Symplectic Reduction

Sam Qunell

Abstract

In classical mechanics, the state of a physical system is determined by the positions and momenta of the involved particles. This allows us to define coordinates (x_i, p_i) which describe the "phase space" of our system. Symplectic manifolds are a mathematical model for these systems. With this mathematical language, the elimination of continuous symmetries of our system to obtain a space with fewer coordinates is referred to as "symplectic reduction". In this writeup, I will discuss and prove the Marsden-Weinstein-Meyer theorem that describes how to walk through this symplectic reduction process.

1 Introduction

As is known in undergraduate physics, the dynamical behavior of a classical mechanical system can be studied via the symmetries of its equations of motion. Suppose a massive particle moves in an n dimensional smooth manifold M (Euclidean space, a hypersphere, etc.) with local generalized coordinates $\{q_i\}_{1\leq i\leq n}$. The infinitesimal behavior of the system can be represented by a point in the cotangent bundle T^*M, where (q,p) represents the particle at position q moving in direction p. This cotangent bundle is sometimes called the **phase space** of our system. Due to natural construction of the cotangent bundle, we may define a canonical 2-form on T^*M locally. We write $\omega = \sum_{i=1}^n dq_i \wedge dp_i$. This form is independent of the choice of generalized coordinates, and has several useful properties. Most importantly, its restriction to any tangent space is a nondegenerate bilinear form, and so, the map $TM \to T^*M$, $v_p \to \omega_p(v_p, -)$ is an isomorphism of vector bundles. (T^*M, ω) is the canonical example of a symplectic manifold.

This isomorphism means that, given the Hamiltonian function H for our physical system, its differential $dH \in \Gamma(M, T^*M)$ is naturally associated to some vector field $X_H \in \Gamma(M, TM)$ that describes the motion of our particle. This correspondence makes symplectic manifolds a natural setting to study symmetries of equations of motion. In particular, it gives us a mathematical language to describe how symmetries in the physical setup of our configuration space correspond to simplifications that can be made in the equations of motion. For example, a particle in 3 dimensional space travelling only along the x axis has cylindrical symmetry, and hence, our configuration space may be viewed as only one dimensional.

Symplectic reduction is the mathematical formulation of this principle. Symmetries in our phase space may be described as a smooth group action of a Lie group G on T^*M. Under sufficiently nice conditions of this action, we may produce a quotient manifold from T^*M that is also symplectic, and is essentially our "reduced" version of the phase space. The Marsden-Weinstein-Meyer theorem states this precisely.

Theorem 1.1 (Marsden-Weinstein-Meyer)**.** *Let (M,ω) be a symplectic manifold, and let G be a Lie group with a Hamiltonian action on (M,ω), along with its associated moment map $\mu : M \to \mathfrak{g}^*$. Suppose that G acts freely and properly on $\mu^{-1}(0)$. Then the orbit space $\mu^{-1}(0)/G$ (sometimes written $M//G$) is a symplectic manifold.*

This writeup is an exposition of the formalism used in this theorem, along with an organized proof. I will first define several key concepts used in the proof, such as Hamiltonian actions and moment maps. I will then prove Theorem 1 in several steps, describing both the manifold structure and symplectic structure on $\mu^{-1}(0)/G$. To conclude, I will discuss some variations and applications of this symplectic reduction process. Throughout the paper, I will highlight the physical intuition of each definition and proposition as much as is possible, as well as provide examples from physical systems. The organization roughly follows [1], although with more details and examples.

2 Preliminaries

Despite the technical language used in the statement of the main theorem, the underlying ideas are easy to describe, especially through the lens of classical mechanics.

2.1 Symplectic Linear Algebra and Geometry

Definition. A **symplectic manifold** (M, ω) is a smooth manifold M along with a closed, nondegenerate 2 form ω.

By nondegenerate, we mean that ω_p is a nondegenerate, bilinear alternating form on each tangent space T_pM. This manifold often represents the phase space of a physical system, and ω allows us to "link" a position coordinate with its corresponding momentum in such a way as to allow for the aforementioned duality from vector fields to 1 forms. Vector spaces with such a bilinear form will be important to our study, as they will be used to represent the infinitesimal motion of our particle. This theory will be outlined briefly.

Definition. A **symplectic vector space** (V, ω) is a finite dimensional vector space V over $\mathbb{R}$ with a nondegenerate, alternating bilinear form ω.

By this definition, every tangent space of a symplectic manifold is a symplectic vector space. We will use the following fact as a sanity check later on.

Fact. Symplectic vector spaces, and therefore symplectic manifolds, are even dimensional.

This is usually proven by explicitly producing an even dimensional basis via nondegeneracy of ω. We will want to study subspaces of symplectic vector spaces that also respect this symplectic structure.

Definition. For (V, ω) a symplectic vector space and W a vector subspace, we define the **symplectic complement**

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

Note that V^ω is itself a vector subspace.

