Opening the Door to Stable Hamiltonian Topology

A gentle introduction

Lay de Jong

M.Sc. thesis presented for the programme Mathematical Sciences

Supervisor: Dr. Álvaro del Pino Gomez

Universiteit Utrecht

July 7, 2024

Utrecht, The Netherlands

Acknowledgements

A master's thesis falls into the category of written works where one has the unique opportunity to personally thank a significant portion of the ever-to-exist readers, of which I will gladly make use of.

First, I would like to thank Álvaro for his patience and guidance during this thesis. The process become thoroughly enjoyable due to the encouragement to embrace learning, which I have most certainly done. I promise to not apply the $\mathbb{Z}_2$ or $SO(2)$ group action on the diacritic in your name again.

Secondly, I want to express my appreciation for the Utrecht mathematics department; both students and professors, in particular Gil. I think there are few places where the environment is both so supportive as inspiring.

Lastly, a big collection of thanks to my friends and family. The supremum of which is directed towards my mother and brother; Liesbeth and Tyan, for always in any circumstance being there for me.

I leave with the following quote, which I believe describes all mathematicians.

> "*Owl explained about the Necessary Dorsal Muscles. He had explained this to Pooh and Christopher Robin once before and had been waiting for a chance to do it again, because it is a thing you can easily explain twice before anybody knows what you are talking about.*"
>
> A.A. Milne

Contents

1 Introduction

If we were to kick down the door to the office of a stable Hamiltonian topologist and ask them what a stable Hamiltonian structure is, we might receive the answer that they are a simultaneous generalization of contact structures and taut foliations defined by closed 1-forms. To study their topology is to study their structures up to homotopy. They might walk up to their always nearby blackboard and chalk down:

Stable Hamiltonian Structure.
A stable hamiltonian structure is a closed 2-form ω on an odd-dimensional manifold M^{2n+1}, together with a 1-form λ satisfying:

$$\lambda \wedge \omega^n \neq 0, \text{and } \ker(\omega) \subset \ker(d\lambda).$$

It is now that one's familiarity with mathematical vocabulary maps this information to a point on the spectrum ranging from utter nonsense to crystal clear nigh tautologies. We will assume a point well in between these two extremities and provide this thesis to translate the map towards the latter.

1.1 A Hamiltonian Knock

Stable Hamiltonian structures naturally arise in the context of dynamical systems. We will briefly describe the historical development leading to stable Hamiltonian structures. The astute reader might have noticed the capitalized 'Hamiltonian', and it should come as no surprise that the origins are to be traced back to dynamical Hamiltonian systems. The study of periodic solutions, aptly named orbits, of differential equations may well find its most prominent roots in celestial mechanics. Within this realm one can dwell arbitrarily far into the past. But we will satisfy ourselves with going back to a series of papers by Rabinowitz [Rab78] where he proves the existence of periodic orbits on a prescribed energy surface of a Hamiltonian system of ordinary differential equations, subject to some conditions. Here an energy surface is defined as the level set of a Hamiltonian, and naturally arises in the context of mathematical physics.

Following this observation, Weinstein publishes [Wei79]. It is here that the departure from analysis to geometry takes place. In this paper is first posed the now prominent Weinstein conjecture.

Conjecture [Weinstein]. *If $S \subset (W, \Omega)$ is a compact hypersurface of contact type with $H^1(S;\mathbb{R}) = 0$, then $\ker(\Omega)|_S$ has a closed orbit.*

The Weinstein conjecture is posed from a natural question: to what level can we relax Rabinowitz' requirements in order to find periodic orbits. Hypersurfaces of contact type are the direct precursor of stable Hamiltonian structures. Such a hypersurface has the interesting property that they are stable; there exists a family of diffeomorphic hypersurfaces which possess the same dynamics. Meaning if we have a found such a closed orbit, we are able to lift it to a nearby family of hypersurfaces which locally foliate our symplectic manifold. It shows that we only have to prove Weinstein conjecture nearby S in order to have proven it on S. Weinstein further poses that the particular Hamiltonian function used to define such a hypersurface is irrelevant, it depends on geometric properties.

Stable Hamiltonian systems on one hand generalize this concept of stability, which forms the foundation of many studies of the Weinstein conjecture. On the other hand they lean on the fact that these hypersurfaces locally foliate the symplectic manifold.

The first mention of a stable Hamiltonian structure does not yet go by this name. It was by Hofer and Zehnder [HZ12]. They remark that any contact form λ on a closed orientable manifold W has an associated unique vector field, the Reeb vector field, adhering to the properties

$$i_X d\lambda = 0 \text{ and } i_X \lambda = 1.$$

And they rephrase the Weinstein conjecture; dropping the necessity of embedding W into a symplectic manifold.

Conjecture [Weinstein II]. *For every closed odd-dimensional manifold M with contact form λ, its Reeb vector field has a periodic orbit.*

Weinstein's conjecture was, and still is, one of the driving forces behind symplectic topology. The full conjecture remains an open problem, but it has been proven in different capacities. Weinstein's conjecture has fairly recently been proven to be true in the three-dimensional case by Taubes [Tau07]. A logical next step has been to generalize its three-dimensional statement to stable Hamiltonian structures.

1.2 A Peek into the Homotopy Principle

Stable Hamiltonian structures are also a natural framework when studying the h-principle on contact structures and taut foliations. The h-principle is a way to look at different geometric structures, where the 'h' stands for homotopy. It studies weak homotopy equivalences between spaces of geometric structures bestowed with the C^∞-topology. We have now encountered four different structures for a three-dimensional manifold M to be interested in:

- Contact structures, denoted by $\mathcal{CS}(M)$,
- Stable Hamiltonian structures, denoted by $\mathcal{SHS}(M)$,
- Foliations, denoted by $\mathcal{F}(M)$,
- Taut foliations, denoted by $\mathcal{TF}(M)$.

Historically speaking we are interested in the following diagram

$$\begin{array}{ccccccc}
\mathcal{CS}(M) & & & & & & \mathcal{TF}(M) \\
\downarrow & \searrow & & & & \swarrow & \downarrow \\
& & \mathcal{SHS}(M) & & \mathcal{F}(M) & & \\
\downarrow & \swarrow & & & & \searrow & \downarrow \\
\mathcal{CS}^f(M) & & \longrightarrow & & \longrightarrow & & \mathcal{F}^f(M).
\end{array} \tag{1.1}$$

The arrows denote some sort of inclusion; contact structures naturally include into the space of stable Hamiltonian structures, and taut foliations naturally include into the space of foliations. Without going into too much detail, the superscript f stands for a formalization of the respective structure. The important facts are first that $\mathcal{CS}^f$ deals with a pair λ, ω: a 1-form and 2-form, such that $\lambda \wedge \omega \neq 0$. So it generalizes both $\mathcal{SHS}$ and $\mathcal{CS}$. And second that homotopically speaking

$$\mathcal{CS}^f \cong \mathrm{Map}(M, S^2) \times \mathbb{Z}_2$$

is a nice space. This is intuitively due to the fact that a contact form is a nowhere-vanishing section, hence trivialization, of the cotangent bundle.

It turns out that regarding weak homotopy equivalences, we have:

- $\mathcal{CS}(M) \not\simeq \mathcal{CS}^f(M)$,
- $\mathcal{TF}(M) \not\simeq \mathcal{F}(M)$,
- $\mathcal{F}(M) \simeq \mathcal{F}^f(M)$,

respectively due to Bennequin, Novikov and Thurston.

How does this tie into stable Hamiltonian structures? Observe so-called confoliations defined in [ET98] as

$$\lambda \wedge d\lambda \geq 0,$$

which generalizes both foliations and contact structures. Note that in the space of confoliations, taut foliations defined by a closed 1-form agree only with $\lambda \wedge d\lambda = 0$, whilst all others agree with a contact structure. Now it turns out foliations can be perturbed to obtain contact structures. Specifically, taut foliations perturb to tight contact structures, a certain type of contact structure defined as the complement of overtwisted contact structures.

Theorem [Eliashberg & Thurston]. *The space of contact structures is open and dense in the space of confoliations bestowed with C^0-topology.*

We can visualize the space of confoliations as in figure 1. The green lines represent the complement of contact structures, these have measure 0 in the space. Red dots and crosses are respectively normal and taut foliations. These

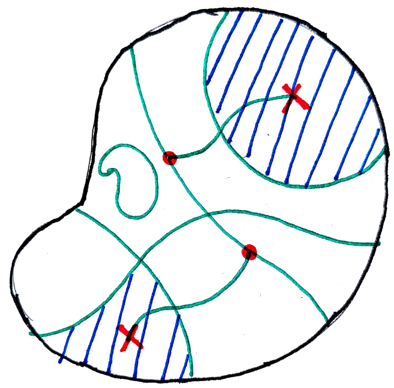

Figure 1: The Space of Confolations

can be perturbed to respectively contact structures or tight contact structures, the latter represented by the blue shaded areas.

Every taut foliation has a close neighbourhood of tight contact structures, which is bounded by a line of foliations. Now, every contact structure λ naturally induces an exact stable Hamiltonian structure $(\lambda, d\lambda)$, and homotopies of contact forms induce homotopies of stable Hamiltonian structures. However, Cieliebak and Volkov prove this is not a bijection of homotopy classes.

It is these types of observations which motivate the study of stable Hamiltonian topology in the context of h-principles. They seem to be the natural structure which arises if one wants to study the homotopy principle in relation to contact structures and foliations.

1.3 A Thesis Roadmap

The paper is meant to pave the way in order to understand the main concepts behind stable Hamiltonian topology.

In section 2 we will explain the theory of distributions, and provide the necessary vocabulary of these within the context of differential geometry. The section leads through the establishment of Frobenius' theorem, which is fundamental to all fields treated in this paper. We will also extensively treat Liouville integrability, which ties into the symplectic nature of stable Hamiltonian structures.

In section 3 we will adapt the theory of distributions to establish the framework of foliations. And study these more generally as well in their own right. We will give several constructions of foliations, and several invariants of foliations. Central will be the statement and proof of Tischler's theorem and the concept of taut foliations.

In section 4 we will approach stable Hamiltonian structures. This section will start with in-depth treatment of the theory of contact structures, which are penultimate in arriving to stable Hamiltonian structures. We will then treat stable Hamiltonian structures in generality. This main theorems of this section is the actual stability for both contact and stable Hamiltonian structures.

In section 5 we will make a foray into stable Hamiltonian topology by reducing the dimension to three and analysing some topological properties. The thesis will conclude with an exposition of all-important structure theorem by Cieliebak and Volkov.

Lastly, in section 6 we will provide a very brief introduction into the field of symplectic field theory, which we believe is a very interesting segue into further research.

The impatient reader familiar with most concepts may skip to the second part of section 4, and treat the rest of the paper as a reference work.

2 Distributions

Contact geometry concerns the study of maximally non-integrable distributions and stable Hamiltonian structures are a generalization of that. Therefore, we will do well to develop a solid grasp of what an integrable distribution is before moving on.

A system of linear first-order homogenous partial differential equations may or may not have a solution. One can intuitively image this that at any given point we have a set of vectors pointing in different directions. If a solution of this system exists at that point, that means that there exists a function whose partial derivatives satisfy the requirements set vectors. From solving partial differential equations, one knows that visually such an equation, given initial conditions of a point, will carve out some level set.

In differential geometry, these analytical concepts have found a geometric formalization. Giving a homogenous linear system of partial differential equations, amounts to giving a set of vectors at each point, which will become a so-called distribution. Finding a solution of this system of equations will amount to finding an integral manifold which at each point is spanned by these vectors. If all of this is done is a smooth manner, we can speak of smooth integrable distributions. If such a system of equations is not solvable, then such an integral manifold cannot be found, and the distribution is said to be non-integrable.

Although the theory of finding necessary and sufficient condition for a system of equations to be solvable predates its geometric formalization. It was Frobenius who applied these methods heavily to differential geometry, and hence the geometric realization of the above has been dubbed Frobenius' theorem.

2.1 Distributions and Integrability

In the following we will start by defining the basic concepts of distributions and integral manifolds. We will then keep on adding onto these in order to eventually have proven Frobenius' theorem. Notationally we will almost always assume everything to be smooth unless mentioned otherwise.

Definition 2.1.1. Given a smooth manifold M and denote by TM its tangent bundle. A **distribution** $\mathcal{D}$ is pointwise a collection of vector subspaces

$$\{\mathcal{D}_x \subset T_xM | x \in M\}.$$

Furthermore, around each point of M we can find an open neighbourhood U and a collection of k linearly independent smooth vector fields. Such that we have

$$\mathcal{D}_p = \text{span}(X_1(p), \ldots, X_k(p)).$$

It is called **regular of rank k** if the rank of $\mathcal{D}_x$ is constant k for all $x \in M$.

It follows that a regular distribution is a vector subbundle

$$\mathcal{D} \subset TM.$$

▲

We note that as a vector subbundle of the tangent bundle, the smooth sections of a distribution, denoted from now on by $\Gamma(\mathcal{D})$, are also given by smooth vector fields defined on M. With an added restriction on the amount of directions one can choose.

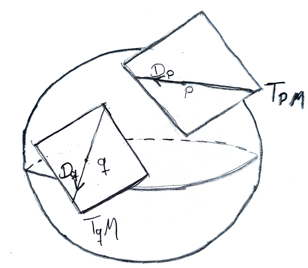

Figure 2: A 1 distribution on a 2-manifold

Now it is a well known fact that the space of all vector fields on M, often denoted by $\mathcal{X}(M)$ form a Lie algebra, meaning the space is closed under the operation of the Lie bracket. We can of course wonder if a similar thing is true for $\mathcal{D} \subset \mathcal{X}(M)$.

Definition 2.1.2. A distribution $\mathcal{D}$ is called **involutive** if for all $X, Y \in \Gamma(\mathcal{D})$ we have that $[X, Y] \in \Gamma(\mathcal{D})$. In this case $\Gamma(\mathcal{D})$ is a Lie subalgebra of $\mathcal{X}(M)$. ▲

If one takes the geometric interpretation of the Lie bracket as the difference vector which arises between when one first flows along X and then Y, or vice versa. Then failure of the Lie bracket to vanish means a failure of the flows to close, and conversely the vanishing of the Lie bracket means the flows commute. Visually it is not hard to imagine that if this difference of flows is not contained in the distribution itself, then it might be impossible to find a nice smooth manifold on which these flow: one could flow out of the manifold. Conversely, if the vector fields of a distribution are indeed a Lie algebra, then there should be some submanifold $N \subset M$ on which these vector fields are defined. This is the core of Frobenius's theorem.

For now, it will motivate to introduce this as a separate concept. In other words: given a distribution $\mathcal{D}$, can there actually be an immersed submanifold on which all the flows of $\mathcal{D}$ are well-defined and commute?

Definition 2.1.3. A regular distribution $\mathcal{D}$ of rank k is called **integrable** if for all $x \in M$ there exists an immersed smooth k-submanifold N such that

$$T_x N = \mathcal{D}_x.$$

In turn, N is called an **integral manifold of** $\mathcal{D}$. It is **maximal** if there does not exist another integral manifold N' of $\mathcal{D}$ such that $N \subset N'$. ▲

If we recall the intuition that giving a distribution $\mathcal{D}$ was equivalent to giving a system of partial differential equations, then by uniqueness of a solution it must follow that two different maximal integral manifolds N, N' of a distribution $\mathcal{D}$ must necessarily have $N \cap N' = \varnothing$. Hence, maximal integral manifolds through a point are unique. It follows directly that an integrable distribution thus gives rise to a collection of maximal integral manifolds which do not intersect each other and together cover M. This is a special class of submanifolds called foliations which we will treat later on.

Example 2.1.3.1. Any non-vanishing vector field X on M defines a rank 1 distribution by setting $\mathcal{D}_x = \text{span}(X_x)$. This is because any vector field commutes with itself. An integral manifold γ of this distribution is precisely given by the more familiar notion of an integral curve of a vector field. Each point in M lies on a unique maximal integral curve. ⧫

Example 2.1.3.2. The following example hints at a deeper connection of differential forms and distributions. Let $f \in C^\infty(\mathbb{R}^n)$, then recall that its differential was given by

$$df = \sum_{i=1}^{n} \frac{\partial f}{\partial x_i} dx_i.$$

We see that finding a vector field in $\ker(df)$ amounts solving a single equation with n variables. Hence, $\ker(df)$ is $(n-1)$-dimensional. If we define the distribution as span by solutions

$$\mathcal{D}_x := \text{ span}\{X \mid df_x(X) = 0\},$$

then we have an integrable distribution whose integral manifolds are given by level sets as

$$L_X f = i_X df = 0.$$

⧫

Example 2.1.3.3. This example is central in contact geometry, meaning maximally non-integrable distributions. Let $M = \mathbb{R}^3$ with its usual (x, y, z) coordinates. Define $\mathcal{D}$ to be the pointwise distribution

$$\mathcal{D}_x = \text{span}(\partial_x + y\partial_z, \partial_y).$$

It is clear this distribution is regular of rank 2. Visually it is non-integrable. Note that $\mathcal{D}$ is invariant when moving along the x-axis. But when moving along the y-axis the planes start tilting. Moving from the origin a little along the y-axis would tilt the manifold. Then moving in the x-direction one would continuously gain height in the z-coordinate. However, moving in the x-direction from the origin, one would remain flat in the plane. For integrability to hold, there must be a manifold connecting this point smoothly with the one described before which lies higher. This cannot possibly exist. ⧫

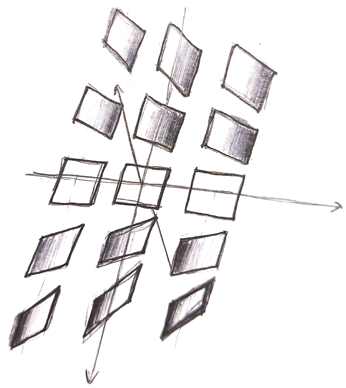

Figure 3: A non-integrable distribution

2.2 Distributions via Differential Forms

Continuing the narrative that integrable distributions reflect in some sense a solution of a system of partial differential equations, then it comes at no surprise differential forms are closely related to these. Indeed, the differential of a function, as it is so aptly named, captures its differential properties. Furthermore, as distributions are subbundles of the tangent bundle, we possess a lot of tools from differential geometry to generalize the concept of the differential of a function, to higher order differential forms.

In the following section we will carefully dissect how distributions can be constructed from differential forms, and also what role they play in the subsequent distribution. The connection between these two will form the foundation of Frobenius' theorem and motivate contact and stable Hamiltonian structures.

We assume the reader to be familiar with the concept of differential forms, however, we will for completeness' sake recall the following algebraic structure on differential forms to aid us in terminology later.

Definition 2.2.1. The local anti-commutative graded **algebra of smooth differential forms** on M over the ring $C^\infty(M)$ is defined as

$$\Omega^\bullet(M) := \bigoplus_{i=0}^{n} \Omega^i(M),$$

where multiplication is given by the wedge product. ▲

As a note of subtlety: the word local in the preamble refers to the fact that differential forms are often not defined globally, but locally. Choosing an atlas of M with opens U_α, we technically only have the anti-commutative graded algebra structure over local differential forms $\Omega^\bullet(U_\alpha)$. However, as the atlas with the transition maps cover M, it makes sense to talk about $\Omega^\bullet(M)$. We will omit writing down the local part in notation, but please remember it is implicit.

With this algebraic structure, we can use algebraic terminology. In particular, we will be interested in different types of ideals. We will see that it are these ideals which are closely related to the integrability of distributions. Let us first make good on our promise to show how to construct smooth distributions using differential forms.

Lemma 2.2.2. *A regular rank k distribution $\mathcal{D}$ is smooth if and only if for each point in M we have an open neighbourhood U, and a collection of smooth 1-forms $\omega_1, \ldots, \omega_{n-k}$ defined on U, such that for each $p \in U$ we have that*

$$\mathcal{D}_p = \bigcap_{i=1}^{n-k} \ker((\omega_i)_p).$$

Proof. Assume such a collection of 1-forms with the above property exists. By assumption

$$\dim\left(\bigcap_{i=1}^{n-k} \ker((\omega_i)_p)\right) = k,$$

so by dimensionality it must follow each ω_i is linearly independent at all $p \in U$. Extend it to a local co-frame of $T^*_U M$, denoted by $(\omega_1, \ldots, \omega_n)$. Where the additional ω_j for $j > (n-k)$ are constructed. This gives rise to a dual frame $(\alpha_1, \ldots, \alpha_n)$ of $T_U M$. Because $\mathcal{D}$ was defined by the intersection of the kernels of the first $(n-k)$ 1-forms, it follows that locally $\mathcal{D}$ is spanned by the last k vectors of this frame $(\alpha_{n-k+1}, \ldots, \alpha_n)$. Thus, as for each x this distribution locally admits a local smooth frame, it is a smooth distribution.

The reverse implication follows by reversing the steps under the assumption that given $\mathcal{D}$ is smooth, there is a smoothly varying local frame of k vector fields. ∎

Definition 2.2.3. We call each collection of $(n-k)$ smooth 1-forms defining a regular rank k distribution $\mathcal{D}$ via

$$\mathcal{D}_p = \bigcap_{i=1}^{n-k} \ker((\omega_i)_p),$$

the **local defining forms for** $\mathcal{D}$. ▲

In the language of the algebraic structure we have on $\Omega^\bullet(M)$, it follows rather directly that a submodule of rank $(n-k)$ is by definition generated by a $(n-k)$-collection of 1-forms which are linearly independent. Lemma 2.2.2 has shown their common kernel defines the rank k distribution. This proves the following lemma.

Lemma 2.2.4. *Locally a regular rank k distribution $\mathcal{D}$ is equivalent to giving a module rank $(n-k)$ submodule $F \subset \Omega^1(U)$.*

As a point of technicality for the above we remark that $\mathcal{D}$ is defined pointwise, hence it is not necessarily individual 1-forms we are interested in, but rather stalks F_x at $x \in M$ of germs of differential form $[\eta]$, which are classes of differential forms which agree in some neighbourhood of x. However, all the proofs given work for such an entire class. Thus, we write η for any representative of the class $[\eta] \in F_x$.

We will rephrase the property of being involutive in terms of differential forms. We first define a special class of forms.

Definition 2.2.5. A p-form η is said to **annihilate** $\mathcal{D}$ if

$$\eta(X_1, \dots, X_p) = 0$$

whenever each $X_i \in \Gamma(\mathcal{D})$. We denote by $I(\mathcal{D}|_{\mathcal{U}}) \subset \Omega^p(U)$ the set of all such locally annihilating forms, and $I(\mathcal{D})$ its extension to $\Omega^p(M)$. ▲

The usual visualization of a p-form is by imagining it via $\ker(\eta)$, which at a point look like a codimension p plane of M. If such a form annihilates $\mathcal{D}$ then this intuitively agrees with the fact that $\ker(\eta)$ is tangent to $\mathcal{D}$.

Recall that there was an algebraic structure on the space of differential forms. The choice of I in the notation is no coincidence, it requires little effort to remark that the set of such locally annihilating forms constitute an ideal in $\Omega^\bullet(M)$. In fact, such an ideal is always locally generated by such forms.

Lemma 2.2.6. *The ideal $I(\mathcal{D})$ is locally generated by a collection of locally defining forms of $\mathcal{D}$.*

Proof. The statement boils down to showing that any annihilating p-form η can be written as

$$\eta = \sum_{i=1}^{n-k} \omega_i \wedge \beta_i,$$

where the ω_i are taken to be locally defining forms and β_i are arbitrary $(p-1)$ forms. So assume η locally annihilates $\mathcal{D}$, and $\mathcal{D}$ has locally defining forms ω_i.

As done earlier, extend the local defining frame to a full local co-frame of T^*M, with corresponding local frame $(\alpha_1, \dots, \alpha_n)$ of TM. It follows that locally we can write

$$\eta = \sum_I \eta_I(\omega_{i_1} \wedge \cdots \wedge \omega_{i_p}),$$

where I is some multi-index $(i_1, \dots, i_p)$ with $i_a < i_b$ if $a < b$, and $\eta_I := \eta(\alpha_{i_1}, \dots, \alpha_{i_p})$, which is some constant. Now by hypothesis this η annihilates $\mathcal{D}$, which is only the case if $\eta_I = 0$ whenever each $i_m \in I$ is in the interval$[(n-k+1), n]$. Thus, the sum can be simplified and rewritten to

$$\eta = \sum_{I: i_1 \le n-k} \eta_I(\omega_{i_1} \wedge \cdots \wedge \omega_{i_p}) = \sum_{i_1=1}^{n-k} \omega_{i_1} \wedge \sum_{I'} \eta_{I'}(\omega_{i_2} \wedge \cdots \wedge \omega_{i_p})$$

where I' is the multi-index where we take the first entry to be fixed by i_1. Now setting

$$\beta_{i_1} := \sum_{I'} \eta_{I'}(\omega_{i_2} \wedge \cdots \wedge \omega_{i_p})$$

proves the statement. ∎

We have shown that the ideal of annihilators of a distribution is closed under the wedge product. However, on this graded algebra we have the important operation of exterior differentiation:

$$d : \Omega^m(M) \to \Omega^{m+1}(M).$$

It is natural to wonder; if η annihilates $\mathcal{D}$, does $d\eta$ also annihilate $\mathcal{D}$? Meaning $I(\mathcal{D})$ would be closed under differentiation. We can rephrase this with the following definition.

Definition 2.2.7. An ideal $I \subset \Omega^\bullet(M)$ is called a **differential ideal** if whenever $\eta \in I$ then $d\eta \in I$. ▲

The question then becomes: is the annihilator ideal $I(\mathcal{D})$ also a differential ideal? Note that for any $\eta \in I(\mathcal{D})$ and $X, Y \in \Gamma(\mathcal{D})$ we get

$$d\eta(X, Y) = -\eta([X, Y]).$$

So being a differential ideal corresponds very closely to $\Gamma(\mathcal{D})$ being involutive.

Lemma 2.2.8. *$\mathcal{D}$ is involutive if and only if $I(\mathcal{D})$ is a differential ideal.*

Proof. Assume $\mathcal{D}$ is involutive. Then for X_i smooth vector fields of $\mathcal{D}$, and an annihilating p-form η, it follows that;

$$d\eta(X_0, \ldots, X_p) = \sum_{i=0}^{p}(-1)^i X_i(\eta(\hat{X}_i)) + \sum_{i<j}^{p}(-1)^{i+j}\eta([X_i, X_j], (\hat{X}_i, \hat{X}_j)) = 0,$$

where $(\hat{X}_j)$ denotes the p-tuple of X_i where X_j is removed. It is immediate that $d\eta$ hence also annihilates $\mathcal{D}$.

Conversely, choose local defining forms ω_i for $\mathcal{D}$. By lemma 2.2.6 we can write any annihilating 1-form as $\eta = \sum_{i=1}^{n-k} f_i \omega_i$. Remark our β_i are now 0-forms, hence smooth functions. Now some algebra shows that for two vector fields X, Y of $\mathcal{D}$ we have;

$$\begin{aligned} d\eta(X, Y) &= \sum_{i=1}^{n-k} df_i(X) \wedge \omega_i(Y) + \sum_{i=1}^{n-k} f_i d\omega_i(X, Y) \\ &= \sum_{i=1}^{n-k} f_i \omega_i([X, Y]). \end{aligned}$$

Now we can choose f_i such that each sum is positive, whence the assumption $d\eta = 0$ implies $\omega_i([X,Y]) = 0$ for each local defining form. Thus,

$$[X,Y] \in \bigcap_i^{n-k} \ker(\omega_i) = \mathcal{D},$$

proving the distribution is involutive. ■

Example 2.2.8.1. Let M be a smooth manifold, and $f \in C^\infty(M)$ such that df never vanishes. It follows readily that $df \in I(\ker(df))$. Furthermore, $d^2 f = 0 \in I(\ker(df))$. We conclude that any smooth non-constant function without extrema defines an involutive codimension one distribution. ♦

2.3 Frobenius' Integrability

In the above discussion we have continuously given definitions and descriptions of involutivity and integrability in tandem, whilst simultaneously giving examples and constructions in which the one implies the other. In fact, these two notions are equivalent; this is Frobenius' integrability theorem which is a major theorem in differential geometry.

Theorem 2.3.1 (Frobenius)**.**
A distribution is integrable if and only if it is involutive.

Most of the groundwork and the proof has already been given, so let us finalize this discussion. However, the last few steps are not fully trivial, and the proof of Frobenius' theorem is done via a concept which at first sight might seem stronger that just integrability.

Definition 2.3.2. A rank k distribution $\mathcal{D}$ is called **completely integrable**, if there exists an atlas of M consisting of charts (U, ϕ) such that for all $p \in U$ we have

$$\mathcal{D}_p = \text{span}(\partial_1(p), \ldots, \partial_k(p)).$$

In other words: the distribution is spanned by the first k coordinate vector fields. The coordinate chart (U, ϕ) is said to be **flat for** $\mathcal{D}$. ▲

Observe that in such a chart, each level set where we take the last $(n-k)$ coordinates to be constant forms an integral manifold. One can visualize these integral manifolds to lie flat with respect to a collection of height-coordinates determined by the last $(n-k)$ entries.

By definition of an atlas, all these charts cover M, and so for each $p \in M$ we find an integral manifold. The following lemma captures the first half of Frobenius' integrability theorem.

Lemma 2.3.3. *For a regular rank k distribution $\mathcal{D}$ complete integrability implies integrability, which implies involutivity.*

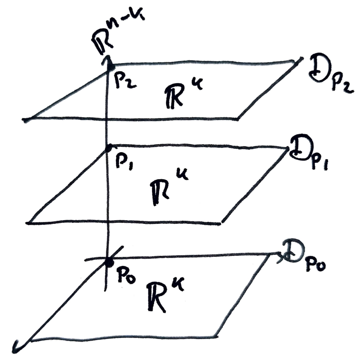

Figure 4: Flat Coordinate chart for $\mathcal{D}$

Proof. The first implication is trivial.

So assume $\mathcal{D}$ is integrable, let $X, Y \in \Gamma(\mathcal{D})$ be defined on some neighbourhood of x. Because $\mathcal{D}$ is integrable, it follows X, Y are both vector fields on some sub-manifold $N \subset M$ containing x. It follows that $[X, Y]$ is also a vector field of N, thus it is involutive. ∎

Now the crux of Frobenius' integrability theorem is that seemingly the weakest property of being involutive, implies seemingly the strongest property of being completely integrable.

Lemma 2.3.4. *If a smooth regular k-rank distribution $\mathcal{D}$ is involutive then it is completely integrable.*

Proof. Assume $\mathcal{D}$ is involutive, as complete integrability is a local property, we will prove such locally.

Choose a local frame $\{X_1, \ldots, X_k\}$ for $\mathcal{D}$. We can choose a coordinate chart (U, ϕ) centred around x, such that in local coordinates $X_i|_x = \partial_i|_x$. Hence, $\mathcal{D}_x$ is locally the complement to

$$\text{span}(\partial_{k+1}|_x, \ldots, \partial_n|_x) \subset T_x\mathbb{R}^n.$$

Now define $\pi : \mathbb{R}^n \to \mathbb{R}^k$ as the projection of the first k coordinates. It follows the differential $d\pi : \mathbb{R}^n \to \mathbb{R}^k$ maps

$$d\pi(\sum_{i=1}^{n} v_i \partial_i|_x) = \sum_{i=1}^{k} v_i \partial_i|_{\pi(x)}.$$

Note that $\mathcal{D}_x$ is the complement to $\ker(d\pi|_x)$. By property of linear maps it follows that the restriction $(d\pi|_x)|_{\mathcal{D}_x}$ is a bijection. And from continuity it follows this is true for the entirety of U.

Using this bijection we define a new local frame for $\mathcal{D}_U$ by setting for all $p \in U$ the following

$$V_i|_p := (d\pi|_{\mathcal{D}_U})^{-1}(\partial_i|_{\pi(x)}).$$

As the Lie bracket is natural with respect to differentials and maps, it follows that

$$d\pi_p([V_i, V_j]_p) = [\partial_i|_{\pi(p)}, \partial_j|_{\pi(p)}] = 0,$$

whenever $i \neq j$. By injectivity of $d\pi_p$ one obtains that $[V_i, V_j]_p$ is a vertical vector field for all p. By assumption the distribution was involutive, and hence $[V_i, V_j]_p \in \mathcal{D}|_p$, for which we had $(d\pi|_p)|_{\mathcal{D}_p}$ was a bijection, particularly it is injective. Thus, $[V_i, V_j] = 0$. We get for each p a local frame of commuting vector fields.

Now given k linearly independent commuting vector fields in a k-dimensional neighbourhood of a point, we can utilize their flows to define a new coordinate chart. These coordinate charts are precisely flat charts for $\mathcal{D}$. Hence, $\mathcal{D}$ is everywhere locally completely integrable locally, hence globally. ∎

Using the lemma together with the one-way implication chain we had already obtained, we have proven Frobenius' theorem.

Example 2.3.4.1. Recall the example of the tilting planes 2.1.3.3. We had deduced this distribution to be non-integrable visually speaking. Now we also see algebraically

$$[\partial_x + y\partial_z, \partial_y] = -\partial_z \notin \mathcal{D}_x.$$

The distribution is not involutive, hence non-integrable. ♦

Although we have proven Frobenius' theorem, we have not yet made good on the promise to link this to differential forms.

Recall that according to lemma 2.2.8, a distribution was involutive if and only if its annihilator ideal was a closed differential ideal. As a corollary we immediately deduce that.

Corollary 2.3.5. *A distribution $\mathcal{D}$ is integrable if and only if $I(\mathcal{D})$ is closed under d.*

However, one can do much better. First, by complete integrability we can choose around any point a flat coordinate chart for $\mathcal{D}$, meaning it is locally spanned by the first k coordinate vector fields ∂_i. We then dually get a smooth local coframe $(dx_i)_{1\leq i\leq n}$. It is readily deduced that

$$\ker(dx_j) = \operatorname{span}(\partial_i \mid i \neq j).$$

From this, one immediately sees

$$\mathcal{D}_x = \bigcap_{i=k+1}^{n} \ker(dx_i).$$

This means the last $(n-k)$ coordinate 1-forms form a local defining frame of $\mathcal{D}$.

Recalling the equivalence given in lemma 2.2.4 between giving a rank $(n-k)$ submodule $F \subset \Omega^1(U)$ and a rank k distribution, it follows that for an $\eta \in F$ we get $\eta = \sum_{i=k+1}^{n} f_i dx_i$. Now note that this rank $(n-k)$ submodule generated by these (dx_i) is exactly $I(\mathcal{D})$.

Corollary 2.3.6. *A rank k distribution $\mathcal{D}$ is integrable if and only if $I(\mathcal{D})$ is generated by $(n-k)$ exact differential forms.*

Putting this all together, we arrive at the following phrasing of Frobenius' theorem, which we will use throughout this paper.

Theorem 2.3.7 (Frobenius)**.**
For a smooth manifold M and a smooth regular rank k distribution $\mathcal{D}$ the following are equivalent.

1. *$\mathcal{D}$ is integrable.*
2. *$I(\mathcal{D})$ is generated by $(n-k)$ exact differential 1-forms.*
3. *$\mathcal{D}$ is involutive.*

This paper will focus heavily on hyperplane distributions, that is to say distributions of codimension one. It follows that in this case the theorem achieves the particularly nice form.

Corollary 2.3.8 (Codimension one Frobenius)**.** *Let α be a non-vanishing 1-form on M. Then the following are equivalent.*

1. *$\xi := \ker(\alpha)$ is an integrable distribution.*
2. *$\alpha \wedge d\alpha = 0$.*

Proof. First assume $\ker(\alpha)$ to be integrable, as a consequence of Frobenius' $\ker(\alpha)$ is involutive, so it follows for $X, Y \in \ker(\alpha)$ that

$$d\alpha(X,Y) = X(\alpha(Y)) + Y(\alpha(X)) - \alpha([X,Y]) = 0.$$

Thus, we see $d\alpha \in I(\ker(\alpha))$. And so it follows by lemma 2.2.2 that $d\alpha = \alpha \wedge \beta$ for some 1-form β, and hence $\alpha \wedge d\alpha = 0$.

