Bachelor's project

NTNU
Norwegian University of Science and Technology
Faculty of Information Technology and Electrical Engineering
Department of Mathematical Sciences

Trygve Poppe Oldervoll

The Arnold Conjecture

An introduction to symplectic topology

Bachelor's project in Mathematical Sciences
Supervisor: Glen M. Wilson

May 2020

NTNU
Kunnskap for en bedre verden

Trygve Poppe Oldervoll

The Arnold Conjecture

An introduction to symplectic topology

Bachelor's project in Mathematical Sciences
Supervisor: Glen M. Wilson
May 2020

Norwegian University of Science and Technology
Faculty of Information Technology and Electrical Engineering
Department of Mathematical Sciences

NTNU
Norwegian University of
Science and Technology

Contents

Acknowledgements

This thesis would not have been possible without my supervisor Glen M. Wilson. I thank him as a lecturer, for piquing my interest in the field of differential topology, and as a supervisor for pushing me in the right directions while still allowing me the freedom to explore.

1 Introduction

Many problems in mathematics boil down to finding fixed points of certain maps. Therefore we are always interested in tools for finding, or just guaranteeing the existence of such points. To this end we have many important results like Banach's, and Brouwer's fixed point theorems. In the case of differential topology, the most standard tool to study fixed points is the Lefschetz fixed point theorem. It states that if $f : M \to M$ is a smooth map of the smooth manifold M,

$$\Lambda_f = \sum_{x \in \operatorname{Fix}(f)} i(x, f) \tag{1.1}$$

where $i(x, f)$ is the index of the fixed point, and Λ_f is the Lefschetz number of f. The Lefschetz number can be computed as the alternating sum of traces of the matrix representations of f_* on the rational homology spaces. In particular, if f is homotopic to the identity, the formula above becomes

$$\Lambda_f = \Lambda_{id} = \chi(M) = \sum_{k=0}^{n} (-1)^k b_k \tag{1.2}$$

where b_k are the Betti numbers of M, and $\dim M = n$. This is very useful, but if all we care about is the number of fixed points, the best we can do is give the lower bound

$$\Lambda_f \neq 0 \implies \# \operatorname{Fix}(f) \geq 1. \tag{1.3}$$

Under what circumstances can we do better? One important special case of smooth maps of manifolds are symplectomorphisms. This special class of diffeomorphisms arise naturally as the time evolutions and symmetries of Hamiltonian systems in physics, and are at the core of the field of symplectic topology. So what can we say about the number of fixed points of symplectomorphisms? Quite a lot actually, especially under some mild extra conditions. Our hopes are summarized in the following conjecture by Vladimir Arnold.

Conjecture 1.1 (Arnold)**.** *If $\psi : M \to M$ is a Hamiltonian symplectomorphism of a symplectic manifold (M, ω), then ψ must have at least as many fixed points as a function on M must have critical points. If all the fixed points of ψ are nondegenerate, ψ must have at least as many fixed points as a Morse function on M must have critical points.*

To see the power of this conjecture, remember that Morse theory estimates that the number of critical points of a Morse function is at least the sum of the Betti numbers of M. So in the nondegenerate case, Arnold's conjecture implies that

$$\# \operatorname{Fix}(\psi) \geq \sum_{k=1}^{m} b_k. \tag{1.4}$$

Comparing this to (1.2) and (1.3), we can see that this really is a powerful estimate. This large estimate hints at the fact that the structure of symplectic geometry is in fact quite rigid.

Unfortunately, the Arnold conjecture does not hold in full generality, but many slightly weaker results have been proven. The major breakthrough came with the development of Floer homology. Building on Floer's work, symplectic geometers were able to prove that in the nondegenerate case on a closed manifold, the sum of the Betti numbers give a lower bound for the number of fixed points. Further references and details can be found in [Sal99]. In this thesis we introduce the theory of symplectic geometry, and highlight some of its important features. We then go on to prove two special cases of the Arnold conjecture, the case where $M = \mathbb{T}^{2n}$ and the case where ψ is sufficiently close to the identity map in a particular C^1 topology on the space of symplectomorphisms.

This thesis is mainly based on [MS98], both when it comes to structure of the chapters, and statement of theorems and definitions. In particular, we include the relevant material from chapters 1,2,3,9 and 11. Wherever a different source has been used, we will provide a reference.

2 Classical mechanics

The field of symplectic geometry arises as a generalization of concepts from classical mechanics. In this section we show how the Hamiltonian equations on $\mathbb{R}^{2n}$ arise from a variational principle, and how this is related to the modern theory of symplectic geometry. This will provide important motivation for the variational techniques employed later.

2.1 The Legendre transform

In Lagrangian mechanics, we think of $\mathbb{R}^{2n}$ as the tangent-bundle of $\mathbb{R}^n$ – that is, we use coordinates $(x_1, .., x_n, v_1, ..., v_n)$, where the x's describe position, and the v's velocity. A Lagrangian system is specified by a twice differentiable function

$$L = L(t, x, v) : \mathbb{R}^{2n+1} \to \mathbb{R}.$$

The system evolves from the state (x_0, t_0) to the state (x_1, t_1) along a path $x \in C^1([t_0, t_1], \mathbb{R}^{2n})$ minimizing the action integral

$$I(x) = \int_{t_0}^{t_1} L(t, x(t), \dot{x}(t))\mathrm{d}t \tag{2.1}$$

with respect to variations fixing the endpoints. Using simple variational techniques, we show that the problem of minimizing this integral is related to solving the Euler–Lagrange equation:

$$\frac{\mathrm{d}}{\mathrm{d}t}\frac{\partial L}{\partial v}(t, x, \dot{x}) = \frac{\partial L}{\partial x}(t, x, \dot{x}) \tag{2.2}$$

where

$$\frac{\partial L}{\partial x} = (\frac{\partial L}{\partial x_1}, ..., \frac{\partial L}{\partial x_n}), \qquad \frac{\partial L}{\partial v} = (\frac{\partial L}{\partial v_1}, ..., \frac{\partial L}{\partial v_n}).$$

A path $x \in C^1([t_0, t_1], \mathbb{R}^{2n})$ satisfying the boundary conditions

$$x(t_0) = x_0, x(t_1) = x_1 \tag{2.3}$$

is called minimal with respect to variations fixing the endpoints if

$$I(x) \leq I(x + \xi) \tag{2.4}$$

for all differentiable paths $\xi \in C^1([t_0, t_1], \mathbb{R}^{2n})$ such that $\xi(t_0) = \xi(t_1) = 0$.

Lemma 2.1. *If a path $x \in C^1([t_0, t_1], \mathbb{R}^{2n})$ satisfying* (2.3) *is minimal with respect to variations fixing the endpoints, it is a solution to the Euler–Lagrange equation* (2.2)*.*

Proof. If x minimizes I, all the directional derivatives of I vanish at x, so we have

$$\begin{aligned}
0 &= \frac{\mathrm{d}}{\mathrm{d}s}\Big|_{s=0} I(x+s\xi) \\
&= \frac{\mathrm{d}}{\mathrm{d}s}\Big|_{s=0} \int_{t_0}^{t_1} L(t, x(t)+s\xi(t), \dot{x}(t)+s\dot{\xi}(t))\mathrm{d}t \\
&= \int_{t_0}^{t_1} \langle \frac{\partial L}{\partial x}, \xi\rangle + \langle \frac{\partial L}{\partial v}, \dot{\xi}\rangle \mathrm{d}t \\
&= \int_{t_0}^{t_1} \left(\langle \frac{\partial L}{\partial x}, \xi\rangle - \langle \frac{\mathrm{d}}{\mathrm{d}t}\frac{\partial L}{\partial v}, \xi\rangle \right) \mathrm{d}t + \langle \frac{\partial L}{\partial v}, \xi(t_0)\rangle - \langle \frac{\partial L}{\partial v}, \xi(t_1)\rangle \\
&= \int_{t_0}^{t_1} \langle \frac{\partial L}{\partial x} - \frac{\mathrm{d}}{\mathrm{d}t}\frac{\partial L}{\partial v}, \xi\rangle \mathrm{d}t.
\end{aligned}$$

Note that all the partial derivatives are evaluated at $(t, x(t), \dot{x}(t))$. The third equality holds if we assume we can interchange differentiation and integration, the fourth via integration by parts, and the last because of the boundary conditions. Since this holds for all ξ, the fundamental lemma of calculus of variations implies that x solves the Euler–Lagrange equation. □

Remark 2.2. *It should be noted that the converse of this lemma does not hold in general. A solution to the Euler–Lagrange equation is a critical point of $I(x)$, but not necessarily a global minimum. Sometimes a global minimum may not even exist. We will leave these problems to the physicists for now, and focus on solutions to the Euler–Lagrange equations.*

The Euler Lagrange equations determine a set of second order differential equations in the n variables $(x_1, ..., x_n)$. If the Legendre condition

$$\det\left(\frac{\partial^2 L}{\partial v_j \partial v_k}\right) \neq 0 \tag{2.5}$$

is satisfied, we can introduce a new set of variables to create a set of first order differential equations in $2n$ variables. Let

$$y_k = \frac{\partial L}{\partial v_k}(x, v). \tag{2.6}$$

The Legendre condition implies that the mapping $(x, v) \mapsto (x, y)$ has an inverse, so we can think of x and y as independent variables. In other words, we can think of v as a function $v = G(t, x, y)$. Whenever x solves (2.2), we have

$$\dot{y} = \frac{\mathrm{d}}{\mathrm{d}t}\frac{\partial L}{\partial v} = \frac{\partial L}{\partial x}.$$

We now define a new function $H : \mathbb{R}^{2n+1} \to \mathbb{R}$ by

$$H(t, x, v) = \langle y, v\rangle - L(t, x, v).$$

This is the corresponding Hamiltonian function of the system. Using G as a Green's function, we can consider H as a function of t, x and y. The partial derivatives are

$$\frac{\partial H}{\partial x}(t,x,y) = -\frac{\partial L}{\partial x}(t,x,G(t,x,y)), \qquad \frac{\partial H}{\partial y}(t,x,y) = G(t,x,y).$$

It now follows that whenever the Euler–Lagrange equations (2.2) are satisfied, we have

$$\dot{x} = \frac{\partial H}{\partial y}, \qquad \dot{y} = -\frac{\partial H}{\partial x}. \tag{2.7}$$

This pair of equations is known as the Hamiltonian differential equations, and will be crucial to our further study. The process we just described is known as the Legendre transform. We have changed coordinates from tangent vectors, v, to linear functions of tangent vectors, $\frac{\partial L}{\partial v}$. In a sense, we are now thinking of $\mathbb{R}^{2n}$ as the cotangent-bundle, $T^*\mathbb{R}^n$. Note that given any Hamiltonian function $H \in C^2(\mathbb{R}^{2n+1}, \mathbb{R})$ satisfying a nondegeneracy condition

$$\det\left(\frac{\partial^2 H}{\partial y_j \partial y_k}\right) \neq 0, \tag{2.8}$$

an inverse Legendre transform can be performed, retrieving a corresponding Lagrangian function L and the variables (x, v).

2.2 Symplectic action

We have seen that for a system satisfying a nondegeneracy condition (2.5), the Hamiltonian differential equations (2.7) can be expressed in terms of a variational principle. The Legendre condition (2.5) can be quite restrictive, so we wish to avoid it. This turns out to be possible if we formulate a different variational principle: Given a curve $z = (x, y) \in C^2([t_0, t_1], \mathbb{R}^{2n})$ and a Hamiltonian function $H(t, x, y)$, define the symplectic action integral as

$$\Phi_H(z) = \int_{t_0}^{t_1} \langle y, \dot{x}\rangle - H(t,x,y)\mathrm{d}t. \tag{2.9}$$

It is not hard to check that whenever H arose from L via a Legendre-transform, the integral (2.9) agrees with $I(x)$ in (2.1). The upshot is that the action integral Φ_H is defined for any curve in $C^2([t_0, t_1], \mathbb{R}^{2n})$, and any Hamiltonian function. The next lemma shows that the Hamiltonian differential equations (2.7) can be formulated as a variational principle in terms of Φ_H.

Lemma 2.3. *A curve $z \in C^2([t_0, t_1], \mathbb{R}^{2n})$ is a critical point for Φ_H with respect to variations with fixed endpoints if and only if $z = (x, y)$ is a solution to the Hamiltonian differential equations* (2.7)*.*

Proof. Let $z_s = (x_s, y_s)$ be a smooth one-parameter family of curves with $z_0 = z$. We denote the directional derivatives at zero by

$$\xi = \frac{\partial}{\partial s} x_s \Big|_{s=0}, \qquad \eta = \frac{\partial}{\partial s} y_s \Big|_{s=0}, \qquad \hat{\Phi}_H = \frac{\partial}{\partial s} \Phi_H(z_s) \Big|_{s=0}.$$

Then, by differentiating under the integration sign we get

$$\begin{aligned}
\hat{\Phi}_H &= \int_{t_0}^{t_1} \frac{\mathrm{d}}{\mathrm{d}s} \left(\langle y_s, \dot{x}_s \rangle - H(t, x_s, y_s) \right) \mathrm{d}t \Big|_{s=0} \\
&= \int_{t_0}^{t_1} \langle \frac{\mathrm{d}}{\mathrm{d}s} y_s, \dot{x}_s \rangle + \langle y_s, \frac{\mathrm{d}}{\mathrm{d}s} \dot{x}_s \rangle - \langle \frac{\partial H}{\partial x}, \frac{\mathrm{d}}{\mathrm{d}s} x_s \rangle - \langle \frac{\partial H}{\partial y}, \frac{\mathrm{d}}{\mathrm{d}s} y_s \rangle \mathrm{d}t \Big|_{s=0} \\
&= \int_{t_0}^{t_1} \langle \eta, \dot{x} \rangle + \langle y, \dot{\xi} \rangle - \langle \frac{\partial H}{\partial x}, \xi \rangle - \langle \frac{\partial H}{\partial y}, \eta \rangle \mathrm{d}t \\
&= \int_{t_0}^{t_1} \langle \eta, \dot{x} - \frac{\partial H}{\partial y} \rangle \mathrm{d}t + \int_{t_0}^{t_1} \langle \xi, -\dot{y} - \frac{\partial H}{\partial x} \rangle \mathrm{d}t + \langle y(t_1), \xi(t_1) \rangle - \langle y(t_0), \xi(t_0) \rangle.
\end{aligned} \tag{2.10}$$

Where we have used integration by parts. Due to our boundary conditions $\xi(t_1) = \xi(t_0) = 0$, the last two terms vanish. Again using the fundamental lemma of calculus of variations, we get our equivalence. $\square$

2.3 Hamiltonian flows and symplectomorphisms

From now on, our standard coordinates on $\mathbb{R}^{2n}$ will be

$$z = (x_1, x_2, ..., x_n, y_1, ..., y_n).$$

In these coordinates the Hamiltonian equations can be reformulated as

$$J_0 \dot{z} = \nabla H_t(z), \tag{2.11}$$

where $\nabla H_t(z)$ denotes the gradient of H_t at z, and J_0 is the standard complex-structure on $\mathbb{R}^{2n}$:

$$J_0 = \begin{pmatrix} 0 & -I_n \\ I_n & 0 \end{pmatrix} \tag{2.12}$$

Note that $J_0^2 = -I_{2n}$, and that if we associate $\mathbb{R}^{2n}$ with $\mathbb{C}^n$ by letting $z_j = x_j + iy_j$, multiplication by J_0 corresponds to multiplication by i. The vector field

$$X_{H_t} = -J_0 \nabla H_t : \mathbb{R}^{2n} \to \mathbb{R}^{2n} \tag{2.13}$$

is called the Hamiltonian vector field generated by H_t. This vector field determines a flow: For suitable $t_0, t \in \mathbb{R}$ let $\phi_H^{t,t_0} : \mathbb{R}^{2n} \to \mathbb{R}^{2n}$ be the solutions to the first order differential equation

$$\frac{\mathrm{d}}{\mathrm{d}t} \phi_H^{t,t_0} = X_{H_t} \circ \phi_H^{t,t_0}, \qquad \phi_H^{t_0,t_0} = id. \tag{2.14}$$

The family of diffeomorphisms ϕ_H^{t,t_0} is called the Hamiltonian flow generated by H_t, and it satisfies

$$\phi_H^{t_2,t_1} \circ \phi_H^{t_1,t_0} = \phi_H^{t_2,t_0}, \qquad \phi_H^{t,t} = id.$$

These diffeomorphisms are prototypical examples of symplectomorphisms:

Definition 2.4. A symplectomorphism $\psi : \mathbb{R}^{2n} \to \mathbb{R}^{2n}$ is a diffeomorphism such that

$$\mathrm{d}\psi^T J_0 \, \mathrm{d}\psi = J_0. \tag{2.15}$$

In the modern theory, symplectomorphisms will be exactly the maps preserving "symplectic structure." We are yet to define a notion of symplectic structure, but we do have Hamiltonian flows. We start by checking that Hamiltonian flows actually are symplectomorphisms.

Lemma 2.5. *The Hamiltonian flow ϕ_H^{t,t_0} is a symplectomorphism wherever defined.*

Proof. Let $z_0 \in \mathbb{R}^{2n}$ and define

$$z(t) = \phi_H^{t,t_0}(z_0), \qquad \Phi(t) = \mathrm{d}\phi_H^{t,t_0}(z_0).$$

Our goal is to show that $\Phi(t)$ satisfies the symplectomorphism condition

$$\Phi(t)^T J_0 \Phi(t) = J_0$$

for all t. This is clearly satisfied for $\Phi(t_0) = id$, so if we could show that the time derivative of the left hand side is zero, we would be finished. For every ζ_0 in $\mathbb{R}^{2n}$, the function $\zeta(t) = \Phi(t)\zeta_0$ satisfies the linearized Hamiltonian equations:

$$\dot{\zeta} = \mathrm{d}X_H(z) \circ \Phi.$$

Multiplying both sides with J_0 and using the fact that $J_0 X_H = \nabla H$, we get

$$J_0 \dot{\Phi}(t) = S(t)\phi(t),$$

where $S(t)$ is the Hessian of H at $z(t)$. Using the product rule, the time derivative is

$$\begin{aligned}\frac{\mathrm{d}}{\mathrm{d}t}\left(\Phi(t)^T J_0 \Phi(t)\right) &= \dot{\Phi}(t)^T J_0 \Phi(t) + \Phi(t)^T J_0 \dot{\Phi}(t) \\ &= \Phi(t)^T S(t) \Phi(t) - \Phi(t)^T S(t)^T \Phi(t) \\ &= \Phi(t)^T \left(S(t) - S(t)^T\right) \Phi(t),\end{aligned}$$

which is zero since the Hessian is symmetric for all t. □

Hamiltonian flows are not the only examples of symplectomorphisms. In general they represent the symmetries of the Hamiltonian system. This notion of symmetry is captured by the following lemma.

Lemma 2.6. *If $\psi : \mathbb{R}^{2n} \to \mathbb{R}^{2n}$ is a symplectomorphism, and $\zeta \in C^1([t_0, t_1], \mathbb{R}^{2n})$ is a solution of the Hamiltonian differential equation*

$$\dot{\zeta} = X_{H\circ\psi}(\zeta),$$

then $z = \psi \circ \zeta$ is a solution to the standard Hamiltonian differential equation. In other words,

$$\psi^* X_H = X_{H\circ\psi}.$$

Proof. Using the relationship between ∇H_t and $\mathrm{d}H_t$, and the chain rule, we obtain

$$\nabla(H \circ \psi)(p) = \mathrm{d}\psi^T(p)\nabla H(\psi(p)).$$

Using our hypothesis, and the chain rule $\dot{z} = \mathrm{d}\psi(\zeta)\dot{\zeta}$, we get

$$\begin{aligned} \mathrm{d}\psi^T(\zeta)\nabla H(\psi(\zeta)) &= \nabla(H \circ \psi)(\zeta) \\ &= J_0\dot{\zeta} \\ &= \mathrm{d}\psi(\zeta)^T J_0\, \mathrm{d}\psi(\zeta)\dot{\zeta} \\ &= \mathrm{d}\psi(\zeta)^T J_0 \dot{z}. \end{aligned}$$

Since $\mathrm{d}\psi$ is non-singular, this implies that $J_0\dot{z} = \nabla H(z)$ as required. □

This is all good for dynamics on $\mathbb{R}^{2n}$, but we want to generalize the theory to suitable manifolds. The way to do this is to rephrase our conditions in the language of differential forms. The standard symplectic form on $\mathbb{R}^{2n}$ is the two-form

$$\omega_0 = \sum_{j=1}^{n} \mathrm{d}x_j \wedge \mathrm{d}y_j. \tag{2.16}$$

A quick calculation shows that for any vectors $z, z' \in \mathbb{R}^{2n}$,

$$\omega_0(z, z') = \sum_{j=1}^{n} x_j y_j' - x_j' y_j = -z^T J_0 z. \tag{2.17}$$

In this language, the condition for being a symplectomorphism (2.15) is equivalent to

$$\psi^*\omega_0 = \omega_0, \tag{2.18}$$

and the equation for the Hamiltonian vector field can be rewritten as

$$\iota_{X_t}(\omega_0) = \mathrm{d}H_t. \tag{2.19}$$

The symplectic action integral can also be reformulated in this language as

$$\Phi_H(z) = \int_\gamma \lambda + H_t \mathrm{d}t, \tag{2.20}$$

where γ is the curve $z([0,1])$, and $\lambda = \sum y_j \mathrm{d}x_j$. This gives a quick glimpse into the language of the modern theory which we will devote the rest of this thesis to.

3 Linear Symplectic Geometry

In this section we study the linear theory, which will be the model for the smooth theory we develop later.

3.1 Symplectic vector spaces

The prototypical example of a symplectic vector space is $\mathbb{R}^{2n}$ with the standard form ω_0 from (2.16). This form has two important properties that we wish to keep.

Definition 3.1. A bilinear form $\beta : V \otimes V \to \mathbb{R}$ is called

1. alternating if for all $v, w \in V, \beta(v, w) = -\beta(w, v)$.
2. nondegenerate if for all $0 \neq v \in V, \exists w \in V$ such that $\beta(v, w) \neq 0$.