Definition. The subspace W is **isotropic** if $W \subset W^\omega$, and is **coisotropic** if $W^\omega \subset W$.

Definition. A submanifold N of the symplectic manifold (M, ω) is **(co)isotropic** if, for each $p \in N$, T_pN is a (co)isotropic subspace of T_pM.

The dimensions of W and W^ω are closely related.

Fact. $\dim(W) + \dim(W^\omega) = \dim(V)$.

This should not be too surprising, since elementary linear algebra gives that $\dim(W) + \dim(W^\perp) = \dim(V)$. The proofs of both equations use that the associated bilinear form (ω or any inner product) are nondegenerate, and such maps yield an isomorphism $V \to V^*$, $v \to \omega(v, -)$.

2.2 Hamiltonian Vector Fields

In the context of physical systems, our equations of motion are generally governed by kinetic and potential energies. In the most basic physical systems, total energy is conserved in motion. More generally, we have a Hamiltonian function $H : M \to \mathbb{R}$ that often coincides with a system's total energy, and is (usually!) conserved. If we assume H is a smooth function, then we can define a 1-form dH. In order to describe the motion of a particle in our manifold, we need to present a flowline, or vector field, following its trajectory. Fortunately, due to the symplectic structure, this is actually possible.

Fact. In a symplectic manifold M, smooth vector fields correspond bijectively with smooth 1-forms under the map $X \to \iota_X \omega = \omega(X, -)$.

This was discussed briefly in the intro, and this fact follows from the nondegeneracy of ω. So, for any Hamiltonian H, we can produce a vector field X_H with $\omega(X_H, -) = dH$. We can also go in reverse.

Definition. A vector field X on M is called **Hamiltonian** if $\omega(X, -)$ is exact.

The physical connection with this definition is that the Hamiltonian function H is actually *preserved* along the flowlines of X_H. This fact is essential to our physical reasoning, so we present a proof.

Lemma 2.1. *If X_H is the Hamiltonian vector field corresponding to H, then H is constant on any flow line (integral curve) of X_H.*

Proof. The integral curves of X_H yield a one parameter family of diffeomorphisms $\rho_t : M \to M$, where fixing x gives the integral curve through x. We rewrite the value of H along our curve algebraically as $H(\rho_t(x)) = \rho_t^* H(x)$. Since ρ_t is the flow corresponding to the Hamiltonian vector field X_H, the definition of Lie derivative gives that $\frac{d}{dt}(\rho_t^* H)|_{t=0} = \mathcal{L}_{X_H} H$. Cartan's magic formula gives that this is

$$\mathcal{L}_{X_H} H = d\iota_{X_H} H + \iota_{X_H} dH = \iota_{X_H} dH = \iota_{X_H} \omega(X_H, -) = \omega(X_H, X_H) = 0.$$

The time derivative of $H(\rho_t(x))$ is zero, so $H(\rho_t(x))$ is equal to $H(x)$ for all t. $\square$

So, it is always possible to define a vector field that conserves a single smooth function on M. These flows are, in fact, those adopted by physical systems via Hamiltonian mechanics, although the full proof of this fact uses variational calculus, more of the compatible complex structures on (M, ω), and physical assumptions that are independent from ZFC.

We can change our perspective on the matter to make it more easily generalized. This one parameter family of diffeomorphisms ρ_t corresponding to H is equivalent to a smooth group action of $\mathbb{R}$ on M, where $r * x = \rho_r(x)$. This is an extremely basic kind of symmetry on M. Unfortunately, physical systems generally exhibit multiple kinds of symmetry, and it is not obvious how to define several simultaneous vector fields on M that are somehow mutually compatible with each other and with the symplectic structure of M. We do so with *Hamiltonian group actions*.

2.3 Hamiltonian Group Actions

As before, we represent the general symmetries of a physical system with a (real) Lie group G acting on M. Without further restrictions, a group action can be incredibly pathological, so we must enforce several conditions.

Two of these conditions are more obvious for an action on a symplectic manifold. Firstly, we require that the action of G on M be smooth, meaning that $ev : G \times M \to M$, $ev(g, p) = g * p$ is smooth at all $(g, p) \in G \times M$. This gives us a group homomorphism $\psi : G \to Diff(M)$, taking each G element to its corresponding diffeomorphism. We also require that this action respect the symplectic structure on M (we say that G acts symplectically). This is expressed as, for all $g \in G$, we have the equality on the pullback $\psi(g)^* \omega = \omega$. This means that ψ has image in $Symp(M, \omega)$, the group of symplectomorphisms of (M, ω). For the rest of this writeup, we assume implicitly that all group actions are smooth, and when appropriate, symplectic.

In the context of physical systems, our equations of motion are generally governed by several distinct components of momentum and angular momentum. In a generalized set of coordinates, these momenta can take strange forms. We appeal to $\mathfrak{g}$, the Lie algebra of G, to make this precise.

