Notes on Giroux's 1991 paper, "Convexité en topologie de contact"

Daniel Mathews

March 22, 2007

Contents

1 Introduction **2**

2 So what is a convex surface? **2**

3 Why? **2**

4 What a convex surface looks like I **3**

4.1 First example: a contactization 3

4.2 Second (not just an) example: vertically invariant contact structures 4

4.2.1 The form of the form 4

4.2.2 It's mostly contactizations 5

4.2.3 Where do the contact planes go? 6

4.2.4 Where does the characteristic foliation go? 6

4.3 Characterising vertically invariant contact structures 8

5 The power of characteristic foliations **10**

5.1 What singularities can occur in characteristic foliations? 10

5.2 Contact structures on $S \times \mathbb{R}$ and germs 11

5.3 A characterisation of isochore singularities 13

5.4 Which foliations are characteristic foliations? 13

5.5 The characteristic foliation determines the germ of the contact structure... up to isotopy... 14

5.6 ... and up to isomorphism 15

6 Singularities and their orientations **16**

7 What a convex surface looks like II **17**

7.1 It has expanding vector fields directing the characteristic foliation, and dividing sets... 18

7.2 ...if it's closed, it can't have a characteristic foliation defined by a closed form... 18

7.3 ...or a closed leaf with return map tangent to the identity... . . . 19
7.4 ... the dividing set avoids certain separatrices... 19
7.5 ...it's sufficient that the characteristic foliation is almost Morse-Smale... 20

8 What Giroux says about eliminating singularities 22

9 The dividing set is all powerful 23

1 Introduction

This paper of Giroux is absolutely seminal in the study of contact geometry. Convexity at first glance may not seem so crucial or natural, but it is! Giroux showed why, and along the way proved a number of basic results that are standard today.

Much of Giroux's paper deals with higher-dimensional contact manifolds. I will stick only to 3-manifolds. Also, I am more interested in convex *surfaces* than the other notions of convexity discussed by Giroux: namely, those relating to convex contact structures.

Let (M, ξ) be a closed oriented contact 3-manifold throughout. Let α be a contact form for ξ.

2 So what is a convex surface?

Definition 2.1 *A* convex surface *S in (M, ξ) is an embedded surface (possibly with boundary) for which there exists a transverse contact vector field X, i.e. a vector field transverse to S for which the flow of X preserves ξ.*

So, the crucial thing appears to be a contact vector field. Are there many contact vector fields on ξ? It turns out there are very many — as many as there are sections of the line bundle TM/ξ on M. Giroux refers to Arnold for this, but it's not too difficult to prove. So, it seems that it should be fairly generic that this should be transverse to S, though maybe you might think there are topological obstructions. It turns out that there are no obstructions, and a generic surface *is* convex. This is special to 3 dimensions.

A generic surface is convex. This is one of Giroux's greatest achievements. Maybe it wasn't that hard to prove once he set his mind to it, but discovering this was a great breakthrough.

3 Why?

Indeed. Why convex surfaces.

By this one might mean: what is the motivation for this definition? Well, it comes from pseudoconvex embeddings in complex geometry: these often give you contact structures. They are all about, relevant to our context, Morse functions and gradient vector fields keeping things invariant. “Convex” in the complex context says something about a Morse function “pointing out” of the manifold, for instance as a gradient vector field for a Morse function points out of a level set. So this is where the invariant vector field and the transversality come from.

But by the question one might mean: Why do we bother with these things?

And the most basic layperson’s reason could be: “Because humans are so bad at visualising 3 dimensions!”

It’s true. A field of planes moving about in 3-dimensional space is potentially something you can visualise. It’s not an unacceptably high number of dimensions. Planes are not that complicated. If we could see them better, we could probably prove a lot more directly.

Why are they useful? We will soon see that the contact structure near a surface is determined by its characteristic foliation; and a foliation on a surface is much easier to keep track of than a whole contact structure. Some lines on a surface are much easier to visualise than planes which rotate all over the place! At the price of only seeing the contact structure in the neighbourhood of a surface, we gain the pleasure of not having to tax our 3-dimensional visualisation abilities. This was known before Giroux’s paper, it seems; but Giroux’s paper is the earliest location I know of where there is a written proof of this result.

So: the contact structure near a surface is determined by its characteristic foliation. But this has nothing to do with convexity, yet. However a generic surface is convex. The crucial blow is struck by the following result: the characteristic foliation on a convex surface is more or less determined by its *dividing set*. We will define what a dividing set is in due course: it is a certain finite set of curves on a surface. What the “more or less” means will become clear shortly. But note the obvious point: as easy as a characteristic foliation is to understand, a finite set of curves is much simpler again.