Conversely, if $\alpha \wedge d\alpha = 0$, locally we can extend α to a basis $(\alpha, \alpha_2, \ldots, \alpha_n)$ of T^*M. It follows

$$d\alpha = \sum_{i<j} f_{ij} \alpha_i \wedge \alpha_j.$$

As $\alpha \wedge d\alpha = 0$, it must follow that $d\alpha = \sum_{j=2}^{n} f_j \alpha \wedge \alpha_j$. Now defining a new 1-form $\eta := \sum_{j=2}^{n} f_j \alpha_j$, it follows for each $X, Y \in \ker(\alpha)$ that

$$d\alpha(X,Y) = (\alpha \wedge \eta)(X,Y) = \alpha(X)\eta(Y) - \alpha(Y)\eta(X) = 0.$$

As such it follows $\alpha[X,Y] = 0$, and so $[X,Y] \in \ker(\alpha)$. One concludes $\ker(\alpha)$ is involutive, hence is it integrable by Frobenius. ■

Example 2.3.8.1. For any smooth function f for which $df \neq 0$, we have that df is an exact 1-form, hence closed. It follows that the level sets of a submersion always define a codimension 1 integral submanifold. This is in fact similar to the submersion theorem. ♦

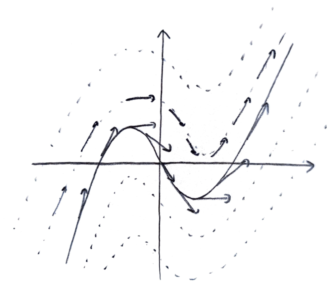

Figure 5: The level sets of f agree with the integrable distribution $\ker(df)$

Example 2.3.8.2. Let $M = S^2$. By dimensionality, it follows that any non-vanishing 1-form on S^2 defines an integrable distribution. In this case the dimension of the ambient manifold is 2, so it defines a nowhere vanishing vector field on M which at each point is spanned by $\ker(\alpha)$. It follows from a well-known fact that such vector fields cannot exist on S^2, that S^2 cannot admit any non-vanishing 1-forms. As a corollary any function $f : S^2 \to \mathbb{R}$ must have extrema. ♦

2.4 Coorientability

There is a property of distributions of hyperplanes, meaning codimension 1 distributions, which we will make use of frequently. Certainly in the realm of contact geometry. It is a dual notion to orientability, aptly named coorientability. To define this let first recall a familiar object.

Definition 2.4.1. Given an immersion $i : N \hookrightarrow M$ then the **normal bundle** to N in M is the quotient bundle $\nu_N := TM/TN$. ▲

If such is the case, we had the short exact sequence of vector bundles

$$0 \to TN \to TM|_N \to \nu_N \to 0.$$

One can think of the normal bundle to N in M the complement of TN in TM; if one chooses a Riemannian metric g on M and combines this with the splitting from the short exact sequence, one obtains

$$TM_N \cong TN \oplus \nu_N \cong TN \oplus TN^{\perp}.$$

Now as integrable distributions ξ have associated to it their integrable submanifold, we can extend this notion by defining the normal bundle to an integrable distribution as $\nu_\xi := TM/\xi$. Now in the specific case that ξ is a distribution of hyperplanes, hence the associated integrable submanifold is of codimension one, we can define our notion of coorientability.

Definition 2.4.2. A codimension one distribution ξ of TM is **coorientable** if ν_ξ is trivial. ▲

As before we can construct a short exact sequence

$$0 \to \xi \to TM \to \left(\nu_\xi \cong M \times \mathbb{R}\right) \to 0.$$

To gain some geometric intuition of coorientability, recall that a bundle was trivial if and only if we could choose a global frame. And by dimensionality this is equivalent to being able to choose a non-vanishing section of the normal bundle. If we choose a metric g on TM, it is equivalent to choosing a non-vanishing section of $\nu_\xi \cong \xi^\perp$. Thus, we obtain a geometric notion of coorientability: it is equivalent to being able to choose a smooth nowhere vanishing vector field X on M which is always perpendicular to our distribution. For this reason coorientability is sometimes also referred to as transverse orientability. In fact, keeping our chosen metric, we obtain the following equivalent statement of coorientability.

Lemma 2.4.3. *A codimension one distribution ξ of TM is coorientable if and only if $\xi = \ker(\alpha)$, where $\alpha \in \Omega^1(M)$ nowhere vanishing.*

Proof. Assume ξ is coorientable, as a result we can choose a non-vanishing vector field always perpendicular to ξ. Define a 1-form $\alpha := g(X, -)$, by definition $g(X,Y) = 0$ for $Y \neq 0$ if $Y \perp X$. As by construction $\xi \perp X$, we see $\xi = \ker(\alpha)$, as a single 1-form uniquely defines a regular codimension one distribution via its kernel.

The other direction is simply a reversion of steps. Given an α such that $\ker(\alpha) = \xi$, we define an orthogonal X to be a unit vector field with $\alpha(X) > 0$. ■

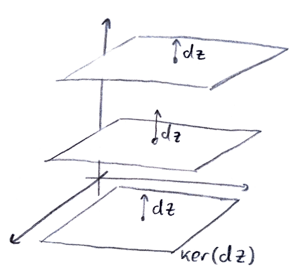

Figure 6: Coorientation is a transverse orientation by a non-vanishing 1-form

The notion of coorientability is of course closely related to orientability, in fact, they are equivalent if the ambient manifold M is orientable.

Lemma 2.4.4. *If M is an orientable manifold, and ξ is a codimension one distribution. Then ξ is coorientable if and only if ξ is orientable.*

Proof. Recall a manifold M is orientable if and only if TM is orientable as a bundle. From the splitting of the exact sequence we get that

$$TM \cong \xi \oplus \nu_\xi.$$

By this isomorphism it follows that if M is orientable, then so is $\xi \oplus \nu_\xi$.

If ξ is coorientable, then the normal bundle is trivial; $\nu_\xi \cong M \times \mathbb{R}$, hence clearly orientable, thus ξ must also be orientable.

Conversely, if ξ is orientable, then ν_ξ is orientable as well. This is a line bundle, and so orientability is equivalent to being trivial, hence ξ is coorientable. ■

2.5 Liouville Integrability

The concept of Frobenius integrability is applicable to any smooth manifold and denotes a deep connection between 1-forms and distributions on general smooth manifolds. However, we will not be studying general manifolds, but those equipped with a stable Hamiltonian structure. As mentioned before, stable Hamiltonian structures are a generalization of contact structures, which are themselves an odd-dimensional counterpart of symplectic structures. On symplectic manifolds there is the notion of Liouville integrability, and by relation is often applicable in a modified form to stable Hamiltonian structures.

We will explain the concept of Liouville integrability, and derive some essential results. For the remainder of this section (M, ω) is a closed symplectic $2m$-dimensional smooth manifold, unless specified differently.

2.5.1 Poisson Brackets

Recall that a symplectic 2-form was closed, and non-degenerate. As a result of its non-degeneracy, there is an induced fibre-wise isomorphism:

$$\begin{aligned} \iota : T_pM &\to T_p^*M \\ v &\mapsto \omega_p(v, -), \end{aligned}$$

sometimes also called the left flat map $\flat_\omega$ of ω. This gives a 1-1 correspondence between 1-forms and vector fields. There is also a straightforward way to construct 1-forms from smooth functions h by taking the differential dh, this gives rise to the following construction.

Definition 2.5.1. Given a symplectic manifold (M, ω) and a smooth function $h : M \to \mathbb{R}$. We define the **Hamiltonian vector field** of h to be the unique vector field X_h satisfying

$$\omega(X_h, -) = dh,$$

at all points of M. ▲

There is a nice geometric interpretation of the Hamiltonian vector field. Recall, in parallel, that if we had chosen a metric g, then the gradient ∇h was the vector field which denoted the greatest rate of change, hence is perpendicular to the level sets of h. It was defined as the unique vector field satisfying $g(\nabla h, -) = dh$. Similarly, for the Hamiltonian vector field we see that

$$L_{X_h} h = i_{X_h} dh = \omega(X_h, X_h) = 0,$$

so we see the Hamiltonian vector field runs parallel to the level sets of h. Note that the integral curves of X_h therefore run tangent to a certain connected component of $h^{-1}(c)$, often called an energy level. An extra calculation shows us that if we have chosen a metric g, then we see that

$$g(\nabla h, X_h) = dh(X_h) = \omega(X_h, X_h) = 0,$$

and so we arrive at the following.

Lemma 2.5.2. *If one chooses a metric g on (M, ω), then $\nabla h \perp X_h$.*

Hamiltonian vector fields also arise naturally from symplectic vector fields, which are defined by the property $L_X \omega = 0$. If X, Y are both symplectic vector fields, we have that

$$\begin{aligned} i_{[X,Y]}\omega &= (L_X i_Y \omega - i_Y L_X \omega) \\ &= d(i_X i_Y \omega), \end{aligned}$$

realizing that $\omega(Y, X)$ defines a smooth function on M, we conclude the following.

Lemma 2.5.3. *If X, Y are symplectic vector fields, then $[X, Y]$ is a Hamiltonian vector field.*

Now we have three facts which together appeal to a generalization of the Lie bracket definition: every Hamiltonian vector field is symplectic, symplectic vector fields are closed under the operation of the Lie bracket, and there is a way to construct Hamiltonian vector fields from smooth functions. So whereas vector fields formed a Lie algebra under the Lie derivative, we have something similar for smooth functions.

Definition 2.5.4. Given two smooth functions on M we define the **Poisson bracket** as

$$\{f, g\} := \omega(X_f, X_g).$$

▲

It is closely related to the Lie bracket as we can see:

$$\omega(X_{\{f,g\}}, -) = d\{f, g\} = i_{[X_g, X_f]}\omega,$$

and thus

$$X_{\{f,g\}} = -\left[X_f, X_g\right]. \tag{2.1}$$

It is easy to check that the Poisson bracket $\{-,-\}$ defines a Lie bracket on the space $C^\infty(M)$ as it is bilinear, antisymmetric, and satisfies the Jacobi identity. However, note that it also satisfies the Leibniz rule;

$$\{fg,-\} = d(fg) = fdg + gdf = f\{g,-\} + g\{f,-\}.$$

Definition 2.5.5. A **Poisson algebra** is a Lie algebra with Lie bracket $\{-,-\}$ which also satisfies the Leibniz rule. We say $\{-,-\}$ is the **Poisson bracket**. ▲

What is the geometric intuition behind this new object? Recall that two vector fields were commuting if $[X,Y] = 0$. It is a general fact that X, Y commute if and only if their flows also commute. If follows directly from (2.1) that if $\{f,g\} = 0$ then $[X_f, X_g] = 0$ and hence if f, g commute in the Poisson bracket sense, their Hamiltonian vector fields also commute. However, the Poisson bracket being zero is in fact stronger than just saying their respective Hamiltonian vector fields commute. From direct computation it follows:

$$L_{X_f} g = -\omega(X_f, X_g) = -L_{X_g} f = 0.$$

Interpreting the Hamiltonian vector fields as being parallel to the level sets of the function, we formulate a geometric interpretation of the Poisson bracket as follows:

Definition 2.5.6. If $\{f,g\} = 0$, then f is constant along the levels of g and vice versa. We also say g **is an integral of** f. ▲

Geometrically we have two vector fields X_f and X_g on M which flow parallel to both a level set of f and a level set of g.

We note the amount integrals up to scaling is limited by the dimension of the manifold M. We can derive this from the geometric image constructed above. Assume the simplest case that M is 2-dimensional and $\{f,g\} = 0$. It follows X_f flows parallel to both a level set of f, and a level set of g. It must follow by dimensionality that in the case X_f and X_g are independent that f is constant on the entirety of M, in which case $df = 0$ and thus non-degeneracy of ω implies $X_f = 0$. Or if X_f is parallel to X_g, then $g = c \cdot f$ for some constant c, in which case the level sets completely coincide. So if M is 2-dimensional, there can only be one integral of f up to scaling, which is f itself. This argument also applies in higher dimensions.

Lemma 2.5.7. *If (M, ω) is $2m$-dimensional, and $df \neq 0$. Then, up to scaling with a constant, there can be at most m integrals of f.*

2.5.2 Hamiltonian Systems

Let us now introduce the concept of a completely integrable Hamiltonian system. Recall that for a smooth function $f : M \to \mathbb{R}$ a point p is regular if df does not vanish at that point, and is of constant rank if df does not vanish anywhere. Let

$U \subset M$ a region containing p on which f is of constant rank, and furthermore let $f(p) = r_1$. Then it follows by the implicit function theorem that, on a possibly smaller open contained in U, a connected component of $f^{-1}(r)$ is a $(2m-1)$-dimensional smooth submanifold.

Now as before let $\{f, g\} = 0$ and let X_f and X_g be linearly independent on this possible smaller open. Let $g(p) = r_2$, and apply the same reasoning as above. Then by dimensionality it must follow for connected components of the level sets

$$\dim\left(f^{-1}(r_1) \cap g^{-1}(r_2)\right) = 2m - 1 + 2m - 1 - 4m = 2m - 2. \tag{2.2}$$

In fact this construction continues for any n-tuple of Poisson commuting smooth functions. If we take the maximally allowed number unique integrals m we have the following definition.

Definition 2.5.8. Let (M, ω) be a $2m$-dimensional symplectic manifold, and let $U \subset M$ an open neighbourhood. Furthermore, let $h_i : M \to \mathbb{R}$ be m smooth functions and denote by $X_i := X_{h_i}$.

If each X_i is linearly independent for all $p \in U$, and if all h_i Poisson commute, then $h := (h_1, \ldots, h_m)$ is called the **moment map** of a **completely integrable Hamiltonian system** (U, ω, h), and U is called an **integrable region**. We can often take $U = M$, and so we obtain (M, ω, h). ▲

Now $h : M \to \mathbb{R}^m$ is a smooth function, moreover by the demand that each X_i is linearly independent at a point p, and so in particular non-zero, we get that each $(dh_i)_p$ is non-zero, and hence $h(p) = (r_1, \ldots, r_m)$ is a regular value. Again as in (2.2) by implicit function theorem we obtain for a connected component T_r of a level set that

$$\dim(T_r) = \dim(h^{-1}(r)) = m,$$

and so each T_r is a closed m-dimensional smooth submanifold. This result also follows from Frobenius' integrability theorem. As Poisson commuting functions imply commuting Hamiltonian vector fields, the distribution defined by the span of the Hamiltonian vector fields has an integral manifold of dimension m.

We can in fact say more about T_r than just its dimension.

Lemma 2.5.9. *Let (M, ω, h) be a completely integrable Hamiltonian system. Then each T_r is diffeomorphic to $\mathbb{T}^m$.*

Proof. By hypothesis M is a compact manifold, it follows that $T_r := h^{-1}(r)$ is a compact m-dimensional smooth submanifold. Now pick a base-point $p \in T_r$, note that for each h_i there is a unique integral curve $\gamma_{i,p}(t) : I \to T_r$ passing through p, and note $\gamma_{i,p}(I) \subset T_r$. If all these integral curves are closed, then we are done, as the coordinate chart can be given by the flows. Unfortunately, this need not be the case.

However, denote by ϕ_i the flow of the Hamiltonian vector field X_i. As by assumption the flows commute, and they flow within T_r, we have a well-defined action of $\mathbb{R}^m$ on T_r given by composing flows

$$\Phi : \mathbb{R}^m \times T_r \to T_r,$$
$$(t_1, \dots, t_m, p) \mapsto \phi_m^{t_m} \circ \dots \circ \phi_1^{t_1}(p).$$

This action is well-defined as by assumption the order of composition does not matter. Let us denote Φ^t for the action of $t \in \mathbb{R}^m$ on M. Now if we keep p fixed, then we have a well-defined continuous surjective smooth map

$$\Phi_p : \mathbb{R}^m \to T_r,$$
$$t \mapsto \Phi^t(p).$$

Note Φ_p is continuous and surjective, and T_r is compact without boundary. If Φ_p was also injective, the preimage $\Phi_p^{-1}(T_r)$ would be a closed and bounded interval in $\mathbb{R}^m$. It follows that Φ_p cannot be injective, hence Φ_p has fixed points. The specific times which leave p fixed, form a subgroup, denote it by

$$\Gamma_p := \{\tau \in \mathbb{R}^m \mid \Phi_p(\tau) = p\} < (\mathbb{R}^m, +).$$

Recall that all Hamiltonian vector fields were linearly independent, in particular they are non-vanishing, and so locally Φ_p is a diffeomorphism onto its image. The consequence is that locally around each $\tau \in \Gamma_p$ we have a neighbourhood U such that $U \cap \Gamma_p = \{\tau\}$. So in fact Γ_p is a discrete subgroup.

We will show Γ_p actually defines an integral lattice in $\mathbb{R}^m$. Meaning there are $\{\tau_1, \dots, \tau_n\} \in \Gamma_p$ with $n \leq m$, such that for each $\tau \in \Gamma_p$ we can write

$$\tau = \sum_{i=1}^{n} z_i \tau_i,$$

with $z_i \in \mathbb{Z}$. Observe the free $\mathbb{R}$ vector space $\langle \Gamma_p \rangle$. This is a linear subspace of $\mathbb{R}^m$, and so at most the dimension of this is m. Hence, pick a basis $B := \{\tau_1, \dots, \tau_n\}$ of $\langle \Gamma_p \rangle$ with $n \leq m$. Now look at the free Abelian group Γ_p^0 generated by B. The index $[\Gamma_p : \Gamma_p^0]$ denotes the number of left cosets of Γ_p^0 in Γ_p, we claim it is finite. Indeed, let τ_j represent the coset $\tau_j \Gamma_p^0$. Define some sort of fundamental domain $B_0 := [0, 1)B$. As this fundamental domain tiles $\langle \Gamma_p \rangle$ it follows that there exists a $\beta_j \in B_0$ and a $\tau_{0j} \in \Gamma_p^0$ defining

$$\tau_j = \beta_j + \tau_{0j}.$$

As by assumption each representative τ_j is discrete, it follows each corresponding β_j is discrete in B_0. Because B_0 is bounded, it follows the collection of β_j, hence the amount of representatives τ_j must be finite. Thus, we can define a finite integer $k := [\Gamma_p : \Gamma_p^0]$. Now it follows $k\Gamma_p = \Gamma_p^0$, and so $\Gamma_p = \frac{1}{k}\Gamma_p^0$. As Γ_p^0 was free Abelian, it follows Γ_p has an integral basis $\frac{1}{k}B$.

Let $\theta := \{\theta_1, \ldots, \theta_n\}$ now denote this integral basis of $\Gamma_p \subset \mathbb{R}^n$. The integral domain defines a torus, so we get a diffeomorphism

$$\mathbb{T}^n \cong \mathbb{R}^n/\theta.$$

Moreover, we remark that Φ_p descends to a map $\overline{\Phi_p} : \mathbb{T}^n \times \mathbb{R}^{m-n} \to T_r$, indeed one can see

$$\Phi_p(a\theta_i + b\theta_j) = \Phi_p(b\theta_j) \circ \Phi_p(a\theta_i) = \Phi_p([b]\theta_j) \circ \Phi([a]\theta_i) = \overline{\Phi_p}([a]\theta_i + [b]\theta_j).$$

But note by definition Γ_p was defined as those times which left p fixed, so in particular $\overline{\Phi_p}$ is a smooth bijection. By compactness of T_r, it must follow $n = m$, and hence we have $\mathbb{T}^m \cong T_r$.

In fact, it will be useful to remark that we have constructed so-called angle coordinates on T_r given by

$$\begin{aligned} T_r &\to \mathbb{R}^m \\ q &\mapsto \Phi_p^{-1}(q), \end{aligned}$$

denote these by $\theta_1, \ldots, \theta_m$. ■

Note these angle coordinates are far from unique, any other choice of base point p defines different angle-coordinates. Furthermore, the angle coordinates are not constant through the t action of Φ. Indeed, if one has flowed already for a certain time, the remaining flow time to reach the point has reduced. However, keeping this in mind, they do satisfy the following

$$\frac{d\theta_i}{dt} = \frac{d}{dt}(\Phi(t,p))_i \implies \theta_i(t) = \theta_i(0) - t_i.$$

Moreover, as $T_p(T_r)$ is span by $(X_i)_p$, it follows that $\omega_p = 0$ for all $p \in T_r$, so we additionally conclude the following.

Lemma 2.5.10. *Each T_r is a Lagrangian submanifold of (M, ω).*

This notion of the level sets being Lagrangian m-tori is a very rigid concept of integrability, called **Liouville integrability**. The Lagrangian tori are called **Liouville tori**. It is obvious that Liouville integrability implies Frobenius' integrability, but not the other way around.

Note that if $U \subset M$ is a completely integrable system, then for a point $p \in U$ its coordinates can be expressed by an m-tuple $(p_1, \ldots, p_m)$ determining on which specific connected component of a level set $h^{-1}(r)$ it lies, and then an m-tuple $(\theta_1, \ldots, \theta_m)$ describing its position within this level-set. Owing to the usability of this framework in mechanics, if we interpret r as the energy level, and keeping in mind that canonically $\mathbb{T}^m \cong S^1 \times \cdots \times S^1$, we call these **action-angle coordinates**. It turns out that completely integrable systems have a very nice and homogeneous structure, that is to say: there is a coordinate change such that the Hamiltonian vector fields are invariant under the torus action on each Liouville torus, this is the major Arnold-Liouville theorem.

Theorem 2.5.11 ([Arn78, Arnold-Liouville]).
Given a completely integrable Hamiltonian system (M, ω, h). For a level set $T_r := h^{-1}(r)$, there exists an open neighbourhood U of T_r and an open neighbourhood D of $0 \in \mathbb{R}^m$. Such that there is a symplectomorphism

$$\psi : (U, \omega|_U) \to (\mathbb{T}^m \times D, \omega_{\text{std}}),$$

with

$$\psi(T_r) \cong \mathbb{T}^m \times \{0\}.$$

Moreover, if $(\theta, p) \in \mathbb{T}^m \times D$, then h depends only on its action coordinates:

$$\frac{d}{d\theta} h \circ \psi^{-1}(\theta, p) = 0.$$

We will not give much more of the proof than already given, as it is quite tedious. A full proof can be found in [Arn78]. The following sketches what steps still have to be done, and what the neighbourhoods are.

We have already proven each T_r is diffeomorphic to $\mathbb{T}^m$. Consequently, define $D_r \subset \mathbb{R}^m$ to be an open neighbourhood of r, then we define the open neighbourhood U of T_r as $h^{-1}(U)$, as h was smooth, it is moreover continuous, hence this preimage is open as well. We have that for each $t \in U$, $h^{-1}(t) \cong \mathbb{T}^m$. It is now clear that $U \cong \mathbb{T}^m \times D_r$, composing with a translation of D_r gives us the open neighbourhood D of the theorem.

To prove the diffeomorphism can be adapted to a symplectomorphism, one starts by using an adaptation of Darboux's theorem. This states that locally every symplectic manifold is symplectomorphic. As a result each point $p \in M$ has an open neighbourhood W which is symplectomorphic to an open neighbourhood V of 0 in $(\mathbb{R}^{2m}, \omega_0)$ with standard symplectic form. Moreover, given a Hamiltonian h we can construct such a symplectomorphism σ to satisfy

$$h \circ \sigma^{-1}(x, y) = y.$$

We can apply a symplectic change of coordinates to the already found coordinates (θ, h) on U to become (θ, I). These are the so-called angle-action coordinates. It is with respect to these coordinates that

$$(\psi^{-1})^* \omega = \sum_{i=1}^{m} dI_i \wedge d\theta_i.$$

By the relative Poincaré lemma, there is a possibly smaller tubular neighbourhood $\bar{U}$ of T_r such that ω is exact so let $d\lambda = \omega_{\bar{U}}$. Let γ_i form a basis of $H_1(T_r)$. Then define the action variables by calculating

$$p_i : T_r \to \mathbb{R},$$
$$p_i(p) = \oint_{\gamma_i} \lambda.$$

In fact, these are invariant of the point $p \in T_r$ chosen.

3 Foliations

Recall that it was earlier postulated that if a regular rank k distribution $\mathcal{D}$ was everywhere defined and integrable, then two maximal integrable manifolds of $\mathcal{D}$ could not intersect or be tangent to each other. Thus, each point $p \in M$ lies on a unique submanifold L of dimension k. The geometric picture is then that all these k dimensional manifolds together cover M in much the same way gills line the underside of a mushroom cap. This notion is formalized in the concept of a foliation of a manifold, which we will define in this section and explore some initial useful properties.

3.1 Foliations and Distributions

Definition 3.1.1. A **foliation of dimension** k **on** M is a collection of disjoint, connected, non-empty, immersed k-dimensional submanifolds $L \subset M$ which together cover M. We denote the foliation by $\mathcal{F}$. Every individual $L \subset \mathcal{F}$ is called a **leaf**. ▲

As each point lies on a unique leaf, we denote by L_x the unique leaf passing through a point x, remark that there are multiple points denoting the same leaf.

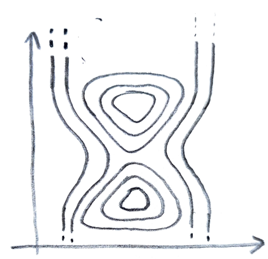

Figure 7: A foliation partitions, but leaves are not necessarily diffeomorphic, nor embedded

It is clear that the concept of foliations readily ties in into the theory of distributions we have discussed earlier. Any smooth foliation $\mathcal{F}$ defines a distribution $\mathcal{D}$ by straightforwardly setting $\mathcal{D}_x := T_x L$, and conversely any integrable distribution $\mathcal{D}$ defines a foliation by defining $\mathcal{F}$ to be the union of maximally integrable submanifolds.

Lemma 3.1.2. *A smooth distribution $\mathcal{D}$ is integrable if and only if each $\mathcal{D}_x = T_x L$ for some leaf L of a foliation $\mathcal{F}$.*

In fact, we can readily use this to our advantage already by describing the following nice looking local picture of a foliation.

Lemma 3.1.3. *Locally for a $1 \leq k < n$ dimensional foliation $\mathcal{F}$ of an n-manifold M, each leaf $L \subset \mathcal{F}$ is an embedded $\mathbb{R}^k$ plane in $\mathbb{R}^n$ where the last $(n-k)$ are taken to be constant.*

Proof. This follows readily from the fact that each leaf of the foliation is an integrable manifold, and integrable was equivalent to being completely integrable. Hence, for each point x there is a chart (U, ϕ) such that

$$\mathcal{D}_x = T_x L = \text{span}(\partial_1(x), \dots, \partial_k(x)).$$

It follows the coordinates $x_{k+1}, \dots, x_n$ must be constant. ∎

The existence of such charts warrant a different way to prove a dissection of M into k-dimensional submanifolds defines a foliation. If there exists an atlas of M such that in each chart the submanifolds look as described above, then these submanifolds define a foliation $\mathcal{F}$ of M.

Definition 3.1.4. Let (U, ϕ) a chart of M, endowed with ak-dimensional foliation $\mathcal{F}$. It is called a **foliated chart** if $\phi : U \to D^k \times D^{n-k} \subset \mathbb{R}^k \times \mathbb{R}^{n-k}$ such that

$$\phi(L_p \cap U) \cong D^k \times \{\phi(x)\}.$$

An atlas consisting of foliated charts is called a **foliated atlas**. See figure 4. ▲

Observing foliations as a result from the existence of a foliated atlas, also allows us to readily introduce the concept of local foliation. In this case we might not find a foliated atlas for the entirety of M, but we may find opens $U \subset M$ which might be foliated, and even some partial covering of M with foliated charts.

Recall that any non-vanishing vector field X defined a 1-dimensional foliation of M via its integral curves. Slightly weaker is that any locally non-vanishing vector field X at least locally defines a 1-dimensional foliation if its flow is locally defined. As a consequence of the existence of the local flow theorem, this means that for a point $p \in M$ where $X_p \neq 0$ there is a small neighbourhood U' of p in which the flow is defined. We conclude there is a foliated chart (U, ϕ) of a possible even smaller neighbourhood of p, we can of course centre this chart around p such that $\phi(p) = 0$. In this chart we have that the integral curves, leaves, look like straight lines keeping the last $m-1$ coordinates constant. In fact, this is precisely the statement of the flow box theorem, which says we can always choose a coordinate chart which straightens the flow of a vector field.

Lemma 3.1.5 (Flow box)**.** *Let p be a point in M, and X a vector field such that $X_p \neq 0$, then there exists a chart (U, ϕ) on M such that*

$$\phi(p) = 0,$$

and moreover

$$d\phi(X) = e_1.$$

3.2 Fibrations and Submersions

Foliation also naturally arise in the context of submersions, the following is the well-known submersion theorem.

Lemma 3.2.1. *Let M an m-manifold, N an n-manifold, and $f : M \to N$ be a smooth surjective submersion. Then the connected components of the level sets $L_q := f^{-1}(q)$ together form a codimension n foliation of M.*

Proof. Denote $p := f^{-1}(q)$, then by surjectivity of df_p we have a short exact sequence

$$0 \to \ker(df_p) \to T_pM \to T_qN \to 0,$$

it follows $\dim\big(\ker(df_p)\big) = m - n$. Furthermore, df is an exact 1-form, so readily $df \wedge d^2 f = 0$ and so by Frobenius' 2.3.8 it follows $\ker(df)$ is integrable. As $\ker(df)$ is locally precisely span by those directions in which f does not vary, it coincides with the level sets. We conclude each L_q is a $(m-n)$ dimensional leaf of the foliation by level sets $\ker(df)$. ∎

It is tempting to think that in the above case this defines some fibre bundle of some sort, where M is the total space, N the base space, the fibres are given by the level sets L_q, and f functions as the projection to create the fibre bundle (M, N, f, L). This is not far from the truth, be we do need some additional requirements on f for this to hold. This is in fact a theorem by Ehresmann.

Lemma 3.2.2. *If $f : M \to N$ is a proper, surjective, smooth submersion, then (M, N, f, L) defines a fibre bundle.*

Proof. Because fibres are defined locally, we will assume we are working in a local trivialization of N centred around q, and hence use $\mathbb{R}^n$ and $q = 0$. Because f is a surjective smooth submersion, we can find vector fields X_i for $1 \leq i \leq n$ on M such that

$$(df)_p(X_i) = (\partial_i)_q,$$

where ∂_i denote the usual coordinate vector fields on $\mathbb{R}^n$, we say these vector fields are f-related. Denote by ϕ_X^t the flow of a vector field X. A property of f-related vector fields is that the flow commutes with the map f, that is to say

$$f \circ \phi^t_{X_i} = \phi^t_{\partial_i} \circ f.$$

By properness of f it follows that the flow of X_i is defined on an interval I as long as the flow of ∂_i is defined on that same interval I. Because the flows of the coordinate vector fields are clearly complete, so are the flows of X_i.

Now we will use these complete flows of X_i to define a diffeomorphism:

$$\Phi : f^{-1}(\mathbb{R}^n) \to \mathbb{R}^n \times f^{-1}(0).$$

As we are using the coordinate vector fields on $\mathbb{R}^n$, we can say for a point $r \in \mathbb{R}^n$ that its coordinates are given by $(t_1, \ldots, t_n)$, where we interpret these as the t_i flow needed from 0 via the coordinate vector field ∂_i to reach r. This is a well-defined operation as coordinate vector fields commute. The diffeomorphism is then given by reversing the flows of the f-related vector fields

$$f^{-1}(r) \ni s \mapsto (r, \phi^{-t_n}_{X_n}, \ldots, \phi^{-t_1}_{X_1}(s)).$$

■

In fact, fibre bundles are more rigid than foliations. With foliations the requirements mostly only concern the dimensionality and submanifold structure of each leaf. As a result each leaf can look vastly different from one another, as in figure 7. Whereas a fibre bundle has the additional demand, amongst others, that the fibres are diffeomorphic to each other. It is clear to see that fibre bundles define foliations in which the leaves are diffeomorphic to each other.

Lemma 3.2.3. *Each fibre bundle* (E, B, π, F) *defines a foliation* $\mathcal{F}$ *by setting* $L_x := F_{\pi(x)}$.

A particular type of foliation one often encounters is the suspension foliation. In fact, these are closely related to the Poincaré map for a periodic vector field.

Definition 3.2.4. Let M a smooth manifold, and $\phi : M \to M$ a diffeomorphism. Then its mapping torus

$$M_\phi := \frac{I \times M}{(0, x) \sim (1, \phi(x))},$$

is a fibre bundle over S^1. The **suspension foliation of** ϕ is then obtained by foliation over orbits of ϕ. That is to say a leaf is defined as

$$L_p := \{(t, \phi^n(p)) \mid n \in \mathbb{Z}\},$$

meaning any points we can reach after iteratively applying ϕ. ▲

Foliations arise in a variety of ways. In our case we will often look at foliations by either level sets, vector fields or distributions. In order to gain some more intuitions on foliations, observe the following three examples.

Example 3.2.4.1. Define $\mathbb{T}^2 := S^1 \times S^1$ with angle coordinates (θ, ϕ). Then the level sets of $f : \mathbb{T}^2 \to S^1$ defined as $f(\theta, \phi) \mapsto \theta$ defines a 1-dimensional foliation of the torus. The leaves are embedded copies of S^1. We see that f is a proper map, by definition and by the fact that $\mathbb{T}^2$ is compact. Furthermore, it is a surjective, smooth submersion, so it defines a fibre bundle $(\mathbb{T}^2, S^1, f, S^1)$. ◆

Example 3.2.4.2. Define a vector field Y on $\mathbb{T}^2$ as follows. Let X be the vector field on $\mathbb{R}^2$ defined as $X_p = \partial_x + r\partial_y$ with r an irrational number. Let $Y = \pi_* X$ where $\pi : \mathbb{R}^2 \to \mathbb{T}^2$ is the universal cover. We know that any non-vanishing vector field defines a foliation, and hence we have a 1-dimensional foliation of $\mathbb{T}^2$. Remark however that each leaf of the foliation is dense in $\mathbb{T}^2$. And hence every leaf comes arbitrarily close to any other leaf, but still does never intersect it. This is intuitively a result of each leaf being an immersed copy of non-compact $\mathbb{R}$ into compact $\mathbb{T}^2$. There are multiple obstructions to this being an embedding of $\mathbb{R}$ into $\mathbb{T}^2$, one being for example that in the subspace topology the leaf is not locally path connected. It is also not a fibre bundle. Note a single leaf will come arbitrarily close to itself, an infinite amount of times. So the leaf passes any neighbourhood U of p on the torus an infinite amount of times, hence $\pi^{-1}(U) \ncong U \times F$. ♦

Example 3.2.4.3 (Reeb Foliation)**.** The following is a very important example of a certain type of foliation which play a central role in the general theory of foliations. Consider the punctured upper-half plane $\mathbb{R}^2_+ \setminus \{0\}$ and its foliation by level sets $y = c$. These level sets are generally single copies of $\mathbb{R}$, except of course for $y = 0$, where we have two disjoint copies of $\mathbb{R}$. Observe the automorphism ψ on $\mathbb{R}^2_+ \setminus \{0\}$ given in polar coordinates by scaling $(x, y) \mapsto (\lambda x, \lambda y)$ with some positive scalar λ. Note that ψ maps leaves onto leaves: $L_{(x,y)} \to L_{(\lambda x, \lambda y)}$, and hence this foliation descends to a foliation on the quotient, for which there is an isomorphism

$$\Psi : \mathbb{R}^2_+ \setminus \{0\} / \psi \cong S^1 \times D^1,$$

to the two-dimensional annulus. Note that this isomorphism is slightly counter-intuitive; if we denote ψ in polar coordinates, it becomes $\psi : (r, \phi) \mapsto (\lambda r, \phi)$. The isomorphism Ψ then maps the radius coordinate r to S^1, rather than the angle coordinate, which itself maps to $[0, \pi] \cong D^1$. Let us observe what this foliation looks like by looking at the images of the leaves. The two copies of $\mathbb{R}$ in the preimage corresponding to $y = 0$ each get mapped to a copy of S^1. In a sketch of coordinates, the equivalence class of the interval $[1, \lambda] \subset \mathbb{R}$ corresponds to $S^1 \times \{0\}$, and $[-1, -\lambda]$ to $S^1 \times \{\pi\}$. Now look at the image of the level set $y = c$. Starting at $(0, c)$ we must be equally spaced between the two boundary circles. Now moving into the positive x-direction, we increase the radius, hence we move along S^1 after applying Ψ. We are simultaneously also decreasing the angular coordinate, hence we are moving towards the inner boundary circle. If we moved in the negative x-direction, we would still move along S^1 in the same direction, but now we would move towards the outer boundary circle. The result is that this foliation on the annulus can be sketched as some infinitely swirled two coloured turban cake before it is put in the oven, where the two copies of S^1 are its annulus cake mould. One sees for instance directly that it is not a fibre bundle, as two of the leaves are compact submanifolds diffeomorphic to S^1, all others are immersed copies of $\mathbb{R}$. ♦

As alluded to before, this example is a specific case of a general collection of foliations. In fact, nothing is stopping us to replace $\mathbb{R}^2_+ \setminus \{0\}$ by a general $\mathbb{R}^n_+ \setminus \{0\}$ and repeating the construction produce a codimension 1 foliation of $S^1 \times D^{n-1}$. We end up that the interior, which is homeomorphic to $S^1 \times (0, 1)^{n-1}$, is foliated by immersed copies of $\mathbb{R}^{n-1}$, and the boundary is a copy of $S^1 \times S^{n-2}$. Remark that in the annulus case, S^0 consists of two points.

