A tourist's account of characteristic foliations on convex surfaces in 3-D contact geometry

Luke Volk

Supervised by
Dr. Maia Fraser

A thesis submitted in partial fulfilment of the requirements for the
Master of Science in Mathematics[1]

uOttawa
Department of Mathematics and Statistics
Faculty of Science
University of Ottawa

[1] The Master's program is a joint program with Carleton University, administered by the Ottawa-Carleton Institute of Mathematics and Statistics

Acknowledgements

It is not uncommon for tourists to need tour guides—especially when they are likely to get lost—and so I would like to thank Maia Fraser for taking on the role of guiding me through this region of the awe-inspiring mathematical jungle. To her I owe many explanations, simplifications, and geometric insights which have found their way into this document.

I also extend my appreciation to Michelle, who probably has heard the phrase "an overtwisted disk is like a pizza" too many times to count. At the risk of making the acknowledgements of a thesis in pure mathematics into something syrupy—your eagerness to listen made a difference.

Abstract

We begin with a rapid introduction to the theory of contact topology, first spending more time than you would probably want on developing the notion of contact manifold before launching right into the thick of the theory. The tools of characteristic foliations and convex surfaces are introduced next, concluding with an overview of Legendrian knots in contact 3-manifolds.

Next, we develop a number of lemmas as tools for dealing with characteristic foliations, concluding with some sightseeing with regards to the theory of so-called "movies", allowing a glimpse into the workings of a theorem due to Colin:

> Two smoothly isotopic embeddings of S^2 into a tight contact 3-manifold inducing the same characteristic foliation are necessarily contact isotopic.

We finish with an original observation that Colin's theorem can be used to replace a key step in Eliashberg and Fraser's classification of topologically trivial knots, thus providing an alternate proof of that result and thereby highlighting the power of the aforementioned theorem. We provide a simplification of this proof using intermediate results we encountered along the way.

Contents

Notational conventions

S_ξ	Characteristic foliation of the hypersurface S induced by ξ.
$\operatorname{div}_\Omega \mu$	Divergence of μ with respect to the volume form Ω.
$\iota_v \mu$	Interior product of v with μ (contraction).
$\Omega^k(M)$	Set of smooth k-forms on M.
$[\cdot,\cdot]$	Lie bracket of vector fields.
$\mathcal{L}_v\, \mu$	Lie derivative of μ in the direction of v.
$f_{*,p}$	Pushforward (differential) of a map f at a point p.
μ^*	Pullback of the tensor (field) μ.
$\mathrm{r}(K)$	Rotation number of a knot K.
$C^\infty(M)$	Algebra of smooth functions on M.
ξ_0	Standard contact structure on $\mathbb{R}^3$, $\xi_0 = \ker(dz - y\,dx)$.
TM	Tangent bundle of the manifold M.
T_pM	Tangent space of the manifold M at a point $p \in M$.
$\mathrm{tb}(K)$	Thurston-Bennequin invariant of a knot K.

Magic tricks

$\mathcal{L}_v = \iota_v \circ d + d \circ \iota_v$	Cartan's formula.

1 Introduction

1.1 Overview

What follows is a more nuanced overview of each section of the thesis.

1.1.1 Introduction to contact geometry

Some time is taken initially to develop the notions of integrability and to give an adequate background to Frobenius's theorem, the result to which we owe in part the definition of a contact manifold.

The development of integrability culminates with the definition of a contact manifold, an object whose structure is characterised by its extreme lack of integrability. Following the basic definitions related to the notion of contact manifolds, there is an interlude giving some of the various connections to contact geometry's sister-subject symplectic geometry. We then launch into a discussion of the characteristic foliation, a geometric tool of surfaces used to study contact structures. We finish off with Giroux's incredibly useful elimination lemma for manipulating characteristic foliations.

We introduce the study of convex surfaces, a generic property for surfaces in 3-dimensional contact geometry that was introduced by Giroux and has proven itself to be a powerful tool. We show how convex surfaces correspond to discrete combinatorial objects called dividing curves which greatly simplify the study of characteristic foliations. A complete algorithm for drawing dividing curves is given.

Finally, we take a closer look at Legendrian knots. Some discussion of the wavefront projection is given, followed by the explanation and discussion of some classic topological notions in knot theory which are sufficient for defining the so-called classic invariants of Legendrian knots.

1.1.2 Machinery

The first tool we develop is a deformation lemma which allows us to locally "bend" the leaves of the characteristic foliation of a surface. Moreover, the resulting porism asserts that this deformation is continuous in the sense that we can fine-tune it as much as we would like. An application of this deformation lemma shows how one can break saddle connections.

The second tool developed is the corner lemma which allows us, with an application of our deformation lemma, to use piece-wise linear models for surfaces, and then correct the foliation after smoothing so that conclusions of the foliation of the piecewise linear surface may be carried over to the smoothed version.

The third tool is the pivot lemma, allowing us flexibility with regards to the leaves of an elliptic point which are smoothly connected. This allows more control of the introduction of singularities. The proof of the pivot lemma makes use of piecewise linear constructions with a subsequent application of the corner lemma.

The tools we developed are then applied to an observation that a closed leaf in a characteristic foliation of a sphere necessitates the existence of an overtwisted disk. We use a result of Sotomayor's and the build up of previous observations to describe a generic "movie" of foliations of a surface.

Finally, we survey some results of movies which allow a glimpse into the proof of Colin's theorem:

> Two smoothly isotopic embeddings of S^2 into a tight contact 3-manifold inducing the same characteristic foliation are contact isotopic.

This makes use the previously built up tools and observations, as well as techniques from earlier proofs.

1.1.3 Trivial Legendrian knots

The proof of Eliashberg and Fraser's classification of topologically trivial Legendrian knots begins with the standardization of Seifert disks. We demonstrate a process for giving a Seifert disk a normal form, allowing us to readily see the ways in which Seifert disks depend on the values $(\mathrm{tb}, \mathrm{r})$ of the classical invariants, and how disks corresponding to identical values of $(\mathrm{tb}, \mathrm{r})$ might differ.

A summary of the proof in Eliashberg and Fraser is then given in stepwise fashion. We show how Colin's theorem can be used to replace one step in the proof, and then conclude with yet another alternative which only uses intermediate results developed towards understanding Colin's proposition in previous sections.

2 Introduction to contact geometry

2.1 Integrability

2.1.1 Integrability and vector fields

Recall that we can view a vector field v as a smooth section $v\colon M \to TM$ of the tangent bundle of a manifold M. To each point $p \in M$, v associates a tangent vector $v_p := v(p) \in T_pM$. An **integral curve** to v is a smooth curve $\gamma\colon (a,b) \to M$ such that for each $s \in (a,b)$, we have:

$$\gamma'(s) = v_{\gamma(s)}.$$

In other words, v is tangent to $\gamma(s)$ at each time $s \in (a,b)$.

Example 2.1.1. Consider $\mathbb{R}^2$ with the vector field v defined by $v(x,y) = (-y,x)$. The integral curves are concentric circles around the origin.

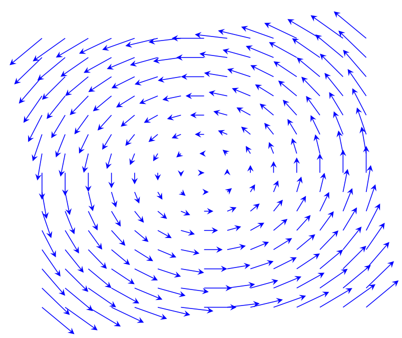

Any vector field $v\colon M \to TM$ on a manifold M gives rise locally to a system of first order ordinary differential equations. To see this, consider a chart $(U \subseteq M, \varphi)$ on M, and consider the induced map $f\colon \varphi(U) \to T\varphi(U) \cong T\mathbb{R}^n$ which makes the following diagram commute:

$$\begin{tikzcd} TM \arrow[r, "\varphi_*"] & T\mathbb{R}^n \\ U \subseteq M \arrow[u, "v"] \arrow[r, "\varphi"'] & \varphi(U) \subseteq \mathbb{R}^n \arrow[u, dashed, "f"'] \end{tikzcd} \qquad f(\varphi(p)) := \varphi_{*,p}(v(p)), \quad p \in U.$$

This f is—in some sense—a linearised version of v. As a result, any curve $\gamma\colon (a,b) \to U$ is an integral curve for v if and only if

$$(\varphi \circ \gamma)'(t) = f(\varphi(\gamma(t))).$$

With the identification of $T\varphi(U)$ with $\mathbb{R}^n$, we get a first order ODE of the form:

$$\frac{dy}{dt}(t) = f(y(t)), \qquad y = \varphi \circ \gamma.$$

The Picard-Lindelöf theorem guarantees the existence and uniqueness of a solution to the system inherent to f. In particular, for each $p \in \varphi(U)$ there exists real numbers $a_p < 0 < b_p$ and a unique integral curve $F_p(t)\colon (a_p, b_p) \to \varphi(U)$ of maximal length passing through p. The unique maximal **local flow** of f is the function $F\colon V \to \varphi(U)$ whose domain V is the open subset

$$V := \bigcup_{p\in\varphi(U)} (a_p, b_p) \times \varphi(U) \quad \subseteq \mathbb{R} \times \varphi(U),$$

where F is defined by:

$$F(s,p) := F_p(s).$$

The function F satisfies the following properties:

(i) F is smooth.

(ii) For all $p \in \varphi(U)$ we have $F(0,p) = p$.

(iii) For all $(s,p) \in V$ we have:

$$\frac{d}{dt} F(s,p) = f(F(s,p)).$$

Now any local flow for f pulls back along φ to a local flow Φ for v, satisfying an analogous set of conditions as F. An argument can be made for patching the local flows of charts together on the manifold. The conclusion is the following:

Observation 1. Every vector field integrates to a unique maximal flow[a].

[a] As it turns out, on a compact manifold, the flow is globally defined on $\mathbb{R} \times M$ (v thus being said to be **complete**), and hence the flow forms a one-parameter group of diffeomorphisms, $\{\Phi_t := \Phi(t,\cdot) \mid t \in \mathbb{R}\}$.

This result does not seem initially remarkable until you begin to consider a more general object of study, called d-fields.

2.1.2 Integrability and d-fields

A smooth, nowhere-vanishing vector field v on M corresponds to a rank 1 smooth subbundle of TM. That is, $\eta := \text{span}\{v\}$ is a subspace of $\text{vect}(M)$ such that:

$$\eta_p := \{u_p \in T_pM \mid u \in \eta\}$$

is 1-dimensional for all p, and so

$$\eta = \bigsqcup_{p\in M} \eta_p$$

is a smooth vector bundle in its own right, as well as an embedded submanifold of TM. The notion of integrability for vector fields passes to the case of rank 1 smooth subbundles η of TM where we say a 1-submanifold $\Gamma \subseteq M$ is an integral submanifold of η if for each $p \in \Gamma$, we have $T_p\Gamma = \eta_p$. That is to say, η is "everywhere tangent" to Γ.

Example 2.1.2. Note that if η is the rank 1 subbundle of TM corresponding to the nowhere-vanishing vector field v, then the image of any integral curve of v constitutes an integral submanifold of η.

Define a **d-field** to be a smooth subbundle of TM of rank d. The notion of integral submanifolds extends to d-fields. We will define an **integral submanifold** of M to a d-field η to be any k-submanifold $\Gamma \subseteq M$ such that for each $p \in \Gamma$ we have:

$$T_p\Gamma \subseteq \eta_p,$$

for some $0 < k \leqq d$, and in particular we say that η is **integrable** if each point $p \in M$ is contained in an integral d-submanifold[2] to η.

Example 2.1.3. Consider $\mathbb{R}^3$ and a 2-field (that is, a plane field), $\eta \subset T\mathbb{R}^3$. To any $p \in \mathbb{R}^3$, you can associate to the plane η_p a line which is normal to η_p. Inherent to this line is a projection map which takes any tangent vector in $T_p\mathbb{R}^3$ and projects it to the line normal to η_p, making a map

$$\alpha_p \colon T_p\mathbb{R}^3 \to \mathbb{R},$$

which gives a 1-form α on $\mathbb{R}^3$ with the property that:

$$v_p \in \eta_p \iff \alpha_p(v_p) = 0.$$

That is, a 1-form α gives a 2-field as its kernel, $\ker \alpha$!

This strategy gives us a plethora of examples of 2-fields in $\mathbb{R}^3$, but also gives us examples[3] of $(n-1)$-fields (so-called hyperplane fields) on n-manifolds! It is also important to note that multiplying a 1-form by a nowhere-vanishing function will not change the (hyper)plane field associated to it.

Example 2.1.4. Consider $\mathbb{R}^3$ with the 2-field ξ_0 given as the kernel of the 1-form $\alpha := dz - y\,dx$. It is easy to calculate that for $p = (p_1, p_2, p_3) \in \mathbb{R}^3$:

$$(\xi_0)_p = \ker \alpha_p = \operatorname{span}\left\{ \frac{\partial}{\partial y}(p), \frac{\partial}{\partial x}(p) + p_2 \frac{\partial}{\partial z}(p) \right\}.$$

Along the x,z-plane where $p_2 = 0$, $(\xi_0)_p$ can be identified as the x,y-plane. If we look at the lines in the x,y-plane (where $p_3 = 0$) which are parallel to the y-axis, $(\xi_0)_p$ rotates (if you are viewing from the origin) counter-clockwise as $p_2 \to \infty$ to get closer and closer to the y,z-plane.

[2]That is, we allow integral submanifolds of dimension $1 \leqq k \leqq d$, however we say that a d-field is "integrable" if and only if each point of M is contained in an integral manifold of maximal dimension d. This is not exactly standard, but helps simplify the terminology.

[3]But not *every* example! Because $\mathbb{R}^3$ does not admit non-orientable line bundles, $T\mathbb{R}^3/\xi$ cannot be non-orientable for a plane field ξ. Thus ξ itself must be orientable. Picking a positive normal for ξ_p for each $p \in \mathbb{R}^3$, we can define a globally 1-form. However, for a general M, this bundle quotient might admit non-orientable line bundles.

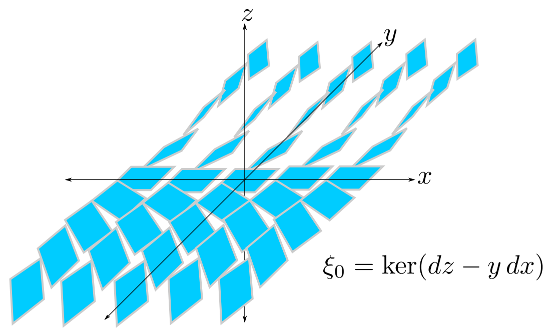

Suppose that Γ is an integral surface of ξ_0 which contains the origin. Because Γ is tangent to the vector fields $v = \frac{\partial}{\partial y}$ and $w = \frac{\partial}{\partial x} + y\frac{\partial}{\partial z}$ at the origin, there is a small neighbourhood containing an integral curve to v through 0 (being a small segment of the y-axis), and another small neighbourhood around the origin containing an integral curve to w through 0 (being a small segment of the x-axis). Thus Γ must contain a tiny open set of the x, y-plane, but then Γ would not be integral as it would transversely intersect with ξ_p for p off of the x axis! So ξ_0 is apparently not integrable!

Observation 2. Despite line fields always being integrable, d-fields are in general not necessarily integrable for $d \geqq 2$.

Thus we have learned that the case of line fields always being integrable *is indeed* a rather remarkable happening.

2.1.3 Frobenius's theorem

Now we will work towards understanding integrability a bit more, since apparently we did not grok it from considering the case of line fields.

Suppose that $\eta \subseteq TM$ is an integrable d-field on M and let $v, w\colon M \to \eta$ be sections of η. As η is integrable, for each $p \in M$ there exists an integral submanifold Γ of dimension d such that:

$$T_p\Gamma = \eta_p.$$

Note that we can view Γ as an embedded submanifold $e\colon \Gamma \to M$ where $e_{*,q}$ is an isomorphism for each $q \in \Gamma$, so we rephrase the defining property of Γ as:

$$e_{*,q}(T_q\Gamma) = \eta_p, \qquad \text{with } e(q) = p. \tag{$\star$}$$

In particular, we have the existence of sections $\tilde{v}, \tilde{w}\colon \Gamma \to \eta$ such that:

$$e_{*,q}(\tilde{v}_q) = v_{e(q)}, \quad \text{and} \quad e_{*,q}(\tilde{w}_q) = w_{e(q)},$$

and so (see Proposition A.1.1) if $e(q) = p$ then we have:

$$[v, w]_p = e_{*,q}([\tilde{v}, \tilde{w}]_q),$$

which is in η_p by ($\star$). Thus for any sections v, w of η, we have $[v, w]$ is a section of η. A d-field like η with this property is called **involutive**.

Observation 3. Every integrable d-field is involutive.

Example 2.1.5. Considering the plane field $\xi_0 := \ker(dz - y\,dx)$ in $\mathbb{R}^3$ we explored in Example 2.1.4, we see that:

$$\left[\frac{\partial}{\partial y}, \frac{\partial}{\partial x} + y\frac{\partial}{\partial z}\right] = \frac{\partial}{\partial z},$$

and thus it is not involutive. If you were not convinced before that it was non-integrable, now you should be!

Is there perhaps a deeper connection[4] between integrable and involutive d-fields?

Example 2.1.6. Consider a plane field η on a 3-manifold M. Locally in a neighbourhood U of $p \in M$, we have that there exists non-vanishing vector fields v, w such that:

$$\eta|_U = \mathrm{span}\{v, w\}.$$

Moreover, we can[5] then find a coordinate chart $V \subseteq U$ containing p with coordinates $\{y_1, y_2, y_3\}$ such that:

$$v|_V = \frac{\partial}{\partial y_1}.$$

Picking a $C := y_1(p) \in \mathbb{R}$, we define a 2-submanifold of M:

$$N := \{x \in V \mid y_1(x) = C\}.$$

In general, w might not take values in TN, and we would like to force it to. To this end, project w to TN via taking:

$$u_1 := v, \quad \text{and} \quad u_2 := w - v(y_1)v,$$

noting that these still give a local frame, $\eta|_V = \mathrm{span}\{u_1, u_2\}$. Moreover, note that $u_2(y_1) = 0$ and so the nowhere-vanishing vector field u_2 takes values in TN. It thus defines a line field on N:

$$\xi = \mathrm{span}\{u_2\},$$

[4]Perhaps a bit more geometric than the example which follows: if a d-field is involutive, it is locally represented by a Lie subalgebra of $\mathrm{vect}(M)$. Using the integrability of line-fields, one can use the flow of the vector fields as a Lie group which acts (as $\mathbb{R}^d$) on M, inducing a foliation on M, the leaves being the orbits of this action. This action (and hence the foliation into leaves) is not well-defined unless the flows commute, which is the case for the involutive d-fields. Thus we have a hunch that perhaps involutive implies integrable.

[5]Just use the flow of v to "straighten" it out.

which by Picard-Lindelöf, we know is integrable. Thus there is a 1-dimensional integral submanifold $S \subset N$ to the line field η. This additionally gives us a coordinate system z_2, z_3 in a neighbourhood of $p \in N$ and a $C' \in \mathbb{R}$ such that:

$$S = \{x \in N \mid z_3(x) = C'\}.$$

Define the projection $\pi \colon V \to N$ via:

$$\pi(y_1(x), y_2(x), y_3(x)) = (y_2(x), y_3(x)),$$

and then define new coordinates on V by:

$$x_1 := y_1, \quad x_2 := y_2 \circ \pi, \quad x_3 := y_3 \circ \pi.$$

Now we extend S along x_1 to a 2-submanifold $\Gamma \subset V$, setting $C' := x_3(p) \in \mathbb{R}$:

$$\Gamma := \{x \in V \mid (x_1(x), x_2(x)) \in N, \quad x_3 = C'\}.$$

By construction, u_1 is tangent to Γ, as $u_1 = \partial/\partial x_1$. It is not necessarily the case that u_2 is also tangent to Γ (if it were, then we would immediately have an integral submanifold passing through p.

What conditions can we impose to force Γ to be an integral manifold? Let us consider the derivative:

$$\frac{\partial}{\partial x_1} u_2(x_j) = u_1(u_2(x_j)) - u_2(u_1(x_j)) = [u_1, u_2](x_j), \qquad 2 \leqq j \leqq 3.$$

If we by chance had that η was involutive, then we can write:

$$\frac{\partial}{\partial x_1} u_2(x_j) = a\, u_1(x_j) + b\, u_2(x_j) = b\, u_2(x_j),$$

which is an ODE with an obvious unique solution. Moreover, we know that on N, $u_2(x_j)$ vanishes, and so this initial value problem gives us that:

$$u_2(x_j) = 0,$$

and thus u_2 is also tangent to Γ. Hence we have shown that Γ is an integral submanifold to η, containing p!

The construction in the example is by no means restricted to plane fields. We can reap a magnificent crop from the seed Picard-Lindelöf has sowed. An inductive argument (see [21]) much alike the preceding example gives the following theorem:

Theorem 2.1.7 (Frobenius's theorem). *A d-field is integrable if and only if it is involutive.*

Let us consider only 3-manifolds for now. Suppose ξ is an integrable plane field on a 3-manifold M. For each p, we can pick a Riemannian metric g and a nowhere-vanishing section x of $\xi^g := \{v_q \in TM \mid \forall w_q \in \xi_q \,.\, g(v_q, w_q) = 0\}$ defined in (at least[6]) a neighbourhood U of p. Then on U, ξ is locally represented as the kernel of the nowhere-vanishing 1-form α defined by:

$$\alpha(v) := g(x, v).$$

Note that this argument is not unique to 3-manifolds.

> **Observation 4.** Any hyperplane field is given locally as the kernel of a 1-form.

If v and w are vector fields which take values in an integrable plane field $\xi = \ker \alpha$, then (see Proposition A.1.2):

$$d\alpha(v, w) = v(\alpha(w)) - w(\alpha(v)) - \alpha([v, w]), \qquad (\star\star)$$

where the right-hand side is zero, as $[v, w]$ also takes values in ξ by Frobenius's theorem. Hence $d\alpha$ is 0 on ξ as well. Conversely (without assuming integrability), if $d\alpha$ annihilates ξ, then $\alpha([v, w]) = 0$ for any vector fields taking values in ξ, thus the hyperplane field is involutive.

> **Observation 5.** If a hyperplane field ξ is given locally as the kernel of α, then ξ is involutive (and thus integrable) if and only if $d\alpha$ annihilates ξ.

In particular for when M is a 3-manifold, at $p \in M$, take a local frame of $\{v, w\}$ of $\xi = \ker \alpha$ in a neighbourhood U of p and extend it with a local section y such that $TU = \xi \oplus \mathbb{R}y$. If ξ is integrable, then Observation 5 tells us that:

$$(\alpha \wedge d\alpha)(y, v, w) = \alpha(y) d\alpha(v, w) = 0.$$

Hence $\alpha \wedge d\alpha = 0$.

Conversely, if $\alpha \wedge d\alpha = 0$, then for any vector fields v, w which take values in ξ, we have from ($\star\star$) that:

$$d\alpha(v, w) = -\alpha([v, w]),$$

and we have a vector field y with $\alpha(y) \neq 0$ such that:

$$0 = (\alpha \wedge d\alpha)(y, v, w) = \alpha(y) d\alpha(v, w) = -\alpha(y)\alpha([v, w]),$$

so $\alpha([v, w]) = 0$, and hence ξ is involutive, and thus integrable!

[6]The obstacle to admitting a global section would be the orientability of ξ^g.

Observation 6. If M is a 3-manifold with a hyperplane field ξ, then if locally $\xi = \ker\alpha$ for a 1-form α, then:

$$\xi \text{ is integrable } \iff \alpha \wedge d\alpha = 0.$$

Example 2.1.8. Again, consider the example of the plane field $\xi_0 := \ker\alpha$ on $\mathbb{R}^3$ given by $\alpha := dz - y\,dx$. We see:

$$\alpha \wedge d\alpha = (dz - y\,dx) \wedge (-dy \wedge dx) = dx \wedge dy \wedge dz,$$

which is clearly non-zero! If you are not convinced that this particular plane field is non-integrable, I now resign.

2.1.4 Complete non-integrability

If it would seem unfair to the reader to have spent all this time so far building up the theory of integrability only to now tell you we are turning it all upside down to consider its antithesis, non-integrability, then the author has some bad news for you. However, the theory we have developed so far is anything but useless to us with regards to non-integrability!

Recall that we observed that a hyperplane field ξ on an m-manifold M, given locally as the kernel of α, is integrable if and only $d\alpha|_\xi = 0$. It follows that if we desired ξ to not have a snowball's chance in hell of being integrable, then it is an obvious choice to guarantee that $d\alpha$ is non-degenerate on ξ.

This would make $d\alpha|_\xi$ a non-degenerate, closed 2-form—that is, a symplectic[7] form on ξ! Non-degeneracy implies that ξ must have even dimension (see Corollary A.2.2), and thus forces M to be of odd dimension, $m = 2n + 1$.

Observation 7. If M is a manifold with a hyperplane field ξ given locally as the kernel of α, then $d\alpha|_\xi$ being non-degenerate implies M is odd-dimensional.

A closed 2-form ω being symplectic (read: non-degenerate, since we assume ω is closed) is equivalent to ω^n being a volume form (see Proposition A.2.5) and hence the non-degeneracy of $d\alpha|_\xi$ is equivalent to $(d\alpha)^n|_\xi \neq 0$ being a volume form on ξ. Using a little linear algebra magic, Proposition A.1.3 tells us that $(d\alpha)^n|_\xi \neq 0$ is equivalent to $\alpha \wedge (d\alpha)^n \neq 0$.

Definition 2.1.9. If M is a manifold of dimension $2n+1$ with a hyperplane field ξ given locally as the kernel of α, then we say ξ is **completely non-integrable** if and only if:

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

[7] In the case the reader is unfamiliar with the basic definitions of symplectic geometry, we have provided everything we will need for this paper in Appendix A.2.