For anyone who has ever struggled to draw a picture even of the standard contact structure and figure out which way all those planes were wiggling, convex surfaces, then, offer a simplification of contact structures that is truly awesome.

4 What a convex surface looks like I

Before we go proving all these major results, we would like to get a nice picture of a convex surface S in (M, ξ). Note that, having a transverse contact vector field, S is automatically oriented.

4.1 First example: a contactization

An example: a *contactization* of a symplectic manifold. Yes, symplectic and contact are *that* closely related — you can symplectize a contact manifold, and contactize a symplectic manifold. In both cases, you “-ize” it by crossing it with $\mathbb{R}$, and defining an appropriate contact/symplectic form on the product.

Actually, you can't contactize *any* symplectic manifold. It has to be *exact.* Well, how else were you going to get a contact 1-form canonically out of a symplectic 2-form? So take an exact symplectic manifold (W, ω) where $\omega = d\beta$. Let W have dimension $2n$ Then on $W \times \mathbb{R}$ we want to write a contact 1-form. What about β? The simplest possibility, perhaps, but that's not going to work; for β doesn't change in the $\mathbb{R}$ direction. So, you might try the next simplest possibility, say $\beta + dt$, where t is the coordinate in the $\mathbb{R}$ direction. Is that a contact form? Yes:

$$(\beta + dt) \wedge (d(\beta + dt))^n = \beta \wedge (d\beta)^n + dt \wedge (d\beta)^n = dt \wedge (d\beta)^n = dt \wedge \omega^n \neq 0.$$

Why does $\beta \wedge (d\beta)^n = 0$? Because it's a $(2n+1)$-form but is only nontrivial in the W direction, which is $2n$-dimensional. And why is $dt \wedge \omega^n \neq 0$? From the definition of a symplectic form, ω^n is a volume form on W; and dt is a volume form on $\mathbb{R}$; so their product is a volume form on $W \times \mathbb{R}$.

Consider now $M = W \times \mathbb{R}$ as a 3-manifold (so W is a 2-manifold); however the discussion all still works in higher dimensions, it we define everything properly.

Our contact form is very simple: $\beta + dt$. In particular it is invariant under translations in the $\mathbb{R}$ direction. That is, it is invariant under the flow of the vector field $X = \partial/\partial t$. So $X = \partial/\partial t$ is a contact vector field. Any horizontal surface $W \times \{t\}$ is therefore a convex surface. In fact, any slice of $W \times \mathbb{R}$ which is "never vertical" is going to give us a convex surface. More precisely, take any function $f : W \longrightarrow \mathbb{R}$ and consider its graph, which is a surface in $W \times \mathbb{R}$. Being the graph of a function is what me mean by "never vertical"; this will also be a convex surface.

4.2 Second (not just an) example: vertically invariant contact structures

But in fact *every* convex structure looks a bit (not exactly!) like this, and this is the key to understanding what's going on. Any convex surface S certainly has a neighbourhood diffeomorphic to $S \times \mathbb{R}$; and we can choose the $\mathbb{R}$ coordinate so that our transverse contact vector field is $X = \partial/\partial t$; then in these coordinates, the contact structure is invariant under the flow of X, i.e. invariant under translations in the $\mathbb{R}$ direction. So if we can understand the situation of $S \times \mathbb{R}$ with $X = \partial/\partial t$ as a contact vector field, we will understand what any convex surface looks like.

4.2.1 The form of the form

So, we have $S \times \mathbb{R}$, and want to consider contact structures which are invariant under translations in the $\mathbb{R}$ direction. What does the 1-form α look like? If we write x, y for some local coordinates on S and t for the coordinate on $\mathbb{R}$, we have

$$\alpha = \alpha_x \, dx + \alpha_y \, dy + \alpha_t \, dt$$

where $\alpha_x, \alpha_y, \alpha_t$ are functions $S \times \mathbb{R} \longrightarrow \mathbb{R}$. However if the contact structure is to be invariant under translations in the t direction, we must have $L_X \alpha$ being a multiple of α; actually, by rescaling α, we can choose α to satisfy $L_X \alpha = \alpha$. That is, α is to be invariant under translations in the $\mathbb{R}$ direction. So $\alpha_x, \alpha_y, \alpha_t$ are all functions on S alone. From now on we write β for the local expression $\alpha_x \, dx + \alpha_y \, dy$ and have a 1-form on S; we write u for α_t and u is a function on S. So we have

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

This is the general form of a vertically invariant 1-form on $S \times \mathbb{R}$.