Definition. Let $\mu : M \to \mathfrak{g}^*$ be any function. If $\langle -, - \rangle$ is an inner product for the pairing $\mathfrak{g}^* \times \mathfrak{g}$, then for $\zeta \in \mathfrak{g}$, define the **ζ-component of** μ,

$$\mu^\zeta : M \to \mathbb{R}, \mu^\zeta(p) = \langle \mu(p), \zeta \rangle.$$

Definition. For $\zeta \in \mathfrak{g}$, let X^ζ be the vector field of M given by

$$(X^\zeta)_p = \frac{d}{dt}(e^{t\zeta} * p)|_{t=0}.$$

Definition. The symplectic action $\psi : G \to Symp(M, \omega)$ is a **Hamiltonian action** with **moment map** μ if the following two conditions hold.

1.
$$d(\mu^\zeta) = \iota_{X^\zeta}\omega = \omega(X^\zeta, -),$$
i.e., "μ respects the symplectic structure on M". Note that this says that each component $\langle \mu, \zeta \rangle$ is a Hamiltonian function for the vector field X^ζ. This encodes the fact that several different quantities are conserved by the infinitesimal action of G.

2. μ is G-equivariant. More formally,
$$\forall g \in G, \mu \circ \psi(g) = Ad_g^* \circ \mu.$$
This means that "μ respects the action of G on M".

μ is named evocatively; it is our mathematical parallel to the momenta or other conserved quantities of our physical system. A basic physical example will make this connection clear.

Example. Let $SE(3) = \mathbb{R}^3 \rtimes SO(3)$ be the special Euclidean group in three dimensions. The product structure is $(v, A) * (w, B) = (Aw + v, AB)$. The canonical action on $\mathbb{R}^3$ is $(v, A) * x = Ax + v$. The special Euclidean group is the group of all (determinant 1) isometries on $\mathbb{R}^3$, and it includes all rotations and translations. This makes it a natural group to act on a physical system.

Suppose a single particle moves in $\mathbb{R}^3$, so that our phase space is the cotangent bundle $T^*\mathbb{R}^3 = \mathbb{R}^6$. The action of $SE(3)$ on $\mathbb{R}^3$ lifts to an action on this cotangent bundle, with $(v, A) * (x, y) = (Ax + v, Ay)$. It is easily verified that this action is smooth and symplectic.

To evaluate the corresponding moment map, we first look at the Lie algebra. The Lie algebra of $SO(3)$ is the set of all 3 by 3 skew-symmetric matrices, which is a 3 dimensional vector space. In fact, $\mathfrak{so}(3)$ is isomorphic as a Lie algebra to $(\mathbb{R}^3, \times)$, for $\times$ the cross product. One can show that $\mathfrak{se}(3)$ is $\mathbb{R}^3 \times \mathbb{R}^3$ as a vector space, with bracket given by $[(\vec{v}, \vec{x}), (\vec{w}, \vec{y})] = (\vec{x} \times \vec{w} - \vec{y} \times \vec{v}, [\vec{x}, \vec{y}]) = (\vec{x} \times \vec{w} - \vec{y} \times \vec{v}, \vec{x} \times \vec{y})$ [2].

Consider the map
$$\mu : T^*M = \mathbb{R}^6 \to \mathfrak{g}^* = \mathbb{R}^6, \mu(x, v) = (\vec{v}, \vec{x} \times \vec{v}).$$
The vector notation is only used to distinguish between T^*M elements and $\mathfrak{g}^*$ elements. One can check that this is indeed a moment map for our $SE(3)$ action. Now, for $\zeta = (\vec{a}, \vec{0}) \in \mathfrak{g}$, we compute that
$$\mu^\zeta(x, v) = \langle \mu(x, v), (\vec{a}, \vec{0}) \rangle = \langle (\vec{v}, \vec{x} \times \vec{v}), (\vec{a}, \vec{0}) \rangle = \vec{v} * \vec{a}.$$
This is precisely the component of the linear momentum of our particle in the direction $\vec{a}$. Similarly, for $\zeta = (\vec{0}, \vec{b}) \in \mathfrak{g}$, we compute that
$$\mu^\zeta(x, v) = \langle \mu(x, v), (\vec{0}, \vec{b}) \rangle = \langle (\vec{v}, \vec{x} \times \vec{v}), (\vec{0}, \vec{b}) \rangle = (\vec{x} \times \vec{v}) * \vec{b}.$$
This is exactly the component of angular momentum of our particle in the direction $\vec{b}$.

We show a generalization of Lemma 2.1 that is well-studied in undergraduate physics. Let (M, ω) be a symplectic manifold with (G, μ) a Hamiltonian action and corresponding moment map. Then $f \in C^\infty(M)$ is called a **symmetry** if f is G invariant, and f is called an **integral of motion** if μ is constant on the integral curves of X_f. The Noether principle shows that these are the same.