Observation 8. Let ξ be a hyperplane field on a $(2n+1)$-dimensional manifold M given locally as the kernel of α. Then:

$$\begin{aligned} \xi \text{ completely non-integrable} &\iff \alpha \wedge (d\alpha)^n \neq 0, \\ &\iff (d\alpha)^n|_\xi \neq 0, \\ &\iff d\alpha|_\xi \text{ is non-degenerate (i.e. symplectic).} \end{aligned}$$

When $n = 1$ and M is a 3-manifold, the notion of ξ being completely non-integrable is equivalent to the property that for each $p \in M$, there does not exist an integral surface containing p. However, this is not in general true.

Proposition 2.1.10. *If M is a $(2n+1)$-manifold with a completely non-integrable hyperplane field ξ, then any integral submanifold to ξ has dimension at most n.*

Proof. Suppose that ξ be given locally as the kernel of α, so by the preceding observation, $d\alpha|_\xi$ is a symplectic form on ξ. If Γ is an integral submanifold to ξ, then for each $p \in \Gamma$ we have that $d\alpha|_{T_p\Gamma} = 0$. Hence $T\Gamma$ is an isotropic submanifold of ξ, and thus (see Proposition A.2.6):

$$\dim T\Gamma \leqq \frac{1}{2} \dim \xi = 2n \implies \dim \Gamma \leqq n.$$

■

Note that when $n = 1$, then an integral submanifold Γ to a completely non-integrable hyperplane field ξ is of maximum dimension 1. If we consider $n = 2$, then a completely non-integrable hyperplane field has integral submanifolds Γ where $\dim \Gamma \leqq 2$, however this does not preclude the possibility that your manifold is non-integrable "at each point" (that there does not exist an integral submanifold of maximal dimension 4 containing each point), but it admits an integral submanifold of dimension 3. Thus in general, completely non-integrable is the other extreme of a spectrum to "integrable".

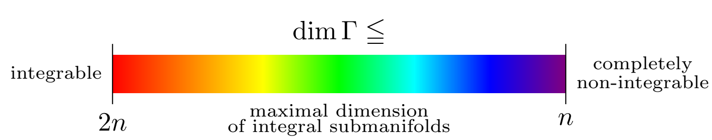

In the case of $n = 1$, this "spectrum" is the shortest possible.

And now we have worked up enough material to state, with full generality, the object of interest: contact manifolds! For more on the general theory of manifolds and a more complete exposition on integrability and Frobenius's theorem, the reader is advised to consult [15] or [21].

2.2 Contact topology

2.2.1 Contact manifolds

Definition 2.2.1. A manifold M endowed with a completely non-integrable hyperplane field ξ is called a **contact manifold**. We refer to ξ as its **contact structure**.

If ξ is given as the kernel of a globally-defined 1-form α, we call α a **contact form** (such a 1-form is not unique). The existence of a globally-defined 1-form is equivalent to the co-orientability of ξ (the existence of a trivial normal bundle to ξ), in which case we will say the contact manifold (or structure) is **co-oriented**.

Recall that by Observation 7, a contact manifold (M, ξ) is necessarily odd-dimensional. If M is $(2n+1)$-dimensional where ξ is given locally by $\ker \alpha$, then we observe the following:

- If n is even, then one can calculate that for $f \in C^\infty(M)$:

  $$(d(f\alpha))^n = f^n (d\alpha)^n,$$

  and so the volume form $(d\alpha)^n|_\xi$ does not depend on the sign of α. Hence there is a natural orientation for ξ.

- If n is odd, then one can calculate that for $f \in C^\infty(M)$:

  $$f\alpha \wedge (d(f\alpha))^n = \lambda^{n+1} \alpha \wedge (d\alpha)^n,$$

  and so the volume form $\alpha \wedge (d\alpha)^n$ on M does not depend on the sign of α. Hence there is a natural orientation for M.

Observation 9. A contact manifold (M, ξ) of dimension $2n+1$ has a globally-defined contact form if and only if its normal bundle is trivial, which is if and only if both M and ξ are orientable.

Moreover, if n is even, then M is naturally oriented, and if n is odd, then ξ is naturally oriented.

It is also important to note that if α locally defines a contact structure, then so does $f\alpha$ for any non-vanishing function $f \in C^\infty(M)$. The condition $\alpha \wedge (d\alpha)^n \neq 0$ on the contact form is sometimes referred to as the **contact condition**, and it implies that $\alpha \wedge (d\alpha)^n$ is either always positive, or always negative. If $\alpha \wedge (d\alpha)^n > 0$, we call $\ker \alpha$ a **positive contact structure**, and in the case $\alpha \wedge (d\alpha)^n < 0$, we call $\ker \alpha$ a **negative contact structure**.

Example 2.2.2 (Standard contact structure on $\mathbb{R}^{2n+1}$)**.** Consider $\mathbb{R}^{2n+1}$ with the co-orientable contact structure ξ_0 given as the kernel of

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

where $\mathbb{R}^{2n+1}$ is given the coordinate system $(x_1, \ldots, x_n, y_1, \ldots, y_n, z)$. This is known as the **standard contact structure of** $\mathbb{R}^n$. Having made it this far, the reader should already be very familiar with the complete non-integrability of the case for $n = 1$.

Example 2.2.3. Define another contact structure ξ_1 on $\mathbb{R}^3$ as the kernel of the 1-form

$$\alpha := dz + r^2\, d\theta,$$

where $\mathbb{R}^3$ is given cylindrical coordinates, (r, θ, z).

We have observed that integral submanifolds of a contact manifold (M, ξ) of dimension $2n + 1$ are at most dimension n. These submanifolds constitute isotropic submanifolds of the symplectic manifold $(\xi, d\alpha)$, appearing "the same in all directions" from the point of view of $d\alpha$. The integral submanifolds $L \subset M$ of maximal dimension n are called **Legendrian submanifolds**.

Just as symplectic topology focuses on isotropic submanifolds of maximal dimension called Lagrangian submanifolds, a fair share of contact topologists study contact manifolds through their Legendrian submanifolds.

Example 2.2.4. Suppose $\mathbb{R}^3$ is given the standard contact structure, $\xi_0 := \ker(dz - y\, dx)$. Consider the embedding $K \colon S^1 \to \mathbb{R}^3$ defined by:

$$K(t) := \begin{cases} \frac{1}{2}\left(1 - \cos(2t), -\sin(2t), \frac{1}{8}\sin(4t) - \frac{1}{2}t\right), & t \in [0, \pi]; \\ \frac{1}{2}\left(\cos(2t) - 1, -\sin(2t), \frac{1}{2}t - \frac{1}{8}\sin(4t) - \pi\right), & t \in [\pi, 2\pi]. \end{cases}$$

It is easily checked[8] that K has a derivative

$$K'(t) = \begin{cases} \left(\sin(2t), -\cos(2t), -\frac{1}{2}\sin^2(2t)\right), & t \in [0, \pi]; \\ \left(-\sin(2t), -\cos(2t), \frac{1}{2}\sin^2(2t)\right), & t \in [\pi, 2\pi], \end{cases}$$

where it is easy to see that:

$$K'(t) = -\cos(2t)\begin{bmatrix} 0 \\ 1 \\ 0 \end{bmatrix} \pm \sin(2t)\begin{bmatrix} 1 \\ 0 \\ -\frac{1}{2}\sin(2t) \end{bmatrix},$$

and thus is everywhere tangent to ξ_0. The embedding K of the circle is an example of a Legendrian **unknot**.

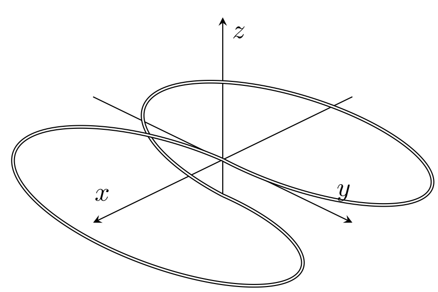

[8]Because I hate trigonometric identities *at least as much* as the reader, I will include that I used $\frac{1}{4}(1 - \cos(4t)) = \frac{1}{4}(1 - (1 - 2\sin^2(2t))) = \frac{1}{2}\sin^2(2t)$.

Now that we have some examples of contact manifolds, we explicitly define the isomorphisms of interest, called contactomorphisms, which are analogous to the symplectomorphisms of symplectic topology (in that the push-forward preserves the structure of interest).

Definition 2.2.5. Let (M_1, η_1) and (M_2, η_2) be two contact manifolds. A diffeomorphism $f\colon M_1 \to M_2$ is a **contactomorphism** if

$$f_*(\eta_1) = \eta_2.$$

If such a map exists, the manifolds are said to be **contactomorphic**. If η_1 and η_2 are given locally as the kernel of α_1 and α_2 respectively, then this is equivalent to requiring that

$$f^*\alpha_2 = g\alpha_1,$$

for a non-vanishing function g on M_1.

Example 2.2.6. Consider $\mathbb{R}^3$ with contact structures ξ_0 and ξ_1 given as the kernels of $\alpha_0 := dz - y\,dx$ and $\alpha_1 := dz + x\,dy - y\,dx$ respectively. Consider the map $f\colon \mathbb{R}^3 \to \mathbb{R}^3$ given by:

$$f(x, y, z) = \frac{1}{2}(x + y, y - x, 2z - xy).$$

Note that for $p = (p_1, p_2, p_3) \in \mathbb{R}^3$, we have that:

$$[f_{*,p}] = \frac{1}{2}\begin{bmatrix} 1 & 1 & 0 \\ -1 & 1 & 0 \\ -p_2 & -p_1 & 2 \end{bmatrix},$$

and so we see:

$$\begin{aligned}
(f^*\alpha_1)_p \begin{bmatrix} v_1 \\ v_2 \\ v_3 \end{bmatrix} &= (\alpha_1)_{f(p)} \left([f_{*,p}] \begin{bmatrix} v_1 \\ v_2 \\ v_3 \end{bmatrix} \right) \\
&= \frac{1}{2}(\alpha_1)_{f(p)} \begin{bmatrix} v_1 + v_2 \\ v_2 - v_1 \\ 2v_3 - p_2 v_1 - p_1 v_2 \end{bmatrix} \\
&= \frac{1}{2}(2v_3 - p_2 v_1 - p_1 v_2 + \frac{1}{2}((p_1 + p_2)(v_2 - v_1) - (p_2 - p_1)(v_1 + v_2))) \\
&= \frac{1}{2}(2v_3 - p_2 v_1 - p_1 v_2 + p_1 v_2 - p_2 v_1) \\
&= v_3 - p_2 v_1 \\
&= (\alpha_0)_p \begin{bmatrix} v_1 \\ v_2 \\ v_3 \end{bmatrix},
\end{aligned}$$

and so $f^*\alpha_1 = \alpha_0$. Thus f is a contactomorphism of $(\mathbb{R}^3, \ker \alpha_0)$ and $(\mathbb{R}^3, \ker \alpha_1)$.

Now let us look closer at the local theory of contact 3-manifolds. The goal is to measure the extent of which a contact 1-form on a 3-manifold differs from the standard 1-form on $\mathbb{R}^3$ (in an effort to examine "invariance of domain" for contact 3-manifolds). That is, we will tweak a coordinate chart so that our contact 1-form in the new chart is just $dz - y\,dx$.

Consider a contact 3-manifold (M, ξ). Fix a $p \in M$ and consider a "germ" of a surface S around p which is transverse to ξ_p. That is, sufficiently close to p we can guarantee that $TS \cap \xi$ forms a line bundle on S. We can take a section v of this line bundle and straighten it by picking coordinates (x, y) on S so that $v = \partial/\partial y$.

Now pick a local section w of ξ near p that is transverse to S (note that w is defined in a neighbourhood of p *in* M, not S). The flow of w gives a new coordinate t, where $t(q)$ is the time it takes for a point in S to flow along w to the point q. The surface S may be flowed along w to extend the coordinates (x, y) to points off of S, but note that because w is tangent to ξ, the extension of $v = \partial/\partial y$ is not everywhere tangent to ξ (otherwise we would cook up an integral surface for ξ). Thus we have coordinates (x, y, t) about p.

Our contact structure ξ is locally given as the kernel of a 1-form:

$$\alpha = a\,dx + b\,dy + c\,dt,$$

where we see that $w = \partial/\partial t$ is a section of ξ, so :

$$\alpha(w) = 0 \implies c = 0.$$

This is analogous to how $\partial/\partial y$ is always in the standard contact structure on $\mathbb{R}^3$, the planes spinning like propellers around the y-axis (so the t-axis for M is like the y-axis in $\mathbb{R}^3$). Moreover, not both of a or b are zero at p, otherwise $\alpha \wedge d\alpha = 0$ at p. Hence without loss of generality we may assume $a \neq 0$ in a sufficiently small neighbourhood of p and that α is of the form:

$$\alpha = dx - h\,dy,$$

for some function h. Note that:

$$\begin{aligned} \alpha \wedge d\alpha &= (dx - h\,dy) \wedge (dx - h\,dy) \\ &= (dx - h\,dy) \wedge dy \wedge dh \\ &= dx \wedge dy \wedge dh \neq 0, \end{aligned}$$

and so then

$$\begin{aligned} dh = \frac{\partial h}{\partial x}\,dx + \frac{\partial h}{\partial y}\,dy + \frac{\partial h}{\partial t}\,dt &\implies \frac{\partial h}{\partial t}\,dx \wedge dy \wedge dt \neq 0 \\ &\implies \frac{\partial h}{\partial t} \neq 0, \end{aligned}$$

and thus h is a local diffeomorphism and in a sufficiently small neighbourhood is invertible. Hence applying the implicit function theorem, (y, h, x) is a set of coordinates[9] on M, and we have an associated coordinate chart $\phi\colon U \to \mathbb{R}^3$ around p such that $\alpha = \phi^*(dz - y\,dx)$.

[9] We picked the order so that it is analogous to $dz - y\,dx$.

Observation 10. Locally, any contact 3-manifold is contactomorphic to an open subset of $\mathbb{R}^3$ endowed with the standard contact structure.

The result might be surprising (or to the experienced, might not!), but this is analogous to how every symplectic structure locally looks like the standard symplectic structure of $\mathbb{R}^{2n}$. For symplectic manifolds, this is a consequence of Darboux's theorem, and contact manifolds have an analogous result that generalises the 3-dimensional case above. The more general proof can either be proven similarly (and less elegantly) as above, or with Moser's trick[10].

Theorem 2.2.7 (Darboux-Pfaff theorem). *Let (M_1, η_1) and (M_2, η_2) be two contact manifolds of the same dimension. For any $p \in M_1$ and $q \in M_2$ there exists a neighbourhood $N_1 \subseteq M_1$ of p and $N_2 \subseteq M_2$ of q such that $(N_1, \eta_1|_{N_1})$ is contactomorphic to $(N_2, \eta_2|_{N_2})$.*

We conclude our introduction to contact manifolds with Gray's theorem which effectively states that deforming a contact structure within a compact set does not yield any novel contact structures, up to contactomorphism. This notion of stability for contact structures led to this theorem also being known as the Gray stability theorem.

Theorem 2.2.8 (Gray's theorem). *If $\{\eta_t \mid t \in [0,1]\}$ is a smooth family of contact structures on M that agree off of a non-empty compact subset of M, then there exists an isotopy consisting of diffeomorphisms $\{\psi_t \mid t \in [0,1]\}$ such that $\psi_0 = id_M$ and $(\psi_t)_*(\eta_0) = \eta_t$.*

A proof of both Gray's and the general Darboux-Pfaff theorem can be found in [1][8.1, 8.2] or in [9][2.2.2, 2.5.1].

2.2.2 Siblings

Perhaps one of the more useful aspects of studying contact topology is the very close connection it has with symplectic topology. Acting as the projective sibling to symplectic topology, contact topology inherits many of the reasons to study symplectic manifolds (e.g. Hamiltonian dynamics). Because of this, we will state a few results and techniques that connect the two.

We have already seen the usefulness of viewing the definition of completely non-integrable with a symplectic lens. As a result, we saw results like the

[10] I am unsure if it is called a trick more because it is quick, or because after seeing it, you feel you are tricked out of a proof. Given a family of k-forms $\{\alpha_t \mid t \in [0,1]\}$, Moser's trick constructs a diffeomorphism ϕ satisfying $\phi^*\alpha_1 = \alpha_0$ by finding a time-dependent vector field v_t so that the flow ϕ_t satisfies $\phi_t^*\alpha_t = \alpha_0$. It can be shown that one needs to solve

$$\frac{d}{dt}\phi_t^*\alpha_t = \phi_t^*(\alpha_t' + \mathcal{L}_{v_t}\alpha_t) = 0,$$

and then we get $\phi := \phi_1$ as our diffeomorphism after this condition is satisfied.

odd-dimensionality of contact manifolds and the precise dimension of Legendrian submanifolds, as well as the analogy between Lagrangian and Legendrian submanifolds.

Example 2.2.9. Suppose that $(M, \omega := d\lambda)$ is an exact symplectic manifold of dimension $2n$, where we call λ a **Liouville 1-form**. A **Liouville vector field** corresponding to λ is a vector field X such that

$$\iota_X \omega = \lambda.$$

Thus in the presence of a Liouville vector field, we get that $\omega = d\lambda$ is equivalent to $\mathcal{L}_X \omega = \omega$, using Cartan's magic formula.

Suppose that $\Sigma \subset M$ is a hypersurface transverse to X (meaning, at any $p \in \Sigma$, $T_pM = T_p\Sigma + \mathbb{R}X_p$). Then we see that:

$$\lambda \wedge (d\lambda)^{n-1} = (\iota_X \omega) \wedge \omega^{n-1} = \frac{1}{n}(\iota_X \omega^n),$$

where we calculate the last equality by noting:

$$\iota_X \omega^n = \iota_X(\omega \wedge \omega^{n-1}) = (\iota_X \omega) \wedge \omega^{n-1} + \omega \wedge (\iota_X \omega^{n-1}) = \cdots = n(\iota_X \omega) \wedge \omega^{n-1}.$$

Hence $\iota_X \omega^n$ is non-zero on Σ. Thus Σ is a contact manifold of dimension $2n-1$ with contact form λ!

These hypersurfaces are termed **hypersurfaces of contact type**. On the other hand, given a co-orientable contact manifold (M, η) with a globally defined contact form α, consider[11] $V = \mathbb{R} \times M$ and define a symplectic form on V by $\omega := d(e^t \alpha)$ for the coordinate on $\mathbb{R}$. The manifold V is called the **symplectization** of M, and M embeds in V as a hypersurface of contact type.

This observation is not just a one-way street: every compact hypersurface Σ of contact type in a symplectic manifold (M, ω) is locally symplectomorphic to $\{0\} \times \Sigma$ in the symplectization of Σ with regards to the Liouville 1-form as the contact form.

Another connection is that the Darboux-Pfaff theorem follows from Darboux's theorem in symplectic geometry. We consider a contact $(2n+1)$-manifold, (M, ξ).

Lemma 2.2.10. *For each $p \in M$, we can find a piece of hypersurface S containing p which is transverse to ξ, such that $d\alpha$ is symplectic on S.*

Proof. Note that $d\alpha|_\xi$ is non-degenerate (maximally so). For each $q \in M$ we have that the $d\alpha$-complement of T_qM is 1-dimensional, giving a line field on M. Let w be a non-zero section of this line field near p and pick a hypersurface S

[11] On his blog [22], Chris Wendl poses an excellent point about the order of the product of the symplectization. Because (M, η) is co-oriented, M has a natural orientation, and thus the product $\mathbb{R} \times M$ is naturally oriented. However, if we re-ordered the product as $M \times \mathbb{R}$, then this natural orientation would not match that given by the symplectic form $d(e^t \alpha)$. A minor, but pedagogical caution about orientations.

which is transverse to w at p. We can shrink S to a neighbourhood of p where it is everywhere-transverse to w.

Now suppose that $d\alpha$ is degenerate on S, that there exists a $q \in S$ and $u_q \in T_qS$ such that for all $v_q \in T_qS$ we have:

$$d\alpha(u_q, v_q) = 0.$$

By definition of w, we have that $d\alpha(u_q, w_q) = 0$, and since w is transverse to S, then $T_qM = T_qS + \operatorname{span}\{w_q\}$ and thus u_q is $d\alpha$-orthogonal to T_qM, being a multiple of w_q. This contradicts our construction of S being transverse to w at q. ■

Thus $d\alpha$ is a symplectic form on the S guaranteed by the lemma, and by the Darboux's theorem (see Theorem A.2.4, the symplectic analogue to Theorem 2.2.7), we have that there are coordinates $(x_1, \dots, x_n, y_1, \dots, y_n, z)$ on M satisfying:

$$d\alpha = \sum_{i=1}^{n} dx_i \wedge dy_i.$$

It is easy to calculate that $\alpha + \sum_{i=1}^{n} y_i\, dx_i$ is closed, so it must be exact, appealing to Poincaré's Lemma[12]. That is, there is some function f such that:

$$df = \alpha + \sum_{i=1}^{n} y_i\, dx_i,$$

and hence we conclude by the condition that $\alpha \wedge (d\alpha)^n \neq 0$ that:

$$\begin{aligned}
\alpha \wedge (d\alpha)^n &= \left(df - \sum_{i=1}^{n} y_i\, dx_i\right) \wedge \left(\sum_{i=1}^{n} dx_i \wedge dy_i\right)^n \\
&= n! \left(df - \sum_{i=1}^{n} y_i\, dx_i\right) \wedge \left(\bigwedge_{i=1}^{n} dx_i \wedge dy_i\right) \\
&= n!\, df \wedge \left(\bigwedge_{i=1}^{n} dx_i \wedge dy_i\right) \neq 0,
\end{aligned}$$

and so $\{df, dx_i, dy_i \mid 1 \leqq i \leqq n\}$ constitutes a linearly independent set, and hence gives us our desired "Darboux chart" for (M, ξ) in the original coordinate neighbourhood.

The intrigued reader can read more about symplectic geometry in [1], [5], or [17].

2.2.3 Characteristic foliations

Not entirely unrelated to our build-up of integrability is the concept of foliations. A foliation of a manifold M is a decomposition (partition) of M into a smooth

[12] If M is contractible, then all closed forms in $\Omega(M)$ are exact.

family of submanifolds, called *leaves* (much akin to the geological notion of foliations). Any integrable d-field η on M gives a foliation whose leaves are the integral submanifolds of maximal dimension. In particular, a 1-dimensional foliation (a foliation whose leaves are all 1-dimensional submanifolds) is defined locally as the span of a non-vanishing vector field.

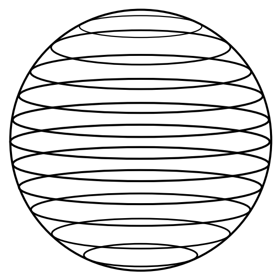

Because the special structure of a contact manifold is not particularly fond of integrability, the foliations we will speak of will manifest themselves in a different way, as there may be certain singularities present.

By **1-dimensional singular foliation** on M, we will mean a closed subset $X \subset M$ (which constitute the set of singularities) and a 1-dimensional foliation $\mathscr{F}$ of $M \setminus X$ (in the sense that it is locally given as a non-vanishing vector field).

Now consider a contact manifold (M, ξ) and an embedded hypersurface $S \subset M$. For a general $p \in S$, T_pS intersects ξ_p transversely as a line. On occasion, T_pS and ξ_p might coincide, and in this case we consider p to be singular. Setting

$$X = \{q \in S \mid T_qS = \xi_q\},$$

we see that locally in a chart U around a $p \in S$, then $U \setminus X$ is associated a line bundle. In particular, if we pick a local orientation for each of TS, $\xi|_S$, and $TM|_S$. The transverse intersection of oriented subspaces of an oriented space happens to be canonically oriented, and hence the line bundle is locally oriented. Orientable line bundles are trivial, and hence there exists a non-vanishing vector field[13] v on U which locally defines our line bundle here. Hence we have an associated singular 1-dimensional foliation of the hypersurface S, called the **characteristic foliation** of S, denoted by S_ξ.

Example 2.2.11. Consider $\mathbb{R}^3$ with the contact structure $\xi = \ker(dz + r^2\, d\theta)$ in cylindrical coordinates. It is easily calculated that:

$$\xi = \operatorname{span}\left\{\frac{\partial}{\partial r}, \frac{\partial}{\partial \theta} - r^2 \frac{\partial}{\partial z}\right\}.$$

Now consider the infinitely long cylinder $C \subset \mathbb{R}^3$ of unit radius (that is, defined by $r = 1$). It is also easily seen that:

$$TC = \operatorname{span}\left\{\frac{\partial}{\partial \theta}, \frac{\partial}{\partial z}\right\}.$$

[13]If $E \to M$ is a trivial line bundle, then $E \cong M \times \mathbb{R}$ and taking a section $p \mapsto (p, 1)$ suffices.

