The derivative $d\pi(u, J) : T_{(u, J)}\mathcal{M}_s^l(A) \rightarrow T_J\mathcal{J}^l$ is projection $(\xi, Y) \rightarrow Y$, so the kernel of $d\pi$ is equal to the kernel of D_u . Thus $d\pi(u, J)$ is Fredholm with the same index as D_u and is onto when D_u is onto. Hence the set of regular J is precisely the set of regular values for π .

By Sard–Smale, the set of regular values of a Fredholm map between separable Banach manifolds is a countable intersection of open and dense subsets (such a set is said to be *residual* in the sense of the Baire category theorem). This is true in \mathcal{J}^l in the C^l topology for each l ; an argument of Taubes ([12], § 3.4) shows that it is true in \mathcal{J} in the C^∞ topology. In other words,

Theorem 4.14. *There is a residual subset $\mathcal{J}_{reg} \subset \mathcal{J}$ so that for any $J \in \mathcal{J}_{reg}$ the space $\mathcal{M}_s(A, J)$ is a smooth oriented manifold of dimension equal to the index of D_u .*

A similar argument applied to the path space of M shows that for any two regular J_1, J_2 there are a residual set of smooth paths in \mathcal{J} whose preimage in $\mathcal{M}_s(A)$ is a smooth oriented cobordism between $\mathcal{M}_s(A, J_1)$ and $\mathcal{M}_s(A, J_2)$.

4.9. Computation of the index. Index of Fredholm operators is a homotopy invariant. Any unitary connection on a complex vector bundle over a Riemann surface determines a holomorphic structure with $\nabla^{0,1} = \bar{\partial}$, so D_u has the same index as $\bar{\partial}_J$. For a holomorphic line bundle L over a genus g Riemann surface, the (complex) index (which is half the real index) is $\deg(L) + 1 - g$ by the Riemann-Roch formula. The degree of a line bundle over a Riemann surface Σ is equal to $c_1(L)([\Sigma])$. Thus, by the splitting principle, the (real) index of a holomorphic bundle E over Σ is $\dim_{\mathbb{C}}(E)(2 - 2g) + 2c_1(E)([\Sigma])$.

Lemma 4.15. *If (M, ω) is a symplectic manifold of dimension $2n$, and $u \in \mathcal{M}(A, J)$ is a J -holomorphic sphere in the homology class A , the index of D_u (and therefore the dimension of $\mathcal{M}(A, J)$ near u if u is simple and J is regular) is $2n + 2c_1(M)([A])$.*

4.10. Compactness. The space $\mathcal{M}(A, J)$ is almost never compact, since the orbit of the automorphism group $\text{Aut}(S^2)$ on any simple curve is proper and noncompact. But even after we quotient out by this action the space $\overline{\mathcal{M}}(A, J)$ can fail to be compact. Let's understand this in some simple examples.

Example 4.16 (Rational maps). A degree d rational map $f : \hat{\mathbb{C}} \rightarrow \hat{\mathbb{C}}$ has a graph Γ_f which is a holomorphic curve in the Kähler manifold $\hat{\mathbb{C}} \times \hat{\mathbb{C}}$ representing the homology class $A := [\hat{\mathbb{C}} \times \text{point}] + d[\text{point} \times \hat{\mathbb{C}}]$.

For simplicity let's suppose $f(\infty) = 1$ so that f is determined by its divisor $\text{div} f$ (supported in \mathbb{C}) which can be written in terms of zeros and poles as $Z - P$ where each of Z and P is an unordered list of d complex numbers. Repetitions are allowed, but Z and P are disjoint. Some of the noncompactness of the space of rational maps arises in families

where elements of Z and P collide. Let $a \in Z$ and $b \in P$ so that f has a factor of the form $(z - a)/(z - b)$. When $|a - b| = \epsilon$ is very small and $|z - a| > \sqrt{\epsilon}$, this factor differs from 1 by $\sim \sqrt{\epsilon}$. If we deform f in a family $f(t)$ by adjusting a so that $a(t) \rightarrow b$, the graphs $\Gamma_{f(t)}$ converge in the Hausdorff topology to a nodal curve $\Gamma_g \cup b \times \hat{\mathbb{C}}$, the union of the holomorphic graph Γ_g (where g of degree $(d - 1)$ is obtained from f by canceling the $(z - a)/(z - b)$ factor) and the ‘vertical’ curve $b \times \hat{\mathbb{C}}$. One calls this phenomenon *bubbling off*, where $b \times \hat{\mathbb{C}}$ is the ‘bubble’.

Here is another way to look at the local picture of the degeneration $\Gamma_{f(t)} \rightarrow \Gamma_g \cup b \times \hat{\mathbb{C}}$ near the singular point $(b, g(b))$. For $\epsilon \in \mathbb{C}$ let S_ϵ be the curve $z_1 z_2 = \epsilon$ in \mathbb{C}^2 . When $\epsilon \neq 0$ the curve S_ϵ is an annulus, with a holomorphic isomorphism to $\mathbb{C} - 0$ given by projection to either the z_1 or the z_2 axis. As $\epsilon \rightarrow 0$ the S_ϵ degenerate to S_0 which (ignoring the embedding) is homeomorphic to the quotient space obtained from an annulus by pinching a meridian curve to a point.

It will turn out that the space of (unparameterized) holomorphic curves in a fixed homology class can be compactified by adding such nodal curves (more commonly called *cusplike curves* in the symplectic topology literature). Here we may allow curves of any (fixed) genus, so throughout this section we do not assume that the genus is zero.

If S is a smooth surface and γ_i is a collection of disjoint simple loops, we obtain a singular surface \bar{S} by collapsing each γ_i to a point p_i . Let \hat{S} be the end completion of $\bar{S} - \cup p_i = S - \cup \gamma_i$. Then \hat{S} is a closed surface, and there is a canonical map $\hat{S} \rightarrow \bar{S}$ which is the identity on $\bar{S} - \cup p_i$, and is 2–1 on the preimage of each p_i . A complex structure on \hat{S} makes \bar{S} into a *nodal Riemann surface*. Equivalently, a complex structure on \bar{S} is just a complex structure on $S - \cup \gamma_i$ for which the modulus of every annular end is infinite.

If M is a smooth manifold with almost-complex structure J and if \bar{S} is a nodal Riemann surface, then $\bar{u} : \bar{S} \rightarrow M$ is (J -)holomorphic if it is holomorphic on $\bar{S} - \cup p_i$ in the usual sense. Such holomorphic curves are called *cusplike curves*.

Definition 4.17. A sequence of holomorphic curves $u_n : S_n \rightarrow M$ is said to *converge weakly* to a cusplike curve $\bar{u} : \bar{S} \rightarrow M$ if

- (1) the areas of $u_n(S_n)$ converge to the area of $\bar{u}(\bar{S})$; and
- (2) there are families of disjoint loops $\gamma_{n,i} \subset S_n$ and diffeomorphisms $\phi_n : \bar{S} - \cup p_i \rightarrow S_n - \cup \gamma_{n,i}$ so that the maps $u_n \phi_n$ converge uniformly to \bar{u} on compact subsets of $\bar{S} - \cup p_i$.