We quickly check that ω_0 actually has these properties. The alternating property follows from the definition of the wedge product. To check for nondegeneracy, take any $z \in \mathbb{R}^{2n}$. It follows from (2.17) that

$$\omega_0(z, J_0 z) = -z^T J_0^2 z = \langle z, z \rangle = ||z||^2,$$

which is positive if $z \neq 0$. This shows that ω_0 is both alternating and nondegenerate. We now use these properties to define the general notion of a symplectic vector space.

Definition 3.2. A symplectic vector space is a pair (V, ω), where V is a finite dimensional vector space, and $\omega : V \otimes V \to \mathbb{R}$ is a nondegenerate alternating bilinear form on V.

As remarked earlier, $(\mathbb{R}^{2n}, \omega_0)$ is the canonical example of a symplectic vector space. In fact, the goal of this subsection will be to show that up to isomorphism, this is the only symplectic vector space. One might wonder why we are not considering spaces of arbitrary dimension, such as $(\mathbb{R}^n, \omega)$. The next proposition excludes this possibility.

Proposition 3.3. *If (V, ω) is a symplectic vector space, V is of even dimension.*

Proof. Assume $\dim V = m$. Fixing some basis of V, we can write ω as

$$\omega(x, y) = x^T A y$$

for some matrix $A \in M_{m \times m}(\mathbb{R})$. The alternating property now gives that for any $x, y \in V$

$$x^T A y = \omega(x, y) = -\omega(y, x) = -y^T A x = -x^T A^T y,$$

which implies that $A^T = -A$. Now

$$\det(A) = \det(A^T) = \det(-A) = (-1)^m \det(A).$$

So if m is odd, A is singular, and ω is degenerate. $\square$

We are now ready to define the notion of equivalence in linear symplectic geometry. Notice the similarity to (2.18).

Definition 3.4. A linear symplectomorphism of symplectic vector spaces $(V_0, \omega_0), (V_1, \omega_1)$ is a vector space isomorphism

$$\Psi : V_0 \to V_1$$

that preserves the symplectic structure in the sense that

$$\Psi^* \omega_1 = \omega_0. \tag{3.1}$$

The (auto)symplectomorphisms of (V, ω) form a group denoted $\mathrm{Sp}(V, \omega)$, and we use the shorthand $\mathrm{Sp}(2n) = \mathrm{Sp}(\mathbb{R}^{2n}, \omega_0)$ for the standard space.

Remark 3.5. *Note that this definition is not the one found in [MS98], which only considers symplectomorphisms in the automorphism sense. Our definition agrees with the one found in [CdS06]. We have chosen it because it is more general, and will make certain statements more concise.*

We also want to study some special types of subspaces of symplectic vector spaces. To this extent we define the symplectic complement of a linear subspace. This is a symplectic analogue of the orthogonal complement.

Definition 3.6. The symplectic complement of a linear subspace $W \subset V$ is the subspace

$$W^\omega = \{v \in V : \omega(v, w) = 0 \ \ \forall w \in W\}. \tag{3.2}$$

A subspace $W \subset V$ is called:

1. Isotropic if $W \subset W^\omega$.
2. Coisotropic if $W \supset W^\omega$.
3. Symplectic if $W \cap W^\omega = \{0\}$.
4. Lagrangian if $W = W^\omega$.

In other words, W is isotropic if and only if ω vanishes on W and W is Lagrangian if and only if it is both isotropic and coisotropic. A subset W is symplectic if and only if $\omega\big|_W$ is nondegenerate, which means that $(W, \omega\big|_W)$ is a symplectic vector space.

The next result highlights the similarity of the symplectic and orthogonal complements. It uses an important fact we will see many times later, namely that ω gives rise to an explicit isomorphism of V and V^*.

Lemma 3.7. *For any subspace $W \subset V$*

$$\dim W + \dim W^\omega = \dim V$$

and

$$(W^\omega)^\omega = W$$

Proof. Define a map

$$\iota_\omega : V \to V^* : v \mapsto \omega(v, -), \tag{3.3}$$

It is not hard to see that the nondegeneracy of ω implies that ι_ω is an isomorphism. We now claim that

$$\iota_\omega(W^\omega) = W^\perp = \{l \in V^* : l(W) = 0\}.$$

To see this, note that if $v \in W^\omega, w \in W$,

$$\begin{aligned} \iota_\omega(v)(w) = \omega(v, w) = 0 \\ \implies \iota_\omega(v) \in W^\perp. \end{aligned}$$

If $l \in W^\perp, v = {\iota_\omega}^{-1}(l)$

$$\begin{aligned} \omega(v, w) = \iota_\omega(v)(w) = l(w) = 0 \qquad \forall w \in W \\ \implies v \in W^\omega. \end{aligned}$$

Thus $\dim W^\perp = \dim W^\omega$. A known result of standard linear algebra is that for finite dimensional vector spaces $\dim W + \dim W^\perp = \dim V$. This proves the first part of the lemma. To prove the second part note that clearly $W \subset (W^\omega)^\omega$ and since the first part of the lemma gives

$$\dim W = \dim V - \dim W^\omega = \dim(W^\omega)^\omega,$$

we must have $W = (W^\omega)^\omega$. $\square$

We are now ready for the main result of this subsection. It is a symplectic analogue of Gram-Schmidt, and we will see that an immediate consequence is that all symplectic vector spaces of the same dimension are symplectomorphic – that is, the only linear symplectic invariant is dimension. This result is highly instructive since it will also hold locally in the smooth case. The nonexistence of local invariants is a defining feature of symplectic geometry, and contrasts it with Riemannian geometry.

Proposition 3.8. *Any symplectic vector space has a symplectic basis. More precisely, let (V, ω) be a symplectic vector space with $\dim(V) = 2n$. Then there exists a basis*

$$u_1, ..., u_n, v_1, ..., v_n$$

of V, such that for all $1 \leq j, k \leq n$

$$\begin{aligned} \omega(u_j, u_k) = \omega(v_j, v_k) = 0 \\ \omega(u_j, v_k) = \delta_{jk}, \end{aligned}$$

where δ_{jk} denotes the Kronecker delta.

Proof. We proceed by induction over n.
Base case, n=1
Taking any $u \in V$, the nondegeneracy of ω guarantees the existence of some $v \in V$ such that $\omega(u, v) = 1$. The alternating property means that u and v must be linearly independent, so they are a basis.
Induction step
Assume any symplectic vector space of dimension $2n-2$ has a symplectic basis. As in the base case, choose any u_1, v_1 such that $\omega(u_1, v_1) = 1$. It is now easy to see that $W = \mathrm{span}(u_1, v_1)$ is a symplectic subspace. Using the previous lemma, we have that

$$W \oplus W^\omega = V,$$

thus $(W^\omega, \omega\big|_{W^\omega})$ is a symplectic vector space of dimension $2n - 2$. By the induction hypothesis, there exists a symplectic basis

$$u_2, ..., u_n, v_2, ..., v_n$$

of $(W^\omega, \omega\big|_{W^\omega})$, and since $w(v, w) = 0$ for all $v \in W^\omega, w \in W$,

$$u_1, u_2, ..., u_n, v_1, v_2, ..., v_n$$

is a symplectic basis of (V, ω). □

Corollary 3.9. *For any symplectic vector space (V, ω) of dimension $2n$, there exists a linear symplectomorphism of $(\mathbb{R}^{2n}, \omega_0)$ with (V, ω).*

Proof. Let $\{u_j, v_j\}_{j=1}^n$ be a symplectic basis of (V, ω). Using the standard symplectic coordinates on $\mathbb{R}^{2n}$, let

$$\Psi(z) = \sum_{j=1}^{n} x_j u_j + y_j v_j.$$

A straightforward calculation shows that this is indeed a symplectomorphism. □

3.2 The symplectic linear group

We now turn our attention to the group $\mathrm{Sp}(V, \omega)$. Since we just showed that any symplectic vector space is symplectomorphic to $(\mathbb{R}^{2n}, \omega_0)$, we just need to study $\mathrm{Sp}(2n)$. Using the standard basis, we can think of an element of $\mathrm{Sp}(2n)$ as a matrix. But what makes a matrix a symplectomorphism? An important property of ω_0 is that its nth exterior power is the volume form on $\mathbb{R}^{2n}$ – that is,

$$\omega_0^n = \left(\sum \mathrm{d}x_j \wedge \mathrm{d}y_j\right)^n = n!\mathrm{d}x_1 \wedge \mathrm{d}y_1 \wedge ... \wedge \mathrm{d}x_n \wedge \mathrm{d}y_n. \tag{3.4}$$

Let $\Psi \in \mathrm{Sp}(2n)$.Since pullback distributes over the wedge product we get

$$\Psi^* \omega_0^n = (\Psi^* \omega_0)^n = \omega_0^n.$$

So Ψ preserves the volume form on $\mathbb{R}^{2n}$, thus $\det(\Psi) = 1$. This means that all symplectomorphisms are volume preserving, but the converse does not hold. Here is a more precise condition:

Lemma 3.10. *A matrix Ψ is a symplectomorphism if and only if*

$$\Psi^T J_0 \Psi = J_0, \tag{3.5}$$

where J_0 is as in (2.12).

Proof. Using (2.17) we get that for all $z, z' \in \mathbb{R}^{2n}$,

$$\begin{aligned} \omega_0(z, z') &= -z^T J_0 z \\ \Psi^* \omega_0(z, z') &= \omega_0(\Psi z, \Psi z') = -z^T \Psi^T J_0 \Psi z. \end{aligned}$$

Hence $\Psi^T J_0 \Psi = J_0$ if and only if $\Psi^* \omega_0 = \omega_0$. □

The above proposition imposes restrictions on the shape of inverses of symplectic matrices. Let

$$\Psi = \begin{pmatrix} A & B \\ C & D \end{pmatrix},$$

where A, B, C and D are real $n \times n$ matrices. Using the proposition and the fact that $J_0^2 = -I_{2n}$, we get

$$\begin{aligned} \Psi^{-1} &= -\, J_0 \Psi^T J_0 \\ \implies \Psi^{-1} &= -\, J_0 \begin{pmatrix} A^T & C^T \\ B^T & D^T \end{pmatrix} J_0 \\ \implies \Psi^{-1} &= \begin{pmatrix} D^T & -B^T \\ -C^T & A^T \end{pmatrix} \end{aligned}$$

In the case $n = 1$, this equation implies that any matrix with determinant 1 is symplectic. I.e. $\mathrm{Sp}(2) = \mathrm{Sl}(2)$ This hints at the fact that symplectomorphisms preserve some kind of two-dimensional area, as well as 2n-dimensional volume.

As mentioned in Section 2, J_0 is the standard complex structure on $\mathbb{R}^{2n}$, meaning that if one associates $\mathbb{R}^{2n}$ with $\mathbb{C}^n$ by letting $z_j = x_j + iy_j$, J_0 corresponds to multiplication with i. Under this identification the standard symplectic form looks like

$$\omega_0(z, z') = \langle iz, z \rangle$$

Unitary matrices on $\mathbb{C}^n$ preserve inner products, so it is clear that under the above identification, $\mathrm{U}(n) \subset \mathrm{Sp}(n)$. The next lemma makes this relationship more clear. We will continue to use the same conventions for notation, i.e $\mathrm{U}(n)$ and $\mathrm{Gl}(n, \mathbb{C})$ will denote the corresponding subsets of $\mathrm{Gl}(2n, \mathbb{R})$ under the standard identification. Note that a complex $n \times n$ matrix $C = A + iB$ corresponds to the real $2n \times 2n$ matrix $\left(\begin{smallmatrix} A & -B \\ B & A \end{smallmatrix}\right)$.

Lemma 3.11 (Two out of three property)**.**

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

Proof. Any real $2n \times 2n$ matrix Ψ satisfies:

1. $\Psi \in \mathrm{Gl}(n, \mathbb{C}) \iff \Psi J_0 = J_0 \Psi$ (Ψ is complex linear).
2. $\Psi \in \mathrm{Sp}(2n) \iff \Psi^T J_0 \Psi = J_0$.
3. $\Psi \in \mathrm{O}(2n) \iff \Psi^T \Psi = I_{2n}$.

Note that any two of the conditions on the right imply the third. This proves the first part of the lemma. We now show that $\mathrm{O}(2n) \cap \mathrm{Gl}(n, \mathbb{C}) = \mathrm{U}(n)$. Under the identification,

$$U^* = (A + iB)^* = A^T - iB^T \backsim \begin{pmatrix} A^T & B^T \\ -B^T & A^T \end{pmatrix} = \begin{pmatrix} A & -B \\ B & A \end{pmatrix}^T.$$

Hence adjoints in $\mathrm{Gl}(n, \mathbb{C})$ correspond to transposes in $\mathrm{Gl}(2n, \mathbb{R})$. This means that unitary matrices are orthogonal, and that orthogonal matrices of the form $\left(\begin{smallmatrix} A & -B \\ B & A \end{smallmatrix}\right)$ are unitary. □

We also have some restrictions on the eigenvalues of symplectic matrices. Note how this also hints at the preservation of area we mentioned earlier.

Lemma 3.12. *If λ is an eigenvalue of $\Psi \in \mathrm{Sp}(2n)$, so is λ^{-1}. Moreover λ and λ^{-1} have the same multiplicities. Both 1 and -1 must have even multiplicities.*

Proof. From proposition 3.10 we know that Ψ^T and Ψ^{-1} are similar, so any eigenvalue λ must have the same multiplicity in Ψ and Ψ^{-1}. Since the multiplicity of λ in Ψ is equal to the multiplicity of λ^{-1} in Ψ^{-1} the first part is finished. Since $\det(\Psi) = 1$, and the determinant is the product of the eigenvalues, -1 must have even multiplicity. Since Ψ has even rank, and we have eliminated an even number of eigenvalues, 1 must also have even multiplicity. □

This can be very restrictive, just remember that for any complex eigenvalue λ of any linear transformation A, the multiplicities of λ and $\overline{\lambda}$ must be equal. This means that the eigenvalues of linear symplectomorphisms occur either in pairs $\lambda, \frac{1}{\lambda}$ with $\lambda \in \mathbb{R}$, in pairs $\lambda, \overline{\lambda}$ with $|\lambda| = 1$, or in quadruplets $\lambda, \overline{\lambda}, \lambda^{-1}, \overline{\lambda^{-1}}$.

The fact that λ and λ^{-1} occurs in pairs can be reformulated informally as "If you squeeze in some direction, you must stretch in some other direction" we now show that the choice of direction is in some sense determined by the symplectic form.

Lemma 3.13. *Let $\Psi \in \mathrm{Sp}(2n)$. If*

$$\Psi z = \lambda z, \Psi z' = \lambda' z',$$

then

$$\lambda\lambda' \neq 1 \implies \omega_0(z, z') = 0.$$

Proof. Since Ψ is a symplectomorphism, we have

$$\begin{aligned}\omega_0(z,z') &= \Psi^*\omega_0(z,z')\\ &= \omega_0(\lambda z,\lambda' z')\\ &= \lambda\lambda'\omega_0(z,z').\end{aligned}$$

Hence, either $(1-\lambda\lambda') = 0$ or $\omega_0(z,z') = 0$. $\square$

We now formulate one last lemma, which we will need to prove the main result of this subsection.

Lemma 3.14. *If $P \in \mathrm{Sp}(2n)$ is a symmetric, positive definite matrix, then $P^\alpha \in Sp(2n)$ for all $\alpha > 0$.*

Proof. Since P is symmetric and nonsingular, we can decompose $\mathbb{R}^{2n}$ as a direct sum of the eigenspaces of P:

$$\mathbb{R}^{2n} \simeq \bigoplus_{\lambda\in\sigma(P)} V_\lambda.$$

Each V_λ is also an eigenspace of P^α corresponding to the eigenvalue λ^α. By the previous lemma, if $\lambda\lambda' \neq 1$, then V_λ and $V_{\lambda'}$ are ω_0 orthogonal. In particular, since P is positive definite, we have $-1 \notin \sigma(P)$. This means that ω_0 vanishes on each V_λ. Now let $z \in V_\lambda, z' \in V_{\lambda'}$. Since we must have either $\omega_0(z,z') = 0$ or $\lambda\lambda' = 1$, we get that

$$\omega_0(P^\alpha z, P^\alpha z') = (\lambda\lambda')^\alpha\omega_0(z,z') = \omega_0(z,z').$$

We can now pick a basis of eigenvectors, and see that by linearity of ω_0, P^α is a symplectomorphism. $\square$

Proposition 3.15. $\mathrm{Sp}(2n)/\mathrm{U}(n)$ *is contractible.*

Proof. Any $\Psi \in \mathrm{Sp}(2n)$ has a polar decomposition, $\Psi = PQ$ where P is symmetric, positive definite and Q is orthogonal. In this construction,

$$P = (\Psi\Psi^T)^{\frac{1}{2}}, Q = P^{-1}\Psi,$$

so by previous lemma, both P and Q are symplectic. We now define a map

$$\begin{aligned}\mathrm{Sp}(2n)\times[0,1] &\to \mathrm{Sp}(2n)\\ (\Psi,t) &\mapsto (\Psi\Psi^T)^{-\frac{t}{2}}\Psi.\end{aligned}$$

At $t = 0$ this map is just the identity. At $t = 1$ it maps any Ψ to Q, the orthogonal part of its polar decomposition. From lemma 3.11 we know that $Q \in \mathrm{U}(n)$. Thus the map gives a deformation retraction of $\mathrm{Sp}(2n)$ into $\mathrm{U}(n)$. $\square$

3.3 Compatible complex structures

We have seen that the form ω_0 on $\mathbb{R}^{2n}$ was generated by the matrix J_0. This was very useful, as it allowed us to relate ω to the standard inner product on $\mathbb{R}^{2n}$. This situation can be generalized.

Definition 3.16. A complex structure on a real, finite dimensional vector space V is a linear automorphism $J : V \to V$ such that

$$J^2 = -id_V.$$

Such a structure allows us to consider V as a complex vector space by letting J correspond to multiplication by i – that is, we define a scalar multiplication map

$$\begin{aligned} \mathbb{C} \otimes V &\to V \\ (a+ib) \otimes v &\mapsto av + bJv. \end{aligned}$$

As one might expect, we can only find such structures for even dimensional spaces.

Lemma 3.17. *If J is a complex structure on V,* $\dim(V) = 2n$.

Proof. Assume $\dim(V) = m$. Then,

$$\det(J)^2 = \det(J^2) = \det(-id_V) = (-1)^m,$$

Since $\det(J) \in \mathbb{R}$, we must have $m = 2n$. □

The complex structure J_0 on $\mathbb{R}^{2n}$ was special in the sense that it was "compatible" with the form ω_0. We now generalize this situation.

Definition 3.18. Let (V, ω) be a symplectic vector space. A complex structure $J : V \to V$ is said to be compatible with ω if

$$J^*\omega = \omega \tag{3.6}$$

$$\omega(v, Jv) > 0 \quad \forall v \in V : v \neq 0. \tag{3.7}$$

If these conditions are satisfied, the bilinear form

$$g_J(v, w) = \omega(v, Jw) \tag{3.8}$$

is a well defined inner product on V since

$$g_J(v,w) = \omega(v,Jw) = \omega(-J^2v, Jw) = -J^*\omega(Jv, w) = \omega(w, Jv) = g_J(w, v)$$
$$g_J(v,v) = \omega(v, Jv) > 0 \quad \forall v \in V : v \neq 0.$$

Given a symplectic structure ω, an inner product g and a complex structure J on a vector space V, we say that (g, ω, J) is a compatible triple if for all $v, w \in V$,

$$g(v, w) = \omega(v, Jw)$$
$$\omega(v, w) = g(Jv, w)$$
$$J(v) = \iota_g^{-1} \circ \iota_\omega(v).$$

The maps $\iota_\omega, \iota_g : V \to V^*$ given by $\iota_\omega(v)(w) = \omega(v, w), \iota_g(v)(w) = g(v, w)$ are isomorphisms due to the nondegeneracy of g and ω. We say that any two such structures are compatible if we can construct a structure of the third type from the above equations. I.e, any two compatible structures extends uniquely to a compatible triple. Note the similarity to the "two out of three" property of the unitary group from lemma 3.11. We now show that any symplectic vector space admits a compatible complex structure.

Proposition 3.19. *Every symplectic vector space (V, ω) has a compatible complex structure J.*

Proof. We know that any vector space has an inner product, but we are not guaranteed that it is compatible with ω. The idea will be to fix an inner product g, and start with the automorphism

$$A = \iota_g^{-1} \circ \iota_\omega.$$

This can be modified this to get a compatible complex structure. A satisfies the following equation by definiton;

$$g(Av, w) = \omega(v, w)$$

Thus the alternating property of ω implies that $A^* = -A$, where A^* denotes the g-adjoint of A. As in the proof of proposition 3.15, take J to be the g-orthogonal part of the polar decomposition of A – that is, $J = A(A^*A)^{-\frac{1}{2}}$. This must also satisfy $J^* = -J$, so now

$$J^2 = -JJ^* = -id_V$$

– that is, J is a complex structure on V. We now check if it is ω compatible. From elementary linear algebra we know that since J is defined in terms of powers and inverses of A^*A, we have $JA = AJ$. Combining this with the orthogonality of J we get that

$$\omega(Jv, Jw) = g(AJv, Jw) = g(JAv, Jw) = g(Av, w) = J(v, w).$$

From the definition of J, we compute that for all nonzero $v \in V$

$$\omega(v, Jv) = g(Av, A(A^*A)^{-\frac{1}{2}}v) = g(v, (A^*A)(A^*A)^{-\frac{1}{2}}v = g(v, (A^*A)^{\frac{1}{2}}v)$$

which is positive since A^*A is positive definite. Hence J is a compatible complex structure. $\square$

Remark 3.20. *Note that most of the theory from this chapter can immediately be generalized to vector bundles* $\pi : E \to B$*. In this setting a symplectic structure is defined as a smooth section* ω *of the tensor bundle* $T^2_0 E$ *such that each fiber* $(\pi^{-1}(q), \omega_q)$ *is a symplectic vector space. Importantly, proposition 3.19 generalizes to say that any symplectic vector bundle admits a smooth section* J *of the tensor bundle* $T^1_1 E$ *such that* J_q *is a compatible complex structure on each* $(\pi^{-1}(q), \omega_q)$*. In the language of structure groups, we can see that proposition 3.19 really is just a corollary to proposition 3.15. This result reduces the structure group* $\mathrm{Sp}(2n)$ *of a symplectic bundle to the unitary group* $\mathrm{U}(n)$*, meaning that all symplectic bundles can be represented as complex bundles.*

4 Symplectic manifolds

We now move on to the smooth theory. This section will be concerned with the first consequences of symplectic structure, generalizing Hamiltonian flows, and showing that the local structure is in fact trivial.