Hence at a point $p = (r_0, \theta_0, z_0) \in C$, we can calculate $T_pC \cap \xi_p$ by calculating the solutions to the system (meanwhile identifying $T_p\mathbb{R}^3 \cong \mathbb{R}^3$):

$$a_1 \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix} + a_2 \begin{bmatrix} 0 \\ 1 \\ -r_0^2 \end{bmatrix} = a_3 \begin{bmatrix} 0 \\ 1 \\ 0 \end{bmatrix} + a_4 \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix},$$

which is equivalent to solving the homogeneous system:

$$\begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & -1 & 0 \\ 0 & -{r_0}^2 & 0 & -1 \end{bmatrix} \longrightarrow \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 1/{r_0}^2 \\ 0 & 0 & 1 & 1/{r_0}^2 \end{bmatrix} \implies \begin{bmatrix} a_1 \\ a_2 \\ a_3 \\ a_4 \end{bmatrix} = \begin{bmatrix} 0 \\ -s/{r_0}^2 \\ -s/{r_0}^2 \\ s \end{bmatrix}, \quad s \in \mathbb{R}.$$

Thus the intersection is spanned by:

$$v_p := 0 \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix} - \frac{1}{{r_0}^2} \begin{bmatrix} 0 \\ 1 \\ -{r_0}^2 \end{bmatrix} = \begin{bmatrix} 0 \\ -1/{r_0}^2 \\ 1 \end{bmatrix}.$$

Of course, C is defined by $r = 1$ so $v_p = \frac{\partial}{\partial z}(p) - \frac{1}{{r_0}^2}\frac{\partial}{\partial \theta}(p) = \frac{\partial}{\partial z}(p) - \frac{\partial}{\partial \theta}(p)$. Thus the characteristic foliation C_ξ foliates the infinite cylinder as a family of helices. If we decreased the radius of C, the inclination of the helices would increase, and likewise an increase in radius of C would decrease the inclination.

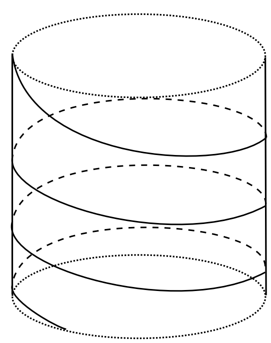

Let us take a moment to recall some basic theory about vector fields. First, consider M with the coordinate chart $(U, x_1, \ldots, x_n)$ about $p \in M$. This induces the basis $\{(\partial/\partial x_1)(p), \ldots, (\partial/\partial x_n)(p)\}$ for T_pM. In particular, any $v_p \in T_pM$ looks like:

$$v_p = \sum_{i=1}^{n} v_i \frac{\partial}{\partial x_i}(p),$$

for $v_i \in \mathbb{R}$, and so we get induced coordinates $y_i(v_p) := v_i$ on the tangent space. Thus $(TU, x_1, \ldots, x_n, y_1, \ldots, y_n)$ is a coordinate chart on TM. The zero section,

$$Z := \{0_q \in TM \mid q \in M\},$$

is a submanifold of TM, and in particular, we see that:

$$T_{(0,p)}Z = \operatorname{span}\left\{\frac{\partial}{\partial x_1}(p,0), \dots, \frac{\partial}{\partial x_n}(p,0)\right\},$$

and so we can see:

$$T_{(p,0)}TM = T_{(p,0)}Z \oplus T_{(p,0)}\pi^{-1}(\{(p,0)\}).$$

Finally, recall that the zero section Z is naturally diffeomorphic to M and so we have a natural linear isomorphism, and $\pi^{-1}(\{(p,0)\}) = T_pM$ is a vector space whose tangent space is naturally[14] isomorphic to itself, hence:

$$T_{(p,0)}TM = T_pM \oplus T_pM.$$

Now if $v\colon M \to TM$ is a vector field on M such that $v(p) = 0$, then we can linearise v at p by merely projecting the push-forward $v_{*,p}\colon T_pM \to T_{(p,0)} \cong T_pM \oplus T_pM$ to the second T_pM, thus having a linear map $T_pM \to T_pM$, called the **linearisation** of v at the zero p. Note that the linearisation manifests itself in a chart similarly to the vector field f defined at the beginning of Section 2.1.1 (taking the projection to the second factor of $T\mathbb{R}^n \cong \mathbb{R}^n \times \mathbb{R}^n$).

We say that a zero p of v is **non-degenerate** if the linearisation of v at p is invertible (that is, an isomorphism). Geometrically, this corresponds precisely to when (viewed as a submanifold of TM) the section v is transverse to the zero section. It follows that non-degenerate zeros are both isolated and generic (in the sense that they are stable under perturbations).

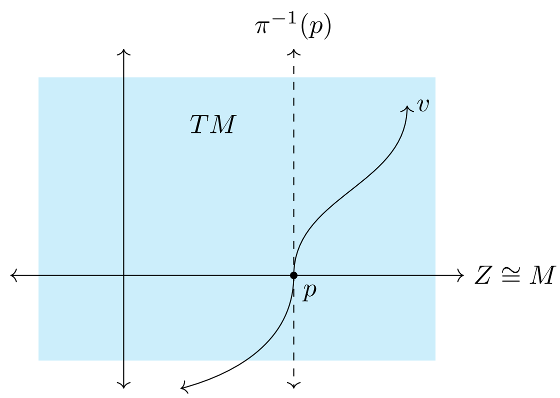

We now restrict our attention to the special case where v is a vector field on a (2-dimensional) surface S. Because the linearisation of v at a non-degenerate zero is invertible, its has no zero eigenvalues. A strengthening of this is to have the linearisation at p have non-zero real parts (a property often called "hyperbolic"

[14]If V is a vector space, consider the isomorphism $\gamma \mapsto \gamma - p$ which maps curves γ that go through $p \in V$ to curves $\gamma - p$ which go through 0. Hence we have $T_pV \cong T_0V$ naturally. Finally note that $T_0V = V$, as any $v \in V$ can be shown to be in T_0V via the curve $\gamma(t) := t \cdot v$.

(or linear) in the study of dynamical systems). We will call the zero p **generic** in this case to avoid other unfortunate naming conventions.

Having constructed the linearisation $L_p\colon \mathbb{R}^2 \to \mathbb{R}^2$ at the zero p of the vector field v on S, we get a system consisting of two first-order ODEs

$$y' = L_p y, \quad y = y(t).$$

Given an initial condition $y(0) = K$, the initial value problem is well known (see [13]) to have the unique solution:

$$y(t) = \exp(tL_p)K.$$

The local geometry of the associated flow,

$$\{\exp(tL_p)\colon \mathbb{R}^2 \to \mathbb{R}^2 \mid t \in [0,1]\},$$

can be distinguished based on the eigenvalues of the map L_p. In particular, if the eigenvalues of L_p:

(i) are real and of opposite sign, then the flow manifests itself as a **saddle**.

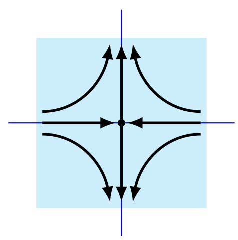

(ii) have real parts of the same sign, then the flow manifests itself as a **node**:

- If they are both positive, the flow manifests itself as a **source**.
- If they are both negative, the flow manifests itself as a **sink**.

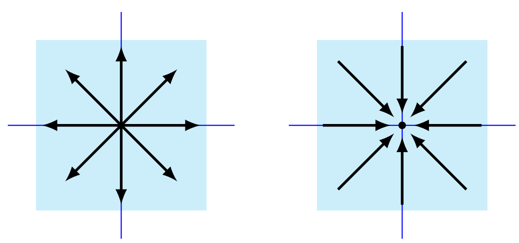

(iii) are pure imaginary, then the flow manifests itself as periodic orbits.

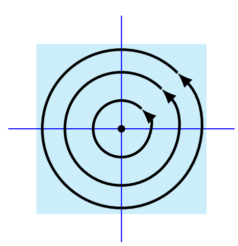

And finally, if the matrix is not diagonalizable, then it has one eigenvalue and:

(iv) there exists[15] a change of coordinates so that L_p is represented as a matrix

$$\begin{bmatrix} \lambda & 0 \\ 1 & \lambda \end{bmatrix},$$

in which case the flow exhibits the geometry of an "improper node".

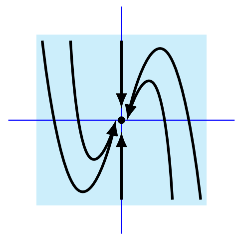

Zeros of vector fields whose local geometry is described by (i) will be termed **hyperbolic**[16], and those described by (ii) or (iv) will be termed **elliptic**. Thus the generic zeros of a vector field consist of the elliptic and hyperbolic types.

We interchangeably refer to the flow lines of a vector field as leaves or orbits depending on whether the context is foliations or dynamical systems. Some orbits of a vector field are periodic, called closed orbits. In the context of singular foliations where we prefer the terminology of leaves, the analogous term is a **closed leaf**.

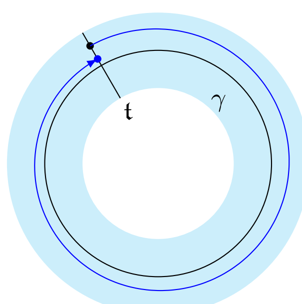

If γ is a closed leaf on an orientable surface S, we may identify a neighbourhood of $\gamma = S^1 \times \{0\}$ with $S^1 \times (-\varepsilon, \varepsilon)$. Fixing a $(p, 0) \in \gamma$, define the transversal

$$\mathfrak{t} := \{p\} \times (-\varepsilon, \varepsilon),$$

and define the **Poincaré return map** to be the map[17] $x \mapsto \pi(x)$ where $(p, \pi(x))$ is the first intersection of the leaf containing (p, x) forward in time with $\mathfrak{t}$. With this notation, we define:

[15]Thank you, Jordan canonical form!

[16]Reader be warned: the word "hyperbolic" is used in dynamical systems, not in this sense, but in the sense we use "generic".

[17]This map is well-defined for small enough $x \in (-\varepsilon, \varepsilon)$.

(i) γ is **degenerate** if $\pi'(0) = 1$.

(ii) γ is **attracting** if $\pi'(0) < 1$.

(iii) γ is **repelling** if $\pi'(0) > 1$.

A **generic** (or non-degenerate) closed leaf is either attracting or repelling.

Now we turn our attention to a particular class of vector fields. A vector field v is said to be **Morse-Smale** if it satisfies the following conditions:

(i) There are finitely many singularities and cycles, each being generic.

(ii) There are no trajectories connecting saddle points.

(iii) The limit set of each half-orbit (forward or backward in time) is either a singularity, a closed leaf, or a union of singularities connected by leaves.

The property (iii) is also known as the **Poincaré-Bendixson** property. A powerful tool when dealing with vector fields and flows in general was proven by Peixoto:

Theorem 2.2.12 (Peixoto [18])**.** *Suppose S is an orientable, compact surface. Then the set of Morse-Smale vector fields on S is an open, dense subset of the set of C^1-vector fields on S.*

Because of Peixoto's result, we will from this point on refer to Morse-Smale vector fields as **generic**.

Observation 11. Generically, the characteristic foliation of a compact surface S will have finitely many singularities and cycles (all generic, and thus isolated), and will not contain any saddle connections.

Suppose that the characteristic foliation is oriented. In this case, a choice of orientation for S_ξ gives the leaves direction (like a vector field, as opposed to a line bundle). If we additionally let S be oriented, then a singularity p of S_ξ may be classified as follows:

- If the orientation of T_pS agrees with that of S_ξ at p, then we say p is a positive singularity.
- If the orientation of T_pS disagrees with that of S_ξ at p, then we say p is a negative singularity.

A positive elliptic point is identical to a source, while a negative elliptic point is identical to a sink. Geometrically, hyperbolic points of different signs are indistinguishable.

The power of the characteristic foliation is that it determines a "germ" of a contact structure, as captured in the following theorem:

Theorem 2.2.13 (**Moser-Weinstein theorem** [1][8.23])**.** *Let (M,ξ) and (N,η) be contact manifolds and $S \subset M$, $S' \subset N$ be closed*[18] *embedded hypersurfaces. If there exists a diffeomorphism $f\colon S_0 \to S_1$ such that f preserves the characteristic foliation (sends leaves of S_ξ to leaves of S'_η), then f may be extended to a contactomorphism in some neighbourhood of S_0.*

We finish our introduction to characteristic foliations by considering a specific example.

Example 2.2.14 (Overtwisted disk)**.** Consider $\mathbb{R}^3$ with the contact structure $\xi = \ker(\cos(r)\, dz + r\sin(r)\, d\theta)$ in cylindrical coordinates. It can be shown for $p = (r_0, \theta_0, z_0) \in \mathbb{R}^3$ that:

$$\xi_p = \operatorname{span}\left\{ \frac{\partial}{\partial r}(p), \cos(r_0)\frac{\partial}{\partial \theta}(p) - r_0 \sin(r_0)\frac{\partial}{\partial z}(p) \right\}.$$

If we take the disk $D_{\text{ot}} = D(0;\pi)$ of radius π centred at the origin in the r,θ-plane as our hypersurface (with boundary), then we note that:

- We have singularities at for all $p \in \partial D_{\text{ot}}$ and at the origin.
- For all other $p \in D_{\text{ot}}$, we have $T_p D_{\text{ot}} \cap \xi_p$ is the span of $(\partial/\partial r)(p)$.

The characteristic foliation of D_{ot} thus looks akin to a pizza, with a closed leaf of singularities (another example of a Legendrian knot!) and an elliptic singularity dead-centre. Note that this characteristic foliation is also not generic.

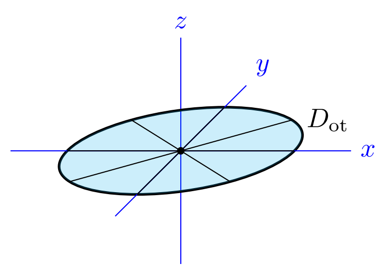

The contact structure in Example 2.2.14 is called the **standard overtwisted contact structure** of $\mathbb{R}^3$. A contact 3-manifold (M,ξ) is called **overtwisted** if there exists an embedding of a disk D in M such that D is diffeomorphic to to the standard overtwisted disk D_{ot}, as presented in Example 2.2.14, carrying D_ξ to the characteristic foliation of D_{ot} induced by the standard overtwisted structure. If a contact 3-manifold is not overtwisted[19], it is called **tight**. Due to the Darboux-Pfaff theorem, the twisted aspect of overtwisted contact structures happens to be a global concern.

[18] There is a version with the hypothesis of compact with Legendrian boundary instead of closed.

[19] For example, the standard contact structure on $\mathbb{R}^3$ as given in Example 2.1.4 is tight, and is *not* contactomorphic to the standard overtwisted structure on $\mathbb{R}^3$.

2.2.4 Giroux's elimination lemma

The characteristic foliation of a surface can sometimes be simplified. The main technique for simplifying the characteristic foliation is by using Giroux's elimination lemma, which allows you to eliminate two singularities of the same sign under certain circumstances.

Consider $\mathbb{R}^3$ endowed with the standard contact structure $\xi := \ker(dz - y\,dx)$, and the surface S_ε, defined to be the graph of

$$z = \varepsilon y,$$

for some $\varepsilon > 0$ (a plane which is tilted in the y, z-plane. At any $p = (p_1, p_2, p_3) \in S_\varepsilon$ we have that:

$$T_p S_\varepsilon = \operatorname{span}\left\{ \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix}, \begin{bmatrix} 0 \\ 1 \\ \varepsilon \end{bmatrix} \right\},$$

using the isomorphism $T_p\mathbb{R}^3 \cong \mathbb{R}^3$. Solving for $T_p S_\varepsilon \cap \xi_p$, we see:

$$\begin{bmatrix} 1 & 0 & 0 & -1 \\ 0 & 1 & -1 & 0 \\ 0 & \varepsilon & 0 & -p_2 \end{bmatrix} \longrightarrow \begin{bmatrix} 1 & 0 & 0 & -1 \\ 0 & 1 & 0 & -p_2/\varepsilon \\ 0 & 0 & 1 & -p_2/\varepsilon \end{bmatrix},$$

so the intersection is spanned by, for example,

$$v_p = \begin{bmatrix} \varepsilon \\ p_2 \\ \varepsilon p_2 \end{bmatrix} = \begin{bmatrix} \varepsilon \\ 0 \\ 0 \end{bmatrix} + p_2 \begin{bmatrix} 0 \\ 1 \\ \varepsilon \end{bmatrix},$$

which foliates S_ε as shown below (projected to the x, y-plane). There are no singularities in $(S_\varepsilon)_\xi$, as (the above calculation shows that) $T_p S_\varepsilon$ never coincides with ξ_p, their intersection always being one-dimensional.

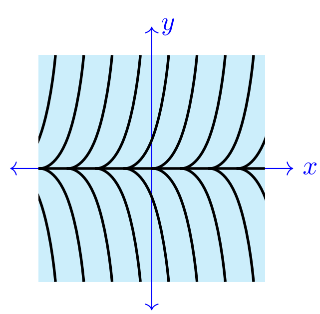

Pick a smooth bump $h\colon \mathbb{R} \to \mathbb{R}$ where $h|_{[-c,c]} \equiv -1$ and $h|_{\mathbb{R}\setminus[-d,d]} \equiv 1$ for $d > c$. The surface defined by the parametrisation

$$\psi(x, y) = (x, y, \varepsilon h(x) y),$$

which we denote by ${}^{\reflectbox{$\varepsilon$}}S$, is S_ε with a twist. The parametrisation ψ gives $T_p{}^{\reflectbox{$\varepsilon$}}S$ for $p = \psi(x, y)$ as $\operatorname{im}\psi_{*,(x,y)}$, where (in coordinates):

$$\psi_{*,(x,y)} = \begin{bmatrix} 1 & 0 \\ 0 & 1 \\ \varepsilon h'(x) & \varepsilon h(x) \end{bmatrix} \implies T_p{}^{\reflectbox{ε}}S = \operatorname{span}\left\{ \begin{bmatrix} 1 \\ 0 \\ \varepsilon h'(x) \end{bmatrix}, \begin{bmatrix} 0 \\ 1 \\ \varepsilon h(x) \end{bmatrix} \right\}$$

If $x \notin [-d, d]$, then $h(x) = 1$ and ${}^{\reflectbox{$\varepsilon$}}S$ and S_ε coincide (as do their tangent spaces). If $x \in [-c, c]$, then the surface ${}^{\reflectbox{$\varepsilon$}}S$ is sloping the opposite way that S_ε is, coinciding with the graph of the function $z = -\varepsilon y$. Solving for the intersection of $T_p{}^{\reflectbox{$\varepsilon$}}S \cap \xi_p$, we get:

$$\begin{bmatrix} 1 & 0 & 0 & -1 \\ 0 & 1 & -1 & 0 \\ \varepsilon h'(x) & \varepsilon h(x) & 0 & -y \end{bmatrix} \longrightarrow \begin{bmatrix} 1 & 0 & 0 & 1 \\ 0 & 1 & -1 & 0 \\ 0 & 0 & \varepsilon h(x) & \varepsilon h'(x) \end{bmatrix},$$

which indicates that $T_p{}^{\reflectbox{$\varepsilon$}}S$ and ξ_p coincide if and only if

$$\begin{cases} h(x) = 0 \\ \varepsilon h'(x) - y = 0 \end{cases}.$$

By the intermediate value theorem, we can find $x_0 \in (-d, -c)$ such that $h(x_0) = 0$, so that the point $\psi(x_0, \varepsilon h'(x_0))$ is a singularity in $({}^{\reflectbox{$\varepsilon$}}S)_\xi$. Similarly, we get a $x_1 \in (c, d)$ such that $\psi(x_1, \varepsilon h'(x_1))$ is a singularity. Assuming that h is monotonically decreasing and increasing on $[-d, -c]$ and $[c, d]$ respectively, these are the only singularities introduced, and moreover interpolating linearly between S_ε and ${}^{\reflectbox{$\varepsilon$}}S$ gives us a smooth isotopy.

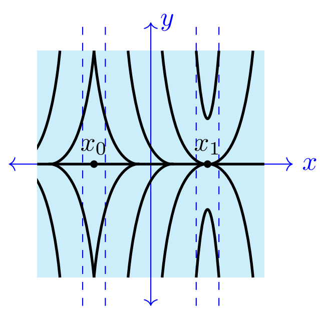

The singularity x_0 is hyperbolic, and x_0 is elliptic, as can be guessed from the local flows seen in the characteristic foliation $({}^{\reflectbox{$\varepsilon$}}S)_\xi$ above, and the singularities are of the same sign (since along the x-axis the contact planes rotate one way, until $x = 0$ when they rotate back the other direction failing to flip between x_0 and x_1).

As it turns out, given a hyperbolic and elliptic points of the same sign like x_0 and x_1 connected by a leaf γ, one can find a deformation (much like the one above) supported in an arbitrarily small neighbourhood of γ which "eliminates" the pair. This result is known as Giroux's elimination lemma.

Theorem 2.2.15 (**Giroux's elimination lemma**). *Let S be a closed, oriented surface embedded in a contact 3-manifold (M, ξ) such that S_ξ is generic. If there exists an elliptic point e and a hyperbolic point h of the same sign, connected by a separatrix γ of h, then given any neighbourhood N of γ, we may find an isotopy, supported in N, of S to a surface S' such that $(S' \cap N)_\xi$ contains no singularities and S_ξ and S'_ξ agree outside of N.*

A full proof of this result may be found in Giroux's own writing in [11], or in Geiges's tome [9][4.6.26]. A proof similar to our introduction is given in [1][8.26].

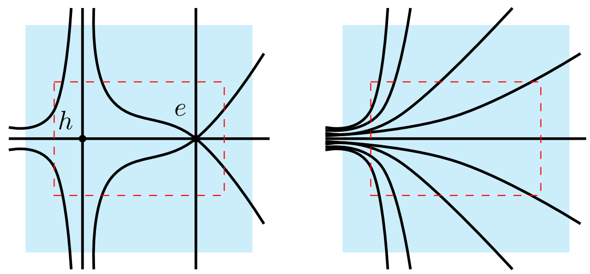

It is important to note that just as you can eliminate hyperbolic-elliptic pairs of the same sign, one can also use Giroux's elimination lemma to introduce the pair of singularities within a neighbourhood, much like we did during the discussion preceding the theorem statement.

The intrigued reader can read more about the general theory of contact geometry and topology in [9]. A delightful gallery of characteristic foliations on various spheres and surfaces can be found in [3].

2.3 Convexity

This section follows closely the presentation of both [7] and [16]. An alternative discussion may be found in [9] or Giroux's original paper [11].

2.3.1 Contact vector fields and convex surfaces

We now introduce the study of convexity and dividing curves, initiated by Emmanuel Giroux [11]. Under a suitable condition known as "ξ-convexity" (not in the naive sense), the essential info of a characteristic foliation can be summarised in an often simpler object called a "dividing curve", encapsulating all the information about the germ of a contact structure. It turns out that a convex surface is generic, and thus we may substitute the study of characteristic foliations with that of dividing curves—at least *generically* so.

Suppose that (M, ξ) is a contact manifold. A vector field v on M is called a **contact vector field** if its flow preserves the contact structure, ξ. That is, if $\phi_t \colon M \to M$ is the local flow, then it is a contactomorphism for all t:

$$(\phi_t)_* \xi = \xi.$$

If ξ is given locally as the kernel of a 1-form α, then for any non-vanishing $g \in C^\infty(M)$, $g\alpha$ also gives ξ locally. Hence for each t, there is a $g_t \in C^\infty(M)$ such that $\phi_t^*\alpha = g_t\alpha$, varying smoothly in t. Taking the Lie derivative:

$$\mathcal{L}_v\,\alpha = \left.\frac{d}{dt}\right|_{t=0} \phi_t^*(\alpha) = \left.\frac{d}{dt}\right|_{t=0} g_t\alpha = g_0'\alpha.$$

In other words, if v is a contact vector field, then there exists a $g \in C^\infty(M)$ such that:

$$\mathcal{L}_v\,\alpha = g\alpha.$$

Note that the linear properties of the Lie derivative imply that the subset of contact vector fields is a subspace of the vector space of vector fields.

Suppose that (M, ξ) is a contact manifold. A hypersurface S is **ξ-convex** if there is a contact vector field v transverse to S, that is:

$$v|_S \oplus TS = TM|_S.$$

Example 2.3.1. Consider $\mathbb{R}^3$ with the standard contact structure $\xi_0 = \ker\alpha$ for $\alpha = dz - y\,dx$. Let $v = x\frac{\partial}{\partial x} + y\frac{\partial}{\partial y} + 2z\frac{\partial}{\partial z}$ and note by direct calculation:

$$\begin{aligned}\mathcal{L}_v(\alpha) &= \iota_v d\alpha + d\iota_v\alpha \\ &= \iota_v(dx \wedge dy) + d\iota_v(dz - y\,dx) \\ &= dx(v) \otimes dy - dy(v) \otimes dx + d(z - yx) \\ &= x\,dy - y\,dx + 2dz - x\,dy - y\,dx \\ &= 2\alpha,\end{aligned}$$

and hence v is a contact vector field. Moreover, v is transverse to the hypersurface $S^2 \subset \mathbb{R}^3$. Hence S^2 is a ξ_0-convex hypersurface in $\mathbb{R}^3$.

2.3.2 Dividing curves

Let M be an orientable 3-manifold, and $S \subset M$ a orientable surface in M. Around any point $p \in S$, we can find a coordinate chart $U \subset M$ about p with coordinates $\{x_1, x_2, x_3\}$ such that $S \cap U$ is defined by the equation $x_3 = 0$, which moreover is orientation preserving with respect to S. Pick $\delta \in \{\pm 1\}$ so that

$$\left\{\frac{\partial}{\partial x_1}(p), \frac{\partial}{\partial x_2}(p), \delta\frac{\partial}{\partial x_3}(p)\right\},$$

is positively oriented with respect to M. It follows that we may extend $\delta\frac{\partial}{\partial x_3}$ to a vector field v on M, transverse to S.

Now given this vector field v transverse to S, we may flow S within the time interval $(-\varepsilon, \varepsilon)$ for a sufficiently small $\varepsilon > 0$, in order to find an oriented neighbourhood of S diffeomorphic to $S \times \mathbb{R}$. Denote the $\mathbb{R}$-coordinate by t.