But it's not the general form of a vertically invariant *contact* 1-form on $S \times \mathbb{R}$. For that we need to do a computation:

$$\begin{aligned} \alpha \wedge d\alpha &= (\beta + u \, dt) \wedge (d\beta + du \wedge dt) \\ &= \beta \wedge d\beta + \beta \wedge du \wedge dt + u \, dt \wedge d\beta \\ &= \beta \wedge du \wedge dt + u \, dt \wedge d\beta \\ &= (\beta \wedge du + u \, d\beta) \wedge dt \end{aligned}$$

Again $\beta \wedge d\beta = 0$ since this is a 3-form on S. So the contact-ness of α depends on the form

$$\theta = \beta \wedge du + u \, d\beta.$$

This is a 2-form on S: if it is nondegenerate everywhere, then wedging it with dt will give us an everywhere nondegenerate 3-form on $S \times \mathbb{R}$; if it is degenerate anywhere, then wedging it with dt will give us a degenerate 3-form on $S \times \mathbb{R}$. So α is contact if and only if θ is nondegenerate, i.e. an area form on S.

Clearly for any such α, provided θ is nondegenerate, any surface $S \times \{t\}$ is a convex surface in $S \times \mathbb{R}$ with contact vector field $X = \partial/\partial t$. In fact, again, for any function $S \longrightarrow \mathbb{R}$, the graph of the function inside $S \times \mathbb{R}$ is a convex surface in $(S \times \mathbb{R}, \alpha)$.

4.2.2 It's mostly contactizations

Note that the expression for α, namely $\beta + u \, dt$, is actually quite close to the contact form on a contactization, as given above. In fact, we usually do have mostly a contactization! The u is the only difference. If we had $u \neq 0$, then we could divide through by u and obtain a different form — of the contactization type — giving the same contact structure. So: whenever u is nonzero, we really have a contactization. The region of S where $u = 0$ is generically a set of curves which we will call Γ; away from Γ, on $(S \backslash \Gamma) \times \mathbb{R}$, we can take the contact form as $\beta/u + dt$, so we have the contactization of the exact symplectic manifold $(S \backslash \Gamma, d(\beta/u))$.

A priori, of course, Γ could be a nasty set, but generically it will be a set of curves. In fact more is true: if α is a contact form, so that θ is nondegenerate, then Γ *must be* a set of curves. For whenever $u = 0$, to have $\theta \neq 0$ means $\beta \wedge du + u \, d\beta = \beta \wedge du \neq 0$. Hence $du \neq 0$, and so the tangent space to Γ is precisely the one-dimensional kernel of du; so Γ is a 1-manifold.

Actually, note that we now have *two* area/symplectic forms on S. From considering the contact-ness of $\beta + u\ dt$ we have the form $\theta = \beta \wedge du + u\ d\beta$. And from considering the exact symplectic manifold $(S \backslash \Gamma)$ 1-form we have $d(\beta/u)$. These are both area forms, so differ by a function which is nowhere zero. It turns out that function is u^2, which is certainly positive away from Γ. In particular, $\theta = u^2 d(\beta/u)$:

$$d\left(\frac{\beta}{u}\right) = \frac{d\beta}{u} + d\left(\frac{1}{u}\right) \wedge \beta = \frac{1}{u}\ d\beta - \frac{1}{u^2} du \wedge \beta = \frac{1}{u^2}\left(u\ d\beta + \beta \wedge du\right) = \frac{1}{u^2}\ \theta.$$

4.2.3 Where do the contact planes go?

Let's consider the surface $S \times \{0\}$ (or S times any number) in more detail. What does the contact structure look like here? The contact plane is the kernel of $\alpha = \beta + u\ dt$.

The plane is usually transverse to $S \times \{0\}$; the contact plane is only tangent at isolated points of $S \times \{0\}$. At these points, α is of the form $u\ dt$; these are the points where $\beta = 0$. (It's possible to have $\beta = 0$ but still have $\theta \neq 0$.)

The plane can sometimes be vertical. At such points, the contact plane contains $\partial/\partial t$, so $u = 0$, and α is simply of the form β; the points where ξ is vertical are the points where $u = 0$. If you like, the value of u tells us how non-vertical the plane is, well, not really, but it may be useful to think this way.