4.1 Basic concepts

Definition 4.1. A symplectic structure on a smooth manifold M is a nondegenerate, closed two-form $\omega \in \Omega^2(M)$ – that is, each tangent space (T_pM, ω_p) is a symplectic vector space, and $\mathrm{d}\omega = 0$.

Remark 4.2. *Note that ω is a symplectic structure on the vector bundle TM.*

The form ω must satisfy two conditions. Non-degeneracy is an algebraic condition, and is local in nature. Closedness on the other hand, is a geometric condition that is global in nature. Many of the algebraic properties implied by non-degeneracy carry over from the linear case to the smooth case:

Proposition 4.3. *If (M, ω) is a symplectic manifold the following hold:*

1. $\dim(M) = 2n$.
2. *M is orientable.*
3. *$\iota_\omega : TM \to T^*M : (p, v) \mapsto (p, \iota_{\omega_p}(x)) = (p, \omega_p(x, -))$ is an isomorphism of vector bundles.*

Proof. By proposition 3.3, we must have $\dim(T_pM) = 2n$. It follows that $\dim(M) = 2n$. From equation (3.4), we know that ω^n is a volume form on M, so M must be orientable. The map ι_ω is just the identity on the base-space, and as remarked before, the nondegeneracy clearly implies that each ι_{ω_p} is an isomorphism. □

This shows that the class of manifolds that admit a symplectic structure is quite restricted. Closedness imposes some further restrictions.

Proposition 4.4. *Let (M, ω) be a symplectic manifold. If M is a closed manifold (A manifold M is closed if it is compact and $\partial M = \emptyset$), then*

$$H^2(M, \mathbb{R}) \neq 0.$$

Proof. Since ω is closed, it represents some cohomology class

$$a = [\omega] \in H^2(M; \mathbb{R}).$$

Now a^n is represented by ω^n, and since w^n is a volume form,

$$\int_M \omega^n \neq 0.$$

This implies that ω^n is not exact, which implies that ω is not exact. Hence $0 \neq a \in H^2(M; \mathbb{R})$ □

Proposition 4.4 eliminates the possibility of having a symplectic structure for many closed even-dimensional manifolds, for instance S^2 is the only sphere that admits a symplectic structure.

Example 4.5. The $2n$-torus $\mathbb{T}^{2n}$ has a natural symplectic structure induced by the universal cover $q : \mathbb{R}^{2n} \to \mathbb{R}^{2n}\mathbb{Z}^{2n} \simeq \mathbb{T}^{2n}$. Since the standard form ω_0 on $\mathbb{R}^{2n}$ is invariant under translations, we can use the local inverses of q to define a two-form ω on $\mathbb{T}^{2n}$ such that $q^*\omega = \omega_0$. Using the coordinates $\theta_j, \varphi_j : \mathbb{T}^{2n} \to S^1$, we may express ω as

$$\omega = \sum_{j=0}^{n} \mathrm{d}\theta_j \wedge \mathrm{d}\varphi_j.$$

Note the similarity to the standard form ω_0.

4.2 Symplectomorphisms

As in the linear theory, we define symplectomorphisms as isomorphisms that preserve the symplectic structure. As in the linear case, our definition agrees with [CdS06] rather than [MS98].

Definition 4.6. A symplectomorphism of symplectic manifolds (M_0, ω_0) and (M_1, ω_1) is a diffeomorphism

$$\psi : M_0 \to M_1,$$

such that the symplectic structure is preserved – that is,

$$\psi^*\omega_1 = \omega_0. \tag{4.1}$$

Note that this is equivalent to requiring that ψ is a diffeomorphism such that

$$\mathrm{d}\psi : T_p(M_0) \to T_{\psi(p)}(M_1) \tag{4.2}$$

is a linear symplectomorphism for each $p \in M_0$. We denote the group of symplectomorphisms of (M, ω) with itself by $\mathrm{Symp}(M, \omega)$.

How can we construct such symplectomorphisms? One approach is to recall the Hamiltonian flows of section 2. The key there was to construct a certain vector field and solve the first order differential equation associated to this field. Note that the map ι_ω from proposition 4.3 can be thought of as an isomorphism of the vector spaces of vector fields and one-forms on M – that is,

$$\begin{aligned} \iota_\omega : \chi(M) &\to \Omega^1(M) \\ X &\mapsto \iota_X(\omega) = \omega(X, -), \end{aligned}$$

where ι_X denotes the interior product with X.

Definition 4.7. A vector field $X \in \chi(M)$ is called symplectic if $\iota_X(\omega)$ is closed. We denote the space of symplectic vector fields by $\chi(M, \omega) \subset \chi(M)$.

The next proposition combines the closedness and non-degeneracy conditions to show that when M is closed, $\chi(M,\omega)$ is exactly the Lie-algebra of $\mathrm{Symp}(M,\omega)$.

Proposition 4.8. *Let (M,ω) be a closed, symplectic manifold. If $\psi_t \in \mathrm{Diff}(M)$ is a smooth 1-parameter family of diffeomorphisms generated by the smooth family of vector fields $X_t \in \chi(M)$ via*

$$\frac{\mathrm{d}}{\mathrm{d}t}\psi_t = X_t \circ \psi_t, \qquad \psi_0 = id_M, \tag{4.3}$$

then

$$\psi_t \in \mathrm{Symp}(M,\omega) \quad \forall t \iff X_t \in \chi(M,\omega) \quad \forall t.$$

In other words, symplectic isotopies are precisely the flows of symplectic vector fields.

Proof. The Cartan formula for the Lie-derivative of a differential form is

$$\mathcal{L}_X(\omega) = \iota_X(\mathrm{d}\omega) + \mathrm{d}\iota_X(\omega). \tag{4.4}$$

This can be generalized to time dependent families of vector fields via the identity

$$\frac{\mathrm{d}}{\mathrm{d}t}\psi_t^*\omega = \psi_t^*\,\mathcal{L}_{X_t}(\omega), \tag{4.5}$$

which holds whenever X_t and ψ_t satisfy (4.3). Outlines for proofs of both these identities can be found in [CdS06]. In short, the proof boils down to showing both sides are derivations of the algebra $(\Omega^\bullet(M), \wedge)$, which both commute with d, and agree on 0-forms. Combining them, and using the fact that ω is closed, we get that

$$\frac{\mathrm{d}}{\mathrm{d}t}\psi_t^*\omega = \psi_t^*\mathrm{d}\iota_{X_t}(\omega).$$

Now, note that ψ_t is a symplectomorphism for all t if and only if the left hand side is zero. And that $\iota_{X_t}(\omega)$ is closed for all t if and only if the right hand side is zero. One can also use these identities to show that $\chi(M,\omega)$ is closed under the Lie-bracket, but this will not be particularly relevant to us, so we omit the proof. ☐

For any symplectic vector field X_t, the form $\iota_{X_t}(\omega)$ is closed for all t. If this form is also exact for all t – that is, there exists some $H_t \in C^\infty(M,\mathbb{R})$ such that

$$\iota_{X_t}(\omega) = \mathrm{d}H_t, \tag{4.6}$$

we call $X_t = X_{H_t}$ the Hamiltonian vector field generated by H_t. Comparing this to (2.19) we see that this is a direct generalization of the classical case.

Definition 4.9. The one parameter family of diffeomorphisms ψ_t is called a Hamiltonian isotopy if it is generated by the Hamiltonian vector fields X_{H_t} via (4.3). A symplectomorphism $\psi : M \to M$ is called a Hamiltonian symplectomorphism if there exist a Hamiltonian isotopy ψ_t such that $\psi_0 = id_M, \psi_1 = \psi$. We denote by $\mathrm{Ham}(M,\omega)$ the subgroup of Hamiltonian symplectomorphisms.

It is by definition that the subspace of Hamiltonian vector fields is the Lie-algebra of $\mathrm{Ham}(M, \omega)$. It can be shown that the Lie bracket on this algebra is given by the Poisson bracket of the generating Hamiltonian functions.

4.3 Cotangent bundles

One natural generalization of the $\mathbb{R}^{2n}$ case is to cotangent bundles T^*L of smooth manifolds L. In physics one would interpret this as L describing the "position" of a system, while the cotangent vector describes momentum. This interpretation has a natural formulation in symplectic geometry through the canonical one-form λ_{can}. At any point $(x, \sigma) \in T^*L$, define

$$\lambda_{can(x,\sigma)} = \pi^*\sigma, \tag{4.7}$$

where $\pi : T^*L \to L$ is the projection. In other terms, if

$$(v, \tau) \in T_{(x,\sigma)}T^*L \simeq T_xL \oplus T_x^*L,$$

then

$$\lambda_{can(x,\sigma)}(v, \tau) = \sigma(v).$$

This canonical one-form is useful as it gives rise to a canonical closed two-form $\omega_{can} = -\mathrm{d}\lambda_{can}$. If we can show that this form is nondegenerate, we would have a canonical symplectic structure on the cotangent bundle. To this extent it is useful to describe λ_{can} in terms of local coordinates. If $x : U \to L$ is a chart on the open set $U \subset L$, any $\sigma \in T_q^*L$ can be described uniquely in terms of the basis vectors $\mathrm{d}x_j$ as

$$\sigma = \sum_{j=1}^{n} y_j \mathrm{d}x_j.$$

The y_j's are uniquely determined by, and smoothly dependent on σ. This gives rise to a new coordinate function $(x, y) : T^*U \to \mathbb{R}^{2n}$. In these coordinates, it is not hard to see that the canonical one-form is

$$\lambda_{can} = \sum_{j=1}^{n} y_j \mathrm{d}x_j. \tag{4.8}$$

A quick computation now shows that on T^*U,

$$\omega_{can} = -\mathrm{d}\lambda_{can} = \sum_{j=1}^{n} \mathrm{d}x_j \wedge \mathrm{d}y_j \tag{4.9}$$

Note that this is analogous to the definition of ω_0 on $\mathbb{R}^{2n}$ given in (2.16). In fact, ω_0 is exactly the canonical two-form on $T^*\mathbb{R}^n$. Non-degeneracy is a local matter, so the argument for nondegeneracy of w_0 found in section 2 goes through for ω_{can} as well. Note the striking similarity of the expressions for σ and λ_{can}. This similarity can be formalized as follows.

Proposition 4.10. *The one-form $\lambda_{can} \in \Omega^1(T^*L)$ is characterized by*

$$\sigma^*\lambda_{can} = \sigma \tag{4.10}$$

for every one-form σ. Note that on the right we consider σ as a form, while on the left we think of it as a map $\sigma : L \to T^*L$.

Proof. On some coordinate patch $x : U \to \mathbb{R}^n, U \subset L$, any one-form σ is given uniquely by

$$\sigma = \sum_{j=1}^{n} a_j(x)\mathrm{d}x_j$$

Where a_j are smooth functions on U. In these coordinates the map $\sigma : L \to T^*L$ is given by

$$x = (x_1, .., x_n) \mapsto (x_1, .., x_n, a_1(x), .., a_n(x)) = (x, a(x)),$$

and its derivative is

$$\mathrm{d}\sigma = \begin{pmatrix} I_n \\ \mathrm{d}a \end{pmatrix}.$$

We now evaluate at the basis vector $\frac{\partial}{\partial x_j} \in T_xL$ to get

$$\begin{aligned}
(\sigma^*\lambda_{can})_x(\frac{\partial}{\partial x_j}) &= \lambda_{can,\sigma(x)}(\mathrm{d}\sigma_x(\frac{\partial}{\partial x_j})) \\
&= \sum_{k=1}^{n} a_k(x)\mathrm{d}x_k(\mathrm{d}\sigma_x(\frac{\partial}{\partial x_j})) \\
&= a_j(x) = \sigma(\frac{\partial}{\partial x_j}).
\end{aligned}$$

The forms $(\sigma^*\lambda_{can})_x$ and σ_x agree on a basis, so by linearity they are equal. □

It would be useful to have a more compact representation of ω_{can}. The following lemma gives us just that.

Lemma 4.11. *Let $v = (v_0, v_1^*), w = (w_0, w_1^*) \in T_{(q,\tau)}T^*L \simeq T_qL \oplus T_q^*L$. Then*

$$\omega_{can}(v, w) = w_1^*(v_0) - v_1^*(w_0). \tag{4.11}$$

Proof. In terms of the local coordinates $x, y : T^*U \to \mathbb{R}^{2n}$, the linear combinations

$$v = \sum_{j=1}^{n} a_j \frac{\partial}{\partial x_j} + a_j^*\mathrm{d}x_j, \qquad w = \sum_{j=1}^{n} b_j \frac{\partial}{\partial x_j} + b_j^*\mathrm{d}x_j$$

are represented by the coordinate vectors

$$v = (a_1, .., a_n, a_1^*, ..., a_n^*), \qquad w = (b_1, .., b_n, b_1^*, ..., b_n^*).$$

So

$$
\begin{aligned}
(\omega_{can})_q(v,w) &= v^T J_0 w \\
&= \sum_{j=1}^n a_j b_j^* - a_j^* b_j \\
&= \sum_{j=1}^n b_j^* \mathrm{d}x_j \left(\sum_{k=1}^n a_k \frac{\partial}{\partial x_k} \right) - \sum_{j=1}^n a_j^* \mathrm{d}x_j \left(\sum_{k=1}^n b_k \frac{\partial}{\partial x_k} \right) \\
&= w_1^*(v_0) - v_1^*(w_0). \qquad \square
\end{aligned}
$$

We have seen that cotangent bundles have a natural symplectic structure. We will exploit this structure later, but first we need to relate it to arbitrary symplectic manifolds. One of the nice features is that cotangent bundles have a clear difference between position and momentum; position is the base space, and momentum the fibers. This may not be well defined on an arbitrary manifold, but we will show in the next section that compact Lagrangian submanifolds give a local notion of position. Another nice feature of cotangent bundles is the relationship $-\mathrm{d}\lambda_{can} = \omega_{can}$. This allows one to define an analogue of the action integral (2.20), which means that we can employ variational techniques similar to those in section 2.

4.4 Local theory

The goal of this subsection is to classify the local structure of symplectic manifolds. The motivation will be Darboux's theorem stating that all symplectic manifolds of the same dimension are locally symplectomorphic. To this end we will develop the Moser argument of homotopies of forms. This method will turn out to be very useful, as it will also allow us to classify the structure of neighbourhoods of certain submanifolds.

Lemma 4.12 (Moser argument)**.** *Let M be a smooth manifold. If ω_t is some time dependent family of forms satisfying*

$$\frac{\mathrm{d}}{\mathrm{d}t}\omega_t = \mathrm{d}\sigma_t \tag{4.12}$$

for some family of forms $\sigma_t \in \Omega^1(M)$, there exists a family of symplectomorphisms ψ_t of (M, ω_t) with (M, ω_0) – that is, a family of diffeomorphisms such that

$$\psi_t^* \omega_t = \omega_0.$$

Proof. For this proof we will use the following formula for the Lie derivative of a time dependent family of forms along a time dependent family of vector fields. See proposition 6.4 in [CdS06]. If X_t and ψ_t satisfy the first order differential equation (4.3), we have

$$\frac{\mathrm{d}}{\mathrm{d}t}\psi_t^*\omega_t = \psi_t^*(\mathcal{L}_{X_t}(\omega_t) + \frac{\mathrm{d}}{\mathrm{d}t}\omega_t). \tag{4.13}$$

We will construct a vector field X_t such that the flow ψ_t generated by X_t satisfies

$$\begin{aligned}0 &= \frac{\mathrm{d}}{\mathrm{d}t}\psi_t^*\omega_t \\ &= \psi_t^*(\mathcal{L}_{X_t}(\omega_t) + \frac{\mathrm{d}}{\mathrm{d}t}\omega_t) \\ &= \psi_t^*(\mathrm{d}\iota_{X_t}(\omega_t) + \mathrm{d}\sigma_t) \\ &= \mathrm{d}\psi_t^*(\iota_{X_t}(\omega_t) + \sigma_t).\end{aligned}$$

Taking $X_t = \iota_{\omega_t}^{-1}(-\sigma_t)$, the above equation holds. So $\psi_t^*\omega_t$ is constant, and at $t = 0$ it is $\psi_0^*\omega_0 = id^*\omega_0 = \omega_0$ □

We are interested in a special case of this. The goal is to find a symplectic isotopy fixing some compact submanifold, but making two forms agree on a neighbourhood of the submanifold.

Proposition 4.13 (Relative Moser argument)**.** *Let M be a smooth manifold with $\dim(M) = 2n, Q \subset M$ a compact submanifold. Suppose $\omega_0, \omega_1 \in \Omega^2(M)$ are nondegenerate and equal on T_qM whenever $q \in Q$. Then there exist open neighbourhoods N_0, N_1 of Q, and a diffeomorphism $\psi : N_0 \to N_1$ such that*

$$\psi\big|_Q = id, \qquad \psi^*(\omega_1\big|_{N_1}) = \omega_0\big|_{N_0}.$$

Proof. If we can find a one-form $\sigma \in \Omega^1(N_0)$, where N_0 is some open neighbourhood of Q, such that

$$\forall q \in Q : \sigma\big|_{T_qM} = 0, \qquad \mathrm{d}\sigma = \omega_1 - \omega_0, \tag{4.14}$$

we can apply the Moser argument to the family

$$\omega_t = \omega_0 + t(\omega_1 - \omega_0) = \omega_0 + t\mathrm{d}\sigma.$$

If necessary, shrink N_0 so that ω_t is nondegenerate on N_0, and so that the resulting family of diffeomorphisms ψ_t is defined on N_0 for all $t \in [0, 1]$. Looking at the construction in the Moser argument, it is clear that the time derivative of ψ_t will be zero on Q since our σ is zero on Q. Since $\psi_0 = id$, we must have $\psi_t\big|_Q = id$, and our result would follow.

To find a form σ satisfying the above conditions, consider the exponential map from the normal bundle of Q,

$$\exp : NQ \to M.$$

Consider an ϵ neighbourhood of the zero section in NQ,

$$U_\epsilon = \{(q, v) \in NQ : |v| < \epsilon\}.$$

By the tubular neighbourhood theorem and compactness, there exist an $\epsilon > 0$ such that $\exp\big|_{U_\epsilon}$ is a diffeomorphism to its image, which we define to be $N_0 =$

$\exp(U_\epsilon)$. For $t \in [0,1]$, define $\phi_t : N_0 \to N_0$ by the following diagram

$$\begin{tikzcd} U_\epsilon \arrow[r, "t"] \arrow[d, "\exp"] & U_\epsilon \arrow[d, "\exp"] \\ N_0 \arrow[r, "\phi_t"] & N_0 \end{tikzcd},$$

where t represents multiplication by t in each fiber. Then ϕ_t is a diffeomorphism for $t > 0$, and

$$\phi_0(N_0) \subset Q, \qquad \phi_1 = id_{N_0}, \qquad \phi_t\big|_Q = id_Q.$$

The form $\omega_1 - \omega_0$ is closed, and

$$\phi_0^*(\omega_1 - \omega_0) = 0, \qquad \phi_1^*(\omega_1 - \omega_0) = \omega_1 - \omega_0.$$

For $t > 0$, define $X_t \in \chi(N_0)$ to be the vector field generating the flow ϕ_t – that is,

$$X_t = \left(\frac{\mathrm{d}}{\mathrm{d}t}\phi_t\right) \circ \phi_t^{-1}.$$

Defining $\sigma_t = \phi_t^* \iota_{X_t}(\omega_1 - \omega_0) \in \Omega^1(N_0)$, we have $\sigma_t\big|_Q = 0$. Using (4.5), we get that

$$\begin{aligned} \frac{\mathrm{d}}{\mathrm{d}t}\phi_t^*(\omega_1 - \omega_0) &= \phi_t^* \, \mathcal{L}_{X_t}(\omega_1 - \omega_0) \\ &= \phi_t^*(\mathrm{d}\iota_{X_t}(\omega_1 - \omega_0)) \\ &= \mathrm{d}\sigma_t. \end{aligned}$$

Note that the form σ_t is given by

$$(\sigma_t)_q(v) = (\omega_1 - \omega_0)_{\phi_t(q)}\left(\frac{\mathrm{d}}{\mathrm{d}t}\phi_t(q), (\mathrm{d}\phi_t)_q(v)\right),$$

which is smooth and well defined for all $t \in [0,1]$, not just for $t > 0$ as we originally had. Assuming exchange of integration and differentiation, we get

$$\begin{aligned} \omega_1 - \omega_0 &= \phi_1^*(\omega_1 - \omega_0) - \phi_0^*(\omega_1 - \omega_0) \\ &= \int_0^1 \frac{\mathrm{d}}{\mathrm{d}t}\phi_t^*(\omega_1 - \omega_0)\mathrm{d}t \\ &= \int_0^1 \mathrm{d}\sigma_t \mathrm{d}t = \mathrm{d}\int_0^1 \sigma_t \mathrm{d}t. \end{aligned}$$

Taking

$$\sigma = \int_0^1 \sigma_t \mathrm{d}t,$$

we have satified the conditions in (4.14), and the proof is finished. □

After quite a bit of work, we are immediately rewarded with the following result.

Proposition 4.14 (Darboux' Theorem)**.** *Every symplectic manifold is locally symplectomorphic to* $(\mathbb{R}^{2n}, \omega_0)$*.*

Proof. Let (M, ω_1) be a symplectic manifold with $\dim(M) = 2n$. Any $q \in M$ has an open neighbourhood U with coordinates $(x, y) : U \to \mathbb{R}^{2n}$. Composing with the linear isomorphism from corollary 3.9 if necessary, we may assume that $\frac{\partial}{\partial x_j}, \frac{\partial}{\partial y_j}$ is a symplectic basis for the symplectic vector space $(T_q M, \omega_1)$. In these coordinates, we can define $\omega_0 \in \Omega^2(U)$ by

$$\omega_0 = \sum_{j=1}^{n} \mathrm{d}x_j \wedge \mathrm{d}y_j$$

Now the coordinate functions give a symplectomorphism of (U, ω_0) with $(\mathbb{R}^{2n}, \omega_0)$. The symplectic forms $\omega_1, \omega_0 \in \Omega^2(U)$ agree on the compact submanifold $\{q\} \subset U$, so by the relative Moser argument, shrinking U if necessary, there exists a symplectomorphism ψ of (U, ω_1) with (U, ω_0). $\square$

We can already see that the Moser argument is pretty powerful. So far we have only used it for the simplest compact submanifold, the singelton $\{q\}$. We would like to use it for other compact submanifolds Q. To do this, we need some extra information about what a neighbourhood of Q looks like. The tubular neighbourhood theorem relates such neighbourhoods to the normal bundle NQ, but this bundle does not always have a natural symplectic structure. To impose some structure on this bundle, we extend definition 3.6 to submanifolds.