Lemma 2.2 (Noether principle)**.** *In the conditions described above, f is a symmetry iff it is an integral of motion.*

Proof. We can reduce the given definitions to computations along components of $\mathfrak{g}$ elements. So, we have that
$$f \text{ is a symmetry} \iff \mathcal{L}_{X^\zeta} f = 0 \; \forall \zeta \in \mathfrak{g},$$

$$f \text{ is an integral of motion } \iff \mathcal{L}_{X_f}\mu^\zeta = 0 \ \forall \zeta \in \mathfrak{g}.$$

For a fixed $\zeta \in \mathfrak{g}$, we compute that

$$\mathcal{L}_{X^\zeta} f = \iota_{X^\zeta} df = \iota_{X^\zeta}\iota_{X_f}\omega = -\iota_{X_f}\iota_{X^\zeta}\omega = -\iota_{X_f} d\mu^\zeta = -\mathcal{L}_{X_f}\mu^\zeta,$$

so one is zero whenever the other is. □

This type of relationship between conserved quantities of motion and symmetries of our system underlies the spirit of symplectic reduction. While it is not used explicitly in the proof of the main theorem, one should keep in mind the picture that the orbits of the G action are describing a coordinate that is unused and so conserved dynamically.

3 Marsden-Weinstein-Meyer Theorem

We now have the tools needed to proceed with the proof. The goal of the theorem is to show that the quotient $\mu^{-1}(0)/G$ is a symplectic manifold, although there are several intermediate constructions that are needed. It is not guaranteed, in general, that either of $\mu^{-1}(0)$ or $\mu^{-1}(0)/G$ are smooth manifolds at all.

To start, we need to restrict our attention to $\mu^{-1}(0)$. We do this for two reasons. Firstly, this captures the notion of our physical system having some conserved quantity during its motion. Secondly, it is not always possible that M/G can even be symplectic,. Symplectic manifolds must always be even dimensional, but when M/G is a manifold, it has dimension $\dim(M) - \dim(G)$, which is even only when $\dim(G)$ is.

Lemma 3.1. *Let M be a symplectic manifold, and let G be a Lie group with a Hamiltonian action on (M, ω) with associated moment map $\mu : M \to \mathfrak{g}^*$. Then $\mu^{-1}(0)$ is closed under the action of G.*

Proof. Fix $p \in \mu^{-1}(0)$. Since the action of G is Hamiltonian, μ is G-equivariant. So,

$$\forall g \in G, \mu(g * p) = Ad_g^* \circ \mu(p) = Ad_g^* \circ 0.$$

Ad_g^* is linear, so this is zero, i.e., $g * p \in \mu^{-1}(0)$, as desired. □

This lemma, and in fact, the entirety of our main theorem, only focuses on $\mu^{-1}(0)$ among all other level sets because 0 is always a fixed point of the coadjoint representation. In fact, all of these results will work for $\mu^{-1}(\zeta)$ for ζ any such fixed point. Nontrivial ζ are less commonly studied, although in the case of G a torus, $\mathfrak{g}$ is abelian, so all points of $\mathfrak{g}^*$ are fixed points. More generally, if O is a coadjoint orbit in $\mathfrak{g}^*$, then under similar assumptions, $\mu^{-1}(O)/G$ can be made into a symplectic manifold [6].

So, $\mu^{-1}(0)$ is a union of G orbits. However, in order to better understand the G action restricted to $\mu^{-1}(0)$, we will need to study these orbits more carefully.

Lemma 3.2. *Let M be a smooth manifold and let G be a Lie group acting freely and properly on M. Then for any $p \in M$, the orbit $G * p$ is an embedded submanifold of M that is homeomorphic to G.*

Proof. The map $f_p : G \to M$, $g \to g*p$ is injective since our action is free, and it is also proper by assumption. Consider the map

$$(df_p)_e : \mathfrak{g} \to T_pM.$$

By definition, $(df_p)_e(\zeta) = 0$ iff $e^\zeta * p = p$. But the action is free, and so e^ζ must be the identity element. Since the Lie subalgebra of $\mathfrak{g}$ corresponding to $\mathrm{Stab}(p) = \{e\}$ must be zero dimensional, we have that $\zeta = 0$ as well. Similarly, for a general $g \in G$, we note that right multiplication by g gives us a diffeomorphism of M, written R_g, and hence, an explicit isomorphism of tangent spaces

$$\mathfrak{g} = T_eG = d(R_{g^{-1}})_g(T_gG).$$

Then

$$(df_p)_g = d(f_p \circ R_g)_e \circ (dR_{g^{-1}})_g.$$

We see that $f_p \circ R_g = f_{g*p}$, and so

$$d(f_p \circ R_g)_e = d(f_{g*p})_e,$$

which is injective by our previous work. Furthermore, $dR_{g^{-1}}$ is an isomorphism of vector spaces, so also injective. This gives that the composite $(df_p)_g$ is injective. So, (df_p) is injective, and so f_p is an embedding of smooth manifolds. □

This allows us to study the function $d\mu_p$ in more detail, since we can now describe the orbits of any point in $\mu^{-1}(0)$.