If a plane field in M is given as the kernel of a 1-form α, then we may take α modulo dt to get:

$$\alpha = \beta_t + f_t\,dt,$$

for a family of functions $f_t \in C^\infty(S)$ and 1-forms β_t on S. Note that the exterior derivative on $S \times \mathbb{R}$ is given by:

$$d\mu = d_S\mu + \mu' \wedge dt,$$

where the exterior derivative d_S on the right that of S, and μ' is the derivative of the k-form μ with respect to t. Thus we may calculate:

$$\begin{aligned} d\alpha &= d\beta_t + d(f_t\, dt) \\ &= d_S\beta_t + \beta_t' \wedge dt + d_S f_t \wedge dt \\ &= d_S\beta_t + (\beta_t' + d_S f_t) \wedge dt, \end{aligned}$$

and thus if $\ker\alpha$ is completely non-integrable, then:

$$\begin{aligned} \alpha \wedge d\alpha \neq 0 &\iff (\beta_t + f_t\, dt) \wedge (d_S\beta_t + (\beta_t' + d_S f_t) \wedge dt) \neq 0 \\ &\iff \beta_t \wedge d_S\beta_t + \beta_t \wedge (\beta_t' + d_S f_t) \wedge dt + f_t\, dt \wedge d_S\beta_t \neq 0 \\ &\iff \beta_t \wedge d_S\beta_t + (\beta_t \wedge (\beta_t' + d_S f_t) + f_t\, d_S\beta_t) \wedge dt \neq 0 \\ &\iff (\beta_t \wedge (\beta_t' + d_S f_t) + f_t\, d_S\beta_t) \wedge dt \neq 0, \end{aligned}$$

noting that $\beta_t \wedge d_S\beta_t$ is a 3-form on S, so it vanishes. Hence we are left with the condition

$$\beta_t \wedge (\beta_t' + d_S f_t) + f_t\, d_S\beta_t \neq 0, \tag{†}$$

for when $\ker\alpha$ defines a contact structure.

Now if we assume both f_t and β_t are independent of t, then the contact structure is invariant under translations in the $\mathbb{R}$ direction. That is to say, the vector field $\partial/\partial t$ is transverse to $S \times \{t\}$ for all t, being a contact vector field, and so the surfaces $S \times \{t\}$ are ξ-convex!

On the other hand, if we have an orientable ξ-convex surface $S \subseteq (M, \xi)$ with respect to the contact vector field v, then we can define[20] a new contact vector field $\tilde{v}$ which is zero outside of an arbitrarily small open neighbourhood N of S and satisfying $\tilde{v}|_S = v|_S$. The flow of $\tilde{v}$ gives an oriented tubular neighbourhood, $S \times \mathbb{R}$ for which ξ is $\mathbb{R}$-invariant.

Observation 12. An orientable surface $S \subset M$ in a contact manifold (M, ξ) is ξ-convex if and only if it admits a tubular neighbourhood $S \times \mathbb{R}$ such that $\xi|_{S\times\mathbb{R}}$ is $\mathbb{R}$-invariant.

Thus in the case where S is ξ-convex, we have that the condition (†) is equivalent to:

$$\beta \wedge d_S f + f\, d_S\beta \neq 0,$$

[20]This follows from an isomorphism Φ between the contact vector fields and the sections of the normal bundle TM/ξ to ξ, a theorem by Libermann. If we pick a function $f\colon M \to \mathbb{R}$ with support in N which is the identity on S, then taking $\tilde{v} := \Phi^{-1}(f \cdot \Phi(v))$ is our "cut off" of v away from S.

where we have explicitly dropped the dependency of β and f on t. Because S is orientable, we can assume the existence of an area form ω, and thus we have a vector field u which is ω-dual to β, defined by the equation $\beta = \iota_u\omega$. Here, u and β have the same vanishing set and u directs[21] $\ker\beta$, so β defines the characteristic foliation S_ξ (the $f\,dt$ part of α being negligible as the contact structure is $\mathbb{R}$-invariant). So the vector field u directs S_ξ.

Observation 13. The vector field u which is ω-dual to β, defined by $\beta = \iota_u\omega$, directs S_ξ.

Recall that if ω is an area form, then we define $\operatorname{div}_\omega u$ as the function satisfying the equation:

$$\mathcal{L}_u\,\omega = (\operatorname{div}_\omega u)\omega.$$

Phrasing the contact condition in terms of u and appealing to divergence, $\operatorname{div}_\omega \cdot$, being the ω-dual concept to the exterior derivative $d\cdot$, we see:

$$\begin{aligned}
\beta\wedge d_S f + f\,d_S\beta &= (\iota_u\omega)\wedge d_S f + (f\operatorname{div}_\omega u)\,\omega\\
&= (\iota_u(\omega\wedge d_S f) - \omega\wedge\iota_u(d_S f)) + (f\operatorname{div}_\omega u)\,\omega\\
&= (f\operatorname{div}_\omega u)\,\omega - \omega\wedge\iota_u(d_S f)\\
&= (f\operatorname{div}_\omega u - d_S f(u))\,\omega,
\end{aligned}$$

and hence our condition becomes:

$$f\operatorname{div}_\omega u - d_S f(u) \neq 0. \tag{††}$$

Note that if we fix our contact structure to be positive (or negative), the set of functions f satisfying the condition (††) forms a convex set, as does the analogous set for β.

Observation 14. The set of functions f satisfying the positive contact condition

$$f\operatorname{div}_\omega u - d_S f(u) > 0,$$

form a convex set. An analogous statement holds for the negative version.

We now turn our attention to the set:

$$\Gamma := f^{-1}(\{0\}) = \{x\in S \mid f(x) = 0\},$$

where we see that $d_S f(u)\neq 0$ on Γ, and hence $d_S f$ is surjective. Thus Γ is a regular 1-dimensional submanifold of S, and it follows that Γ manifests itself as a

[21] Locally in an open set U of S, the foliation S_ξ is defined by a vector field v. A vector field u **directs** S_ξ if in each U there exists a non-vanishing function f such that $v = fu$.

family of disjoint, closed curves in S. Let $p \in \Gamma$ and note that if $T_p\Gamma = \text{span}\{v_p\}$, then:

$$\begin{aligned}
\alpha_p(v_p) &= \beta_p(v_p) + f(p)\,(d_S t)_p(v_p) \\
&= \beta_p(v_p), && \text{since } f(p) = 0; \\
&= \omega_p(u_p, v_p) \neq 0,
\end{aligned}$$

and so it follows that Γ is transverse to S_ξ. In addition, because $d_S f(u) \neq 0$, the vector field u is always heading along Γ from where $f > 0$ to where $f < 0$, or conversely $f < 0$ to $f > 0$, depending on the sign of $\alpha \wedge d\alpha$ (former positive, latter negative).

Defining the sets,

$$S_+ := \{p \in S \mid f(p) > 0\}, \qquad S_- := \{p \in S \mid f(p) < 0\},$$

it is possible to modify the area form ω into Ω such that $\pm \operatorname{div}_\Omega u > 0$ on $S_\pm$, vanishing along Γ. The full details of this process can be found in [9][pp.230-231].

Given a positive contact 3-manifold (M, ξ) and an embedded hypersurface $S \subset M$, we say that a family of curves Γ **divides** the characteristic foliation S_ξ if there is an area form Ω on S and a vector field u directing S_ξ such that:

- $\operatorname{div}_\Omega u$ vanishes only on Γ.
- u goes transversely out of S_+ and into S_- along Γ, where

$$S_\pm = \{p \in S \mid \pm(\operatorname{div}_\Omega u)(p) > 0\}.$$

Such a curve Γ is called a **dividing curve**. Note that dividing curves are in no way uniquely determined. We observe from our discussion preceding the definition:

Observation 15. Every convex surface admits a dividing curve.

Example 2.3.2 ([9][4.8.4])**.** Consider $S^2 \subset \mathbb{R}^3$ with the standard contact structure. We have already seen in Example 2.3.1 that S^2 is convex. Consider:

$$u^{22} := (xz - y)\frac{\partial}{\partial x} + (yz + x)\frac{\partial}{\partial y} - (x^2 + y^2)\frac{\partial}{\partial z},$$

$$\Omega := x\,dy \wedge dz + y\,dz \wedge dx + z\,dx \wedge dy.$$

[22]Note that $u(x^2 + y^2 + z^2) = 0$, so it is indeed a vector field on the sphere.

Note that:

$$\begin{aligned}
\iota_u \Omega = \iota_u(&x(dy \otimes dz - dz \otimes dy) \\
&+ y(dz \otimes dx - dx \otimes dz) \\
&+ z(dx \otimes dy - dy \otimes dx)) \\
= x(&(yz + x)\, dz + (x^2 + y^2)\, dy) \\
&+ y(-(x^2 + y^2)\, dx - (xz - y)\, dz) \\
&+ z((xz - y)\, dy - (yz + x)\, dx) \\
= -&(x^2 y + y^3 + yz^2 + xz)\, dx \\
&+ (x^3 + xy^2 + xz^2 - yz)\, dy \\
&+ (x^2 + y^2)\, dz, \\
= (&-y - xz)\, dx + (x - yz)\, dy + (1 - z^2)\, dz,
\end{aligned}$$

and so taking the exterior derivative:

$$\begin{aligned}
d(\iota_u \Omega) &= d((-y - xz)\, dx + (x - yz)\, dy + (1 - z^2)\, dz,) \\
&= -dy \wedge dx - d(xz) \wedge dx + dx \wedge dy - d(yz) \wedge dy + d(1 - z^2) \wedge dz \\
&= 2\, dx \wedge dy + dx \wedge (z\, dx + x\, dz) + dy \wedge (z\, dy + y\, dz) \\
&= 2\, dx \wedge dy + x\, dx \wedge dz + y\, dy \wedge dz \\
&= 2\, dx \wedge dy,
\end{aligned}$$

where we apply the relation $x\,dx + y\,dy + z\,dz = 0$ to get the last line. Note that:

$$d(\iota_u \Omega) = (\operatorname{div}_\Omega u)\Omega,$$

and so it is easily checked that at the north and south poles of the sphere, the divergence, $\operatorname{div}_\Omega u$ is respectively positive and negative. The 2-form $2\, dx \wedge dy$ vanishes only when the tangent plane is perpendicular to the x, y-plane, and thus is a non-zero multiple of Ω off of the equator. Hence $\operatorname{div}_\Omega u$ is positive on the upper hemisphere, and negative on the lower hemisphere. Thus setting Γ to be the equator, $\{z = 0\}$, gives a dividing curve!

Example 2.3.3. With regards to Observation 15, if the hypersurface is ξ-convex with respect to the contact vector field v, then

$$\Gamma_v := \{p \in S \mid v_p \in \xi\}$$

divides S_ξ. This can be shown by selecting local coordinates such that v corresponds to $\partial/\partial t$ and then applying the discussion preceding the definition of a dividing curve.

Of course, us discovering a one way implication and thinking it is nice enough to share with the reader heightens the expectations of the reader that we are about to show that this indeed also goes the other direction. To let the reader cool down, we present the following lemma:

Lemma 2.3.4 (Local reconstruction)**.** *Given two positive (or negative, as long as they agree) contact structures which give identical characteristic foliations on a closed surface S, there is a neighbourhood of S on which the contact structures are isotopic via an isotopy which preserves S.*

Proof. Consider the tubular neighbourhood $S \times \mathbb{R}$ where our two contact structures are $\xi := \ker(\beta_t + f_t\, dt)$ and $\tilde{\xi} := \ker(\tilde{\beta}_t + \tilde{f}_t\, dt)$ respectively satisfying the following versions of the contact condition:

$$\beta_t \wedge (\beta_t' + d_S f_t) + f_t\, d_S \beta_t \neq 0,$$

$$\tilde{\beta}_t \wedge (\tilde{\beta}_t' + d_S \tilde{f}_t) + \tilde{f}_t\, d_S \tilde{\beta}_t \neq 0.$$

Because the characteristic foliations on S are identical, this means that $\tilde{\beta}_0 = g\beta_0$ for some function g. The contact conditions at $S = S \times \{0\}$ are convex in f_0 and $\tilde{f}_0$, as well as β_0' and $\tilde{\beta}_0'$. Thus we may consider the alternative contact form $\beta_t + (\tilde{f}_t/g)\, dt$ which gives the same characteristic foliation on S. We then interpolate linearly, setting:

$$h_0^t := (1-t)f_0 + t(\tilde{f}_t/g), \quad t \in [0,1],$$

which gives us a family of plane fields $\ker(\beta_0 + h_0^t\, dt)_{t\in[0,1]}$ preserving the foliation on S between ξ and $\tilde{\xi}$. Nearby S these will remain as contact structures because the contact condition is open. The isotopy then can be found by applying Gray's theorem, 2.2.8. ■

The idea will be to set up an $\mathbb{R}$-invariant contact structure defined by the 1-form $\alpha = \beta \pm dt$, tracing the same characteristic foliation on the surface. After an application of the local reconstruction lemma, we get the other direction.

So to this end, consider a surface $S \subset (M, \xi)$ which admits a dividing curve, Γ, along with the area form, Ω, on S and vector field, v, directing S_ξ as per the definition of dividing curve. Set

$$\beta := \iota_v \Omega,$$

and recall from Observation 13 that β is wholly responsible for defining the characteristic foliation on $S \setminus \Gamma$ (where it does not vanish), and thus, setting $\tilde{\xi} := \ker(\beta \pm dt)$ (which is a contact structure!), the foliations S_ξ and $S_{\tilde{\xi}}$ are identical off of Γ. It is easy then to extend them so they are identical on the entirety of S. Applying the reconstruction lemma, we get a contact isotopy in a neighbourhood of S between ξ and $\tilde{\xi}$. Thus S is ξ-convex.

Observation 16. A surface whose characteristic foliation is divided by a dividing curve is necessarily convex.

We state the equivalence as a theorem for good measure. Convexity of a surface reduces to the existence of a dividing curve.

Theorem 2.3.5. *An orientable, closed*[23] *surface $S \subset M$ in a contact manifold (M, ξ) is ξ-convex if and only if it admits a dividing curve.*

Recall the Moser-Weinstein theorem (Theorem 2.2.13) which indicates that the characteristic foliation determines a germ of a contact structure about the surface. Simplifying this matter even further, Giroux proved an analogous theorem for dividing curves:

Lemma 2.3.6 (**Giroux's Flexibility / Realisation lemma**). *Let S be an embedded ξ-convex surface in a contact 3-manifold (M, ξ) which is divided by Γ. For any singular 1-dimensional foliation $\mathscr{F}$ which is also divided by Γ, there exists an isotopy $\{\psi_t \mid t \in [0,1]\}$ with support in an arbitrarily small neighbourhood of S such that:*

$$S_{\psi_1^* \xi} = \mathscr{F},$$

or equivalently,

$$(\psi_1(S))_\xi = \psi_1(\mathcal{F}).$$

Proof. A proof follows from Gray's theorem, noting that using a contact form $\beta \pm dt$ much like the lead up to Observation 16 yields the contact condition:

$$\pm d\beta \neq 0.$$

This condition is convex in β and a similar argument as in the proof of Lemma 2.3.4 may be used to get a family of contact structures on which we may apply Gray's theorem. ■

This reduces the very geometric tool of characteristic foliations of convex surfaces down to the very *discrete* tool of dividing curves, in the sense that instead of dealing with the uncountably many leaves of the foliation, it suffices to consider a countable family of curves.

The forward-thinking reader might see a problem: "But this simplification to dividing curves only helps for convex surfaces—what's the catch?", she might ask, "For all we know, convex surfaces might be rare!". The answer, which might be expected, is the following:

Theorem 2.3.7. *Any closed embedded surface in a contact 3-manifold is C^∞-close to a convex surface.*

It is important to note that this is unique to the case of surfaces in contact 3-manifolds, and does not generalise to higher dimensional spaces. When discussing the 3-dimensional case, convex surfaces are generic, the property of not being convex being an unstable phenomenon with regards to perturbation. We will return to discuss this result in the next section.

[23] There is a variation of this theorem that admits a compact surface with Legendrian boundary, instead of closed.

2.3.3 Geometry of dividing curves

Let $S \subset M$ be an orientable, ξ-convex surface in a contact manifold (M, ξ). As in the previous section, we write ξ near S as the kernel of a 1-form:

$$\alpha = \beta + f\,dt.$$

We had shown prior that the condition $\alpha \wedge d\alpha \neq 0$ is equivalent to:

$$\beta \wedge d_S f + f\, d_S \beta \neq 0,$$

for a 1-form $\beta \in \Omega^1(S)$ and a function $f \in C^\infty(S)$. If we pick a volume form ω, there is a vector field $u = \ker \beta$ such that $\beta = \iota_u \omega$. The contact condition becomes:

$$f\ \mathrm{div}_\omega\, u - d_S f(u) \neq 0.$$

Let $p \in S$ be a singularity of S_ξ. Because S_ξ is defined by the vector field u, we have that $u_p = 0$, and hence we get immediately that:

$$f(p)\,(\mathrm{div}_\omega\, u)(p) \neq 0.$$

If we take ξ to be a positive[24] contact structure (that is, $\alpha \wedge d\alpha > 0$), then we conclude that the sign of the divergence of u at p agrees with the sign of the singularity, defined by its orientation.

Observation 17. If Γ divides the characteristic foliation of a convex surface S, then $S_\pm$ contains all singularities of sign $\pm$.

Consider two singularities $p, q \in S$ of opposite sign, connected by a leaf ℓ directed from $p \to q$. Because the singularities are of opposite sign, our observation indicates that if Γ divides S_ξ, Γ separates S so that p and q are in different components. In particular, Γ and ℓ intersect transversely, and the direction of ℓ is from S_+ into S_- when moving across Γ. Hence p must be positive and q must be negative. When elliptic points are involved, this is not a surprising conclusion, however considering the connection of hyperbolic points, we see that on convex surfaces they will only ever connect *from* a positive hyperbolic *to* a negative hyperbolic point. A saddle connection between differently-signed hyperbolic points which is directed from the negative to the positive point is called a **retrograde saddle connection**.

Observation 18. On a convex surface, retrograde saddle connections cannot occur.

[24] Taking a negative contact structure would give analogous results, but we will just focus on the positive case where the divergence matches the signs, rather than is the opposite.

Considering a single positive hyperbolic point, the dividing curve Γ cannot cross the stable separatrices as that would contradict the vector field u flowing across Γ from S_+ into S_-. Thus Γ may only cross the unstable separatrices of positive hyperbolic points. We can make an analogous observation about negative hyperbolic points.

Observation 19. If Γ divides the characteristic foliation of a convex surface S, then:

- Γ cannot cross the stable separatrices of a positive hyperbolic point.
- Γ cannot cross the unstable separatrices of a negative hyperbolic point.

Similarly, we now consider a closed leaf γ on S. Note that γ and Γ are transverse. If they intersect, $f|_\gamma$ has a critical point p (γ is compact) at which the contact condition becomes

$$f(p)\,(\operatorname{div}_\omega u)(p) > 0,$$

and thus repelling orbits are in S_+ and attracting orbits are in S_-. Thus Γ and γ have empty intersection.

Observation 20. If Γ divides the characteristic foliation of a convex surface S, then S_+ contains all repelling closed leaves, and S_- all attracting closed leaves.

Now note that if $d_S\beta = 0$, the contact condition becomes:

$$\beta \wedge d_S f \neq 0.$$

If S is closed, then f has a critical point and thus the surface S cannot satisfy the contact condition. Hence S would fail to be ξ-convex. Considering the reformulation of the contact condition in terms of the ω-dual of β, the condition $d_S\beta = 0$ becomes $\operatorname{div}_\omega u = 0$, and the contact condition $-d_S f(u) \neq 0$. The manner above in which S failed to be convex only requires a compact leaf with vanishing divergence, and any compact leaf of S_ξ is necessarily a closed leaf. As before, given a closed leaf γ such that the divergence vanishes, compactness implies that $f|_\gamma$ consists of critical points, and thus a similar conclusion follows: S cannot be ξ-convex. A leaf γ with this property happens to be a degenerate closed leaf.

Observation 21. On a convex surface, degenerate closed leaves cannot occur.

With the observations we have made, we will now algorithmically arrange a dividing curve on an orientable surface whose foliation satisfies the Poincaré-Bendixson property[25]. It will also highlight the qualitative properties of dividing curves.

Example 2.3.8. Suppose that $S \subset M$ is an ξ-convex surface in a contact manifold (M, ξ) whose characteristic foliation S_ξ satisfies the Poincaré-Bendixson property. Then we can draw a dividing curve Γ on S defined by the following regions:

- Around each elliptic point, embed a sufficiently small disk (containing no other singularities) with boundary transverse to S_ξ.
- Along the stable (respectively, unstable) separatrices of each positive (respectively, negative) hyperbolic points, embed a sufficiently thin neighbourhood (containing no other singularities except maybe at end points) with boundary transverse to S_ξ.
- Sufficiently close to a closed leaf, embed an annulus (containing no other singularities) transverse to S_ξ.

The dividing curve Γ is then realised by taking the union of disks, bands, and annuli containing positive singularities or repelling disks as a set A_+ and the others containing negative singularities or attracting disks as a set A_-. The connected components of $S \setminus (A_+ \cup A_-)$ consist of annuli which do not contain singularities or closed leaves of any sort.

The Poincaré-Bendixson property of S_ξ guarantees that each leaf entering (respectively, exiting) the outer boundary then exits (respectively, enters) the inner boundary, and so taking an interior circle in each annulus transverse to S_ξ as a component of Γ suffices. It is then not a surprising fact that you can define a volume form and vector field which has positive divergence on A_+ and negative on A_-, which then can be extended to one which is zero only on Γ, the regions $A_\pm$ being retracts of $S_\pm$.

Note that due to lack of retrograde saddle connections, the approach to handling hyperbolic points does not run into trouble. Due to lack of degenerate closed leaves, there is no possibility of a sequence of closed leaves with an accumulation point, and thus the approach with annuli for closed leaves does not run into trouble either.

The cumulative observations of this section, concluding with the above example, is the following theorem.

Theorem 2.3.9. *Suppose that $S \subset M$ is an oriented, closed[26] surface in a contact 3-manifold (M, ξ). If S_ξ satisfies the Poincaré-Bendixson property, then S is ξ-convex if and only if S_ξ has neither retrograde saddle connections nor degenerate closed leaves.*

[25]Recall that the Poincaré-Bendixson property of a vector field sates that the limit set of each half-orbit is either a singularity, a closed leaf, or a union of singularities connected by leaves.

[26]Again, an analogous theorem holds for compact surfaces with Legendrian boundary.

Thus Theorem 2.3.7 is merely a consequence of the above and Peixoto's theorem (Theorem 2.2.12) which implies that a generic foliation will satisfy the Poincaré-Bendixson property, have no saddle connections, and have only non-degenerate closed leaves.

We end this section on convexity by stating (without proof) Giroux's criterion for tight neighbourhoods in terms of dividing curves.

Theorem 2.3.10 (Giroux's criterion). *Suppose that $S \subset M$ is an ξ-convex surface in a contact 3-manifold (M, ξ) whose characteristic foliation is divided by Γ. Then S admits a tight neighbourhood if and only if one of the following holds:*

(i) *$S = S^2$ and Γ is a single circle.*

(ii) *$S \neq S^2$ and no component of Γ bounds a disk in S (that is, no component of Γ is contractible in S).*

2.4 Legendrian knots

For dimension 3, the Legendrian submanifolds of a contact manifold (M, ξ) are 1-dimensional. One particular class of 1-submanifolds that has its own prominent research field is the class of knots: embeddings of S^1 into 3-manifolds.

In Example 2.2.4 we have seen a particular example of a(n un)knot which is also a Legendrian submanifold. Such objects are aptly named **Legendrian knots**, embeddings $e\colon S^1 \to M$ such that

$$e'(\theta) \in \xi_{e(\theta)},$$

for each $\theta \in S^1$.

Classically, knot theorists try to classify knots up to smooth isotopy, characterising knots by algebraic invariants. With regards to Legendrian knots, we are interested not in ordinary isotopy but Legendrian isotopy (isotopy within the family of Legendrian knots). Classical knot invariants fail to recognize the nuances of the contact structure and the Legendrian condition, and so we will turn our attention to two particular invariants which happen to take into account these aspects.

Two Legendrian knots K_0 and K_1 are **Legendrian isotopic** if and only if there exists a smooth map $F\colon S^1 \times [0,1] \to M$ such that:

(i) $F_0(S^1) = K_0$.

(ii) $F_1(S^1) = K_1$.

(iii) $F_t(S^1)$ is Legendrian for each $t \in [0,1]$.

Using Gray's theorem, this is equivalent to there existing a 1-parameter family of contact diffeomorphisms $\{\phi_t \mid t \in [0,1]\}$ such that $\phi_0 = \mathrm{id}_M$ and $\phi_1(K_0) = K_1$. Hence a Legendrian isotopy is equivalently a contact isotopy.

2.4.1 Projections and constructions

Let $K \subset \mathbb{R}^3$ be a Legendrian knot embedded in the standard contact 3-manifold $(\mathbb{R}^3, \xi_0)$. Much like in the more classic study of knots, it is convenient to study knots by their projections into the plane. Recall that ξ_0 is spanned by the vector fields $\frac{\partial}{\partial y}$ and $\frac{\partial}{\partial x} + y\frac{\partial}{\partial z}$. Having in mind that the geometry of ξ_0 plays a huge role in the geometry of K, we consider the different projections of ξ_0 to various planes.

Example 2.4.1. Projecting ξ_0 into the x, y-plane gives us that the projection of K does not *obviously* have any obvious restrictions as compared to knots in the classical sense[27]. This is because a vector tangent to the projection of K is spanned by $\frac{\partial}{\partial x}$ and $\frac{\partial}{\partial y}$. For example, the knot K from Example 2.2.4 is projected to:

$$\pi_{x,y}K(t) = \begin{cases} \frac{1}{2}(1 - \cos(2t), -\sin(2t)), & t \in [0, \pi]; \\ \frac{1}{2}(\cos(2t) - 1, -\sin(2t)), & t \in [\pi, 2\pi]; \end{cases}$$

which does not look like anything unusual[28].