Definition 4.15. A submanifold Q of a symplectic manifold (M, ω) is called symplectic (isotropic, coisotropic, Lagrangian) if at each $q \in Q$, $T_q Q$ is a symplectic (isotropic, coisotropic, Lagrangian) subspace of the symplectic vector space $(T_q M, \omega_q)$.

In our work, the by far most important of these concepts is that of a Lagrangian submanifold. As we shall see later many natural and important questions turn out to be about the intersection of two Lagrangian submanifolds. As Alan Weinstein puts it in [Wei81], "Everything is a Lagrangian submanifold". We will give a couple of important examples, but first we formulate an easier way to test if a submanifold is of a certain type.

Lemma 4.16. *Let (M, ω) be a symplectic manifold with dim$(M) = 2n$, $i : Q \to M$ an embedding.*

1. *The submanifold $i(Q)$ is isotropic if and only if $i^*\omega = 0$.*
2. *... Lagrangian if and only if it is isotropic and* $\dim(Q) = n$*.*
3. *... symplectic if and only if $i^*\omega$ is nondegenerate.*

Proof. The first and third claims follow immediately from the fact that $\mathrm{d}i : T_qQ \to T_{i(q)}M$ is an isomorphism onto its image, which is $T_{i(q)}i(Q)$. For 2, combine 1 with lemma 3.7. $\square$

Example 4.17. Let L be a smooth n-dimensional manifold, (T^*L, λ_{can}) the canonical symplectic structure on its cotangent bundle. Then the zero section L_0, as well as any fibre $\pi^{-1}(q)$, is a Lagrangian submanifold since

$$\begin{aligned}\lambda_{can}(0^*, v) &= 0^*(v) = 0\\ \lambda_{can}(v^*, 0) &= v^*(0) = 0.\end{aligned}$$

Here we have again used the splitting $TT^*L \simeq TL \oplus T^*L$. In general, for any submanifold $Q \subset L$, the annihilator

$$TQ^{\perp} = \{(q, v^*) \in T^*Q : v^*(T_qQ) = 0\} \subset T^*L$$

is a Lagrangian submanifold. Our examples above can be reformulated as

$$L_0 = TL^{\perp}, \qquad \pi^{-1}(q) = T\{q\}^{\perp}.$$

Example 4.18. Keeping the notation from the previous example, consider a one-form $\sigma \in \Omega^1(L)$ as a map $\sigma : L \to T^*L$. The graph $\sigma(L)$ is a Lagrangian submanifold of T^*L if and only if σ is closed. To see this, just recall that by proposition 4.10,

$$\begin{aligned}\sigma^*\omega_{can} &= \sigma^*(-\mathrm{d}\lambda_{can})\\ &= -\mathrm{d}\sigma^*\lambda_{can}\\ &= -\mathrm{d}\sigma.\end{aligned}$$

Example 4.19. If (M, ω) is a symplectic manifold, we define a symplectic structure on the product manifold $M \times M$ by $(-\omega) \times \omega := \pi_2^*\omega - \pi_1^*\omega$, where π_j denotes the Cartesian projections. The submanifolds $\{p\} \times M$ and $M \times \{p\}$ are symplectic submanifolds of $(M \times M, (-\omega) \times \omega)$ To see this, let $v_p : M \to M \times M$ and $h_p : M \to M \times M$ denote the vertical and horizontal embeddings of M at p respectively – that is, $v_p(M) = \{p\} \times M, h_p(M) = M \times \{p\}$. Then a quick calculation shows that

$$\begin{aligned}h_p^*\left((-\omega) \times \omega\right) &= -\omega\\ v_p^*\left((-\omega) \times \omega\right) &= \omega.\end{aligned}$$

Both of which are nondegenerate, so the submanifolds are symplectic.

Example 4.20. Given any diffeomorphism $\psi \in \mathrm{Symp}(M, \omega)$, we denote its graph by $gr_\psi(M) \subset M \times M$. This is a Lagrangian submanifold of $(M \times M, (-\omega) \times \omega)$ if and only if ψ is a symplectomorphism. To see this, just note that

$$\begin{aligned}gr_\psi^*((-\omega) \times \omega) &= gr_\psi^*(\pi_2^*\omega - \pi_1^*\omega)\\ &= \psi^*\omega - \omega,\end{aligned}$$

which is zero if and only if ψ is a symplectomorphism. In particular, the diagonal $\Delta(M) \subset M \times M$ is Lagrangian. With this in mind, we can actually restate the Arnold conjecture in terms of the Lagrangian intersection $gr_\psi(M) \cap \Delta(M)$. This motivated the generalization of the Arnold conjecture to arbitrary Lagrangian intersections $L_0 \cap L_1$ where L_0 and L_0 are related by a Hamiltonian isotopy.

Let us first consider neighbourhoods of symplectic submanifolds. If $Q \subset M$ is a symplectic submanifold, each tangent space splits as

$$T_qM \simeq T_qQ \oplus T_qQ^\omega.$$

This gives us a natural description of each normal space N_qQ as follows,

$$N_qQ = T_qM/T_qQ = \frac{T_qQ \oplus T_qQ^\omega}{T_qQ} \simeq T_qQ^\omega.$$

The next result uses this isomorphism to show that for compact symplectic submanifolds Q, the symplectomorphism class of the bundle TQ^ω determines the symplectomorphism class of any sufficiently small open neighbourhood $N(Q)$ of Q.

Proposition 4.21 (Symplectic Neighbourhood Theorem)**.** *Let (M_0, ω_0) and (M_1, ω_1) be symplectic manifolds with compact symplectic submanifolds $Q_j \subset M_j, j = 0, 1$. Let $N(Q_j)$ denote any sufficiently small open neighbourhood of Q_j. If there exists a symplectomorphism of the bundles $\Phi : TQ_0^{\omega_0} \to TQ_1^{\omega_1}$ covering a symplectomorphism of the base spaces $\phi : Q_0 \to Q_1$, then ϕ extends to a symplectomorphism ψ of $(N(Q_0), \omega_0)$ with $(N(Q_1), \omega_1)$.*

Proof. We may extend ϕ to neighbourhoods by taking $N(Q_0)$ small enough that $\exp^{-1}$ is defined, then taking ϕ' to be the composition

$$\begin{tikzcd} N(Q_0) \arrow[r, "\exp^{-1}"] \arrow[rrr, bend right, "\phi'"'] & TQ_0^{\omega_0} \arrow[r, "\Phi"] & TQ_1^{\omega_1} \arrow[r, "\exp"] & M_1 \end{tikzcd},$$

and defining $N(Q_1) = \phi'(N(Q_0))$. Now the two forms $\omega_0, \phi'^*\omega_1 \in \Omega^2(N(Q_0))$ are closed, nondegenerate, and agree on the compact submanifold Q_0 since $\phi'\big|_{Q_0} = \phi$ is a symplectomorphism. Hence, we may apply the relative Moser argument to get a symplectomorphism ψ' of $(N(Q_0), \omega_0)$ with $(N(Q_0), \phi'^*\omega_1)$ fixing Q. Taking $\psi = \phi' \circ \psi'$ we get the desired symplectomorphism. $\square$

We now consider the Lagrangian case, which is in fact even easier; the diffeomorphism class of L determines the symplectomorphism class of $N(L)$. This theorem will be essential in proving the "C^1-close to the identity"-case of the Arnold conjecture.

Proposition 4.22 (Lagrangian neighbourhood theorem)**.** *Let (M, ω) be a symplectic manifold with $\dim(M) = 2n$, $L \subset M$ a compact Lagrangian submanifold.*

Let L_0 denote the zero section in the cotangent bundle T^*L. Then there exist open neighbourhoods $N(L) \subset M, N(L_0) \subset T^*L$, and a symplectomorphism ϕ of $(N(L_0), \omega_{can})$ with $(N(L), \omega)$.

Proof. By remark 3.20, there exists a compatible complex structure J on TM, At each $q \in L, T_qL$ is a Lagrangian subspace of T_qM, so we know that

$$J_q(T_qL) = T_qL^{\perp} \simeq N_qL,$$

where $T_qL^{\perp}$ denotes the orthogonal complement with respect to the inner product g_J induced by J. Hence $J : TL \to NL$ is a bundle isomorphism. The inner product g_J induces an isomorphism $\Phi : T^*L \to TL$ satisfying the equation

$$(g_J)_q \left(\Phi_q(v^*), v\right) = v^*(v), \quad \forall v \in T_qL, v^* \in T_q^*L.$$

We define ϕ to be the following compositon,

$$\begin{tikzcd} T^*L \arrow[r, "\Phi"] \arrow[rrr, bend right, "\phi"'] & TL \arrow[r, "J"] & NL \arrow[r, "\exp"] & M \end{tikzcd} .$$

Restricting to a suitable neighbourhood $N(L)$ where exp is a diffeomorphism makes $\phi : N(L_0) \to N(L)$ a diffeomorphism. Note that for any $v = (v_0, v_1^*) \in T_{(q,0)}T^*L \simeq T_qL \oplus T_q^*L$,

$$(\mathrm{d}\phi)_{(q,0)}(v_0, v_1^*) = v_0 + J_q \circ \Phi_q(v_1^*). \tag{4.15}$$

This holds since J and Φ are linear, and the derivative of exp at the zero section is the identity. So if $v = (v_0, v_1^*), w = (w_0, w_1^*) \in T_{(q,0)}T^*L$,

$$\begin{aligned} (\phi^*\omega)_{(q,0)}(v, w) &= \omega_q \left(v_0 + J_q\Phi_q(v_1^*), w_0 + J_q\Phi_q(w_1^*)\right) \\ &= \omega_q \left(v_0, J_q\Phi_q(w_1^*)\right) - \omega_q \left(w_0, J_q\Phi_q(v_1^*)\right) \\ &= (g_J)_q \left(v_0, \Phi_q(w_1^*)\right) - (g_J)_q \left(w_0, \Phi_q(v_1^*)\right) \\ &= w_1^*(v_0) - v_1^*(w_0) \\ &= (\omega_{can})_q(v, w). \end{aligned}$$

Where we have used (4.11). This shows that the symplectic forms $\phi^*\omega$ and ω_{can} agree on the compact submanifold L_0. By the relative Moser argument, we can modify ϕ to make it a symplectomorphism on sufficiently small neighbourhoods as desired. $\square$

We are now ready to make precise what we mean by C^1-close to the identity, and see why it will be useful in proving the Arnold conjecture. The following lemma follows from the properties of the Whitney topology for compact manifolds. For an introduction to this theory, see appendix A.

Lemma 4.23. *Let M be a compact smooth manifold.*

1. *Let $g : M \to T^*M$ be an embedding sufficiently C^1-close to the embedding of the zero section $M_0 : m \mapsto (m, 0)$. Then $g(M) = \sigma(M)$ for some one-form $\sigma \in \Omega^1(M)$.*

2. *Let $g : M \to M \times M$ be an embedding sufficiently C^1-close to the diagonal embedding $\Delta : m \mapsto (m, m)$. Then $g(M) = gr_\phi(M)$ for some diffeomorphism $\phi : M \to M$. As before, $gr_\phi : M \to M \times M$ denotes the embedding of the graph of ϕ.*

Proof. By proposition A.5, the set $\mathrm{Diff}(M, M)$ of diffeomorphisms of M is open in $C^1(M, M)$. Let $\pi : T^*M \to M$ be the bundle projection, and $\pi_2 : M \times M \to M$ projection to the second factor. Then the diagrams

$$\begin{tikzcd} & T^*M \arrow[dr, "\pi"] & \\ M \arrow[ur, shift left, "g"] \arrow[ur, shift right, "M_0"'] \arrow[rr, shift left, "\pi\circ g"] \arrow[rr, shift right, "id"'] & & M \end{tikzcd} \qquad \begin{tikzcd} & M \times M \arrow[dr, "\pi_2"] & \\ M \arrow[ur, shift left, "g"] \arrow[ur, shift right, "\Delta"'] \arrow[rr, shift left, "\pi_2\circ g"] \arrow[rr, shift right, "id"'] & & M \end{tikzcd}$$

commute, and by proposition A.6, the maps

$$\pi_* : C^r(M, T^*M) \to C^r(M, M),$$
$$(\pi_2)_* : C^r(M, M \times M) \to C^r(M, M)$$

are continuous. The preimages of $\mathrm{Diff}(M, M)$ are thus open neighbourhoods of 0 and Δ in their respective spaces. If $g \in \pi_*^{-1}(\mathrm{Diff}(M, M))$, the map $g \circ (\pi \circ g)^{-1}$ defines a one-form with graph $g(M)$. If $g \in (\pi_2)_*^{-1}(\mathrm{Diff}(M, M))$, the map $\pi_2 \circ g$ is a diffeomorphism with graph $g(M)$. $\square$

The following proposition combines this lemma with the Lagrangian neighbourhood theorem and our results about the structure of the cotangent bundle.

Proposition 4.24. *Let (M, ω) be a compact symplectic manifold. Then there exists an open neighbourhood $N(id_M)$ of id_M in $\mathrm{Symp}(M, \omega)$ which can be identified with an open neighbourhood $N(0)$ of zero in the vector space $Z^1(M)$ of closed one-forms on M.*

Proof. We let

$$M_0 : M \to T^*M$$
$$\Delta : M \to M \times M$$

denote the embeddings of the zero section and diagonal respectively. We showed in example 4.20 that $\Delta(M)$ is a Lagrangian submanifold of $(M \times M, (-\omega) \times \omega)$. The diagonal $\Delta(M)$ is also compact by assumption, so by the Lagrangian neighbourhood theorem, there exist neighbourhoods $N(\Delta(M))$ and $N(M_0(M))$ of the diagonal and the zero section in $M \times M$ and T^*M respectively, and a symplectomorphism

$$\Psi : N(\Delta(M)) \to N(M_0(M)).$$

Since M is compact, we can find a C^0-neighbourhood $N(\Delta)$ of the diagonal embedding in $C^1(M, M \times M)$ such that every map in $N(\Delta)$ has its image contained in $N(\Delta(M))$. Since $C^1(M, M \times M) \simeq C^1(M, M) \times C^1(M, M)$ by proposition A.7, it is easy to see that the graph map

$$\begin{aligned} gr_- : C^1(M, M) &\to C^1(M, M \times M) \\ f &\mapsto (gr_f : x \mapsto (x, f(x)) \end{aligned}$$

is continuous. The preimage of $N(\Delta)$ under this map is a neighbourhood $N(id_M)$ of the identity in $C^1(M, M)$. Shrinking this to the preimage of the neighbourhood $N(M_0)$ of the zero embedding in $C^1(M, T^*M)$ obtained in lemma 4.23, we get the following diagram of neighbourhoods, continuous maps, and open inclusions

$$\begin{tikzcd} N(id_M) \arrow[r, "gr_-"] \arrow[d, hook] & N(\Delta) \arrow[r, "\Psi_*"] \arrow[d, hook] & N(M_0) \arrow[d, hook] \\ C^1(M, M) & C^1(M, M \times M) & C^1(M, T^*M) \end{tikzcd}.$$

Now for any map $\psi \in N(id_M)$, there exists a one-form $\sigma : M \to T^*M$ such that $\Psi \circ gr_\psi(M) = \sigma(M)$. Since $gr_\psi(M)$ is Lagrangian, and Ψ is a symplectomorphism, $\sigma(M)$ must also be Lagranagian. Thus σ is closed by example 4.18. It is not hard to see that this process can be reversed; just take the neighbourhood $N(\Delta) \subset C^1(M, M \times M)$ obtained in lemma 4.23, and perform appropriate restrictions to get the diagram

$$\begin{tikzcd} N(M_0) \arrow[r, "\Psi_*^{-1}"] \arrow[d, hook] & N(\Delta) \arrow[r, "gr_-^{-1}"] \arrow[d, hook] & N(id_M) \arrow[d, hook] \\ C^1(M, T^*M) & C^1(M, M \times M) & C^1(M, M) \end{tikzcd}.$$

Again using examples 4.18 and 4.20 we get that a closed one-form is identified with a symplectomorphism. $\square$

As mentioned, we are now close to proving the C^1-case of the Arnold conjecture. To see why, just note that if a given symplectomorphism $\psi \in N(id_M)$ corresponds not only to a closed, but to an exact one-form, we get the following equation.

$$\Psi \circ gr_\psi(M) = \mathrm{d}F(M) \tag{4.16}$$

Now since $\Psi \circ \Delta(M) = M_0(M)$, any critical point of $F : M \to \mathbb{R}$ corresponds to an intersection of $gr_\psi(M)$ with $\Delta(M)$, which is exactely a fixed point of ψ. Since Ψ is a diffeomorphism, transversality of the intersection $\mathrm{d}F(M) \cap M_0(M)$ is equivalent to transversality of the intersection $gr_\psi(M) \cap \Delta(M)$. Therefore, if ψ has only nondegenerate fixed points, F must be a Morse function. But for what sort of symplectomorphisms can we expect that the corresponding form is exact? This will be answered in the next section.

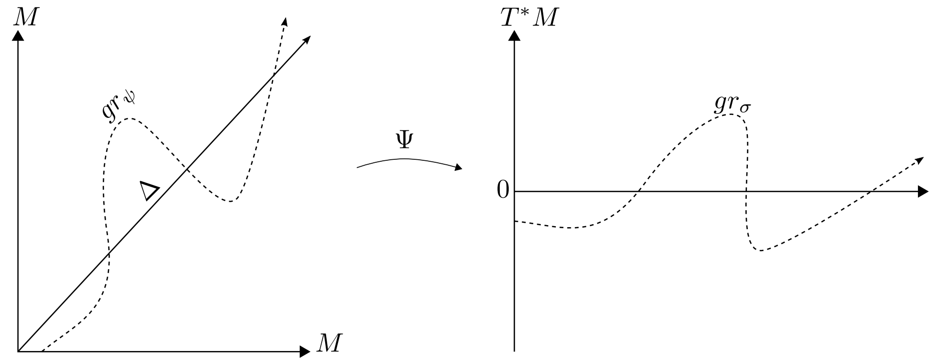

Figure 1: Mapping a neighbourhood of the diagonal to a neighbourhood of the zero section.

5 Generating functions

We have seen that there is a connection between closed forms and symplectomorphisms. In this section, we will expand this connection, and examine when we get exact forms.

5.1 Generating functions of type S

Cotangent bundles have the nice property that we can easily construct a Lagrangian submanifold by picking a closed one-form on the base space. For a general symplectic manifold M, the issue is that we only have the structure of a cotangent bundle locally around the diagonal $\Delta(M) \subset M \times M$. If we restrict ourselves to a cotangent bundle T^*L, we can easily get a Lagrangian submanifold of the symplectic manifold $(T^*L \times T^*L, \omega \times \omega)$. Letting $\pi : T^*L \to L$ denote the bundle projection, the map

$$\begin{aligned} T^*L \times T^*L &\to T^*(L \times L) \\ ((x_0, y_0), (x_1, y_1)) &\mapsto (x_0, x_1, \pi_0^* y_0 + \pi_1^* y_1) \end{aligned}$$

is a diffeomorphism, and when we consider both bundles with their symplectic structure,[1] it is also a symplectomorphism. Hence we can get a Lagrangian submanifold of the product by choosing a closed one-form $\sigma \in H^1(L \times L)$. In particular we can choose any smooth function $S : L \times L \to \mathbb{R}$ and look at the graph of $\mathrm{d}S$. The problem is that the graph of a symplectomorphism is only Lagrangian when we consider the twisted symplectic structure $(-\omega) \times \omega$. Therefore, we define a twist on the first factor by the map

$$\begin{aligned} \tau = -1 \times id : T^*L \times T^*L &\to T^*L \times T^*L \\ (x_0, y_0, x_1, y_1) &\mapsto (x_0, -y_0, x_1, y_1), \end{aligned}$$

[1] a quick computation shows that the canonical one-form on $T^*(L \times L)$ is in fact $\lambda \times \lambda := \pi_0^*\lambda + \pi_1^*\lambda$ when λ is the canonical one-form on T^*L.

where $y_j \in T^*_{x_j}L$. A quick computation shows that

$$\tau^*((-\omega) \times \omega) = \omega \times \omega.$$

Hence for any Lagrangian submanifold Y of $(T^*L \times T^*L, \omega \times \omega)$, we get a Lagrangian submanifold $Y^\tau := \tau(Y)$ of $(T^*L \times T^*L, -(\omega) \times \omega)$. If Y^τ is the graph of a diffeomorphism ψ, we know by example 4.20 that ψ must also be a symplectomorphism. Looking at the case where $Y_S = \mathrm{d}S(L \times L)$ is the graph of the one-form $\mathrm{d}S$, we can see that ψ must satisfy

$$Y_S^\tau = \{(x_0, -\mathrm{d}S_0(x_0, x_1), x_1, \mathrm{d}S_1(x_0, x_1))\} = \{x_0, y_0, u(x_0, y_0), v(x_0, y_0)\}$$

where we use the notation $\psi(x_0, y_0) = (u(x_0, y_0), v(x_0, y_0))$, and $\mathrm{d}f_j = \pi_j^*\mathrm{d}f$. This is equivalent to the generalized discrete Euler–Lagrange equations

$$y_0 = -\frac{\partial S}{\partial x_0}(x_0, x_1), \quad y_1 = \frac{\partial S}{\partial x_1}(x_0, x_1) \iff \psi(x_0, y_0) = (x_1, y_1). \tag{5.1}$$

Given a function S, we can solve the above equations if the map

$$G_S(x_0, x_1) = \left(x_0, \frac{\partial S}{\partial x_0}(x_0, x_1)\right) \tag{5.2}$$

is a diffeomorphism. By the implicit function theorem, we can solve this locally at x_0, x_1 if S satisfies the Legendre-condition

$$\det\left(\frac{\partial^2 S}{\partial x_1 \partial x_1}\right) \neq 0. \tag{5.3}$$

Conversely, given a symplectomorphism $\psi : T^*L \to T^*L$, we know that if the graph of ψ is the graph of a one-form σ, this one-form must also be closed. This happens exactly when the map

$$G_u(x_0, y_0) = (x_0, u(x_0, y_0)) \tag{5.4}$$

is a diffeomorphism. Intuitively, we can understand this condition by a physical analogy. If $\psi = \phi_H^{1,0}$ is the time-one map of a Hamiltonian flow, we might ask the question: Given any two locations $x_0, x_1 \in L$, is it possible to travel from x_0 to x_1 in one unit of time? If so, what are the possible starting momenta? We approach this question by considering all possible initial momenta, $y_0 \in T^*_{x_0}L$, and all possible terminal momenta, $y_1 \in T^*_{x_1}L$. We are then interested in the intersection $\psi(T^*_{x_0}L) \cap T^*_{x_1}L$. (See figure 2.) This intersection is transverse (and the equation is locally solvable by the implicit function theorem) if ψ satisfies the inverse-Legendre condition

$$\det\left(\frac{\partial u}{\partial y_0}(x_0, y_0)\right) \neq 0.$$

If L is simply connected, we know that the form σ is not just closed, but exact. Hence we have proved the following lemma.