We may now state Gromov’s Compactness Theorem ([8], Thm. 1.5.B):

Theorem 4.18 (Compactness). *Let M be a smooth closed manifold with an almost-complex structure J and compatible metric g (i.e. an almost-Hermitian manifold). Then any sequence of holomorphic curves of uniformly bounded genus and uniformly bounded area has a subsequence which converges weakly to a cusplike curve.*

The proof of this theorem shall take up the rest of this subsection and the next. First we observe that there is a uniform upper bound on the curvature of a holomorphic curve.

Lemma 4.19. *For any holomorphic curve S the sectional curvature K_S satisfies $K_S \leq K$ pointwise where K is the maximum of the sectional curvature of M .*

Proof. Holomorphic curves are minimal for a compatible metric. Gauss's equation for a minimal surface says $K_S = K_M - |\mathbb{I}|^2/2$ where K_M is the sectional curvature of M along the tangent plane to S , and \mathbb{I} is the second fundamental form. \square

Surfaces with an upper curvature bound satisfy an isoperimetric inequality that can be expressed in terms of the geometry of *comparison surfaces*.

Definition 4.20 (Comparison Surface). Let S be a Riemannian surface. A rotationally symmetric Riemannian surface S_0 is a *comparison surface* for S if, for any domain $E \subset S$, if $E_0 \subset S_0$ is a rotationally symmetric disk with the same area, then $\text{length}(\partial E) \geq \text{length}(\partial E_0)$.

Example 4.21. If S has curvature bounded above by K and area bounded above by $4\pi/K$, the sphere S_K^2 of constant curvature K is a comparison surface for S .

Lemma 4.22. *Let S be a holomorphic curve, and let S_0 be a comparison surface for S . Let $u : \mathbb{D} \rightarrow S$ be an injective holomorphic map with area A and let $u_0 : \mathbb{D} \rightarrow S_0$ be a comparison map — i.e. a conformal map onto a symmetric disk with the same area. Then $|du(0)| \leq |du_0(0)|$.*

Proof. For any $r \leq 1$ let $\mathbb{D}_r \subset \mathbb{D}$ be the disk of radius r . Let $A(r) := \text{area}(u(\mathbb{D}_r))$ and $A_0(r) := \text{area}(u_0(\mathbb{D}_r))$ and likewise $L(r) := \text{length}(u(\partial\mathbb{D}_r))$ and $L_0(r) := \text{length}(u_0(\partial\mathbb{D}_r))$. By the Cauchy–Schwarz inequality,

$$A'(r) = \int_{\partial\mathbb{D}_r} |du|^2 \geq \frac{1}{2\pi r} \left(\int_{\partial\mathbb{D}_r} |du| \right)^2 = \frac{L(r)^2}{2\pi r}$$

whereas for u_0 we have (by the same calculation) equality: $A_0'(r) = L_0^2(r)/2\pi r$.

Let $\alpha : [0, 1] \rightarrow [0, 1]$ be the function for which $A(r) = A_0(\alpha(r))$ so that $A'(r) = \alpha'(r)A_0'(\alpha(r))$ for any r . Then the isoperimetric inequality gives $L(r) \geq L_0(\alpha(r))$ so

$$A'(r) \geq \frac{L_0^2(\alpha(r))}{2\pi r} = \frac{\alpha(r)}{r} \frac{L_0^2(\alpha(r))}{2\pi\alpha(r)} = \frac{\alpha(r)}{r} A_0'(\alpha(r)) = \frac{\alpha(r)}{r\alpha'(r)} A'(r)$$

so we get $r\alpha'(r) \geq \alpha(r)$ and by integrating $\log(r) - \log(\epsilon) \leq \log(\alpha(r)) - \log(\alpha(\epsilon))$ for any positive ϵ . But $\alpha(1) = 1$ so $\alpha(\epsilon) \leq \epsilon$ for any $\epsilon \leq 1$ and therefore $\alpha'(0) \leq 1$. Since $|du(0)| = \alpha'(0)|du_0(0)|$ the lemma follows. \square

This Lemma and the isoperimetric inequality in Example 4.21 lets us control the derivative at 0 of a holomorphic disk with area less than $4\pi/K$. Remarkably, it is possible to trade this area bound for a *diameter* bound. This is the so-called Gromov–Schwarz Lemma:

Theorem 4.23 (Gromov–Schwarz Lemma). *There exist constants $\epsilon > 0$ and $C > 0$ so that any holomorphic map $u : \mathbb{D} \rightarrow M$ whose image is contained some ball $B_\epsilon(p)$ of radius ϵ is C -Lipschitz with respect to the hyperbolic metric on \mathbb{D} and the g -metric on M .*

Proof. By compactness of M , there are positive constants $C_1 > 0$ and $\epsilon > 0$ so that on any ball of radius ϵ we can find a 1-form β with $\|\beta\| \leq C_1$ for which $d\beta = \omega$. For any

$u : \mathbb{D} \rightarrow B_\epsilon(p)$ we have

$$\text{area}(u(\mathbb{D})) = \int_{u(\mathbb{D})} \omega = \int_{u(\partial\mathbb{D})} \beta \leq C_1 \text{length}(u(\partial\mathbb{D}))$$

We may therefore construct a comparison surface S_0 , valid for *all* holomorphic curves with image contained in any ball of radius $\leq \epsilon$, by inserting a round bubble of positive curvature K into a hyperbolic disk scaled to have constant curvature $-1/C_1^2$. By Lemma 4.22 a comparison map $u_0 : \mathbb{D} \rightarrow S_0$ has $|du(0)| \leq |du_0(0)|$. On the other hand, S_0 is bi-Lipschitz to the ordinary hyperbolic plane so by the (usual) Schwarz Lemma we get a uniform bound on $|du(0)|$ independent of u . \square

4.11. Completion of the proof. The Gromov–Schwarz Lemma and a bootstrap argument allows us to promote C^0 convergence of holomorphic curves to C^∞ convergence (this is Lemma 4.24). The bootstrap argument treats the (higher) jets of a holomorphic curve as holomorphic curves in their own right in the space of *bundle maps*.

This requires a few words of explanation. If we fix S and M , there is a complex vector bundle E over $S \times M$ with fiber over a point (p, q) equal to $\text{Hom}_{\mathbb{C}}(T_p S, T_q M)$. The (almost)-complex structures on S and M determine an almost complex structure J_E on E , and if $u : S \rightarrow M$ is a holomorphic curve in M then evidently the graph of (u, du) (i.e. the 1-jet of u) is a holomorphic curve in E with respect to J_E . The symplectic structure on M and an area form on S together give a symplectic structure on $S \times M$; if we pull this back to E and add a fiberwise symplectic form compatible with the fiberwise complex structure we obtain a symplectic structure on E compatible with J_E . For details see e.g. [11], Chapter 3.