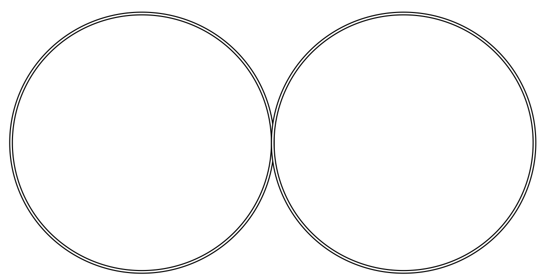

Example 2.4.2. On the other hand, projecting ξ_0 onto the x, z-plane means that the vectors tangent to the projection of K are spanned by $\frac{\partial}{\partial x} + y\frac{\partial}{\partial z}$. Now this is a special restriction on the vectors tangent to the projection as it will never point vertically along the z-axis. For example, the knot K from Example 2.2.4 is projected to:

$$\pi_{x,z}K(t) = \begin{cases} \frac{1}{2}(1 - \cos(2t), \frac{1}{8}\sin(4t) - \frac{1}{2}t), & t \in [0, \pi]; \\ \frac{1}{2}(\cos(2t) - 1, \frac{1}{2}t - \frac{1}{8}\sin(4t) - \pi), & t \in [\pi, 2\pi]; \end{cases}$$

which presents **cusps** instead of vertical tangencies.

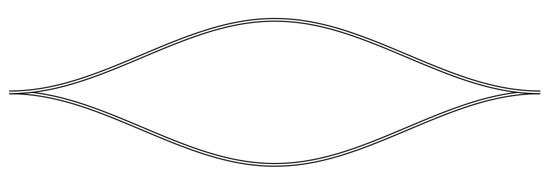

[27] Spoiler: there actually are, but this is not the projection we are after in this thesis. See [2] for more.

[28] Not to say that indicates that this projection is not useful, but just bear with me for a moment.

In the case for the projection to the y, z-plane, the vectors tangent to the projection of K are spanned by $\frac{\partial}{\partial y}$ and $y\frac{\partial}{\partial z}$, and so it does not seem to have as much of a restriction as the x, z-projection.

Observation 22. The projection of a Legendrian knot in $(\mathbb{R}^3, \xi_0)$ to the x, z-plane presents cusps instead of vertical tangencies.

The projection $\pi_{x,z} \colon \mathbb{R}^3 \to \mathbb{R}^2$ is called the **wavefront projection** when used in context with Legendrian knots in $(\mathbb{R}^3, \xi_0)$ due to Observation 22. We can perturb any Legendrian knot so that its projection has isolated double points.

On the other hand, if we are given a closed curve $\gamma(t) := (x(t), z(t))$ in the plane with isolated, ordinary double points[29], satisfying the condition that it has no vertical tangencies but instead cusps, then we may lift γ to a Legendrian knot. To do this, set $y(t) := \frac{z'(t)}{x'(t)}$, noting that the cusp condition makes this well-defined, and thus:

$$K(t) := \left(x(t), \frac{z'(t)}{x'(t)}, z(t)\right)$$

satisfies:

$$K'(t) = \left(\frac{d}{dt}\frac{z'(t)}{x'(t)}\right)\frac{\partial}{\partial y} + x'(t)\left(\frac{\partial}{\partial x} + \frac{z'(t)}{x'(t)}\frac{\partial}{\partial z}\right),$$

and hence $K'(t) \in \xi_0$. We will call a closed plane curve satisfying the same properties as γ a **Legendrian knot diagram**. Unlike usual knot diagrams, Legendrian knot diagrams are unambiguous at double points.

Observation 23. Any Legendrian knot diagram in $\mathbb{R}^2$ may be lifted to a Legendrian knot in (M, ξ). The lift is unique into $\mathbb{R}^3$, using coordinates (x, z) in $\mathbb{R}^2$ and defining a third coordinate y by:

$$y(t) := \frac{z'(t)}{x'(t)}.$$

The Darboux-Pfaff theorem then gives a knot in M.

Corollary 2.4.3. *A corollary of this observation is that any homotopy of Legendrian knot diagrams in $\mathbb{R}^2$ lifts uniquely to a Legendrian isotopy of Legendrian knots in $\mathbb{R}^3$.*

The proof of the following proposition has the rather unfortunate ability to be easily sketched using vague hand gestures and the frequent reliance on statements like "easy to see", "obvious", or "clearly".

[29] Also known as **crunodes**, we mean that the double point arises as a transverse intersection of two curves in some neighbourhood of the double point.

Proposition 2.4.4. *Any arc in a contact 3-manifold (M, ξ) may be C^0-approximated by a Legendrian arc, smoothly isotopic to the original relative to endpoints.*

A proof of Proposition 2.4.4 (which is sure to leave those who are disgruntled by the more common proofs with a smile on their faces[30]) can be found in the cute paper [14].

A direct corollary is the following life-altering theorem (assuming the reader owns a car of negligible length):

Theorem 2.4.5 (the parking theorem [10])**.** *Suppose one has a vehicle of negligible length which can only turn or move forward. That is, its movements are described by the phase space $M := \mathbb{R}^2 \times S^1$ with coordinates (x, y, θ) (position denoted by (x, y), and steering denoted by θ) and the limitations on steering gives a contact structure on M,*

$$\xi := \operatorname{span}\left\{\cos\theta \frac{\partial}{\partial x} + \sin\theta \frac{\partial}{\partial y}, \frac{\partial}{\partial \theta}\right\} = \ker(\sin\theta\, dx - \cos\theta\, dy),$$

such that any path the car takes is described by a Legendrian curve in M. Then no matter the circumstances, you can rest assured that you can park your car.

It is worth mentioning that though we placed particular stress on the wavefront projection of Legendrian arcs (owing in part to the unambiguous Legendrian knot diagrams), many analogous observations may be made about the x, y-projection. For example, one may recover a Legendrian knot from the x, y-projection as:

$$z(t) := z_0 + \int_0^t y(t)x'(t)\, dt,$$

where $z_0 = z(0)$ is an initial condition. Note that this initial condition is not needed for the wavefront projection.

2.4.2 The Thurston-Bennequin invariant

Consider a knot $K \subset M$ of an orientable contact 3-manifold (M, ξ) with normal bundle NK. Because M is orientable, the normal bundle to its Seifert surface S is trivial, giving us a non-vanishing section of NS. Because K is orientable in S, its normal bundle (taken within TS) is trivial, defining another non-vanishing section, this time of NK. These sections on NK never coincide and hence define a trivialisation of NK (that is, a canonical identification of NK with the trivial bundle $K \times \mathbb{R}^2$).

Example 2.4.6. Let $K \subset M$ be a Legendrian knot. Due to the Legendrian condition, at each $p \in K$ the contact plane ξ_p intersects transversally with N_pK, giving a line bundle ℓ, defined by $\ell_p := \xi_p \cap N_pK$. Using the tangent vector field τ to K, we may complete a non-vanishing section v of ℓ to a frame of NK by extending (τ, v) to an oriented frame (τ, v, w) of M. The frame of NK is given by the sections v and w.

[30] Count me as one of the disgruntled.

Intuitively, a non-vanishing section of NK makes K into an orientable "ribbon", giving it "breadth".

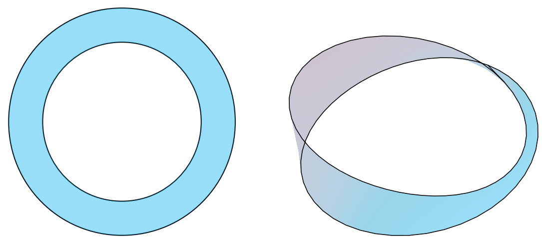

However, we would like to say that the "ribbons" are equivalent if we can isotope one into the other. We call an isotopy class of a trivialisation of NK a **framing** of K. The above drawings illustrate two different framings of the unknot. Geometrically, framings differ from each other by number of twists, and—after fixing a convention for positive and negative twists—one can discuss the difference in number of twists between two framings. Generally, one cannot canonically decide which framing has zero twists, and so without such a "zero framing", the set of all framings forms a set similar to an affine space called a $\mathbb{Z}$-torsor, where no "zero" element is chosen, but "differences" between elements take values in $\mathbb{Z}$.

Though no zero framing may be canonically determined, given extra information, some framings are natural. For example, in the presence of a contact structure, the framing of the Legendrian knot K defined in Example 2.4.6 is called the **contact framing** of K which one gets for free from the Legendrian condition (though we stress that a framing can be defined for any type of knot, not just Legendrian knots).

Another example of a natural framing arises from the presence of a particular spanning surface for K. Given an embedding of a null-homologous knot K, we can discuss various surfaces $S \subset M$ with boundary $\partial S = K$. If S is connected, compact, and orientable, it is called a **Seifert surface** of K. It is well known that if $M = \mathbb{R}^3$ we can always find[31] a Seifert surface whose orientation induces that of K (using what is known as the Seifert algorithm). Any Seifert surface S of K induces its own framing, resulting from shrinking L into the interior of S (effectively taking a sufficiently small closed collar neighbourhood of K in S). This framing is called the **Seifert framing** of K with respect to the Seifert surface S. Because every Seifert surface is "tube equivalent" (equivalent up to isotopy and a finite addition/removal of tubes), the Seifert framing is independent of the choice of Seifert surface.

We now take a moment to formalise the $\mathbb{Z}$-torsor of framings. Consider two framings $\mathcal{F}$ and $\mathcal{F}'$ of a knot K with associated non-vanishing sections $v, v' \colon K \to NK$. At each $p \in K$, there exists an invertible linear map $\Phi_p \in$

[31]It is not unique, just add any number of "tubes" to it that does not obstruct orientability. Additionally, Seifert's algorithm extends to arbitrary M given that K is null-homologous.

$\mathrm{GL}^+(N_pK) \cong \mathrm{GL}^+(2,\mathbb{R})$ such that:

$$\Phi_p(v_p) = v'_p,$$

and thus we have a resulting map $\Phi\colon K \to \mathrm{GL}^+(2,\mathbb{R})$ defined by $p \mapsto \Phi_p$. The framings themselves are isotopy classes and so the difference between $\mathcal{F}$ and $\mathcal{F}'$ is captured in the homotopy class of the map Φ. Because of the deformation retract (using Gram-Schmidt) from $\mathrm{GL}^+(n,\mathbb{R}) \to \mathrm{SO}(n)$, the homotopy class of Φ can be seen as a map $[\Phi]\colon K \to SO(2) \cong S^1$, and the **difference** $\mathcal{F} - \mathcal{F}'$ between the framings is defined to be the degree of this map, seen as a map $S^1 \to S^1$.

Now given a null-homologous Legendrian knot $K \subset M$ in a contact 3-manifold (M,ξ), we get two natural framings:

- The contact framing, $\mathcal{C}$.
- The Seifert framing, $\mathcal{S}$.

It is now quite natural to ask what the difference between these framings is. This gives us our first knot invariant for Legendrian knots, called the **Thurston-Bennequin invariant**, defined to be:

$$\mathrm{tb}(K) := \mathcal{C} - \mathcal{S},$$

the difference between the contact and Seifert framings of K. So the invariant $\mathrm{tb}(K)$ measures the signed rotation of the contact framing relative to the Seifert framing.

Considering the contact framing of K as an orientable ribbon, there are two edges, one being the knot K, and the other "framing strand" twisting around K itself, defining the breadth of the ribbon. If S is a Seifert surface for K, then as the framing strand twists about K, it (generically) intersects transversally with S, a full twist in one direction constituting a new intersection.

If the framing strand intersects from the negative to the positive side of S, we call this a positive intersection (corresponding to right-hand twists, thumb in the direction of the knot, hand curling in the direction of the twist), the opposite being a negative intersection. The signed sum of these intersections gives the net number of twists of the contact frame relative to the Seifert surface S—this is exactly the Thurston-Bennequin number of K!

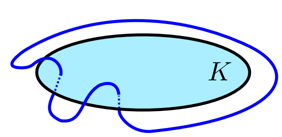

More generally, if K and K' are two knots with Seifert surfaces S and S' respectively, the **linking number** $\mathrm{lk}(K,K')$ of the knots is defined to be the signed intersection between K and S'. It follows from tube-equivalence of Seifert surfaces that this is again invariant under choice of Seifert surface, after considering the ways a knot might intersect with a tube.

It also turns out that lk is symmetric. Without loss of generality we may assume that S and S' intersect transversally. Any intersection is either an embedding of a circle, or an arc. The arcs either have one end point in K and the other in K', or both end points are on the same knot. In the former case, the signs will be the same, while in the latter case they will be opposite and hence contribute net 0 to the linking number. Hence $\mathrm{lk}(K, K') = \mathrm{lk}(K', K)$.

With this new notation, we may push points $p \in K$ along a vector field v such that $v_p \in \xi_p \cap N_p K$ for each p, and get a new knot K' such that:

$$\mathrm{tb}(K) = \mathrm{lk}(K, K').$$

From a link diagram (a knot diagram of two knots, K and K', usually entangled), the linking number $\mathrm{lk}(K, K')$ can be calculated by enumerating certain types of crossings in the diagram. In particular, consider the following conventions[32]:

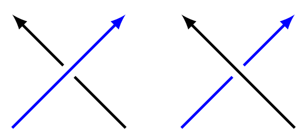

$+$ $\qquad$ $-$

A **positive** and **negative** crossing, respectively.

Heuristically, it takes two positive crossings to make a positive, right hand twist (and hence contribute $+1$ to $\mathrm{lk}(K, K')$), and similarly for negative crossings in terms of negative left hand twists. With this convention, we have that:

$$\mathrm{lk}(K, K') = \frac{1}{2}\left(\#\text{ positive crossings} - \#\text{ negative crossings}\right).$$

Example 2.4.7. Consider a knot $K \subset \mathbb{R}^3$, $\mathbb{R}^3$ being endowed with the standard contact structure, $\xi_0 = \ker(dz - y\,dx)$. The vector field $\frac{\partial}{\partial z}$ is transverse to the contact structure and so we may push off any knot vertically to a new knot K'.

Using the wavefront projection, this is very easy to imagine, and we can analyse the types of crossings that arise. In particular:

- For each crossing in K, we get two crossings between K and K' of the same sign as the original.
- For each cusp in K, we get a new negative crossing between K and K'.

32 A mnemonic to recall these conventions is to use the right hand rule to remember a positive crossing: Your thumb directs the top strand, and your fingers curl "underneath" (indeed it does not matter which side you look from) your thumb in the direction of the bottom strand. The negative crossing thus would be remembered as *not* being the right hand rule, or perhaps just the left hand rule.

This means that, calculating the linking number using the signed crossings in the link diagram, we have:

$$\mathrm{tb}(K) = -\frac{1}{2}\#\text{ cusps} + \#\text{ positive crossings} - \#\text{ negative crossings}.$$

2.4.3 The rotation number

Consider a null-homologous, oriented Legendrian knot K in a contact manifold (M, ξ). Given a Seifert surface S of K, we have that $\xi|_S$ is trivial (as it is orientable as per Observation 9), and so

$$\xi|_{\partial S} = K \times \mathbb{R}^2.$$

Consider a non-vanishing vector field v tangent to K whose direction matches the orientation of K, we view it as residing in the trivialisation induced by ξ, indeed as a map $t\colon S^1 \to S^1$. The degree of this map,

$$\mathrm{r}(K) := \deg t,$$

is called the **rotation number** of K (a.k.a. the **Maslov number**). Reversing the orientation of K changes the sign of $\mathrm{r}(K)$.

Example 2.4.8. Consider a null-homologous knot $K \subset \mathbb{R}^3$, $\mathbb{R}^3$ being endowed with the standard contact structure $\xi_0 := \ker(dz - y\,dx)$. Trivialize ξ_0 using the oriented frame

$$\left(\frac{\partial}{\partial y}, \frac{\partial}{\partial x} + y\frac{\partial}{\partial z}\right),$$

and then note that $\mathrm{r}(K)$ will be the number of times the tangent vector crosses $\frac{\partial}{\partial y}$ while going around K. With respect to the wavefront projection, half rotations around $\frac{\partial}{\partial y}$ occur at the cusps, and so it follows that $\mathrm{r}(K)$ may be calculated as:

$$\mathrm{r}(K) = \frac{1}{2}(\#\text{ downward cusps} - \#\text{ upward cusps}),$$

where upward cusps are those where the z-coordinate is increasing as you traverse the cusp in the orientation of K, and similarly downward cusps are those where it is decreasing.

The relation between a null-homologous Legendrian knot, a Seifert surface, and the invariants tb and r is given by the Thurston-Bennequin inequality:

Theorem 2.4.9 (**Thurston-Bennequin inequality** [9][4.6.36])**.** *Let K be a Legendrian knot in $(\mathbb{R}^3, \xi_0)$ with Seifert surface S. Then:*

$$\mathrm{tb}(K) + |\,\mathrm{r}(K)| \leqq -\chi(S).$$

Fascinatingly enough, (M, ξ) is tight if and only if the Thurston-Bennequin inequality holds for every null-homologous Legendrian knot in M.

The intrigued reader is encouraged to consult [2] for more on both knot theory as a field of its own, and as it relates to contact topology. The paper [8] also contains a good introduction to Legendrian knots.

3 Machinery

Throughout this section, all contact structures are assumed to be positive and co-oriented. All surfaces are assumed to be oriented, and (unless otherwise noted) closed.

3.1 Deformation lemma

Given an open disk D whose characteristic foliation contains no singularities, Lemma 3.1.1 helps us show that we can change the leaf that any entering leaf exits on. As with many of the lemmas the reader will encounter in the next few subsections, this lemma explicitly makes use of a model surface in the standard contact manifold $(\mathbb{R}^3, \ker(dz - y\,dx))$. The result then can be carried over to the general case with the help of the Moser-Weinstein theorem. But first a little notation!

Suppose that S is an embedded hypersurface in a contact manifold (M, ξ). Define the **leaf function of** S to be the map leaf_S on S which returns the leaf, $\mathrm{leaf}_S(p)$, of the foliation S_ξ for which $p \in M$ belongs to.

Next, define a **smooth deformation** of a surface $S \subset M$ to be a family $\{\psi_t\colon S \to M \mid t \in [0,1]\}$ of smooth maps, diffeomorphic onto their images, varying smoothly in t, such that ϕ_0 is the inclusion.

Lemma 3.1.1 (**Deformation lemma**). *Consider $\mathbb{R}^3$ with the contact structure $\xi_0 = \ker(dz - y\,dx)$, and let P denote the embedded x,z-plane. If $U \subseteq \mathbb{R}^3$ is an open neighbourhood of $0 \in \mathbb{R}^3$ and $D \subseteq P \cap U$ an open disk of radius $r > 0$ centred at 0, then there exists a smooth deformation $\{\psi_t \mid t \in [0,1]\}$ of $P \cap U$ such that:*

- *(i) ψ_0 is the inclusion $i_{P\cap U}\colon P \cap U \hookrightarrow \mathbb{R}^3$.*
- *(ii) For each $t \in [0,1]$, $\psi_t(p) = p$ for each $p \in (P \cap U) \setminus D$.*
- *(iii) For each $t \in [0,1]$, the characteristic foliation of $P_t := \psi_t(P \cap U)$ contains no singularities.*
- *(iv) There exists an $\varepsilon \neq 0$ such that $\mathrm{leaf}_{P_1}((-r,0,0)) = \mathrm{leaf}_{P_1}((r,0,\varepsilon))$.*

Moreover, there is another smooth deformation $\{\psi'_t \mid t \in [0,1]\}$ satisfying (i), (ii), and (iii) such that $\mathrm{leaf}_{P_1}((-r,0,0)) = \mathrm{leaf}_{P_1}((r,0,-\varepsilon))$.

Proof. Let $p_0 = (-r,0,0)$ and $p_1 = (r,0,0)$. Pick $0 < \delta < r$ small enough so such that for all $p \in D$, we have $p \pm (0,\delta,0) \in U$.

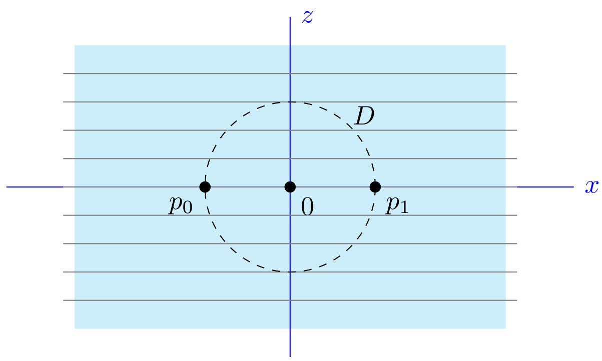

Define a function $f\colon \mathbb{R}^2 \to \mathbb{R}$ by:

$$f(x,y) := \begin{cases} \frac{\delta}{\exp(-1/r)} \exp(-(r - \|(x,y)\|)^{-1}), & (x,y) \in D; \\ 0, & (x,y) \notin D, \end{cases}$$

(a bump function in the positive y-direction) and for each $t \in [0,1]$ define $\psi_t\colon P \cap U \to \mathbb{R}^3$:

$$\psi_t(x,0,z) = (x, tf(x,z), z)$$

Clearly ψ_t is a diffeomorphism onto its image, and both of (i) and (ii) are satisfied.

For each $t \in [0,1]$, we will explicitly calculate the characteristic foliation of P_t. Use ψ_t as a parametrisation $\phi\colon \mathbb{R}^2 \to P_t$ of P_t:

$$\phi(x,z) := \psi_t(x,0,z) = (x, tf(x,z), z),$$

and hence if $q := \phi(x,z) \in P_t$ we have $\operatorname{im}\phi_{*,q} = T_q P_t$. In coordinates, we see that,

$$\phi_{*,q} = \begin{bmatrix} 1 & 0 \\ \partial_x(tf(x,z)) & \partial_z(tf(x,z)) \\ 0 & 1 \end{bmatrix}$$

and so when using the isomorphism $T_q\mathbb{R}^n \cong \mathbb{R}^n$,

$$T_q P_t = \operatorname{span}\left\{ \begin{bmatrix} 1 \\ \partial_x(tf(x,z)) \\ 0 \end{bmatrix}, \begin{bmatrix} 0 \\ \partial_z(tf(x,z)) \\ 1 \end{bmatrix} \right\},$$

while ξ_0 is defined as the kernel of $dz - y\,dx$, so:

$$(\xi_0)_q = \operatorname{span}\left\{ \begin{bmatrix} 0 \\ 1 \\ 0 \end{bmatrix}, \begin{bmatrix} 1 \\ 0 \\ tf(x,z) \end{bmatrix} \right\}.$$

Clearly, if $t = 0$ or $q \notin D$, then $tf(x,z) = 0$ and the intersection $T_q P_t \cap (\xi_0)_q$ is spanned by the vector $\begin{bmatrix} 1 & 0 & 0 \end{bmatrix}^T$. If $q \in D$ and $t > 0$, then we solve for the intersection as a system of linear equations as follows:

$$\begin{bmatrix} 0 & 1 & -1 & 0 \\ 1 & 0 & -\partial_x(tf(x,z)) & -\partial_z(tf(x,z)) \\ 0 & tf(x,z) & 0 & -1 \end{bmatrix}$$
$$\longrightarrow \begin{bmatrix} 1 & 0 & 0 & -\partial_z(tf(x,z)) - \frac{\partial_x(tf(x,z))}{tf(x,z)} \\ 0 & 1 & 0 & -1/(tf(x,z)) \\ 0 & 0 & 1 & -1/(tf(x,z)) \end{bmatrix}$$

which has the solution:

$$\begin{bmatrix} \partial_z(tf(x,z)) + \frac{\partial_x(tf(x,z))}{tf(x,z)} \\ 1/(tf(x,z)) \\ 1/(tf(x,z)) \\ 1 \end{bmatrix} s, \qquad s \in \mathbb{R}.$$

This gives the intersection as the line with direction vector

$$v_q = \begin{bmatrix} 1/(tf(x,z)) \\ \partial_z(tf(x,z)) + \frac{\partial_x(tf(x,z))}{tf(x,z)} \\ 1 \end{bmatrix},$$

and so the characteristic foliation of P_t is without singularities (TP_t never coincides with ξ_0), hence (iii) is satisfied.

To prove (iii), the point is to see that the slope of the foliation with respect to the x, z-plane will always be non-negative, and in particular, strictly positive on D. Thus $\mathrm{leaf}_{P_1}(p_0) \neq \mathrm{leaf}_{P_1}(p_1)$!

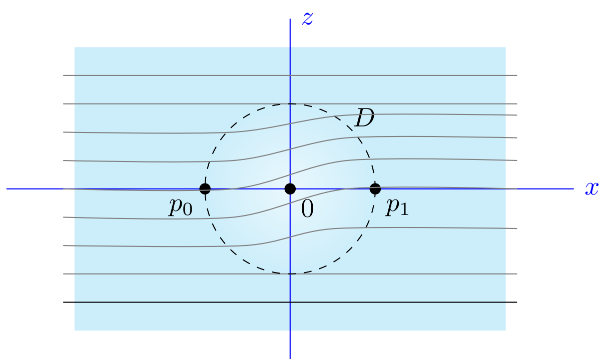

To this end, let $t \neq 0$ and note that $tf(x,z) > 0$ on D as a result of our parametrisation. The slope of the projection of the direction vector v_q to the x, z-plane is $\begin{bmatrix} 1/(tf(x,z)) \\ 1 \end{bmatrix}$ which has a positive slope on D. Meanwhile, on $P_t \setminus D$, the foliation is always horizontal lines of slope 0, and hence p_0 and p_1 cannot possibly be on the same leaf in the characteristic foliation of P_t. Hence $\mathrm{leaf}_{P_t}(p_0) = \mathrm{leaf}_{P_t}(r, 0, \varepsilon)$ for some $\varepsilon > 0$. Hence (iii) is satisfied. Moreover, note that by utilizing $g := -f$ in place of f results in $\mathrm{leaf}_{P_t}(p_0) = \mathrm{leaf}_{P_t}((r, 0, -\varepsilon))$. ■

Porism 3.1.2. *For each $s \in [0, \varepsilon]$ there exists a $t \in [0, 1]$ such that:*

$$\text{leaf}_{P_t}(p_0) = \text{leaf}_{P_t}((r, 0, s)).$$

Proof. Let $0 < t' < t \leqq 1$ and $p \in B(0; r) \subset \mathbb{R}^2$. We wish to prove that there exists a $R > 0$ such that for all $x \in B(p; R) \subset \mathbb{R}^2$ we have that:

$$tf(p) > t' f(x).$$

That is to say, proving this would mean that the slope of leaves in $(P_{t'})_{\xi_0}$ at time t' who project sufficiently close to a point p in the x, z-plane can be guaranteed to have a smaller slope than $\text{leaf}_{P_t}(\psi_t(p))$ at time t.

