SYMPLECTIC MANIFOLDS

ALEX EASTMAN

ABSTRACT. This paper will cover the basics of symplectic structures on manifolds. We will begin by basic covering results when the manifold in question is a real valued vector space. Other topics covered in the paper will be the cotangent bundle, which is a natural example of a symplectic manifold, Hamiltonian vector fields, the Hamiltonian phase flow, and integral invariants.

CONTENTS

1. SYMPLECTIC STRUCTURES

Definition 1.1. A symplectic structure on a manifold M is a closed non-degenerate differential 2-form.

Explicitly unpacking this definition we have that a symplectic structure is a 2-form ω^2 on M which satisfies both of the following.

$$d\omega^2 = 0. \tag{1.2}$$

(1.3) For all $p \in M$, $\xi \in T_pM$ with $\xi \neq 0$, there exists η such that $\omega^2(\xi, \eta) \neq 0$.

The pair (M, ω^2) is called a symplectic manifold.

Lemma 1.4. *If (M, ω^2) is a symplectic manifold of finite dimension, then M is even dimensional.*

Proof. Since the non-degeneracy condition is a local property, it is enough to show that symplectic vector spaces are even dimensional. Assume we have a vector space V with a non-degenerate 2-form λ^2. First we will pick a vector $v_0 \in V, v \neq 0$. Since λ^2 is non-degenerate, there exists a vector u_0 such that $\lambda^2(v_0, u_0) = 1$. By skew-symmetry v_0 and u_0 are linearly independent. We can define the subspaces

$$V_0 = \text{span}\{v_0, u_0\}. \tag{1.5}$$

Date: AUGUST 31, 2019.

$$V_{\perp} = \{v_1 \in V : \text{for all } v \in V_0, \lambda^2(v, v_1) = 0\}. \tag{1.6}$$

We wish to show that $V = V_0 \bigoplus V_{\perp}$. The definitions of V_0 and $V_{\perp}$ give that their intersection is trivial. Suppose that $w_0 \in V$. Let

$$a = \lambda^2(v_0, w_0). \tag{1.7}$$

$$b = \lambda^2(u_0, w_0). \tag{1.8}$$

$$w_1 = au_0 - bv_0. \tag{1.9}$$

It follows that $\lambda^2(v_0, w_1) = a$ and $\lambda^2(u_0, w_1) = b$. If we let $w_{\perp} = w_0 - w_1$, then we have $w_0 = w_1 + w_{\perp}$. Since $w_1 \in V_0$, all that is left to show is that $w_{\perp} \in V_{\perp}$. To show that $\lambda^2(v_0, w_{\perp}) = 0$ we perform the following manipulations.

$$\lambda^2(v_0, w_{\perp}) = \lambda^2(v_0, w_0 - w_1). \tag{1.10}$$

$$= \lambda^2(v_0, w_0) - \lambda^2(v_0, w_1). \tag{1.11}$$

$$= a - a. \tag{1.12}$$

$$= 0. \tag{1.13}$$

Likewise we have $\lambda^2(u_0, w_{\perp}) = 0$. This gives $w_{\perp} \in V_{\perp}$ and thus $V = V_0 \bigoplus V_{\perp}$. We can then repeat this process on the restriction of λ^2 to $V_{\perp}$ which will decompose V into finitely many dimension 2 vector spaces. Therefore V has even dimension.

□

From now on we will always take M to be an even dimensional manifold and ω^2 to be a symplectic structure on M.

We will now cover a special case where $M = V = \mathbb{R}^{2n}$ is an even dimensional vector space. This case is very relevant because for any symplectic manifold the restriction of the symplectic structure to a point x is a symplectic form on the tangent space at that point.

Definition 1.14. A symplectic structure on $\mathbb{R}^{2n}$ is a non-degenerate 2-form on $\mathbb{R}^{2n}$. This is commonly denoted by brackets $[\cdot, \cdot]$. $\mathbb{R}^{2n}$ together with $[\cdot, \cdot]$ is a symplectic vector space.

If we denote the $2n$ coordinate functions on $\mathbb{R}^{2n}$ by $(q^1, \dots q^n, p_1, \dots p_n)$, then there is a canonical symplectic structure.

$$\omega^2(\xi, \eta) = [\xi, \eta]. \tag{1.15}$$

$$\omega^2 = q^1 \wedge p_1 + \cdots + q^n \wedge p_n. \tag{1.16}$$

We can interpret the value of $[\xi, \eta]$ as being the sum of the areas of the parallelograms formed by projecting ξ and η into the planes formed by the pairs (q^i, p_i).