Lemma 4.24. *If $u_n : S \rightarrow M$ is a sequence of holomorphic maps converging in the C^0 topology to $u : S \rightarrow M$ then in fact the u_n converge in the C^∞ topology and the limit is a holomorphic map.*

Proof. Let $p \in S$ and let D be a neighborhood of p in S for which $u_n(D)$ is contained in the ϵ -neighborhood of $u(p)$. Then the u_n are uniformly Lipschitz on D in the hyperbolic metric. Thus the u_n take a sufficiently small neighborhood of the zero section in TD to a small neighborhood of the zero section in $TB_\epsilon(u(p))$.

Thus the 1-jets of u_n have relatively compact image in the space of bundle maps. Applying the Gromov–Schwarz Lemma to the 1-jets of the u_n gives control on the norms of the 2-jets, and by induction on all higher derivatives. \square

It is also important to be able to extend holomorphic curves over punctures:

Lemma 4.25. *Let $u : S - p \rightarrow M$ be holomorphic with relatively compact image. If either*

- (1) *area($u(S - p)$) is finite; or*
- (2) *$u(S - p)$ is contained in $B_\epsilon(q)$ for some $q \in M$*

then u extends to a holomorphic map $\bar{u} : S \rightarrow M$.

Proof. By restricting to a subset of S if necessary we may assume $S = \mathbb{D}$ and $p = 0$. In the hyperbolic metric on $\mathbb{D} - 0$ a neighborhood of 0 has arbitrarily small area, so by Gromov–Schwarz the second case reduces to the first.

As in the proof of Lemma 4.22 we use the notation $A(r) := \text{area}(u(\mathbb{D}_r - 0))$ and $L(r) := \text{length}(u(\partial\mathbb{D}_r))$. Since $\text{area}(u(\mathbb{D} - 0))$ is finite and $u : S \rightarrow M$ is conformal it follows exactly as in Lemma 4.22 that

$$\infty > A(r) \geq \int_0^r \frac{L(s)^2}{2\pi s} ds$$

and therefore there is a sequence of radii $r_j \rightarrow 0$ with $L(r_j) \rightarrow 0$. By relative compactness of the image in M we may extract a subsequence of radii so that $u(\partial\mathbb{D}_{r_j}) \rightarrow q$. Suppose that there is a sequence of points $w_j \in \mathbb{D}_{r_j} - \mathbb{D}_{r_{j+1}}$ with a subsequence converging to $q' \neq q$ and let $\epsilon = d(q', q)/2$. Then for each index j in the subsequence, $u(\mathbb{D}_{r_j} - \mathbb{D}_{r_{j+1}})$ contains a point $u(w_j)$ arbitrarily close to q' , but its boundary lies outside $B_\epsilon(q')$. By the monotonicity formula for minimal surfaces there is a uniform positive lower bound on the area of each $u(\mathbb{D}_{r_j} - \mathbb{D}_{r_{j+1}})$. But this implies $\text{area}(u(\mathbb{D} - 0))$ is infinite, contrary to assumption.

Thus u extends to a continuous map \bar{u} , and by Gromov–Schwarz u is uniformly Lipschitz in the hyperbolic metric. Thus the map on 1-jets has finite area and relatively compact image on $\mathbb{D}_r - 0$ and by induction \bar{u} has continuous partial derivatives of all orders and is therefore holomorphic. \square

We are now ready to conclude the proof of Theorem 4.18.

Proof. Fix small constants $A < 2\pi/K$ and $\epsilon > 0$ so that every minimal surface F in M intersects every ball $B_\epsilon(p)$ with $p \in F$ in a subsurface of area at least A . If $u : S \rightarrow M$ is holomorphic we can find a maximal subset of points $Q \subset S$ so that the ϵ -balls about the points of $u(Q)$ are disjoint (and therefore also the 2ϵ -balls about $u(Q)$ cover $u(S)$). Then $|Q| \leq \text{area}(u(S))/A$ and (if we take ϵ small enough) also $|Q| \geq 3$. Thus under the assumptions of the theorem $S - Q$ is hyperbolic of uniformly bounded area and by Gromov–Schwarz, the norm of du in the hyperbolic metric is uniformly bounded on $S - Q$, independent of u .

If $u_n : S_n \rightarrow M$ is a sequence of holomorphic curves, and $Q_n \subset S_n$ points as above, then either the hyperbolic metrics on $S_n - Q_n$ have a convergent subsequence in some moduli space, or there is a subsequence for which these metrics degenerate by stretching necks centered at boundedly many essential simple closed loops $\cup \gamma_{n,i}$. Then $u_n : S_n - Q_n - \cup \gamma_{n,i}$ converge on some subsequence to $u : \bar{S} - Q - P$ for some nodal Riemann surface \bar{S} and for some finite collections of points Q, P . By Lemma 4.25 this map extends to a cusp curve $\bar{u} : \bar{S} \rightarrow M$, and by the equicontinuity of the u_n in the hyperbolic metrics on $S_n - Q_n$ the areas of the $u_n(S_n)$ converge to the area of $\bar{u}(\bar{S})$. \square

This completes the proof of the Compactness theorem.

One immediate application is as follows.

Corollary 4.26. *Let M, ω be symplectic, and let $A \in H_2(M)$ be a spherical class. Suppose there is no spherical class B with $0 < \omega(B) < \omega(A)$. Then for any ω -compatible almost-complex structure J , the unparameterized moduli space $\overline{\mathcal{M}}(A, J)$ is compact.*

Proof. Any cusp curve $\bar{u} : \bar{S} \rightarrow M$ compactifying $\overline{\mathcal{M}}(A, J)$ has domain a nodal Riemann surface \bar{S} whose irreducible components S_j (of which there are at least 2) are 2-spheres. But $\omega(A) = \text{area}(\bar{u}(\bar{S})) = \sum_j \text{area}(\bar{u}(S_j)) = \sum_j \omega(A_j)$ where A_j is the (spherical) homology class represented by $\bar{u} : S_j \rightarrow M$. Since every $\omega(A_j) > 0$ we violate the hypothesis. \square

4.12. **The evaluation map.** The group of orientation-preserving conformal automorphisms of S^2 is $\mathrm{PSL}(2, \mathbb{C})$ which acts by linear fractional automorphisms after identifying S^2 with the Riemann sphere. This gives rise to an action of $\mathrm{PSL}(2, \mathbb{C})$ on the product $\mathcal{M}(A, J) \times S^2$ by $g \cdot (u, z) = (ug^{-1}, g(z))$.