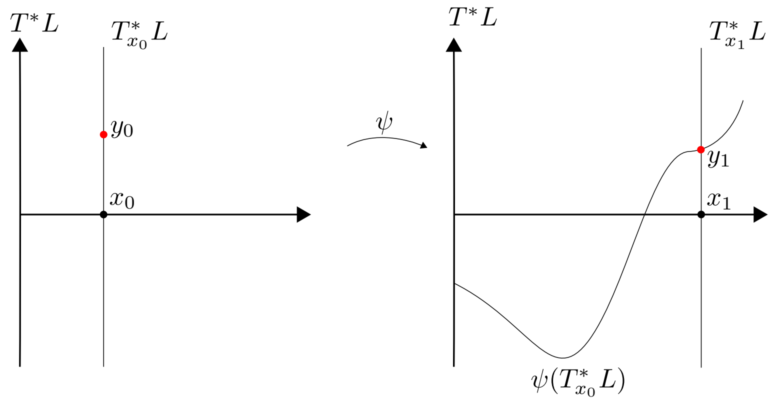

Figure 2: A question of intersection between two Lagrangian submanifolds.

Lemma 5.1 (S-Lemma)**.** *Equip all cotangent-bundles with the canoncial symplectic structure.*

1. *Let L be a simply connected manifold. Given any symplectomorphism $\psi = (u, v) : T^*L \to T^*L$ such that the map G_u from* (5.4) *is a diffeomorphism, there exists a function $S : L \times L \to \mathbb{R}$ such that ψ and S satisfy* (5.1)*.*

2. *Let L be a manifold, $S : L \times L \to \mathbb{R}$ a smooth function such that the map G_S from* (5.2) *is a diffeomorphism. Then there exists a symplectomorphism $\psi : T^*L \to T^*L$ such that ψ and S satisfy* (5.1)*.*

Any function S with these properties is called a generating function of type S for ψ.

One interesting special case of this is $T^*\mathbb{R}^n \simeq \mathbb{R}^{2n}$. In particular, we know that $\mathbb{R}^n$ is simply connected, so we expect all symplectomorphisms satisfying the non-degeneracy condition to have a type S generating function. For Hamiltonian symplectomorphisms, we can find an explicit form of this function via the action integral

$$\Phi_H(z) = \int_{t_0}^{t_1} \langle \dot{x}, y \rangle - H_t(z) \mathrm{d}t.$$

For suitable points (x_0, x_1), we define $(x, y) = z : [t_0, t_1] \to \mathbb{R}^{2n}$ to be the unique curve solving the Hamiltonian equations (2.7) with boundary conditions $x(t_j) = x_j$. Equivalently, if $\phi_H^{t_1,t_0}$ satisfies the non-degeneracy condition, we may take $z(t) = \phi_H^{t,t_0}(G_u^{-1}(x_0, x_1))$. We then define S_H as the composition

$$S_H(x_0, x_1) = \Phi_H(z). \tag{5.5}$$

Lemma 5.2. *With the above definitions, S_H is a generating function of type S for the symplectomorphism $\psi = \phi_H^{t_0,t_1}$.*

Proof. Choose $(x_0, x_1) \in U$, and a perturbation (ξ_0, ξ_1). Then for sufficiently small $s > 0$, let $z_s : [t_0, t_1] \to \mathbb{R}^{2n}$ be the unique path solving (2.7) with boundary conditions

$$x_s(t_0) = x_0 + s\xi_0, \qquad x_s(t_1) = x_1 + s\xi_1. \tag{5.6}$$

Now by definition, $\Phi(z_s) = S_H(x_0 + s\xi_0, x_1 + s\xi_1)$. Differentiating both sides with respect to s, and using the calculation in (2.10), we get that

$$\langle y_1, \xi_1 \rangle - \langle y_0, \xi_0 \rangle = \frac{\partial S_H}{\partial x_0} \cdot \xi_0 + \frac{\partial S_H}{\partial x_1} \cdot \xi_1.$$

Since this holds for all $(\xi_0, \xi_1) \in \mathbb{R}^{2n}$, the result follows. $\square$

Example 5.3. To get a feel for what these equations look like, we consider the example of $S : \mathbb{R}^n \times \mathbb{R}^n \to \mathbb{R}$ given by

$$S(x_0, x_1) = \frac{||x_1 - x_0||^2}{2}.$$

The solution to (5.1) is clearly

$$y_0 = y_1 = x_1 - x_0.$$

Hence the corresponding symplectomorphism must be

$$\psi(x, y) = (x + y, y),$$

which is precisely free translational motion in $\mathbb{R}^n$. In fact this system is generated by the Hamiltonian $H(x, y, t) = \frac{||y||^2}{2}$ on the time-interval $[0, 1]$, which can be interpreted as a system with no potential energy. This can be generalized to cotangent bundles of geodesically convex Riemannian manifolds by replacing the euclidean norm with the induced metric. The generated symplectomorphism then corresponds to geodesic flow. See [CdS06].

The Euler–Lagrange equations (2.2) arose from a variational principle on the space of paths. Since equations (5.1) are a discrete version of these equations, we expect there to be a discrete variational principle generating the latter as well. Let $\psi : \mathbb{R}^{2n} \to \mathbb{R}^{2n}$ be a symplectomorphism with generating function S. We denote an orbit of length l of (x_0, y_0) by

$$(x_j, y_j) = \psi^j(x_0, y_0), \qquad 0 \le j \le l.$$

From the S-lemma 5.1, we know that the orbit is uniquely determined by the sequence

$$X = (x_0, x_1, ..., x_l),$$

since the y_j's can be recovered as

$$y_j = \partial_2(x_{j-1}, x_j) = -\partial_1(x_j, x_{j+1}),$$

where ∂_1 and ∂_2 denote the derivatives of S with respect to the first and second variable respectively. The two expressions on the right agree for all $0 < j < l$ by lemma 5.1, so y_j is well defined. For $j = 0$ and $j = l$, only one of the expressions is well defined, so we set y_j equal to that one. It is not hard to see that if X satisfies the discrete Euler Lagrange equation

$$\partial_2(x_{j-1}, x_j) + \partial_1(x_j, x_{j+1}) = 0$$

then, X is also a critical point of the discrete action functional

$$I(X) = \sum_{j=0}^{l-1} S(x_j, x_j + 1).$$

Unfortunately this approach is not as fruitful as we hoped. First of all we can only expect generating functions for simply-connected base spaces. Second, we don't have a general way to avoid the non-degeneracy condition of lemma 5.1. This prompts us to modify our assumptions, and try a new approach.

5.2 Hamiltonian symplectomorphisms of exact manifolds

Cotangent bundles have a very specific symplectic structure. In particular, they have a canonical one-form λ with the two properties

$$\sigma^*\lambda = \sigma, \qquad \omega = -\mathrm{d}\lambda.$$

In this subsection, we get rid of the first of these, and show that if we add the assumption that our symplectomorphism is Hamiltonian, we still get generating functions.

Definition 5.4. A symplectic manifold (M, ω) is called exact if the closed two-form ω is exact – that is, if there exist some one-form λ with $\omega = -\mathrm{d}\lambda$. The form λ is called an action form of (M, ω).

We now show that on an exact symplectic manifold, Hamiltonian and symplectic isotopies are characterized by how they act on an action form λ. Recall from section 4.2 that an isotopy $\phi_t \in \mathrm{Diff}(M)$ generated by a vector field X_t via (4.3) is called symplectic if the form $\iota_\omega(X_t)$ is closed, and Hamiltonian if it is exact.

Proposition 5.5. *Let ϕ_t be an isotopy of an exact symplectic manifold $(M, -\mathrm{d}\lambda)$ generated by the vector field X_t. Then ϕ_t is a symplectic isotopy if and only if the form*

$$\phi_t^*\lambda - \lambda \tag{5.7}$$

is closed for all t. Moreover, ϕ_t is a Hamiltonian isotopy if and only if the form is also exact for all t – that is, if there exist some smooth one parameter family of functions $F_t : M \to \mathbb{R}$ such that

$$\phi_t^*\lambda - \lambda = \mathrm{d}F_t. \tag{5.8}$$

One choice of F_t is

$$F_t = \int_0^t (\iota_{X_s}(\lambda) - H_s) \circ \phi_s \, \mathrm{d}s, \tag{5.9}$$

where H_t is the Hamiltonian function generating ϕ_t. This is unique up to adding a constant.

Proof. For the first part, just note that

$$0 = \mathrm{d}(\phi_t^*\lambda - \lambda) = \omega - \phi_t^*\omega \iff \phi_t^*\omega = \omega.$$

Which is equivalent to ϕ_t being a symplectomorphism for all t. If equation (5.8) holds, let

$$H_t = \iota_{X_t}(\lambda) - \left(\frac{\mathrm{d}}{\mathrm{d}t}F_t\right) \circ \phi_t^{-1}.$$

Then the integral in equation (5.9) becomes

$$\int_0^s \left(\iota_{X_s}(\lambda) - \iota_{X_s}(\lambda) + \left(\frac{\mathrm{d}}{\mathrm{d}t}F_s\right) \circ \phi_s^{-1}\right) \circ \phi_s \, \mathrm{d}s = \int_0^t \frac{\mathrm{d}}{\mathrm{d}t}F_s \, \mathrm{d}s = F_t - F_0.$$

Since $dF_0 = id^*\lambda - \lambda = 0, F_0$ must be constant, so our claim holds. To see that H_t generates ϕ_t, we combine equations (4.4) and (4.5), to get that for all t,

$$\begin{aligned}
\frac{\mathrm{d}}{\mathrm{d}t}\phi_t^*\lambda &= \phi_t^*\left(\iota_{X_t}(\mathrm{d}\lambda) + \mathrm{d}\iota_{X_t}(\lambda)\right) \\
&= \phi_t^*\left(\iota_{X_t}(\mathrm{d}\lambda) + \mathrm{d}\left(H_t + \left(\frac{\mathrm{d}}{\mathrm{d}t}F_t\right) \circ \phi_t^{-1}\right)\right) \\
&= -\phi_t^*\iota_{X_t}(\omega) + \phi_t^*\mathrm{d}H_t + \frac{\mathrm{d}}{\mathrm{d}t}\mathrm{d}F_t \\
&= -\phi_t^*\iota_\omega(X_t) + \phi_t^*\mathrm{d}H_t + \frac{\mathrm{d}}{\mathrm{d}t}(\phi_t^*\lambda - \lambda) \\
&= -\phi_t^*\iota_\omega(X_t) + \phi_t^*\mathrm{d}H_t + \frac{\mathrm{d}}{\mathrm{d}t}\phi_t^*\lambda.
\end{aligned}$$

This implies that $\iota_\omega(X_t) = \mathrm{d}H_t$, which means that H_t generates ϕ_t. Conversely, if $\iota_\omega(X_t) = \mathrm{d}H_t$, we define F_t by equation (5.9). The above calculation becomes

$$\begin{aligned}
\frac{\mathrm{d}}{\mathrm{d}t}\phi_t^*\lambda &= \phi^*(\mathrm{d}\iota_{X_t}(\lambda) - \mathrm{d}H_t) \\
&= \mathrm{d}\left((\iota_{X_t}(\lambda) - H_t) \circ \phi_t\right).
\end{aligned}$$

Integrating this on both sides, we use the fundamental lemma of calculus on the left, and again assume interchange of differentiation and integration on the right. This gives

$$\phi_t^*\lambda - \lambda = \mathrm{d}\int_0^t \left((\iota_{X_s}(\lambda) - H_s) \circ \phi_s\right) \, \mathrm{d}s = \mathrm{d}F_t.$$

□

Remark 5.6. *For $x \in M, F_t(x)$ is the integral of the action form $\lambda - H_s \mathrm{d}s$ along the curve*

$$\begin{aligned} [0,t] &\to M \\ s &\mapsto \phi_s(x). \end{aligned}$$

Notice the similarity to the generating function S_H.

We say that any symplectomorphism ϕ is exact with respect to λ if it satisfies the analogue of (5.8),

$$\phi^*\lambda - \lambda = \mathrm{d}F. \tag{5.10}$$

One issue now is that there might be several choices of one-form λ satisfying $-\mathrm{d}\lambda = \omega$, and that the definition of exact depends on this choice unless $\phi^* : H^1(M;\mathbb{R}) \to H^1(M;\mathbb{R})$ is the identity. One way to deal with this to consider two different one-forms; one on the source an one on the target. We then consider the equation

$$\phi^*\lambda_1 - \lambda_0 = \mathrm{d}F.$$

Even more generally, we might consider one-forms α on the product manifold such that $-\mathrm{d}\alpha = (-\omega) \times \omega$. The notion of exactness in this setting is that of an exact Lagrangian embedding.

Definition 5.7. Let $(M, -\mathrm{d}\lambda)$ be an exact symplectic manifold. A Lagrangian embedding $\iota : L \to M$ is called exact with respect to λ if the one-form $\iota^*\lambda$ is exact. (The form $\iota^*\lambda$ will always be closed since the condition for being a Lagrangian embedding is $\iota^*\omega = 0$.)

This definition might also be dependant on choice of λ, unless $\iota^* : H^1(M;\mathbb{R}) \to H^1(L;\mathbb{R})$ is the zero morphism. We can immediately see that the notion of exact Lagrangian embeddings generalizes equations (5.10) and (5.2), since they are equivalent to saying that $gr_\phi : M \to M \times M$ is an exact Lagrangian embedding with respect to the forms $(-\lambda) \times \lambda$ and $(-\lambda_1) \times \lambda_0$ respectively. We now use the fact that Hamiltonian symplectomorphisms are exact with respect to any form to prove existence of a new kind of generating function.

Proposition 5.8. *For a symplectic manifold (M, ω), let $\alpha \in \Omega^1(M \times M)$ be any one-form such that*

$$-\mathrm{d}\alpha = (-\omega) \times \omega, \qquad \alpha\big|_{\Delta(M)} = 0,$$

where $\Delta : M \to M \times M$ denotes the diagonal embedding. Then for any Hamiltonian Symplectomorphism ϕ, the embedding of the graph, $gr_\phi : M \to M \times M$ is an exact Lagrangian embedding with respect to α. Any function $S_{\alpha,\phi} : M \to \mathbb{R}$ satisfying

$$\mathrm{d}S_{\alpha,\phi} = gr_\phi^*\alpha$$

is called an α-generating function for ϕ.

Proof. We can express gr_ϕ as the composition

$$gr_\phi = (id_M \times \phi) \circ \Delta.$$

Clearly, $\phi \times id$ is a Hamiltonian symplectomorphism since ϕ is one. Hence, by proposition 5.5, there exists a function $F : M \times M \to \mathbb{R}$ such that

$$(id \times \phi)^* \alpha - \alpha = \mathrm{d}F.$$

Since $\Delta^* \alpha = 0$ by assumption, we get that

$$\mathrm{d}(F \circ \Delta) = \Delta^* \mathrm{d}F = \Delta^*(id_M \times \phi)^* \alpha - \Delta^* \alpha = gr_\phi^* \alpha.$$

Hence $F \circ \Delta : M \to \mathbb{R}$ is an α-generating function for ϕ. $\square$

5.3 Proof of the C^1-close to the identity case

Now for general symplectic manifolds, there might not exist any global choice of α. In particular, any closed symplectic manifold cannot be exact by proposition 4.4. However, the last calculation of the previous section offers another approach; we can use the symplectomorphism $\Psi : N(\Delta(M)) \to N(M_0(M))$ from the Lagrangian neighbourhood theorem 4.22 to define a local action form $\alpha = \Psi^* \lambda_{can}$. While this form does not technically satisfy all the conditions of proposition 5.8, it is not hard to see how the proof can be adapted so that this form admits an α generating function $S_{\alpha,\phi}$ on $N(\Delta(M))$. Combining this with the calculations of proposition 4.24, we are in a position where we can easily prove a special case of the Arnold conjecture. First we remind ourselves of the definitions of nondegeneracy.

Definition 5.9. Let X be a manifold. A fixed point x of some C^1-map $f : X \to X$ is called nondegenerate if $||\mathrm{d}f - I||$ is nonsingular, or equivalently if the intersection of the graph $gr_f(X)$ with the diagonal $\Delta(X)$ is transverse at (x, x) in $X \times X$.

Definition 5.10. Let X be a manifold. A critical point x of a C^1-function $F : X \to \mathbb{R}$ is called nondegenerate if the Hessian $H_F(x)$ is nonsingular. Equivalently if the intersection of the graph $\mathrm{d}F(X)$ with the graph of the zero-section $X_0(X)$ is transverse at $(x, 0)$ in T^*X. A function with only nondegenerate critical points is called a Morse function.

Proposition 5.11. *For any compact symplectic manifold (M, ω), there exists a C^1-neighbourhood $N(id_M)$ of the identity such that if ϕ_t is a Hamiltonian isotopy with $\phi_t \in N(id_M)$ for $0 \leq t \leq 1$, then ϕ_1 has at least as many fixed points as a function on M has critical points. If all the critical points are nondegenerate, there are at least as many as a Morse function on M has critical points.*

Proof. By proposition 4.24, there exists a C^1–neighbourhood $N(id_M)$ of the identity such that if $\phi_t \in N(id_M)$, the image gr_{ϕ_t} is contained in the domain of $\Psi : N(\Delta(M)) \to N(M_0(M))$, and the image of $\Psi \circ gr_{\phi_t}$ is the image of a one-form σ_t. Hence there exists a diffeomorphism $f_t : M \to M$ such that

$$\begin{tikzcd} & T^*M \\ M \arrow[ur, "\Psi \circ gr_{\phi_t}"] \arrow[r, "f_t"'] & M \arrow[u, "\sigma_t"'] \end{tikzcd}$$

commutes. We also know that the form $\alpha = \Psi^*\lambda_{can}$ admits an α-generating function F_t satisfying

$$gr^*_{\phi_t}\Psi^*\lambda_{can} = \mathrm{d}F_t.$$

Combining these facts with proposition 4.10, we calculate

$$\begin{aligned} f_t^*\sigma_t^*\lambda_{can} &= gr^*_{\phi_t}\Psi^*\lambda_{can} \\ f_t^*\sigma_t &= \mathrm{d}F_t \\ \implies \qquad \sigma_t(M) &= \mathrm{d}F_t(M) \\ \implies \Psi \circ gr_{\phi_t}(M) &= \mathrm{d}F_t(M) \end{aligned} \tag{5.11}$$

Since Ψ maps the diagonal to the zero-section, we know that a critical point of F_1 corresponds uniquely to a fixed point of ϕ_1. As we have remarked before, Ψ preserves transversality of intersections, so if all the fixed points are nondegenerate, all the critical points are also nondegenerate. $\square$

We have managed to prove one special case of the Arnold conjecture. This result however is not too interesting. In a sense, it is a local result, and can be seen as an consequence of the local theory being trivial. On the other hand, it will give us a new kind of generating function on $\mathbb{R}^{2n}$, which will be useful in our further study.

5.4 Generating functions of type V

As we have already seen, α-generating functions give rise to the generating function S_H as the special case $\alpha = (-\lambda) \times \lambda$. We can get a different kind of generating function on $\mathbb{R}^{2n}$ by using a one-form α that is not of the form $-\lambda_0 \times \lambda_1$, namely the form $\Psi^*\lambda_{can}$ from the proof of proposition 5.11. In the special case of $\mathbb{R}^{2n}$, we can extend the symplectomorphism Ψ from the Lagrangian neighbourhood theorem 4.22 to a global symplectomorphism

$$\Psi : T^*\mathbb{R}^n \times T^*\mathbb{R}^n \to T^*\mathbb{R}^{2n}$$

as follows. On $T^*\mathbb{R}^{2n}$, we use coordinates (q_0, q_1, p_0, p_1), where the q's describe the base-space, and the p's describe the fibers. In these coordinates the canonical symplectic form is

$$\omega_{can} = \sum_{j=1}^{n} \mathrm{d}q_{0j} \wedge \mathrm{d}p_{0j} + \mathrm{d}q_{1j} \wedge p_{1j} = \mathrm{d}q_0 \wedge \mathrm{d}p_0 + \mathrm{d}q_1 \wedge \mathrm{d}p_1.$$

Using the coordinates (x_0, y_0, x_1, y_1) on $T^*\mathbb{R}^n \times T^*\mathbb{R}^n$, the twisted standard form is

$$(-\omega_0) \times (\omega_0) = \mathrm{d}x_1 \wedge \mathrm{d}y_1 - \mathrm{d}x_0 \wedge \mathrm{d}y_0.$$

If we then define

$$\Psi(x_0, y_0, x_1, y_1) = (x_1, y_0, y_1 - y_0, x_0 - x_1),$$

It is not hard to check that we get a symplectomorphism, and that the diagonal is mapped to the zero-section. Thus we define an action form

$$\alpha = \Psi^* \lambda_{can} = (y_1 - y_0)\mathrm{d}x_0 + (x_0 - x_1)\mathrm{d}y_0$$

on $T^*\mathbb{R}^n \times T^*\mathbb{R}^n$. The graph of any symplectomorphism $\psi = (u, v) : T^*\mathbb{R}^n \to T^*\mathbb{R}^n$ is an exact Lagrangian embedding,[2] so there exists an α-generating function $S_{\alpha,\psi}$ such that

$$gr_\psi^* \alpha = \mathrm{d}S_{\alpha,\psi}.$$

We know from proposition 4.24 that there should be some C^1-condition such that if ψ satisfies this condition, $\Psi \circ gr_\psi$ is the graph of a one-form. Specifically, this condition is that the composition

$$\begin{aligned} \pi \circ \Psi \circ gr_\psi &: M \to M \\ (x_0, y_0) &\mapsto (u(x_0, y_0), y_0) \end{aligned}$$

is a diffeomorphism. If it is, let f denote its inverse. This satisfies

$$f(x_1, y_0) = (x_0, y_0) \iff u(x_0, y_0) = x_1.$$

We then define $V = f^* S_{\alpha,\psi}$, which has differential

$$\mathrm{d}V = f^* \mathrm{d}S_{\alpha,\psi} = f^* gr_\psi^* \alpha.$$

Expanding this equation, we get that

$$y_1 - y_0 = -\frac{\partial V}{\partial x_1}(x_1, y_0), \quad x_1 - x_0 = \frac{\partial V}{\partial y_0}(x_1, y_0)$$

if and only if $\psi(x_0, y_0) = (x_1, y_1)$. We call any function satisfying this equation a generating function of type V for ψ. Just as for generating functions of type S, there is a corresponding non-degeneracy condition such that any function V satisfying this condition defines a symplectomorphism via the above equations. It is not hard to see that the condition is that the map

$$(x_1, y_0) \mapsto (x_1 - \frac{\partial V}{\partial y_0}(x_1, y_0), y_0)$$

should be a diffeomorphism. We have thus proved the following lemma.