Example 1.17. If we use (1.16) to evaluate $[\cdot, \cdot]$ on pairs of basis vectors, then we get the following result.

$$[e_{q^i}, e_{q^j}] = [e_{p_i}, e_{p_j}] = 0. \tag{1.18}$$

$$[e_{q^i}, e_{p_i}] = \delta_{ij}. \tag{1.19}$$

Definition 1.20. Two vectors ξ and η are said to be skew-orthogonal if $[\xi, \eta] = 0$.

Definition 1.21. The skew-orthogonal complement of a vector η is the set of all ξ such that ξ and η are skew-orthogonal. If we have a linear subspace V, the skew-orthogonal complement of V is the set of all ξ such that ξ and η are skew-orthogonal for all $\eta \in V$.

By the anti-symmetry of the 2-form, every vector is skew-orthogonal to itself. The skew-orthogonal complement is also a subspace by the linearity of the 2-form.

Theorem 1.22. *Let V be a linear subspace of dimension k. Then the dimension of the skew-orthogonal complement is $2n - k$.*

Proof. Let $(\xi_1, \dots, \xi_k)$ be a basis for V. Then the skew-orthogonal complement is the solution to the following system.

$$\begin{aligned} &[\xi_1, \eta] = 0. \\ &. \\ &. \\ &[\xi_k, \eta] = 0. \end{aligned}$$

The non-degeneracy conditions then gives that this system eliminates a space of dimension k, leaving the solution as a subspace of dimension $2n - k$. □

Definition 1.23. A symplectic basis is a set of basis vectors $(e_{q^1}, \dots e_{q^n}, e_{p_1}, \dots e_{p_n})$ which have the form of Example 1.17.

We will now prove an analogue of the theorem which states that in every Euclidean space there is an orthonormal basis, and any non-zero vector can be chosen as a basis element.

Theorem 1.24. *Every symplectic space has a symplectic basis. If e_{q^1} is an arbitrary non-zero vector in V, then e_{q^1} can be one of the basis vectors.*

Proof. By the non-degeneracy of $[\cdot, \cdot]$, there is a vector $e_{p_1} \in V$ such that $[e_{q^1}, e_{p_1}] = 1$. We will denote by U the skew-orthogonal complement of $\text{span}(e_{q^1}, e_{p_1})$. By Theorem 1.22, U has dimension $2n - 2$. We want to show that $[\cdot, \cdot]$ restricted to U gives a symplectic subspace. This requires that $[\cdot, \cdot]$ restricted to U is non-degenerate. Let $\xi \in U$. We have that $[\xi, e_{q^1}] = 0$ and $[\xi, e_{p_1}] = 0$. Since $[e_{q^1}, e_{p_1}] = 1$, neither e_{q^1} nor e_{p_1} are contained in U. If ξ were skew-orthogonal to all vectors in U, then ξ would be skew-orthogonal to all vectors in V which contradicts the non-degeneracy of $[\cdot, \cdot]$ over V. Now we have reduced the problem to a symplectic vector space U of dimension $2n - 2$. Repeating this process gives the theorem by induction. □

Corollary 1.25. *All symplectic spaces of the same dimension are isomorphic.*

2. The Cotangent Bundle

The cotangent bundle on a manifold is a convenient way of generating an even dimensional manifold with a natural symplectic structure.

To define the cotangent bundle let V be an n-dimensional differentiable manifold and x be a point in V. Then we have the following definitions.

Definition 2.1. A cotangent vector to V at x is a 1-form on the tangent space to V at x.

Definition 2.2. The set of all cotangent vectors to V at x forms an n-dimensional vector space. This space is called the cotangent space and is denoted by $T^*_x V$. The union of all tangent spaces is called the cotangent bundle and is denoted by T^*V.

The cotangent bundle can be given the structure of a differentiable manifold of dimension $2n$. The dimension follows from the fact that points in T^*V are 1-forms on the tangent space at some point in V. If $\mathbf{q}$ is a choice of n local coordinates in V and a 1-form is given by its n components $\mathbf{p}$, then the $2n$ numbers $\mathbf{p}$, $\mathbf{q}$ form local coordinates for points in T^*V.

Theorem 2.3. *This choice of local coordinates allows us to assign a natural symplectic structure to the cotangent bundle T^*V. This structure is given by the following formula.*

$$\omega^2 = d\boldsymbol{q} \wedge d\boldsymbol{p} = dq^1 \wedge dp_1 + \cdots + dq^n \wedge dp_n.$$

Proof. First we will note that there is a canonical projection $\pi : T^*V \to V$ which sends every 1-form on $T^*_x V$ to x. This map is both differentiable and surjective.