Definition 4.27. Let $\mathcal{W}(A, J)$ denote the quotient of $\mathcal{M}(A, J) \times S^2$ by the action of $\mathrm{PSL}(2, \mathbb{C})$. This is the domain of the *evaluation map*

$$e : \mathcal{W}(A, J) \rightarrow M, \quad e(u, z) = u(z)$$

Say that a class $A \in H_2(M)$ is *spherical* if it is in the image of the Hurewicz map $\pi_2(M) \rightarrow H_2(M)$. Recall: if M, ω is symplectic then \mathcal{J} denotes the space of smooth almost-complex structures J compatible with ω , and $\mathcal{J}_{\mathrm{reg}}$ denotes the space of regular compatible J . Then with this terminology we have the following theorem:

Theorem 4.28. *Let M, ω be symplectic, and let $A \in H_2(M)$ be a spherical class. Suppose there is no spherical class $B \in H_2(M)$ with $0 < \omega(B) < \omega(A)$. Then:*

- (1) *for any $J \in \mathcal{J}$ the space $\overline{\mathcal{M}}(A, J)$ is compact;*
- (2) *for any $J \in \mathcal{J}_{\mathrm{reg}}$ the space $\mathcal{W}(A, J)$ is a smooth oriented compact manifold of dimension $2n + 2c_1(M)([A]) - 4$;*
- (3) *for any $J_1, J_2 \in \mathcal{J}_{\mathrm{reg}}$ the spaces $\mathcal{W}(A, J_1)$ and $\mathcal{W}(A, J_2)$ are oriented cobordant and their images under e in M are homologous.*

Proof. First we show that the unparameterized moduli space $\overline{\mathcal{M}}(A, J)$ is compact. Suppose not, and suppose $u_j : S^2 \rightarrow M$ is any sequence of holomorphic curves weakly converging to a nontrivial cusp curve $\bar{u} : \bar{S} \rightarrow M$. The irreducible components S_k of \bar{S} are spheres whose images represent spherical homology classes B_k in M with

$$\sum \omega(B_k) = \sum \mathrm{area}(\bar{u}(S_k)) = \mathrm{area}(u_j(S^2)) = \omega(A)$$

thereby violating the hypothesis.

Second, the hypothesis of the theorem implies that A is primitive and nonzero, from which it follows that any J -holomorphic curve is simple. Thus $\mathcal{M}_s(A) = \mathcal{M}(A)$ and for any regular J , the space $\mathcal{M}(A, J)$ is a smooth manifold, and therefore so is $\mathcal{M}(A, J) \times S^2$. The group $\mathrm{PSL}(2, \mathbb{C})$ acts on this space freely and properly so the quotient space $\mathcal{W}(A, J)$ is a smooth oriented manifold of dimension

$$\dim(\mathcal{M}(A, J)) + \dim(S^2) - \dim(\mathrm{PSL}(2, \mathbb{C})) = 2n + 2c_1(M)([A]) - 4$$

Furthermore, $\mathcal{W}(A, J)$ is compact, because it is an S^2 bundle over $\overline{\mathcal{M}}(A, J)$.

Lastly, suppose J_t is a smooth path of almost complex structure interpolating between J_1 and J_2 whose preimage $\mathcal{M}(A, J_t)$ in $\mathcal{M}(A)$ (which is equal to $\mathcal{M}_s(A)$) is a smooth oriented cobordism between $\mathcal{M}(A, J_1)$ and $\mathcal{M}(A, J_2)$. Then we can form $\mathcal{M}(A, J_t) \times S^2$ and $\mathcal{W}(A, J_t)$ in the obvious way, and observe that this last space is a smooth oriented compact cobordism between $\mathcal{W}(A, J_1)$ and $\mathcal{W}(A, J_2)$ realizing a homology between their images under the evaluation map. \square

4.13. The Nonsqueezing Theorem. We are now ready to state and prove the Nonsqueezing Theorem. For any $r > 0$ let $B^2(r)$ denote the ball of radius r in symplectic \mathbb{R}^2 , and let $Z(r) \subset \mathbb{R}^{2n}$ denote the product of $B^2(r)$ with symplectic \mathbb{R}^{2n-2} . In other words,

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

Theorem 4.29 (Gromov's Nonsqueezing Theorem). *Let \mathbb{R}^{2n}, ω denote standard symplectic \mathbb{R}^{2n} . Let $B^{2n}(1) \subset \mathbb{R}^{2n}$ denote the unit ball. If there is a symplectomorphism ψ from $B^{2n}(1)$ into $Z(r)$ then $r \geq 1$.*

Proof. Suppose $\psi : B^{2n}(1) \rightarrow Z(r)$ is symplectic. Its image is contained in some compact set $B^2(r) \times K$ which may be symplectically embedded in the product $M := S^2 \times T^{2n-2}$ where S^2 is a symplectic sphere of area $\pi r^2 + \epsilon$, and T^{2n-2} is some Euclidean flat torus big enough to contain a translate of K . By abuse of notation we denote the symplectic form on M by ω .

Let J_0 denote the standard complex structure on $\mathbb{R}^{2n} = \mathbb{C}^n$ and let J be any ω -compatible almost-complex structure on M that restricts to $\psi_*(J_0)$ on $\psi(B^{2n}(1))$. We claim there is at least one (possibly singular) J -holomorphic sphere through each point of M in the homology class $A = [S^2 \times \text{point}]$. Let's see how the theorem follows from this claim.

Let S be such a sphere passing through $\psi(0)$. The preimage $\psi^{-1}(S)$ is a proper minimal surface in $B^{2n}(1)$ (in the Euclidean metric) passing through 0; thus by the monotonicity formula, $\text{area}(\psi^{-1}(S)) \geq \pi$. Since ψ is symplectic it is area preserving; thus $\text{area}(\psi^{-1}(S)) \leq \text{area}(S)$. But then

$$\pi \leq \text{area}(\psi^{-1}(S)) \leq \text{area}(S) = \omega(A) = \pi r^2 + \epsilon$$

Now let's prove the claim. Let J_{split} denote the product complex structure on M . Any map $S^2 \rightarrow S^2 \times T^n$ lifts to $S^2 \rightarrow S^2 \times \mathbb{R}^n$ so a holomorphic map must be constant on the second factor, and $\overline{\mathcal{M}}(A, J_{\text{split}})$ consists exactly of horizontal curves $S^2 \times \text{point}$. Thus $e : \mathcal{W}(A, J_{\text{split}}) \rightarrow M$ is a diffeomorphism. Since $\mathcal{W}(A, J_{\text{split}})$ is a smooth manifold of dimension equal to the formal dimension $2n + 2c_1(S^2) - 4 = 2n$ it follows that $J_{\text{split}} \in \mathcal{J}_{\text{reg}}$. In particular, $e : \mathcal{W}(A, J') \rightarrow M$ represents the fundamental class $[M]$ (and is therefore surjective) for every $J' \in \mathcal{J}_{\text{reg}}$. Since the space of regular almost-complex structures is dense in \mathcal{J} and path-connected, there are holomorphic curves in the class of A through every point of M for a sequence of almost-complex structures converging to J , and by the compactness theorem, there is some (possibly singular) holomorphic curve through every point for the almost-complex structure J . This proves the claim and also the theorem. \square

4.14. Dimension 4. Holomorphic curves in dimension 4 satisfy extra rigidity properties arising from properties of the intersection form on homology. A compact symplectic 4-manifold X^4 is oriented, and there is a symmetric nondegenerate unimodular intersection form on $H_2(X)$ Poincaré dual to the cup product on $H^2(X)$. If $A, B \in H_2(X)$ are represented by smooth oriented surfaces S_A, S_B in general position, then $A \cdot B$ is equal to the signed count of intersections of S_A with S_B .