[2]Since the space is simply connected, all Lagrangian embeddings are exact, therefore it is not necessary to assume ψ is Hamiltonian.

Lemma 5.12. *Any symplectomorphism $\psi = (u, v)$ of $\mathbb{R}^{2n}$ such that the map*

$$\begin{aligned} G_u : \mathbb{R}^{2n} &\to \mathbb{R}^{2n} \\ (x_0, y_0) &\mapsto (u(x_0, y_0), y_0) \end{aligned}$$

is a diffeomorphism admits a generating function of type V. Any function V on $\mathbb{R}^{2n}$ such that the map

$$\begin{aligned} G_V : \mathbb{R}^{2n} &\to \mathbb{R}^{2n} \\ (x_1, y_0) &\mapsto \left(x_1 - \frac{\partial V}{\partial y_0}(x_1, y_0), y_0 \right) \end{aligned}$$

is a diffeomorphism generates a unique symplectomorphism ψ of $\mathbb{R}^{2n}$ via (5.4).

We claimed earlier that the diffeomorphism condition on G_u is a C^1-condition. Therefore, there should be some way of rephrasing it in terms of conditions on the first derivative of ψ. This is the idea of the following lemma.

Lemma 5.13. *The map G_u defined above is a diffeomorphism if ψ satisfies*

$$||\mathrm{d}\psi(z) - I_{2n}|| \leq \frac{1}{2} \qquad \forall z \in \mathbb{R}^{2n}. \tag{5.12}$$

Proof. The proof has two parts. First we show that if ψ satisfies (5.12), then so does G_u. Then we show that any map with this property must be a diffeomorphism. It follows from the definition of G_u that

$$\mathrm{d}\psi = \begin{pmatrix} A & B \\ C & D \end{pmatrix} \implies \mathrm{d}G_u = \begin{pmatrix} A & B \\ 0 & I_n \end{pmatrix}.$$

Now a quick calculation shows that

$$\begin{aligned} ||(\mathrm{d}G_u - I_{2n})(u, v)|| &= \left\| \begin{pmatrix} A - I_n & B \\ 0 & 0 \end{pmatrix} \begin{pmatrix} u \\ v \end{pmatrix} \right\| \\ &= \left\| \begin{pmatrix} (A - I_n)u + Bv \\ 0 \end{pmatrix} \right\| \\ &\leq \left\| \begin{pmatrix} (A - I_n)u + Bv \\ Cu + (D - I_n)v \end{pmatrix} \right\| \\ &= ||\mathrm{d}\psi - I_{2n}(u, v)||. \end{aligned}$$

It is clear that $\mathrm{d}G_u$ is invertible, since if $w \neq 0, \mathrm{d}G_u(w) = 0$, we have

$$||(\mathrm{d}G - I_{2n})(w)|| = ||w|| > \frac{1}{2}||w||.$$

We now use a fixed point argument to show that G is bijective. For any $\zeta \in \mathbb{R}^{2n}$, define a map $T_\zeta : \mathbb{R}^{2n} \to \mathbb{R}^{2n}$ by

$$T_\zeta(z) = z + \zeta - G_u(z)$$

and consider the fixed point problem $T_\zeta(z) = z$. We have that

$$||\mathrm{d}T_\zeta|| = ||I_{2n} - \mathrm{d}G|| \leq \frac{1}{2},$$

so by a mean value theorem argument, T_ζ is a contraction. Hence, by the Banach fixed point theorem, there exists a unique fixed point. This is the same as a unique solution to the equation $G_u(z) = \zeta$. $\square$

Symplectomorphisms satisfying this condition will admit generating functions of type V. It might seem like the C^1-condition is at least as strict as the Legendre condition of the S-Lemma, but we have a trick up our sleeve. If $\psi = \phi_H^{t_1,t_0}$ is a Hamiltonian Symplectomorphism, we can discretize the flow by sampling it as follows. For some sufficiently large N, we define

$$\psi_j = \phi_H^{\tau_{j+1},\tau_j}, \qquad \tau_j = t_0 + \frac{j}{N}(t_1 - t_0), \qquad j = 0, ..N-1. \tag{5.13}$$

Then we can retrieve ψ as the composition

$$\psi = \psi_{N-1} \circ ... \circ \psi_1 \circ \psi_0.$$

If the Hamiltonian has bounded first and second order derivatives, that is if

$$\sup_{t\in\mathbb{R}} \sup_{z\in\mathbb{R}^{2n}} \left(||\mathrm{d}^2 H_t(z)|| + ||\mathrm{d}H_t(z)||\right) = L < \infty, \tag{5.14}$$

we can Taylor expand the map $t \mapsto \phi_H^{t,s}(z)$ at $t = s$ and get

$$\begin{aligned} ||\phi_H^{t,t_0}(z) - z|| &= ||\phi_H^{t_0,t_0}(z) + (t - t_0)\frac{\mathrm{d}}{\mathrm{d}t}\phi_H^{t_0,t_0}(z) + o(|t-s|^2) - z|| \\ &\leq ||(t-s)\cdot X_{H_t}(z) + o(|t-s|^2)|| \leq L(t-s). \end{aligned}$$

Hence, by picking N large enough, each ψ_j satisfies the C^1-condition (5.12), so there exist functions $V_j : \mathbb{R}^{2n} \to \mathbb{R}$ such that

$$x_{j+1} - x_j = \frac{\partial V_j}{\partial y}(x_{j+1}, y_j), \qquad y_{j+1} - y_j = -\frac{\partial V_j}{\partial x}(x_{j+1}, y_j) \tag{5.15}$$

if and only if $(x_{j+1}, y_{j+1}) = \psi_j(x_j, y_j)$. This is a discrete version of Hamilton's equations (2.7). Just as before, we are interested in finding a corresponding variational principle. Let $\mathcal{P} \simeq \mathbb{R}^{2nN+n}$ denote the set of discrete paths $z = (x_0, .., x_N, y_0, ...y_{N-1})$ of length N in $\mathbb{R}^{2n}$. (Note that y_N is uniquely determined by the rest of z by (5.15)). We now define a symplectic action functional on this space by

$$\Phi(z) = \sum_{j=0}^{N-1} \langle y_j, x_{j+1} - x_j \rangle - V_j(x_{j+1}, y_j). \tag{5.16}$$

As expected, this is just a discrete analogue of the action integral (2.9). As in lemma 2.10, we fix the boundary conditions x_0 and x_N and consider variations fixing these endpoints.

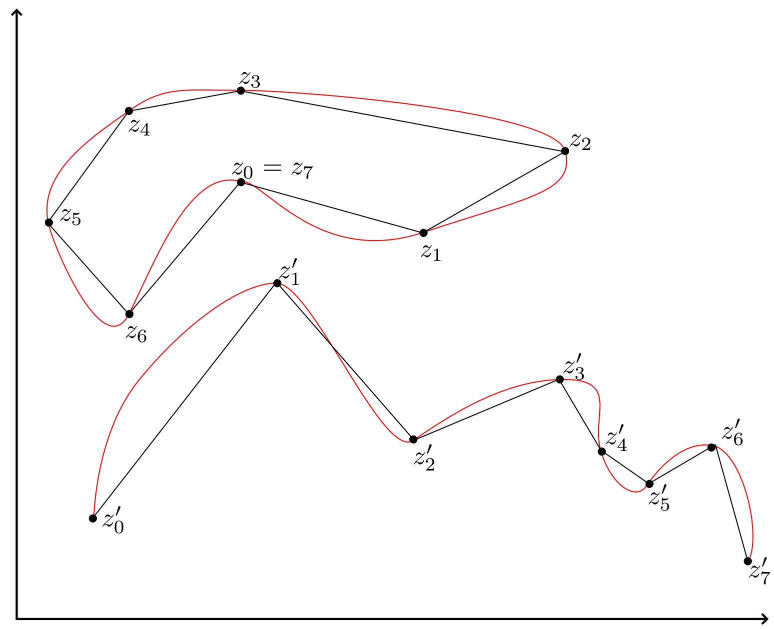

Figure 3: The discrete paths z and z' are critical points of the action functional approximating the red Hamiltonian flow. Note that $\psi(z_0) = z_0$ is a fixed point.

Lemma 5.14. *A point $z \in \mathcal{P}$ is a critical point of Φ with respect to variations fixing the endpoints if and only if z satisfies* (5.15) *(the first equation for $0 \leq j \leq N-2$, the second for $0 \leq j \leq N-1$).*

Proof. We just need to calculate that the partial derivatives are

$$\begin{aligned} \frac{\partial \Phi}{\partial x_j} &= y_{j-1} - y_j - \frac{\partial V_{j-1}}{\partial x}(x_j, y_{j-1}) && 1 \leq j \leq N-1 \\ \frac{\partial \Phi}{\partial y_j} &= x_{j+1} - x_j - \frac{\partial V_j}{\partial y}(x_{j+1}, y_j) && 0 \leq j \leq N-1 \end{aligned} \tag{5.17}$$

Setting all of these to zero is equivalent to (5.15). □

In light of the Arnold conjecture, we are specifically interested in the special case of periodic boundary conditions $x_0 = x_N$. The space of such paths can be identified with the space $\mathcal{P}_{per}$ of N-periodic sequences with terms in $\mathbb{R}^{2n}$. For notational convenience we extend ψ_j and V_j to all $j \in \mathbb{Z}$ by

$$V_{j+N} = V_j, \qquad \psi_{j+N} = \psi_j.$$

It is not to hard to see that fixed points of ψ correspond to critical points of

$\Phi : \mathcal{P}_{per} \to \mathbb{R}$, but if we want to use Morse theory, we need to show that nondegeneracy carries over as well. This is the essence of the following lemma.

Lemma 5.15. *A periodic sequence z is a critical point of Φ if and only if z_0 is a fixed point of ψ. Moreover, z_0 is nondegenerate as a fixed point if and only if z is non degenerate as a critical point.*

Proof. The calculations of (5.17) continue to hold for the periodic boundary conditions, so the first part of the lemma holds. For the second part, we need to compute the Hessian at a critical point z, which we denote $H_z\Phi$. Differentiating equation (5.17) yields

$$\frac{\partial^2\Phi}{\partial x_k \partial x_j} = -\delta_{jk} \cdot \frac{\partial^2 V_j}{\partial y^2} \qquad \frac{\partial^2\Phi}{\partial y_k \partial y_j} = -\delta_{jk} \cdot \frac{\partial^2 V_j}{\partial x^2}$$

$$\frac{\partial^2\Phi}{\partial y_j \partial x_k} = \delta_{(j+1)k}\left(1 - \frac{\partial^2 V_j}{\partial y \partial x}\right) - \delta_{jk}.$$

Where δ denotes the Kronecker-delta. We now consider the equation

$$H_z\Phi(\xi, \eta) = 0.$$

This gives the following recursive relations on ξ and η.

$$\begin{aligned} \eta_{j+1} + \frac{\partial^2 V_j}{\partial^2 x}\xi_{j+1} &= \left(1 - \frac{\partial^2 V_j}{\partial x \partial y}\right)\eta_j \\ \xi_j + \frac{\partial^2 V_j}{\partial y^2}\eta_j &= \left(1 - \frac{\partial^2 V_j}{\partial x \partial y}\right)\xi_{j+1}. \end{aligned} \tag{5.18}$$

We shall see that this is related to the action of $\mathrm{d}\psi$. Since z is a critical point, it satisfies the equations

$$\begin{aligned} u_j(x_j, y_j) &= x_j + \frac{\partial V_j}{\partial y}(u_j(x_j, y_j), y_j) \\ v_j(x_j, y_j) &= y_j - \frac{\partial V_j}{\partial x}(u_j(x_j, y_j), y_j), \end{aligned} \tag{5.19}$$

where $\psi_j(x_j, y_j) = (u_j(x_j, y_j), v_j(x_j, y_j)$. Differentiating these equations, one can solve for the components of $\mathrm{d}\psi_j$, and find that the recursive relationship (5.18) is equivalent to

$$(\xi_{j+1}, \eta_{j+1}) = \mathrm{d}\psi_j(\xi_j, \eta_j), \qquad 0 \le j \le N-1$$

Using the chain rule to combine all these, and the boundary condition $z_0 = z_N$, we get

$$(\xi_0, \eta_0) = \mathrm{d}\psi_{z_0}(\xi_0, \eta_0).$$

Hence (ξ, η) is in the kernel of $H_z\Phi$ if and only if (ξ_0, η_0) is in the kernel of $I - \mathrm{d}\psi_{z_0}$. In particular, triviality of one kernel implies trivialtiy of the other. $\square$

This is a powerful result about symplectomorphisms of $\mathbb{R}^{2n}$, and it will be the key to proving the $\mathbb{T}^{2n}$ case of the Arnold Conjecture.

6 The Arnold Conjecture

Our goal is to relate fixed points of a Hamiltonian symplectomorphism to critical points of some function. In the case where ϕ_H^{t,t_0} is generated by some time-independent Hamiltonian $H_t = H$, the connection is obvious. If however we have a 1-periodic Hamiltonian, $H_{t+1} = H_t$, the 1-periodic solutions of the flow ϕ_H^{t,t_0} will correspond to fixed points of the time-1 map $\psi = \phi_H^{1,0}$. The periodic orbits will also correspond to critical points of the action integral

$$\Phi_H(z) = \int_{S^1} z^*\lambda - \int_0^1 H(t, z(t))\mathrm{d}t,$$

where $z : S^1 \to M$ is some loop in M. This is only well defined for exact manifolds. For contractible loops in non-exact manifolds, we can generalize by taking $u : B \to M$ to be some extension of z to the closed unit ball, and define

$$a_H(z, u) = \int_B u^*\omega - \int_0^1 H(t, z(t))\mathrm{d}t.^3$$

Counting critical points of this functional on an abstract infinite dimensional space was the motivation for developing Floer theory, which eventually solved the Arnold conjecture for closed manifolds. The construction of Floer homology, unfortunately, is too involved for this thesis. Our goal is to replace this infinite dimensional variational problem with a finite dimensional one. The key to this reduction is to exploit the linear structure of the universal cover $\mathbb{R}^{2n} \to \mathbb{T}^{2n}$. We use this to prove the following special case of the Arnold conjecture known as the Conley–Zehnder theorem.

Proposition 6.1 (Conley–Zhender)**.** *A Hamiltonian symplectomorphism $\psi : \mathbb{T}^{2n} \to \mathbb{T}^{2n}$ of the $2n$-torus with the standard symplectic structure has at least $2n + 1$ fixed points. If all the fixed points are nondegenerate, there are at least 2^{2n} fixed points.*

We will content ourselves with proving the nondegenerate case. The idea will be to use the discrete action functional Φ defined in (5.16). We know that $q : \mathbb{R}^{2n} \to \mathbb{R}^{2n}/\mathbb{Z}^{2n} \simeq \mathbb{T}^{2n}$ is a universal cover of the torus, so if $\phi_H^{1,0} : \mathbb{T}^{2n} \to \mathbb{T}^{2n}$ is a Hamiltonian symplectomorphism generated by the Hamiltonian isotopy ϕ_H^{t,t_0}, we get a lift to a Hamiltonian isotopy $\psi^{t,t_0} : \mathbb{R}^{2n} \to \mathbb{R}^{2n}$ such that

$$\psi^{t,t_0}(z + a) = \psi^{t,t_0}(z) + a \quad \forall a \in \mathbb{Z}^{2n}, z \in \mathbb{R}^{2n}.$$

We say that the flow is invariant under the action $(a, z) \mapsto z + a$ of $\mathbb{Z}^{2n}$ on $\mathbb{R}^{2n}$. As in the previous section, we can sample this flow at N points and get

$$\psi^{1,0} = \psi_N - 1 \circ ... \circ \psi_0$$

[3] Even with extra conditions, this map is only independent of choice of extension up to an integer. See [Sal99]

Since ψ is invariant under the action of $\mathbb{Z}^{2n}$, the Hamiltonian H' generating it must also be invariant. Since it is also periodic in t, we know that H' satisfies the boundedness condition (5.14). Hence we can find N such that each ψ_j has a generating function V_j of type V. Since the flow is invariant under the action of $\mathbb{Z}^{2n}$, it is easy to see from equation (5.15) that each V_j is also invariant under this action. We denote the space of periodic sequences $z = (x, y)$ with terms in $\mathbb{R}^{2n}$ satisfying $z_j = z_{N+j}$ by $X \simeq \mathbb{R}^{2nN}$. We can define the discrete symplectic action

$$\Phi(z) = \sum_{j=0}^{N-1} \langle y_j, x_{j+1} - x_j \rangle - V_j(x_{j+1}, y_j)$$

on X. In light of lemma 5.15, we know that a critical point of Φ on the space of periodic sequences corresponds to a fixed point of $\psi_H^{1,0}$, and that nondegeneracy carries over. We call two such fixed points geometrically distinct if they correspond to different points on the torus. This can be described in terms of the action

$$(a, \{z_n\}_{n\in\mathbb{N}}) \mapsto \{z_n + a\}_{n\in\mathbb{N}}$$

of $\mathbb{Z}^{2n}$ on X; two fixed points are geometrically distinct if and only if they are not related by this action. A quick calculation shows that Φ is invariant under the action, and hence descends to a function $\Phi : X/\mathbb{Z}^{2n} \to \mathbb{R}$. Now the critical points of this function corresponds to geometrically distinct fixed points of $\phi_H^{1,0}$, so all that remains to show is that if Φ is a Morse function, it has 2^{2n} critical points on $X/\mathbb{Z}^{2n}$. The standard Morse theory will not be of much use, however, since this space is not compact. In the next section we develop a localized version of Morse theory that will fix this issue.

6.1 Morse theory and the Conley index

Assume that ϕ^t is a flow on a locally compact metric space M with

$$\phi^{t+s} = \phi^t \circ \phi^s, \qquad \phi^0 = id_M.$$

A subset $\Lambda \subset M$ is called invariant if $\phi^t(\Lambda) = \Lambda$ for all t. We call a neighbourhood N of Λ an isolating neighbourhood if

$$\Lambda = I(N) = \bigcap_{t\in\mathbb{R}} \phi^t(N).$$

In this case we say that Λ is the largest invariant subset in N. Compact invariant subsets for which there exist isolating neighbourhoods are called isolated. We want to compute some sort of Morse-index-like invariant for a given isolated invariant set. The technical tool for this calculation is an index pair. See figure 4 for examples.

Definition 6.2. Let cl and int denote the topological closure and interior of sets respectively. An index pair for an isolated invariant subset Λ is a pair of compact subsets $L \subset N \subset M$ such that:

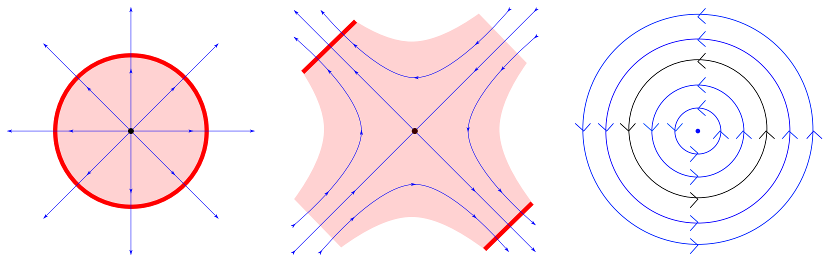

Figure 4: Three compact invariant sets, of which the two first are isolated, while the third is not. For the two hyperbolic critical points, N is marked in light red, and L in red. The circulation is not isolated since any neighbourhood will contain another circle, which is also an invariant set.

1. $\mathrm{cl}(N - L)$ is an isolating neighbourhood for Λ satisfying

$$\Lambda = I(\mathrm{cl}(N - L)) \subset \mathrm{int}(N - L).$$

2. L is positively invariant in N, that is,

$$x \in L, \ \ \phi^{[0,t]}(x) \subset N \implies \phi^t(x) \in L.$$

3. Every orbit which leaves N must pass through L, that is,

$$x \in N, \ \ \phi^t(x) \notin N \implies \exists t_0 \in [0,t] : \phi^{t_0} \in L.$$

A pair such that L is a deformation retract of one of its neighbourhoods in N is called regular. In this case the reduced homology of the quotient N/L is the same as the relative homology of the pair (N, L).