Definition 3.2.5. The foliation described above for general n is called the **Reeb foliation** in dimension n. ▲

In the case of the example where $n = 2$ we speak of the Reeb annulus 8. In the case of $n = 3$ we speak of the Reeb foliation of a solid torus.

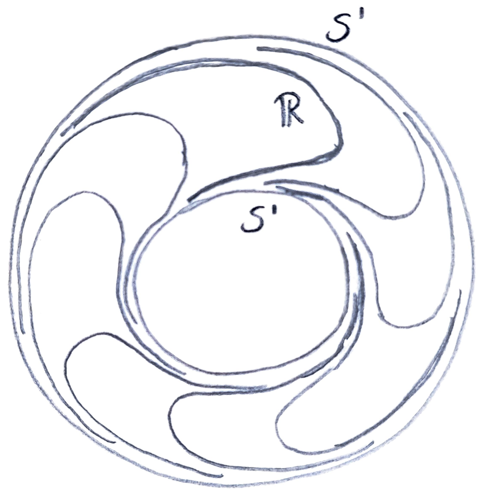

Figure 8: Reeb annulus

3.3 Foliations by Lie Group Actions

Given a vector field X, we have used the flow ϕ_X to determine integral curves. These integral curves then defined a 1-dimensional foliation of our manifold M. Recall that a complete flow is technically a group action of $\mathbb{R}$ on M, meaning $\phi_X : \mathbb{R} \times M \to M$ with ϕ_X^t a diffeomorphism from M onto itself. An integral curve passing through a point $p \in M$ is then precisely the orbit of the group $G = \mathbb{R}$ of point p, denoted by

$$Gp := \{gp \in M \,|\, g \in G\},$$

where we use the shorthand $\phi(g, p) =: gp$ which we will use more often. Thus, in essence we have foliated M by the orbits of the group $\mathbb{R}$. This construction can be generalized if we replace $\mathbb{R}$ by any general Lie group G which satisfies the following mild condition.

Definition 3.3.1. A Lie group action $\phi : G \times M \to M$ is called a **foliated action** if for every $p \in M$ we have that the tangent space of its orbit has constant dimension, meaning $\dim(T_x(Gp)) = k$ for all $x \in Gp$. ▲

Note a group action was free if $gp = p$ implied $g = id_G$, in other words: only the identity element fixes any point p. A less stringent condition is to be locally free, which is only possible for topological groups, in which there is an open neighbourhood U of id_G whose action is free. We then remark that this is equivalent to a foliated action having fixed tangent orbit dimension of the original group G. An intuitive way to think of the dimension k of the tangent space of the orbit at a point p, is the amount of infinitesimal actions of G which move p. As a result the dimension of this tangent space cannot be higher than the dimension of the Lie algebra $\dim(\mathfrak{g})$.

Now if the action is not locally free, then by definition we will have some complementary part $I \subset T_pM$ of dimension $(m-k)$ to the tangent space of the orbit $T_p(Gp)$. Intuitively the dimension of this subspace reflects the amount of infinitesimal actions of G which do not move p. Recall we had the so-called isotropy group of p being the subgroup of G fixing p and denoted by

$$G_p := \{g \in G \mid gp = p\}.$$

For any Lie group G, an isotropy group G_p is also a Lie group. As a result the dimension of I equals the dimension of the Lie algebra of the isotropy group $\dim(T_e(G_p)) = (m-k)$.

Armed with this intuition, let us prove that a foliated Lie group action, actually produces a foliation.

Lemma 3.3.2. *Let G a Lie group and $\phi : G \times M \to M$ be a foliated Lie group action. Then it defines a foliation $\mathcal{F}$ on M by setting $L_p := Gp$.*

Proof. Let G define such a foliated Lie group action, denote for each $p \in M$ the dimension by $\dim(T_p(Gp)) = k$. As above, decompose the tangent space

$$T_pM = I \oplus T_p(Gp),$$

into a part of dimensionality $(m-k)$ tangent to the orbit, and a part complementary to it. Similarly, decompose the tangent space at the identity of G into

$$T_eG = O \oplus T_eG_p,$$

where O is complementary to the isotropic subgroup and has dimension k. Note I reflects T_eG_p, and O reflects $T_p(Gp)$. Now embed two discs reflecting these directions as follows

$$i_1 : D^k \hookrightarrow G \text{ with } i_1(0) = id_G,$$

such that $di_1(T_0D^k) = O$. And similarly let

$$i_2\ D^{m-k} \hookrightarrow M \text{ with } i_2(0) = p,$$

such that $di_2(T_0D^{m-k}) = I$. Now define a map

$$\begin{aligned} \psi : D^k \times D^{m-k} &\to M, \\ (\gamma, \xi) &\mapsto i_1(\gamma) i_2(\xi). \end{aligned}$$

Note that by construction $\psi(0,0) = id_G p = p$ and by surjectivity and dimensionality we get that

$$(d\psi)_{(0,0)} : T_0D^k \times T_0D^{m-k} \to T_pM,$$

is a linear isomorphism, hence ψ is a local diffeomorphism. Choose an open neighbourhood U of p for which ψ is a diffeomorphism. This defines a foliated chart for the orbits of G. Indeed, let $q = i_2(x)$ be in U, by definition of the embeddings the orbit Gq is distinct from Gp. Now note that for a point in this orbit we can write $gq \in Gq$, and we obtain

$$\psi^{-1}(gq) = (i_1^{-1}(g), i_2^{-1}i_2(x)) = (i_1^{-1}(g), x) \in D^k \times \{x\}.$$

And thus we have a foliated chart of M foliated by orbits of G. ■

3.4 Subfoliations

Remark that if one starts with the Reeb foliation of the solid torus, and takes a top-down cross-section, one obtains the Reeb annulus. It will be more often that a foliation descends onto a foliation of a submanifold. For this the submanifold has to satisfy fairly relaxed prerequisites.

Definition 3.4.1. A map $f : N \to M$ is said to be **transverse** to a submanifold $L \subset M$ if whenever $f(p) = q \in L$ we have that

$$T_qM = T_qL + df_p(T_pN).$$

▲

This notion is readily extended for immersed submanifolds $i : N \hookrightarrow M$ and leaves L of a foliation $\mathcal{F}$. We say a submanifold is transverse to a foliation if it is transverse to all leaves it intersects of that foliation. Note that geometrically the different leaves of a foliation never intersect, and by transversality of the submanifold to the leaves, the intersections with the leaves each individually have measure 0. We see that the foliation hence readily descends onto a foliation of the submanifold.

Lemma 3.4.2. *If N is a transverse manifold to $\mathcal{F}$, then there is an induced foliation $\mathcal{F}_N$ of codimension $(m-k)$ on N, where the leaves are given by $\lambda_p := N \cap L_p$ for $p \in N$.*

Proof. We shall look pointwise, so pick two vector fields X, Y defined on $T\lambda_p$. Clearly $T\lambda_p \subset TN$ and as N certainly is a submanifold we have $[X, Y] \in TN$. However, similarly $T\lambda_p \subset TL_p$, which also by virtue of being a submanifold satisfies $[X, Y] \in TL_p$. It follows $[X, Y] \in TN \cap TL_p = T\lambda_p$. Hence, λ_p is a submanifold of N. ■

The above is an example of a subfoliation, the notion of which is most easily formalized using the language of distributions.

Definition 3.4.3. Let F_1 and F_2 be the distributions associated with two foliations $\mathcal{F}_1$ and $\mathcal{F}_2$. Then $\mathcal{F}_2$ is a **subfoliation** of $\mathcal{F}_1$ if F_2 is a vector subbundle of F_1. ▲

This formalizes the geometric notion that the integrable manifolds of F_2, hence the leaves of $\mathcal{F}_2$ are nicely contained within the integrable manifolds of F_1, hence the leaves of $\mathcal{F}_1$.

Example 3.4.3.1. Observe $\mathbb{R}^3$ and foliate it by the level sets of the following smooth function

$$f(x, y, z) = \sin(\pi y)(x - \tfrac{1}{2}) - z,$$

and denote the corresponding distribution by ξ. We can visualize this as a 2-dimensional foliation by planes with a little wave in them. Readily we see that it induces a subfoliation on $\mathbb{R}^2$ defined as taking the x, z, denote each corresponding distribution by the y-coordinate of the plane ξ_y. Note these are line sub-bundles, hence are span at a point by a single vector field. Now via the universal covering of $\pi : \mathbb{R} \to S^1$ we can try to descend this foliation on products of S^1.

We can easily take the quotient in the z-direction, as we see the foliation is invariant in this direction. It is also clear we can take the quotient in the y-direction, as f is periodic in the y-direction with period 2. We have now already descended to a foliation, and corresponding distribution ξ, on $\mathbb{R} \times S^1 \times S^1$. And correspondingly a subfoliation, and hence subbundle ξ_y, of $\mathbb{R} \times S^1$. Where this subfoliation is given by line bundles with slope $\sin(\pi y)$.

However, note such a quotient cannot possibly exist in the x-direction. If such a quotient did exist, then we would have induced sub-foliations of copies of tori where some ξ_y are pointwise span by irrational vector fields, and other are span by rational vector fields. However, via the foliation on the ambient space, we have induced flows, hence diffeomorphisms, between subfoliations. This would imply some diffeomorphism between S^1, the periodic orbits, and $\mathbb{R}$, the aperiodic orbits. This cannot possibly exist. ♦

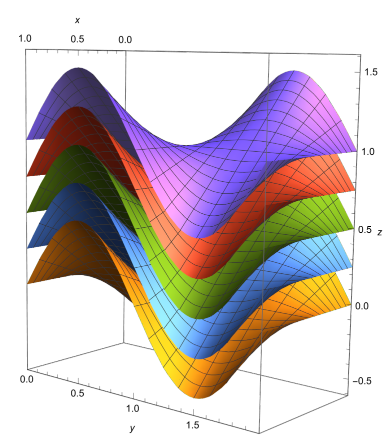

Figure 9: Foliation of $\mathbb{R}^3$ subfoliating $\mathbb{R} \times S^1$ but not $\mathbb{T}^2$

3.5 Coorientation and Transversals

So far we have given ways of constructing foliations. Foliating a manifold is in itself be a useful operation: it simplifies higher dimensional manifolds by breaking it up into lower dimensional parts. Furthermore, foliations often arise from some other structure, and these in turn apply some rigidity to the foliation, which translates to some structure on the entire foliated manifold. Let us define some properties of foliations.

First remark that like distributions, codimension 1 foliations can be cooriented.

Definition 3.5.1. A codimension 1 foliation $\mathcal{F}$ of a smooth manifold M is **coorientable** if for each leaf $L \subset \mathcal{F}$ the normal bundle ν_L is trivial. ▲

Similar to distributions we have the corollary of lemma 2.4.4 of orientability applying to foliations.

Corollary 3.5.2. *If M itself is already orientable, then coorientability is equivalent to orientability.*

Similarly, as a corollary of lemma, 2.4.3 we can apply Frobenius' to foliations and obtain the following result.

Corollary 3.5.3. *A codimension 1 foliation $\mathcal{F}$ is coorientable if and only if there exists a non-vanishing 1-form $\alpha \in \Omega^1(M)$ such that $\alpha|_L = 0$ for each leaf $L \subset \mathcal{F}$.*

We recall earlier that choosing a coorientation was geometrically equivalent to choosing a well-defined transverse direction of the submanifold. For codimension 1 foliations, this is necessarily a one-dimensional story. There is some use to extending this notion to higher dimensions.

Definition 3.5.4. Let M be an m-dimensional manifold, and $\mathcal{F}$ a k-dimensional foliation. A **transversal** τ to a foliation $\mathcal{F}$ is a submanifold of M of dimension $n := m - k$ such that at each $x \in \tau \cap L$ we have $T_xN \oplus T_xL = T_xM$. ▲

Note there is subtle difference between a transversal and being transverse. A transversal is a transverse submanifold with the added property the dimensions add up to the ambient space. If we again choose a local coordinate chart in which the foliations are flat, meaning a coordinate chart of M in which the leaves are k-dimensional planes by keeping the last $(m-k)$ coordinates constant, then a transversal is given by an n-dimensional plane keeping the first k coordinates constant. We can centre the chart around a transverse leaf to obtain the first n coordinates to be 0.

3.6 Distinguishing Foliations

Almost any manifold M can be foliated in a multitude of different ways. However, every foliation of the same dimension k locally looks like flat non-intersecting k-planes of $\mathbb{R}^m$. We would like to have some more tools in order to discern different types of foliations.

3.6.1 Holonomy

If one imagines moving along the surface of a leaf, then other leaves of the foliation can either diverge or asymptotically approach the leaf we are in. This is roughly what the notion of holonomy is.

For the sake of visualization let $\mathcal{F}$ be a codimension 1 foliation. Transversals τ of the foliation are then given by lines which locally intersect a leaf L at a single point. Denote by τ_p the transversal intersecting L_p at the point p, and denote by τ_q another transversal intersecting the same leaf L_p at some other point q.

We can join the points p and q by some path $\gamma : I \to M$ laying completely in the leaf L_p. Denote the points in γ by their time coordinates. Now as $\mathcal{F}$ by definition is accompanied by a foliated atlas, we can use a collection of foliated charts $\{(U_t, \phi_t)\}$ for all $t \in \gamma$ to cover γ. As γ is a proper map as I is compact, the path itself is compact. Because of this compactness we can extract a finite subcover indexed by i for $i \in [0, N] \subset \mathbb{N}$ with the properties that

$$[\gamma(t_i), \gamma(t_{i+1})] \subset U_i,$$

furthermore they are ordered in the order in which they cover γ such that consecutive charts overlap:

$$U_i \cap U_{i+1} \neq \varnothing.$$

We call this a chain subordinate to γ.

Denote a chosen transversal passing through a chosen point $t_i \in U_i \cap U_{i+1}$ by τ_i and set $\tau_0 = \tau_p$ and $\tau_N = \tau_q$. We can now use the foliated charts, together with the path defined on the leaf, to define transport of transversals to transversals as follows. Choose a point $x \in \tau_0 \cap U_0$, a point on the transversal which is sufficiently close to the starting point of the path p, such that the x projects onto p via the transversal. Now follow the path γ till a point $t_1 \in \tau_1 \subset U_0 \cap U_1$, where τ_1 a chosen transversal. Locally the transversal τ_1 intersects the leaf L_x at a unique point, denote it by y. As such a unique point can be found for any point $x \in \tau_0$ sufficiently close to p, we define a map $h_{0,1} : \tau_0 \to \tau_1$.

Now as $h_{0,1}(x) := y$ is in the overlap of charts, we can repeat the same process to map this point to another unique point on the transversal τ_2 in the next overlap. We have now mapped x to some point in the overlap $U_1 \cap U_2$, let's call this point

$$h_{0,2}(x) := h_{1,2}(y) = h_{1,2} \circ h_{0,1}(x).$$

It is clear to see this process can be done iteratively for the entire path, and moreover, the starting transversal can also be changed. Also remark we never truly used the dimension of the manifold, nor the foliation, only some basic assumptions which apply to any connected manifold and foliation of higher dimension than 1. Please do note that the transversals are now no longer given by lines, but rather by $(m - k)$ dimensional transverse manifolds, which locally look like $D^{(m-k)}$ discs. Keeping all else the same we define a collection of maps:

$$h_{i,j} : U_i \to U_j,$$

for a general foliated chain subordinate to γ on a manifold M, and compose them together to obtain the following definition.

Definition 3.6.1. Let M a manifold foliated by $\mathcal{F}$, let γ be a path connecting two point p, q of a leaf, and lastly let $h_{i,j}$ be defined as above. Then the **holonomy of $\mathcal{F}$ along** γ is the total composition

$$h_\gamma = h_{N-1,N} \circ \dots h_{0,1}.$$

h_γ is called a holonomy map. ▲

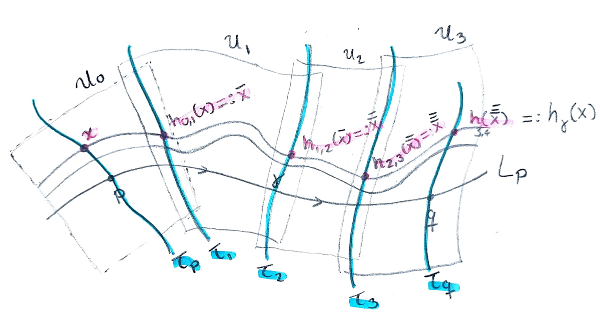

Figure 10: Holonomy map h_γ

It is a fact, we leave unproven, that the holonomy both does not depend on the chosen foliated charts. Moreover, we will show it does not depend on the chosen transversals τ_i for $1 \leq i < N$ up to conjugation.

Now as mentioned in the introduction, we are interested in where we are transported to, when we move along the leaf all the way up to when we re-encounter our starting neighbourhood of a point. This is the case when the starting and ending transversal are the same, in fact, when this is the case we have the following interesting property of holonomy maps.

Lemma 3.6.2. *If f_γ is defined using transversals $\tau_0 = \tau_N$, and g_γ is defined using transversals $\sigma_0 = \sigma_N$ then they differ by conjugacy by diffeomorphisms.*

Proof. Let us first assume for generality that $\tau_0 \neq \tau_N$ and $\sigma_0 \neq \sigma_N$. We can still use the foliated charts in possibly smaller neighbourhoods U_0 of t_0 and U_N of t_N. In these foliated charts we can project σ onto τ via $\pi_0 : \sigma_0 \to \tau_0$ and similarly for π_N. Locally these are diffeomorphisms. It is then clear to see:

$$g_\gamma = \pi_N^{-1} \circ f_\gamma \circ \pi_0.$$

It follows readily that if the starting and ending transversal are the same that this becomes

$$g_\gamma = \pi_0^{-1} \circ f_\gamma \circ \pi_0.$$

■

Now calling this property a conjugacy property alludes to group theory. In fact, as preparation for defining a so-called holonomy group, we would like to have some kind of group structure on holonomy maps. Note f_γ in the above defines a map from the starting transversal τ_0 to itself. Moreover, it is clear to see that for any holonomy map for any closed loop γ, it maps the base point of the transversal back to itself. Denoting the base point by x, we have the property

$$f_\gamma(x) = x.$$

As all holonomy maps agree at this point, an obvious equivalence class would be to look at germs around this point.

Definition 3.6.3. For a chosen transverse τ of $\mathcal{F}$ containing x we define a map germ $G(\tau, x)$ as equivalence classes of holonomy maps $f : \text{Dom}(f) \subset \tau \to \tau$ such that they agree on some neighbourhood of x in τ. ▲

We have a readily defined group structure on this space.

Lemma 3.6.4. *The space $G(\tau, x)$ has a group structure given by composition.*

Proof. This is a quick check of properties. Note that $f \circ g$ is a well-defined map. Furthermore, if some map g' agrees with g on some neighbourhood around x, and similarly for f and f'. Then $f \circ g$ agrees with $f' \circ g'$ on some neighbourhood contained in both. Hence, $[f] \circ [g] = [f \circ g]$. Note we have an identity holonomy map given by the null-loop, and an inverse of each holonomy map given by the reverse loop. ■

If we define the transversals to be the same, then we have implicitly defined γ is a loop. As a consequence that both the choice of chains subordinate to γ and intermediate transversals do not matter, we get the following result.

Lemma 3.6.5. *If γ_0 and γ_1 are homotopic paths. Then $f_{\gamma_0} = f_{\gamma_1}$.*

Sketch of Proof. Denote by H the homotopy relative to their endpoints. Given a chain subordinate to γ_0 it follows that this chain is also subordinate to a small enough perturbation γ_ε, hence they define the same holonomy map. Iteratively, γ_ε has its own subordinate chain, and we can use this chain to move even further up to homotopy. It follows that the holonomy maps of γ_0 and γ_1 agree. Setting the starting and end point to be equal applies this proof to loops. ■

We have a group structure on the fundamental group of a leaf with fixed base point given by concatenation of loops. We have shown above that homotopic loops define the same holonomy map germ. Lastly it is not difficult to see that

$$f_{\gamma_1 * \gamma_2} = f_{\gamma_2} \circ f_{\gamma_1},$$

meaning the holonomy map of a concatenation agrees with the composition of individual holonomy maps. We arrive at the following.

Lemma 3.6.6. *There is a group homomorphism*

$$\begin{aligned} \Phi : \pi_1(L, x) &\to G(\tau, x), \\ [\gamma] &\mapsto [f_\gamma]. \end{aligned}$$

Now we define the image of this map

$$\mathrm{Hol}(L, \tau, x) := \Phi(\pi_1(L, x))$$

as the holonomy group of L with respect to τ and x. However, by lemma 3.6.2, the chosen transversal did not matter up to conjugacy by diffeomorphisms. So we prefer to rather look at the conjugacy classes of this image

$$\mathrm{Hol}(L, x) := \mathrm{Cl}(\mathrm{Hol}(L, \tau, x)).$$

Moreover, if we have chosen another x' in L, any path δ from x to x' defines an isomorphism

$$\begin{aligned} \delta^* : \mathrm{Hol}(L, x) &\to \mathrm{Hol}(L, x'), \\ \Phi([\gamma]) &\mapsto f_\delta \circ \Phi([\gamma]) \circ f_{\delta^{-1}}. \end{aligned}$$

Thus, in a very similar construction as to how we generalize $\pi_1(L, x)$ to just $\pi_1(L)$, we get the following general definition.

Definition 3.6.7. The **holonomy group** $\mathrm{Hol}(L)$ is defined as any group isomorphic to $\mathrm{Hol}(L, x)$. ▲

When looking at $G(\tau, x)$, we are looking at germs of diffeomorphisms leaving x fixed. We are free to look at germs of diffeomorphisms up to a certain degree, equivalent to how we look at k-jets in A. A specifically recurring type is that of linear holonomy, where we take map-germs up to first degree equivalence; $k = 1$-jets. In this case it might be useful to interpret the spaces as follows. If $\mathcal{F}$ has codimension q, then the transversals τ have dimension q. Via charts centred around x we obtain a diffeomorphism $G(\tau, x) \cong G(\mathbb{R}^q, 0)$. Now note that the linearization of $G(\mathbb{R}^q, 0)$ at 0 is precisely $\mathrm{GL}_q(\mathbb{R})$.

Definition 3.6.8. The **linear holonomy** group $\mathrm{dHol}(L)$ of L is defined as any group isomorphic to $\mathrm{Cl}(T_0 \circ \Phi(\pi_1(L, x)))$, where $T_0 \circ \Phi(\pi_1(L, x))$ is diffeomorphic to some subgroup of $\mathrm{GL}_q(\mathbb{R})$. ▲

Unpacking the definition; conjugation classes fix the issue of choosing a base-point, and linearization is done by taking equivalence up to first derivative at a base-point identified with 0 under a chart.

In practice, we will often work with any $\mathrm{Hol}(L, \tau, x)$ representing $\mathrm{Hol}(L)$. Furthermore, it is often very useful to keep the group homomorphism Φ derived in lemma 3.6.6 in mind when calculating the holonomy group. In fact, it is clear that any contractible leaf has trivial fundamental group π_1, and hence the holonomy group of any such leaf will be trivial.

Definition 3.6.9. We say a leaf L of a foliation $\mathcal{F}$ has **trivial holonomy** if $\mathrm{Hol}(L) = [id]$. If all leaves of the foliation have trivial holonomy, then $\mathcal{F}$ is a foliation without holonomy. ▲

Corollary 3.6.10. *Any foliation with contractible leaves, or any foliation of a contractible space, has trivial holonomy.*

Let us gain some more intuition for the holonomy of a foliation by treating some examples.

Example 3.6.10.1. Observe the torus foliated by vector fields of irrational slope. Because the leaves of this foliation are immersed copies of $\mathbb{R}$, it follows this particular foliation has no holonomy. ♦

Example 3.6.10.2. Observe the torus foliated by vector fields with rational slope. In this case the leaves of the foliation are embedded copies of S^1, of which we know $\pi_1(S^1) = \mathbb{Z}$. So we might actually find some holonomy. But there is really only one possible unique loop: moving around the same orbit once or multiple times. Now choose a convenient the longitudinal. It is clear the holonomy map of this path on this transversal is simply the identity. Hence, we conclude this foliation, too, has no holonomy. ♦

Example 3.6.10.3. Let M be the mapping torus of the identity on the torus:

$$M := \frac{I \times \mathbb{T}^2}{(0, x) \sim (1, x)}.$$

This is in fact a fibre bundle over S^1 with fibres $\mathbb{T}^2$. As in lemma 3.2.3 define a 2-dimensional foliation where the leaves are given by copies of $\mathbb{T}^2$. We know $\pi_1(\mathbb{T}^2) = \mathbb{Z}^2$, so there is again a possibility of holonomy. As a transversal choose the base space S^1. It is clear that by property of a fibre bundle $\pi : \mathbb{T}^2 \to S^1$ is a single point for each torus. So either generator of $\pi_1(\mathbb{T})^2$ has as image the trivial holonomy map. We conclude that this foliation, too, has trivial holonomy, even though both the ambient space and the leaves are non-trivial. ♦

The first two examples show that, even though holonomy of a foliation is an invariant of the foliation, it may not always be enough to distinguish two different foliations which might be fundamentally different. Whilst the last example shows it might generally be difficult to guess what the holonomy is by just looking at the space. But let us revisit our example of the Reeb foliation the visit one of the first non-trivial examples of foliation with holonomy.

Example 3.6.10.4. In example 3.2.4.3 we studied and derived a way to foliate the general product of $S^1 \times D^{n-1}$. The interior of $S^1 \times D^{n-1}$ was foliation by $\mathbb{R}^{n-1}$ planes, whilst the boundary was a copy of $S^1 \times S^{n-2}$. In the case of $n = 3$ we get a foliation of the solid torus.

Observe S^3 as $\mathbb{R}^3 \cup \{\infty\}$. This gives us a way to split S^3 as the gluing of two solid tori along a boundary regular torus. For a quick geometric intuition this is done as follows. First embed the first solid torus with its usual geometric representation in $\mathbb{R}^3$. The strictly positive and negative complement of this embedding, minus the complement in the xy-plane, can be fibred by discs projecting on longitudinal circles of the first embedded torus. So now we have decomposed S^3 in three parts: a solid torus, an interval of discs $I \times D^2$, and the remaining plane plus infinity $\mathbb{R}^2 \backslash D^2 \cup \{\infty\}$. Note this last part is homeomorphic to $(\mathbb{R}^2 \cup \{\infty\}) \backslash D^2 \cong S^2 \backslash D^2 \cong D^2$. It is exactly this D^2 which is the boundary of the interval of discs $I \times D^2$, and hence we obtain another copy of $\mathbb{T}^2_S$, glued to the other copy of $\mathbb{T}^2_S$ via the boundary torus.

Now Reeb foliate both tori. As there is no interaction of their leaves, except their boundary leaf, which is nicely glued together, this foliation of $\mathbb{T}^2_S \sqcup \mathbb{T}^2_S$ descends onto a foliation of S^3. We now have a 2 dimensional foliation of S^3. We recall that the interior of the Reeb foliation where copies of euclidean planes. So the only leaf possibly admitting a non-trivial holonomy, is exactly this special $\mathbb{T}^2$ leaf. The Reeb foliation of the inner torus looks as expected. The Reeb foliation of the outer torus has its $\mathbb{R}^2$ leaves wrapping asymptotically ever closer towards the inner torus. The tops of the swirls stack through the hole of the inner torus towards infinity.

Recall again that $\pi_1(\mathbb{T}^2) = \mathbb{Z}^2$. And let the generators be a meridian and a longitude. Now pick a transversal τ as follows: for a point p on the torus, move upwards towards into $\mathbb{R}^3_+$ all the way up to $\{\infty\}$, after passing to $\{\infty\}$, we can move on upwards, but now we are moving away from $\{-\infty\}$ through $\mathbb{R}^3_-$, back to p. Now call the meridian m and the longitudinal l. Let us start with the holonomy of m. If we pick a point x outside the common torus, then the leaves appear to wrap ever closer if we move along m, so the holonomy map f_m is non-trivial, and maps x closer and closer to p. However, if we pick a point x inside the common torus, then the leaves are constant level when moving along m, so we have trivial holonomy. Conversely, if we investigate the holonomy map f_l the opposite is true. The leaves outside the common torus seem to be level when moving along l, so have trivial holonomy, but the leaves inside the common torus wrap ever closer and closer, so have non-trivial holonomy. So we have two independent generators of $\mathrm{Hol}(\mathbb{T}^2) \cong \mathbb{Z}^2$. ⧫

Now we have a firmer grasp on what the holonomy of a foliation is, and how to calculate them. Let us derive two more situations in which we can readily say the holonomy of a foliation is trivial.

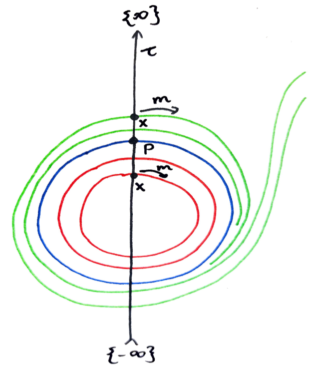

Figure 11: Meridian cross-section with transversal

First, recall that by theorem 2.3.8 any non-vanishing closed 1-form α on a smooth manifold M defines a smooth codimension 1 foliation via its kernel distribution $\ker(\alpha)$. Then the following lemma says it has trivial holonomy.

Lemma 3.6.11. *The smooth coorientable codimension one foliation defined via a non-vanishing closed 1-form α has trivial holonomy.*

Proof. As our codimension one foliation $\mathcal{F}$ is coorientable, one we have chosen a Riemannian metric on M, we can readily define a nowhere vanishing perpendicular, hence transverse, vector field to the leaves of the foliation. Now fix a leaf L of the foliation, together with an immersed loop γ in L. Use the flow of X to define an immersed cylinder, that is to say: define $\psi : S^1 \times (-\varepsilon, \varepsilon) =: C \to M$ as

$$(\theta, t) \mapsto \phi_X^t(\theta),$$

remark that this is also an immersion.

Now define the pullback form on C by pulling back α via ψ. And note that this pullback form is closed via

$$d(\psi^*\alpha) = \psi^* d\alpha = 0.$$

This again defines a codimension one foliation $\psi^*\mathcal{F}$, but now of C via $\ker(\psi^*\alpha)$. Not that this foliation is exactly given by the pullback of the cuts of $\mathcal{F}$ with our immersed cylinder.

Now choose a transverse $\tau = \theta \times (-\varepsilon, \varepsilon)$ to calculate the holonomy. The base leaf of this holonomy is S^1, and we know $\pi_1(S^1) = \mathbb{Z}$, so we only have one possible loop along this leaf, and it is generated by this element. Denote by f the holonomy map of this generator, meaning that for a $p \in \mathrm{Dom}(f) \subset \tau$ we have $f(p) = q$. If f is trivial, then $f(p) = p$ for all p. Now choose a $p \neq \theta$ in τ. Observe the region $R \subset C$ enclosed by first the base leaf $b := S^1 \times \{0\}$, second the segment $l \subset L_p$ connecting p and q along L_p, and lastly the straight segment δ connecting p and q along τ. We then see

$$\begin{aligned} 0 &= \int_R d(\psi^*\alpha) \\ &= \int_{\partial R} \psi^*\alpha \\ &= \int_b \psi^*\alpha + \int_l \psi^*\alpha + \int_\delta \psi^*\alpha \\ &= \int_\delta \psi^*\alpha, \end{aligned}$$

where in the second the last step we use the fact that both b and l are pullbacks of $\ker(\alpha)$. We conclude this equality only holds if $\delta = 0$, as τ was transverse to $\ker(\alpha)$ and α was nowhere vanishing. As a result it follows $f(p) = p$, hence we have trivial holonomy. ■

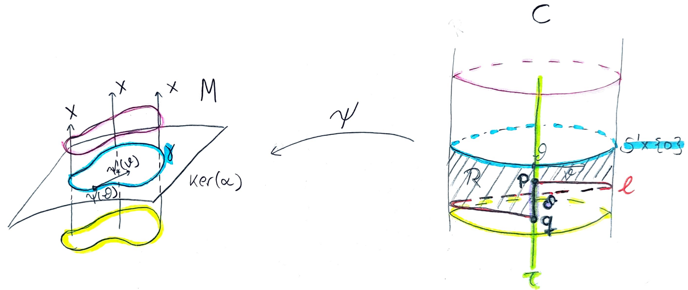

Figure 12: Trivial holonomy on $\ker(\alpha)$

The converse of this statement is a deep theorem by Sacksteder, which we will not prove, but mention.

Lemma 3.6.12 (Sacksteder)**.** *Let $\mathcal{F}$ be a smooth codimension one foliation of a smooth manifold M. Then $\mathcal{F}$ has trivial holonomy if and only if $\mathcal{F}$ is topologically equivalent to a foliation defined by a closed 1-form.*

Now the other situation where we can readily say the foliation is trivial, follows directly as a corollary of this theorem.

Corollary 3.6.13. *Any codimension 1 foliation which fibres over S^1 has trivial holonomy.*

Proof. Choose a non-vanishing 1-form θ on S^1, and let $\pi : M \to S^1$ be the projection of the fibration. Then as θ is closed by dimensionality, it follows $\pi^*\theta$ is a closed non-vanishing 1-form on S^1, hence defines a codimension 1 foliation with trivial holonomy. ■

We can now in fact combine several results together in the following big theorem due to Tischler.

Theorem 3.6.14 ([Tis70, Tischler]).
Let M be a closed smooth manifold, then the following are equivalent:
1. *M has a closed nowhere vanishing smooth 1-form.*
2. *M admits a codimension one smooth foliation without holonomy.*
3. *M admits a codimension one smooth foliation which is invariant by a transverse vector field.*
4. *M fibres over S^1.*

Proof. The theorem by Sacksteder 3.6.12 already implies the equivalence between statements 1 and 2.

Now we will show statement 3 implies statement 1. So assume such a transverse vector field X to a codimension one foliation $\mathcal{F}$ exists and denote its flow by ϕ_X^t. Now given a point $p \in M$, choose a foliated chart (U, ψ) centred around p. Note that locally $\psi(\phi_X^t) \in \mathbb{R}^{n-1} \times \{t\}$, so locally the leaves of $\mathcal{F}$ are taken to be $\mathbb{R}^{n-1}$ planes having the last coordinate constant, and it is along this coordinate the flow of X points. Define a projection onto the time coordinate of the flow by defining $\pi : \mathbb{R}^{n-1} \times \mathbb{R} \to \mathbb{R}$ as projection onto the last coordinate. Then $\pi \circ \psi(q) \in \mathbb{R}$ projects a point q in M to the time coordinate needed to flow from the leaf L_p upwards to q. We can define an exact 1-form on U via

$$\omega_U := d(\pi \circ \psi).$$

We can extend this to a smooth nowhere vanishing 1-form ω on M. Given two foliated charts (U_i, ψ_i) for $i = 1, 2$ which overlap, and respectively define $\omega_{U_i} := \omega_i$ then we see

$$\begin{aligned}
\omega_1|_{U_1 \cap U_2} &= d(\pi \circ \psi_1)|_{U_1 \cap U_2} \\
&= d(\pi \circ \psi_1 \circ \psi_1^{-1} \circ \psi_2)|_{U_1 \cap U_2} \\
&= d(\pi \circ \psi_2)|_{U_1 \cap U_2} \\
&= \omega_2|_{U_1 \cap U_2}.
\end{aligned}$$

Now we see ω is a smooth nowhere vanishing 1-form on M, it is even exact, so in particular it is closed. Now lastly because $\pi \circ \psi \circ \phi_X^t = t$ by definition, it satisfies the following property.

$$((\phi_X^t)^* \omega)(X) = \frac{d}{dt}((\pi \circ \psi)(\phi_X^t)) = 1,$$

and thus $L_X \omega = 0$ by Cartan, thus X preserves the foliation of ω.

Now note that 4 readily implies 2 via lemma 3.6.13, as we can use the fibres, which are necessarily of codimension one, as a foliation, and have trivial holonomy as they fibre over S^1.

Now Tischler proved that statement 1 implies statement 4, and hence we have an equivalence of all statements. The proof relies on some concepts related to the de Rham isomorphism, a quick recollection of which is given in appendix B.

The crux of Tischler's proof lies in the fact that any non-integral de Rham class can be approximated by at least a rational de Rham class. This is essential as a priori the resulting foliation of a closed 1-form need not be closed, hence cannot fibre over S^1. As is for example the case for the irrational foliation of the torus.