Note for $0 = t' < t \leqq 1$ this is trivial, and if $p = (0, 0)$, then any $R > 0$ works. From now on, we suppose that $p \neq (0, 0)$.

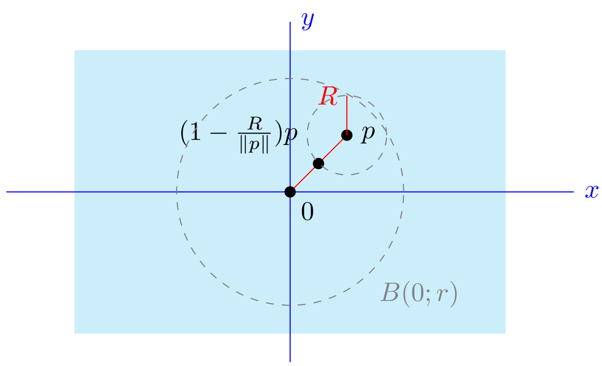

Note that for any $0 < R < r$, we have that:

$$\sup_{x \in B(p;R)} f(x) = f((1 - R/\|p\|)p),$$

and so it suffices to find an R such that:

$$tf(p) > t' f((1 - R/\|p\|)p).$$

First, note that the condition that:

$$\ln(t/t') \geqq \frac{1}{r - \|p\|}, \tag{$\star$}$$

means we can pick any $R > 0$ (even $R = 2r$, as

$$\ln(t/t') > \frac{1}{r - \|p\|} - \frac{1}{r} \implies tf(p) > t' f(0, 0).$$

Otherwise, we assume the negation of ($\star$) and note that we need to seek out an

R such that:

$$\begin{aligned}
&tf(p) > t'f((1-R/\|p\|)p) \\
\iff\ &t\exp(-(r-\|p\|)^{-1}) > t'\exp(-(r-\|p\|+R)^{-1}) \\
\iff\ &\frac{1}{r-\|p\|} - \ln(t/t') < \frac{1}{r-\|p\|+R} \\
\iff\ &C(r-\|p\|) + CR < 1, \qquad C := \frac{1}{r-\|p\|} - \ln(t/t') > 0; \\
\iff\ &R < \frac{1-C(r-\|p\|)}{C},
\end{aligned}$$

and so we can pick $R = \frac{1-C(r-\|p\|)}{2C}$ as our neighbourhood radius. Now, to actually prove the porism, we simply note that if

$$\mathrm{leaf}_{P_t}(p_0) = \mathrm{leaf}_{P_t}((r,0,s)), \qquad \mathrm{leaf}_{P_{t'}}(p_0) = \mathrm{leaf}_{P_{t'}}((r,0,s')),$$

then what we have shown above implies $0 \leqq s' < s \leqq \varepsilon$. Hence using a Bolzano-Weierstrass-esque argument, we get a sequence $(t_n) \subset [0,1]$ which converges to a t such that $\mathrm{leaf}_{P_t}(p_0) = \mathrm{leaf}_{P_t}((r,0,s))$. ■

3.2 Breaking a saddle connection

As a result of the deformation lemma (Lemma 3.1.1), we have an immediate application. The following lemma says that we may deform a surface which contains saddle connections to one which does not.

Lemma 3.2.1. *Let (M,ξ) be a co-orientable contact 3-manifold, and suppose S is a surface in M. Let S_ξ be its characteristic foliation. If we have a saddle connection on S between two hyperbolic singularities $x_0, x_1 \in S$, where x_0 flows into x_1, then there exists a deformation $\{\psi_t\colon S \to M \mid t \in [0,1]\}$ of S such that the points $\psi_t(x_0)$ and $\psi_t(x_1)$ in $\psi_t(S)$ are no longer connected.*

Moreover, given that there are finitely many singularities in the foliation, we can do this in a way that produces no other saddle connections.

Proof. Let (M,ξ) be a co-orientable contact 3-manifold, and suppose S is a surface in M. Let S_ξ be its characteristic foliation. Suppose that we have a saddle connection on S between two hyperbolic singularities $x_0, x_1 \in S$, where x_0 flows into x_1.

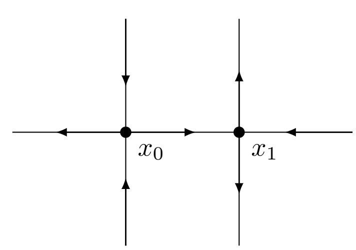

We want to "break" it so that the separatrix of x_0 no longer flows into x_1. Let p be a point on the separatrix connecting x_0 to x_1, and let B' be an open neighbourhood in M around p which does not contain the singularities x_0 and x_1. Define $B := S \cap B'$.

Now consider the standard contact structure ξ_0 on $\mathbb{R}^3$ given by $\alpha_0 := dz - y\,dx$, and the x, z-plane P as a surface embedded in $\mathbb{R}^3$. There exists an open neighbourhood $U \subseteq \mathbb{R}^3$ of 0 and a diffeomorphism $\eta\colon B \to P \cap U$ such that $\eta(B_\xi) = (P \cap U)_{\xi_0}$ (as foliations) and $\eta(p) = 0$.

By the Moser-Weinstein theorem (Theorem 2.2.13) there exists a contactomorphism φ between neighbourhoods of B and $P \cap U$ (respecting the characteristic foliation) such that $\varphi(B) = P \cap U$ and $\varphi|_B = \eta$. Without loss of generality, we assume the neighbourhoods are B' of B and U of $P \cap U$.

The characteristic foliation $\varphi(B)_{\xi_0}$ is that of lines parallel to the x-axis on $P \cap U$. Pick some open disk $D \subseteq \varphi(B)$ centred at $\varphi(p) = 0$. Let L_0 denote the leaf of the characteristic foliation D_{ξ_0} containing 0, then denote the two lonely elements of $\partial D \cap L_0$ by $y_0, y_1 \in \varphi(B)$.

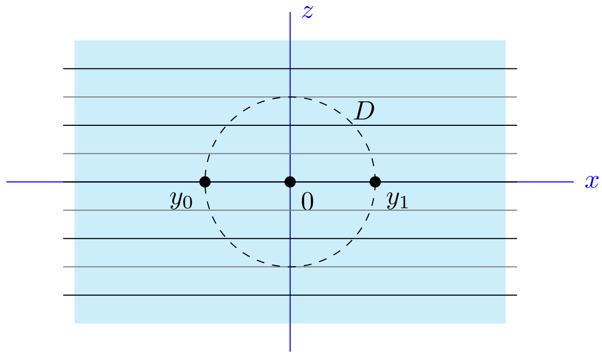

By Lemma 3.1.1, there is a smooth isotopy of $P \cap U$ such that the leaf containing y_0 does not contain y_1. Because this deformation preserves the position of y_0 and y_1, then we may compose it appropriately with the contactomorphism φ which took B to $P \cap U$. As a result, we have a deformation of S where x_0 and x_1 are no longer a saddle connection.

If the characteristic foliation contains finitely many singularities, then Porism 3.1.2 allows us to choose a sufficiently small deformation so that no intermediate steps in the deformation produces saddle connections. ■

Observation 24. Note that by Peixoto's theorem, it is already a generic property of characteristic foliations on compact surfaces that there are no saddle connections, but this can be simply seen as a corollary of the deformation lemma (Lemma 3.1.1) without resorting to Peixoto's tank[a].

[a]Sometimes it is fun to crush bugs with a tank, but times are changing and we ought to consider the fuel economy.

3.3 Corner lemma

The goal of the lemma we prove in this section is to be able to use a piece-wise linear deformation, and then smooth the sharp "corners" to a smooth deformation. The benefit of doing this is that, as seen in the proof of Lemma 3.1.1, use of smooth bump functions and the like tend to present verbose arguments[33] to show geometrically intuitive phenomena.

Consider the piece-wise linear surface in $\mathbb{R}^3$ (endowed with the standard contact structure) defined as:

$$C := (\{y = 0\} \cap \{x \leq 0\}) \cup (\{x = 0\} \cap \{y \geqq 0\}).$$

Visually, it is a corner along the z-axis.

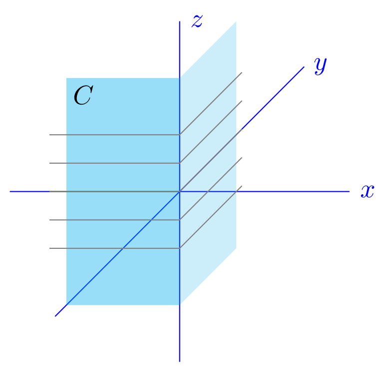

We consider smoothings of the corner C of various radii $r > 0$, defining:

$$\begin{aligned} C_r :=& (\{y = 0\} \cap \{x \leqq -r\}) \\ &\cup (\{(x + r)^2 + (y - r)^2 = r^2\} \cap \{-r < x < 0 \text{ and } 0 < y < r\}) \\ &\cup (\{x = 0\} \cap \{y \geqq r\}) \end{aligned}$$

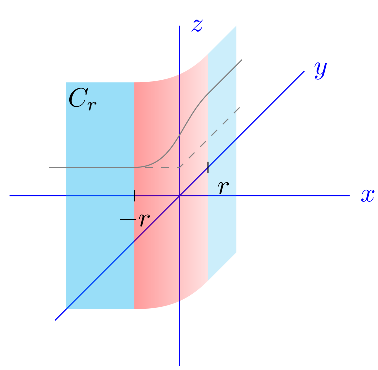

[33]There is of course always the possibility that the author has failed to note significantly shorter arguments.

The idea is that the leaf containing the point $(-r,0,0)$ will no longer correspond to the leaf containing $(0,r,0)$ as it does in the characteristic foliation of C, however we can control the degree to which the z-coordinate varies by choosing an appropriately small radius r for the smoothing.

Lemma 3.3.1 (**Corner lemma**). *For each $\varepsilon > 0$ there exists a $\delta > 0$ such that if $0 < r < \delta$ and*

$$\mathrm{leaf}_{C_r}(-r,0,0) = \mathrm{leaf}_{C_r}(0,r,z),$$

then $0 < z < \varepsilon$.

Proof. The route we will take is the following: the "slope" of the leaf $\ell := \mathrm{leaf}_{C_r}(-r,0,0)$ is positive and bounded, and so we will pick δ so that the worst case scenario of a leaf (albeit only piece-wise smooth) of constant maximum "slope" on the "elbow" of C_r has its z-coordinate bounded by ε.

First, abbreviate the "elbow" of C_r as

$$E := \{(x+r)^2 + (y-r)^2 = r^2\} \cap \{-r < x < 0 \text{ and } 0 < y < r\}.$$

Recall that the standard contact structure on $\mathbb{R}^3$ is given by $\xi_0 = \ker\alpha$ where $\alpha := dz - y\,dx$. We will now calculate the direction vector of leaves on E by calculating $(\xi_0)_p \cap T_pE$. After identifying $T_p\mathbb{R}^3 \cong \mathbb{R}^3$, for $p = (p_1,p_2,p_3) \in E$, we have:

$$\begin{aligned} \ker\alpha_p &= \mathrm{span}\{(0,1,0),(1,0,p_2)\}, \\ T_pE &= \mathrm{span}\{(r-p_2, p_1+r, 0),(0,0,1)\}. \end{aligned}$$

We solve for the intersection as a system:

$$\begin{bmatrix} 0 & 1 & p_2 - r & 0 \\ 1 & 0 & -p_1 - r & 0 \\ 0 & p_2 & 0 & -1 \end{bmatrix} \longrightarrow \begin{bmatrix} 1 & 0 & 0 & \frac{p_1+r}{p_2(p_2-r)} \\ 0 & 1 & 0 & -1/p_2 \\ 0 & 0 & 1 & \frac{1}{p_2(p_2-r)} \end{bmatrix}$$

and so the system has a solution:

$$\begin{bmatrix} -\frac{p_1+r}{p_2(p_2-r)} \\ 1/p_2 \\ -\frac{1}{p_2(p_2-r)} \\ 1 \end{bmatrix} s, \qquad s \in \mathbb{R},$$

and so at $p \in E$, the leaf passing through p has direction vector:

$$v_p = -\frac{1}{p_2(p_2-r)} \begin{bmatrix} r - p_2 \\ p_1 + r \\ 0 \end{bmatrix} + \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix} = \begin{bmatrix} 1/p_2 \\ \frac{p_1+r}{p_2(r-p_2)} \\ 1 \end{bmatrix}.$$

Now, let $w_p := (r-p_2, p_1+r, 0)$, and $\hat{z} := (0,0,1)$. We want to calculate the slope m_p of v_p with respect to the ordered orthonormal basis $\left(\frac{w_p}{r}, \hat{z}\right)$. This is given as the following quotient of the Fourier coefficients:

$$m_p := \frac{\hat{z}\cdot v_p}{\frac{w_p}{r}\cdot v_p} = \frac{r}{\frac{r-p_2}{p_2} + \frac{(p_1+r)^2}{p_2(r-p_2)}} = \frac{p_2(r-p_2)}{r}.$$

Note well that $m_p > 0$ because $p_2 < r$ on E. Then we note that:

$$\sup\left\{ \frac{p_2(r-p_2)}{r} \middle| p_2 \in (0,r) \right\} = \frac{r}{4},$$

obtained by basic calculus, the steepest possible slope of a leaf at any given point. Our worst-case scenario is thus the leaf which has a slope $r/4$ on E, and a slope of 0 off of E.

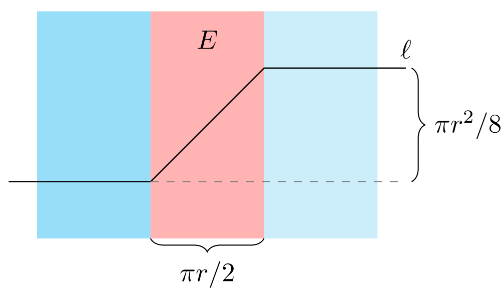

This means that, if given $r > 0$, then $\ell = \mathrm{leaf}_{C_r}(0, r, z)$ for some $0 < z \leqq \pi r^2/8$. We thus pick $\delta > 0$ such that:

$$\frac{\pi^2\delta^2}{8} < \varepsilon \Longrightarrow \delta < 2\sqrt{2\varepsilon/\pi}.$$

The delta $\delta := \sqrt{2\varepsilon/\pi}$ will do, and hence the proof follows from the discussion and calculations above.

■

3.4 Pivot lemma

One of the fine details of Giroux's elimination lemma is that the merging of the hyperbolic-elliptic pair results in the merging of certain leaves and separatrices.

For example, if γ is the stable separatrix of a positive hyperbolic point, p, which flows from a positive elliptic point, q, then use of Giroux's elimination lemma would result in the leaf opposite and smoothly connected to q to merge with the stable separatrix of p other than γ.

At first glance, we might just have to live our life with this constraint. However with the so-called pivot lemma which we prove below, one can show that at any two leaves of an elliptic point may be smoothly connected after a smooth deformation. This gives us extra control when using Giroux's elimination lemma, not only for elimination, but also for introduction of singularities.

It is also worth noting that during the proof of the lemma, we will make use of piecewise linear deformations and then use the corner lemma (Lemma 3.3.1) to smooth it. This is in contrast to our approach in the proof of Lemma 3.1.1.

Lemma 3.4.1 (**Pivot Lemma**). *Let (M,ξ) be a co-orientable contact 3-manifold, and $S \subset M$ a surface with characteristic foliation S_ξ. If $p \in S$*

is an elliptic point in S_ξ and ℓ_0 and ℓ_1 are any two leaves of p, then there exists a smooth deformation of S so that ℓ_0 and ℓ_1 are smoothly connected.

Proof. Let $U = S \cap U'$ be an open neighbourhood of $p \in S$ containing no other singularities of S_ξ. There exists an open disk D of radius r in the x, y-plane about $0 \in (\mathbb{R}^3, \xi_1)$ (for $\xi_1 := \ker(dz + r^2\, d\theta)$) and a diffeomorphism $\eta\colon U \to D$ such that $\eta(U_\xi) = D_{\xi_1}$ as foliations, and $\eta(p) = 0$. We also assume without loss of generality that the leaf ℓ_0 coincides (under the contactomorphism) with the leaf of D_{ξ_1} along the positive x-axis, and that ℓ_1 coincides with the leaf along the ray of p making an angle of $\theta_1 \in (\pi, 2\pi)$ with the x-axis. Note that if $\theta_1 = \pi$, we are already done, so we exclude this triviality.

There exists a contactomorphism φ between open sets $U' \subseteq M$ and some $B \subseteq \mathbb{R}^3$ containing D such that $\varphi(U) = \varphi(U' \cap S) = D$ and $\varphi|_U = \eta$. Without loss of generality, we assume that B is the open ball of radius r about $0 \in \mathbb{R}^3$.

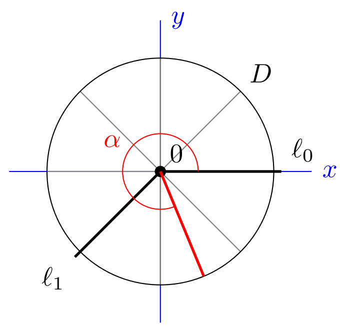

Let $\alpha := \theta_1/2 + \pi$. We prove the following short lemma:

Lemma 3.4.2. *There is a piece-wise linear deformation of D into a "top hat" of radius η and height h within B such that the leaf corresponding to ℓ_0 flows into the spoke α of the elliptic singularity 0. Moreover, it can be chosen so that the θ coordinate of the leaf corresponding to ℓ_0 is always within $[\alpha, 0]$ (modulo 2π).*

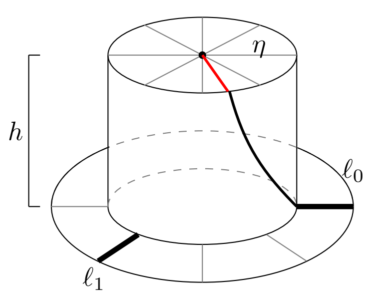

Proof. Pick $h' := r - \eta$ for some $0 < \eta < r$ such that

$$\alpha \equiv -\frac{h'}{\eta^2} \pmod{2\pi}.$$

Consider the piece-wise linear deformation of D where the subdisk $D' \subset D$ of radius η is lifted to height h', the surface deformed into a "top hat". Note it is still contained in B.

We see that the helix from the base of the cylinder is parametrized[a] by:

$$H_1 : \begin{cases} r(t) = \eta \\ \theta(t) = -t \\ z(t) = \eta^2 t \end{cases} \qquad t \in [0, h'/\eta^2]$$

which then hits $z(t) = h'$ at $t = h'/\eta^2$, hence enters the disk D' at the angle $\alpha \equiv -h'/\eta^2 \pmod{2\pi}$. Moreover, we can lower the height of the cylinder to a suitable height $h = h' - 2\pi\eta^2 n$, where H_1 achieves $\theta(t) = \alpha$ for $n+1$ distinct t. ■

[a] Note that it can be shown that the intersection of the tangent plane of the cylinder at q and $(\xi_1)_q$ is spanned by $(0, -1/\eta^2, 1)$.

Normally, the leaf corresponding to ℓ_0 would then pass through the singularity at $(0,0,h)$ back to the edge of the cylinder at $(\eta, \alpha+\pi, h)$, where it spirals down to $(\eta, \pi, 0)$ and exits along the negative x-axis. We instead shrink the radius of the $\theta \in [\pi/2, \alpha - \varepsilon]$ portion of the cylinder for some $\varepsilon > 0$ such that $\theta_1 < \alpha - \varepsilon < \alpha \pmod{2\pi}$. In particular, we shrink the radius to the largest $0 < \rho < \eta$ such that:

$$\frac{\theta_1}{2} \equiv \frac{h}{\rho^2} \pmod{2\pi},$$

so that the helix on the way down is defined by:

$$H_2 : \begin{cases} r(t) = \rho \\ \theta(t) = \alpha + \pi + t \equiv \theta_1/2 + t \\ z(t) = h - \rho^2 t \end{cases} \qquad t \in [0, h/\rho^2]$$

and meets at time $t = h/\rho^2$ the point

$$(\rho, \theta_1/2 + h/\rho^2, h - \rho^2(h/\rho^2)) = (\rho, \theta_1, 0).$$

The leaf then continues radially on the leaf corresponding to ℓ_1.

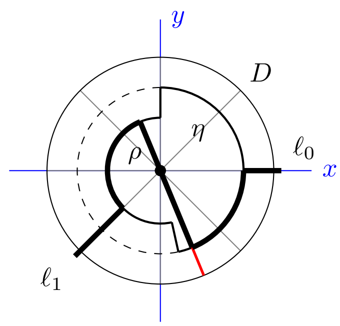

Finally, to smooth this piece-wise linear construction, we know by Lemma 3.3.1 that for any $\varepsilon > 0$, we can choose a sufficient δ-small smoothing of the corners of the surface such that the leaf in question only varies from the piece-wise linear on the surface by at most ε distance. We then reverse the effects of the smoothing by applying Lemma 3.1.1 and Porism 3.1.2 to cancel any change to where the leaf enters and exits the edges of the cylinder. ■

3.5 Closed leaf gives an overtwisted disk

The following lemma suggests that whenever an embedded sphere (or disk) has a closed leaf in its characteristic foliation, the ambient contact structure is necessarily overtwisted. Indeed, if you have a closed leaf that bounds a disk on any surface, the result follows as well. Here, we will see an application of the pivot lemma for both introduction and elimination of singularities.

Lemma 3.5.1. *Let (M, ξ) be a co-orientable contact 3-manifold and let S be an embedding of S^2 in M. If the characteristic foliation S_ξ admits a closed leaf γ, then (M, ξ) is overtwisted.*

Proof. Let (M, ξ) be a co-orientable contact 3-manifold and let S be a sphere in M. Generically (by Peixoto's theorem, Theorem 2.2.12), the characteristic foliation S_ξ will have finitely many singularities and cycles, each of which is generic, and also no saddle connections.

Suppose that γ is a closed leaf in the characteristic foliation S_ξ. Without loss of generality we can suppose that no closed leaves exist inside the region bounded by γ and that it is absorbing. If there is a hyperbolic point in the region bounded by γ, then its stable separatrices must come from positive elliptic points (since there are no saddle connections), and if there is a negative elliptic point, only finitely many leaves which enter it come from hyperbolic points, and so the remaining leaves must originate from positive elliptic points. Thus[34], there must be a positive elliptic point p within the region bounded by γ.

If p is alone in the region bounded by γ, then we have our overtwisted disk, so we now assume this is not the case. We will show that we can deform S to find a leaf with strictly fewer singularities than in the region bounded by γ, and

[34] You could appeal also to the Poincaré-Hopf index theorem to get a similar conclusion.

thus can eventually create an overtwisted disk. There are two possibilities for singularities connected to p, and we treat these cases separately now.

Suppose first that p is connected to a negative elliptic point q by a leaf ℓ. Denote the unstable manifold of p as $\mathrm{Un}(p)$ and the finite set of singularities in the region bounded by γ as Ω. Note that if $\overline{\mathrm{Un}(p)} \cap \Omega = \{p, q\}$, then we get a contradiction as their flows would form a sphere. Thus there is another singularity r which we assume to be the singularity in $\overline{\mathrm{Un}(p)}$ different from q who receives a leaf from p nearest radially (on at least one side) to ℓ. Because r must also be attached to q, r is a hyperbolic point:

- If r is a positive hyperbolic point, then we may use Giroux's elimination lemma to remove both p and r.
- If r is a negative hyperbolic point, then we may use Giroux's elimination lemma to remove both q and r.

And hence the number of singularities in Ω is strictly smaller.

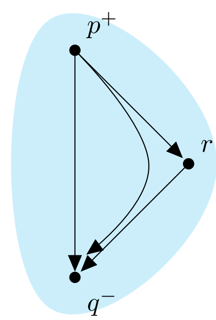

So suppose that p is connected to a hyperbolic point q. If q is positive, use Giroux's elimination lemma to remove both singularities, so we can assume q is negative. Follow its unstable separatrices:

- There are no saddle connections, and so neither separatrix connects to a hyperbolic point.
- If a separatrix connects to a negative elliptic point, we may use Giroux's elimination lemma to remove both negative singularities.

Thus we may assume that both of the unstable separatrices of q go towards the boundary, γ.

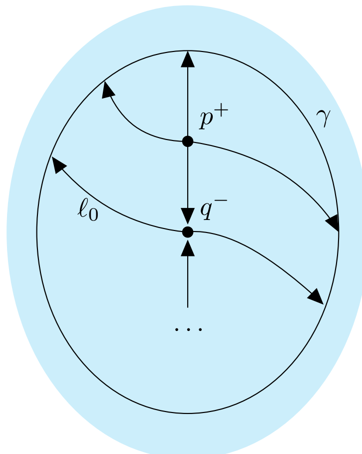

We may now use the pivot lemma in conjunction with Giroux's elimination lemma to introduce two new singularities:

- A negative elliptic point r into which an unstable separatrix ℓ_0 of q flows.
- A negative hyperbolic point located on the boundary of γ, not in the basin of p, whose one stable separatrix flows from the aforementioned elliptic point.