In [Con78], Conley proves that any isolated invariant set has a regular index pair, and that if M is a manifold, the quotient N/L has finite homology. Moreover, he shows that the homotopy type of the quotient N/L depends only on the invariant set Λ, not on the choice of index pair. This homotopy type is the Conley index of the isolated invariant set Λ. If the homology is finite, we define the index polynomial of Λ to be

$$p_\Lambda(s) = \sum_k \dim(H_k(N, L))s^k.$$

Note that throughout this thesis, we use homology with rational coefficients, so the notion of dimension is well defined. The index polynomial is additive in the sense that if Λ is the disjoint union of Λ_1 and Λ_2, then

$$p_\Lambda(s) = p_{\Lambda_1}(s) + p_{\Lambda_2}(s).$$

We wish to relate this new index theory to the more familiar Morse index. Let $\Phi : M \to \mathbb{R}$ be a Morse function on an n-dimensional Riemannian manifold M. Consider the gradient field given by

$$v(x) = -\nabla\Phi(x).$$

Assuming that this is complete, we denote it's flow by ϕ^t. Since Φ is Morse, all the fixed points of the flow are hyperbolic, so the unstable manifold

$$W^u(x) = \{y \in M | \lim_{t\to-\infty} \phi^t(y) = x\}$$

is well defined for any critical point x of Φ. The Morse index of x is defined as

$$\mathrm{ind}(x) := \dim W^u(x),$$

which can be computed as the number of eigenvalues with negative real part of the Hessian $H\Phi(x)$. As mentioned before, a hyperbolic fixed point is also an isolated invariant set, so we can compute its Conley index. Assume 0 is a hyperbolic fixed point of the flow

$$\dot{x} = v(x).^4$$

To find an index pair, consider the linearized system

$$\dot{\xi} = \mathrm{d}v(x)\xi$$

on $\mathbb{R}^n$. Denote the stable and unstable eigenspaces of $\mathrm{d}v$ by E^s and E^u respectively.[5] There exist some $r > 0$ such that the sets

$$\begin{aligned} N &= \{x_s + x_u : x_s \in E_s, x_u \in E_u, \ \ ||x_s||, ||x_u|| \leq r\} \\ L &= \{x_s + x_u \in N : ||x_u|| = r\} \end{aligned}$$

form an index pair for $\{0\}$. It is not hard to see that N/L is homotopy equivalent to S^k where $k = \dim E_u$. In the special case where $v(x) = -\nabla\Phi(x)$, the linearization is given by $\dot{\xi} = -H\Phi(0)\xi$, so the dimension of the unstable space is exactly the number of negative eigenvalues of $H\Phi(0)$. It follows that

$$p_{\{0\}}(s) = s^{\mathrm{ind}(0)}.$$

This can be generalized to arbitrary points on Riemannian manifolds by using the exponential map $\exp : T_xM \to M$ to make the approximation

$$x(t) \sim \exp(\xi(t)).$$

Where ξ is the solution to the linearized system $\dot{\xi} = \mathrm{d}v(x)$ on T_xM.

---

[4] This abuse of notation is typical in dynamical systems, and we really mean the flow of the first order ODE $\frac{\mathrm{d}}{\mathrm{d}t}\phi^t = v \circ \phi^t(x)$.

[5] The unstable space is spanned by the eigenvectors whose eigenvalue has positive real part, the stable by the ones with negative real part.

For a general isolated invariant set Λ of the gradient flow of Φ, we wish to relate the number of critical points

$$c_k(\Lambda) = \#\{x \in \Lambda : \mathrm{d}\Phi(x) = 0, \mathrm{ind}(x) = k\}$$

to the Conley–Betti numbers

$$b_k(\Lambda) = \dim H_k(N, L).$$

This is the essence of the Morse inequalities.

Proposition 6.3 (Morse inequalities)**.** *Let Φ be a Morse function on an n-dimensional Riemannian manifold M such that the gradient field is complete. If Λ is an isolated invariant set of the gradient flow, the following inequality holds for all $k \leq n$, with equality at $k = n$*

$$c_k(\Lambda) - c_{k-1}(\Lambda) + \ldots \pm c_0(\Lambda) \geq b_k(\Lambda) - b_{k-1}(\Lambda) + \ldots \pm b_0(\Lambda).$$

Proof. We fix a regular index pair (N, L) for Λ. For any regular value $a \in \mathbb{R}$ of $\Phi\big|_N$, we define

$$N^a = \{x \in N : \Phi(x) \leq a\} \cup L$$

For any critical value $c \in \mathbb{R}$, the set

$$\Lambda^c = \{x \in N : \Phi(x) = c, \mathrm{d}\Phi(x) = 0\}$$

is an isolated invariant set. In particular it is the disjoint union of a finite number of hyperbolic fixed points. Since critical values have measure zero, we can choose $a < c < b$ such that c is the only critical value in $[a, b]$. Then (N^b, N^a) is a regular index pair for Λ_c, and since Λ_c is the disjoint union of hyperbolic fixed points, we get

$$\sum_{k=0}^{n} H_k(N^b, N^a) s^k = \sum_{x \in \Lambda_c} s^{\mathrm{ind}(x)}. \tag{6.1}$$

We introduce the coefficients

$$\begin{aligned} b_k^a &= b_k^a(\Lambda) = \dim H_k(N^a, L), \\ c_k^a &= c_k^a(\Lambda) = \#\{x \in \Lambda \cap N^a : \mathrm{d}\Phi(x) = 0,\ \mathrm{ind}(x) = k\}. \end{aligned}$$

In these terms, equation (6.1) becomes

$$\dim H_k(N^b, N^a) = c_k^b - c_k^a. \tag{6.2}$$

It is a standard result of algebraic topology that the inclusion of pairs

$$(N^a, L) \longhookrightarrow (N^b, L) \longhookrightarrow (N^b, N^a)$$

gives rise to a long exact sequence in relative homology.

$$\cdots \longrightarrow H_{k+1}(N^b, N^a) \xrightarrow{\partial_k} H_k(N^a, L) \longrightarrow H_k(N^b, L) \longrightarrow H_k(N^b, N^a) \xrightarrow{\partial_{k-1}} \cdots$$

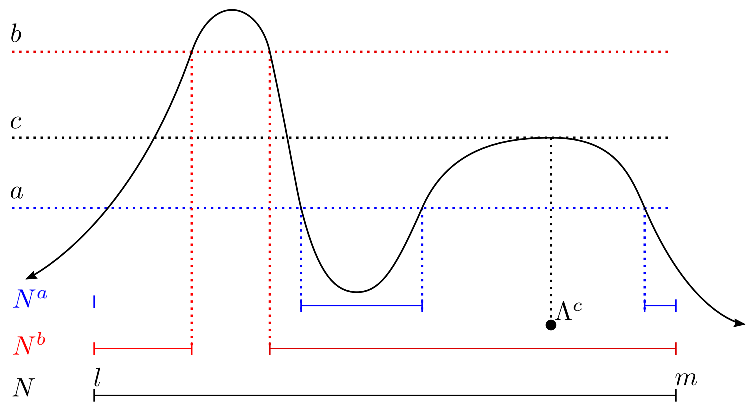

Figure 5: A one-dimensional example. The three critical ponts form an isolated invariant set which has regular index pair given by $N = [l, m], L = \{l, m\}$. It is not hard to see that the relative homology $H_k(N^b, N^a)$ is the same as the reduced homology $\tilde{H}_k(S^1)$.

We denote the rank of the connecting homomorphism ∂_k by $d_k^{a,b}$. Then combining the exactness of this sequence with (6.2), we get that

$$d_{k-1}^{a,b} + d_k^{a,b} = c_k^b - c_k^a - b_k^b + b_k^a. \tag{6.3}$$

We reformulate this by introducing the polynomials

$$p_{\text{crit}}^a(s) = \sum_k c_k^a s^k, \quad p_\Lambda^a(s) = \sum_k b_k^a s^k, \quad q^{a,b}(s) = \sum_k d^{a,b} s^k.$$

Then equation (6.3) is equivalent to

$$\begin{aligned}(1+s)q^{a,b}(s) &= p_{\text{crit}}^b(s) - p_{\text{crit}}^a(s) - (p_\Lambda^b(s) - p_\Lambda^a(s)) \\ &= P_{\text{crit}}^b(s) - p_\Lambda^b(s) - (p_{\text{crit}}^a(s) - p_\Lambda^a(s)).\end{aligned} \tag{6.4}$$

The coefficients of $q^{a,b}$ are non negative, intuitively, this means that when we move "up" from a to b, the number of fixed points increases at least as much as the Conley–Betti numbers. Since N is compact, and critical values have measure zero, we can find a finite cover $\{[a_j, a_{j+1}] : j = 0, .., l\}$ of $\Phi(N)$ such that there is exactly one critical value c_j in each $[a_j, a_{j+1}]$. We use (6.4) to combine all of the intervals. Since there are no critical points in N^{a_0}, we have $p_{\text{crit}}^{a_0}(s) = p_\Lambda^{a_0}(s) = 0$. Since $N = N^{a_l}$, we have that

$$p_\Lambda^{a_l}(s) = p_\Lambda(s) = \sum b_k s^k, \qquad p_{\text{crit}}^{a_l}(s) = p_{\text{crit}}(s) = \sum c_k s^k.$$

It follows by induction on j that

$$p_{\text{crit}}(s) - p_\Lambda(s) = (1+s)q(s)$$

where $q(s)$ has non negative coefficients. Doing polynomial division, we calculate that the coefficients d_k of q must be

$$d_k = \sum_{i=0}^{k} (-1)^j (c_{k-1} - b_{k-j}) \geq 0$$

Which is precisely the Morse inequalities. For equality at $k = n$, note that exactness of the sequence implies that the alternating sum of dimensions from 0 to n is zero. So with a_j as above,

$$\sum_{k=0}^{n} (-1)^k (c_k - b_k) = \sum_{j=0}^{l} \sum_{k=0}^{n} (-1)^k (c_k^{a_{j+1}} - c_k^{a_j} - b_k^{a_{j+1}} + b_k^{a_j}) = 0.$$

□

We immediately get the following corollary.

Corollary 6.4. *The number of critical points of Φ in Λ is bounded below by the sum of the Conley–Betti numbers.*

$$\#\mathrm{Crit}(\Phi\big|_{\Lambda}) = \sum_{k=0}^{n} c_k(\Lambda) \geq \sum_{k=0}^{n} b_k(\Lambda). \tag{6.5}$$

Proof. The Morse inequality at $k = 0$ yields the base-case. Inducting over k using the corresponding Morse inequality at each step gives the result. □

We are now ready to apply this theory to our special case.

6.2 Proof of the nondegenerate case

We identify the space of geometrically distinct periodic sequences $X/\mathbb{Z}^{2n}$ with the space $\mathbb{T}^{2n} \times \mathbb{R}^{2n(N-1)}$ by introducing the coordinates $(z_0, \zeta) \in \mathbb{T}^{2n} \times \mathbb{R}^{2n(N-1)}$ satisfying

$$\zeta_j = z_j - z_{j-1}.$$

We now claim that in these coordinates, Φ is given as by

$$\Phi(z) = \Psi(z_0, \zeta) = \langle \zeta, P\zeta \rangle + W(z_0, \zeta) \tag{6.6}$$

Where P is the $2n(N-1) \times 2n(N-1)$ matrix determined by

$$P = -\frac{1}{2} \begin{pmatrix} 0 & B \\ B^T & 0 \end{pmatrix}, \qquad B = \begin{pmatrix} I_n & \cdots & \cdots & I_n \\ 0 & \ddots & & \vdots \\ \vdots & \ddots & \ddots & \vdots \\ 0 & \cdots & 0 & I_n \end{pmatrix},$$

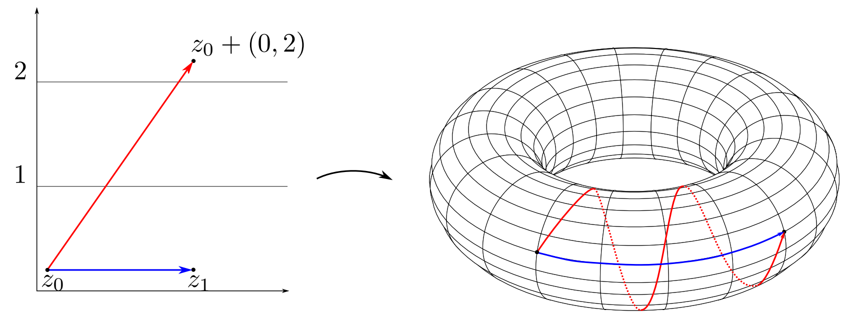

Figure 6: Different paths with same endpoint in $\mathbb{R}^2/\mathbb{Z}^2$. The red path clearly has higher kinetic energy.

and W descends to a function on $\mathbb{T}^{2nN} \simeq \mathbb{R}^{2nN}/\mathbb{Z}^{2nN}$. Too see why this holds, let $\zeta = \binom{\xi}{\eta}$, and calculate

$$
\begin{aligned}
\langle \zeta, P\zeta \rangle &= -\frac{1}{2} \begin{pmatrix} \xi^T & \eta^T \end{pmatrix} \begin{pmatrix} 0 & B \\ B^T & 0 \end{pmatrix} \begin{pmatrix} \xi \\ \eta \end{pmatrix} = -\xi^T B \eta \\
&= -\sum_{j=1}^{N-1} \langle \xi_j, \sum_{k=j}^{N-1} \eta_j \rangle \\
&= -\sum_{j=1}^{N-1} \langle x_j - x_{j-1}, y_{N-1} - y_{j-1} \rangle \\
&= \sum_{j=1}^{N-1} \langle x_j - x_{j-1}, y_{j-1} \rangle - \langle \, x_{N-1} - x_0, y_{N-1} \rangle \\
&= \sum_{j=0}^{N-1} \langle x_{j+1} - x_j, y_j \rangle. \qquad (6.7)
\end{aligned}
$$

We have used the periodic boundary conditions to exchange x_0 with x_N. Comparing this to the expression for $\Phi(z)$ found in (5.16), we see that taking

$$
W(z_0, \zeta) = \sum_{j=0}^{N-1} V_j(x_{j+1}, y_j),
$$

equation (6.6) is satisfied. Since each V_j is invariant under the action of $\mathbb{Z}^{2n}$, it is easily seen that W is invariant under the action of Z^{2nN}. Let us now break down this result intuitively. The components of ζ correspond to the "velocity" of the system through phase space. The integral of kinetic energy over a path is represented by (6.7), which we can see is a quadratic function of this velocity. This makes sense, since the average velocity of a path on the torus not only depends on its endpoints, but on how many times the path "wraps around" the torus. See figure 6. Therefore we expect that the action grows as this velocity grows. The W-term represents potential energy. Since potential energy is only dependant on position on the torus, it makes sense that it is invariant under $\mathbb{Z}^{2nN}$. Note that this also means that it is periodic, so that both W and all its derivatives are bounded. All of this shows that the gradient flow of Ψ is a compact perturbation of the quadratic flow determined by $\langle \zeta, P\zeta \rangle$, so we expect their large scale behaviour to agree. The gradient flow[6] ϕ associated to Ψ on $\mathbb{T}^{2n} \times \mathbb{R}^{2n(N-1)}$ is given by

$$\dot{z}_0 = -\frac{\partial W}{\partial z_0}, \qquad \dot{\zeta} = -P\zeta - \frac{\partial W}{\partial \zeta}, \tag{6.8}$$

The large scale qualitative behaviour of the system is determined by the eigenvalues and eigenvectors of P. Since P is symmetric, we know that all the eigenvalues are real. Furthermore, if

$$P \begin{pmatrix} \xi \\ \eta \end{pmatrix} = \frac{1}{2} \begin{pmatrix} -B\eta \\ -B^T\xi \end{pmatrix} = \lambda \begin{pmatrix} \xi \\ \eta \end{pmatrix}$$

for some $\lambda > 0$, then

$$P \begin{pmatrix} \xi \\ -\eta \end{pmatrix} = \frac{1}{2} \begin{pmatrix} B\eta \\ -B^T\xi \end{pmatrix} = \lambda \begin{pmatrix} -\xi \\ \eta \end{pmatrix} = -\lambda \begin{pmatrix} \xi \\ -\eta \end{pmatrix}.$$

This shows that $(\xi, \eta) \mapsto (\xi, -\eta)$ is a linear bijection between the spaces E^+ and E^- spanned by the stable and unstable eigenvectors respectively. Since P is nonsingular, we get a splitting

$$\mathbb{R}^{2n(N-1)} = E^+ \oplus E^-$$

where the summands each have dimension $n(N-1)$. We will now use this splitting to define an index pair.

Lemma 6.5. *There exists some $R \in \mathbb{R}$ such that the sets*

$$N = \{(z_0, \zeta^- + \zeta^+) : \zeta^\pm \in E^\pm, \ \ ||\zeta^\pm|| \leq R\}$$
$$L = \{(z_0, \zeta^- + \zeta^+) \in N : ||\zeta^-|| = R\}$$

form an index pair for the isolated invariant set

$$\Lambda = I(\mathrm{cl}(N - L))$$

and that all the critical points of Ψ are contained in Λ.

Intuitively, this lemma should not be hard to believe; outside some large compact set, the flow is very close to the linear flow $(0, -P\zeta)$, for which (N, L) is clearly an index pair. In particular, all the critical points of Ψ must be

[6] Note that we are avoiding questions of completeness and existence here. This is justified, since the flow exists on any compact set, and we only really need some large compact set to define our index pair.

contained within this large compact set. We postpone the full proof of this lemma to appendix B. All that remains now is to compute the homotopy type of the quotient N/L. We denote the closed $n(N-1)$-disk by $D^{n(N-1)}$. Note that we get a homotopy equivalence

$$\frac{N}{L} = \frac{\mathbb{T}^{2n} \times D^{n(N-1)} \times D^{n(N-1)}}{\mathbb{T}^{2n} \times D^{n(N-1)} \times \partial D^{n(N-1)}} \simeq \frac{\mathbb{T}^{2n} \times D^{n(N-1)}}{\mathbb{T}^{2n} \times \partial D^{n(N-1)}}$$

by contracting one copy of $D^{n(N-1)}$. The next lemma shows that this quotient is actually a suspension.

Lemma 6.6. *For any compact manifold X, there exists a homeomorphism*

$$\frac{X \times D^n}{X \times \partial D^n} \simeq \Sigma^n X_+,$$

where $X_+ = X \coprod \{*\}$ denotes X with a disjoint basepoint, and Σ^n the n-fold suspension (see appendix C).

Proof. The following diagram commutes

$$\begin{tikzcd}
X \times D^n \arrow[r, hook] \arrow[d, two heads] & X_+ \times D^n \arrow[r, two heads] & X_+ \times \frac{D^n}{\partial D^n} \arrow[r, two heads] & \frac{X_+ \times (D^n/\partial D^n)}{(X_+ \times \partial D^n) \cup (\{*\} \times (D^n/\partial D^n))} \arrow[d, "g"] \\
\frac{X \times D^n}{X \times \partial D^n} \arrow[rrr, two heads, "f"] & & & \frac{X \times (D^n/\partial D^n)}{X \times \partial D^n}
\end{tikzcd}.$$

The map f is clearly a bijection; it just collapses a set that is already collapsed. Since its domain is compact and its target is a Hausdorff space, it must be a homeomorphism. The map g is also a homeomorphism since distributing the product over the disjoint union gives

$$\frac{X_+ \times (D^n/\partial D^n)}{(X_+ \times \partial D^n) \cup (\{*\} \times D^n/\partial D^n)} = \frac{X \times (D^n/\partial D^n) \coprod (\{*\} \times (D^n/\partial D^n))}{(X \times \partial D^n) \coprod (\{*\} \times (D^n/\partial D^n))}.$$

The map g is just forgetting the second summand, which is killed off by the quotient anyways. The top-right quotient in the diagram is the definition of $\Sigma^n(X_+)$ when we associate $S^n \simeq (D^n/\partial D^n)$. $\square$

Using this lemma and the suspension isomorphism (C.1), we get isomorphisms

$$H_{k+n(N-1)}(n, L) \simeq \tilde{H}_{k+n(N-1)}(N/L) \simeq \tilde{H}_k(\mathbb{T}^{2n}_+) \simeq H_k(\mathbb{T}^{2n}).$$

In particular, the sum of the Conley–Betti numbers of (N, L) are bounded below by the sum of the Betti numbers of the $2n$-torus, which is the sum

$$\sum_{k=0}^{2n} \dim H_k(\mathbb{T}^{2n}) = \sum_{k=0}^{2n} \binom{2n}{k} = 2^{2n}.$$

Combining this with the Morse inequalities, the nondegenerate case of the Conley–Zhender theorem follows. $\square$

Appendix

A Mapping spaces

In this section we give a quick introduction to the Whitney topology on the space $C^r(M,N)$ of r times continuously differentiable maps from M to N. This approach is due to [Hir76], but one can also use the equivalent approach via jet-bundles found in for instance [Mat69]. Let $\phi : U \to \mathbb{R}^m$ and $\psi : V \to \mathbb{R}^n$ be coordinate functions on the open subsets $U \subset M$ and $V \subset N$. For any compact $K \subset U$, any C^r map g such that $g(K) \subset V$ and any $\varepsilon > 0$, we define a weak subbasic neighbourhood

$$N^r(g; U, V, K, \varepsilon)$$

as the set of all maps $f \in C^r(M,N)$ such that $f(K) \subset V$ and

$$||D^k(\psi \circ g \circ \phi^{-1})(x) - D^k(\psi \circ f \circ \phi^{-1})(x)|| < \varepsilon$$

for all $0 \leq k \leq r, x \in K$. In short, a map $f \in N^r(g; U, V, K, \varepsilon)$ is subject to both a compact open condition, as well as $r+1$ C^k conditions. The C^0-condition ensures that $f(x)$ is "close" to $g(x)$ for $x \in K$, the higher order C^k conditions ensure that the derivatives $D^k f$ and $D^k g$ are "close" on K. Taking these sets as a subbasis, we get a well defined topology on $C^r(M,N)$ which we call the weak topology. We denote the resulting space $C^r_W(M,N)$. The weak topology has a lot of nice features, specifically it has a complete metric and a countable base. The problem is that for non-compact manifolds it does not control behaviour "at infinity." Therefore, if we want global results about non-compact manifolds, we need a finer topology. Let $\Phi = \{\phi_j : U_j \to \mathbb{R}^m\}_{j \in \mathbb{N}}$ and $\Psi = \{\psi_j : V_j \to \mathbb{R}^n\}_{j \in \mathbb{N}}$ be collections of coordinate functions on M and N respectively such that the collections $\{U_j\}$ and $\{V_j\}$ are locally finite. Let $K = \{K_j\}_{j \in \mathbb{N}}$ be a sequence of compact sets such that $K_j \subset U_j$. Now for any map $g \in C^r(M,N)$ such that $g(K_j) \subset V_j$ for all j, and any positive sequence ε of real numbers, we define the strong basic neighbourhood

$$N^r(g; \Phi, \Psi, K, \varepsilon)$$

to be the set of maps $f \in C^r(M,N)$ such that for all $j \in \mathbb{N}, 0 \leq k \leq r$, and that for all $x \in K_j$, we have $g(K_j) \subset V_j$ and

$$||D^k(\psi_j \circ g \circ \phi_j^{-1})(x) - D^k(\psi_j \circ f \circ \phi_j^{-1})(x)|| < \varepsilon_j.$$

It is an easy exercise to see that the collection of such sets is closed under finite intersections. The topology generated by this basis is called the strong, or Whitney topology on $C^r(M,N)$. We denote the resulting topological space $C^r_s(M,N)$. It differs from the weak topology in the fact that we can control behaviour at an infinite number of compact sets at a time, not just a finite one. It is not hard to see that when M is compact, the strong and weak topologies agree. For proof of the following results about the strong topology, we refer to [Hir76].