Let α be the closed nowhere vanishing smooth 1-form as in the assumption. Now as M is a closed manifold, each $H^n_{dR}(M)$ is of finite rank, related to the familiar Betti numbers. Let k be the rank for $n = 1$. Choosing a basis ω_i generating $H^1_{dR}(M;\mathbb{Z})$, we can write

$$\alpha = \sum_{i=1}^{k} c_i \omega_i + dh$$

for real numbers c_i and some $h \in C^\infty(M)$.

As described in B, we can now choose a collection of maps $f_i : M \to S^1$ each with the property

$$f_i^* \mu = \omega_i - dh_i.$$

Now we rewrite α in the following form:

$$\alpha = \sum_{i=1}^{k} c_i f_i^*(\mu) + \sum_{i=1}^{k} c_i dh_i + dh.$$

Now remark that as M is by assumption compact, the image of any smooth function $g(M)$ is some closed interval in $\mathbb{R}$. As additionally M is without boundary, g can pass to a quotient S^1 of $\mathbb{R}$. Denote by $\pi : \mathbb{R} \to S^1$ the projection, now we can write for any sum of pullback form and exact form

$$f_i^*(\mu) + dg = (f_i^* + \pi \circ g)^*(\mu).$$

Now by slight abuse of notation we define

$$f_i^*(\mu) = \left(f_i + \pi \circ (h_i + \frac{h}{C}) \right)^* (\mu),$$

where $C = \sum_i c_i$. Use this to rewrite α in the following appealing form

$$\alpha = \sum_{i=1}^{k} c_i f_i^*(\mu).$$

Now choose a Riemannian metric on M with an induced norm. Then as $\mathbb{Q}$ is dense in $\mathbb{R}$, we can choose rational numbers with a common denominator $\frac{n_i}{m}$ such that

$$\|\alpha - \frac{1}{m} \sum_{i=1}^{k} n_i f_i^*(\mu)\|$$

becomes arbitrarily small. Define

$$\beta := \sum_{i=1}^{k} n_i f_i^*(\mu).$$

By non-vanishing of α, we have that β is non-vanishing as well. Furthermore, μ was closed, hence $f_i^*\mu$ is closed, thus β is a closed, non-vanishing 1-form on M. More importantly, as the result of all the above machination, we see that

$$\mathrm{Per}(\beta) = \left\langle \sum_{i=1}^{k} \int_\gamma n_i f_i^*(\mu) \mid \gamma \in H_1(M;\mathbb{Z}) \right\rangle = \langle \lambda_\gamma \in \mathbb{Z} \rangle < (\mathbb{Z}, +),$$

and so β is in fact an integral 1-form. Recalling the bijection between functions $M \to S^1$ and integral 1-forms, we define $f : M \to S^1$ as the smooth function such that $f^*\mu = \beta$, and we readily see it must follow

$$f = \sum_{i}^{k} n_i f_i.$$

We conclude $f : M \to S^1$ is non-vanishing, and hence a proper, surjective, smooth submersion. By lemma 3.2.2 this defines a fibre bundle over S^1.

Lastly, statement 4 implies statement 3. Assume M fibres over S^1. Then define ∂_t as the vector field transverse to the fibres. The codimension 1 foliation of the fibres is naturally defined by $\ker(\pi^* d\theta)$, where $\pi : M \to S^1$ and $[d\theta] \in H^1_{dR}(S^1)$. It follows that

$$L_{\partial_t}(p^* d\theta) = d i_{\partial t}(p^* d\theta) = 0.$$

∎

The following corollary shows that having multiple such closed non-vanishing 1-forms gives a fibration over a torus.

Corollary 3.6.15. *If M admits k closed linearly independent non-vanishing smooth 1-forms $\omega_1, \ldots, \omega_k$, then M fibres over the torus $\mathbb{T}^k$. Equivalently M admits a k-dimensional foliation whose leaves are diffeomorphic to $\mathbb{T}^k$*

Proof. If we denote the induced map by ω_i by f_i then we get a proper, smooth surjective submersion

$$\begin{aligned} F : M &\to S^1 \times \cdots \times S^1 \cong \mathbb{T}^k, \\ p &\mapsto (f_1(p), \ldots, f_k(p)). \end{aligned}$$

Hence, by lemma 3.2.2 we have a fibre bundle. And by 3.2.3 we have a foliation where the leaves are $L_p \cong \mathbb{T}^k$ ∎

In fact, one might readily see the resemblance to the Arnold-Liouville theorem 2.5.11, which postulates that the compact and connected level sets of any integrable Hamiltonian system are diffeomorphic to k-tori. This was done by using the symplectic form ω to define k linearly independent closed nowhere vanishing 1-forms $df_i = \omega_i := i_{X_i}\omega$, where X_i was a Hamiltonian vector field of a smooth function $f_i : M \to \mathbb{R}$.

3.6.2 Godbillon-Vey Invariant

Another invariant of foliations is the Godbillon-Vey invariant. Recall from lemma 2.4.3 that any codimension one distribution ξ was coorientable if and only if it was the kernel distribution of some nowhere vanishing 1-form. By Frobenius' theorem 2.3.8 we get the corollary that if $\xi = \ker(\alpha)$ defines a codimension one foliation, then $\alpha \wedge d\alpha = 0$. As a consequence it must follow that we can write

$$d\alpha = \alpha \wedge \beta$$

for some non-vanishing 1-form β. Note by differentiating we get

$$0 = d\alpha \wedge \beta - \alpha \wedge d\beta = -\alpha \wedge d\beta.$$

And so we similarly deduce that

$$d\beta = \alpha \wedge \theta,$$

for some non-vanishing 1-form θ. This is where this process stops, however, let us observe the 3-form $\beta \wedge d\beta$. Differentiating this we obtain

$$d(\beta \wedge d\beta) = d\beta \wedge d\beta - \beta \wedge d^2\beta = 0,$$

so this 3-form is in fact closed. So we conclude

$$[\beta \wedge d\beta] \in H^3_{dR}(M).$$

Definition 3.6.16. The 3-form $\beta \wedge d\beta$ as defined above is called the **Godbillon-Vey form**. And $[\beta \wedge d\beta]$ is the **Godbillon-Vey class** denoted $gv(\mathcal{F})$. ▲

Now the interesting part is that this is an invariant of the foliation $\mathcal{F}$.

Lemma 3.6.17. *Let $\mathcal{F}$ be a coorientable codimension 1 foliation defined by a non-vanishing 1-form α. Then the cohomology class of the Godbillon-Vey form $[\beta \wedge d\beta] \in H^3_{dR}$ is an invariant of $\mathcal{F}$.*

Proof. The proof boil down to showing it is independent of choice of α and β. Let us first show that the class $[\beta \wedge d\beta]$ is invariant under choosing a different

$\beta' = \beta + f\alpha$, for some $f \in C^\infty(M)$ which still has the property $d\alpha = \alpha \wedge \beta'$. From some algebraic manipulation it follows rather readily that

$$\begin{aligned}\beta' \wedge d\beta' &= (\beta + f\alpha) \wedge d(\beta + f\alpha) \\ &= \beta \wedge d\beta + \beta \wedge df \wedge \alpha \\ &= \beta \wedge d\beta + d\alpha \wedge df \\ &= \beta \wedge d\beta + d(\alpha \wedge df),\end{aligned}$$

so $\beta' \wedge d\beta' \in [\beta \wedge d\beta]$. Now if we scale α by some function g which is non-vanishing on M, then it still defines the same kernel, and hence foliation. So let $\alpha' = g\alpha$, then we see:

$$\begin{aligned}d\alpha' &= dg \wedge \alpha + gd\alpha \\ &= dg \wedge \alpha + g\alpha \wedge \beta \\ &= dg \wedge \tfrac{1}{g}\alpha' + \alpha' \wedge \beta \\ &= d\log(g) \wedge \alpha' + \alpha' \wedge \beta \\ &= (d\log(g) - \beta) \wedge \alpha'.\end{aligned}$$

Note if we define $\beta' = d\log(g) - \beta$, we see $\beta' \in [\beta]$. If β and β' are cohomologous, then $d\beta = d\beta'$. As a consequence $[\beta' \wedge d\beta'] = [\beta \wedge d\beta]$.

We conclude the Godbillon-Vey class is an invariant of $\mathcal{F}$. ∎

Even though the Godbillon-Vey class is readily calculated and proved, it is more unintuitive to find a geometric interpretation of this invariant. We will presently give such a visualization in the form of defining a concept of infinitesimal holonomy.

We have shown that in the case $\ker(\alpha)$ defines a foliation, we get for free a 1-form β with the above properties. One visualizes $\ker(\alpha)$ as globally defined hyperplanes, and one can locally visualize $\ker(\beta)$ similarly. Note that it need not be globally, as $\ker(\beta)$ need not define a regular foliation. However, locally, by property of $d\alpha = \alpha \wedge \beta$, we obtain than whenever β does not vanish, the kernel of $d\alpha$ can be seen as the intersection of the hyperplanes

$$\ker(d\alpha) = \ker(\alpha) \cap \ker(\beta).$$

Now to define the holonomy of a foliation, we had to make a choice of transversal. We have shown before in lemma 3.6.2 that the choice of transversal did not matter up to conjugation by diffeomorphisms. However, having the codimension one foliation defined by the kernel of a non-vanishing 1-form gives us a canonical choice of transversal. The non-vanishing 1-form α defines a coorientation, aptly also called a transverse orientation, as is outlined in section 2.4. This is specifically a choice of linear transversal, namely the fibre of the trivialized normal bundle:

$$\phi : \pi_1(L, x) \to G((\nu_L)_x, 0).$$

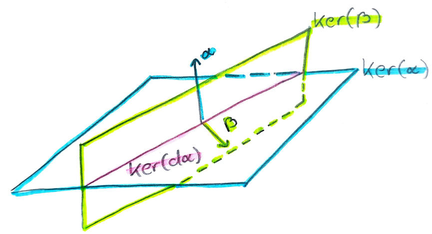

Figure 13: Interaction of kernels

Moreover, as we are working with a codimension 1-foliation, we obtain that the linearization of $G((\nu_L)_x, 0)$ is diffeomorphic to $\mathrm{GL}_1(\mathbb{R}) = \mathbb{R}^*$, the multiplicative group of the reals. Hence, the linear holonomy group $\mathrm{dHol}(L \subset \ker(\alpha))$ canonically acts on fibres of the normal bundle of L by non-zero scalar multiplication. Thus, we obtain that α induces a map α^*, which provides us a way to canonically measure linear holonomy. It associates to a linear holonomy map $T_0 f_\gamma$ for each $\gamma \in \pi_1(L, x)$ an element in the scalar multiplicative group:

$$\alpha^* : \pi_1(L, x) \to \mathbb{R}^*,$$
$$\gamma \mapsto \alpha^*(\gamma).$$

An added benefit of having canonically defined transversals is that we can extend this notion to any infinitesimal subsection of γ. Let $\dot{\gamma}$ be the velocity vector of γ. Via Cartan's formula we deduce

$$L_{\dot{\gamma}}\alpha = di_{\dot{\gamma}}\alpha + i_{\dot{\gamma}}d\alpha = i_{\dot{\gamma}}(\alpha \wedge \beta) = -\beta(\dot{\gamma}) \cdot \alpha.$$

This clarifies the notion we set to introduce: β defines some infinitesimal holonomy along γ. The interpretation of this is if $\beta(\dot{\gamma}) < 0$ then the leaves have diverged along γ, similarly $\beta(\dot{\gamma}) > 0$ implies the leaves have converged. If $\beta(\dot{\gamma}) = 0$ then the leaves have stayed constant on this part. Moreover, recall we could visualize $\ker(d\alpha) = \ker(\alpha) \cap \ker(\beta)$, so the line field on $\ker(\alpha)$ defined by this intersection, is precisely the directions along which the leaves reflect parallel to each other.

Now as a result it must follow that when integrated over the entire path γ, the infinitesimal holonomy has to add up to the linear holonomy

$$\int_\gamma \beta = \alpha^*(\gamma).$$

Also note β is leaf-wise closed, as $d\beta$ is some multiple of α, whose kernel exactly defines the leaves. Moreover, this also nicely agrees with our preconception that null-homotopic leaves should have trivial holonomy. Indeed, if each leaf is null-homotopic, then by definition $\gamma = \partial D$ for some contractible disc fully contained in the leaf. We observe that

$$\int_{\gamma := \partial D} \beta = \int_D d\beta = 0,$$

by closedness of β leaf-wise. Now by closedness of β on each leaf, we get that although β may not define a codimension one foliation of M, it does define a codimension one, possibly singular, subfoliation of $\ker(\alpha)$. This line field coincides precisely with the aforementioned $\ker(d\alpha) = \ker(\alpha) \cap \ker(\beta)$.

Now we will finally conclude how the above discussion will aid us in visualizing the Godbillon-Vey invariant. Recall these lines on $\ker(\alpha)$ precisely define paths along which the leaves run perfectly parallel to each other: there was no infinitesimal holonomy. Moving into the direction defined by the natural coorientation of $\ker(\beta)$ means $\beta(\dot{\gamma}) = 1$, or some other positive constant, thus the leaves pinch together. Conversely, moving in the opposite direction, the leaves open up.

Now, and this concept will be central in contact geometry, if $\beta \wedge d\beta \neq 0$, then $\ker(\beta)$ cannot possibly define an integral manifold. It must follow the line fields on nearby leaves must constantly be changing, and hence the way the leaves are pinched together must constantly be changing. As we will explore further in section 4.1, this change is visualized by twisting of the hyperplanes defined by $\ker(\beta)$, hence the leaves of $\ker(\alpha)$ have some helical pinching [Thu72].

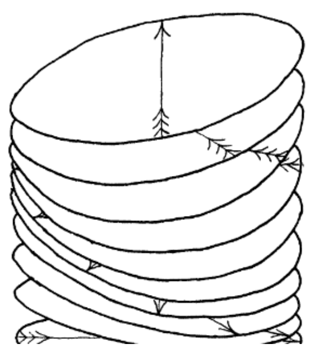

Figure 14: Helical Wobble [Thu72]

Which gives us a fairly geometric intuition of the Godbillon-Vey invariant.

Let derive some properties of this invariant. Recall that by Sacksteder's theorem 3.6.12 we have that any codimension one foliation which has trivial holonomy is topologically equivalent to a foliation defined by a closed 1 form. And so as a corollary we get the following.

Corollary 3.6.18. *Let M a closed smooth manifold and $\mathcal{F}$ is smooth codimension one foliation. If $\mathcal{F}$ has trivial holonomy, then $gv(\mathcal{F}) = 0$.*

Now let M be a smooth 3-manifold, and $\mathcal{F}$ a coorientable codimension one foliation. Because the Godbillon-Vey class in this case is a top form, it can be integrated over M. Recall cohomologous volume forms define the same volume on M, so if we associate to $gv(\mathcal{F})$ its induced volume

$$gv(\mathcal{F})(M) := \int_M gv(\mathcal{F}),$$

then we set the following lemma.

Lemma 3.6.19. *For a three-dimensional manifold M;*

- *If M is closed, then $gv(\mathcal{F})(M) \in \mathbb{R}$ need not be zero.*
- *If M is open, then $gv(\mathcal{F})(M) = 0$ for any $\mathcal{F}$.*

Proof. The first statement coincides with the fact that a closed manifold is, by definition, compact, hence there can be a defined volume. The second statement follows that an open manifold if by definition non-compact. By Poincaré duality the top singular cohomology group is isomorphic to the 0th compactly supported cohomology group $H^n_M \cong H^0_c(M)$. The latter is 0 if and only if M is non-compact, hence every top-form H^n_M is exact, thus by Stokes' theorem any integral over M is zero. ■

3.7 Taut Foliations

In the following section we introduce a special class of foliations with a certain type of rigidity introduced to them called tautness. Even though the concept of tautness can be generalized to both higher codimensions and dimensions, we will focus our interests on the case where the ambient manifold is of dimension three, and the foliation of codimension one, which is the form in which they were first introduced. In this case we have the following definition.

Definition 3.7.1. A codimension one, cooriented, oriented foliation $\mathcal{F}$ of a connected closed 3-manifold M is **taut** if for each leaf $L \subset \mathcal{F}$ there exists an immersed smooth closed curve intersecting the leaf, and doing so only transversely. ▲

We will from now on always assume $\mathcal{F}$ to be smooth of codimension one, cooriented and oriented, and M to be a smooth closed connected 3-manifold. Now directly from the definition we can readily describe the following construction of taut foliations.

Example 3.7.1.1. Let M be a closed connected 3-manifold, and let M fibre over S^1 with fibres L. This produces a foliation of M with leaves indexed by S^1. The monodromy map obtained by traversing S^1 once, defines a homeomorphism $\phi : L \to L$. It is with respect to this map that M can be seen as the mapping torus $M \cong L_\phi$. Now consider the immersed closed C^0 curve c' formed by concatenating a path $\gamma \subset L$ connecting a point $p \in L$ with $\phi(p) \in L$ together with the path $\tau \subset M$ defined by $I \times \{p\}/\sim$, of p passing through the mapping torus L_ϕ. It is clear τ is transverse to every leaf of the foliation except L_p. Via a perturbation of c' we produce c which is transverse to every leaf. ♦

There are many equivalent definitions of taut foliations, we will not give all of them, but will focus on a few which we will briefly introduce without much of a proof.

The first of which we will introduce is that such perturbations as done in the above example, can actually produce a single immersed closed curve passing through all the leaves.

Lemma 3.7.2. *If $\mathcal{F}$ is taut, there is a single immersed closed curve γ intersecting every leaf transversely.*

Keeping this very geometric picture in mind, the following equivalent definition is not hard to intuitively imagine.

Lemma 3.7.3. *There is no proper embedding $i : N \hookrightarrow M$ of a compact N for which the boundary ∂N is tangent to $\mathcal{F}$, and whose coorientation points everywhere inwards or outwards of M. Such a submanifold is called a dead end.*

Example 3.7.3.1. The Reeb foliation of the two torus cannot be taut, as the boundary torus leaf is an embedded submanifold satisfying the hypotheses. ⧫

In fact this is generalized to a definition-lemma combination.

Definition 3.7.4. Any properly embedded $S^1 \times D^{n-1}$ together with its Reeb foliation, such that its boundary $S^1 \times S^{n-2}$ is tangent to $\mathcal{F}$ is called a **Reeb component** of $\mathcal{F}$. ▲

Lemma 3.7.5. *Any foliation admitting a Reeb component cannot be taut.*

The following equivalent definition serves more so as an explanation of the name taut, and we will only mention it together with a slight sketch of the proof to illustrate the moniker.

Lemma 3.7.6. *A foliation $\mathcal{F}$ is taut if and only if there exists a smooth metric g on M for which the leaves are minimal surface area surfaces.*

We have not given a formal definition of what being a minimal surface area means, but we will illustrate the intuitive part. If the induced volume form on M restricts to an area form λ on each leaf $L \subset \mathcal{F}$, then for any $D \subset L$ we have that for any other $D' \subset M$ with a shared boundary ∂D, and isotopic to D relative to this boundary ∂D. We have that $\int_D \lambda \leq \int_{D'} \lambda'$. Thus, one can imagine the taut foliation to be the foliation pulled the most taut: any isotopy of the leaves will produce leaves with a greater surface area, and hence we would have to somehow stretch the leaves.

The last equivalence we will give, may be ever so slightly more abstract, but we will make use of this in the context of contact structure and stable Hamiltonian structures.

Lemma 3.7.7. *A foliation $\mathcal{F}$ is taut if and only if there is a transverse flow which preserves some volume form on M*

Proof. By assumption $\mathcal{F}$ was coorientable, so there exists some nowhere vanishing 1-form which defines $\mathcal{F}$ via $\ker(\alpha)$. Assume $\mathcal{F}$ is taut, and hence there exist a single immersed transversal closed curve intersecting each leaf at a point p. By dimensionality, we can perturb such an immersed curve to be an actual embedding. Call each such curve τ_p. Construct a tubular open neighbourhood T_p for each of these. By compactness of M it follows a finite sub-collection of these cover M, call these T_i indexed by some I. Each such T_i is now an embedded copy of the solid open torus $t_i : \mathring{D}^2 \times S^1 \hookrightarrow M$.

Now given radius-angle coordinates (r, ϕ) define a closed non-vanishing 2-form on D^2 via $\beta := (1-r)dr \wedge d\phi$, and note that it has the property $\beta = 0$ on ∂D^2 but non-vanishing everywhere else. The projection $\pi : \mathring{D}^2 \times S^1 \to D^2$ pulls back this form to a non-vanishing closed 2-form $\omega' := \pi^* \beta$. Now using the embeddings to define a family on non-vanishing closed 2 forms on each open torus T_i via $\omega_i := t_i^* \omega'$. And define a non-vanishing closed smooth 2-form on M by taking the sum of all these

$$\omega := \sum_{i \in I} \omega_i.$$

Now we define a volume form on M by setting

$$\mu := \omega \wedge \alpha.$$

Note that by dimensionality the non-vanishing 2-form ω must have a one dimensional kernel, so define the rank 1 distribution to be $\xi = \ker(\omega)$. Necessarily the kernel of ω must point along τ_i the closed embedded curve transverse to the leaves of $\mathcal{F}$. So the distribution defined by ξ is transverse to $\mathcal{F}$. Define a constant non-vanishing vector field such that

$$\langle X \rangle = \xi,$$

and $\alpha(X) = 1$. We see this preserves the volume form μ as we have

$$\begin{aligned} L_x \mu &= i_X d\mu + d i_X \mu \\ &= d i_X (\omega \otimes \alpha - \alpha \otimes \omega) \\ &= 0. \end{aligned}$$

And so any taut foliation preserves some volume form on M.

Conversely, let X be such a vector field. Assume on the contrary $\mathcal{F}$ is not taut, and hence there is some properly embedded dead end N. The flow of X must necessarily flow inside N. The integral curves defined by the flows of X are compact, and hence must be compact in N. It follows N would compactly embed into itself. As a result the area form cannot be preserved. ∎

Example 3.7.7.1. Similar to the previous example. Let Σ be a closed surface together with an area form σ, and let $\phi : \Sigma \to \Sigma$ be some homeomorphism.

Define the mapping torus Σ_ϕ as usual. Then the pullback form of dt on S^1 via the usual projection π, together with the area form σ defines a volume form on Σ_ϕ. Defining the vector field on Σ_ϕ as ∂_t, we see

$$L_{\partial_t}(\pi^* dt \wedge \omega) = 0.$$

Which is a different way to prove this fibre foliation is taut. ◆

4 Stable Hamiltonian Structures

Equipped with a thorough understanding of distributions, integrability and foliations, we are ready treat to direct precursor of stable Hamiltonian structures; contact structures, which capture the idea of being maximally non-integrable. After this has been done, we will finally introduce the definition of stable Hamiltonian structures, and explain how they unify these seemingly opposing of taut foliations defines by a closed 1-form and maximally non-integrable distributions.

4.1 Contact Geometry

The natural mathematical framework for classical mechanics has been found to be within the field of symplectic geometry. The underlying intuition is that most mechanical system can be split up in a momentum-position configuration which fully describes the system. One then defines a symplectic form as a closed nondegenerate differential 2-form wedging differential positions with differential momenta. By virtue of non-degeneracy of this form, symplectic geometry naturally can only be applied to even dimensional manifolds.

Now given such a classical system, one defines the Hamiltonian h on it. The Hamiltonian naturally can be seen as an energy function, and level sets of the Hamiltonian correspond with configuration space of a certain energy level. By virtue of the Hamiltonian, the level sets of h are symplectomorphic to embedded $\mathbb{T}^m$ in our symplectic manifold, tangent to these surfaces were the Hamiltonian vector fields.

To transform from one embedded torus to another is to intuitively transform an integral manifold tangent to the Hamiltonian vector fields to another one with the same properties. Recalling the narrative that integral manifolds represented solutions to a system of homogeneous linear partial differential equations, this means we are transforming one system of equations to another.

In its original terms a contact transformation is then a map one Pfaffian equation to another:

$$dz - \sum_{i=1}^{p} y_i dx_i = 0 \mapsto f \cdot (dz - \sum_{i=1}^{p} y_i dx_i) = 0. \tag{4.1}$$

Where $f \in C^\infty(\mathbb{R}^{(2n+1)})$ does not vanish, at least locally. If we set α to be the left-hand side of either equation, then the p is maximal in its property that $\alpha \wedge (d\alpha)^p \neq 0$.

In order to gain some visualization, let us work locally. Now in a foliated chart, a foliation locally looks like a collection of hyperplanes. Originally a contact element was exactly that: a collection of hyperplanes. Observe the collection of hyperplanes in $\mathbb{R}^m \times \mathbb{R}$ given by the equations

$$z = \sum_{i=1}^{m} p_i x_i + C,$$

where now p_i are real numbers. We see alternatively the entire collection is defined by kernel of the 1-form

$$\beta := dz - \sum_{i=1}^{m} p_i dx_i.$$

The space of all such contact elements can be identified with $\mathbb{R}^{2m+1}$. In this case β would have coordinates (x_i, p_i, z). We arrive back at the initially introduced shape of equation (4.1). An element such as β are the precursor to contact forms.

Now a contact transformation, is in a sense equivalent to giving the flow of a non-vanishing vector field X on M transverse to each hyperplane. In order for X to be transverse to the kernel, one needs $i_X\beta \neq 0$. A consequence of this property together with the fact that X has to preserve $\ker(\beta)$, is the Liouville-esque property

$$L_X\beta = di_X\beta = f \cdot \beta.$$

Now look back on equation (4.1), the contact element β, and the aforementioned role of p. Note these are 1-forms, so by Frobenius' theorem if and only if $p = 0$, one has $\alpha \wedge d\alpha = 0$. Only in this case do we have a codimension one foliation of our space of contact elements. Conversely, one sees that for β we have $p = m$, so it is maximally away from Frobenius' as can be. Geometrically this means $\beta \neq \beta'$ cannot be connected by a non-transverse vector field, which intuitively is a desired property as otherwise the hyperplanes were not really distinct in the first place.

One can observe we have an odd amount of variables in contact element space, so this would be even codimension one foliation of an odd dimensional manifold. Upon closer inspection one also sees that $d\alpha$ looks suspiciously like ω_0 if we restrict it to a (x, y)-plane. It is in this context contact geometry arises naturally as the odd-dimensional counterpart of symplectic geometry.

The theory of contact geometry is rich, we will give a brief oversight of basic theory in order to ease the transition into stable Hamiltonian structures in the following section.

4.1.1 Contact Manifolds

In the above we derived contact forms in the way they were initially stumbled upon. However, as with many mathematical objects, this might not be the most general, nor the most insightful way. In the following we will define contact structures in their own right. For this section we will assume M to be a $(2m + 1)$ dimensional manifold.

Let $\alpha \in \Omega^1(M)$ be a nowhere vanishing 1-form. Then we can define a rank $2m$ distribution $\xi := \ker(\alpha)$ by virtue of Frobenius' theorem. Remark that this distribution is necessarily coorientable via lemma 2.4.3. Now recall that ξ was integrable if and only if ξ was involutive. This necessarily implied for two vector fields $X, Y \in \xi$, that

$$d\alpha(X,Y) = X(\alpha(Y)) - Y(\alpha(X)) - \alpha([X,Y]) = -\alpha([X,Y]) = 0.$$

By the contrapositive we see exactly that if $[X,Y] \notin \xi$ implies ξ is non-integrable.

Now $d\alpha|_\xi \neq 0$ still implies there might be a smaller kernel of $d\alpha$ contained in ξ. Meaning even though $d\alpha$ is not identically 0 on ξ, there still might be a certain direction X within ξ such that $\alpha([X,-]) = 0$. This would imply that although ξ might not be completely integrable, there might be some subdistribution of ξ which is integrable. To make sure that ξ is non-integrable and any possible subdistribution is also non-integrable, we have to place the demand

$$d\alpha|_\xi(X,Y) = 0 \text{ for all } Y \in \xi \text{ if and only if } X = 0,$$

in other words $d\alpha$ is non-degenerate when restricted to ξ. This property of ξ being non-integrable, together with any subdistribution may as well be interpreted as being maximally non-integrable.

It is a well-known fact that for any 2-form, non-degeneracy implies it is of maximal rank, so this condition is equivalent to demanding $d\alpha$ is of maximal rank when restricted to ξ. As ξ is of rank $2m$ it follows $(d\alpha)^m$ is nowhere vanishing on ξ. Finally, we also have the corollary from theorem 2.3.8 that $\alpha \wedge d\alpha \neq 0$ as ξ was non-integrable. We combine the above discussion into the following definition.

Definition 4.1.1. A **contact form** on a $(2m+1)$-manifold M is a nowhere vanishing 1-form α with the property that

$$\alpha \wedge (d\alpha)^m \neq 0.$$

The **contact structure** is the resulting maximally non-integrable coorientable distribution $\xi := \ker(\alpha)$. The pair (M, ξ) is called a **contact manifold**, sometimes also denoted (M, α). ▲

Note that any scaling of α by a non-vanishing smooth function $f \in C^\infty(M)$ leaves ξ unchanged, this is motivates denoting only ξ instead of the specific 1-form used to construct it. Also remark that by virtue of $\alpha \wedge (d\alpha)^m$ being a volume form, a contact form α defines an orientation on M and $-\alpha$ defines the opposite orientation. As a corollary of lemma 2.4.4, ξ is necessarily coorientable and orientable. Please take caution that literature often allows the contact form to vanish, in which case the above conclusion do not hold in this generality.

Having defined what contact manifolds are, we are interested in isomorphisms in between them, not surprisingly they are simply diffeomorphisms which preserve the contact structures.

Definition 4.1.2. Given two contact manifolds (M_1, ξ_1) and (M_2, ξ_2) a **contactomorphism** is a diffeomorphism $\phi : M_1 \to M_2$ such that $\phi_* \xi_1 = \xi_2$. ▲

Note the non-degeneracy of $d\alpha$ on a large even-dimensional part of the tangent space, is the reason why contact geometry is often dubbed the odd-dimensional counter-part of symplectic geometry. In fact, to drive home this point, observe the following two lemmas.

Lemma 4.1.3. *If α is a contact form on M, then $d\alpha$ is a symplectic form on ξ.*

Note here we observe $d\alpha$ as a symplectic form of the symplectic vector space ξ_p. Indeed, ξ is by definition maximally non-integrable, so it does not make sense to look for a smooth structure.

Now observe that as α is the only defining form of ξ, it is a rank $2n$ distribution. It follows TM/ξ is a rank 1 distribution. Furthermore, as $d\alpha|_\xi$ is a non-degenerate 2-form, it follows it has 1-dimensional kernel. In fact, its kernel precisely coincides with TM/ξ by property of α being a contact form. Thus, we can define vector fields which precisely lie in $\ker(d\alpha)$. Applying a normalization condition, we obtain a unique vector field.

Definition 4.1.4. The **Reeb vector field** R_α of a contact form α, is the unique vector field R_α satisfying:

(i) $d\alpha(R_\alpha, -) = 0$,

(ii) $\alpha(R_\alpha) = 1$.

▲

There are obviously many parallel vector fields to the Reeb vector field in which the second property is scaled by any non-vanishing scalar function. Each of these also satisfy the following property upon which the Reeb is defined:

$$\langle R_\alpha \rangle = TM/\xi.$$

As a direct consequence of the way it is defined, the Reeb vector field preserves the contact structure in the following sense.

Lemma 4.1.5. *$L_{R_\alpha}(f\alpha) = R_\alpha(f)\alpha$, where $f \in C^\infty(M)$ is non-vanishing.*

Proof. It is a one line calculation to show

$$L_{R_\alpha}(f\alpha) = di_{R_a}(f\alpha) + i_{R_\alpha}d(f\alpha) = R_\alpha(f)\alpha.$$

Note we just scale α, so ξ is fully preserved. ■

There are many similarities and results on contact structures which follow parallel arguments of symplectic structures. If one keeps these in mind, the coming results are recognizable.

Recall that for a symplectic manifold (W, ω) the symplectic form ω was non-degenerate, and hence induced an isomorphism $TM \cong T^*M$ via the flat map

$$\flat_\omega(X) := \omega(X, -).$$

We then used this to construct Hamiltonian vector fields $X_f = \flat_\omega^{-1}(df)$, which also were symplectic vector fields: $L_{X_f}\omega = 0$. We obtain a similar construction in the contact case.

Definition 4.1.6. A **contact vector field** X is a vector field preserving the contact structure. Meaning

$$L_X\alpha = f\alpha,$$

or equivalently; its flow ϕ_X^t is a contactomorphism

$$(\phi_X^t)_*\xi = \xi.$$

▲

We have shown that the Reeb vector field is a very specific contact vector field. Furthermore, we have a subfamily of the contact vector fields which are defined using the following.

Lemma 4.1.7. *Let (M, ξ) be a contact manifold with a contact form α. Then for every smooth function $f \in C^\infty(M)$ we have a unique vector field X_f satisfying*

$$\alpha(X_f) = -f \text{ and } (i_{X_f}d\alpha)|_\xi = (df)|_\xi.$$

Proof. We will construct the isomorphism

$$\begin{aligned} \flat_\alpha : TM &\to T^*M, \\ X &\mapsto i_X d\alpha + i_X\alpha \cdot \alpha. \end{aligned} \tag{4.2}$$

We first construct the splitting $TM = \xi \oplus TM/\xi$, and the corresponding dual splitting. As such any vector field $X \in TM$ can be written as

$$X = Y + gR_\alpha,$$

where $Y \in \xi$ and $gR_\alpha \in TM/\xi$ is some scaling of the Reeb vector field by a non-vanishing function. It follows from non-degeneracy of $d\alpha|_\xi$ that we have an isomorphism

$$\begin{aligned} \psi_1 : \xi &\to \xi^*, \\ Y &\mapsto i_Y d\alpha. \end{aligned}$$

Thus, we can define $\widehat{X_f}$ as the unique vector field satisfying

$$\psi(\widehat{X_f}) = (df)|_\xi.$$

Now any $\theta \in T^*M/\xi^*$ is fully determined by how it acts on R_α, the generator of TM/ξ. Furthermore, it should vanish on ξ. Thus, we can define an isomorphism

$$\begin{aligned}\psi_2 : TM/\xi &\to T^*M/\xi^*,\\ gR_\alpha &\mapsto i_{gR_\alpha}\alpha \cdot \alpha = g\alpha.\end{aligned}$$

It follows from $\psi_1|_{TM/\xi} = 0$ and $\psi_2|_\xi = 0$, that

$$\flat_\alpha = \psi_1 + \psi_2,$$

is an isomorphism. Note that $\flat_\alpha(R_\alpha) = \alpha$. Now define the unique vector field

$$X_f = \widehat{X_f} - fR_\alpha.$$

We remark that for this vector field we readily obtain $\alpha(X_f) = -f$, and

$$(\flat_\alpha(X_f))|_\xi = (i_{\widehat{X_f}}d\alpha)|_\xi = (df)|_\xi.$$

■

The lemma proves the existence of the following definition.

Definition 4.1.8. For a smooth function $f \in C^\infty(M)$, we define the **Hamiltonian contact vector field** as the unique vector field X_f satisfying

$$\alpha(X_f) = -f \text{ and } (df)|_\xi = (\flat_\alpha(X_f))|_\xi.$$

▲

Remark that the Reeb vector field is the contact Hamiltonian vector field for the constant function $h = 1$. Note we use the exact some notation as for a Hamiltonian vector field. If confusion might arise, we will denote the contact Hamiltonian vector field with $\mathcal{X}_h$. A quick check using Cartan's formula leads to the following corollary.

Corollary 4.1.9. *The contact Hamiltonian vector field X_f is a contact vector field which strictly preserves the contact structure*

$$L_{X_f}\alpha = 0.$$

Now note that scaling the contact form with a non-vanishing smooth function $f \in C^\infty(M)$, should preserve the contact structure, but does change the Reeb vector field. Actually in more generality one can wonder how the Reeb vector field transforms under contactomorphisms.