Holomorphic curves are typically neither nonsingular nor in general position with respect to each other, and therefore we must invoke Theorem 4.8 which describes the local structure of (possibly singular) holomorphic curves and their (possibly self-) intersections. Let's make the following definitions:

Definition 4.30 (Local intersection number). Let $u : \Sigma \rightarrow X^4$ be holomorphic (not necessarily connected) with $p_j \in u^{-1}(x)$ for $j = 1, 2$. Write $\psi u \phi_j^{-1}(z) = (z^{Q_j}, f_j(z))$ as in Theorem 4.8. Let Q be the least common multiple of Q_1, Q_2 and write $Q = m_j Q_j$. Then we may define the *local intersection number*

$$\delta(p_1, p_2) := \frac{1}{m_1 m_2} \sum_{\nu^{Q=1}} \text{ord}_0 (f_1(\nu z^{m_1}) - f_2(z^{m_2}))$$

where ord_0 means order of vanishing at 0.

Likewise we may define:

Definition 4.31 (Local degree). With notation as above let x be an (isolated) singular point of u , and $p \in u^{-1}(x)$ a preimage. Write $\psi u \phi^{-1}(z) = (z^Q, f(z))$ for some $Q > 1$ where $f(z)$ vanishes at 0 to order $Q' \geq Q$. Define the *local degree*

$$\delta(p) := \sum_{\nu^{Q=1}, \nu \neq 1} \text{ord}_0 \left(\frac{f(\nu z) - f(z)}{z} \right)$$

Example 4.32. For a single curve $u : \Sigma \rightarrow X$ we may define

$$\delta(u) = \sum_{(p_1, p_2)} \delta(p_1, p_2) + \sum_p \delta(p)$$

where the sum is taken over all singularities and multiple points. The quantity $\delta(u)$ measures the difference $\chi(u(\Sigma)) - \chi(\Sigma')$ where Σ' is a suitable ‘desingularization’ of $u(\Sigma)$.

This is easiest to explain in the integrable case. Consider the (projective) elliptic curve given in an affine chart by the equation $y^2 = x(x-a)(x-b)$. When $0, a, b$ are distinct, this elliptic curve is a nonsingular torus and has $\chi = 0$. The curve $y^2 = x^3 + x^2$ is genus zero and has a transverse double point at 0. Thus $\delta(u) = 1$ and $\chi = 1$. The curve $y^2 = x^3$ is genus zero and embedded (though not smoothly at the cusp point 0) so $\delta(u) = 0$ and $\chi = 0$.

In every case $\delta(p_1, p_2)$ is a positive integer $\geq Q_1 Q_2$. Likewise $\delta(p)$ is always $\geq (Q - 1)(Q' - 1)$. In fact, since $\delta(p)$ counts the contribution to χ from a change of genus, it is always *even*.

Using these definitions we may give formulae for the intersection product in terms of geometry for holomorphic curves.

Lemma 4.33 (Intersection Formula). *Let $u_j : \Sigma_j \rightarrow X$ be simple holomorphic curves in homology classes $A_1, A_2 \in H_2(X)$. Then either the u_j both have the same image, or*

$$A_1 \cdot A_2 = \sum_{(p_1, p_2)} \delta(p_1, p_2)$$

In particular, $A_1 \cdot A_2 \geq 0$ with equality if and only if the curves are disjoint.

Proof. If the curves are nonsingular and the intersections are transverse, this is equivalent to saying that every intersection is positive. This is because the tangent spaces are complex subspaces of TX . The general case follows from Theorem 4.8. \square

Lemma 4.34 (Adjunction Formula). *Let $u : \Sigma \rightarrow X^4$ be a simple holomorphic curve in the homology class $A \in H_2(X)$. Then*

$$c_1(X)(A) = \chi(\Sigma) + A \cdot A - 2 \sum_{(p_1, p_2)} \delta(p_1, p_2) - 2 \sum_p \delta(p)$$

Proof. For simplicity we assume the singularities of u are transverse double points. We have $u^*TX = T\Sigma \oplus \nu$ where ν is the normal bundle, so $c_1(u^*TX) = c_1(T\Sigma) + c_1(\nu)$ in $H^2(\Sigma)$. Thus $c_1(X)(A) = \chi(\Sigma) + c_1(\nu)[\Sigma]$.

Now, $c_1(\nu)[\Sigma]$ is the number of intersections of Σ with a push-off of itself in $u^*\nu$, whereas A is the number of intersections of $u(\Sigma)$ with a push-off of itself; the difference between these numbers is twice the number of double points.

The general case follows from Theorem 4.8. \square

For brevity in the sequel we denote the contributions δ from all singularities of u by $\delta(u)$. As corollaries of the Adjunction Formula we obtain:

Corollary 4.35 (Dimension count). *Let $A \cdot A \leq -2$. Then for regular J the space $\mathcal{M}(A, J)$ is empty. Likewise, if $A \cdot A = -1$, for regular J the space $\overline{\mathcal{M}}(A, J)$ is 0-dimensional and its points correspond to embedded smooth curves.*

Proof. By the Adjunction Formula, the dimension of $\mathcal{M}(A, J)$ for regular J is

$$\dim \mathcal{M}(A, J) = 4 + 2c_1(X)(A) = 8 + 2A \cdot A - 4\delta(u)$$

where $\delta(u) > 0$ unless u is embedded and nonsingular. Since $\dim \overline{\mathcal{M}} = \dim \mathcal{M} - 6$ the proof follows. \square

Corollary 4.36 (Smooth embeddedness persists). *If u, u' are J, J' holomorphic spheres in the same homology class A , then $\delta(u) = \delta(u')$. In particular*

- (1) *if for some J there is some u which is a smooth embedding then for every J' every u' is a smooth embedding;*
- (2) *if for some J there is some u which is smooth with exactly one transverse double point then for every J' every u is smooth with exactly one transverse double point.*

Proof. By the Adjunction Formula $\delta(u) = \delta(u')$ because the other terms depend only on the homology class A . \square

Another phenomenon special to dimension 4 is that one can show $\mathcal{M}(A, J)$ is a smooth manifold near a nonsingular embedded curve u under purely homological conditions on A :

Lemma 4.37 (Automatic Regularity). *Let $u : S^2 \rightarrow X^4$ be a smoothly immersed simple holomorphic curve. Then D_u is onto (so that \mathcal{M} is a smooth oriented manifold of the correct dimension near u) if and only if $c_1(X^4)(A) \geq 1$.*

Proof. First let's suppose that J is integrable. Then u^*TX is a holomorphic \mathbb{C}^2 bundle over S^2 , which splits (because u is nonsingular) as a sum of holomorphic line bundles $TS^2 \oplus \nu$. The cokernel of D_u is therefore isomorphic to the space of holomorphic sections of $\nu^* \otimes K$. But $c_1(X^4)(A) = c_1(TS^2) + c_1(\nu)$ so $c_1(\nu) \geq -1$ so $c_1(\nu^* \otimes K) \leq -1$ and therefore D_u is onto.

If J is not integrable there is an argument due to Hofer–Lizan–Sikorav [10] that we summarize. Showing that D_u is onto is equivalent to showing that the adjoint D_u^* has trivial kernel. Ignoring the tangential part of u^*TX , we may write $D_u^* = \bar{\partial} + a$ for some $a \in \Omega^{0,1}(\text{End}_{\mathbb{R}}\nu^* \otimes K)$ and we want to show that if $\bar{\partial}f + af = 0$ then (under the homological condition on c_1) we have $f = 0$.