Lemma 3.3. *Let (M, ω) be a symplectic manifold, and let G be a Lie group with a Hamiltonian action on (M, ω) with associated moment map $\mu : M \to \mathfrak{g}^*$. Suppose G acts freely and properly on $\mu^{-1}(0)$. Then $\mu^{-1}(0)$ is a submanifold of M with dimension* $\dim(M) - \dim(G)$.

Proof. We will apply the regular value theorem after showing that 0 is a regular value. Fix $p \in \mu^{-1}(0)$ and $\zeta \in \mathfrak{g}$. Since our action is Hamiltonian, we have $\omega_p(X_p^\zeta, v) = d(\mu^\zeta)(v)$ for any $v \in T_pM$. By the definition of μ^ζ, we compute that

$$d(\mu^\zeta)(v) = d\langle \mu_p, \zeta\rangle(v) = \langle d\mu_p(v), \zeta\rangle.$$

We wish to show that $d\mu_p$ is full rank. By Rank-Nullity Theorem, it is sufficient to show that its kernel has dimension

$$\text{null}(d\mu_p) = \dim(M) - \dim(\mathfrak{g}^*) = \dim(M) - \dim(G).$$

Since the bilinear pairing $\langle -, -\rangle$ is nondegenerate, we have that $d\mu_p(v) = 0$ iff $\langle d\mu_p(v), \zeta\rangle = 0$ for all $\zeta \in \mathfrak{g}$ iff $\omega_p(X_p^\zeta, v) = 0$ for all $\zeta \in \mathfrak{g}$. Let V be the subspace of T_pM generated by all X_p^ζ. Then $d\mu_p(v) = 0$ iff $v \in V^\omega$, the symplectic complement of V. Elementary symplectic linear algebra shows that

$$\dim(V^\omega) = \dim(T_pM) - \dim(V) = \dim(M) - \dim(V).$$

So, we have that $d\mu_p$ is surjective iff $\dim(V) = \dim(G)$. It remains to characterize V and show this equality.

Consider $G * p$, the orbit of p in M under the action of G. By Lemma 3.2, $G * p$ is a submanifold of M homeomorphic to G. Note that in $G * p$, the tangent directions are exactly those vectors determined by the infinitesimal action of $\mathfrak{g}$ on p. More formally, $v \in T_p(G * p)$ iff there exists $\zeta \in \mathfrak{g}$ with $\frac{d}{dt}(e^{t\zeta} * p)|_{t=0} = v$ iff $v = X_p^\zeta$ for some $\zeta \in \mathfrak{g}$ iff $v \in V$. Therefore, $\dim(V) = \dim(G * p) = \dim(G)$, as desired. □

Note that, via the dimension computations in this lemma, we have shown that, at every point $p \in \mu^{-1}(0)$, $T_pG * p \subset T_p\mu^{-1}(0) = \ker(d\mu_p) = (T_pG * p)^\omega$. Therefore,

Corollary 3.4. *In the conditions above, $G * p$ is an isotropic submanifold of M, and $\mu^{-1}(0)$ is a coisotropic submanifold of M.*

The fact that $\mu^{-1}(0)$ is coisotropic will allow us to push the form ω to the quotient $\mu^{-1}(0)/G$. Coisotropy essentially says that the tangent spaces within the quotient are "big enough" for the resulting form to be nondegenerate.

We now need to quotient out the action of G. Formally, we define the equivalence relation $\sim$ on $\mu^{-1}(0)$, with $p \sim q$ iff $q \in G * p$, and then define the quotient topological space $\mu^{-1}(0)/\sim := \mu^{-1}(0)/G$, sometimes called the **orbit space**. Note that the equivalence classes of this relation are exactly the G orbits, each of which is an embedded submanifold of $\mu^{-1}(0)$ homeomorphic to G by Lemma 3.2. Furthermore, one can show that, in a suibtable sense, nearby G orbits are parallel, meaning that the connected components of these equivalence relations form what is called a **foliation** of $\mu^{-1}(0)$. One would expect that these conditions imply that $\mu^{-1}(0)/G$ is well-behaved enough to inherit a smooth structure from $\mu^{-1}(0)$. In fact, this is the case, but the proof of this fact is extremely long and technical, and uses little more of the G action and none of the symplectic structure. For these reasons, only an outline of the proof is provided. A complete proof can be found in [4].

Theorem 3.5 (Quotient Manifold Theorem). *Let M be a smooth manifold, and let G be a Lie group acting freely and properly on M. Then the orbit space M/G is a smooth manifold of dimension* $\dim(M) - \dim(G)$*, with topology and smooth structure induced by the projection $\pi : M \to M/G$.*

Sketch.