Lemma 4.1.10. *Let $\phi : (M, \xi) \to (M', \xi')$ be a contactomorphism. Denote the contact forms by α and α', and additionally let $\phi^*\alpha' = f\alpha$ for some non-vanishing function $f \in C^\infty(M)$. Then*

$$R_{\alpha'} = \phi_*(X_{-\frac{1}{f}}).$$

Proof. Because ϕ is an isomorphism, there must exist a vector field X such that $\phi_* X = R_{\alpha'}$ Having first to satisfy the normalization principle, it follows

$$1 = \alpha'(\phi_* X) = \phi^* \alpha'(X) = f\alpha(X),$$

which implies that $X = \frac{1}{f} R_\alpha + Y$ where $Y \in \xi$. Secondly it has to satisfy being part of the kernel of $d\alpha'$, so it follows

$$\begin{aligned} 0 &= i_X \phi^* d\alpha', \\ 0 &= i_X d(fa), \\ 0 &= df(X)\alpha - \alpha(X) df + f i_X d\alpha, \\ i_X d\alpha &= \frac{1}{f^2} df - \frac{1}{f} df(X)\alpha. \end{aligned}$$

Lastly contracting with R_α we obtain $f^{-2} R_\alpha(f) = f^{-1} df(X)$. And so X is characterized by

$$\alpha(X) = \frac{1}{f}, \text{ and } i_X d\alpha = \frac{1}{f^2}\left(df - R_\alpha(f)\alpha\right).$$

But this is exactly equal to the unique Hamiltonian contact vector field of $X_{-\frac{1}{f}}$, hence we conclude $X = X_{-\frac{1}{f}}$. ■

Corollary 4.1.11. *Scaling α by a non-vanishing function f transforms the Reeb vector field via $R_{f\alpha} = X_{-\frac{1}{f}}$.*

Example 4.1.11.1. Observe $\mathbb{R}^3$ with the usual (x, y, z) coordinates, and define a 1-form as in the introduction by setting

$$\beta = dz - y dx.$$

This is clearly a nowhere vanishing. Furthermore, $d\beta = -dy \wedge dx$ and so

$$\beta \wedge d\beta = -dz \wedge dy \wedge dx = dx \wedge dy \wedge dz,$$

which is the standard volume form on $\mathbb{R}^3$. Note

$$\xi_p := \ker(\beta_p) = \text{span}(y\partial_z + \partial_x, \partial_y).$$

So the contact structure can be visualized as a collection of planes which angle more and more vertically when travelling along the y-axis, and being invariant when travelling along over the (x, z)-plane see figure 3. Note this is similar to examples 2.1.3.3 and 2.3.4.1.

Calculate the corresponding Reeb vector field setting $R := (R_x, R_y, R_z)$ with respect to the basis ∂_i at each point. Note it should satisfy $d\alpha(R, -) = 0$. As such $-R_x dy + R_y dx = 0$, thus these components must identically be zero and the choice of R_z does not matter. We conclude the Reeb vector field looks like $R_\alpha|_p = (R_z(p)\partial_z)$. Then the normalization condition (ii) sets $R_z = 1$. And hence $R_\alpha|_p = \partial_z$, the constant vector field pointing in the z-direction. ◆

Example 4.1.11.2. We have the so-called **standard contact structure** on $\mathbb{R}^{2n+1}$. Using symplectic coordinates, define

$$\alpha := dz - \sum_{i=1}^{n} y_i dx_i.$$

It follows $\xi_p = \text{span}(y_i\partial_z + \partial_{x_i}, \partial_{y_i})_{1\leq i\leq n}$. The Reeb vector field in this coordinate choice still reduces to ∂_z. It agrees with the twisting planes picture shown in figure 3. Note the Reeb vector fields above precisely agrees with the notion of a contact transformation as we introduced in the example. The Reeb vector field precisely takes contact forms, seen as $\mathbb{R}^{(2m+1)}$ representatives of hyperplanes in $\mathbb{R}^{m+1}$, to contact forms. ♦

Example 4.1.11.3. We can also formulate a radially symmetric contact structure. Let (z, r, θ) be cylindrical coordinates on $\mathbb{R}^3$. Then $\alpha := dz + \frac{1}{2}r^2 d\theta$ is a contact form and $\xi = \text{span}(\partial_r, \frac{1}{2}r^2\partial_z) - \partial_\theta$, apart from being radial, it does not differ much from the standard contact structure.

More interesting is to give a fundamentally different contact structure; the so-called standard overtwisted structure.

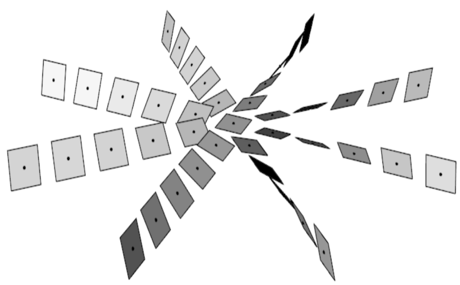

Figure 15: Standard Cylindrical [Mas14]

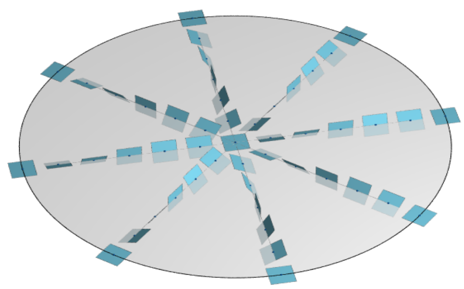

Figure 16: Overtwisted [Mas14]

This is done by defining

$$\lambda := \cos(r)dz + r\sin(r)d\theta,$$

and we obtain the overtwisted contact structure

$$\xi = \text{span}(\partial_r, r\sin(r)\partial_z - \cos(r)\partial_\theta).$$

We will not delve into much detail in what consequences being overtwisted has. However, note that the standard contact structure asymptotically twists towards a vertical plane; it is tight, whereas the overtwisted structure over twists. The formal definition of an overtwisted structure is to find an embedded disc of positive radius which is tangent to the plane fields at the boundary. ♦

Example 4.1.11.4. Observe $\mathbb{R}^m \times \mathbb{RP}^{m-1}$. Geometrically a point in $\mathbb{RP}^{m-1}$ is a line in $\mathbb{R}^m$ passing through the origin. Taking the Cartesian product of this line with $\mathbb{R}^m$ naturally defines a collection of hyperplanes of $\mathbb{R}^m$. Indeed, if coordinates are given by $(p_1, \ldots, p_m, [q_1 : \ldots : q_m])$, then a collection of hyperplanes of $\mathbb{R}^m$ is defined by

$$\sum_{i=1}^{m} q_i dp_i = 0.$$

Now uniform scaling in q_i coordinates is allowed. So define $q_i' := -\frac{q_i}{q_1}$. And we obtain that the kernel of

$$dp_1 + \sum_{i=2}^{m} q_i' dp_i,$$

is the recognizable standard contact structure of $\mathbb{R}^{2m-1}$. ♦

Example 4.1.11.5. The above construction generalizes. Note that a contact structure is given by a non-vanishing contact form $\alpha \in T^*M$. We also remarked that scaling by a non-vanishing function, still defined the same contact structure, so one is more so looking at an equivalence of contact forms up to scaling. We can formalize this by defining the so called projectivized tangent bundle

$$(\mathbb{P}T^*M)_p := (T_p^*M \backslash \{0\}) / \sim,$$

where $\alpha_p \sim \beta_p$ if there is a $\lambda \in \mathbb{R}$ such that $\beta = \lambda\alpha$. It is a fact this defines a vector bundle. Similar to above one restricts the tautological 1-form $(\theta_0)_x : T_x(T^*M) \to \mathbb{R}$, given in local coordinates by

$$\theta_0 = \sum_{i=1}^{m} q_i \pi^* dp_i,$$

where q_i and p_i are defined parallel to the previous example. This defines a contact structure via $\ker(\theta_0)$. ♦

Example 4.1.11.6. We will give the following example which also follows naturally in the theory of Jet spaces. See section A and example A.1.7.3 for a detailed explanation on these concepts.

We have a canonical identification of $J^1(M, \mathbb{R}) \cong T^*M \times \mathbb{R}$ given by

$$j_p^1 f \mapsto (df_p, f(p)).$$

Note $J^1(M, \mathbb{R})$ was itself a smooth manifold, as well as a bundle over M. Hence, admits a tangent bundle $TJ^1(M, \mathbb{R})$, and similarly admits distributions. There is a particular distribution on it called the Cartan distribution, which is defined as follows. Observe holonomic sections of $J^1(M, \mathbb{R})$, which are defined as 1-jet prolongations, see A.1.5, of $f \in C^\infty(M)$ a smooth function

$$j^1 f : M \to J^1(M, \mathbb{R}).$$

A holonomic section gives rise to an embedding of M into $J^1(M, \mathbb{R})$ via its graph

$$\Gamma(f) := \{(df_p, f(p)) \mid p \in M\}.$$

Now if the local coordinates on $J^1(M, \mathbb{R})$ are given via the isomorphism with $T^*M \times \mathbb{R}$ to be $(p_1, \ldots, p_m, q_1, \ldots, q_m, z)$. Then define the Cartan distribution as the kernel of the standard contact structure

$$\xi = \ker(dz - \sum_{i=1}^{m} q_i dp_i).$$

It is clear that $T_x\Gamma(f) \in \xi$ for $x \in \Gamma(f)$ as a local calculation boils down to $df_p - df_p = 0$. Indeed, one can prove that, equivalently, the Cartan distribution can be defined by the span of all tangent spaces to graphs of holonomic sections. ♦

For symplectic structures, one had the very nice Darboux's theorem, which postulates that symplectic manifolds locally all look alike. Meaning that we could always locally find a symplectomorphism to the standard contact structures in Euclidean space. This makes studying symplectic manifolds locally completely different from studying say Riemannian manifolds, whereas the latter can permit many local invariants, such invariants do not exist for symplectic manifolds. Or rather; they do not readily aid one in distinguishing different symplectic manifolds. It turns out the same is true for contact manifolds. In fact, Darboux's theorem for symplectic manifold is a special case of his more general theorem, which applies equally to contact manifolds.

Lemma 4.1.12 (Contact Darboux)**.** *Given two contact manifolds (M_1, ξ_1) and (M_2, ξ_2). Given any two points $p_1 \in M_1$ and $p_2 \in M_2$, then they both respectively have neighbourhoods U_1 and U_2 such that there exists a contactomorphism*

$$\psi : (U_1, \xi_1) \to (U_1, \xi_2).$$

Corollary 4.1.13. *For any point p in a contact manifold (M, ξ), and let $\xi = \ker(\alpha)$. Then there is a chart (U, ψ) around p such that*

$$(\psi^{-1})^*\alpha = dz - \sum_{i=1}^{m} y_i dx_i.$$

One can see this being manifested in the examples given. Indeed, almost every example reduced to finding local coordinates in which we could define the local standard contact structure.

4.1.2 Contact Stability

The introduction of this section alluded to a fact that contact forms naturally arose when looking at level sets of energy functions, and contact transformations were precisely the transformations needed to move from one energy level to the other. In so far only examples of stand-alone contact manifolds have been given. In the examples given, a multitude of references were given that the standard contact form does look alike the tautological 1-form, and one can readily see that a restriction of coordinates gives back the symplectic form. The following section will make more concrete the idea that indeed contact manifolds appear naturally when talking about hypersurfaces in symplectic manifolds.

Recall that the entire construction of Liouville integrability in section 2.5 fundamentally relies on the construction of so-called Hamiltonian vector fields for smooth functions $h : W \to \mathbb{R}$ for a symplectic manifold W. It were these Hamiltonian vector fields, flowing tangent along the level sets of their corresponding, which defined the coordinates on the Liouville tori.

Now recall that Hamiltonian vector fields, by construction, flowed along the level sets of a hamiltonian function, so these cannot be used to flow from one level set to another. This warrants a definition of a different type of vector field which must be able to be transverse to the level sets.

Definition 4.1.14. A vector field X on a symplectic manifold (W, ω) with the property that

$$L_X \omega = \omega,$$

is called a **Liouville vector field**. ▲

Whereas Hamiltonian vector fields left the symplectic form invariant, Liouville vector fields do the opposite. If we denote the flow of a Liouville vector field by ϕ_X^t, then the above property combined with the definition of the Lie derivative and Cartan's formula leads to

$$L_X \omega = \frac{d}{dt}((\phi_X^t)^* \omega) = (\phi_X^t)^*(L_X \omega + \frac{d}{dt}\omega) = (\phi_X^t)^* \omega.$$

Solving this differential equation with initial condition $(\phi_X^0)^* \omega = \omega$, we obtain that X exponentially scales the form:

$$(\phi_X^t)^* \omega = e^t \omega. \tag{4.3}$$

Now whereas a Hamiltonian vector field by construction produced an exact 1-form $i_{X_h} \omega = dh$, taking the interior product with a Liouville vector field must define a different 1-form. This 1-form is non-exact, in fact the opposite; it has the property that it is a primitive of the symplectic form

$$\omega = L_X \omega = d i_X \omega.$$

Now if ω was a symplectic form, it was by definition of maximal rank let us say $(m+1)$, and thus $\omega^{(m+1)} \neq 0$. If it has a primitive λ, then we can write

$$\omega \wedge (d\lambda)^m \neq 0.$$

This form cannot define a volume on hypersurfaces; it is too high of a form. However, one now feels utterly motivated to indeed lay a link between contact forms on hypersurfaces and symplectic forms of the ambient space.

Lemma 4.1.15. *If X is a Liouville vector field on a $(2m+2)$-dimensional symplectic manifold (W, ω). Then for any hypersurface M transverse to X. We have that*

$$\lambda := i_X \omega$$

is a primitive of ω and defines a contact form on M.

Proof. Being a primitive was shown above. Showing it defines a contact form follows from direct algebraic manipulation. For ease of calculations, let us first derive the following identity which applies to any 2-form

$$\begin{aligned} i_X(\omega^{m+1}) &= i_X(\omega \wedge \omega^m) \\ &= i_X\omega \wedge \omega^m + \omega \wedge i_X(\omega^m) \\ &= 2(i_X\omega \wedge \omega^m) + \omega^2 \wedge i_X(\omega^{m-1}) \\ &\vdots \\ &= (m+1)(i_X\omega \wedge \omega^m). \end{aligned}$$

Then applying this to the suspected contact form one obtains

$$\begin{aligned} \lambda \wedge (d\lambda)^m &= i_X\omega \wedge \omega^m \\ &= \frac{1}{m+1} i_X(\omega^{m+1}) \end{aligned}$$

We know $\omega^{m+1} \neq 0$ on W, and so $\lambda \wedge (d\lambda)^m$ is a volume form when restricted to M, as long as M is transverse to X. ■

One might readily believe this to be a very nice way to produce all kinds of contact forms and hypersurfaces of symplectic manifolds. However, we both assumed such a vector field exists, and assumed our hypersurface to be transverse to this. It turns out that these requirements are in fact quite stringent. For example, a globally defined Liouville vector field already implies (W, ω) to be an exact symplectic manifold. And there cannot be any closed exact symplectic manifolds as

$$\int_W \omega^{m+1} = \int_W d(\lambda \wedge \omega^m) = \int_{\partial W} \lambda \wedge \omega^m = 0,$$

which is a contradiction with the non-degeneracy of ω.

In fact, close study of such contact forms and hypersurfaces arose from whether the Hamiltonian vector fields permitted closed orbits. We saw both in our proof of the Arnold-Liouville theorem 2.5.11 and Tischler's theorem 3.6.14, that considerable effort had to be done exactly because, generally, those closed orbits do not exist. In the study of Weinstein's conjecture, contact structures naturally arose as a sufficient condition for those orbits to exist.

Definition 4.1.16. A compact hypersurface $j : M \hookrightarrow (W, \omega)$ is of **contact type**, if there exists a contact form reflecting ω on M. Meaning:

(i) $d\lambda = j^*\omega$,

(ii) $\lambda(\nu) \neq 0$ for $\nu \in \ker(\omega|_M)\backslash\{0\}$.

▲

Remark the second requirement does make sense: as M is odd dimensional by virtue of being a hypersurface of an even dimensional space. It follows $\omega|_M$, a 2-form, must have a non-trivial kernel when restricted. This kernel defines a line bundle on M. As this construction is canonical, this is often called the characteristic line bundle of M. Furthermore, it guarantees that $\ker(\lambda) \cap \ker(\omega|_M) = \varnothing$, as such there is a splitting

$$TM = \ker(\lambda) \oplus \ker(\omega|_M).$$

It follows $\ker(\lambda)$ defines pointwise a $(2m)$-dimensional subspace of T_xM on which $\omega|_M = d\lambda$ is non-degenerate by the splitting. It follows readily that $\lambda \wedge (d\lambda)^n \neq 0$ on M. Thus, λ is a contact form on M with contact structure $\ker(\lambda)$.

Now not only does the existence of such a transverse Liouville vector field produce a contact form. The converse is also true: if M is of contact type, it has such a transverse Liouville vector field. In order to prove this equivalence of definitions we first introduce a poincaré-esque lemma such that closed forms become exact in a neighbourhood of our hypersurface.

Lemma 4.1.17. *Let $r : U \to M$ be a tubular neighbourhood of a submanifold M. Let $\alpha \in \Omega^k(U)$ such that $d\alpha = 0$ and $j^*\alpha = 0$, where $j : M \hookrightarrow U$ is the inclusion. Then there exists $\beta \in \Omega^{k-1}(U)$ such that $d\beta = \alpha$ and $j^*\beta = 0$.*

Proof. Let U be a tubular neighbourhood of M in W. Then U strongly deformation retracts onto M. Denote by $h_t : U \to U$ this strong deformation retract. Meaning h_t is a smooth family of diffeomorphisms satisfying $h_0 = id_U$, $h_1 = j \circ \pi$, and $h_t \circ j = j$. Define a time-dependent vector field X_t on U whose flow is exactly generated by these diffeomorphisms: $X_t \circ h_t := \frac{d}{dt} h_t$. Now by definition of flow, using Cartan, and closedness of α, we obtain:

$$L_{X_t}\alpha = \frac{d}{dt} h_t^*\alpha = h_t^*(L_{X_t}\alpha + \frac{d}{dt}\alpha) = h_t^*(di_{X_t}\alpha).$$

Also observe that

$$\lim_{t\to 1} h_t^*\alpha = (j \circ \pi)^*\alpha = \pi^* j^*\alpha = 0,$$

where the last equality is by hypothesis. And so:

$$\begin{aligned}\alpha &= h_0^*\alpha - \lim_{t\to 1} h_t^*\alpha \\ &= \lim_{t\to 1}\int_t^0 \frac{d}{ds}\bigg|_{s=t} h_s^*\alpha \, ds \\ &= \int_1^0 h_s^*(di_{X_s}\alpha)\, ds \\ &= d\int_1^0 i_{X_t}\alpha\, ds.\end{aligned}$$

And so setting $\beta = \int_1^0 i_{X_t}\alpha\, dt$ has the desired properties: $d\beta = \alpha$ and $j^*\beta = 0$, as $X_t|_M = 0$. ∎

Now the point of this lemma is to be able to extend the contact form on a hypersurface to a 1-form on a neighbourhood of the hypersurface which restricts to the contact form.

Corollary 4.1.18. *Let $(M, \lambda) \subset (W, \omega)$ be of contact type, then there exists a 1-form β on a neighbourhood U of M such that:*

1. $(d\beta)|_U = \omega|_U$
2. $j^*\beta = \lambda$

Proof. Again define a tubular neighbourhood $r : U \to M$ with inclusion $j : M \hookrightarrow U$. Now let us define a 1-form $\mu := r^*\lambda$, and a 2-form $\kappa := \omega - d\mu$. We note the hypotheses of the previous lemma are satisfied: $d\kappa = 0$ and

$$j^*\kappa = j^*\omega - j^*d\mu = d\lambda - d((r\circ j)^*\lambda) = d\lambda - d\lambda = 0.$$

The previous lemma implies the existence of a $\theta \in \Omega^1(U)$ such that $d\theta = \kappa$ and $j^*\theta = 0$. Now as a final step define $\beta = \theta + \mu$. This new form has the properties

$$d\beta = d\theta + d\mu = \omega - d\mu + d\mu = \omega,$$

so is the symplectic form on a neighbourhood of our hypersurface. And furthermore

$$j^*\beta = j^*\theta + j^*\mu = \lambda,$$

so it restricts to our contact form on M. ∎

Now we are ready to prove the equivalence between being of contact type and having a transverse Liouville vector field:

Theorem 4.1.19 ([HZ12, Hofer & Zehnder]).
Let M be a compact hypersurface of a $(2m+2)$-dimensional symplectic manifold (W, ω), with $j : M \hookrightarrow W$ the usual inclusion. Then the following are equivalent:

(a) M is of contact type; there exists $\lambda \in \Omega^1(M)$ such that:
   (i) $d\lambda = j^*\omega$,
   (ii) $\lambda(\nu) \neq 0$ for $\nu \in \ker(\omega|_M)\backslash\{0\}$.

(b) There exists a transverse Liouville vector field X defined in a neighbourhood U of M;
   (i) $L_X\omega = \omega$,
   (ii) $X_x \notin T_xM$ for all $x \in M$.

Proof. Start by assuming M is of contact type. Then take β as proven to exist in the previous lemma. Via the usual isomorphism of the flat map $\flat_\omega$, define a vector field X_β on U via $\beta = i_{X_\beta}\omega$. It follows this vector field is Liouville:

$$L_{X_\beta}\omega = di_{X_\beta}\omega = d\beta = \omega.$$

Furthermore, X_β is transverse to the hypersurface M. Take a non-zero $\nu \in \ker(\omega|_M)$. Then note

$$0 \neq \lambda(\nu) = j^*\beta(\nu) = \beta(\nu) = \omega(X_\beta, \nu).$$

However, $\nu \in \ker(\omega|_M)$, and so it must follow $X_\beta \notin TM$.

Conversely, assume M has such a transverse Liouville vector field X. Then define

$$\lambda := j^*(i_X\omega).$$

It follows λ is a primitive of ω on M:

$$d\lambda = dj^*(i_X\omega) = j^*L_X\omega = j^*\omega.$$

Lastly, let $\nu \in \ker(\omega|_M)\backslash\{0\}$. Then it follows $\omega(\nu, X) \neq 0$ as $X \notin TM$ by assumption. By antisymmetry

$$0 \neq \omega(X, \nu) = (j^*i_X\omega)(\nu) = \lambda(\nu),$$

and so the kernel of λ is disjoint from the characteristic line bundle on M. ■

Before, it was claimed that it is unfortunate that a Liouville vector field might not exist, and so now, equivalently, hypersurfaces of contact type. Even though one still cannot necessarily conjure up hypersurfaces of contact type to their liking, theorem 4.1.19 does lay bare an important property of hypersurfaces of contact type. Namely, if given a hypersurface of contact type, then we can readily produce a whole family of hypersurfaces of contact type, using the transverse Liouville vector field X which will always exist.

Indeed, recall that for a smooth vector field X on M, its flow $\phi : \mathbb{R} \times M \to M$ is locally a diffeomorphism. Meaning that

$$\phi : (-\varepsilon, \varepsilon) \times M \to M,$$

defines a diffeomorphism for a small enough ε. As hinted before, we use the existence of the Liouville vector field, to flow from one energy hypersurface to another, and the above fact shows that, locally, these energy hypersurfaces are diffeomorphic. They are even more so than diffeomorphic. Let $\nu \in \ker(\omega|_M)$. Note $j^* i_\nu \omega$ evaluates identically to 0, but via (4.3) we can pull back, and hence scale, along the flow. This results in

$$j^*(i_\nu \omega) = j^*(i_\nu e^t \omega) = j^*(i_\nu((\phi^t)^*\omega)) = j^*(i_{\phi^t_* \nu}\omega \circ \phi^t_*) = 0.$$

And so

$$\phi^t_* \nu \in \ker\left[((\phi^t)^*\omega)|_{\phi^t(M)}\right] = \ker\left[(\omega)|_{\phi^t(M)}\right],$$

as scaling by a non-vanishing function leaves the kernel invariant. Thus, one not only has that the hypersurfaces are diffeomorphic, there is also an isomorphism of characteristic line bundles

$$\phi^t_* : \ker[\omega|_M] \to \ker\left[((\phi^t)^*\omega)|_{\phi^t(M)}\right].$$

This prompts the following definition.

Definition 4.1.20. A compact hypersurface $M \subset (W, \omega)$ is **stable** if there is a family of diffeomorphisms isotopic to the identity

$$\psi : M \times I \to W,$$

where $I \subset \mathbb{R}$. Such that $\psi^t(M) \cap \psi^s(M) = \varnothing$ if $s \neq t$. Denote $M_t := \psi^t(M)$ the resulting family of hypersurfaces. Furthermore, there need be induced family of bundle isomorphisms

$$\psi^t_* : \ker(\omega|_M) \to \ker(\omega|_{M_t}).$$

▲

Now all the lemmas of the above result in us setting.

Theorem 4.1.21.
A hypersurface of contact type $(M, \lambda) \subset (W, \omega)$, is stable. Each $(M_t, \psi^t_* \lambda)$ is a hypersurface of contact type.

Though historically speaking contact transformations, forms and geometry first arose in the context of partial differential equations, hypersurfaces of contact type especially gained prominence within the field of geometry in the scope of Weinstein's conjecture.

Conjecture 4.1.22 (Weinstein). *A hypersurface M of contact type, having trivial first singular cohomology $H^1(M) = 0$, admits a Hamiltonian vector field with a closed orbit.*

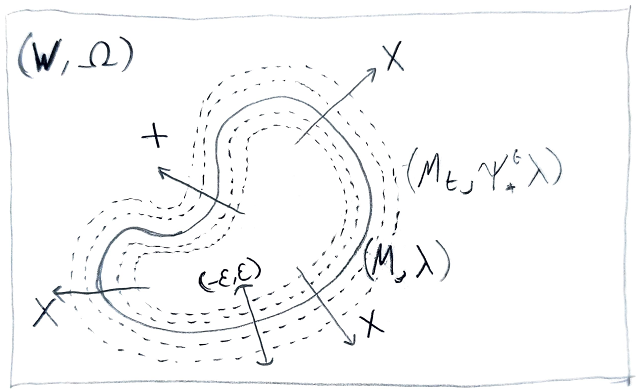

Figure 17: Stable hypersurface of contact type

Which was proven to be true if the hypersurface was stable, in particular it follows it is true for hypersurfaces of contact type. In fact, the nearby hypersurfaces have the same dynamical properties as the original hypersurface of contact type. It allows us to relax the Weinstein conjecture; if we are able to find a closed orbit on a nearby the hypersurface of contact type, we can pull it back to the original hypersurface. Moreover, the existence of a closed orbit is clearly some property dependent on the contact form, not necessarily of the ambient symplectic manifold. It follows there is a more true to nature reformulation of Weinstein's conjecture.

Conjecture 4.1.23 (Contact Weinstein)**.** *For every closed odd-dimensional manifold M with contact form λ, its Reeb vector field has a periodic orbit.*

Example 4.1.23.1. Let (M, ξ) be a contact manifold, with contact form α. And consider the symplectic manifold $(M \times \mathbb{R}, -d(e^t\alpha))$, where t is the coordinate on $\mathbb{R}$. Remark that $-d(e^t\alpha)$ is indeed a non-degenerate, closed 2-form. Consider the vector field ∂_t transverse to each hypersurface. And note

$$L_{\partial_t}\omega = -di_{\partial_t}(e^t dt \wedge \alpha + e^t d\alpha) = -d(e^t\alpha),$$

so ∂_t is a transverse Liouville vector field. We conclude that each $M \times \{t\}$ is a hypersurface of contact type in $(M \times \mathbb{R}, -d(e^t\alpha))$. ♦

Example 4.1.23.2. Consider $D^2 := \{(r, \theta) \in \mathbb{R}^2 \mid 0 \leq r \leq 1, 0 < \theta \leq 2\pi\}$ with standard symplectic form $\omega_0 = dr \wedge d\theta$. We can define the vector field $X := r\partial_r$. Remark

$$L_X\omega_0 = dr \wedge d\theta,$$

so it is indeed Liouville. Now observe $\partial D^2 = S^1$ the boundary circle and j the usual inclusion. As X is a Liouville vector field defined in a neighbourhood of S^1, we readily see that

$$(S^1, j^*(i_X\omega_0)) = ((S^1, d\theta))$$

is a contact manifold. ⧫

Both examples are specific cases of a more general construction, the first example goes by the name of symplectization, which will be expanded upon in more generality further in this paper. The second is an example of the following definition.

Definition 4.1.24. If (W, ω) is a compact symplectic manifold with non-empty boundary ∂W, and its boundary is a hypersurface of contact type with outwards pointing transverse Liouville vector field. Then (W, ω) is a symplectic manifold with **convex contact type boundary**. ▲

4.2 Hamiltonian Structures

One might remark the subtlety in the statement following the Weinstein conjecture 4.1.22. Where we posed that the Weinstein conjecture was true in the particular case of hypersurfaces of contact type, because these were stable. This raises the question: to what extent is this stability related to being precisely of contact type, or is there some sort of generality in play?

Consider example 4.1.23.1, where we were given a contact manifold and produced an ambient symplectic manifold in which it was stable. The fact that it was a symplectic manifold was not exploited to the fullest. It was mostly the fact that it was a closed 2-form, which was non-degenerate simply in a neighbourhood of the hypersurface. This was done in order to, at least locally, define the transverse Liouville vector field, which implies the contact type via theorem 4.1.19. In this example the Liouville vector field coincidentally exists on the whole of the ambient manifold. Nowhere was it truly necessary for it to be non-degenerate on the whole of the ambient manifold. In fact, we have already shown that demanding such a transverse Liouville vector field to exist on the entirety of the manifold, was a very stringent condition which, for example, already excluded all closed symplectic manifolds.

Also observe the following interesting property of the example. Because the Liouville vector field was defined on the whole of the ambient symplectic manifold, we in fact obtain a codimension one foliation by hypersurfaces of contact type. Furthermore, an integral curve of this vector field intersects each leaf transversally. If we had some way to close up this manifold in such a way that the Liouville vector field has closed orbits, we exactly obtain a taut foliation.

This is fascinating; whereas we presented contact structures as a dichotomy to integrability, hence foliations, there seems to be some readily made connection between contact structures and the very specific type of taut foliations defined by a closed 1-form. In fact, observe that although the conceptuality is completely different, in dimension three, both are captured by observing $\lambda \in \Omega^1(M)$ such that

$$\lambda \wedge (d\lambda) \geq 0,$$

giving rise to so-called positive confoliatons [ET98].

The motivation behind Hamiltonian structures, and their stability, rest on these two observations: first; that in example 4.1.23.1 the whole manifold need not be symplectic, owing to the fact that the stability need only to be local, and second; that there seems to be some relation between contact structures and taut foliations defined by a closed 1-form.

4.2.1 Defining Hamiltonian Structures

Let us first define what Hamiltonian structures are, which, like contact structures, only reside on odd-dimensional manifolds.

Definition 4.2.1. A **Hamiltonian structure** on an oriented $(2m+1)$-manifold M is a closed 2-form ω such that

$$\omega^m \neq 0.$$

By dimensionality this form is degenerate. Its kernel defines a rank one distribution called the **Hamiltonian line field**

$$\xi := \ker(\omega).$$

As M is oriented, we orient ξ via the orientation given by ω^m on the transversals to ξ. ▲

Remark that scaling ω by any non-vanishing function leaves the Hamiltonian structure intact. Similar to contact structures, Hamiltonian structures are referred to by (M, ξ) instead of (M, ω). But the latter is common when one wants to fix a defining form. Hamiltonian structures naturally generalize symplectic structures in the context of hypersurfaces.

Example 4.2.1.1. Consider a symplectic $(2m+2)$- manifold (W, Ω) and a closed hypersurface $M \subset W$. By property of being symplectic $\Omega^{(m+1)} \neq 0$. It follows that $(\Omega|_M)^m \neq 0$. Thus, we observe that $(M, \omega|_M)$ is a Hamiltonian structure. ♦

Now recall that in the context of Weinstein's conjecture and theorem 4.1.21, we were interested in when such hypersurfaces were in a sense stable.

Definition 4.2.2. A **stabilizing 1-form** for the Hamiltonian structure ω is a $\lambda \in \Omega^1(M)$ satisfying:

(i) $\lambda \wedge \omega^m \neq 0$.

(ii) $\xi \subset \ker(d\lambda)$.

In this case ω is **stabilizable**. We call the pair (ω, λ) a **stable Hamiltonian structure**. ▲

Stable Hamiltonian structures on their turn are a generalization of contact structures.

Example 4.2.2.1. Observe a contact manifold (M, ξ) with defining 1-form α. Note that setting $\omega := \pm d\alpha$ both define a closed 2-form such that

$$(\pm d\alpha)^m \neq 0,$$

thus defines a Hamiltonian structure. Furthermore, the stabilizing 1-form is precisely α, as by hypothesis

$$\alpha \wedge (\pm d\alpha)^m \neq 0.$$

Hence, it is a stable Hamiltonian structure. ♦

Note the first condition in the definition says that we have a non-vanishing volume form on M. In particular, it implies that λ cannot vanish on ξ. As before, as ξ is 1 dimensional, we can canonically define a vector field generating it. Similar to the contact case one might also define the Reeb vector field for a stable Hamiltonian structure.

Definition 4.2.3. The **Reeb vector field** $R_{(\omega,\lambda)}$ of a stable Hamiltonian (ω, λ) is defined as the unique vector field satisfying:

(i) $\lambda(R_{(\omega,\lambda)}) = 1$.

(ii) $i_{R_{(\omega,\lambda)}}\omega = 0$.

▲

From now on we will denote both Reeb vector fields by just R and assume from context whether it refers to the contact or stable Hamiltonian case.

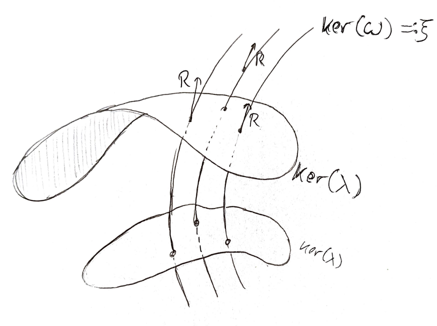

Figure 18: A local picture of a stable Hamiltonian structure

Example 4.2.3.1. Let (W, Ω) be a closed symplectic manifold of dimension $2m$. Construct the following fibre bundle M over S^1 with the manifold as fibres

$$M := \pi : W \times S^1 \to S^1.$$

We can define a Hamiltonian structure on M. Recall that $(S^1, d\theta)$ defined a contact manifold, so observe $\lambda := \pi^*(d\theta)$, and define

$$\omega := \Omega - (\pi^*(d\theta)) = \Omega - \lambda.$$

Then note

$$\lambda \wedge \omega^m = \lambda \wedge \Omega^m \neq 0.$$

Furthermore, ξ is span by the vector field corresponding to $\pi^*(\partial_\theta)$, moreover λ is in fact closed, so $\xi \subset \ker(d\lambda)$ is satisfied trivially.

In fact, note that $\ker(\lambda)$ defines a codimension one foliation of M via the corollary of Frobenius' theorem 2.3.8, given by the hypersurfaces of contact type $\Omega \times \{\theta\}$. ♦

Example 4.2.3.2. The above construction generalizes to the following. Let (W, Ω) be symplectic manifold, and let $\phi : W \to W$ be a symplectomorphism. Define the mapping torus as usual

$$W_\phi := \frac{[0,1] \times W}{(0,x) \sim (1, \phi(x))}.$$

The Hamiltonian structure is derived from ω, the stabilizing one form is again given by $d\theta$ and the Hamiltonian line field by ∂_θ. ♦

Note that the above two example reveal the relationship between taut foliations and stable Hamiltonian structures. If we reduce the dimension of the symplectic manifold to two, hence of the ambient manifold to three, then $\ker(\lambda)$ in fact defines a taut foliation in both cases. In fact, reducing the dimension to three will be explored further in section 5.

4.2.2 Symplectization and Stability

Recall that our segue to stable Hamiltonian structures was prefaced by example 4.1.23.1. Here we demonstrated a process which is sometimes called the extrinsic symplectization of a contact manifold. However, we have hypothesized that perhaps we can fine-tune our construction to product a smaller ambient manifold. Via a process called symplectization we can readily embed any stable Hamiltonian structure in a symplectic manifold $i : M \hookrightarrow W$, such that

$$i^*\Omega = \omega.$$

Locally this symplectic manifold is foliated by hypersurfaces of contact type.

Definition 4.2.4. The **symplectization** of a stable Hamiltonian structure (ω, λ) on M, is a symplectic manifold (W, Ω) defined as:

(i) $W := (-\varepsilon, \varepsilon) \times M$;

(ii) $\Omega := \omega + d(t\lambda)$,

where t is the coordinate in $(-\varepsilon, \varepsilon)$ and $\varepsilon \in \mathbb{R}$ is chosen small enough. ▲

Lemma 4.2.5. *The symplectization (W, Ω) always exists.*

Proof. It is easy to see $d\Omega = 0$, hence a closed form. Furthermore, $i^*\Omega = \omega$ by construction, which can easily be checked. Thus, we will only prove non-degeneracy.