4.2.4 Where does the characteristic foliation go?

We will obtain a characteristic foliation on $S \times \{0\}$. It is the intersection of $\ker \alpha$ with the horizontal surface; but since the $u\ dt$ term has nothing to say about horizontal vectors, the characteristic foliation is completely determined by β. In fact, the characteristic foliation is simply given by $\ker \beta$.

We can even describe a vector Y directing the characteristic foliation on $S \times \{0\}$. Well, we could just take Y to be an arbitrary vector in $\ker \beta$, but that's not so canonical. Rather, we have a 1-form on S and an area/symplectic 2-form θ. Recall that in symplectic manifolds, 1-forms have dual vectors and vice versa; and on a 2-dimensional surface, the dual vector to a 1-form lies in its kernel. So, we can take Y to be the vector dual to the 1-form β with respect to the area/symplectic form θ on S. That is: we define the vector field Y on S by

$$\beta = i(Y)\ \theta.$$

Then Y directs the characteristic foliation on $S \times \{0\}$.

Actually, there was another potential, highly canonical, choice for a vector field directing the characteristic foliation; at least, away from Γ, that is, where $u = 0$. Because recall that $(S \backslash \Gamma) \times \mathbb{R}$ is the contactization of the exact symplectic surface $(S \backslash \Gamma, d(\beta/u))$. And the characteristic foliation can be taken as given not just by the kernel of β, but by the kernel of the highly canonical Liouville 1-form β/u. So by the same reasoning as the previous paragraph, we can define a vector field Z directing the characteristic foliation on $S \times \{0\}$ by being dual to the Liouville form β/u with respect to $d(\beta/u)$:

$$\frac{\beta}{u} = i(Z)\, d\left(\frac{\beta}{u}\right).$$

Are these closely related? One would hope so! Well, of course Y and Z both direct the same characteristic foliation, so they must be scalar multiples of each other. In fact, that multiple is just the function u: $Z = uY$. This follows from a computation: plugging both Y and Z into the same area form $d(\beta/u)$ we can see what the factor is.

$$\begin{aligned} i(Z)d\left(\frac{\beta}{u}\right) &= \frac{\beta}{u} \\ i(Y)d\left(\frac{\beta}{u}\right) &= \frac{1}{u^2} i(Y)\theta = \frac{1}{u^2}\beta = \frac{1}{u} \cdot \frac{\beta}{u}. \end{aligned}$$

So Y, Z point in the same direction along the characteristic foliation when $u > 0$; and in the opposite direction when $u < 0$.

Now, since Y is defined on the whole surface and is a smooth vector field, Y should proceed smoothly across the set Γ, where $u = 0$. From our proceeding comment, as we cross Γ, u changes sign. Thus Z abruptly changes direction every time you pass through Γ. Actually, not *so* abruptly, since $Z = uY$; so Z approaches 0 near Γ, and Z extends continuously to all of S, with singularities along Γ.

So, which way do these vector fields point as we cross Γ? Well, Z isn't defined on Γ, but Y is. Recall that being on Γ means $u = 0$. And we want to know the change in u in the direction of Y, that is, $du(Y) = i(Y)\, du$. Well, from the equation (the only one we really have!)

$$\beta = i(Y)\, \theta = i(Y) \left(\beta \wedge du + u d\beta\right),$$

noting that here $u = 0$ and $i(Y)\beta = 0$ (because Y directs the characteristic foliation, which is the kernel of β), we have

$$\beta = -\beta i(Y)\, du, \quad \text{hence} \quad i(Y)\, du = -1.$$

We conclude that Y points from positive towards negative u. Since Z agrees with Y when $u > 0$, Z points out of positive u pieces of S. And since Z disagrees with Y when $u < 0$, Z points out of negative u pieces of S also. So Z points out of each piece of $S \backslash \Gamma$, towards Γ.

Actually, we could have seen this more quickly and more slickly. Think about the effect of our two vector fields on their respective area forms. As you flow along them, do they dilate or compress area? Well, flowing along Y produces a messy result on θ. But because $S \backslash \Gamma$ is exact and has a Liouville form, the effect of flowing Z on the area form $d(\beta/u)$ is very simple:

$$L_Z\, d\left(\frac{\beta}{u}\right) = di(Z)d\left(\frac{\beta}{u}\right) = d\left(\frac{\beta}{u}\right).$$

Here all we used was $\beta/u = i(Z)d(\beta/u)$. So Z dilates the area form $d(\beta/u)$. If you expand an area form, you certainly can't be coming *inwards* through the entire boundary!