The operator D_u^* is not necessarily complex linear because a isn't. Thus the first step is to replace a by some $b \in \Omega^{0,1}$ so that $Lf = 0$ where $L = \bar{\partial} + b$. This is elementary: for $z \in S^2$ and $v \in T_z S^2$ we may just take $b(z)v = (a(z)v)f(z)/f(z)$ where $f(z) \neq 0$ and $b(z)v = 0$ where $f(z) = 0$. Now the operator L is complex-linear, and therefore defines a complex connection on $\nu^* \otimes K$. Any complex connection on a bundle over a Riemann surface is integrable, so L defines a new *holomorphic* structure on $\nu^* \otimes K$. Anything in the kernel of L would be a holomorphic section of this bundle, which must vanish because the degree is negative. Thus the kernel of L is trivial and so therefore is the kernel of D_u^* . More work is needed to apply this argument in the weaker regularity in the setting of § 4.7. \square

Remark 4.38. If $u : S^2 \rightarrow M^{2n}$ is nonsingular and holomorphic then u^*TM always splits as $TS^2 \oplus \nu$. If J is integrable, ν is holomorphic and (by a theorem of Grothendieck) splits as a sum of holomorphic line bundles whose first Chern classes sum to $c_1(\nu)$. If every holomorphic summand L of ν has $c_1(L) \geq -1$ then (exactly as above) D_u is onto and u is a smooth point on $\mathcal{M}(A, J)$. However, even if J is integrable, the Chern classes of the summands of ν may vary from point to point in $\mathcal{M}(A, J)$, and D_u may be surjective at some points but not others.

4.15. Symplectic rigidity of \mathbb{R}^4 . Putting together the results from the previous section, we deduce the following remarkable theorem of Gromov:

Theorem 4.39 (Symplectic rigidity of \mathbb{R}^4). *Let (X^4, ω) be a symplectic manifold with $H_*(X) = H_*(\text{point})$. Suppose there exist compact subsets $K \subset X$ and $L \subset \mathbb{R}^4$ and a symplectomorphism*

$$\phi : (X - K, \omega) \rightarrow (\mathbb{R}^4 - L, \omega_{std})$$

Then there is a compact set $K' \subset X$ and a symplectomorphism

$$\psi : (X, \omega) \rightarrow (\mathbb{R}^4, \omega_{std})$$

so that ψ agrees with ϕ on $X - K'$.

Remark 4.40. Actually, the hypothesis $H_*(X) = H_*(\text{point})$ is superfluous; all that is needed is that the image of $\pi_2(X)$ in $H_2(X; \mathbb{R})$ is trivial.

First we prove a lemma:

Lemma 4.41. *Let (Y, ω) be a compact symplectic 4-manifold whose homology satisfies $H_*(Y) = H_*(S^2 \times S^2)$. Suppose there are compact subsets $K \subset Y$ and $L \subset \mathbb{R}^2 \times \mathbb{R}^2 \subset S^2 \times S^2$ and a symplectomorphism*

$$\phi : (Y - K, \omega) \rightarrow (S^2 \times S^2 - L, \omega_0)$$

where ω_0 is the product symplectic structure coming from two area forms on S^2 of the same area.

Let J be any ω -compatible almost complex structure on Y . Then Y has a pair of J -holomorphic foliations by embedded holomorphic 2-spheres in the homology classes $A := (0, 1)$ and $B := (1, 0)$ respectively, and each pair of leaves from the different foliations intersect transversely at a single point.

Proof. The symplectomorphism from $Y - K$ to $S^2 \times S^2 - L$ implies several topological facts that we enumerate:

- (1) the classes A and B are spherical;
- (2) $\omega(A) = \omega(B)$;
- (3) $A \cdot A = B \cdot B = 0$ and $A \cdot B = 1$;
- (4) $c_1(Y)(A) = 2$.

These facts all follow by representing A and B by symplectic spheres in $Y - K$ that are carried by ψ to product spheres $a \times S^2$ and $S^2 \times b$ in $S^2 \times S^2 - L$.

Because $H_*(Y) = H_*(S^2 \times S^2)$ the classes A and B generate $H_*(Y)$ and therefore the class A (and likewise B) satisfies the hypothesis of Theorem 4.28 (here it is very important that the two S^2 factors have the same symplectic area!) and therefore $\mathcal{W}_Y(A, J)$ is compact. By the Adjunction Formula,

$$2 = c_1(Y)(A) = 2 - 2\delta(u)$$

so every holomorphic curve $u \in \mathcal{M}_Y(A, J)$ is smooth and embedded.

Therefore by Lemma 4.37 $\mathcal{W}_Y(A, J)$ is smooth of dimension $4 + 2c_1(Y)(A) - 4 = 4$ for every J .

By the intersection formula any two holomorphic curves in $\mathcal{M}_Y(A, J)$ are disjoint. Therefore $\mathcal{W}_Y(A, J) \rightarrow X$ is an *embedding*. So if $\mathcal{W}_Y(A, J)$ is nonempty for some J (which it is e.g. for the product complex structure on the domain of ϕ) it follows that $\mathcal{W}_Y(A, J) \rightarrow Y$ is a diffeomorphism for all J , and Y is foliated by embedded holomorphic 2-spheres.

The same argument applied to B gives another foliation. Since $A \cdot B = 1$, the spheres in either foliation intersect each other transversely in a single point. \square

Now we give the proof of Theorem 4.39.

Proof. By enlarging L (and K) if necessary, we can assume $L = D^2 \times D^2$ for some product of disks of the same (symplectic) area. We can enlarge (L, ω_{std}) to $(S^2 \times S^2, \omega_0)$ and likewise (K, ω) to (Y, ω) with ϕ as in the hypothesis of Lemma 4.41. Let J_{split} be the product complex structure on $S^2 \times S^2$ and let J be an almost complex structure on Y that agrees with ϕ^*J on the complement of K .

By Lemma 4.41 we get a pair of transverse foliations that give Y a product structure by J -holomorphic spheres, and using this we obtain a diffeomorphism $Y \rightarrow S^2 \times S^2$ extending ϕ . By construction the symplectic forms ω and $\phi^*\omega_0$ agree outside K (and are therefore homologous), and are both positive on the spheres of the foliations, so $\omega_t := t\omega + (1 - t)\phi^*\omega_0$ are symplectic and homologous for all t . Therefore by Moser's Theorem they are symplectomorphic, by a symplectomorphism which is the identity outside K . \square

4.16. Symplectic Rigidity of $\mathbb{C}\mathbb{P}^2$. The following theorem combines Theorem 4.39 with deep work of Taubes:

Theorem 4.42 (Symplectic Rigidity of $\mathbb{C}\mathbb{P}^2$). *Let X^4, ω be any symplectic 4-manifold with $H_*(X) = H_*(\mathbb{C}\mathbb{P}^2)$ that contains an embedded symplectic 2-sphere C with $C \cdot C = 1$. Then X^4, ω is symplectomorphic to $\mathbb{C}\mathbb{P}^2$ with a multiple of the standard (Fubini–Study) form ω_0 .*

In particular, if $X^4 = \mathbb{C}\mathbb{P}^2$ then any symplectic form ω is diffeomorphic to a multiple of the standard form.