Recall that a Hamiltonian structure gave rise to the Hamiltonian line field $\xi := \ker(\omega)$ and that we had the volume form $\lambda \wedge \omega^n$, as a result we can decompose $TM = \ker(\omega) \oplus \ker(\lambda)$, which gives rise to the following decomposition of the symplectic manifold W

$$TW = \langle \partial_t \rangle \oplus \ker(\omega) \oplus \ker(\lambda).$$

Furthermore, expanding Ω into

$$\Omega = \omega + dt \wedge \lambda + td\lambda,$$

we can write Ω in block-matrix form with respect to this decomposition to gain insight in how it behaves:

$$\begin{pmatrix} 0 & \lambda & \dots \\ -\lambda & 0 & \dots \\ \vdots & \vdots & \omega + td\lambda \end{pmatrix},$$

where the dots indicate there only to be zeroes. We see this matrix is of maximal rank, hence Ω non-degenerate, if and only if $\omega + td\lambda$ is of maximal rank for all $t \in (-\varepsilon, \varepsilon)$. Assume that there we have such a t_0 such that $\omega + t_0 d\lambda$ is degenerate. Then simply take $0 < \varepsilon < |t_0|$, making Ω non-degenerate.

Hence, we conclude if ε is chosen small enough (W, Ω) is a symplectic manifold. ■

This precisely makes concrete what was meant before; we need not define an entire non-compact symplectic manifold, together with a globally non-vanishing Liouville vector field, to embed contact manifolds in a symplectic manifold.

Now recall that theorems 4.1.19 and 4.1.21 showed that hypersurfaces of contact type were stable, and this stability was the link with Weinstein's conjecture. We have not yet shown such a stability for stable Hamiltonian structures. Indeed, the moniker stable better be well-deserved. This stability is proven using the process of symplectization just outlined.

Theorem 4.2.6.
Given a closed hypersurface M of a $(2m+2)$-dimensional symplectic manifold (X, Ω) the following are equivalent:

1. *Let $\omega := \Omega|_M$, then (M, ω) is a stabilizable Hamiltonian structure.*
2. *There is a tubular neighbourhood $r : U \to M$ and a family of hypersurfaces $\{M_r\} \subset U$ diffeomorphic to M via a family of diffeomorphisms ψ_r. Which restrict to diffeomorphisms*
$$(\psi_r)_* : \ker(\Omega|_M) \to \ker(\Omega|_{M_r}).$$
3. *There exists a vector field T transverse to M such that*
$$\ker(\Omega|_M) \subset \ker(L_T \Omega|_M).$$

Proof.

($1 \implies 2$): Assume λ is the stabilizing 1-form of (M, ω) and construct the symplectization similar to done previously:
$$\Omega' := \omega + d(t\lambda),$$
which is symplectic on $W' := (-\varepsilon, \varepsilon) \times M \subset W$ for ε small enough.

Let T be the vector field ∂_t, where t is the coordinate of $(-\varepsilon, \varepsilon)$. Applying Cartan's formula we get
$$L_T \Omega' = d i_T \Omega' = d\lambda.$$

Now using the flow of the vector field T define the diffeomorphisms $\phi_\tau : W' \to W'$ where $(t, x) \mapsto (t + \tau, x)$. Thus, one can define a family of hypersurfaces modelled on M:
$$M'_r := \phi_r(\{0\} \times M).$$
Denote by $\Omega'_r := \Omega'|_{M_r}$. Note that
$$\ker(\Omega'_r) = \ker(\omega + r d\lambda) \supset \ker(\omega) = \ker(\Omega'_0).$$

Note the inclusion follows from the fact that ε was chosen small enough such that $\omega + t d\lambda$ was non-degenerate on $\ker(\lambda)$ for all t. Thus, we see that the kernel distribution is independent of the t-coordinate. Now finally remark that $\omega := \Omega|_M = \Omega'|_{\{0\} \times M}$ and so (M, ω) coisotropically embeds in both (W, Ω) and (W', Ω'). By the classification of coistropic embeddings it follows that there exists open neighbourhoods $U \subset W$ and $U' \subset W'$ of M which are symplectomorphic. Via this symplectomorphism the family $\{M'_r\}$ gives rise to the family $\{M_r\}$.

($2 \implies 3$): Define T_r as the time-dependent vector field whose flow is exactly generated by the diffeomorphisms ψ_r. It is clear that T_r is transverse to M. By assumption the kernel foliations are constant in t. So we see obtain
$$L_{T_r} \Omega|_M = \frac{d}{dr}\bigg|_{r=0} \psi_r^* \Omega|_M = \frac{d}{dr}\bigg|_{r=0} \Omega|_M((\psi_r)_*, (\psi_r)_*),$$

and so
$$\ker(\Omega|_M) \subset \ker(L_{T_r}\Omega|_M).$$

($3 \implies 1$): Setting $\lambda := i_T\Omega|_M$ we immediately see:
$$\ker(d\lambda) = \ker(di_T\Omega|_M) = \ker(L_T\Omega_M) \supset \ker(\Omega|_M) = \ker(\omega),$$
and
$$i_T\Omega|_M \wedge (\Omega|_M)^n \neq 0.$$
∎

One can wonder how continuous or smooth this construction is. Intuitively one would expect this construction to be fairly smooth, as it is some linear combination of 2-forms and 1-forms. However, let us formalize this statement within the framework of jets. We refer to section A for a detailed explanation.

Lemma 4.2.7. *Let (M, ω, λ) a stable Hamiltonian structure, and construct its symplectization (X, Ω) as above. The mapping*
$$J^0\left(\bigwedge^2 T^*M\right) \times J^1\left(T^*M\right) \to J^1\left(\bigwedge^2 T^*X\right),$$
$$(\omega, \lambda) \mapsto \omega + d(t\lambda) =: \Omega,$$
is continuous.

Proof. We have already shown this is a well-defined mapping
$$\Omega^2(M) \times \Omega^1(M) \to \Omega^2(X),$$
so we are left to proof continuity. By hypothesis choose a stabilizable ω' such that $\langle j^0\omega, j^0\omega'\rangle < \delta$ for some metric $\langle\cdot,\cdot\rangle$ compatible with the topology, and whose stabilizing 1-form λ' is such that $\langle j^1\lambda, j^1\lambda'\rangle < \delta'$. If we return to our local story, this is interpreted as ω being C^0 close to ω' at all points in M, and both λ, λ' and $d\lambda, d\lambda'$ being C^0 close at all points in M. It is clear from
$$\begin{aligned}\Omega - \Omega' &= \omega + dt \wedge \lambda + td\lambda - \omega' - dt \wedge \lambda' - td\lambda' \\ &= (\omega - \omega') + dt \wedge (\lambda - \lambda') + t(d\lambda - d\lambda'),\end{aligned}$$
that Ω and Ω' are C^0 close. Now for free, by closedness of ω, ω', we get that as $d\Omega = d\omega = 0$ and $d\Omega' = d\omega' = 0$ and so Ω and Ω' are in fact even C^1-close, proving the lemma. ∎

4.2.3 Stability and Geodesibility

As we have defined it now, stabilizability seems to be a property of a 2-form ω. However, note that in this definition it only really depended on the interaction of λ with the one dimensional kernel foliation $\xi := \ker(\omega)$. It did not depend heavily on the particularities of ω itself. Indeed, we were free to scale ω with any non-vanishing function. As such we can define stabilizability as a property of 1-dimensional foliations rather than of differential forms.

Definition 4.2.8. An orientable one dimensional foliation $\mathcal{F}$ of a smooth $(2n+1)$-dimensional manifold M is said to be **stabilizable** if there exists a 1-form λ on M such that:

$$\lambda|_{\mathcal{F}} = c \neq 0 \text{ and } (d\lambda)|_{\mathcal{F}} = 0.$$

We often take a specific vector field X generating $\mathcal{F}$ such that $\lambda(X) = 1$ and $i_X d\lambda = 0$. ▲

As before with the 1-dimensional kernel distribution, this results in a splitting of the tangent space

$$TM = \ker(\lambda) \oplus \langle X \rangle.$$

If such a foliation is stabilizable, then λ is constant along the integral curves of X. As we are dealing with smooth manifolds, we can always construct a Riemannian metric g on M, and thus we can introduce the notion of orthogonality as a more specific embodiment of being transverse. Now we can wonder if it is possible to construct a metric g with respect to which we have

$$\ker(\lambda) \perp \langle X \rangle, \quad g(X, X) = 1, \text{ and } \nabla_X(X) = 0.$$

In other words, X is orthogonal, so in particular transverse, to the hypersurfaces defined by $\ker(\lambda)$ and its flow lines are normalized geodesics.

Definition 4.2.9. An orientable one dimensional foliation $\mathcal{F}$ is **geodesible** if there exists a metric g such that a vector field X generating $\mathcal{F}$ has geodesic integral curves. ▲

Remark the similarities to both a stable Hamiltonian structure and a stable contact structure. Both were stable if there existed some transverse vector field under which the structure was in a sense preserved. Note this problem is also closely related to Frobenius' theorem. If g is a metric, we can define a 1-form via

$$\mu := i_X g,$$

the kernel of which is then precisely orthogonal to X. This defines a hypersurface if and only if $\mu \wedge d\mu = 0$ according to theorem 2.3.8.

The following lemma shows that being geodesible is an equivalent notion of being stabilizable.

Lemma 4.2.10. *An orientable one dimensional foliation $\mathcal{F}$ is stabilizable if and only if it is geodesible.*

Proof. First assume $\mathcal{F}$ is geodesible; we are given a g such that $g(X, X) = 1$ and $\nabla_X X = 0$. Then define the 1-form $\mu := i_X g$, we will show this is the stabilizing 1-form. It follows immediately that

$$\mu(X) = g(X, X) = 1.$$

Now pick another vector field Y, we calculate

$$\begin{aligned} d\mu(X,Y) &= X(\mu(Y)) - Y(\mu(X)) - \mu([X,Y]) \\ &= X(g(X,Y)) - Y(g(X,X)) - g(X,[X,Y]) \\ &= g(X, \nabla_Y X) \\ &= \frac{1}{2} Y(g(X,X)) \\ &= 0. \end{aligned} \tag{4.4}$$

And so $i_X d\mu = 0$, hence μ is a stabilizing one form of $\mathcal{F}$.

Conversely, assume $\mathcal{F}$ is stabilizable; there exists a λ such that $\lambda(X) = 1$ and $i_X d\lambda = 0$. Then define g such that $g(X, X) = 1$ and $\ker(\lambda) \perp X$, this can be done freely by property of Riemannian metrics on smooth, finite dimensional, manifolds. As $\ker(\lambda) \perp X$, we readily see that

$$(i_X g)(Y) = \begin{cases} 1 \text{ if } Y = X, \\ 0 \text{ if } Y \in \ker(\lambda), \end{cases}$$

and so $i_X g = \lambda$. Now similar to equation (4.4) we obtain

$$0 = (i_X d\lambda)(Y) = g(\nabla_X X, Y),$$

by non-degeneracy of g it follows $\nabla_X X = 0$. Thus, $\mathcal{F}$ is geodesible. ∎

4.2.4 Obstructions to Stability

We have now defined these different ways to characterize stable Hamiltonian structures. And also demonstrated that being a stable Hamiltonian structure is more lenient than being a contact structure. However, let us explore some scenarios where there might be some obstruction for a Hamiltonian structure to be stabilizable.

Let us recall a visualization of being stabilizable. An orientable $\mathcal{F}$ was stabilizable if there exists a λ which evaluates positively along the foliation, whose differential $d\lambda$ evaluates to zero on planes tangent to the foliation. The following lemma provides an obstruction to this.

Lemma 4.2.11. *An orientable one dimensional foliation $\mathcal{F}$ is non-stabilizable if there exists a closed leaf which can be infinitely well approximated by the boundaries of singular 2-chains tangent to $\mathcal{F}$.*

Proof. This is just a straightforward application of Stokes' theorem. Denote by γ the closed leaf, and let c_n be such a sequence of 2-chains whose boundaries converge to γ. Assume on the contrary $\mathcal{F}$ is stabilizable by a 1-form λ, then it follows

$$0 = \int_{c_n} d\lambda = \int_{\partial c_n} \lambda \text{ which converges to } \int_\gamma \lambda.$$

This implies $\lambda|_\mathcal{F}$ must be zero somewhere, a contradiction with the definition on being stabilizable. ∎

This may seem as a rather artificial construction, however; it readily arises when dealing with Reeb components of foliations. Recall the definition of a Reeb foliation 3.2.5 and a Reeb component of a foliation 3.7.4. When $n = 2$ we spoke of the Reeb annulus. In this case we have a properly embedded $S^1 \times D^1$ with a Reeb foliation, such that the boundary leaf S^1 is tangent to $\mathcal{F}$. Similar to how foliations with a Reeb component could not be taut, they can also not be stabilizable, further hinting at the deeper connection of taut foliations and stable Hamiltonian structures.

Corollary 4.2.12. *Any oriented one dimensional foliation $\mathcal{F}$ containing a Reeb component is non-geodesible*

Proof. Denote by A the annulus which is the Reeb component of $\mathcal{F}$. We see

$$0 = \int_A d\lambda = \int_{\partial A} \lambda,$$

which implies $\lambda|_\mathcal{F} = 0$ somewhere. ∎

Recall how there could not be any closed exact symplectic manifolds, we have a similar result for closed exact stable Hamiltonian structures as a consequence of Stokes' theorem.

Lemma 4.2.13. *Let ω an exact Hamiltonian structure on a closed manifold M. If its primitive α satisfies*

$$\alpha \wedge (d\alpha)^m = 0,$$

then ω is not stabilizable.

Proof. Assume on the contrary ω is stabilizable by λ. Then it follows:

$$\begin{aligned} 0 &= \int_{\partial M} \alpha \wedge \lambda \wedge (d\alpha)^{m-1} \\ &= \int_M d(\alpha \wedge \lambda \wedge (d\alpha)^{m-1}) \\ &= \int_M (\lambda \wedge (d\alpha)^m) - \int_M (\alpha \wedge d\lambda \wedge (da)^{m-1}). \end{aligned}$$

Note the first term is exactly the volume form defined by a stable Hamiltonian structure, and hence is non-vanishing. For the second term it follows that as we can decompose the space $TM = \ker(\omega) \oplus \ker(\lambda)$ we have for $X \in \ker(\omega)$ that $i_X(\omega) = 0$, $i_X(d\lambda) = 0$, and lastly $i_X\alpha = 0$ by vanishing of $\alpha \wedge (d\alpha)^m$ everywhere. We arrive at a contradiction. ■

5 The Three Dimensional Case

So far we have treated stable Hamiltonian structures in the utmost generality regarding dimensions. We will from now focus on the low-dimensional case. Remark that the basic definition of a stable Hamiltonian structure

$$\lambda \wedge \omega^m \neq 0$$

puts a lowest bound on the dimension. If $m = 1$, we require a 3-form to be non-vanishing, which is only possible if the space it resides in is at least three-dimensional.

We remind ourselves that one of the narratives one can employ to arrive at stable Hamiltonian structures, and the narrative we have been using in this paper, is via the route of Weinstein's conjecture and stable hypersurfaces of contact type. It is in fact so that Weinstein's conjecture is still an open problem, but many specific cases have been proven. Fairly recently the Weinstein conjecture has been proven in generality for the three-dimensional case.

Theorem 5.1 ([Tau07, Taubes]).
Let M denote a compact, oriented 3-dimensional manifold, and let α be a contact form on M. Then the Reeb vector field R_α has a closed orbit.

As stable Hamiltonian structures generalize contact structures, there is a next logical step to try to prove the three-dimensional Weinstein conjecture for stable Hamiltonian structures. This problem is still open, however the formulation of Weinstein's conjecture has been generalized to the stable Hamiltonian case and some work has been done.

Theorem 5.2 ([HT09b, Hutchings & Taubes]).
Let (ω, λ) be a stable Hamiltonian structure on a closed oriented connected three-dimensional manifold M. If M is not a $\mathbb{T}^2$ bundle over S^1, the Reeb vector field has a closed orbit.

We will not delve into these proofs in this paper. However, they do motivate a reason to study stable Hamiltonian structures in dimension three. In the following section we will explore a number of interesting phenomena in the three-dimensional case. Owing particularly to the fact that there is so little space to move around, many structures become more rigid. We will start by presenting these properties, which will ready the reader in general to start reading the contemporary literature. We will conclude the section with an exposition of a seminal structure theorem on three-dimensional manifolds endowed with a stable Hamiltonian structure, which originally appeared in [CV15].

5.1 Geodesibility in Dimension Three

It follows directly that $\lambda \wedge \omega$ is a volume form. Note that the Reeb vector field R now satisfies

$$L_R(\lambda \wedge \omega) = d((i_R\lambda)\omega - \lambda \wedge i_R\omega) = 0,$$

so it preserves a volume form. Recall that a stable Hamiltonian structure by definition defined a stable, and hence geodesible, 1-foliation via $\ker(\omega)$. However, in dimension three we obtain that in fact it is precisely only the Reeb vector fields which are geodesible and preserve some volume form.

Corollary 5.1.1. *Let M an oriented 3-manifold. A vector field X is the Reeb vector field of a stable Hamiltonian structure if and only if $\langle X \rangle$ is geodesible and X preserves a volume form μ.*

Proof. Given a Reeb vector field R of (ω, λ), we have shown it preserved the volume form $\lambda \wedge \omega$. Moreover, as in the proof of lemma 4.2.10, we can pick a metric such that $g(X, X) = 1$ and $\nabla_X(X) = 0$.

Conversely, if $\langle X \rangle$ is geodesible and preserves a volume form μ. Then let $\lambda := i_X g$ and $\omega := i_X \mu$. It follows ω is a nowhere vanishing closed 2-form, thus a Hamiltonian structure. Note $i_X g \wedge i_X \mu$ is nowhere vanishing, thus this defines a stable Hamiltonian structure. Moreover, $\ker(\omega) = \langle X \rangle$ and $\lambda(X) = g(X, X) = 1$, so X is the Reeb vector field of the aforementioned structure.

Lastly we remark that if R is the Reeb vector field of (ω, λ). Then each $\ker(\lambda)$, is an oriented 2-dimensional manifold, thus admits a complex structure J. We can define a metric g compatible with J and ω on $\ker(\lambda)$ by setting

$$g(v, w) := \omega(v, Jw).$$

We remark that as ω is equal to the volume form on $\ker(\lambda)$ induced by g, and $\lambda(X) = 1$ for X transverse to $\ker(\lambda)$, it follows $\lambda \wedge \omega$ is the volume form on M induced by extending g. ■

It is because of the above corollary that Reeb vector fields in dimension three can also be characterized as the geodesic, volume preserving vector fields of M.

5.2 Taut Foliations

As the maximal rank of ω has now been reduced to 1, it follows that this demand reduces to requiring it is simply a nowhere vanishing 2-form on M. Now one of the most interesting phenomena has to do with the second demand placed on stable Hamiltonian structures

$$\ker(\omega) \subset \ker(d\lambda),$$

and will finally make clear how exactly stable Hamiltonian structures unify the concept of contact structures and taut foliations.

As before we can readily split the space as follows

$$TM = \ker(\lambda) \oplus \ker(\omega).$$

The following discussion will be point-wise. Now by dimensionality $d\lambda(X,Y) \neq 0$ can only occur if both $X, Y \in \ker(\lambda)$. However, algebraically expanding this expression we arrive at

$$d\lambda(X,Y) = X(\lambda(Y)) - Y(\lambda(X)) - \lambda([X,Y]) = -\lambda([X,Y]).$$

Now again by lack of space, we arrive at two possible conclusions. First, $[X,Y] \in \ker(\lambda)$. If this is the case, then $\ker(\lambda)$ is locally involutive. So by Frobenius' theorem 2.3.7 it follows it is locally integrable. Thus, $\ker(\lambda)$ locally defines a codimension one foliation. And, by corollary 2.3.8, it follows

$$\lambda \wedge d\lambda = 0.$$

Or, secondly, it has to be so $[X,Y] \in \ker(\omega)$. By similar reasoning it follows that now

$$\lambda \wedge d\lambda \neq 0.$$

Hence, $\ker(\lambda)$ now locally defines a maximally non-integrable distribution; a local contact structure.

Now note the first case happens exactly when $d\lambda$ vanishes. If this happens for all of M, we are in the situation that λ is a nowhere vanishing closed 1-form. By Tischler's theorem 3.6.14 it follows M fibres over S^1. And by corollary $\ker(\lambda)$ defines a taut foliation of M. We will formulate this particular example of a stable Hamiltonian structure as a lemma.

Lemma 5.2.1. *Let M an oriented 3-manifold. And let (ω, λ) a stable Hamiltonian structure where $d\lambda = 0$. Then* $\ker(\lambda)$ *defines a taut foliation. Its leaves are symplectic manifolds with symplectic form induced by ω.*

Let us briefly recall an earlier given example.

Example 5.2.1.1. Let $\phi : W \to W$ be a symplectomorphism of a two-dimensional symplectic manifold (W, ω). Then its mapping torus W_ϕ had a stable Hamiltonian structures given by $(\omega, d\theta)$. Note $\ker(d\theta)$ defines a taut foliation of W_ϕ. ♦

5.3 The Cutting of M

In the above we set $d\lambda$ identically to 0. However, we can keep some more generality. Note that we derived that locally either $\ker(d\lambda)$ precisely coincided with $\ker(\omega)$, or $d\lambda$ vanished at that particular point. It follows that in full generality we can write

$$d\lambda = f\omega$$

where $f \in C^\infty(M)$. Interestingly f may vanish, and hence f dictates the behaviour of $\ker(\lambda)$ between being a taut foliation, or being a contact structure. The moral is that the behaviour of f can be used to cut M up into regions where the stable Hamiltonian structures exhibits well-known characteristics.

As f is so fundamental to the behaviour of $\ker(\lambda)$, it would be beneficial for us to know what f looks like. Unsurprisingly, as the Reeb vector field R satisfies both $L_R\lambda = 0$ and $L_R\omega = 0$, it must preserve f as well.

Lemma 5.3.1. *For a SHS (M, ω, λ) on an oriented 3-manifold M we have that $d\lambda = f\omega$ for a smooth function f on M. Where f is an integral of motion for the Reeb vector field:*

$$R(f) = 0.$$

Proof. A one-line calculation yields

$$0 = L_R(d\lambda) = L_R(f\omega) = R(f)\omega + fL_R(\omega) = R(f)\omega.$$

As by assumption ω was nowhere vanishing, it must follow $R(f) = 0$. ■

Now as f is an integral of motion for the Reeb vector field R, it means that R is tangent to level sets of f. As by assumption $df \neq 0$, f is a submersion, it follows by the submersion theorem its level sets $f^{-1}(c)$ on these regions are two-dimensional compact submanifolds. Very similar to a completely integrable system, the level sets of f are diffeomorphic 2-tori.

Lemma 5.3.2. *Every compact connected component of a level set of f, where $df \neq 0$, is an embedded 2-torus in M. Which we call a* ***Liouville torus*** *as $\omega|_M = 0$.*

Proof. Note each such component of $f^{-1}(c)$ is a compact two-dimensional submanifold of M. In particular, it is defined by $\ker(df)$, when df is non-vanishing. As such, df defines a coorientation on $f^{-1}(c)$. Because M was orientable, it follows $f^{-1}(c)$ is orientable. Now the Reeb vector field is a nowhere vanishing vector field on this two-dimensional oriented compact surface, hence it must be a 2-torus. As R is tangent to the surface, and R by definition spans $\ker(\omega)$, it follows $\omega|_{f^{-1}(c)} = 0$. ■

Now again observe f as a scalar function of the whole of M. In general f can either be constant or non-constant. This will be our first and main cut.

5.3.1 The Constant

Within regions where f is constant, let us say $f = c$, we obtain $d\lambda = c\omega$. Now as before, we remark that if $f = 0$, then $d\lambda = 0$, and $\ker(\lambda)$ defines a taut foliation. Now we can perform another cut within the regions where $f \neq 0$ to distinguish between the cases where $c < 0$ and $c > 0$. In both scenarios, $d\lambda = c\omega$ so ω exact and $\ker(\lambda)$ defines a contact structure. However, we will now differentiate. If $c > 0$, we will say that with respect the orientation on M induced by $\lambda \wedge \omega$, the contact structure $\ker(\lambda)$ is positively oriented. Vice versa, if $c < 0$, we will say that it is negatively oriented.

In order to make this idea more clear, we need a way to properly identify these level sets and assure the inverse of f is indeed a submanifold. However, previously we could rely on the fact that $df \neq 0$ and apply the submersion theorem, we now cannot. We will rely on the following analytical lemma which we will not prove here, but refer to [CV15].

Lemma 5.3.3. *Let (ω, λ) be a stable Hamiltonian structure on a closed 3-manifold M. And let $d\lambda = f\omega$ the proportionality function. Now let $Z \subset \mathbb{R}$ be any set of Lebesgue measure zero such that $Z \cap \operatorname{im}(f) \neq \varnothing$.*

Then there exists a stabilizing 1-form λ' for ω which is C^r-close for $r \in [1, 2)$ to λ such that $d\lambda' = f'\omega$. Where f' can be written as $\sigma \circ f$ for a function $\sigma : \mathbb{R} \to \mathbb{R}$ which is C^0 close to the identity, and locally constant on some open neighbourhood of Z. Thus, f' is locally constant on some open neighbourhood of $f^{-1}(Z)$. Moreover, $f' = c$ on an open neighbourhood of $f^{-1}(c)$.

The resulting function f' can be interpreted as a thickening of level sets of f around a discrete set of levels, which the following corollary makes clear.

Corollary 5.3.4. *Let the same hypotheses as the lemma be true, and set $Z = \{c\}$. Then there exists a λ' which is C^1-close to λ which stabilizes ω. Furthermore, there exists f' with $d\lambda = f'\omega$ such that $f'(x) = c$ for all x in an open neighbourhood U of $f^{-1}(c)$.*

We see that even though $f^{-1}(c)$ might be thin or even highly pathological, $f'^{-1}(c)$ defines an open neighbourhood of $f^{-1}(c)$ on which at least f' is constant c. However, this is satisfactory, as we have found an accompanying stabilizing form λ' which, too, stabilizes our original ω.

We are particularly interested in whenever Z is the discrete set of critical values of f, meaning $Z := \{c \mid f(p) = c \text{ and } df(p) = 0\}$. Note, by Sard's theorem Z is indeed of Lebesgue measure 0 in $\mathbb{R}$.

Lemma 5.3.5. *Let (ω, λ) be a stable Hamiltonian structure on a closed 3-manifold M and $d\lambda = f\omega$.*

Then there exists a possibly disconnected, and possible with boundary, compact three-dimensional submanifold N of M which is invariant under the flow of the Reeb vector field. There also exists another stabilizing 1-form λ' of ω which is C^1-close to λ satisfying $d\lambda' = f'\omega$. Where f' is constant on each connected component of N.

Moreover, f' can be written as $f' = \sigma \circ f$ for $\sigma : \mathbb{R} \to \mathbb{R}$ which is C^0 close to the identity.

Lastly we can find a finite family of opens $\{U_i'\}_{1 \leq i \leq m}$ such that $U' := \cup_i U_i'$ satisfies $U' \cup N = M$.

Proof. Construct f' and λ' as in lemma 5.3.3. As M is closed, Z as a subset of its image, must be compact. Thus, we can extract from the open neighbourhood U of Z on which σ is constant, a finite cover

$$J_i := (a_i, b_i) \text{ for } 1 \leq i \leq n.$$

We can assume without loss of generality that a_i, b_i are regular values, if not; simply shrink the interval an infinitesimal amount. Define J as the closure

$$J := \cup_i \overline{J_i},$$

and set

$$N := f^{-1}(J).$$

Now denote by $I_i := (c_i, d_i)$ connected sets of regular values. Note $I_i \cap J$ need not be empty. However, it does follow that either $I_i \cap J = \varnothing$ or $I_i \cap J = I_i$. So define

$$I := (\cup_i I_i) \backslash J.$$

Denote now by I_i only those $I_i \in I$. Moreover, as J is finite, and $\text{im}(f)$ is compact, it follows that we can renumber I_i so $1 \leq i \leq m$. Now define

$$U' := f^{-1}(I),$$

which are precisely those regions in M where $df \neq 0$, we will return to those shortly under the moniker integrable regions.

As $\text{im}(f) \subset I \cup J$, it follows $U \cup N = M$. To summarize, as σ is constant on each J_i, we get f' is constant on each connected component of N, invariance under R follows from $L_R f = 0$. Moreover, by construction U is built up of a finite union of $U_i := f^{-1}(I_i)$. ■

Remark that for $I_i = (c_i, d_i)$ we have that c_i and d_i are themselves critical values. Thus, they must be contained in Z, hence in some $\overline{J_j}$. Hence, the boundary of U' are contained within N.

Now the following definition formalizes the second cut within regions where f' was constant.

Definition 5.3.6. Let M an oriented closed 3-manifold with stable Hamiltonian structure (ω, λ). Write $d\lambda = f\omega$ for $f \in C^\infty M$. Let $N \subset M$ and f' be as in lemma 5.3.5, then for c a strictly positive real, we define three possibly disconnected, possibly with boundary, submanifolds of N as follows.

- $N_0 \subset N$ is defined by $f' = 0$, thus $d\lambda' = 0$. Hence, $\ker(\lambda')$ defines a taut foliation.
- $N_- \subset N$ is defined by $f' < 0$, thus $d\lambda' = -c\omega$. Hence, λ' defines a negative contact structure.
- $N_+ \subset N$ is defined by $f' > 0$, thus $d\lambda' = +c\omega$. Hence, λ defines a positive contact structure.

It follows $N = N_0 \cup N_- \cup N_+$. ▲

As a result of Tischler's theorem 3.6.14 for the regions where $d\lambda = 0$, where λ defines a taut foliation, we obtain;

Corollary 5.3.7. *The regions N_0 fibre over S^1.*

Proof. It is a direct corollary from Tischler's theorem 3.6.14 as λ is now a non-vanishing closed 1-form. We will recall the geometric intuition: λ can be slightly perturbed and rescaled to obtain an integral 1-form

$$\lambda' \in H^1_{dR}(N_0; \mathbb{Z})$$
$$\lambda' = \sum_{i=1}^{k} z_i f_i^*(d\theta),$$

where $z_i \in \mathbb{Z}$ integer coefficients, k was the rank of $H^1_{dR}(N_0)$, and $f_i : N_0 \to S^1$ were smooth maps. Then $f := \sum_{i=1}^{k} z_i f_i$ defined a smooth proper surjective submersion $f : N_0 \to S^1$, thus results in a fibration. ∎

Note that this slightly perturbed λ' still stabilizes ω. The following construction has been done multiple times in this paper. Recall ω restricts to a non-degenerate closed 2-form, a symplectic form, on the fibres of this fibration. It follows the Reeb vector field must be transverse to the fibres. And indeed, as N_0 fibres over S^1, this Reeb vector field has closed orbits parallel to the base space.

5.3.2 The Integrable

Now we will treat the different regions of M. It logically follows for $M \backslash N$ that f is non-constant, hence $df \neq 0$. On these regions we obtain the situation of lemma 5.3.2. Thus, these regions are foliated by level sets of f diffeomorphic to 2-tori.

Definition 5.3.8. Let M an oriented 3-manifold with stable Hamiltonian structure (ω, λ). Let $U_i \subset M$ a compact connected region where $df \neq 0$. Then $U_i \cong I \times \mathbb{T}^2$ and is called an **integrable region**. ▲

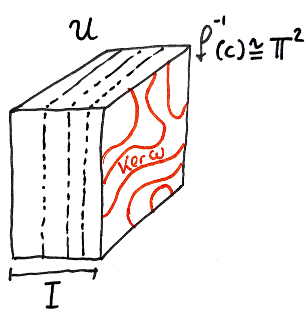

Figure 19: Integrable Region

Recall that in lemma 5.3.5 we already found these regions. To clear up notation, the U_i' of lemma 5.3.5 are $\text{int}(U_i)$ of the above lemma. In fact, we have also already shown the union of these together with N, comprises the whole of M, and the boundary of its closure overlaps with N.

In the following we will denote by U a single compact connected component U_i. Now note that as the Reeb vector field R is tangent to these tori, and by definition R generates $\ker(\omega)$, it follows that each $\ker(\omega)|_{\{r\}\times\mathbb{T}^2}$ defines a line field on each torus. As line-fields are always integrable, these are subfoliations of the tori. We know these line fields fall in either of two categories; they either have periodic orbits or not. More often that not, it is the latter. Moreover, $\ker(\omega)|_{\{r\}\times\mathbb{T}^2}$ may have any form on each torus.

Now in a similar argument to either Tischler 3.6.14 or Arnold-Liouville 2.5.11, we can straighten out these line fields, so they are linear on each torus.

Lemma 5.3.9. *Let U be an integrable region for a stable Hamiltonian structure (ω, λ). There exists a diffeomorphism*

$$\psi : I \times \mathbb{T}^2 \to U \cong I \times \mathbb{T}^2.$$

If we denote by $T_r \subset U$ the tori of the integrable region indexed by coordinate $r \in I$, and similarly denote $\lambda_r := \lambda|_{\{r\}\times\mathbb{T}^2}$, R_r and ω_r. Then ψ has the properties:

- *ψ preserves tori.*
- *$\psi^*\omega$, ψ^*R, and $(\psi^*\lambda)_r$ are linear on $\psi^{-1}(T_r)$.*

Proof. The outline of the proof is to symplectize the integrable region and then apply Arnold-Liouville. Construct the symplectization as in 4.2.4.

$$(W, \Omega) := ((-\varepsilon, \varepsilon) \times U, \omega + d(t\lambda)).$$

So the hypersurfaces of contact type are copies of the integrable region U.

Now recall that Hamiltonian vector fields were defined by using the non-degeneracy of a symplectic form to arrive at the unique vector field satisfying

$$dH = i_{X_H}\Omega.$$

Hamiltonian vector fields flowed along the level set of such a function and preserved the symplectic form. Now if we define

$$\begin{aligned} T : W &\to \mathbb{R}, \\ (t, r, z) &\mapsto t, \end{aligned}$$

we obtain the Hamiltonian vector field X_T satisfying $dT = dt = i_{X_T}\Omega$. The flow of X_T is along its level sets, and these are precisely given $T^{-1}(t) = \{t\} \times U$. Similarly define another Hamiltonian vector field via

$$\begin{aligned} F : W &\to \mathbb{R}, \\ (t, r, z) &\mapsto r, \end{aligned}$$

which defines the Hamiltonian vector field X_F satisfying $dF = dr = i_{X_F}\Omega$. It flows along level sets $F^{-1}(r) = (-\varepsilon, \varepsilon) \times \{r\} \times \mathbb{T}^2$. Now the pair of these functions define the Hamiltonian

$$H := (T, F) : W \to \mathbb{R}^2$$
$$(t, r, z) \mapsto (t, r),$$

whose level sets are precisely given by individual copies of tori $H^{-1}(t, r) = \{(t, r)\} \times \mathbb{T}^2$.

Now we will show this is a completely integrable Hamiltonian system. Note that X_T has to satisfy

$$i_{X_T}\Omega = dt = i_{X_T}\omega + t i_{X_T} d\lambda + i_{X_T}(dt \wedge \lambda),$$

which is only the case if $i_{X_T}\omega = i_{X_T} d\lambda = i_{X_T} dt = 0$, and $i_{X_T}\lambda = -1$. This precisely defines the negative Reeb vector field, so $X_T = -R$ of (ω, λ). The Reeb vector field flowed tangent to each $\{(t, r)\} \times \mathbb{T}^2$, and then so does X_T. Note that X_T still generates $\ker(\omega)$.

As a result we see that

$$L_{X_T} F = i_{X_T} dF = \Omega(X_F, X_T) = 0,$$

so $\{T, F\} = 0$ meaning they Poisson-commute. All the usual results of Poisson commutativity now apply: X_F too flows along the tori $\{(t, r)\} \times \mathbb{T}^2$, both X_T and X_F preserve either T or F, and $[X_T, X_F] = 0$, so their flows commute.

So we have a completely integrable system on X. As in the proof of lemma 2.5.9, needed for the Arnold Liouville theorem, we obtain a well-defined action of $\mathbb{R}^2$ on each level set of $H : X \to \mathbb{R}^2$, which were just individual tori $\{(t, r)\} \times \mathbb{T}^2$. We could then find a complete and discrete lattice of points

$$\Gamma^{(\omega,\lambda)}_{(t,r)} \subset \mathbb{R}^2,$$

which had some integral basis $\{t_1, t_2\}$. Here each $(\tau_1, \tau_2) \in \Gamma^{(\omega,\lambda)}_{(t,r)}$ denotes a 2-flow time of X_T and X_F respectively which leaves a point in $\{(t, r)\} \times \mathbb{T}^2$ fixed.