We can then define the tautological 1-form on T^*V called θ. In order to define this 1-form we will need a canonical way of mapping the tangent vectors of the cotangent bundle into $\mathbb{R}$. The immediate first step is to map the complicated elements of $T(T^*V)$ into a simpler space. To do this we use the derivative of the canonical projection $d\pi$. $d\pi$ maps $T(T^*V)$ into TV. We can then recall that each element in $T(T^*V)$ is contained inside $T_p(T^*V)$ for some $p \in T^*V$. Breaking it down further each point $p \in T^*V$ is a pair (x, ω) for $x \in V$ and $\omega \in T^*_x V$. We can then compose ω and $d\pi$ to create the tautological 1-form. If we write this 1-form in the coordinates of $(\mathbf{p}, \mathbf{q})$ above we get

$$\theta = \sum_{i=1}^{n} p_i dq^i. \tag{2.4}$$

Applying the definition of the exterior derivative to θ gives that

$$-d\theta = dq^1 \wedge dp_1 + \cdots + dq^n \wedge dp_n = \omega^2. \tag{2.5}$$

Since exact forms are necessarily closed, this gives immediately that ω^2 is closed. The non-degeneracy of ω^2 follows by noting that since non-degeneracy is a local property we need only check that this construction works on $M = \mathbb{R}^{2n}$. That check was completed in the previous section when we defined the canonical symplectic structure on a vector space. □

Example 2.6. A natural example arises in physics. We can consider a Lagrangian mechanical system with a configuration manifold V and a Lagrangian function L. We want to be able to understand the "generalized velocity" and "generalized momentum" of Lagrangian mechanics in the language of differential geometry. If we continue the convention of denoting by $\mathbf{q}$ the coordinates of a given configuration, then we can see that the "generalized velocity" vector $\dot{\mathbf{q}}$ is a tangent vector to V. The "generalized momentum", which is defined as $\mathbf{p} = \partial L / \partial \dot{\mathbf{q}}$, is a cotangent vector. Therefore, the cotangent bundle of the configuration space V gives the phase space of the mechanical system. Theorem 2.3 shows that the phase space of any Lagrangian mechanical system is a manifold with a natural symplectic structure.

3. Hamiltonian Vector Fields

A symplectic structure establishes an isomorphism between the spaces of tangent vectors and 1-forms.

Definition 3.1. To each vector ξ, tangent to a symplectic manifold (M, ω^2) at x, we wish to find a 1-form ω_ξ on T_xM. To define the value of ω_ξ on a vector $\eta \in T_xM$ we use the formula

$$\omega_\xi(\eta) = \omega^2(\eta, \xi).$$

That this formula actually gives a 1-form follows from the linearity and differentiability of the symplectic structure ω^2.

Next we will show that Definition 3.1 actually gives an isomorphism.

Lemma 3.2. *$\xi \to \omega_\xi$ is an isomorphism between tangent vectors and 1-forms.*

Proof. First we will show the map is linear and thus defines a homomorphism of vector spaces.

$$\omega_\xi(\alpha\eta_1 + \beta\eta_2) = \omega^2(\alpha\eta_1 + \beta\eta_2, \xi). \tag{3.3}$$

$$= \alpha\omega^2(\eta_1, \xi) + \beta\omega^2(\eta_2, \xi). \tag{3.4}$$

$$= \alpha\omega_\xi(\eta_1) + \beta\omega_\xi(\eta_2). \tag{3.5}$$

Next we need to demonstrate that the map is a bijection. The spaces are of the same dimension (n-dimensional for a fixed point x or $2n$-dimensional if we consider all tangent vectors and all 1-forms on (M, ω^2)). Since the spaces are finite dimensional, it suffices to show that the kernel of the map is trivial and thus the map is an injection. The kernel is all ξ such that $\omega_\xi(\eta) = 0$ for all η. However, the non-degeneracy condition of the symplectic structure ω^2 guarantees that $\omega^2(\eta, \xi) = 0$ for all η if and only if $\xi = 0$. Thus, the kernel is trivial and the map is an isomorphism. □

Remark 3.6. We will find it more useful to work with the inverse of the above isomorphism. We will denote by I the isomorphism mapping cotangent vectors to tangent vectors.

If we have a differentiable function $H : M \to \mathbb{R}$ on our symplectic manifold (M, ω^2), then dH is a differential 1-form on M. At every point on M the isomorphism I assigns a tangent vector. This yields a vector field IdH on M. The vector field IdH is called a Hamiltonian vector field and H is called the Hamiltonian function.