Proof. Let C be an embedded symplectic 2-sphere C in a homology class A with $A \cdot A = 1$. Evidently A is primitive, so the complement of a regular neighborhood $W := \mathbb{C}\mathbb{P}^2 - N(C)$ is a homology 4-ball with boundary $\partial W = S^3$, the total space of a circle bundle over C with Euler class 1. By the symplectic neighborhood theorem, ω near C looks like the standard symplectic form on $\mathbb{C}\mathbb{P}^2$ near a $\mathbb{C}\mathbb{P}^1$, which is isomorphic to the standard symplectic form on \mathbb{R}^4 on an annulus

$$A := \{x \in \mathbb{R}^4 : \delta - \epsilon < \|x\|^2 \leq \delta\}$$

for suitable $0 < \epsilon < \delta$ (depending on $\omega(C)$). By Theorem 4.39 we may obtain a symplectomorphism ϕ from W to a round ball in \mathbb{R}^4 of radius δ ; gluing $N(C)$ back gives $\mathbb{C}\mathbb{P}^2$ with its standard symplectic structure (up to scale).

A deep theorem of Taubes shows that for any symplectic structure on $\mathbb{C}\mathbb{P}^2$ there is an embedded symplectic 2-sphere with $C \cdot C = 1$. The theorem follows. \square

5. FLOER HOMOLOGY

5.1. The Arnold Conjecture.

5.2. Morse Theory. Let M be a smooth compact n -manifold. A smooth function $f : M \rightarrow \mathbb{R}$ is *Morse* if the critical points (i.e. points p where $df(p)$ is the zero map from $T_p M$ to $T_{f(p)} \mathbb{R} = \mathbb{R}$) are nondegenerate. This means there is a neighborhood U of p and smooth local coordinates x_i on U vanishing at p such that throughout U ,

$$f(x) = f(p) - x_1^2 - \cdots - x_i^2 + x_{i+1}^2 + \cdots + x_n^2$$

for some i , the *index* of the critical point p .

A Morse function is *self-indexing* if $f(p) = i$ for every critical point of index i . Local minima resp. maxima of f are critical points of index 0 resp. n so for a self-indexing Morse function, $f(M) = [0, n]$. It is a fact that every smooth compact manifold admits (many) self-indexing Morse functions.

Let f be such a function. For each $t \in \mathbb{R}$ define $M_t := f^{-1}(-\infty, t]$. Then M_t is empty for $t < 0$ and is equal to M for $t > n$. Furthermore, for any t not equal to $0, 1, \dots, n$, the space M_t is a smooth manifold with boundary $\partial M_t := f^{-1}(t)$.

Choose a smooth metric g on M . The vector field $\text{grad}(f)$ is defined by $\langle \text{grad}(f), Y \rangle = df(Y)$ for all vectors Y . It is perpendicular to the level sets of f and has magnitude at every point equal to that of df . Thus, $\text{grad}(f)$ vanishes only at the critical points, and the flow it generates gives a diffeomorphism from M_s to M_t whenever $i < s < t < t + 1$.

When we transition from M_s to M_t for $s < i < t$ the topology changes in a neighborhood of each critical point by attaching an i -handle $D^i \times D^{n-i}$. An i -handle is thought of as a thickened neighborhood of its core $D^i \times 0$, which is attached to ∂M_s along its boundary $\partial D^i \times 0$. The cocore is $0 \times D^{n-i}$ with boundary $0 \times \partial D^{n-i}$. The core of the i -handle associated to a critical point p is the closure of the set of flowlines of $\text{grad}(f)$ (in a neighborhood

of p) asymptotic to p in the future, while the cocore is the closure of the set of flowlines asymptotic to p in the past.

Thus each M_i has the homotopy type of an i -dimensional CW complex, which is the closure of the union of the flowlines of $\text{grad}(f)$ asymptotic in the future to some critical point of index $\leq i$. For each pair of critical points p, q of indexes $i > j$ let $F(p, q)$ denote the space of flowlines of $\text{grad}(f)$ whose closures run from q to p . For generic f , the space $F(p, q)$ is an open oriented manifold of dimension $i - j - 1$. It may be compactified to a compact oriented manifold with corners $\overline{F}(p, q)$ by adding products $F(p, r_1) \times F(r_1, r_2) \times \cdots \times F(r_k, q)$ for intermediate critical points r_1, \dots, r_k .

When $i - j = 1$ the space $F(p, q)$ is a finite set of points, and when $i - j = 2$ the space $F(p, q)$ is a finite union of circles, and open intervals that may be compactified by points of the form $F(p, r) \times F(r, q)$ where the index of r is $i - 1 = j + 1$.

The manifolds $F(p, q)$ may be oriented by thinking of them as intersections of the (oriented) manifolds of *all* flowlines asymptotic to p (resp. q) in the future (resp. past); thus for $i - j = 1$ the set $F(p, q)$ is a finite set of *signed* points, and counting with sign gives an integer $n(p, q)$.

We may define a graded chain complex generated in dimension i by critical points of index i , and differential

$$\partial p := \sum_{\text{index}(q)=i-1} n(p, q)q$$

Thus

$$\partial \partial p = \sum_q \sum_r n(p, r)n(r, q)q$$

But for each q , this sum $\sum n(p, r)n(r, q)$ is equal to the number of boundary points of the compact 1-manifold $\overline{F}(p, q)$, counted with sign. Thus $\partial \partial = 0$ and ∂ is the differential of a chain complex, which by construction has homology isomorphic to $H_*(M; \mathbb{Z})$.

5.2.1. Products.

5.3. Floer Homology. Now let (M^{2n}, ω) denote a symplectic manifold with Lagrangian submanifolds L_0, L_1 and for simplicity let's suppose they intersect in general position. Let Ω denote the space of smooth maps $z : I \rightarrow M$ with $z(0) \in L_0$ and $z(1) \in L_1$.

The idea of Floer Homology is to define a smooth function a on Ω and compute the ‘‘Morse homology’’ associated to the critical points of a and the gradient flowlines joining them. The critical points will be the intersections $L_0 \cap L_1$ and the gradient flowlines will be holomorphic bigons with edges on $L_0 \cup L_1$ running between two intersection points.