Note the superscript (ω, λ) denotes a dependency on (ω, λ). Indeed, in constructing the symplectization, we had chosen a stable Hamiltonian structure. We can wonder how our construction varies by different choice of stable Hamiltonian pair. Furthermore, the subscript (t, r) denotes a dependency on which specific torus we are on, which makes sense as Ω does indeed vary between tori, we can wonder if this dependency is smooth or erratic. First, as (ω, λ) were respectively an honest smooth 2-form and smooth 1-form, it follows that the lattice $\Gamma^{(\omega,\lambda)}_{(t,r)}$ varies smoothly between tori $\{(t, r)\} \times \mathbb{T}^2$ when leaving (ω, λ) fixed. Second, in lemma 4.2.7, we have shown the process of symplectization is

C^1 continuous. Meaning if we denote $\mathcal{SHS}(M)$ the space of stable Hamiltonian structure on M bestowed with the C^1-topology, then the map

$$\mathcal{SHS} \to \Lambda(\mathbb{R}^2),$$
$$(\omega, \lambda) \mapsto \Gamma^{(\omega,\lambda)}_{(t,r)},$$

where $\Lambda(\mathbb{R}^2)$ denotes lattices in $\mathbb{R}^2$, is C^1-continuous with respect to pairs (ω, λ). It is even C^2-continous in just λ.

Now let (ω, λ) be fixed again. Each lattice defines a torus via the quotient

$$T^{(\omega,\lambda)}_{(t,r)} := \mathbb{R}^2 / \Gamma^{(\omega,\lambda)}_{(t,r)}.$$

Denote by $\tau^{(\omega,\lambda)}_{(t,r)}$ the image of $\tau \in \mathbb{T}^2$ under the orientation preserving diffeomorphism $\mathbb{T}^2 \cong T^{(\omega,\lambda)}_{(t,r)}$.

Now denote by ϕ^τ the 2-flow of X_T and X_F. Then all the above gives a C^1-continuous collection of free $\mathbb{T}^2$ actions indexed by (ω, λ) as follows

$$\rho : \mathcal{SHS} \times \mathbb{T}^2 \to \mathrm{Diff}_+(X),$$
$$((\omega, \lambda), \tau) \mapsto \phi^{\tau^{(\omega,\lambda)}_{(t,r)}},$$

where the subscript $+$ denotes the diffeomorphisms are orientation preserving. Now we can use the 2-flow of the vector field to define a change of coordinates. Recall this necessitated a choice of base point from which to calculate the flow-times, which then in turn determined the coordinates. Choosing a different base point, changes the appearance of the coordinates, and a base point has to be chosen for each individual torus $\{(t,r)\} \times \mathbb{T}^2$. To formulate this as general as possible, we can formulate the choice of base point as a smooth section s of the trivial fibration

$$(-\varepsilon, \varepsilon) \times I \times \mathbb{T}^2 \to (-\varepsilon, \varepsilon) \times I.$$

As this is a smooth section, it ensures the choice of base point varies smoothly along the different tori. Then the change of coordinates is defined as

$$\Psi \in \mathrm{Diff}_+((-\varepsilon, \varepsilon) \times I \times \mathbb{T}^2),$$
$$\Psi(t, r, \tau) = \phi^{\tau^{(\omega,\lambda)}_{(t,r)}} (t, r, s(t, r)).$$

Recall the flow was contained within each torus $\{(t,r)\} \times \mathbb{T}^2$, so Ψ descends to a self-diffeomorphism of each individual torus. Note, almost by definition, if $(\theta, \phi) \in \mathbb{T}^2 := \sigma$ then applying the action $\phi^{\sigma^{(\omega,\lambda)}_{(t,r)}}$ to these shifted coordinates, and reverting to the original coordinates, is just the usual $\mathbb{T}^2$ action. Meaning

$$\left(\Psi^{-1} \circ \phi^{\sigma^{(\omega,\lambda)}_{(t,r)}} \circ \Psi \right)(t, r, \tau) = (t, r, \tau + \sigma).$$

We conclude that if either a vector field or a differential form is invariant under ρ, then its pullback under Ψ is invariant under the usual $\mathbb{T}^2$ action.

By construction X_H, X_F are invariant under ρ, indeed; the action is defined using their flows. Thus, their pullbacks $\Psi^* X_H, \Psi^* X_F$ are invariant under the usual torus action, hence they define linear vector fields on each torus $\{t, r\} \times \mathbb{T}^2$. Similarly, Ω was invariant under X_H, X_F, thus $\Psi^*\Omega$ is linear on each torus.

Now observe $\psi := \Psi|_{\{0\}\times I\times\mathbb{T}^2}$ as an orientation preserving diffeomorphism on U. Recall that $\Omega_0 = \omega$ and so $(\Psi^*\Omega)|_{\{0\}\times U} = \psi^*\omega$. Furthermore, remember that X_T was precisely the negative Reeb vector field $-R$. We conclude $\psi^*\omega$ and $\psi^* R$ are also linear. As X_F was invariant under $R = X_T$, it follows also $\phi^* X_F|_{\{0\}\times U}$ is linear.

Now note

$$i_{X_F|_{\{0\}\times U}}\Omega = dF = dr = i_{X_F|_{\{0\}\times U}}\omega + i_{X_F|_{\{0\}\times U}}dt \cdot \lambda - i_{X_F|_{\{0\}\times U}}\lambda \cdot dt,$$

which is only the case if $\lambda(X_F|_{\{0\}\times U}) = 0$ and $i_{X_F|_{\{0\}\times U}}\omega = dr$. The middle term vanishes as X_F was invariant under X_T. Pulling everything back by ψ we obtain the same identities. The invariance of λ_r now follows from the fact that $\psi^*\lambda_r$ is fully defined by

$$(\psi^*\lambda)_r(\psi^* X_F|_{\{0\}\times U}) = 0,$$

and

$$(\psi^*\lambda)_r(\psi^* X_T|_{\{0\}\times U}) = 1,$$

where the second equality follows from identifying the vector field with the Reeb field. As both of these, the pullbacks of the restricted vector fields, were invariant under $\mathbb{T}^2$, it follows $(\psi^*\lambda)_r$ is so too. ∎

The invariance properties of $\psi^*\omega$ and $(\psi^*\lambda)_r$ allow us to expand them as follows. Moreover, we can analyse the invariance of $\psi^*\lambda$ not restricted to a torus.

Corollary 5.3.10. *As $\psi^*\omega$ and $\psi^*\lambda_r$ are linear on the tori. It follows that for an integrable region U we can write*

$$\psi^*\omega = (w_1(r)d\theta + w_2(r)d\phi) \wedge dr,$$

and

$$\psi^*\lambda = l_1(r)d\theta + l_2(r)d\phi + l_3(r, \theta, \phi)dr.$$

Here θ, ϕ are the usual coordinates on $\mathbb{T}^2$.

Furthermore, as $d(\psi^*\lambda) = f \cdot \psi^*\omega$ we obtain

$$\frac{\partial l_1}{\partial r} = \frac{\partial l_3}{\partial \theta} - f w_1,$$
$$\frac{\partial l_2}{\partial r} = \frac{\partial l_3}{\partial \phi} - f w_2.$$

If in addition f is constant on each torus, then $\psi^*\lambda$ is $\mathbb{T}^2$ invariant too.

Proof. Remark that w_1, w_2, l_1, l_2 are all $\mathbb{T}^2$-invariant. If in addition f is constant on the tori $\{r\} \times \mathbb{T}^2$, thus $\mathbb{T}^2$-invariant, then l_3 must be as well. Thus, we obtain invariance of $\psi^*\lambda$. ∎

Remark that if we define the tori as level sets of f, then f is indeed tautologically constant on the tori. The above corollary says that if we foliate the integrable region by different tori along which the Reeb vector field is still tangent, then f might not be constant, and we cannot guarantee invariance of $\psi^*\lambda$. Or differently, if we find another stabilizing 1-form, whose proportionality function is not constant on the tori, the same reasoning applies.

Example 5.3.10.1. Let $(\lambda, d\lambda)$ be a stable Hamiltonian structure where λ is $\mathbb{T}^2$-invariant. Let η be another 1-form such that $\eta \wedge d\lambda = 0$ and $d\eta = g d\lambda$, where g is not constant on the tori. Define a new 1-form $\lambda_\varepsilon := \lambda + \varepsilon \cdot \eta$. Note that for $\varepsilon > 0$ small enough we have;

$$\lambda_\varepsilon \wedge d\lambda = \lambda \wedge d\lambda \neq 0,$$

and

$$\ker(d\lambda_\varepsilon) = \ker(d\lambda + \varepsilon d\eta) = \ker((1 + \varepsilon g)d\lambda) \supset \ker(\omega).$$

So $(\lambda_\varepsilon, d\lambda)$ is also a stable Hamiltonian structure for ε small enough. Note by construction that

$$d\lambda_\varepsilon := f_\varepsilon d\lambda = (1 + \varepsilon g)d\lambda,$$

and clearly $f_\varepsilon = 1 + \varepsilon g$ is not constant on the tori. Thus, λ_ϵ stabilizes $d\lambda$, but we cannot guarantee $\mathbb{T}^2$ invariance of $\psi^*\lambda_\epsilon$. ◆

We will now again use U_i to distinguish a single connected component from the whole of U, and make clear whether we are talking about the open or closed U_i'. Recall from lemma 5.3.5 that each U_i was defined as $f^{-1}((c_i, d_i))$ where (c_i, d_i) was an interval of regular values not contained in a region (a_j, b_j) containing the critical values. So we have the following corollary.

Corollary 5.3.11. *On each $U_i' := f^{-1}(I_i) \cong (c_i', d_i') \times \mathbb{T}^2$ the proportionality function f is given by projection onto the first factor. And for r sufficiently close to either c_i' or d_i', we have $\{r\} \times \mathbb{T}^2 \subset N$.*

Proof. By construction the individual tori are level sets $f^{-1}(c)$, and the intervals (a_i, b_i) are precisely these regular values. Moreover, we have shown that a_i and b_i themselves were critical values. ∎

Additionally, combining the thickened level sets of lemma 5.3.5 together with corollary 5.3.10 we can write down.

Corollary 5.3.12. *On each $U_i \cong [a_i, b_i] \times \mathbb{T}^2$ the stable Hamiltonian structure (ω, λ) is $\mathbb{T}^2$-invariant and $f(r, \tau) = \alpha_i \cdot r + \beta_i$.*

Proof. As f is given by projection onto the first factor, we obtain that $f' = \sigma \circ f$ is constant on the tori $[a_i, b_i] \times \mathbb{T}^2$. This is precisely the demand of 5.3.10 and so (ω, λ') is a $\mathbb{T}^2$-invariant stable Hamiltonian structure. Now λ and λ' are C^1-close, so it follows that $f\omega$ and $f'\omega$ are C^0-close. Thus, $f(r, \tau) = \alpha_i r + \beta_i$ as f' is constant on each torus. Thus, it follows (ω, λ) is $\mathbb{T}^2$-invariant. ∎

5.3.3 The Total

We have now in fact fully proven the seminal structure theorem from Cieliebak and Volkov.

Theorem 5.3.13 ([CV15, Cieliebak & Volkov]).
Let (ω, λ) be a stable Hamiltonian structure on a closed three-dimensional manifold M and let $d\lambda = f\omega$ the proportionality function.
Then there exists a possibly disconnected and possibly with boundary compact three-dimensional submanifold N of M, invariant under the Reeb flow, and U comprised of a finite disjoint union of compact integrable regions U_i.
Moreover, there exists another stabilizing 1-form λ' which is $C^1 -$ close to λ for which we have $f'd\lambda' = \omega$. The above satisfy the following properties:

- *N consists of possibly disconnected, possibly with boundary, compact 3-dimensional submanifolds $N = N_0 \cup N_+ \cup N_-$. On each of which f' is respectively constant 0, positive, or negative.*
  As a result λ defines a taut foliation and fibration over S^1 on N_0, and respectively a positive and negative contact structure on N_+ and N_-.
- *On each $U_i \cong [c_i, d_i] \times \mathbb{T}^2$ the stable Hamiltonian structure (ω, λ) is $\mathbb{T}^2$-invariant. Moreover, $f(r, z) = \alpha_i \cdot r + \beta_i$.*

As said before, this structure theorem is particularly powerful. Given a stable Hamiltonian structure, it enables us to cut up M into regions where λ exhibits a structure for which we already have a lot of tools. Moreover, the amount of connected components of U is also finite and always equal in number.

For example, given the structure theorem and Taubes' theorem 5.1 we already obtain the following.

Corollary 5.3.14. *Let (ω, λ) be a stable Hamiltonian structure on a 3 manifold. If R has no periodic orbits, then f is nowhere vanishing.*

Proof. If f is nowhere vanishing, then λ is a contact form. For which Taubes' proof of the Weinstein theorem holds. Thus, R must have a periodic orbit. Thus, by contradiction f has to vanish somewhere. ∎

We will conclude this section with a slight exposition of the usefulness of the structure theorem in the work [CR23], where they study so-called Birkhoff sections.

Definition 5.3.15. A **section** of a non-singular vector field X on a closed 3-manifold M is an embedded closed surface that is everywhere transverse to

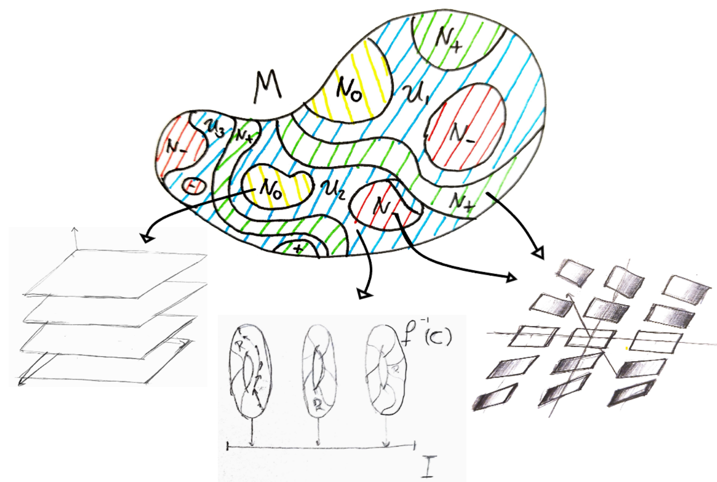

Figure 20: A sketch of Theorem 5.3.13

X and intersects all its orbits. A **transverse surface** is an immersed surface B, whose interior is embedded and transverse to X, and its boundary ∂B is a collection of periodic orbits of X. It is **Birkhoff** if it intersects all orbits of X in bounded time. ▲

Note that finding a section is very restrictive. Indeed, in the three-dimensional case we can apply Tischler and the characterization of taut foliation presented, to observe that a manifold permitting such a section must fibre over S^1. Birkhoff sections are a relaxation of this requirement.

Note that given either a regular or Birkhoff section, of a periodic non-singular Reeb vector field, we can define the first return map; the Poincaré map. Recall that in the three-dimensional case, the Reeb vector field was precisely classified as those volume preserving geodesic vector fields. As it is volume preserving, the determinant of the linearized Poincaré map has to be precisely 1. As it is non-singular, the eigenvalues of the Poincaré map can never be 1. It follows that either the Poincaré map defines an irrational rotation, or it has only two eigenvalues; sections in three-dimensional space are by definition two-dimensional. This already shows that stable Hamiltonian structures in three-dimensional space tell us a lot about the dynamics of the vector fields.

By the structure theorem, given a stable Hamiltonian structure (ω, λ) and its corresponding Reeb vector field R, we can cut up M in three different regions. We had the integrable regions U, the contact regions$N_{\pm}$, and the flat regions N_0. Now we can already readily say that in the flat regions N_0 admits a section as defined above. Moreover, this cutting was done along invariant tori.

Theorem 5.3.16 ([CR23, Cardona & Rechtman]).
Let (ω, λ) be a stable Hamiltonian structure on M^3 and R its Reeb vector field. Assume R has no periodic orbits and $d\lambda = f\omega$ with non-constant f.
Then, cutting along any invariant torus $\mathbb{T}^2$, we obtain $\widehat{M} \cong \mathbb{T}^2 \times I$ a three-manifold with boundary in which R admits an annulus-like section, and R does have an orbit which is orbit equivalent to a suspension flow of an irrational pseudorotation.

Sketch of Proof. For the proof we refer the reader to [CR23]. However, the ideas should be familiar. One can cut up M along invariant tori because of the structure theorem. The closure of the resulting manifold $\overline{M}$ obtained after each cutting has a boundary consisting of one more invariant torus than before the cutting; $\overline{M} \cong \mathbb{T}^2 \times I$.

One the integrable regions the ω, λ could be assumed to be $\mathbb{T}^2$-invariant. In the same paper it is proven this forces R to have an orbit which is orbit equivalent to the suspension flow of an irrational rotation.

In the connected components $N_{\pm}$, we have λ is a contact form. So by Taubes we have periodic orbits on $\overline{M}$. But by assumption R had no periodic orbits on M. By lemma 3.1 in [CR23] it follows that R has an orbit which is orbit equivalent to the suspension flow of an irrational pseudorotation of the annulus. In the connected component N_0 we get the same results.

As a result $\overline{M}$ obtained by gluing these connected components along their boundaries. And on each of these connected components, R has an orbit which is orbit equivalent to the suspension flow of an irrational (pseudo)rotation of the annulus.

Along each invariant which bounds any connected component as described above, we can choose a non-trivial homology class and an annulus-like section of R. The rest of the proof deals with how to carefully glue each section, so it is preserves throughout all the glueings up to recovering $\overline{M}$. ∎

Many more interesting resulting relating the dynamics of Reeb vector fields are derived in [CR23]. The philosophy being that having such a well-defined and relatively easy observed structure on M^3 defined by the proportionality function, puts some very restrictive bounds on the allowed behaviour of the Reeb vector field, and thus allows us to derive many dynamical properties of is by simply cutting up M and observing its very restricted behaviour on each individual piece. Where, moreover, for each individual piece we have a plethora of already available literature and tools available. Both the corollary and the theorem show the power of the structure theorem.

6 Outlook

We will provide a research direction we believe might be interesting to pursue with the aid of stable Hamiltonian structures. The following section is mostly qualitative of nature, and is to provide a rough introduction. We invite the reader to read the sources cited.

6.1 Symplectic Field Theory

Symplectic field theory is a framework introduced in [EGH00] and has the ambition to provide an approach to construct invariants. The field is still very much in its initial development, but has its links to topological quantum field theory. We will give a brief overview of two central topics in symplectic field theory; J-holomorphic curves and symplectic cobordisms, as well as its relation to stable Hamiltonian structures.

We will first introduce J-holomorphic curves.

Definition 6.1.1. Let (M, J) an almost complex manifold and (Σ, j) a Riemann surface. Then a **J-holomorphic curve** is a smooth map

$$u : \Sigma \to M,$$

which satisfies the Cauchy-Riemann equations:

$$du \circ j = J \circ du$$

▲

The moniker curve stems from the fact that Σ as a real two-dimensional manifold, has complex dimension 1, thus $u(\Sigma)$ is in a sense a complex curve in M. Similar to curves we can say a J-holomorphic curve is integrable, which is precisely the case whenever in local coordinates (x, t) of X we satisfy

$$\partial_x u + J(u)\partial_y u = 0.$$

Integrable J-holomorphic curves can provide valuable insights into a symplectic structure on M whenever $\omega(X, JX) > 0$, meaning J is tamed by ω.

The other important concept are symplectic cobordisms. A cobordism between two manifolds M and M' of equal dimension n is roughly speaking a third manifold W of higher dimension $(n + 1)$, whose boundary is the disjoint union $\partial W = M \sqcup M'$. In a sense it is a connecting manifold. Employing the narrative that contact structures naturally arise on boundaries of symplectic manifolds, the concept of a symplectic cobordism is almost self-constructing,

Definition 6.1.2. Given (M_+, ξ_+) and (M_-, ξ_-) two respectively positively and negatively oriented contact $(2m - 1)$- manifolds, with respectively positive and negative contact structure. Then a **symplectic cobordism** between M_+ and M_- is a symplectic compact $(2m)$ manifold (W, Ω) with boundary such that

- $i : M_+ \hookrightarrow \partial W$ is an orientation preserving embedding.
- $j : M_- \hookrightarrow \partial W$ is an orientation reversing embedding.
- $\partial W \cong M_- \sqcup M_+$.
- $i(M_+)$ and $j(M_-)$ are hypersurfaces of contact type.

▲

Note that being a hypersurface of contact type implied the existence of a local transverse Liouville vector field. By the above definition, the Liouville vector field points inwards into W at $j(M_-)$ and outwards at $i(M_+)$.

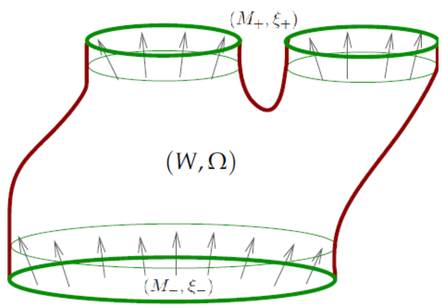

Figure 21: Symplectic cobordism, and Liouville vector fields [Wen16]

As these Liouville vector fields respectively "come from nothing" and "go into nothing", we might be interested in extending such a symplectic cobordism. This is readily done in an intuitive manner.

Definition 6.1.3. Let $(W, M_\pm, \Omega)$ be a symplectic cobordism. Denote by $\lambda_\pm$ the contact forms of $M_\pm$ and denote by

$$\widehat{M_\pm} := [0, \pm\infty) \times (M_\pm, d(e^r \lambda_\pm))$$

the extrinsic symplectization as in 4.1.23.1. Then **symplectic completion** is the non-compact symplectic manifold $(\widehat{W}, \widehat{\Omega})$ obtained by gluing

$$(\widehat{W}, \widehat{\Omega}) = \widehat{M_+} \# W \# \widehat{M_-}.$$

The manifolds $\widehat{M_\pm}$ are also called **cylindrical ends**. ▲

Example 6.1.3.1. In the process of extrinsically symplectizing a contact manifold, as is done in 4.1.23.1, the resulting manifold is the symplectic completion of the trivial symplectic cobordism $([0, 1] \times M, d(e^r \lambda))$. ♦

Now to reintroduce J-holomorphic curves let us observe the following subset of J-holomorphic curves.

Definition 6.1.4. Let Σ a Riemann surface. And let $\Gamma \subset \Sigma$ a finite and discrete subset. Then a **punctured** J-holomorphic curve is the J-holomorphic curve $u : \sigma \to M$, where $\sigma := \Sigma \backslash \Gamma$. ▲

Note there is a natural identification of σ as a Riemann surface with cylindrical ends. Indeed, around each puncture one can choose a biholomorphic coordinate chart to the punctured real disc. Then we can identify each such punctured real disc with the half-cylinder $S^1 \times [0, \infty)$. Thus, each puncture becomes a cylindrical end.

Now without delving into too much detail. One can define the energy of a J-holomorphic curve by integrating over the curve

$$E(u) := \int_\sigma u^* \widetilde{\Omega},$$

where $\widetilde{\Omega}$ is a suitably chosen symplectic form on the symplectic completion $\widehat{W}, \widehat{\Omega}$ such the $E(u)$ is finite. This property of having finite energy is in fact so important, that it forces the $J - holomorphic$ curve to embed in a specific way.

Definition 6.1.5. A puncture ξ of a J-holomorphic curve $u : \sigma \to M$ is **positively or negatively asymptotic** to a periodic Reeb orbit γ with period T of M, if one can choose a holomorphic coordinate chart around ξ identified with the positive or negative half-cylinder, such that in these local coordinates (x, y) we have

$$u(x, y) = \exp_{(Tx, \gamma(Ty))} h(x, y) \text{ for } |x| \text{ sufficiently large.}$$

Here exp is the usual exponential map on Riemannian manifolds sending a vector to the time 1 flow of the unique geodesic, and $h(x, y)$ is a vector field on $\mathbb{R} \times S^1$ along satisfying $h(x, -) \to 0$ as $|x| \to \infty$. ▲

Geometrically what is happening, is each cylindrical end of σ is mapped to a cylindrical end of M of choice, and is done so that asymptotically the cylindrical ends agree with a closed Reeb orbit. Indeed, the finite set of punctures is spit between those positively asymptotic and negatively asymptotic

$$\Gamma = \Gamma_+ \sqcup \Gamma_-.$$

The idea being that this now defines an invariant of M by counting the J-holomorphic with certain positively symplectic punctures and negative symplectic in the symplectic completion of a symplectic cobordism. This is a similar idea as applied in Floer and Morse homology.

Now stable Hamiltonian structures generalize contact structures, so the above discussion may also apply. It is natural to look at symplectic cobordisms between Hamiltonian structures, or Hamiltonian structures with a contact structure. In fact, some work has already been done in section 7 of [CV15]. Where it is shown that the usual notion of symplectic cobordism is too strong to work

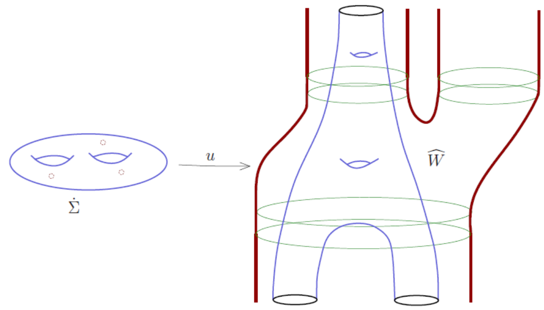

Figure 22: A J-holomorphic curve with one positive, and two negative asymptotic punctures [Wen16]

with in the case of symplectic Hamiltonian structures. And stable Hamiltonian structures too had Reeb vector fields, periodic orbits.

Moreover, symplectic field theory has undergone a growth spurt since their introduction in [EGH00], but the fundamental ideas and theories underlying seem to be rather unstable, see for example the excellent exposition in [Wen16]. Not only can stable Hamiltonian structures aid in laying a more solid foundation for the development of symplectic field theory [CM05]. The invariants themselves arising from look at punctured J-holomorphic curves in the context of symplectic cobordisms between Hamiltonian and contact structures are interesting to research. As is the driving force behind the development of symplectic field theory.

6.2 Conclusion

We hope the reader now has a thorough understanding of stable Hamiltonian structures in order to delve into the still fairly scarce amount of literature on this topic. The exposition in this paper has shown that stable Hamiltonian structures are not a fringe technical construction. They arise very naturally, and almost inevitably, when following the natural flow of symplectic and contact structures along which Weinstein's conjecture floats. Moreover, they are the logical expansion of the already interesting narrative of confoliations; uniting the apparent dichotomy of taut foliations defined by a closed 1 form and contact structures. They solidly manifest themselves as the correct framework to study when researching the h-principle and are a very straightforward geometric construction to make on any three-dimensional manifolds. Lastly, they have made their breakthrough in the budding field of symplectic field theory; an area of mathematics which has developed faster than its foundations can carry it by virtue of its promising outlooks.

Appendices

A Function Space Topology

Though point-set topology readily and visually expands to smooth manifolds, dealing with functions is slightly more abstract. Functions and maps are inherently defined over a domain, and therefore may occupy a wide range in their codomain. To take a poetic angle; two non-parallel lines intersect only once and will subsequently forever drift farther apart, never to see each other again. Two nearby parallel lines will forever be so close, yet so far from intersecting. The obtuse mathematician will of course first claim one has assumed Euclidean space, and secondly define a topology which distinguishes these two types of closeness.

The correct framework to do this in, is the framework of jets. The following section will extensively treat these, and conclude with a definition of the two main topologies used: the strong and weak topology.

A.1 Jets

In differential topology the notion of jets is of utmost importance, and it is the right space to work in for many constructions. Fundamentally, jets formalize the idea of being equivalent up to certain order, in much the same way that Taylor expansions do so for polynomials. The concept of jets is paramount in defining a topology on the space of continuous functions, and we will see that many constructions can be better understood and explained using jets. The following sections will introduce jets, jet spaces and jet bundles.

As a point of technicality, jets of r-differentiable maps are defined point-wise, and hence strictly speaking depend only on the germ of a map $[f]$ at a point. However, for brevity, we will denote such a germ simply by f and speak only of maps, where being a germ of a map in implicit.

Definition A.1.1. Let $f \in C^r(M, N)$ be defined at a point $p \in M$ and choose a local coordinate system around p. The **r-jet of** f **at p** denoted by $j^r_p(f)$ is the equivalence class defined as $f \sim g$, if:

(i) $f(p) = g(p)$;

(ii) all partial derivatives up to and including order r agree.

The collection of all these jets is called the **jet space** $J^r(M, N)$. Note that the 0-jet $j^0_p f$ is just defined by $(p, f(p))$, hence $j^0 f$ is represented by its graph, and the collections of all such "possible graphs" $J^0(M, N)$ is equivalent to $M \times N$. ▲

In fact the jet space $J^r(M, N)$ forms a fibre bundle over $M \times N$. We will sketch the construction in the following.

As an r-jet is based on a map, we have two natural projections onto the source and target of these maps:

$$\begin{aligned}\pi_s : J^r(M,N) &\to M;\\ j^r_p(f) &\mapsto p,\end{aligned}$$

and

$$\begin{aligned}\pi_\tau : J^r(M,N) &\to N;\\ j^r_p(f) &\mapsto f(p).\end{aligned}$$

Using these we can also define subspaces of $J^r(M,N)$ where we denote

$$J^r_x(M,N) := \pi_s^{-1}(x),$$

to mean all jets with source x, and

$$J^r(M,N)_y := \pi_\tau^{-1}(y),$$

to mean all jets with target y. Combining these two we get the subspace contained in both

$$J^r_x(M,N)_y := \pi_s^{-1}(x) \cap \pi_\tau^{-1}(y),$$

which are all jets with source-target pair x, y.

Now if M is m-dimensional and N is n-dimensional, choose local coordinates x_i around p and y_j around $f(p)$, then we can define coordinates on the jet-space $J^r(M,N)$ by using the partial derivatives.

Introduce an m-multi-index $I := (I_1, \dots, I_m)$ with the following operations:

$$\begin{aligned} p^I &:= \prod_{i=1}^m x_i^{I_i}, & |I| &:= \sum_{i=1}^m I_i,\\ \partial_I &:= \prod_{i=1}^m \left(\frac{\partial}{\partial x_i}\right)^{I_i}, & I! &:= \prod_{i=1}^m I_i!. \end{aligned}$$

Then for $|I| \leq r$ we can define the following numbers:

$$u^I_j := \partial_I f_j,$$

we note that if f, g define the same r-jet at p then these numbers agree, thus we use these numbers as coordinates of $J^r(M,N)$. Conversely, if we are given some list of numbers a^I_j, we can define polynomials by using them as coefficients:

$$f_j := \sum_{0 \leq |I| \leq r} \frac{a^I_j x^I}{I!}.$$

Some combinatorics using the definitions leads us to

$$\dim(J^r(M,N)) = m + n\sum_{i=0}^{r}\binom{m+i-1}{i} = m + n\binom{m+r}{r},$$

which, although it gets large quickly, is finite nonetheless and shows locally it is isomorphic to euclidean space. Hence, we conclude:

Lemma A.1.2. *If M, N are C^r-manifolds, then $J^r(M,N)$ is a fibre bundle over $M \times N$ of dimension $m + n\binom{m+r}{r}$.*

The above construction boils down to saying that if we have chosen a coordinate chart (ϕ, U) on M and (ψ, V) on N, then

$$\begin{aligned}\theta : J^r(U,V) &\to J^r(\phi U \subset \mathbb{R}^m, \psi V \subset \mathbb{R}^n),\\ j^r(f) &\mapsto j^r_{\phi(x)}(\psi f \phi^{-1}),\end{aligned}$$

is a bijection, determining a coordinate chart for $J^r(M,N)$. It helps intuition to realize that r-jets of functions between $\mathbb{R}^m$ and $\mathbb{R}^n$ agree with the Taylor series up to order r, thus one might interpret taking r-jets as a coordinate-free generalization of taking r-th order Taylor series. In fact one formally says that in local coordinates the r-th order Taylor polynomial is a local representative of the r-jet of f. Realizing this, one can derive the regularity of the $J^r(M,N)$ manifold readily: if M, N and f are all C^{r+s}, then $J^r(M,N)$ is C^s, meaning: $j^r f : M \to J^r(M,N)$ defined as $x \mapsto j^r f(x)$ is s times differentiable.

Lemma A.1.3. *If M, N are C^{r+s}-manifolds, then $J^r(M,N)$ is a C^s manifold.*

Of course if we assume everything to be smooth we obtain the following.

Lemma A.1.4. *If M, N are smooth manifolds, then $J^\infty(M,N)$ is an infinite dimensional smooth manifold. In particular, it is a smooth fibre bundle over $M \times N$.*

We can see that sections of the jet space are precisely these $j^r f$, these are actually of particular importance in for example algebraic geometry and are defined on their own.

Definition A.1.5. The **s-jet prolongation** of a function $f \in C^r(M,N)$ for $s \le r$ is the map

$$j^s f : M \to J^s(M,N)$$

▲

From now on when taking prolongations or jets, we will assume the functions to be differentiable enough for such an operation to make sense. Prolongation has several nice intuitive algebraic properties. First we can define composition in a very straightforward manner, and this is well-defined.

Lemma A.1.6. *Let $f : K \to M$ and let $g : M \to N$ then*

$$j^r_{f(x)} g \circ j^r_x f = j^r_x(g \circ f).$$

Proof. We see

$$\begin{aligned} j^r_x(g \circ f) &= j^r_{f(x)}(g) \\ &= j^r_{f'(x)}(g') \ . \\ &= j^r_x(g' \circ f') \end{aligned}$$

Intuitively this is exactly saying that if the local Taylor expansions of f, f' and g, g' pairwise agree up to order r, then the Taylor expansions of $g \circ f$ and $g' \circ f'$ agree up to order r. So defining compositions of jets in this way is well-defined ∎

As a corollary we see the following.

Corollary A.1.7. *If $\phi : M' \to M$ and $\psi : N \to N'$ are r differentiable maps, then there is an induced map*

$$J^r(\phi, \psi) : J^r(M, N) \to J^r(M', N')$$

given by

$$j^r_p f \mapsto j^r_p(\psi \circ f \circ \phi).$$

Those familiar with some category theory may recognize that J^r closely resembles the well-known Hom-functor. This should make sense, as we are essentially looking at the Hom-functor with an equivalence up to a certain order. In fact, we will remark for the interested reader, that J^r, too, defines a bifunctor which is contravariant in its first argument, and covariant in its second.

Let us observe some jet spaces to both strengthen our grasp on the concept, and derive some useful results which may aid us later on. In the following we will restart assuming everything is smooth, we will also not concern ourselves with proving the structure groups of the following fibre bundles agree.

Example A.1.7.1. The jet space $J^1_0(\mathbb{R}, M)$ consists of those 1-jets of functions $f : \mathbb{R} \to M$ which have as source 0. Of course $\mathbb{R}$ is already trivial, denote its coordinate with x. Denote (local) coordinates of M by $(y_1, \dots, y_m)$. We get that a local representation of the 1-jet of $f = (f_1, \dots, f_m)$ is an m-tuple of derivatives at 0 prepended by a 0:

$$j^1_0 f = (0, f(0), \partial_x f_1(0), \dots, \partial_x f_m(0)).$$

We form an isomorphism to the tangent space $J^1_0(\mathbb{R}, M) \cong TM$ as follows. First note the target mapping turns this into a smooth fibre bundle over M;

$$\pi_\tau : J^1_0(\mathbb{R}, M) \to M.$$

If $f(0) = p$ then $j^1_0 f \mapsto v \in T_p M$ with

$$v_j = (\partial_x f_j(0))\partial_j,$$

where we had chosen local coordinates y_j around $f(p)$ and hence coordinates ∂_j on the tangent space. This identifies $J^1_0(\mathbb{R}, M)_p \cong T_pM$. So we conclude

$$J^1_0(\mathbb{R}, M) \cong TM$$

as vector bundles over M.

Remark that even though at first sight $J^1_0(\mathbb{R}, M)$ does not carry a vector bundle structure, it does inherit one in much the same way TM inherits one by construction. In fact, upon closer inspection, note that the construction of equivalence of first order tangency of maps $f : \mathbb{R} \to M$ with source 0, is exactly the same as the geometric tangent space definition of equivalence classes of curves $\gamma : \mathbb{R} \to M$ up to velocity at a point $\gamma(0) \in M$. For those interested note that $J^1_0(\mathbb{R}, -)$ defines the same functor as to the tangent functor $T(-)$.