Proposition A.1. *The sets* $\mathrm{Imm}^r(M,N)$ *and* $\mathrm{Sub}^r(M,N)$ *of* r *times continuously differentiable immersions and embeddings respectively are both open in* $C^r_s(M,N)$ *for all* $r \geq 1$.

The proof of this statement boils down to the fact that the subsets of injective and surjective linear transformations are open in the space $L(\mathbb{R}^m,\mathbb{R}^n)$ of linear transformations $l : \mathbb{R}^m \to \mathbb{R}^n$.

Proposition A.2. *The set* $\mathrm{Emb}^r(M,N)$ *of* r *times continuously differentiable embeddings from* M *into* N *is open in* $C^r_s(M,N)$ *for all* $r \geq 1$.

The proof uses the C^1-condition from the previous theorem to ensure maps are local embeddings, and a compact-open condition to ensure injectivity.

Proposition A.3. *The set* $Prop^r(M,N)$ *of* r *times continuously differentiable proper maps from* M *to* N *is open in* $C^r_s(M,N)$ *for all* $r \geq 0$.

Combining this proposition with the fact that an embedding $f : M \to N$ is proper if and only if $f(M)$ is closed we get the following corollary.

Corollary A.4. *The set of closed embeddings is open in* $C^r_s(M,N)$ *for all* $r \geq 1$.

This has an immediate application. Note that for connected manifolds, we now have that any closed embedding must also be surjective since its image is both open and closed. Now if M and N are not connected, we can use a C^0-condition to get a correspondence between the connected components of M and N. Thus we get the following result.

Proposition A.5. *The set* $\mathrm{Diff}^r(M,N)$ *of* r *times continuously differentiable diffeomorphisms of* M *with* N *is open in* $C^r_s(M,N)$ *for all* $r \geq 1$.

In our application to symplectomorphisms, we are interested in compact manifolds. All that remains to prove that in these topologies, pushforwards are continuous. A proof of the general case can be found in [Mat69]. This proof uses the equivalent jet bundle formulation of the topology, so for convenience, we give a proof of the compact case using more elementary techniques.

Proposition A.6. *Let* X, Y *and* Z *be smooth manifolds where* X *is compact. Let* $r \geq 0$*, then given any* $f \in C^r(Y,Z)$*, the pushforward map*

$$\begin{aligned} f_* : C^r_s(X,Y) &\to C^r_s(X,Z) \\ g &\mapsto f \circ g \end{aligned}$$

is continuous.

Proof. Since X is compact, the strong and weak topologies agree. Thus we need only consider weak subbasic neighbourhoods. Given any such neighbourhood $N = N^r(h; U, V, K, \varepsilon)$ in $C^r(X,Z)$, we must show that if $f_*(g) \in N$, then there exist some neighbourhood N' of g in $C^r(X,Y)$ such that $f_*(N') \subset N$. The set $f^{-1}(V)$ is open in Y, but it might not be a coordinate patch. However, for any

$y \in f^{-1}(V)$, there exists a coordinate patch that is diffeomorphic to an open ball of radius one – that is, there exist maps

$$\eta_y : W_y \xrightarrow{\sim} B(0;1) \subset \mathbb{R}^l$$

with $\eta_y(y) = 0$ and $W_j \subset f^{-1}(V)$. If we take $0 < r_y < 1$, then $B(0;r_y) \subset \overline{B(0;r_y)} \subset B(0;1)$. The sets $\eta_y^{-1}(B(0;r_y))$ cover $g(K)$, and by compactness there exists a finite subcover

$$g(K) \subset \bigcup_j \eta_j^{-1}(B(0;r_j)).$$

Now the sets

$$K_j = g^{-1}\left(\eta_j^{-1}\left(\overline{B(0;r_j)}\right) \cap g(K)\right)$$

are closed subsets of K, hence compact subsets of U with $g(K_j) \subset W_j$. For any $\varepsilon' > 0$, the set

$$N'(\varepsilon') = \bigcap_j N(g;U,W_j,K_j,\varepsilon')$$

is a neighbourhood of g such that all $a \in N'(\varepsilon')$ satisfy the compact-open condition $fa(K) \subset V$. We now need to make ε' small enough that the C^k-conditions are also satisfied. Let $\phi : U \to \mathbb{R}^n$ and $\psi : V \to \mathbb{R}^m$ be coordinate functions. We define

$$d = \min_{0 \le k \le r} \left(\inf_{x \in K} \left(\varepsilon - ||D^k(\psi h \phi^{-1})(x) - D^k(\psi f g \phi^{-1})(x)|| \right) \right),$$

which is positive since K is compact. We let $a \in N'(\varepsilon')$ as before and starting with the special case $k = 0$, we calculate that for $x \in K_j$,

$$\begin{aligned}
&||\psi h \phi^{-1}(x) - \psi f a \phi^{-1}(x)|| \\
&\quad \le ||\psi h \phi^{-1}(x) - \psi f g \phi^{-1}|| + ||\psi f g \phi^{-1}(x) - \psi f a \phi^{-1}(x)|| \\
&\quad \le \varepsilon - d + ||(\psi f \eta_j^{-1})(\eta_j g \phi^{-1})(x) - (\psi f \eta_j^{-1})(\eta_j a \phi^{-1})(x)||.
\end{aligned}$$

Since $\psi f \eta_j^{-1}$ is a continuous map between open subset of euclidean space, we have that given any $\varepsilon'' > 0$, there exist $\delta(\varepsilon'', x) > 0$ such that

$$\begin{aligned}
&||(\eta_j g \phi^{-1})(x) - (\eta_j a \phi^{-1})(x)|| < \delta(\varepsilon'', x) \\
\Longrightarrow\ &||(\psi f \eta_j^{-1})(\eta_j g \phi^{-1})(x) - (\psi f \eta_j^{-1})(\eta_j a \phi^{-1}(x))|| < \varepsilon''.
\end{aligned}$$

If we now take ε' smaller than

$$\delta = \min_{x \in K} \delta(d, x),$$

which again is positive by compactness, the C^0-condition is satisfied. For $1 \le k \le r$, we let

$$L = \max_{1 \le k \le r} ||D^k(\psi f \psi^{-1})||,$$

which is finite since the derivatives are bounded linear operators. Then, for any $x \in K_j$ and any $1 \leq k \leq r$,

$$
\begin{aligned}
&||D^k(\psi h\phi^{-1})(x) - D^k(\psi fa\phi^{-1})(x)|| \\
&\quad \leq ||D^k(\psi h\phi^{-1})(x) - D^k(\psi fg\phi^{-1})(x)|| + ||D^k(\psi fg\phi^{-1})(x) - D^k(\psi fa\phi^{-1})(x)|| \\
&\quad \leq \varepsilon - d + ||D^k(\psi f\eta_j^{-1}\eta_j g\phi^{-1})(x) - D^k(\psi f\eta_j^{-1}\eta_j a\phi^{-1})(x)|| \\
&\quad = ||D^k(\psi f\eta_j^{-1})(D^k(\eta_j g\phi^{-1})(x) - D^k(\eta_j a\phi^{-1})(x)|| \\
&\quad < \varepsilon - d + L \cdot \varepsilon'.
\end{aligned}
$$

So if we pick $\varepsilon' < \min(\frac{d}{L}, \delta)$, we get $f_*(N'(\varepsilon')) \subset N$ as desired. $\square$

The proof of the following theorem can also be found in [Mat69], again using the jet bundle formulation. As before we prove the compact case.

Proposition A.7. *For all manifolds X, Y and Z, and integers $r \leq 0$,*

$$C^r_s(X, Y \times Z) \simeq C^r_s(X, Y) \times C^r_s(X, Z)$$

Where both products have the product topology.

Proof. As in the previous proof, it suffices to work with the weak topology since X is compact. Let $\pi_1, \pi_2 : M \times M \to M$ denote the projections to each factor. Then the bijection

$$(\pi_1)_* \times (\pi_2)_* : C^1(M, M \times M) \to C^1(M, M) \times C^1(M, M)$$

is continuous since the factors are continuous by proposition A.6. It remains to show that the inverse is also continuous. We denote the inverse

$$
\begin{aligned}
\mu : C^1(M, M) \times C^1(M, M) &\to C^1(M, M \times M) \\
f_1 \times f_2 &\mapsto (f : x \mapsto (f_1(x), f_2(x)).
\end{aligned}
$$

We show that a subbasis for the weak topology on $C^1(M, M \times M)$ is given by neighbourhoods of the form $N(f; U, V_1 \times V_2, K, \varepsilon)$, where V_1, V_2 both are coordinate patches of M. Take any

$$g \in N = N(f; U, V, K, \varepsilon) \subset C^1(M, M \times M).$$

At $x \in g(K)$, we define $\psi_1 : V_1^x \to \mathbb{R}^n$ and $\psi_2 : V_2^x \to \mathbb{R}^n$ to be coordinate patches around $\pi_1(x)$ and $\pi_2(x)$ respectively, shrinking if necessary so that $V^x = V_1^x \times V_2^x \subset V$. Since $g(K)$ is compact, there exists a finite subcover $g(K) \subset \bigcup_j V^{x_j}$. Dividing K such that $\bigcup_j K_j = K$ and $g(K_j) \subset V^{x_j}$, [7] the subset

$$N_g(\varepsilon_g) = \bigcap_j N(g; U, V^{x_j}, K_j, \varepsilon_g)$$

[7]This is possible by the same argument as in the proof of A.6

is a well defined neighbourhood of g for all $\varepsilon_g > 0$. Now a simple re-centering argument show that there exist $\varepsilon_g > 0$ such that $N_g(\varepsilon_g) \subset N$; just take

$$\varepsilon_g < \min_{0\leq k\leq r} \min_{j} \inf_{x\in K_j} \left(\varepsilon - ||D^k(\psi f \phi^{-1})(x) - D^k(\psi g\phi^{-1})(x)||\right).$$

Now since

$$N = \bigcup_{g\in N} N_g(\varepsilon_g),$$

it suffices to show openness of the preimage $\mu^{-1}(M)$ of any neighbourhood

$$M = N(f; U, V_1 \times V_2, K; \varepsilon),$$

where $\psi_j : V_j \to \mathbb{R}^n, j = 1, 2$ are coordinate patches. Let $g_1 \times g_2 = \mu^{-1}(g), g \in M$. We construct a neighbourhood $M_1 \times M_2$ of $g_1 \times g_2$ as follows; for all $\varepsilon_1, \varepsilon_2 > 0$, we let

$$M_j(\varepsilon_j) = N(g_j; U, V_j, K, \varepsilon_j), \qquad j = 1, 2.$$

Then $M_1(\varepsilon_1) \times M_2(\varepsilon_2)$ is a well defined neighbourhood of $g_1 \times g_2$ since $g(K) \subset V \implies g_j(K) \subset \pi_j(V) = V_j$. We take

$$d = \min_{0\leq k\leq r} \inf_{x\in K} \left(\varepsilon - ||D^k(\psi f\phi^{-1})(x) - D^k(\psi g\phi^{-1})(x)||\right)$$

which is positive since K is compact. Now, if $h_j \in M_j(\varepsilon_j), h = \mu(h_1 \times h_2)$, we have that for all $x \in K, 0 \leq k \leq r$,

$$\begin{aligned}
&||D^k(\psi h\phi^{-1})(x) - D^k(\psi f\phi^{-1})(x)|| \\
&\qquad \leq \varepsilon - d + ||D^k(\psi h\phi^{-1})(x) - D^k(\psi g\phi^{-1})(x)|| \\
&\qquad \leq \varepsilon - d + \left\| \begin{pmatrix} D^k(\psi_1 h\phi^{-1})(x) \\ D^k(\psi_2 h\phi^{-1})(x) \end{pmatrix} - \begin{pmatrix} D^k(\psi_1 g\phi^{-1})(x) \\ D^k(\psi_2 g\phi^{-1})(x) \end{pmatrix} \right\| \\
&\qquad \leq \varepsilon - d + \varepsilon_1 + \varepsilon_2
\end{aligned}$$

So taking $\varepsilon_j < \frac{d}{2}$, we get our result. □

B Proof of the index pair lemma

In this appendix we give a proof of lemma 6.5. We must show that there exist some $R > 0$ such that the sets

$$\begin{aligned}
N &= \{(z_0, \zeta^- + \zeta^+) : \zeta^\pm \in E^\pm, \ \ ||\zeta^\pm|| \leq R\} \\
L &= \{(z_0, \zeta^- + \zeta^+) \in N : ||\zeta^-|| = R\}
\end{aligned}$$

are an index pair for the set $\Lambda = I(\mathrm{cl}(N - L))$ such that all the critical points of Ψ are contained in Λ. So we must prove that for sufficiently large R:

1. $\Lambda \subset \mathrm{int}(N - L)$.

2. L is positively invariant in N – that is,

$$x \in L,\ \ \phi^{[0,t]}(x) \subset N \implies \phi^t(x) \in L.$$

3. Every orbit which leaves N must pass through L – that is,

$$x \in N,\ \ \phi^t(x) \notin N \implies \exists t_0 \in [0,t] : \phi^{t_0} \in L.$$

4. $\mathrm{d}\Psi(x) = 0 \implies x \in \Lambda$.

We denote the orthogonal projection to the stable and unstable eigenspaces of P by $\pi^\pm : \mathbb{R}^{n(N-1)} \to E^\pm$. Let $-\lambda^-_{\mathrm{min}}$ be the negative eigenvalue with smallest absolute value. Choosing some orthogonal basis for E^-, one may easily calculate that

$$\begin{aligned}
||\pi^-(\zeta - P\zeta)||^2 &= ||\zeta^-||^2 - 2\langle\zeta^-, \pi^-(P\zeta)\rangle + ||\pi^-(P\zeta)|| \\
&\geq ||\zeta^-||^2 + 2\lambda^-_{\mathrm{min}}||\zeta^-||^2 + ||\lambda^-_{\mathrm{min}}\zeta||^2 \\
&= (||\zeta^-|| + \lambda^-_{\mathrm{min}}||\zeta^-||)^2.
\end{aligned}$$

Since W is a periodic C^2 function, its first and second derivatives exist, and are bounded. Hence there exist K such that

$$||\frac{\partial W}{\partial \zeta}(z_0, \zeta)|| \leq K, \quad \forall (z_0, \zeta) \in \mathbb{T}^{2n} \times \mathbb{R}^{2n(N-1)},$$

and we may Taylor expand the curve $\zeta(t)$ given by $\zeta(0) = \zeta_0$ as

$$\zeta(t) = \zeta_0 - t\left(P\zeta + \frac{\partial W}{\partial \zeta}(z_0, \zeta)\right) + t^2 e(z_0, \zeta_0)$$

where $||e(z_0, \zeta_0)|| \leq E$ for all (z_0, ζ). We then calculate that for $t > 0$,

$$\begin{aligned}
||\pi^-(\zeta(t))|| &= ||\pi^-(\zeta_0 - t(P\zeta_0 + \frac{\partial W}{\partial \zeta}(z_0, \zeta_0)) + t^2 e(\zeta, z_0)|| \\
&\geq ||\zeta_0^-|| + t(\lambda^-_{\mathrm{min}}||\zeta_0^-|| - K - tE).
\end{aligned}$$

So if we take $||\zeta_0^-|| > R' = (K + E)/\lambda^-_{\mathrm{min}}$, the E^- component of $\zeta(t)$ grows linearly with t for $t \leq 1$. Note that we can perform an analogous argument for $\pi^+(\zeta(t))$. Hence there exist some R' such that outside $N(R')$, the flow increases $||\zeta_-||$ and decreases $||\zeta_+||$ when we move forwards in time, and increases $||\zeta_+||$ when we move backwards in time. Taking $R > R'$, it is now easy to verify 2. since any $\zeta \in L$ has $\phi^t(\zeta) \notin N$. By the intermediate value theorem, any curve leaving N must have some point with either ζ^+ or ζ^- larger than R'. In the region $N(R) - N(R')$, the flow strictly decreases $||\zeta^+||$, so the only way for a curve to exit $N(R)$ is by reaching some $||\zeta^-|| > R$. Using the intermediate value theorem again, we can see that 3. is satisfied. To see why 1. holds, note that

$\Lambda \subset N(R')$, since any point outside $N(R')$ escapes to infinity either as $t \to \infty$ or $t \to -\infty$. A quick triangle-inequality argument shows that

$$\begin{aligned} ||\nabla\Psi(z_0,\zeta)|| &\geq ||P\zeta + \frac{\partial W}{\partial \zeta}(z_0,\zeta)|| \\ &\geq ||P\zeta|| - K \geq \frac{||\zeta||}{||P^{-1}||} - K. \end{aligned}$$

Thus we can take R such that outside $N(R)$, we have $||\zeta|| \geq K||P^{-}||$. Thus all the critical points of Ψ are contained within $N(R)$. Any critical point in $N(R)$ must clearly be in Λ, so 4. holds as well. □

C The suspension isomorphism

In this appendix we briefly cover some properties of homology and suspension. The reference for the material in this section is [May99]. In the category of pointed spaces, the smash product behaves like a tensor product which makes it nice for computing homotopy and homology. It is defined as follows.

Definition C.1. Given two pointed spaces X and Y their smash product is defined as the quotient

$$X \wedge Y = \frac{X \times Y}{X \vee Y} = \frac{X \times Y}{(\{x\} \times Y) \cup (X \times \{y\})}$$

Where x and y are the basepoints of X and Y respectively.

If we work with certain ”convenient”, spaces, the smash product is both symmetric and associative –that is, we have homeomorphisms

$$X \wedge (Y \wedge Z) \simeq (X \wedge Y) \wedge Z, \quad X \wedge Y \simeq Y \wedge X.$$

In particular, the smash product is associative for all locally compact Hausdorff spaces, and hence for all manifolds. Furthermore, one can show that similarly to a tensor product, there exist an isomorphism

$$F(X \wedge Y, Z) \simeq F(X, F(Y, Z))$$

where $F(X,Y)$ denotes the pointed space of basepoint preserving maps from X to Y, with the constant map as basepoint. Since we are extra interested in the space $\Omega X = F(S^1, X)$, it is natural that we are interested in the suspension $\Sigma X := X \wedge S^1$. We define the n-fold suspension by iteration

$$\Sigma^n X = \Sigma^{n-1}\Sigma(X).$$

The reader may verify that $S^1 \wedge S^n \simeq S^{n+1}$. Hence, using the associativity, it is easy to see that an equivalent definition of the n-fold suspension is

$$\Sigma^n X = X \wedge S^n.$$

A somewhat related construction is that of a reduced cone on X

Definition C.2. For a pointed space X, the reduced cone on X is the space CX defined as

$$CX = I \wedge X = \frac{X \times I}{(X \times \{0\}) \cup (\{x_0\} \times I)},$$

where $I = [0, 1]$ is the unit interval with basepoint 0.

It is not hard to see that the inclusion $X \hookrightarrow CX$ at 1, mapping x to $(x, 1)$ gives rise to a quotient map $CX \to CX/X \simeq \Sigma X$. This map glues together the endpoints of I, hence creating a based circle. Now the cone of any space is a contractible space, since we can contract it to the basepoint along I. The long exact sequence of the triple (x_0, X, CX) is

$$\cdots \longrightarrow H_k(CX, x_0) \longrightarrow H_k(CX, X) \xrightarrow{\partial} H_{k-1}(X, x_0) \longrightarrow H_{k-1}(CX, x_0) \longrightarrow \cdots.$$

Using that excision gives an isomorphism $H_*(CX, X) \simeq H_*(CX/X, x_0)$, and denoting the reduced homology by $\tilde{H}_k := H_k(X, x_0)$ where x_0 is the basepoint of X, the above sequence becomes

$$\cdots \longrightarrow 0 \longrightarrow \tilde{H}_k(\Sigma X) \xrightarrow{\partial} \tilde{H}_{k-1}(X) \longrightarrow 0 \longrightarrow \cdots .$$

This means that the boundary map induces an isomorphism

$$\tilde{H}_k(\Sigma X) \simeq \tilde{H}_{k-1}(X). \tag{C.1}$$

This isomorphism is known as the suspension isomorphism, and is such a fundamental concept that it is included as an axiom in the formulation of generalized reduced homology theories.

References

[CdS06] Ana Cannas da Silva. *Lectures on Symplectic Geometry.* Lecture Notes in Mathematics. Springer-Verlag Berlin Heidelberg, 1 edition, 2006.

[Con78] Charles Conley. *Isolated invariant sets and the Morse index*, volume 38 of *Regional conference series in mathematics.* R.I.: American Mathematical Society, 1978.

[Hir76] Morris W. Hirsch. *Differential Topology.* Graduate Texts in Mathematics. Springer-Verlag New York, 1 edition, 1976.

[Mat69] John N. Mather. Stability of C^∞ mappings: II. Infinitesimal stability implies stability. *Annals of Mathematics*, 89(2):254–291, 1969.

[May99] Jon Peter May. *A concise course in algebraic topology.* University of Chicago Press, 1 edition, 1999.

[MS98] Dusa McDuff and Dietmar Salamon. *Introduction to Symplectic Topology.* Oxford Mathematical Monographs. Oxford University Press, 2 edition, 1998.

[Sal99] Dietmar Salamon. Lectures on floer homology. *Symplectic geometry and topology (Park City, UT, 1997)*, 7:143–229, 1999.

[Wei81] Alan Weinstein. symplectic geometry. *Bulletin of the American Mathematical Society*, 5, 1981.

NTNU

Norwegian University of
Science and Technology