If $\omega = d\lambda$ for some 1-form λ we could define a to be the action $a(z) := \int_z -\lambda$ (compare Example 3.4). In fact to do Morse theory we do not really need the function a as such; rather we need its derivative da . If we choose a basepoint z_0 and a sufficiently small neighborhood $z_0 \in U \subset \Omega$ then for any other $z \in U$ we can join z_0 to z by a path z_t in Ω which sweeps out a rectangle $Z : I \times I \rightarrow M$ with left and right edges on L_0 and L_1 respectively. We may then define $a(z) := \int_Z \omega$. If y_t were another path sweeping out another rectangle Y then we could sew Z and Y together to make an annulus with boundary curves on L_0, L_1 . If these curves were sufficiently small, we could cap them off

with small disks in the L_j to make a (null-homotopic) sphere S , and then

$$\int_Z \omega - \int_Y \omega = \int_S \omega = 0$$

The global indeterminacy of a comes from the periods of ω on cylinders whose boundaries are loops in L_1 and L_2 . Such cylinders (homotopically) are determined by intersections of conjugacy classes $\pi_1(L_0) \cap g\pi_1(L_1)g^{-1}$ together with the action of $\pi_2(M)$. Thus (for example) if $H_1(L_0) \cap H_1(L_1) = 0$ and $\pi_2(M) = 0$ then a is globally defined.

In any case da is well-defined. The tangent space $T_z\Omega$ is the space of vector fields ξ along z with $\xi(z(j)) \in T(z(j))L_j$ and

$$da(\xi) = \int_0^1 \omega(z'(t), \xi) dt$$

Thus the critical points of a are precisely the constant maps; i.e. the points of $L_0 \cap L_1$.

A compatible almost-complex structure J determines a metric on Ω by

$$\langle \xi_1, \xi_2 \rangle := \int_0^1 \omega(\xi_1, J\xi_2)$$

and therefore

$$da(\xi) = \int_0^1 \omega(z'(t), \xi) = \int_0^1 \omega(Jz'(t), J\xi) = \langle Jz', \xi \rangle$$

In other words (at least formally) $\text{grad}(a) = Jz'$ and flowlines of $\text{grad}(a)$ (up to sign) are maps $u : I \times I \rightarrow M$ with left and right edges on the L_j , satisfying

$$\frac{\partial u}{\partial s} - J \frac{\partial u}{\partial t} = 0$$

where the horizontal factor on $I \times I$ has coordinate t , and the vertical factor s .

In words: gradient flowlines of a are holomorphic rectangles with edges on the L_j . If $x, y \in L_0 \cap L_1$ are critical points of da , then the flowlines from x to y are holomorphic maps $u : \mathbb{D} \rightarrow M$ sending $-i$ to x and i to y , and so that the arcs of $\partial\mathbb{D}$ with positive (resp. negative) real part maps to L_1 (resp. L_0).

Let $\mathcal{M}(x, y)$ denote the space of such *holomorphic Whitney disks*. The automorphism group of \mathbb{D} fixing i and $-i$ is \mathbb{R} ; this acts freely on $\mathcal{M}(x, y)$ (at least when $x \neq y$ so that $\mathcal{M}(x, y)$ contains no constant maps) and we denote the quotient space $\overline{\mathcal{M}}(x, y)$.

The space $\mathcal{M}(x, y)$ might have many components, and the disks u in \mathcal{M} might lie in different homology classes. For a residual set of almost complex structures J the space $\mathcal{M}(x, y)$ is a smooth manifold of dimension that may be computed by the index formula. Since \mathbb{D} is contractible the pullback u^*TM has a (symplectic) trivialization as $\mathbb{D} \times \mathbb{R}^{2n}$. The circle $S^1 = \partial\mathbb{D}$ factorizes as the union of two arcs $\alpha_0 \cup \alpha_1$ where $u : \alpha_j \rightarrow L_j$, each oriented to run from x to y . Relative to the trivialization each arc $\alpha_j^*TL_j$ may be thought of as a path in \mathcal{L}_n . Join these two paths at the endpoints by paths that do not cross some fixed train. The resulting loop in \mathcal{L}_n has Maslov index $\mu(u)$ and Viterbo [?] shows that the formal dimension of $\mathcal{M}(x, y)$ in the component containing u is $\mu(u)$.

We may now define the Floer Homology of the pair L_0, L_1 as follows. The chain group is the free abelian group generated by intersections $L_0 \cap L_1$. For each x define

$$\partial x := \sum_y n(x, y)y$$

where $n(x, y)$ is the signed count of points in the components of the unparameterized moduli spaces $\overline{\mathcal{M}}(x, y)$ of formal dimension 0 (i.e. for which the corresponding components of $\mathcal{M}(x, y)$ are represented by holomorphic Whitney disks u with Maslov index 1).

This satisfies $\partial^2 = 0$; the homology of this complex is the Floer Homology.

REFERENCES

- [1] V. Arnold, *Sturm theorems and symplectic geometry*, (English translation) *Func. Anal. Appl.* **19** (1985), no. 4, 251–259
- [2] V. Arnold and A. Givental, *Symplectic Geometry*, in *Encyc. Math. Sci. Dynamical Systems IV*, 1–138, Springer, Berlin, 2001.
- [3] M. Atiyah and R. Bott, *The Yang-Mills Equations over Riemann Surfaces*, *Phil. Trans. R. Soc. Lond. A* **308** (1983), no. 1505, 523–615
- [4] J. Cheeger and D. Ebin, *Comparison Theorems in Riemannian Geometry*, North-Holland Math. Library, Vol. 9, Elsevier, New York, 1975.
- [5] C. Chevalley and S. Eilenberg, *Cohomology theory of Lie Groups and Lie Algebras*, *Trans. AMS* **63** (1948), 85–124
- [6] S. Friedl and S. Vidussi, *Twisted Alexander polynomials detect fibered 3-manifolds*, *Ann. Math. (2)* **173** (2011), no. 3, 1587–1643
- [7] W. Goldman, *The symplectic nature of fundamental groups of surfaces*, *Adv. Math.* **54** (1984), no. 2, 200–225
- [8] M. Gromov, *Pseudo holomorphic curves in symplectic manifolds*, *Invent. Math.* **82** (1985), no. 2, 307–347
- [9] S. Helgason, *Differential Geometry and Symmetric spaces*, *Pure Appl. Math.*, Vol. XII. Academic Press, New York-London, 1962.
- [10] H. Hofer, V. Lizan and J.-C. Sikorav, *On Genericity for Holomorphic Curves in Four-Dimensional Almost-Complex Manifolds*,
- [11] C. Hummel, *Gromov’s Compactness Theorem for Pseudo-holomorphic Curves*, *Prog. Math.* **151**, Birkhauser, Basel, 1997.
- [12] D. McDuff and D. Salamon, *J-Holomorphic Curves and Quantum Cohomology*, *AMS University Lecture Series Vol. 6*, AMS, Providence, 1994.
- [13] D. McDuff and D. Salamon, *Introduction to Symplectic Topology*, *Oxford Math. Monographs*, Oxford University Press, Oxford, 1998.
- [14] M. Micalef and B. White, *The structure of branch points in minimal surfaces and in pseudoholomorphic curves*, *Ann. Math. (2)* **141** (1995), no. 1, 35–85
- [15] J.-M. Morvan, *Classe de Maslov d’une immersion lagrangienne et minimalité*, *C. R. Acad. Sci. Paris Sér. I Math.* **292** (1981), no. 13, 633–636
- [16] C. L. Siegel, *Symplectic Geometry*, *Amer. J. Math.* **65** (1943), 1–86

UNIVERSITY OF CHICAGO, CHICAGO, ILL 60637 USA

E-mail address: dannyc@math.uchicago.edu