One can of course generalize this notion to $J^r_0(\mathbb{R}^n, M)$, which in literature are called (n, r)-velocities. For example $(1, 2)$-velocities agree with the geometric notion of equivalence classes of curves having both the same velocity and acceleration. Note that for $r = 1$ one obtains:

$$J^1_0(\mathbb{R}^n, M) \cong TM^n.$$

♦

Example A.1.7.2. Dually the jet space $J^1(M, \mathbb{R})_0$ consists of those 1-jets of functions $f : M \to \mathbb{R}$ which now have as target 0. Choose coordinates as before, so we get a local representative

$$j^1_p f = (p, 0, \partial_1 f(p), \ldots, \partial_m f(p)).$$

Remark that it is now the source map which turns $J^1(M, \mathbb{R})_0$ into a fibre bundle over M. Now as the target space $\mathbb{R}$ is a vector space, this is even a vector bundle. Note the very intuitive mapping

$$j^1_p f \mapsto df_p := (\partial_1 f(p) \ldots \partial_m f(p)),$$

and so we conclude

$$J^1(M, \mathbb{R})_0 \cong T^*M,$$

as vector bundles over M. And those interested might again remark that $J^1(-, \mathbb{R})_0$ defines the same functor as the cotangent functor T^*M.

Similarly, one can generalize this notion again to $J^r(M, \mathbb{R}^n)_0$ which are called (n, r)-covelocities. ♦

Example A.1.7.3. With the above two examples, the result for $J^1(M, N)$ will not be surprising. Given a $j^1_p f \in J^1_p(M, N)_q$ and local coordinates x_i for $0 \leq i \leq m$ around p and y_j for $0 \leq j \leq n$ around q, we then see that a local representative of $j^1_p f$ is given largely by the Jacobian of f at p

$$j^1_p f = (p, q, (Jf)|_p) \text{ with } (Jf)^j_i := \frac{\partial f_j}{\partial x_i}.$$

Using the source and target maps we project

$$(\pi_s \times \pi_\tau)(j^1_p f) = (p, q) \in M \times N.$$

This leaves us with the Jacobian as fibre, which is a map

$$(Jf)|_p : T_pM \to T_qN,$$

locally. The coordinate free global expression of the Jacobian is of course given by the differential df. We derive that there is an isomorphism of vector bundles over $M \times N$ of the bundles

$$J^1(M, N) \cong \mathrm{Hom}(TM, TN) \cong T^*M \otimes TN,$$

given by

$$j^1 f \mapsto df.$$

As a word of caution, one could naively have the following line of reasoning

$$J^1(M, \mathbb{R}) \cong T^*M \otimes \mathbb{R} \cong T^*M \cong J^1(M, \mathbb{R})_0,$$

we would like this to be incorrect as it otherwise moots the point of defining the target in the last jet space. The subtlety herein lies in the fact that these isomorphisms are over different base spaces, namely

$$\pi_s \times \pi_\tau : J^1(M, \mathbb{R}) \to M \times \mathbb{R}$$

as a bundle over the product space, and

$$\pi_s : J^1(M, \mathbb{R})_0 \to M$$

as a bundle over just M. It is true there is a canonical identification of $J^1(M, \mathbb{R})$ as a bundle over M, but this is given by:

$$\begin{aligned} J^1(M, \mathbb{R}) &\cong T^*M \times \mathbb{R}, \\ j^1_p f &\mapsto (df_p, f(p)). \end{aligned}$$

♦

Example A.1.7.4. For a more abstract example, note that jet spaces naturally form a fibre bundle over other jet spaces. For $0 \leq s \leq r$ we have a natural projection

$$\pi^{r,s} : J^r(M, N) \to J^s(M, N),$$

by simply We can of course do this sequentially to get a long sequence of fibre bundles over fibre bundles

$$J^r(M, N) \xrightarrow{\pi^{r,r-1}} J^{r-1}(M, N) \xrightarrow{\pi^{r-1,r-2}} \cdots \to J^1(M, N) \cong \mathrm{Hom}(TM, TN).$$

Now observe the following space

$$\bigodot^r J^1(M,\mathbb{R})_0 \cong \mathrm{Sym}^r(T^*M),$$

which is the symmetric r-tensor product. Its elements can be identified r-tuples of 1-jets $(j^1 f_1, \ldots, j^1 f_r)$ in which the order of jets does not matter. Following the usual product rule of derivatives on local representations $\hat{f}_i$ of these jets, together with the fact that by definition each $\hat{f}_i$ is locally 0, one obtains

$$\partial_x^r \left(\prod_{i=1}^r \hat{f}_i \right) = r! \prod_{i=1}^r \partial_x \hat{f}_i,$$

thus the r-jet of the product of these functions only depends on the 1-jets of the individual functions. As the order of taking the product does not matter, we see we have a natural inclusion

$$\bigodot^r J^1(M,\mathbb{R})_0 \hookrightarrow J^r(M,\mathbb{R})_0$$
$$(j^1 f_1, \ldots, j^1 f_r) \mapsto j^r \left(\prod_{i=1}^r f_i \right).$$

Similarly, we see that $\partial_x^{r-j} \left(\prod_{i=1}^r \hat{f}_i \right) = 0$ for $0 \leq j < r$, and so we conclude

$$j^r \left(\prod_{i=1}^r f_i \right) \in \ker(\pi^{r,r-1}).$$

We can now construct a short exact sequence of vector bundles

$$0 \to \bigodot^r J^1(M,\mathbb{R})_0 \to J^r(M,\mathbb{R})_0 \to J^{r-1}(M,\mathbb{R})_0 \to 0.$$

So we get the vector bundle $J^r(M,\mathbb{R})_0 \to J^{r-1}(M,\mathbb{R})_0$ with fibres $\bigodot^r J^1(M,\mathbb{R})_0$.

Remark that nowhere did we explicitly use that $J^1(M,\mathbb{R})_0$ is the cotangent bundle, only the fact that the target is some well-defined 0 in a vector space, and in fact the above construction generalizes by replacing $\mathbb{R}$ by any vector space N. Now identifying

$$\bigodot^r J^1(M,N)_0 \cong \mathrm{Sym}^r(TM, TN \cong N),$$

we can construct the similar short exact sequence of vector bundles

$$0 \to \mathrm{Sym}^r(TM,N) \to J^r(M,N)_0 \to J^{r-1}(M,N)_0 \to 0,$$

thus constructing the vector bundles $J^r(M,N)_0 \to J^{r-1}(M,N)_0$ with fibres $\mathrm{Sym}^r(TM,N)$.

We will sketch the interesting observation that if one thinks of the exterior derivative d as the antisymmetrization of a form, e.g.: for a 1-form $\alpha = \sum_i f_i dx_i$ we have $d\alpha = \sum_{i,j} \partial_j f_i dx_j \wedge dx_i$, then $d\alpha = 0$ precisely when $\partial_j f_i = \partial_i f_j$, hence it was symmetric. Interpreting lowering the index of an r-jet as taking the exterior derivative once, one sees indeed that it is precisely the N valued symmetric forms which are in the kernel of $J^r(M,N)_0 \to J^{r-1}(M,N)_0$, thus explaining the short exact sequence.

♦

A.2 Jet Bundles

The above framework of jets readily adapts to jets of fibre bundles over M. We would like to observe the jets of sections, in particular of the cotangent bundle, because this is the space where contact structures and stable Hamiltonian structures live in. Given a smooth fibre bundle (E, π, M), a section is a smooth map $s : M \to E$ with the added property that $\pi \circ s = id$. As such we can readily define jets of it:

Definition A.2.1. Given a fibre bundle (E, π, M) we can define the **r-Jet bundle** (E, π, M) as the collection of order r jets $j^r s$ of sections $s : M \to E$. We denote this with $J^r\pi$ or $J^r E$ ▲

The above theory of jet-spaces applies fully, as smooth fibre bundles have coordinate charts which are natural with respect to the projection π. Please remark the subtlety that we in particular look at sections, and not all possible maps: there is a difference between $J^r(M,E)$ and $J^r\pi$. In fact, $J^r\pi$ is a subspace of $J^r(M,E)$.

We have already seen that $J^0(M,N) \cong M \times N$. Similarly, we obtain $J^0 E \cong E$, as we have a restriction $\pi \circ s = id$ on sections.

When talking about jet spaces we derived that $J^1(M,N)$ could be identified with $\mathrm{Hom}(TM, TN)$. As such we expect $J^1 E \subset \mathrm{Hom}(TM, TE)$. This subspace is readily derived by using some intuition behind the Hom-functor. Intrinsic in the definition of a fibre bundle is the projection map $\pi : E \to M$. Its differential is a map $d\pi : TE \to TM$. Similarly, for any function $f : M \to E$ we have its differential $df : TM \to TE$. By property of the Hom-functor we now have an induced map of $d\pi$ acting on df by post-composition

$$\begin{aligned} d\pi_* : \mathrm{Hom}(TM, TE) &\to \mathrm{Hom}(TM, TM), \\ df &\mapsto d\pi \circ df = d(\pi \circ f), \end{aligned}$$

or, with a slight abuse of notation, in the language of jet spaces:

$$j^1 f \mapsto j^1_f \pi \circ j^1 f = j^1(\pi \circ f).$$

Now if instead of a function f we take a section s, we have that $\pi \circ s = id$, and so the differentials of sections project to the identity. This gives us a way to describe the 1-st jet bundle in terms of our previously found 1-jets.

Lemma A.2.2. *Given a smooth fibre bundle (E, π, M), its 1st jet-bundle can be identified as the subspace*

$$J^1\pi \subset J^1(M, E) \cong \mathrm{Hom}(TM, TE)$$

defined as

$$J^1\pi = \{s \mid s \in \mathrm{Hom}(TM, TE),\ d(\pi \circ s) = id_{TM}\}.$$

Or in the language of jet-bundles:

$$J^1\pi = \{j^1_p s \mid j^1_p s \in J^1_p(M, E),\ j^1_p(\pi \circ s) = j^1_p(id)\}.$$

It might be useful to keep the following in mind.

Corollary A.2.3. *$j^1_p s = j^1_p \sigma$ if and only if $ds_p = d\sigma_p$.*

Now let us give some examples of jet bundles, and in particular give a description of the jet bundles of the cotangent bundle.

Example A.2.3.1. Observe the trivial bundle $M \times \mathbb{R}$, smooth sections of the trivial bundle are simply smooth functions $f : M \to \mathbb{R}$. So we get

$$J^1(M \times \mathbb{R}) = \{f^1_p \mid j^1_p f \in J^1_p(M, \mathbb{R}),\ j^1_p(\pi \circ f) = j^1_p(id)\}$$

this is recognizable as the final remark of example A.1.7.3. Hence, we conclude via the same mapping as given previously:

$$J^1(M \times \mathbb{R}) \cong T^*M \times \mathbb{R}.$$

♦

Example A.2.3.2. Let M a smooth manifold and let T^*M be its cotangent bundle. Using the above interpretation of 1-jet bundles we can write

$$J^1(T^*M) = \{\alpha \mid \alpha \in \mathrm{Hom}(M, T^*M),\ d(\pi \circ \alpha) = id_{TM}\}. \tag{A.1}$$

To add some coordinate expression to this fairly abstract expression, let us start with our previously stated results

$$J^1(T^*M) \subset J^1(M, T^*M) \cong T^*M \otimes T(T^*M) \cong \mathrm{Hom}(TM, T(T^*M)).$$

Now if local coordinates for M are again given by $(x_1, \ldots, x_m)$, then at a point $p \in M$ a 1-form λ can be written as

$$\lambda_p = \sum_{i=1}^{m} y_i(p) dx_i \in T^*_p M,$$

where $y_i : M \to \mathbb{R}$. As a result we have a local coordinate system $((x_i),(y_i))$, shorthand for the $2m$-tuple, on the total space T^*M. Which leads to a local expression of $(d\lambda)_p \in T_p^*M \otimes T_{\lambda_p}(T^*M)$ given by:

$$(d\lambda)_p = \sum_{i,j,k=1}^{m} dx_i \otimes (f_{ij}(p)\partial_{x_j} + \hat{g}_{ik}(\lambda_p)\partial_{y_k}), \tag{A.2}$$

where $f_{ij} : M \to \mathbb{R}$ and $\hat{g_{ik}} : T^*M \to \mathbb{R}$. However, by virtue of λ being a section of the cotangent bundle, we have the canonical splitting into its vertical and horizontal part, given by

$$\begin{aligned} T_{\lambda_p}(T^*M) &\cong T_p^*M \oplus T_pM, \\ \sum_k b_k\partial_{y_k} + \sum_i a_j\partial_{x_j} &\mapsto \sum_k b_k dx_k + \sum_i (d\pi)_{\lambda_p}(a_j\partial_{x_j}). \end{aligned}$$

Using this we can rewrite (A.2) as follows:

$$(d\lambda)_p = \sum_{i,j,k=1} dx_i \otimes (f_{ij}(p)(d\pi)_{\lambda_p}(\partial_{x_j}) + g_{ik}(p)dx_k),$$

where we defined $g_{ik} := \hat{g}_{ik} \circ \lambda$. But really by (A.1) the section projects to the identity under the induces map by π. And so the first term must project onto id_{TM}, using that $\sum_{i=1}^{n}(dx_i \otimes \partial_{x_i}) \cong id_{TM}$ we can write

$$(d\lambda)_p = \sum_{i,k=1} dx_i \otimes (f_i(p)\partial_{x_i} + g_{ik}(p)dx_k),$$

combining the found equations, we obtain a local coordinate expression of $j_p^1\lambda = (p, \lambda_p, (d\lambda)_p)$. Now there is a well-defined map from the 1-jet of the 1-form λ to the space of 2-forms by antisymmetrization the above expression, and so we obtain:

$$(p, \lambda_p, (d\lambda)_p) \mapsto \sum_{i,k} \frac{g_{ik} - g_{ki}}{2}(p)(dx_i \wedge dx_k).$$

This is in fact a general construction, there is a mapping:

$$J^1\left(\bigwedge^n T^*M\right) \to \bigwedge^{n+1} T^*M, \tag{A.3}$$

for general jets of n-forms on M. ♦

Example A.2.3.3. It is not hard to see that the kernel of the antisymmetrization mapping (A.3) are precisely those 1-jets in $J^1(\bigwedge^n T^*M)$ which are symmetric. And its image consists of 0 jets of $n+1$-forms, as the construction is analogous to taking the exterior derivative. In fact, this is completely analogous to what we already discussed in the last example of the previous section. So we obtain a short exact sequence of vector bundles

$$0 \to \mathrm{Sym}^r(T^*M) \to J^r\left(\bigwedge^n T^*M\right) \to J^{r-1}\left(\bigwedge^{n+1} T^*M\right) \to 0,$$

which constructs the vector bundle $J^r(\bigwedge^n T^*M) \to J^{r-1}(\bigwedge^{n+1} T^*M)$ with fibre $\mathrm{Sym}^r(T^*M)$.

Again this is a specific example of a more general construction:

$$0 \to \mathrm{Sym}^r(T^*M, E) \to J^r(E) \to J^{r-1}E \to 0,$$

for any vector bundle (E, π, M). ♦

A.3 The Weak and Strong Topology

Having treated jets extensively, the tool central in defining a function space topology, it is time to finally define the topology they induce. As we are dealing with topological spaces, the logical starting point is to define a topology on the space of morphisms between topological spaces, which is the space of continuous functions $C^0(M, N)$ between M and N. There are two common topologies.

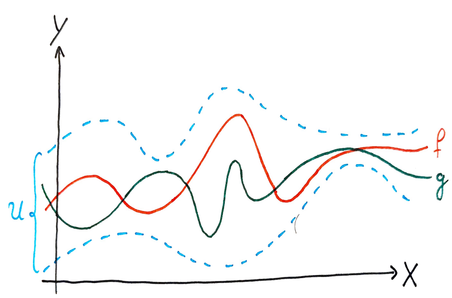

Figure 23: Strong topology base; $f, g \in B(U)$

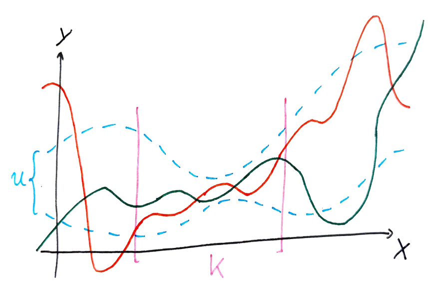

Figure 24: Weak topology sub-base; $f, g \in A(K, U)$

Definition A.3.1. The **weak topology** on $C^0(M, N)$ is defined by taking as a sub-basis the sets:

$$A(K, U) := \{f | f(K) \subset U\},$$

where $K \subset M$ is compact and $U \subset N$ is open. ▲

One can interpret this as the topology of uniform convergence on compact sets. As a result this topology does not describe very well how functions are related to each other "at infinity"; a different topology is to be applied in such a case.

Definition A.3.2. The **strong topology** on $C^0(M,N)$ is defined by taking as a basis the sets:

$$B(U) := \{f|(1 \times f)(X) \subset U\},$$

where now $U \subset X \times Y$ is open in the product topology. ▲

The strong topology also keeps tabs on how functions relate to each other at "infinity". One can readily deduce that in the case that M is compact, these two topologies coincide.

Up to this point we have not yet used jets, however; they are needed to extend these topologies to $C^r(M,N)$ for arbitrary C^r-differentiable manifolds.

Now we are finally ready to define a topology on $C^r(M,N)$. Remark that

$$j^r : C^r(M,N) \to C^0(M, J^r(M,N));$$
$$f \mapsto j^r f,$$

is well-defined and injective. Then using the topology of continuous functions, we have the following two induced topologies on $C^r(M,N)$:

Definition A.3.3. The **weak C^r topology** is the initial topology on $C^r(M,N)$ with respect to the map $j^r f$ and the weak C^0-topology on $C^0(M, J^r(M,N))$. ▲

And:

Definition A.3.4. The **strong C^r topology** is defined as the initial topology on $C^r(M,N)$, with respect to the map $j^r f$, when $C^0(M, J^r(M,N))$ has been endowed with the strong C^0-topology. We also call this the **Whitney C^r-topology** ▲

Not that by injectivity of j^r these topologies are equivalent to viewing $C^r(M,N)$ as a subspace and using the subspace topology. We again remark the topologies agree if M is compact. Note that the topology on $J^r(M,N)$ is induced by its smooth structure, and hence is metrizable. Once chosen a metric d compatible with this topology, we can then write down the weak and strong topology on $C^r(M,N)$ somewhat more explicitly. Namely, the weak topology takes the form of the topology with as sub-basis

$$A(f,K,\varepsilon) := \{g \mid d(j^r_p f, j^r_p g) < \varepsilon \text{ for all } p \in K\},$$

if we take local representatives $\hat{f}$ and $\hat{g}$, and denote by $\|.\|_{\hat{K}}$ the supremum norm over a compact set $\hat{K}$ representing K in euclidean space, we get that the demand is rewritten as

$$A(f,K,\varepsilon) := \{g \mid \|\partial_I \hat{f} - \partial_I \hat{g}\|_{\hat{K}} \leq \varepsilon \text{ for all } |I| \leq r\},$$

that is to say the values of local representatives of f, g and all their possible partial derivatives up to order r are ε-close over each compact set K, convergence in this space can be interpreted as compact convergence up to certain derivative order r. Similarly, the strong topology takes the form of having a basis

$$B(f,\varepsilon) := \{g \mid \|\partial_I \hat{f} - \partial_I \hat{g}\|_{\sup} \text{ for all } |I| \leq r\},$$

that is to say the values of local representatives of f, g and all their possible partial derivatives up to order r are ε-close over the whole of M, convergence in this space can be interpreted as uniform convergence up to certain derivative order r.

Now in the above we assumed M, N had some limit on their differentiability, in this paper however we are mostly interested in the smooth case. We can extend this notion to $r = \infty$ in two different ways.

Definition A.3.5. The C^∞ **topology** on $C^\infty(M, N)$ is the union of all C^r topologies for $r \geq 0$ both weak and strong. ▲

Note that the compact-open topology was defined using sub-bases and the strong topology was defined using bases. If we take the union of only these bases, we get the other topology:

Definition A.3.6. The **strong** C^∞ **topology** on $C^\infty(M, N)$ is the union of all strong C^r topologies for $r \geq 0$. We also call this the **Whitney** C^∞**-topology** denoted by W^∞. ▲

B The de Rham Isomorphism

Recall the de Rham cohomology group $H^n_{dR}(M;\mathbb{R})$ is defined by closed n-forms modulo exact n-forms. Here $\mathbb{R}$ denotes these forms take on real values, we will be particularly interested in whether we will be able to find integer valued 1-forms. Also recall that singular homology $H_n(M;\mathbb{Z})$ was defined as n-cycles modulo n-boundaries. The de Rham isomorphism is a map relating these two groups.

Definition B.1. The **de Rham map** is defined as the map

$$\begin{aligned} \Psi : H^n_{dR}(M;\mathbb{R}) &\to \mathrm{Hom}_{\mathbb{Z}}(H_n(M;\mathbb{Z});\mathbb{R}) \\ [\lambda] &\mapsto \int_\bullet \lambda. \end{aligned}$$

▲

Here $\mathrm{Hom}_{\mathbb{Z}}(H_n(M;\mathbb{Z});\mathbb{R})$ is the space of $\mathbb{Z}$-linear maps from singular homology groups to $\mathbb{R}$. We will not prove this map is an isomorphism, though this is the subject of de Rham's theorem for which many references can be found. We will however give a brief outline of the properties of this map we are interested in, especially in the context of Tischler's theorem.

This map is well-defined, indeed let γ be a representing n-cycle, then we see

$$\int_{n\cdot\gamma} \lambda = n \int_\gamma \lambda \in \mathbb{R},$$

so it is indeed a $\mathbb{Z}$-linear map from and to the right spaces. Now if λ is exact then its cohomology class is 0, so it should map to a zero map. Let us define $d\alpha = \lambda$. Then we can apply Stokes' theorem to obtain

$$\int_\gamma \lambda = \int_\gamma d\alpha = \int_{\partial\gamma} \alpha = \int_0 \alpha = 0.$$

Similarly, if γ is a boundary, let us say $\gamma = \partial\beta$, then its homology class is 0, and so, too, the map should be zero. Again by using Stokes' we see

$$\int_\gamma \lambda = \int_{\partial\beta} \lambda = \int_\beta d\lambda = \int_\beta 0 = 0.$$

So we see that Ψ is indeed a well-defined map.

Now by Universal Coefficient Theorem and the fact that $\mathrm{Ext}(\mathbb{Z}, \mathbb{R}) = 0$, we obtain the following

Lemma B.2.

$$\mathrm{Hom}_\mathbb{Z}(H_n(M;\mathbb{Z});\mathbb{R}) \cong H^n(M;\mathbb{R}).$$

So the de Rham map extends to singular cohomology. As mentioned, de Rham's theorem proves the aforementioned construction with integrals is an isomorphism. Combined with the above, this results in:

$$H^n_{dR}(M;\mathbb{R}) \cong H^n(M;\mathbb{R}).$$

We will from now on drop the $\mathbb{R}$ in notation except for emphasis. Moreover, this isomorphism is natural in the following sense. Given a smooth map $f : M \to N$, we both have an induced map

$$\begin{aligned} f^* : H^n_{dR}(N) &\to H^n_{dR}(M) \\ \beta &\mapsto f^*\beta, \end{aligned}$$

given by the usual pullback of differential forms, and an induced map

$$\begin{aligned} f^* : \mathrm{Hom}_\mathbb{Z}(H_n(N;\mathbb{Z});\mathbb{R}) &\to \mathrm{Hom}_\mathbb{Z}(H_n(M;\mathbb{Z});\mathbb{R}) \\ g &\mapsto g \circ f, \end{aligned}$$

given by the usual pullback of functions. Then it follows for a $\gamma \in H_n(M;\mathbb{Z})$ that

$$(\Psi \circ f^*(\beta))(\gamma) = \int_\gamma f^*\beta = \int_{f\circ\gamma} \beta \circ df = (f^* \circ \Psi(\beta))(\gamma).$$

Now as M is closed manifold, all $n = 1$ (co-)homology groups are finitely generated as $H_1(M;\mathbb{Z})$ is finitely generated, let us define the rank of all spaces by k.

Now recall an Eilenberg-Maclane space denoted by $K := K(G,n)$ was a space whose n'th homotopy group was $\pi_n(K) = G$, and all others equal to 0. For any general CW-complex X there is the following bijection.

Lemma B.3. *For any CW-complex X, and an Eilenberg-Maclane space $K := K(G,n)$ there is a bijection between the nth singular cohomology group and based homotopy classes of maps*

$$H^n(X;G) \leftrightarrow [X,K].$$

Proof. Note first that by definition

$$H^n(K;G) = \mathrm{Hom}(H_n(K;\mathbb{Z}),G).$$

Now by property of Eilenberg-Maclane spaces, $K(G,n)$ is $(n-1)$-connected, and so by the Hurewicz isomorphism we obtain

$$H_n(K;\mathbb{Z}) \cong \pi_n(K) = G.$$

And so $H^n(K,G) \cong \mathrm{Hom}(G,G)$. Similar to how ones argues with the Yoneda lemma, we pick $id_G \in \mathrm{Hom}(G,G)$ and let $u \in H^n(K,G)$ be its preimage under the isomorphism. Now for any map $f : X \to K$ we have an induced map

$$[\mathrm{Hom}(G,G) \cong \mathrm{Hom}(H_n(K;\mathbb{Z}),G)] \xrightarrow{f^*} [\mathrm{Hom}(H_n(X;\mathbb{Z}),G) = H^n(X;G)] .$$

Now the bijection is given by

$$f^*u \leftrightarrow [f].$$

We leave it to the reader to check this is indeed a bijection and is independent of choice of representative. ∎

The following corollary expands this to smooth manifolds.

Corollary B.4. *Any compact smooth manifold admits a CW-complex structure. Thus, for compact smooth M we have $H^n(M;G) \leftrightarrow [M,K]$.*

Now we will use the above to find integer valued 1-forms based on M. Via the de Rham isomorphism $H^1(M;\mathbb{Z}) \cong H^1_{dR}(M;\mathbb{Z})$, we first find a closed 1-form

$$\mu \in H^1_{dR}(S^1;\mathbb{Z})$$

such that for θ, the generator of $H_1(S^1;\mathbb{Z})$, we have

$$\int_\theta \mu = 1.$$

Now applying the corollary together with the fact that S^1 is $K(\mathbb{Z}, 1)$, we can find for any element $\alpha \in H^1_{dR}(M; \mathbb{Z})$ a homotopy class $[f] \in [M, S^1]$ such that

$$f^*\mu = \alpha + dg,$$

where $g \in C^\infty(M)$. On its turn this means that

$$\mathrm{Per}(\alpha) := \left\langle \int_\gamma \alpha \mid \gamma \in H_1(M; \mathbb{Z}) \right\rangle \leq (\mathbb{Z}, +),$$

is a subgroup of the additive integers. Where γ is a closed loop in M.

Definition B.5. The subgroup $\mathrm{Per}(\alpha)$ is called **the group of periods of** α. ▲

Bibliography

[Aeb+13] Beat Aebischer et al. *Symplectic Geometry: an Introduction based on the Seminar in Bern, 1992*. Vol. 124. Birkhäuser, 2013.

[Arn78] VI Arnold. "Methods of classical mechanics". In: *Graduate Texts in Mathematics* 60 (1978).

[BŚ92] Larry Bates and Jedrzej Śniatycki. "On action-angle variables". In: *Archive for rational mechanics and analysis* 120 (1992), pp. 337–343.

[Bur23] Annegret Burtscher. "Introduction to Riemannian Geometry". In: *Lecture Notes: Course module NWI-WB045B* (2023).

[Cal07] Danny Calegari. *Foliations and the geometry of 3-manifolds*. Clarendon press, 2007.

[CN13] César Camacho and Alcides Lins Neto. *Geometric theory of foliations*. Springer Science & Business Media, 2013.

[CR23] Robert Cardona and Ana Rechtman. *Periodic orbits and Birkhoff sections of stable Hamiltonian structures*. 2023. arXiv: 2206.14732.

[Chr17] Austin Christian. "A Short Gallery of Characteristic Foliations". In: 2017.

[CV14] K. Cieliebak and E. Volkov. "Stable Hamiltonian structures in dimension 3 are supported by open books". In: *Journal of Topology* 7.3 (Feb. 2014), pp. 727–770. ISSN: 1753-8416. DOI: 10.1112/jtopol/jtt044. URL: http://dx.doi.org/10.1112/jtopol/jtt044.

[CM05] Kai Cieliebak and Klaus Mohnke. "Compactness for punctured holomorphic curves". In: *Journal of Symplectic Geometry* 3.4 (2005), pp. 589–654.

[CV15] Kai Cieliebak and Evgeny Volkov. "First steps in stable Hamiltonian topology". In: *Journal of the European Mathematical Society* 17.2 (2015), pp. 321–404.

[CDR22] Vincent Colin, Pierre Dehornoy, and Ana Rechtman. *On the existence of supporting broken book decompositions for contact forms in dimension 3*. 2022. arXiv: 2001.01448.

[Col+23] Leonardo Colombo et al. "Liouville-Arnold theorem for contact Hamiltonian systems". In: *arXiv preprint arXiv:2302.12061* (2023).

[CG] Sylvain Courte and Stéphane Guillermou. "Introduction to symplectic and contact geometry". In: ().

[Do 16] Manfredo P Do Carmo. *Differential geometry of curves and surfaces: revised and updated second edition*. Courier Dover Publications, 2016.

[Eli92] Yakov Eliashberg. "Contact 3-manifolds twenty years since J. Martinet's work". In: *Annales de l'institut Fourier*. Vol. 42. 1-2. 1992, pp. 165–192.

[EGH00] Yakov Eliashberg, Alexander Givental, and Helmut Hofer. *Introduction to Symplectic Field Theory*. 2000. arXiv: `math/0010059`.

[ET98] Yakov Eliashberg and William P Thurston. *Confoliations*. Vol. 13. American Mathematical Soc., 1998.

[Etn06] John B Etnyre. "Contact manifolds". In: *Encyclopedia of mathematical physics* 1 (2006), pp. 631–636.

[Etn04a] John B Etnyre. "Convex surfaces in contact geometry: class notes". In: (2004). URL: `http://people.math.gatech.edu/~etnyre/preprints/papers/surfaces.pdf`.

[Etn01] John B Etnyre. "Introductory lectures on contact geometry". In: *arXiv preprint math/0111118* (2001).

[Etn04b] John B. Etnyre. "Lectures on open book decompositions and contact structures". In: *arXiv: Symplectic Geometry* (2004).

[Gei03] Hansjörg Geiges. *h-Principles and Flexibility in Geometry*. American Mathematical Soc., 2003.

[Gei01] Hansjörg Geiges. "A brief history of contact geometry and topology". In: *Expositiones Mathematicae* 19.1 (2001), pp. 25–53.

[Gei08] Hansjörg Geiges. *An introduction to contact topology*. Vol. 109. Cambridge University Press, 2008.

[Gei04] Hansjörg Geiges. *Contact geometry*. 2004. arXiv: `math/0307242`.

[Ghy89] Étienne Ghys. "L'invariant de Godbillon-Vey". In: *Astérisque* 177.178 (1989), pp. 155–181.

[Gir03] Emmanuel Giroux. *Géométrie de contact: de la dimension trois vers les dimensions supérieures*. 2003. arXiv: `math/0305129`.

[GG12] Martin Golubitsky and Victor Guillemin. *Stable mappings and their singularities*. Vol. 14. Springer Science & Business Media, 2012.

[Got82] Mark J Gotay. "On coisotropic imbeddings of presymplectic manifolds". In: *Proceedings of the American Mathematical Society* 84.1 (1982), pp. 111–114.

[Hir12] Morris W Hirsch. *Differential topology*. Vol. 33. Springer Science & Business Media, 2012.

[HZ12] Helmut Hofer and Eduard Zehnder. *Symplectic invariants and Hamiltonian dynamics*. Birkhäuser, 2012.

[Hoh20] Sonja Hohloch. "Integrable Hamiltonian systems". In: *Lecture notes for a course on Integrable Hamiltonian systems* (2020). URL: `https://medialibrary.uantwerpen.be/files/75400/86ec272f-a46c-4d7f-a380-d07a446d48ed.pdf`.

[Hut10] Michael Hutchings. "Taubes's proof of the Weinstein conjecture in dimension three". In: *Bulletin of the American Mathematical Society* 47.1 (2010), pp. 73–125.

[HT09a] Michael Hutchings and Clifford Henry Taubes. "The Weinstein conjecture for stable Hamiltonian structures". In: *Geometry & Topology* 13.2 (2009), pp. 901–941.

[HT09b] Michael Hutchings and Clifford Henry Taubes. "The Weinstein conjecture for stable Hamiltonian structures". In: *Geometry & Topology* 13.2 (Jan. 2009), pp. 901–941. ISSN: 1465-3060. DOI: `10.2140/gt.2009.13.901`. URL: `http://dx.doi.org/10.2140/gt.2009.13.901`.

[KMS13] Ivan Kolár, Peter W Michor, and Jan Slovák. *Natural operations in differential geometry*. Springer Science & Business Media, 2013.

[Lee18] John M Lee. *Introduction to Riemannian manifolds*. Vol. 2. Springer, 2018.

[Lee12] John M Lee. *Smooth manifolds*. Springer, 2012.

[Mas14] Patrick Massot. "Topological methods in 3-dimensional contact geometry". In: *Contact and symplectic topology* 26 (2014), pp. 27–83.

[MS17] Dusa McDuff and Dietmar Salamon. *Introduction to symplectic topology*. Vol. 27. Oxford University Press, 2017.

[MM03] Ieke Moerdijk and Janez Mrcun. *Introduction to foliations and Lie groupoids*. Vol. 91. Cambridge university press, 2003.

[NW23] Ameth NDiaye and Aissa Wade. "Deformations of pairs of codimension one foliations". In: *arXiv preprint arXiv:2306.15060* (2023).

[NBD03] Jet Nestruev, AV Bocharov, and S Duzhin. *Smooth manifolds and observables*. Vol. 220. Springer, 2003.

[Pas12] Federica Pasquotto. "A short history of the Weinstein conjecture". In: *Jahresbericht der Deutschen Mathematiker-Vereinigung* 114 (2012), pp. 119–130.

[Pin22] Álvaro del Pino Gomez. "Syllabus for the course "Topology of geometric structures"". In: (2022).

[Rab78] Paul H Rabinowitz. *Periodic solutions of a Hamiltonian system on a prescribed energy surface*. Mathematics Research Center, University of Wisconsin-Madison, 1978.

[Rut21] Daniel Rutschmann. "Hamiltonian Vector Fields: Definitions and Examples". In: (2021).

[Sau89] David J Saunders. *The geometry of jet bundles*. Vol. 142. Cambridge University Press, 1989.

[SR17] Bruno Scardua and Carlos Arnoldo Morales Rojas. *Geometry, Dynamics and Topology of Foliations: A First Course*. World Scientific, 2017.

[Sil01] Ana Cannas Silva. "Lectures on symplectic geometry". In: *Lecture Notes in Mathematics* 1764 (2001).

[Sul78] Dennis Sullivan. "A foliation of geodesics is characterized by having no "tangent homologies"". In: *Journal of Pure and Applied Algebra* 13.1 (1978), pp. 101–104.

[Tau07] Clifford Henry Taubes. "The Seiberg–Witten equations and the Weinstein conjecture". In: *Geometry & Topology* 11.4 (Oct. 2007), pp. 2117–2202. ISSN: 1465-3060. DOI: 10.2140/gt.2007.11.2117. URL: http://dx.doi.org/10.2140/gt.2007.11.2117.

[Thu72] William P. Thurston. "Noncobordant foliations of S3". In: *Bulletin of the American Mathematical Society* 78 (1972), pp. 511–514.

[Tis70] David Tischler. "On fibering certain foliated manifolds over S1". In: *Topology* 9 (1970), pp. 153–154.

[Tor09] David Martínez Torres. "Contact Geometry and Topology". In: 2009. URL: https://api.semanticscholar.org/CorpusID:209765357.

[Tu17] Loring W Tu. *Differential geometry: connections, curvature, and characteristic classes*. Vol. 275. Springer, 2017.

[Tu11] Loring W Tu. "Manifolds". In: *An Introduction to Manifolds*. Springer, 2011, pp. 47–83.

[Wad75] Andrew W Wadsley. "Geodesic foliations by circles". In: *Journal of Differential Geometry* 10.4 (1975), pp. 541–549.

[Wal16] Charles Terence Clegg Wall. *Differential topology*. Vol. 156. Cambridge University Press, 2016.

[War83] Frank W Warner. *Foundations of differentiable manifolds and Lie groups*. Vol. 94. Springer Science & Business Media, 1983.

[Wei79] Alan Weinstein. "On the hypotheses of Rabinowitz' periodic orbit theorems". In: *Journal of Differential Equations* 33.3 (1979), pp. 353–358. ISSN: 0022-0396.

[Wen16] Chris Wendl. *Lectures on Symplectic Field Theory*. 2016. arXiv: 1612.01009.

[Yaz] Mehdi Yazdi. "Foliations and Three-Dimensional Manifolds". In: ().