4. Hamiltonian Phase Flows and the Symplectic Structure

Definition 4.1. Given 2 manifolds M and N, a diffeomorphism is a smooth map $f : M \to N$ where f is a bijection and the inverse $f^{-1} : N \to M$ is smooth.

We can use the Hamiltonian vector field IdH from the previous section to define a 1-parameter group of diffeomorphisms $g^t : M \to M$. To do this let the velocity at each point x be determined by the value of the Hamiltonian vector field at that point. That is, to see how a point x in M is transforming by g^t at $t = 0$ we use the following formula.

$$\left.\frac{d}{dt}\right|_{t=0} g^t x = (IdH)(x). \tag{4.2}$$

Definition 4.3. We call the 1-parameter group $\{g^t\}$ the Hamiltonian phase flow.

We will now state a theorem which states that Hamiltonian phase flows preserve the symplectic structure.

Theorem 4.4.

$$(g^t)^*\omega^2 = \omega^2$$

I will not provide a proof of this theorem because there is already an excellent proof found in [4]. However, we will quickly note that in the case that $M = \mathbb{R}^2$, we have that ω^2 is a volume form. In this case the statement of Theorem 4.4 gives that the phase flow preserves area. This is commonly referred to as Liouville's theorem for 2 dimensions and is useful in both Hamiltonian and statistical mechanics. Theorem 4.4 is thus a generalization of the 2 dimensional version of Liouville's theorem. We will soon extend this to prove Liouville's theorem on any symplectic manifold.

We will now prove a fundamental result about Hamiltonian phase flows.

Theorem 4.5. *A Hamiltonian function H is constant along the Hamiltonian phase flow.*

Proof. First we will fix a point $x \in M$. We will denote by η the direction of the phase flow at x. To show that H is constant we will show that $dH(\eta) = 0$. We will rewrite $dH(\eta)$ in terms of the symplectic structure ω^2. Since I is an isomorphism we have the following tautology.

$$dH(\eta) = (I^{-1} \circ I)(dH)(\eta) = I^{-1}(IdH)(\eta)$$

Then we apply I^{-1} as defined in 3.1.

$$dH(\eta) = \omega^2(\eta, IdH) \tag{4.6}$$

By (4.2) we see that at every point x IdH is just η.

$$dH(\eta) = \omega^2(\eta, \eta) = 0.$$

□

In physics where H commonly represents the total energy of a mechanical system this is known as the law of conservation of energy. However, the Hamiltonian does not always represent total energy. This theorem shows that nonetheless H is conserved.

5. Integral Invariants

Definition 5.1. Let $f : M \to M$ be a differentiable map. A differential k-form λ^k is an absolute integral invariant if for any k-chain c we have the following.

$$\int_{fc} \lambda^k = \int_c \lambda^k.$$

Here fc is the image of the k-chain c under the map f. If $f : \mathbb{R}^{2n} \to \mathbb{R}^{2n}$ has ω^2 as an integral invariant, then f is called canonical. f is typically referred to as a canonical transformation.

There are a couple of useful facts about absolute integral invariants we will now prove.

Lemma 5.2. *The k-form λ^k is an absolute integral invariant of the map f if and only if $f^*\lambda^k = \lambda^k$.*

Proof. First we will assume that λ^k is an absolute integral invariant of the map f. Then for any chain c we can perform the following manipulations.

$$\begin{aligned}\int_c \lambda^k &= \int_{fc} \lambda^k. \\ &= \int_c f^*\lambda^k.\end{aligned}$$

By subtracting the right hand side from the left hand side we get

$$\int_c \lambda^k - f^*\lambda^k = 0.$$

Since this holds over any chain c we deduce that $\lambda^k - f^*\lambda^k = 0$ and thus $f^*\lambda^k = \lambda^k$. For the other direction we start with the tautology

$$\int_c \lambda^k = \int_c \lambda^k.$$

Then we use the assumption that $f^*\lambda^k = \lambda^k$.

$$\begin{aligned}\int_c \lambda^k &= \int_c f^*\lambda^k. \\ &= \int_{fc} \lambda^k.\end{aligned}$$

Thus both sides are proved.

□

Lemma 5.3. *If λ^k and θ^l are absolute integral invariants of f, then $\lambda^k \wedge \theta^l$ is an integral invariant of f.*

Proof. By the previous lemma we have that $f^*\lambda^k = \lambda^k$ and that $f^*\theta^l = \theta^l$. Since the wedge product commutes with the pullback, we have

$$\begin{aligned}f^*(\lambda^k \wedge \theta^l) &= f^*(\lambda^k) \wedge f^*(\theta^l). \\ &= \lambda^k \wedge \theta^l.\end{aligned}$$

Therefore, by the previous lemma $\lambda^k \wedge \theta^l$ is an absolute integral invariant of f. □

Lemma 5.2 allows us to rewrite Theorem 4.4.