1. The projection π is an open map since, for open $U \subset M$, $\pi^{-1} \circ \pi(U) = G * U$, which is open.
2. Properness of the action implies that $O = \{(p, q) \subset M \times M \mid g * q = p \text{ for some } g \in G\}$ is closed in $M \times M$. Openness of π implies then that the diagonal of $M/G \times M/G$ is closed. This is another characterization of M/G being Hausdorff.
3. π is open and surjective, so M/G is second countable since M is.
4. The connected components of the G orbits are, in fact, leaves of a foliation of M.
5. By definition of a foliation, there exists a cubic atlas (all charts are hypercubes) on M, (U_i, ϕ_i), such that each $G * p \cap U_i \neq \emptyset$ implies that $\phi_i(G * p \cap U_i)$ is a countable union of hyperplanes, where each hyperplane is given by an orthogonal slice $(x_1, ...x_{\dim(M)-\dim(G)}) = (c_1, ...c_{\dim(M)-\dim(G)})$ for some choice of c_i.
6. Because G acts freely and properly, we may refine our cubic atlas to (U_i', ϕ_i') so that $\phi_i'(G * p \cap U_i')$ is at most a single orthogonal slice of the cube. Each such U_i' is therefore essentially a product (choice of orbit, point in orbit), where each component is itself an open cube.
7. Since π is open and surjective, $\pi(U_i')$ is an open Euclidean cover of M/G, where the corresponding chart essentially projects U_i' onto its choice of orbit component.
8. This choice of atlas also inherits the smooth structure from the original cubic atlas on M, and so M/G is a smooth manifold with $\pi : M \to M/G$ a smooth surjective submersion. □

The full proof also makes clear the following corollary.

Corollary 3.6. *In the conditions above, there is a natural isomorphism $T_{[p]}(M/G) \sim T_pM/T_p(G * p)$.*

In the conditions of the main theorem, we therefore have that $\mu^{-1}(0)/G$ is a smooth manifold. Note that we in fact have a bit more; $\mu^{-1}(0)/G$ has the structure of a **principal G-bundle** with total space $\mu^{-1}(0)$. These principal bundles are fairly common in mathematical physics, as they embody the idea of quotienting by smooth symmetries. The main observation of our theorem is that this same quotient will respect the symplectic structure of M since $\mu^{-1}(0)$ is coisotropic. Note that we have $\dim(\mu^{-1}(0)/G) = \dim(\mu^{-1}(0)) - \dim(G) = \dim(M) - 2\dim(G)$, which is always even. So, it is at least possible that this quotient actually is symplectic.

To complete the proof of our theorem, we need to show that the quotient $\mu^{-1}(0)/G$ inherits a version of the symplectic form ω, and is therefore a symplectic manifold in its own right. We first show that this is the case on the level of tangent spaces.

Lemma 3.7. *Suppose that (V, ω) is a symplectic vector space, and W is a coisotropic subspace. Then the quotient W/W^ω inherits a symplectic form from ω.*

Proof. Note that coisotropy is needed to even define the quotient W/W^ω. It is somewhat clear what the form should be. Define the form ω_{red} on W/W^ω, with $\omega_{red}([v], [w]) = \omega(v, w)$. We must first show that this is well-defined. Suppose that $[v] = [v']$ and $[w] = [w']$. Then $v' = v + x$ and $w' = w + y$ for some $x, y \in W^\omega$. Therefore, by bilinearity,

$$\omega(v', w') - \omega(v, w) = \omega(v + x, w + y) - \omega(v, w) = \omega(v, y) + \omega(x, w) + \omega(x, y).$$

Since $v, w \in W$ and $x, y \in W^\omega \subset W$, each of these three terms is zero by definition. Therefore, ω factors through W^ω, and so ω_{red}, the reduction, is a well-defined alternating bilinear form on W/W^ω.

For ω_{red} to be symplectic, we must also show that ω_{red} is nondegenerate. So, suppose that there exists $[v] \in W/W^\omega$ with $\omega_{red}([v], [w]) = 0$ for all $[w] \in W/W^\omega$. Therefore, $\omega(v, w) = 0$ for all $w \in W$. By definition, we have that $v \in W^\omega$, and so $[v] = 0$. $(W/W^\omega, \omega_{red})$ is therefore a symplectic vector space. □

Combining Lemma 3.4, Lemma 3.6, and Lemma 3.7, we have

Corollary 3.8. *In the conditions above, with a choice of p, $T_{[p]}(\mu^{-1}(0)/G)$ inherits a symplectic form from $T_p(\mu^{-1}(0))$.*

We conclude the proof of the main theorem by showing that this form is well-defined on choice of p, and can be extended globally to a symplectic form on $\mu^{-1}(0)/G$.