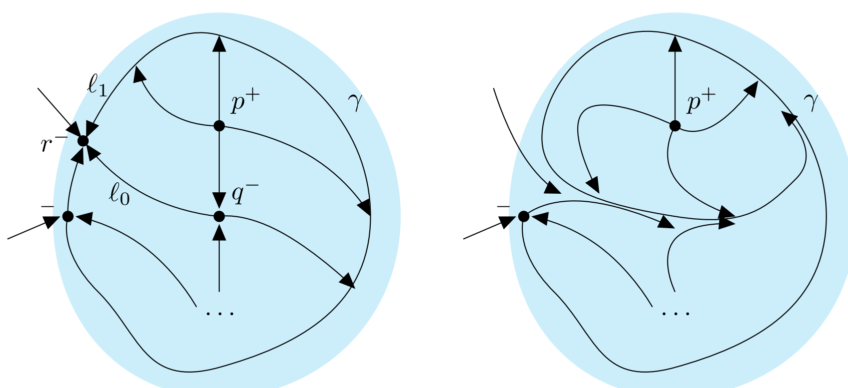

We apply the pivot lemma (Lemma 3.4.1) so that the segment ℓ_1 of γ flowing into r, contained in the closure of the unstable manifold of p, smoothly connects to the separatrix ℓ_0 of q. We then apply Giroux's elimination lemma on q and r which creates a smaller closed leaf γ' around p which bounds a disk containing one less singularity than the disk bounded by γ (as it does not contain q).

■

3.6 Generic movies

We begin by discussing in brief detail a result of Jorge Sotomayor[19] on generic 1-parameter families of vector fields. Sotomayor's result applied more generally to C^4-vector fields on 2-manifolds, but for our purpose, we simplify and assume everything to be C^∞.

Let M be a 2-manifold and $\mathrm{vect}(M)$ its algebra of vector fields. Denote by $\Sigma \subseteq \mathrm{vect}(M)$ the so-called structurally stable, or generic vector fields which, by Peixoto's theorem (Theorem 2.2.12), is exactly the set of vector fields v such that:

(i) v has all its zeros and periodic orbits generic.

(ii) v does not have saddle connections.

(iii) the α and ω-limit sets of each orbit of v are zeros of v or periodic orbits.

Sotomayor defines [K-S] $\subseteq \mathrm{vect}(M)$ as the set of vector fields satisfying only (i) and (ii), the so-called Kupka-Smale vector fields on M.

Sotomayor then spends some time building up the definition of a collection $\Sigma_1 \subseteq \mathrm{vect}(M)$ consisting of vector fields satisfying non-generic properties. In particular, Σ_1 is defined to be the set of vector fields v whose α and ω-limit sets of any orbit are zeros or periodic orbits, and who lies in exactly one of the following subclasses:

(Q_1) v has generic periodic orbits, and has no saddle connections, having a unique, non-generic zero which is either a saddle-node (in Sotomayor, $v \in Q_1^1$) or a "composed" focus (in Sotomayor, $v \in Q_1^2$).

(Q_2) v has only generic zeros and no saddle connections, having a single "quasi-generic" periodic orbit.

(Q_3) v has only generic periodic orbits and generic zeros, having a single saddle connection which—if it is a loop—happens to be a simple loop.

The set $\Sigma_1 = Q_1 \cup Q_2 \cup Q_3$ constitutes a particular set of vector fields with non-generic properties. As it turns out, Σ_1 constitutes a Banach submanifold of $\mathrm{vect}(M)$. Sotomayor's main result (mind you, somewhat rephrased for the purpose of our more specific case) is the following:

Theorem 3.6.1 (J. Sotomayor [19][Theorem 2])**.** *Let Φ be the set of C^1-mappings $F\colon [0,1] \to \mathrm{vect}(M)$ and let Γ be the set of $F \in \Phi$ such that:*

(i) $\mathrm{im}\, F \subset [K\text{-}S] \cup \Sigma_1$.

(ii) F intersects Σ_1 transversally.

(iii) The set of ordinary[35] *values of F is open and dense in $[0,1]$, coinciding with $F^{-1}(\Sigma)$.*

[35]If $F\colon [0,1] \to \mathrm{vect}(M)$ is a 1-parameter family of vector fields, then a value $c \in [0,1]$ is an **ordinary value** of F if there is a neighbourhood J of c such that for each $t \in N$, $F(t)$ is topologically equivalent to $F(c)$.

Then Γ *is dense in* Φ.

Because Γ is dense in the set Φ of 1-parameter families of vector fields, we refer to elements of Γ as "**generic**" 1-parameter families of vector fields.

Consider a closed, orientable surface S. Suppose we endow the "cylinder" $S \times [0,1]$ with a contact structure. For each $t \in [0,1]$, the slice $S_t := S \times \{t\}$ is an embedded surface admitting a characteristic foliation $(S_t)_\xi$. Thus associated to $S \times [0,1]$ is a 1-parameter family of characteristic foliations,

$$\{(S_t)_\xi \mid t \in [0,1]\},$$

which we call a **movie**. In celebration of this motif, for a fixed t we call $(S_t)_\xi$ a **still** of the movie. The notion of movie is due to Giroux [12], a good English reference for the theory of movies being the eponymous chapter in Vogel [20].

Because locally a characteristic foliation is merely the span of a vector field, Theorem 3.6.1 allows us to conclude that the stills of a generic movie of an orientable, closed surface S satisfy the following properties:

(i) They have finitely many singularities.

(ii) If a still has either a saddle node or a composed focus, it is the unique, non-generic singularity.

(iii) If a still has a quasi-generic periodic orbit, then it is the only one, all singularities being generic and no saddle connections being present.

(iv) Any given still can have at most one saddle connection.

So now specialising to the case that our surfaces S_t are spheres, we get the following result.

Lemma 3.6.2. *Let* (M, ξ) *be a tight contact 3-manifold and* S_0 *and* S_1 *be* ξ*-convex embeddings of* S^2*, smoothly isotopic to each other via an isotopy* $\{S_t \mid t \in [0,1]\}$ *of embeddings. There exists an arbitrarily small deformation of* $\{S_t \mid t \in [0,1]\}$ *fixing end points such that for each* $t \in [0,1]$*, the new still* $(S_t)_\xi$ *satisfies the Poincaré-Bendixson property*[36]*, and for finitely many* $t \in [0,1]$*,* S_t *satisfies exactly one of the following:*

- *The characteristic foliation of* S_t *contains a saddle-node or a composed focus.*
- *The characteristic foliation of* S_t *contains a retrograde saddle connection.*

Proof. Pulling back the characteristic foliations on S_t onto S_0, we get a 1-parameter family of vector fields on S_0. By Theorem 3.6.1, we can (after a perturbation) assume that for countably many times, this family has non-generic phenomenon as listed in Theorem 3.6.1, and in particular we may

[36] Indeed, it is enough to note that a singular foliation on the sphere with isolated singularities must satisfy the Poincaré-Bendixson property.

refine this to *finitely* many times by noting that by the transversality condition of Sotomayor's theorem, each time the family $F\colon [0,1] \to \mathrm{vect}(M)$ (viewed as a smooth path in $\mathrm{vect}(M)$) crosses Σ_1 at time t_0, there exists an open neighbourhood $(t_0 - \varepsilon, t_0 + \varepsilon) \subseteq [0,1]$ of t_0 where F only crosses Σ_1 once. By compactness of $[0,1]$, we get only a finite number of crossings, and thus finitely many t are non-generic.

We may bring this result forward to an isotopy of embeddings of S^2 using the fact that the contact condition is open and applying Gray's theorem. To simplify notation, we also refer to this isotopy of embeddings as $\{S_t \mid t \in [0,1]\}$, keeping in mind that the characteristic foliations are generic for all but finitely many times t now.

Using the characterisation of convexity according to Theorem 2.3.9, it is clear that for the at most finitely many time t that the new surface S_t is not convex, either it has a closed leaf, or it has a retrograde saddle connection. Appealing to the assumption that (M, ξ) is tight and Lemma 3.5.1, there are no periodic orbits. Thus there are at most finitely many t with at most one saddle connection, and a subset of which are retrograde, and then at most finitely many t (*different* than those with saddle connections) which have a saddle-node or composed focus.

■

3.7 Colin's theorem

Consider a closed, oriented surface S and the 3-manifold $S \times [0,1]$. As with the tubular neighbourhoods of convex surfaces, we will use the notation $S_t := S \times \{t\}$ for $t \in [0,1]$.

Using a method similar to the local reconstruction lemma (Lemma 2.3.4) if we endow $S \times [0,1]$ with the two contact structures ξ and ξ' which trace the same characteristic foliation on S_t for each $t \in [0,1]$, then ξ and ξ' are isotopic relative to the boundary (see [12][2.1]). Indeed, the contact structures would be—without loss of generality—defined by the contact forms

$$\alpha = \beta_t + u_t\, dt, \qquad \alpha' = \beta_t + u'_t\, dt,$$

where the same β_t is responsible in both for tracing the same characteristic foliations on each S_t. Since β_t is fixed—the contact condition already noted to be convex in the choice of u_t and u'_t—linear interpolation gives a family of contact forms on which we can apply Gray's theorem (Theorem 2.2.8).

The next lemma shows that we can weaken the condition that the characteristic foliation of each still S_t matches across the contact structures to only being a condition on the boundary, given that the S_t are ξ and ξ'-convex.

Lemma 3.7.1 (**Uniqueness lemma**, [9][4.9.2], [20][3.3])**.** *Let S be a closed, oriented surface, and suppose ξ and ξ' are contact structures on $S \times [0,1]$ satisfying:*

(i) The foliation $(S_0)_\xi$ coincides with $(S_0)_{\xi'}$, and $(S_1)_\xi$ with $(S_1)_{\xi'}$.

(ii) For each $t \in [0,1]$, the surface S_t is both ξ and ξ'-convex.

(iii) There is a smooth family $\{\Gamma_t \mid t \in [0,1]\}$ of dividing curves such that Γ_t divides both $(S_t)_\xi$ and $(S_t)_{\xi'}$.

Then ξ and ξ' are isotopic relative to boundary.

The proof of Lemma 3.7.1 uses a similar, but more technically difficult arrangement as the proof of Lemma 2.3.4. Similarly, it concludes with an application of Gray's theorem.

Now if we specify our surface to be the sphere, $S = S^2$, and we consider only tight contact structures, the obstructions to convexity greatly simplify and we (after playing connect-the-dots) get the following corollary:

Corollary 3.7.2. *Let $S := S^2$ and suppose ξ and ξ' are two tight contact structures on $S \times [0,1]$ such that the characteristic foliation of the boundaries S_0 and S_1 respectively coincide. Then ξ and ξ' are isotopic relative to boundary.*

Proof. Appealing to the Moser-Weinstein theorem (Theorem 2.2.13) we may assume that there is an isotopy fixing the endpoints S_0 and S_1 such that the contact structures coincide near the boundary. We may then use the genericity of convex surfaces (Theorem 2.3.7) to find a sufficiently close convex surfaces (in the neighbourhoods where the contact structures coincide) to S_0 and S_1. Thus without loss of generality, we may assume that S_0 and S_1 are convex embeddings.

By Lemma 3.6.2, we have that S_0 and S_1 are isotopic by an isotopy $\{S_t \mid t \in [0,1]\}$ such that for at most finitely many $t \in [0,1]$ the surface S_t is not convex, in particular such an S_t must have a retrograde saddle connection. Using Lemma 3.2.1, we may break the retrograde saddle connection while not producing any more saddle connections.

Now because for each $t \in [0,1]$ we have that the sphere S_t is convex, we can find a dividing curve Γ_t dividing $(S_t)_\xi$. Moreover, by Theorem 2.3.10, Γ_t is an embedded circle. However, the family $\{\Gamma_t \mid t \in [0,1]\}$ of circles is not necessarily smooth with regards to the index t. However, fixing a $t \in [0,1]$, any two dividing curves Γ_t and Γ_t' are isotopic via flowing along the characteristic foliation $(S_t)_\xi$ (which happens to be transverse to each dividing curve). Using this, we can ensure that our family $\{\Gamma_t \mid t \in [0,1]\}$ is indeed smooth in t, and that for the finitely many t where non-generic singularities (see Lemma 3.6.1) may present themselves, Γ_t is bounded sufficiently far away.

Thus, with the set-up as above, the result follows from Lemma 3.7.1. ∎

This puts us up in a very good place to classify tight contact structures on the 3-manifold $S^2 \times [0,1]$. Indeed, given any two contact structures η and η' on $S^2 \times [0,1]$ which give the same characteristic foliation at the endpoints, then they are immediately contact isotopic according to Corollary 3.7.2.

Observation 25. A movie $(S^2 \times [0,1], \eta)$ is determined up to contact isotopy relative to boundary by the characteristic foliation on end points.

However, given two spheres, it is not immediately obvious that they bound a movie. This takes some manipulation, and additionally is an important glimpse into the proof of Colin's theorem. We follow the proof contained in [9][4.12.3].

Lemma 3.7.3 (Colin, [4, Lemma 5])**.** *Let (M, ξ) be a tight contact 3-manifold. Given embeddings $e, f\colon S^2 \to M$, there is a tight contact structure η on $S^2 \times [0,1]$ such that the characteristic foliation on the end points coincide respectively with $(S^2)_{e^*\xi}$ and $(S^2)_{f^*\xi}$, the contact structure η being unique up to isotopy relative to boundary.*

Proof. First note that by Corollary 3.7.2, if η exists, then it is automatically unique up to isotopy relative to boundary.

The first case to consider is when $e(S^2)$ and $f(S^2)$ are ξ-convex (this case will help us prove the non-convex case). The dividing sets Γ_e and Γ_f of the embeddings each consist of a single circle, as per Giroux's Criterion (Theorem 2.3.10). Consider the model embedding of the sphere into $(\mathbb{R}^3, \ker(dz - y\,dx))$ as given in Example 2.3.1 whose dividing curve is the equator (as calculated in Example 2.3.2). First, find orientation preserving diffeomorphisms $\phi_e\colon e(S^2) \to S^2 \subseteq \mathbb{R}^3$ and $\phi_f\colon f(S^2) \to S^2 \subset \mathbb{R}^3$ which take Γ_e and Γ_f to the equator respectively. Using Giroux's flexibility lemma (Lemma 2.3.6) and the flow of the contact vector field

$$v = x\frac{\partial}{\partial x} + y\frac{\partial}{\partial y} + 2z\frac{\partial}{\partial z},$$

we can flow the embeddings $\phi_e \circ e$ and $\phi_f \circ f$ into embeddings $h_0, h_1\colon S^2 \to \mathbb{R}^3$ which are still transverse to v, and such that they extend[37] to an embedding

$$h\colon S^2 \times [0,1] \to \mathbb{R}^3, \qquad h_0 = h(\cdot, 0), \quad h_1 = h(\cdot, 1),$$

preserving the characteristic foliations at the end points. The desired contact structure η is given by:

$$\eta := \ker(h^*(dz - y\,dx)).$$

If the embeddings $e(S^2)$ and $f(S^2)$ are not ξ-convex, then the genericity of convex surfaces (Theorem 2.3.7) allows us to extend e and f to embeddings:

$$\tilde{e}\colon S^2 \times [-\varepsilon, \varepsilon] \to M, \qquad \tilde{e}\colon S^2 \times [1-\varepsilon, 1+\varepsilon] \to M,$$

such that the boundary embeddings are convex, $\varepsilon > 0$ small. Applying the above case where the embeddings are convex (in conjunction with Corollary 3.7.2 to conclude uniqueness as in the statement of the theorem), there exists copies A and A' in $(\mathbb{R}^3, \ker(dz - y\,dx))$ contactomorphic to $(\tilde{e}(S^2 \times [-\varepsilon, \varepsilon]), \xi)$ and $(\tilde{f}(S^2 \times [1-\varepsilon, 1+\varepsilon]), \xi)$ respectively. We may assume that A and A' (viewed as sort of thickened "shells") are disjoint and that A is contained in the bounded component of $\mathbb{R}^3 \setminus A'$.

[37] Because the orientation-preserving diffeomorphism group is connected.

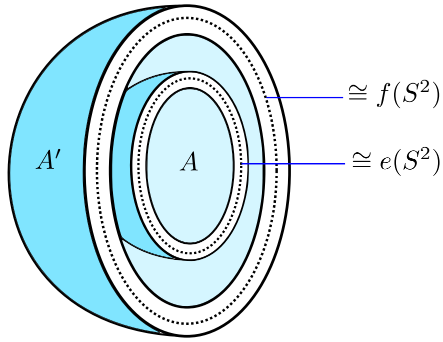

Both A and A' then contain an embedding of S^2 with the desired characteristic foliation, and moreover they bound the desired embedding of $S^2 \times [0, 1] \subset \mathbb{R}^3$ (the space between the shells being filled using the convex case again). ■

This concludes our tourist's account of movies. As mentioned, the above lemma is but a small step towards an important result by Vincent Colin:

Theorem 3.7.4 (Colin [4][Theorem 2])**.** *If we are given two smoothly isotopic embeddings of S^2 into a tight contact 3-manifold (M, ξ) which induce the same characteristic foliation at the end points, then there exists an ambient contact isotopy inducing one from the other.*

As we will see in the next section, this rather powerful theorem can be used as a shortcut in a proof of the classification theorem for topologically trivial Legendrian unknots.

4 Trivial Legendrian knots

In this section we begin by discussing the standardization of the characteristic foliation of Seifert disks (Seifert surfaces which are diffeomorphic to disks) for topologically trivial Legendrian knots, applying some of the tools we have developed so far. This is an important step in the classification of topologically trivial Legendrian knots, as proved in [6], and we outline the proof in the remaining sections, highlighting an application of Colin's theorem 3.7.4 to give a slightly alternative proof. A final simplification can be made so as to only use intermediate results discussed before Colin's proposition.

4.1 Standardizing Seifert disks

Suppose that $L \subset M$ is a Legendrian unknot embedded in a tight, co-orientable contact 3-manifold (M, ξ). Moreover, suppose L has a Seifert disk $D \subset M$, $\partial D = L$. Generically, we may assume that there are finitely many singularities along $L = \partial D$. Twisting D along L, we can assure that the singularities on L are of alternating sign.

Using a similar style to the introduction of singularities outlined before the statement of Giroux's elimination lemma (Theorem 2.2.15) we may introduce and remove singularities in order to "convert" between hyperbolic and elliptic points.

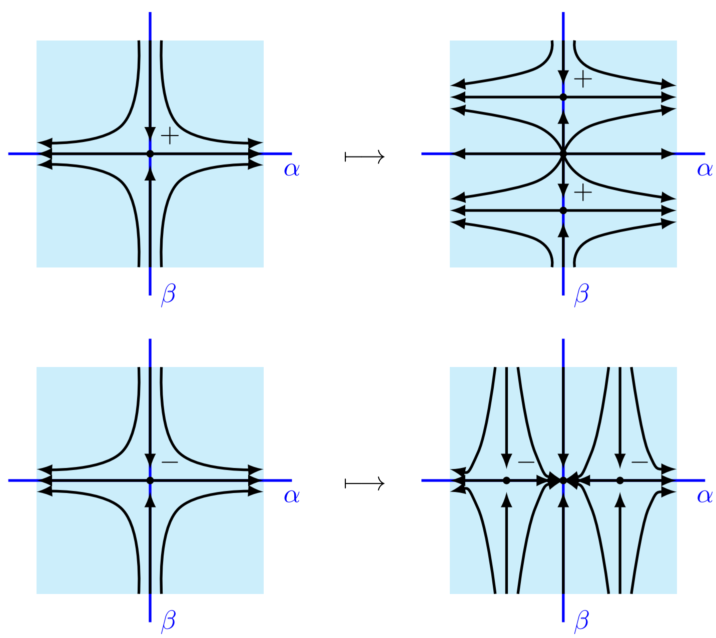

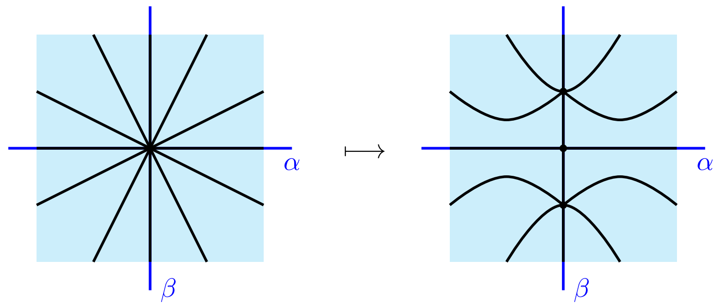

Lemma 4.1.1 (Eliashberg & Fraser[6][2.2])**.** *A hyperbolic singularity may be perturbed into an elliptic singularity and vice versa. Moreover, we may preserve the separatrices/choose which leaves become separatrices.*

Thus the positive elliptic points on L may be perturbed to be positive hyperbolic points (with regards to the above diagrams, L coinciding with α), and similarly the negative hyperbolic points to negative elliptic points (L coinciding with β), and thus the boundary $L = \partial D$ is absorbing. Between adjacent singularities of opposite signs, the contact frame rotates by π, and because we assume the singularities of L alternate, we can guarantee (by twisting D around L) that the contact frame rotates always in the same direction around L. Hence with respect to the Seifert frame, the contact frame makes $\frac{1}{2}$#singularities full flips. Hence, using the definition of $\operatorname{tb}(L)$ we have that the number of singularities along L is precisely:

$$\#\text{singularities} = 2|\operatorname{tb}(L)|.$$

Observation 26. Without loss of generality, we may assume that L has $2|\operatorname{tb}(L)|$ singularities, consisting of only negative elliptic points and positive hyperbolic points, and hence is absorbing.

Just as we have simplified the characteristic foliation of D on its boundary, $L = \partial D$, we may simplify D_ξ on the interior. First, assume that after repeated application of Lemma 3.2.1, there are no saddle connections.

If there is a positive hyperbolic point p in the interior of D, then the stable separatrices of p cannot come from a hyperbolic point, nor from a negative elliptic point. Thus it comes from a positive elliptic point q, and we can apply Giroux's elimination lemma to remove the pair p and q. Repeating this, we conclude that we can remove all positive hyperbolic points from the interior of D.

On the other hand, if there is a negative elliptic point p in the interior of D, then we can consider the closure $\overline{\operatorname{St}(p)}$ of the stable manifold of p, the boundary of which must contain a singularity other than an elliptic point (otherwise the flows of the elliptic points form hemispheres of a sphere, as in the proof of Lemma 3.5.1):

- If all singularities on the boundary of $\overline{\mathrm{St}(p)}$ are positive, then they must alternate positive elliptic and positive hyperbolic. Using Giroux's elimination lemma, we may then create an overtwisted disk with p at the centre, contradicting tightness.
- So there must be a negative singularity on the boundary of $\overline{\mathrm{St}(p)}$, it being necessarily a negative hyperbolic point q. We may then apply Giroux's elimination lemma to eliminate the pair p and q.

Like before, repeating this process concludes with removing all negative elliptic points from the interior of D.

Observation 27. Without loss of generality, we may assume that the interior of D does not contain any negative elliptic or positive hyperbolic points.

Because we may assume that L is absorbing and that there are no saddle connections, the stable separatrices of any negative hyperbolic point originate from positive elliptic points, and the unstable separatrices feed into $\partial D = L$.

Define a graph $\mathcal{G}$ whose vertices consist of positive elliptic points (all of which reside in the interior of D). We draw an edge between two vertices if the corresponding elliptic points are attached by the stable separatrices of a shared negative hyperbolic point.

Within the interior of D, the elliptic points (i.e. vertices of $\mathcal{G}$) must be separated by hyperbolic points, and it follows that $\mathcal{G}$ is connected. Moreover, the unstable separatrices of the interior hyperbolic points must feed into the negative elliptic points on ∂D, and so each edge implicitly gives a chord of ∂D which may not be crossed by any edge of $\mathcal{G}$. This means that $\mathcal{G}$ is necessarily acyclic, and thus $\mathcal{G}$ is a tree, which we will call the **interior tree** of L.

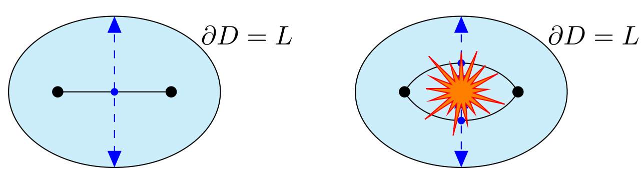

Observation 28. The graph formed by taking positive elliptic points as vertices and the stable separatrix pairs of hyperbolic points as edges is a tree on the interior of D.

Lemma 4.1.2 (Eliashberg & Fraser [6][2.6])**.** *Suppose that $D \subset M$ is a Seifert surface for a Legendrian knot L embedded in a tight contact 3-manifold (M, ξ). If $e_\pm$ and $h_\pm$ denote the number of elliptic and hyperbolic singularities of their respective signs in the interior of D, then:*

$$e_\pm - h_\pm = \frac{1}{2}(1 \mp \mathrm{tb}(L) \pm \mathrm{r}(L)).$$

In particular, we have simplified D_ξ so that $e_- = h_+ = 0$, and so in particular we have that:

$$e_+ = \frac{1}{2}(\mathrm{r}(L) - \mathrm{tb}(L) + 1), \qquad h_- = \frac{1}{2}(\mathrm{r}(L) - \mathrm{tb}(L) - 1).$$

Note that as expected from a tree, we have that the number of vertices is one more than the number of edges[38].

Observation 29. The values of $\mathrm{tb}(L)$ and $\mathrm{r}(L)$ determines the number of vertices of the tree in this "standardized disk foliation". In particular, the number of vertices is:

$$\frac{1}{2}(\mathrm{r}(L) - \mathrm{tb}(L) + 1).$$

Two standardized disk foliations from knots with the same values for $(\mathrm{tb}, \mathrm{r})$ may still differ up to diffeomorphism in two ways:

- The graph isomorphism class of interior trees may differ.
- The interior trees may be "tethered" to the ∂D via the stable separatrices of positive hyperbolic points on the boundary in different ways.