Theorem 5.4. *The symplectic structure ω^2 is an absolute integral invariant of the map g^t for every t.*

Lemma 5.3 gives that the exterior powers of ω^2 are also absolute integral invariants of the Hamiltonian phase flow.

Corollary 5.5. $(\omega^2)^k = \omega^2 \wedge \omega^2 \wedge \cdots \wedge \omega^2$ *is an absolute integral invariant of the map g^t for every t.*

We now want to be able to define a volume form using using the fact that if M is $2n$-dimensional, then $(\omega^2)^n$ is a form of top degree. To do this we need to prove that $(\omega^2)^n \neq 0$.

Lemma 5.6. $(\omega^2)^n \neq 0$

Proof. Theorem 2.3 gives that

$$\omega^2 = dq^1 \wedge dp_1 + \cdots + dq^n \wedge dp_n.$$

Here dq^i and dp_j are basis vectors. Therefore, when ω^2 is raised to the n-th power any term with a multiplicity of basis vectors will be zero. After rearranging the remaining terms we are left with the following equation.

$$(\omega^2)^n = n!(dq^1 \wedge dp_1 \wedge dq^2 \wedge dp_2 \wedge \cdots \wedge dq^n \wedge dp_n). \tag{5.7}$$

This completes the proof because the restriction of $(\omega^2)^n$ at any point will give a volume form on the tangent space. □

Combining Corollary 5.5 and Lemma 5.6 gives the proof of the full $2n$ dimensional version of Liouville's theorem.

We will now look at a weaker form of integral invariant.

Definition 5.8. A differential k-form λ^k is a relative integral invariant of the map $f : M \to M$ if for every closed k-chain c we have the following.

$$\int_{fc} \lambda^k = \int_c \lambda^k.$$

The boundary of any chain c is always closed. This allows us to use Stokes' theorem to prove the following.

Theorem 5.9. *If λ is a relative integral invariant of f, then $d\lambda$ is an absolute integral invariant of f.*

Proof. First we apply Stokes' theorem.

$$\int_c d\lambda = \int_{\partial c} \lambda.$$

Then we can apply the definition of relative integral invariant and use the fact that The boundary of any chain c is always closed.

$$\begin{aligned} \int_{\partial c} \lambda &= \int_{f\partial c} \lambda. \\ &= \int_{\partial fc} \lambda. \end{aligned}$$

Applying Stokes' Theorem one last time gives the result.

$$\int_c d\lambda = \int_{fc} d\lambda.$$

□

Theorem 5.9 has the following application.

Theorem 5.10. *The 1-form dH is an absolute integral invariant of the Hamiltonian phase flow.*

Proof. By Lemma 5.2 it is enough to show that $(g^t)^* dH = dH$. The pullback commutes with the exterior derivative.

$$(g^t)^* dH = d((g^t)^* H). \tag{5.11}$$

By the definition of the pullback on a smooth function we have that $(g^t)^*H = H(g^t)$. In Theorem 4.5 we showed that H is constant along the Hamiltonian phase flow. This gives $(g^t)^*H = H$. We then plug this back into (5.11).

$$(g^t)^* dH = dH. \tag{5.12}$$

This completes the theorem. □

Acknowledgments

Thank you to my mentor, Elia Portnoy, for suggesting this topic and for providing the assistance and reference materials needed to complete the project.

References

[1] http://www.ams.org/publications/authors/tex/amslatex
[2] Michael Downes. Short Math Guide for LaTeX. http://tex.loria.fr/general/downes-short-math-guide.pdf
[3] Tobias Oekiter, Hubert Partl, Irene Hyna and Elisabeth Schlegl. The Not So Short Introduction to LaTeX2e. https://tobi.oetiker.ch/lshort/lshort.pdf
[4] Vladimir Arnold. Mathematical Methods of Classical Mechanics 2e. Springer-Verlag 1978.
[5] Ana Cannas da Silva Lectures on Symplectic Geometry. http://math.mit.edu/classes/18.157/S2012/lecturesonsymplecticgeometryanacannas.pdf
[6] Ana Cannas da Silva Symplectic Geometry https://people.math.ethz.ch/ acannas/Papers/handbook.pdf
[7] Victor Guillemin, Alan Pollack. Differential Topology. American Mathematical Society 1974.