Lemma 3.9. *In the conditions of Theorem 1.1, there exists a symplectic form ω_{red} on $\mu^{-1}(0)/G$ that is equal to the inherited form on each $T_{[p]}(\mu^{-1}(0)/G)$.*

Proof. The intuitive thing to try is to define $\omega_{red[p]}([v],[w]) = \omega_p(v,w)$, since this is clearly smooth when it is well-defined. Furthermore, for $\pi : \mu^{-1}(0) \to \mu^{-1}(0)/G$ the projection, we would have $\pi^*\omega_{red} = \omega$. To see that this form is actually well-defined on a choice of p, choose p, and $v, w \in T_p(\mu^{-1}(0))$. Consider $g * p \in G * p$, and suppose we have $x, y \in T_{g*p}(\mu^{-1}(0))$ such that $D\pi x = d\pi v$ and $D\pi y = D\pi w$. Then

$$\omega_{red[g*p]}([x],[y]) = \omega(x,y) = g * \omega(v,w) = \omega(v,w) = \omega_{red[p]}([v],[w]),$$

where we have used that G acts symplectically. So, ω_{red} is a well-defined 2-form on $\mu^{-1}(0)/G$.

To see that this form is closed, we simply compute that

$$\pi^*(d\omega_{red}) = d\pi^*\omega_{red} = d\omega = 0,$$

since ω on $\mu^{-1}(0)$ is closed. Since π is a surjective submersion, π^* is injective on forms. Therefore, $d\omega_{red} = 0$. We already know that ω_{red} is nondegenerate on each tangent space, so ω_{red} is a symplectic form. □

This finishes the proof of Theorem 1.1, giving us that

Corollary 3.10. *$(\mu^{-1}(0)/G,\ \omega_{red})$ is a symplectic manifold of dimension* $\dim(M) - 2\dim(G)$.

Example. Suppose a particle moves in $\mathbb{R}^2$, and suppose that we know that our system exhibits circular symmetry about the origin. We know S^1 acts on $\mathbb{R}^2$ via rotation: $e^{i\theta} * (x,y) = (\cos(\theta)x + \sin(\theta)y, -\sin(\theta)x + \cos(\theta)y)$. This lifts to an action of S^1 on the cotangent space $T^*\mathbb{R}^2 = \mathbb{R}^4$, where $e^{i\theta} * (x, y, v_x, v_y) = (\cos(\theta)x + \sin(\theta)y, -\sin(\theta)x + \cos(\theta)y, \cos(\theta)v_x + \sin(\theta)v_y, -\sin(\theta)v_x + \cos(\theta)v_y)$. The canonical symplectic form on $\mathbb{R}^4$ is $\omega = dxdv_x + dydv_y$. One checks that the action is in fact symplectic.

What should the moment map be? We turn to physics. Our Lie group is S^1, so $\mathfrak{g} = \mathbb{R}$. The angular momentum for a particle rotating about the origin is proportional to $r\dot{\theta}$. This suggests changing coordinates to $(x, y, v_x, v_y) \to (r, \theta, v_r, v_\theta)$, and defining $\mu = v_\theta$ (v_θ is the linear velocity along $\hat{\theta}$, not the angular velocity $\dot{\theta}$. It is exactly equal to $r\dot{\theta}$). We may omit the origin from our analysis so that this choice of coordinates is well-defined and so that the action is now free.

Our action is now $e^{i\phi} * (r, \theta, v_r, v_\theta) = (r, \theta + \phi \pmod{2\pi}, v_r, v_\theta)$, and $\omega = drdv_r + d\theta dv_\theta$. To see that this is a Hamiltonian S^1 action, note that μ is automatically S^1 equivariant since $\mathfrak{s}^1 = \mathbb{R}$, which is abelian. We need only check one dimension for μ to be Hamiltonian. The flow of our action gives the vector field $\frac{d}{d\theta}$. We check $\iota_{d/d\theta}\omega = dv_\theta = d\mu$. Our action is proper since S^1 is compact, and evidently free. So, we are prepared to apply the main theorem.

We need not restrict ourselves to when $\mu = 0$ since $\mathfrak{s}^1$ is abelian. So, for any $x \in \mathbb{R}$, $\mu^{-1}(x) \sim T^*\mathbb{R}^+ \times S^1$ with coordinates (r, θ, v_r). The orbits are the sets of elements with fixed r and v_r. Therefore, one can see that $\mu^{-1}(x)/S^1 \sim T^*\mathbb{R}^+ = \{(r, v_r) \mid r \in \mathbb{R}^+, v_r \in T_r\mathbb{R}^+\}$. As we expected, the reduced phase space will be essentially a choice of radial position and velocity. As in the theorem, we compute that the symplectic form on this space will be $\omega = drdv_r$.

4 Applications

Symplectic reduction is a tool for refining a system of differential equations to a reduced set describing a physical system's "true" dynamics. The utility within both physics and math is at least clear conceptually, but due to the restrictive conditions imposed on the action of our group G, symplectic reductions are somewhat limited in their applications. That said, several variations of the Marsden-Weinstein-Meyer theorem allow for computing a similar reduction for other types of manifolds, or to actions with less restrictive conditions. A few of these, as well as some other related pieces of research, will be briefly outlined here.