We will call these stable separatrices of positive hyperbolic points on ∂D suspending the interior tree inside L, **tethers**.

Because there are $2|\mathrm{tb}(L)|$ many singularities of alternating sign (positive hyperbolic, negative elliptic) on the boundary and no saddle connections, then the tree has a total of $|\mathrm{tb}(L)|$ many tethers originating from positive hyperbolic points, however it is not clear where exactly the tethers are attached to the interior tree.

One might ask whether each vertex p of the interior tree is attached to the boundary by a tether. Suppose that p is not attached to the boundary. The adjacent edges to p are negative hyperbolic points whose unstable separatrices connect to the negative elliptic points on L. Each boundary elliptic point will receive two separatrices, and this forms a closed Legendrian curve around p (smooth only after using the pivot lemma), all of whose singularities are negative. As in the proof of the Lemma 3.5.1, we can use in conjunction the pivot lemma and Giroux's elimination lemma to create an overtwisted disk, thus contradicting tightness.

Observation 30. Due to tightness of M, each vertex of the interior tree is tethered to the boundary $L = \partial D$ by at least one tether.

[38] Along with connectedness, a graph with $n+1$ vertices and n edges must be a tree. This gives an alternative explanation for Observation 28.

So far, we have given the Seifert disk D of L some kind of normal form satisfying the following properties:

(i) The characteristic foliation D_ξ has no saddle connections or closed leaves, and has finitely many singularities.

(ii) The boundary $\partial D = L$ contains $2|\operatorname{tb}(L)|$ many alternating singularities, consisting of only positive hyperbolic and negative elliptic points.

(iii) The interior of D contains only negative hyperbolic and positive elliptic points.

(iv) The interior tree of D has $\frac{1}{2}(\mathrm{r}(L) - \operatorname{tb}(L) + 1)$ many vertices.

Such a Seifert disk will be said to be **standardized**.

4.2 Model wavefronts

Common in the history of classical knot theory is the tabulation of knots based on projections. In this section we will consider the tabulation of Legendrian knots according to wavefront projections and properties of their associated spanning disks.

In particular, it turns out that there is a **wavefront catalogue** consisting of wavefronts, called **model wavefronts**, such that any of the standardized disk foliations obtained as outlined in Section 4.1 can be realised up to diffeomorphism on a disk spanning the lift of a model wavefront. Eliashberg and Fraser [6] showed the following:

Lemma 4.2.1 (Eliashberg & Fraser [6][3.1])**.** *Any two wavefronts W and W' in the wavefront catalogue sharing common values for* $(\operatorname{tb}, \mathrm{r})$ *are transversally homotopic, possibly after an addition of "swallow-tails". These moves lift to a Legendrian isotopy (see Corollary 2.4.3).*

In particular, we have an **abridged wavefront catalogue** consisting of representatives of the equivalence classes for the wavefronts in the wavefront catalogue under the equivalence relation of transverse homotopy, as given below.

tb$(L) = -(2n+1+s)$
r$(L) = -s$

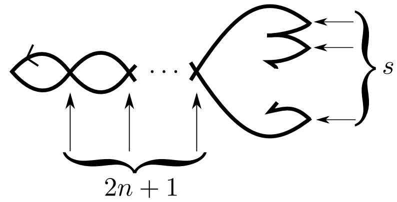

tb$(L) = -(2n+1+s)$
r$(L) = s$

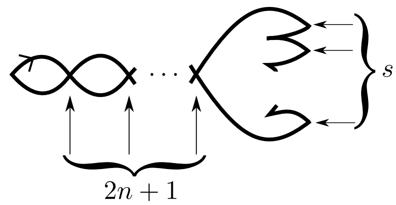

tb$(L) = -(2n+1)$
r$(L) = 0$

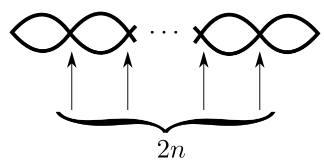

With help of the Thurston-Bennequin inequality (Theorem 2.4.9), it can be shown that this abridged wavefront catalogue is indeed exhaustive.

Example 4.2.2. Given $(\mathrm{tb}(L), \mathrm{r}(L)) = (-5, -2)$, one can construct a standardized disk for the corresponding abridged model wavefront L as below:

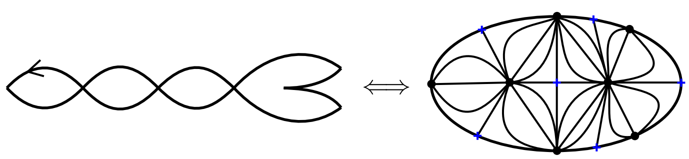

4.3 The proof

Eliashberg and Fraser proved the following classification theorem for topologically trivial Legendrian knots:

Theorem 4.3.1 (**Classification of topologically trivial Legendrian knots**, Eliashberg & Fraser [6])**.** *Let L and L' be two Legendrian unknots bounding disks within a tight contact 3-manifold (M, ξ). If*

$$\mathrm{tb}(L) = \mathrm{tb}(L'), \qquad \mathrm{r}(L) = \mathrm{r}(L'),$$

then the knots L and L' are Legendrian isotopic.

The steps to the proof are as follows:

1. Suppose we have two Legendrian knots L and L' with Seifert disks D and D' respectively, sharing the values for $(\mathrm{tb}, \mathrm{r})$. We may use the results of Section 4.1 to standardize both disks to $\tilde{D}$ and $\tilde{D}'$ respectively.

2. From our catalogue of wavefronts, there are model wavefronts W and W' whose lifts L_W and $L_{W'}$ have Seifert disks D_W and $D_{W'}$, having diffeomorphic foliations to $\tilde{D}$ and $\tilde{D}'$, respectively.

3. Eliashberg and Fraser use an argument to contact isotope L to a neighbourhood of a tree on the foliation, further eventually being able to obtain a contact isotopy between L and the Legendrian boundary L_W of the corresponding model disk D_W. Similarly for L'.

4. Finally, the composition of the aforementioned contact isotopies and the contact isotopy arising from the transverse homotopy of wavefronts of Lemma 4.2.1 (as the model wavefronts have common values $(\mathrm{tb}, \mathrm{r})$) gives us our isotopy.

Colin's theorem, however, presents a zippy replacement (admittedly, somewhat a luxury in terms of fuel economy) for (3). To apply Colin, we need to first find some spheres!

Consider the standardized Seifert disk $\tilde{D}$ for L. The Seifert disk $\tilde{D}$ is convex, and so it admits an $\mathbb{R}$-invariant contact structure for some coordinate system. We obtain a copy of $\tilde{D}$ above and another below, both with identical characteristic foliation. A vertical "belt" at L consists of many copies of L with $2|\,\mathrm{tb}(L)|$ singularities. By rounding corners, one can reduce this to a single equator L with only positive hyperbolic and negative elliptic points, alternating in sign (the same as in the standardization). In this way, the foliation of the top hemisphere consists of the "negative" of the lower hemisphere singularities.

This procedure can also be repeated on D_W, and because the foliations on $\tilde{D}$ and D_W coincided, the foliations on the spheres match.

Observation 31. Applying Colin's theorem (Theorem 3.7.4), there exists a contact isotopy between the spheres, and moreover this must take L to the lift L_W of W because the curve L is uniquely identifiable in the foliation of its sphere[a], and likewise for the lift L_W of W and its sphere.

[a]Uniquely identifiable in the sense that it is the unique closed curve containing negative elliptic and positive hyperbolic points.

Repeating this argument for $\tilde{D}'$ and $D_{W'}$ yields the contact isotopy between L' and the lift of W'. Thus this constitutes an alternative step, completing the proof of Theorem 4.3.1!

The frugal mathematician might wince at the cost of gas for the argument above, but our travels through convex surface theory have not gone completely to waste. Let us call the sphere obtained as above from L and $\tilde{D}$, $S_{\tilde{D}}$, and the analogous sphere from D_W, S_{D_W}. Note that the foliations coincide, as before.

Because M is path connected, there exists an isotopy $\{S_t \mid t \in [0,1]\}$ of the spheres such that $S_0 = S_{\tilde{D}}$ and $S_1 = S_{D_W}$, and a path of diffeomorphisms $\{\psi_t \mid t \in [0,1]\}$ such that:

(i) $\psi_0 = \mathrm{id}$.

(ii) For each $t \in [0,1]$ we have $\psi_t(S_0) = S_t$.

(iii) $(S_1)_\xi = (\psi_1)_*((S_0)_\xi)$.

The idea will be to repeat the argument in the proof of Lemma 3.6.2 to get a generic movie on S_0. We push the resulting generic movie back to an *almost* convex isotopy of surfaces, to which we can begin to apply the argumentation found in the proof of Corollary 3.7.2 to arrange a smooth family of dividing circles.

To this end, consider the foliations drawn on S_0 by:

$$\xi_t := \psi_t^* \xi.$$

By perturbing the surface S_0 within (M, ξ_t), Sotomayor (Theorem 3.6.1) gives a generic movie on S_0, and pushing the movie back to a family of surfaces $\{S_t \mid t \in [0,1]\}$ (we use the same notation for the family as before) by composing the perturbations with ψ_t. Moreover, the respective foliations of the S_t are likewise perturbed, and thus we have an isotopy described as in the statement of Lemma 3.6.2.

We now find ourselves in a similar place as in the proof of Corollary 3.7.2. We first break any retrograde saddle connections while not producing any additional saddle connections. For each $t \in [0,1]$ our sphere S_t is convex, admitting a dividing curve Γ_t (each Γ_t manifests itself as a circle, as per Giroux's criterion, Theorem 2.3.10), and we can arrange for the family $\{\Gamma_t \mid t \in [0,1]\}$ to vary smoothly in t, using the same technique as in Corollary 3.7.2. Let φ_t be the path of diffeomorphisms obtained by composing ψ_t with the flow used to find the smooth family $\{\Gamma_t \mid t \in [0,1]\}$, satisfying $\Gamma_t = \varphi_t(\Gamma_0)$.

To find the desired contact isotopy, apply Giroux's flexibility lemma (Lemma 2.3.6): for each $t \in [0,1]$ we can perturb S_t in an arbitrarily small neighbourhood so that $(S_t)_\xi$ matches the characteristic foliations on S_0 and S_1 (recall: they are diffeomorphic). More formally, the diffeomorphisms φ_t induce a family of contact structures, $\{\eta_t := \varphi_t^* \xi \mid t \in [0,1]\}$, and a movie $\{(S_0)_{\eta_t} \mid t \in [0,1]\}$ on which we can apply the flexibility lemma.

Thus on each S_t, there is a copy L_t of L, varying smoothly in t. This is our Legendrian isotopy taking L to L_W!

Observation 32. A contact isotopy between L and L_W may be found by replicating part of the proof of Corollary 3.7.2 (using the spheres $S_{\tilde{D}}$ and S_{D_W} as end points) to get a convex isotopy of surfaces $\{S_t \mid t \in [0,1]\}$ with a smoothly varying family of dividing circles. Applying Giroux's flexibility lemma (Lemma 2.3.6) yields an identical characteristic foliation on S_t for each time $t \in [0,1]$. Thus we get a smooth family $\{L_t \mid t \in [0,1]\}$ of Legendrian curves with $L_0 = L$ and $L_1 = L_W$.

Postlude

Our tourist's account has come to an end, and we thank the reader for tagging along with this excursion into the exciting and deep theory of characteristic foliations on convex surfaces!

A Miscellany

A.1 General

Proposition A.1.1. *If $v, w\colon M \to TM$ and $\tilde{v}, \tilde{w}\colon N \to TN$ are vector fields on M and N respectively that satisfy for a smooth $\varphi\colon M \to N$:*

$$\varphi_{*,p}(v_p) = \tilde{v}_{\varphi(p)}, \quad \textit{and} \quad \varphi_{*,p}(w_p) = \tilde{w}_{\varphi(p)},$$

for a $p \in M$, then:

$$\varphi_{*,p}([v,w]_p) = [\tilde{v}, \tilde{w}]_{\varphi(p)}.$$

Proof. Let $f \in C^\infty(M)$ and note:

$$\begin{aligned}
\varphi_{*,p}([v,w]_p)(f) &= [v,w]_p(f \circ \varphi) \\
&= v_p(w(f \circ \varphi)) - w_p(v(f \circ \varphi)) \\
&= v_p((\varphi_* \circ w)f) - w_p((\varphi_* \circ w)f) \\
&= v_p(\tilde{w}(f) \circ \varphi) - w_p(\tilde{v}(f) \circ \varphi) \\
&= \varphi_*(v_p)(\tilde{w}(f)) - \varphi_*(w_p)(\tilde{v}(f)) \\
&= \tilde{v}_{\varphi(p)}(\tilde{w}(f)) - \tilde{w}_{\varphi(p)}(\tilde{v}(f)) \\
&= [\tilde{v}, \tilde{w}]_{\varphi(p)}(f).
\end{aligned}$$

■

Proposition A.1.2. *Let α be a 1-form and x, y vector fields. Then:*

$$d\alpha(x,y) = x(\alpha(y)) - y(\alpha(x)) - \alpha([x,y]).$$

Proof. We prove the formula on the 1-form $\alpha = f\,dg$ for $f, g \in C^\infty(M)$. The left-hand side reduces as:

$$\begin{aligned}
d\alpha(x,y) &= d(f\,dg)(x,y) \\
&= (df \wedge dg)(x,y) \\
&= df(x)dg(y) - dg(x)df(y) \\
&= x(f)\,y(g) - x(g)\,y(f),
\end{aligned}$$

while the right reduces as:

$$\begin{aligned}
&x(\alpha(y)) - y(\alpha(x)) - \alpha([x,y]) \\
&= x(f\,dg(y)) - y(f\,dg(x)) - f\,dg([x,y]) \\
&= x(f\,y(g)) - y(f\,x(g)) - f\,[x,y](g) \\
&= x(u)\,y(g) + f\,x(y(g)) - y(f)\,x(g) - f\,y(x(g)) - f\,x(y(g)) + f\,y(x(g)) \\
&= x(u)\,y(g) - y(f)\,x(g),
\end{aligned}$$

and hence we have equality. ■

Proposition A.1.3. *Suppose that V is an n-dimensional vector space and $\alpha \in V^*$, $\omega \in \Lambda^k V^*$, for $k \leqq n$. Then:*

$$\alpha \wedge \omega = 0 \Longleftrightarrow \omega|_{\ker \alpha} = 0.$$

Proof. We prove the $\Rightarrow$ direction first. Suppose that $\alpha \wedge \omega = 0$. Extend α to a basis $\{\alpha_1 = \alpha, \alpha_2, \ldots, \alpha_n\}$ of V^*. Note that $\omega = \sum_I \omega_I \alpha_I$, $|I| = k$, and so:

$$0 = \alpha \wedge \omega = \sum_I \omega_I \alpha \wedge \alpha_I = \sum_{I \not\ni 1} \omega_I \alpha \wedge \alpha_I,$$

where $\{\alpha \wedge \alpha_J \mid 1 \notin J\}$ forms a linearly independent set. Hence $\omega_J = 0$ for each $J \not\ni 1$. Hence:

$$\omega = \alpha \wedge \beta, \qquad \beta := \sum_{|K|=k-1, K \not\ni 1} \omega_K \alpha_K.$$

Thus we see that for all $v_1, v_2, \ldots, v_k \in \ker \alpha$:

$$\omega(v_1, v_2, \ldots, v_k) = \sum_{\sigma \in S_k} \operatorname{sgn}(\sigma) \alpha(x_{\sigma(x_1)}) \beta(x_{\sigma(2)}, \ldots, x_{\sigma(k)}) = 0.$$

Now we turn to the $\Leftarrow$ direction. Suppose that $\omega|_{\ker \alpha} = 0$, then because $V = \ker \alpha \oplus \operatorname{im} \alpha$ and $\dim \operatorname{im} \alpha = 1$, we see that:

$$\alpha \wedge \omega(v_1, v_2, \ldots, v_{k+1}) = \sum_{\sigma \in S_{k+1}} \operatorname{sgn}(\sigma) \alpha(x_{\sigma(1)}) \omega(x_{\sigma(2)}, \ldots, x_{\sigma(k+1)}) = 0.$$

Hence $\alpha \wedge \omega = 0$ and ω decomposes like before as $\omega = \alpha \wedge \beta$. Hence:

$$\alpha \wedge \omega = \alpha \wedge \alpha \wedge \beta = 0.$$

■

A.2 Symplectic geometry

A **symplectic manifold** is a manifold M with a closed, non-degenerate, differential 2-form ω called a **symplectic form**. That is, a section ω of $\bigwedge^2 T^*M$ such that $d\omega = 0$, and for each $p \in M$ and non-zero $v_p \in T_pM$, there exists a $w_p \in T_pM$ such that $\omega_p(v_p, w_p) \neq 0$. This implies that ω_p is a non-degenerate, skew-symmetric *bilinear* form on the tangent space T_pM.

Proposition A.2.1. *Let V be a finite dimensional vector space endowed with a non-degenerate, skew-symmetric bilinear form Ω. Then there exists a basis $\{e_i, f_i \mid 1 \leqq i \leqq n\}$ of V such that:*

$$\Omega(e_i, e_j) = \Omega(f_i, f_j) = 0, \qquad \Omega(e_i, f_j) = \delta_{i,j}.$$

Proof. Pick a non-zero vector in V and call it e_1. Note that Ω is non-degenerate, and so we may find an $f_1 \in V$ such that $\Omega(e_1, f_1) = 1$. Define $V_1 := \operatorname{span}\{e_1, f_1\}$ and define a respective set:

$$V_1^\Omega := \{v \in V \mid \forall w \in V_1 \,.\, \Omega(w, v) = 0\}.$$

It is easy to show that $V = V_1 \oplus V_1^{\Omega}$.

Repeat this process, picking a non-zero $e_2 \in V_1^{\Omega}$, and an f_2 such that $\Omega(e_2, f_2) = 1$, where we then define V_2 and so forth.

This process reduces the dimension by two each iteration and eventually halts as V is finite dimensional. Moreover, V is even-dimensional because otherwise we would conclude this process with a 1-dimensional subspace spanned by a vector u such that $\Omega(u, \cdot) \equiv 0$. ■

Because $\dim M = \dim T_pM$ for each $p \in M$, we automatically get the following corollary:

Corollary A.2.2. *A symplectic manifold (M, ω) is necessarily even-dimensional.*

We call a basis on T_pM as in the statement of Proposition A.2.1 a **symplectic basis**.

Example A.2.3 (Standard symplectic structure on $\mathbb{R}^{2n}$)**.** The standard symplectic structure on $\mathbb{R}^{2n}$ in the coordinates $\{x_i, y_i \mid 1 \leqq i \leqq n\}$ given by the symplectic form:

$$\omega_0 := \sum_{j=1}^{n} dx_i \wedge dx_j.$$

It is easily checked that $\{(\partial/\partial x_i)(p), (\partial/\partial y_i)(p) \mid 1 \leqq i \leqq n\}$ is a symplectic basis for the tangent space $T_p\mathbb{R}^{2n}$.

A **symplectomorphism** between symplectic manifolds (M, ω) and (N, ω') is a diffeomorphism $\phi \colon M \to N$ such that:

$$\phi^* \omega' = \omega.$$

As with the Darboux-Pfaff theorem (see Theorem 2.2.7) in contact geometry, symplectic geometry has its own theorem which asserts the futility of local invariants:

Theorem A.2.4 (Darboux's theorem)**.** *Let (M, ω) be a symplectic $2n$-manifold. Then for each $p \in M$ there exists a coordinate chart (U, x_i, y_i) about p such that in coordinates:*

$$\omega = \sum_{j=1}^{n} dx_j \wedge dy_j.$$

We will not include a proof for the symplectic Darboux's theorem.

Proposition A.2.5. *Let M be a $2n$-manifold and ω a closed 2-form on M. Then:*

$$\omega \text{ is a symplectic form} \iff \omega^n \text{ is a volume form.}$$

Proof. The $\Rightarrow$ direction follows from a proof by induction on the statement that in a symplectic basis[39], p implicit:

$$\omega^k = k! \sum_{1 \leqq i_1 < \ldots < i_k \leqq n} e_{i_1}^* \wedge f_{i_1}^* \wedge \ldots \wedge e_{i_k}^* \wedge f_{i_k}^*.$$

[39] In what follows, v^* denotes the vector in the cotangent space dual to v.

Because $\omega^n(e_1, f_1, \ldots, e_n, f_n) = n!$, clearly ω^n is a never vanishing form of top degree—a volume form! The $\Leftarrow$ direction is clear. ■

Given a symplectic manifold (M, ω), if $U \leq T_pM$ is a linear subspace of the tangent space at $p \in M$, then define:

$$U^\omega := \{v_p \in T_pM \mid \forall u_p \in U \,.\, \omega_p(u_p, v_p) = 0\}.$$

It is relatively straightforward linear algebra to show the following properties:

- $\dim U + \dim^\omega U = \dim T_pM$.
- $(U^\omega)^\omega = U$.
- $U \subseteq W \implies W^\omega \subset U^\omega$.

With this notation, we define the following special submanifolds of M:

1. An **isotropic**[40] submanifold $S \subset M$ is one such that $\omega|_S \equiv 0$, or equivalently, $T_pS \subseteq (T_pS)^\omega$.
2. A **coisotropic** submanifold $S \subset M$ is one such that $(T_pS)^\omega \subseteq T_pS$.
3. A **Lagrangian** submanifold $L \subset M$ is one which is both isotropic and coisotropic, or equivalently, $T_pL = (T_pL)^\omega$.

Proposition A.2.6. *If S is an isotropic submanifold of (M, ω), then:*

$$\dim S \leqq \frac{1}{2} \dim M.$$

Proof. Let $p \in S$. Consider the linear subspace $T := T_pS \subseteq T_pM$. Because S is isotropic, we have $T \subseteq T^\omega$ (so $\dim T \leqq \dim T^\omega$), and so by one of the properties of the ω-complement,

$$\dim M = \dim T + \dim T^\omega \geqq 2 \dim T.$$

Hence the proposition follows. ■

[40] "Isotropic" means that S "looks the same" to the symplectic form, in this case.

References

[1] Beat Aebischer, M. Borer, Markus Kälin, Christoph Leuenberger, and Hans Martin Bach. *Symplectic geometry: an introduction based on the seminar in Bern, 1992*, volume 124. Birkhäuser, 2013.

[2] Daniele Celoria. Relations among topological and contact knot invariants. 2012.

[3] Austin Christian. A short gallery of characteristic foliations, 2017. URL: http://www.math.ucla.edu/~archristian/notes/gallery.pdf. Last visited on 2019/06/24.

[4] Vincent Colin. Chirurgies d'indice un et isotopies de sphères dans les variétés de contact tendues. *Comptes Rendus de l'Académie des Sciences-Series I-Mathematics*, 324(6):659–663, 1997.

[5] Ana Cannas Da Silva and A Cannas Da Salva. *Lectures on symplectic geometry*, volume 3575. Springer, 2001.

[6] Yakov Eliashberg, Maia Fraser, et al. Topologically trivial legendrian knots. *Journal of Symplectic Geometry*, 7(2):77–127, 2009.

[7] John B. Etnyre. Convex surfaces in contact geoometry: Class notes. URL: https://people.math.gatech.edu/~etnyre/preprints/papers/surfaces.pdf. Last visited on 2019/06/24.

[8] John B. Etnyre, Ko Honda, et al. Knots and contact geometry i: torus knots and the figure eight knot. *Journal of Symplectic Geometry*, 1(1):63–120, 2001.

[9] Hansjörg Geiges. *An introduction to contact topology*, volume 109. Cambridge University Press, 2008.

[10] Paolo Ghiggini. Introduction to legendrian knots and legendrian contact homology. 2012.

[11] Emmanuel Giroux. Convexité en topologie de contact. *Commentarii Mathematici Helvetici*, 66(1):637–677, 1991.

[12] Emmanuel Giroux. Structures de contact en dimension trois et bifurcations des feuilletages de surfaces. *Inventiones mathematicae*, 141(3):615–689, 2000.

[13] Morris W. Hirsch, Robert L. Devaney, and Stephen Smale. *Differential equations, dynamical systems, and linear algebra*, volume 60. Academic press, 1974.

[14] Norbert Hungerbühler, Thomas Mettler, and Micha Wasem. Convex integration and legendrian approximation of curves. *arXiv preprint arXiv:1507.07661*, 2015.

[15] John M. Lee. *Introduction to smooth manifolds*. Springer, 2001.

[16] Patrick Massot. Topological methods in 3-dimensional contact geometry. In *Contact and symplectic topology*, pages 27–83. Springer, 2014.

[17] Dusa McDuff and Dietmar Salamon. *Introduction to symplectic topology*. Oxford University Press, 2017.

[18] Mauricio Matos Peixoto. Structural stability on two-dimensional manifolds. *Topology*, 1(2):101–120, 1962.

[19] Jorge Sotomayor. Generic one-parameter families of vector fields on two-dimensional manifolds. *Publications Mathématiques de l'IHÉS*, 43:5–46, 1974.

[20] Thomas Vogel. On the uniqueness of the contact structure approximating a foliation. *Geometry & Topology*, 20(5):2439–2573, 2016.

[21] Frank W. Warner. *Foundations of differentiable manifolds and Lie groups*, volume 94. Springer Science & Business Media, 2013.

[22] Chris Wendl. Signs (or how to annoy a symplectic topologist). https://symplecticfieldtheorist.wordpress.com/2015/08/23/signs-or-how-to-annoy-a-symplectic-topologist/. Accessed: 2019-09-03.