4.1 Orbifolds

As was briefly described after the proof of Lemma 3.1, one way to soften the conditions of the main theorem is to consider $\mu^{-1}(O)$, where O is any coadjoint orbit of G, as opposed to any particular fixed point $\mu^{-1}(\zeta)$. Another technique, which is perhaps most physically valid, is to eliminate the condition that G must act freely on $\mu^{-1}(0)$. One can show that when 0 is a regular value of μ and G is a torus, that the stabilizer of any $p \in M$ under G must be discrete and finite, hence still tractable. While the Quotient Manifold Theorem does not apply as is, a variation can show that $\mu^{-1}(0)/G$ is locally a quotient of some $\mathbb{R}^n$ by part of our group action. This gives $\mu^{-1}(0)/G$ the structure of an **orbifold**. Orbifolds naturally arise in similar situations where a group of symmetries acts on a topological space, and some form of quotient is needed. Visually, an orbifold is a "manifold with singularities", and careful attention must be given to dynamical behavior around these singular points. A good review of the topic and early theory of how to perform reduction in the singular case is given in [5].

4.2 Kähler Manifolds

A Kähler manifold is a special kind of symplectic manifold that contains additional complex and Riemannian structure. They occur naturally in a comlpex algebraic geometric setting, and one of their specializations, the Calabi-Yau manifolds, hold a critical role in mirror symmetry. It is natural to inquire whether our procedure of symplectic reduction can be extended to the Kähler case. One method of doing so is presented in [3], where we consider a complexified version of our group action. The proof and setup are essentially the same, except that one needs also to show that the extra complex and Riemannian structures (or in the case of the paper, a certain polarization) are also inherited by the quotient.

4.3 Geometric Quantization

Symplectic manifolds arise naturally as a means of studying differential equations like those found in classical mechanics. Unfortunately, the physical world is not purely dictated by Newtonian (or Hamiltonian) physics. The Heisenberg uncertainty principle states that it is in fact impossible to know simultaneously a quantum particle's position and momentum, meaning that the bare symplectic structure of our phase space is insufficient for studying quantum behavior. Quantum systems are often mathematically realized as Hilbert spaces along with associated wavefunctions or operators corresponding to observables. In physics, it is an active field of research as to how the theories of classical and quantum mechanics can be unified in a way that is valid at either length scale. In mathematical terms, this roughly means that for a symplectic manifold (M, ω) representing our phase space, there should be a Hilbert space $\mathcal{H}$ that represents the same system. This mathematical procedure is referred to as *quantization* There are several proposed methods of producing such a Hilbert space from a symplectic manifold, although each has its own weaknesses, and there is no general consensus as to any "canonical" method.

One means of assessing such a quantization scheme is whether symmetries of the classical manifold are also symmetries of the quantum Hilbert space. One formulation of this, as described in [3], is that "localization commutes with quantization". More formally, if we have a function Q taking our symplectic manifold to the corresponding Hilbert space, then under an appropriate quotient, $Q(M//G) = Q(M)/G$.

This process was conjectured in [3], along with a few results for a specific scheme of quantization known as geometric quantization. The conjecture was proven in the subsequent decades. One proof of this conjecture is primarily analytic and provides some applications to the case of Kähler manifolds [7].

References

[1] Ana Cannas da Silva, *Lectures on symplectic geometry*, Lecture Notes in Mathematics, vol. 1764, Springer-Verlag, Berlin, 2001. MR 1853077

[2] Jean Gallier and Jocelyn Quaintance, *Differential geometry and Lie groups—a computational perspective*, Geometry and Computing, vol. 12, Springer, Cham, [2020] ©2020. MR 4164718

[3] V. Guillemin and S. Sternberg, *Geometric quantization and multiplicities of group representations*, Invent. Math. **67** (1982), no. 3, 515–538. MR 664118

[4] John M. Lee, *Introduction to smooth manifolds*, second ed., Graduate Texts in Mathematics, vol. 218, Springer, New York, 2013. MR 2954043

[5] Eugene Lerman, Richard Montgomery, and Reyer Sjamaar, *Examples of singular reduction*, Symplectic geometry, London Math. Soc. Lecture Note Ser., vol. 192, Cambridge Univ. Press, Cambridge, 1993, pp. 127–155. MR 1297133

[6] Dusa McDuff and Dietmar Salamon, *Introduction to symplectic topology*, third ed., Oxford Graduate Texts in Mathematics, Oxford University Press, Oxford, 2017. MR 3674984

[7] Youliang Tian and Weiping Zhang, *An analytic proof of the geometric quantization conjecture of Guillemin-Sternberg*, Invent. Math. **132** (1998), no. 2, 229–259. MR 1621428