4.3 Characterising vertically invariant contact structures

Let's summarise what we find on $S \times \mathbb{R}$ with a vertically invariant contact structure.

(i) The contact form has the expression $\beta + u\,dt$ where β is a 1-form on S and u is a function on S.

(ii) The 2-form $\theta = \beta \wedge du + u\,d\beta$ must be an area/symplectic form on S.

(iii) The contact plane is vertical where $u = 0$, and horizontal where $\beta = 0$.

(iv) Away from $\Gamma = \{u = 0\}$, the manifold is actually a contactization. Precisely: $((S\backslash\Gamma) \times \mathbb{R}, \beta/u + dt)$ is the contactization of $(S\backslash\Gamma, d(\beta/u))$.

The characteristic foliation on $S \times \{0\}$ is given by any of:

(i) The kernel of β. (Or, away from Γ, it's the kernel of β/u.)

(ii) It's directed by Y, which is dual to β with respect to the area/symplectic form θ. Across Γ, Y points in the direction of decreasing u, i.e. from positive to negative u regions.

(iii) Away from Γ, it's directed by Z, which is dual to the Liouville form β/u with respect to the exact symplectic form $d(\beta/u)$. The flow of Z expands the area form $d(\beta/u)$; consequently, Z always points outwards across Γ, and changes direction as you cross Γ along a leaf. However since $Z \longrightarrow 0$ as we approach Γ, Z extends continuously to a singular vector field on all of S.

This is more than enough to characterise what's going on here. The question is: if you're given a (singular) foliation $\mathcal{F}$ on S, is it the characteristic foliation on $S \times \{0\}$ for some vertically invariant contact structure on $S \times \mathbb{R}$? Turns out that the set Γ, the dilating an area form property, and the change of direction along Γ, is enough. We can make this precise: this is a 3-dimensional version of Giroux's theorem I.3.2; Giroux's version is stronger and works in higher dimensions.

Proposition 4.1 *Let S be a closed surface and let $\mathcal{F}$ be a dimension-1 singular foliation on S. The following are equivalent:*

(i) There exists on $S \times \mathbb{R}$ a vertically invariant contact structure which induces $\mathcal{F}$ as the characteristic foliation on $S \times \{0\}$.

(ii) There exists on S a 1-manifold Γ (i.e. a finite set of disjoint closed curves) transverse to $\mathcal{F}$ (in particular, avoiding the singularities of $\mathcal{F}$ such that

(a) *the complement S' of an open tubular neighbourhood of Γ, with the fibres in $\mathcal{F}$, is an exact symplectic manifold for which a Liouville vector field directs $\mathcal{F}$ and exits transversely through the boundary; and*

(b) *the involution of the double cover $\partial S' \longrightarrow \Gamma$ obtained by following leaves of $\mathcal{F}$ reverses the orientation on the leaves. (That is, the Liouville vector field changes direction on leaves passing through Γ.)*

Proof That (i) implies (ii) we have already seen: the 1-manifold Γ where $u = 0$ divides S into pieces where we have Liouville vector fields; and the Liouville vector fields change direction on leaves passing through Γ.

In the other direction, a contactization $S \times \mathbb{R}$ of an exact symplectic manifold $(S, d\beta)$ has a vertically invariant contact structure with contact form $\beta + dt$ with the Liouville vector field Z (defined via $\beta = i(Z)d\beta$) directing the characteristic foliation, which is the kernel of β. So we have a vertically invariant contact structure $\beta + dt$ on $S' \times \mathbb{R}$. All we need to do is glue the pieces together, knowing that the Liouville vector field changes direction as we pass through the boundary Γ.

This is a silly little fiddle, but the idea is clear: you can glue it all together. You have to fiddle because the vector field goes to zero on Γ; you need to rescale somehow to get it to work.

So, we consider a tubular neighbourhood of Γ fibred by leaves of $\mathcal{F}$. Consider it as $\Gamma \times (-1-\epsilon, 1+\epsilon)$, where at the ± 1 points it joins with S'. We know that Z points outwards near $\partial S'$, so we can choose a coordinate s on $S' \cap (\Gamma \times (-1-\epsilon, 1+\epsilon)) = \Gamma \times (-1-\epsilon, -1] \cup \Gamma \times [1, 1+\epsilon)$ such that there $Z = -s(\partial/\partial s)$. We need to extend our contact form $\beta + dt$ over $\Gamma \times (-1-\epsilon, 1+\epsilon) \times \mathbb{R}$, remaining vertically invariant.

What does β look like on $\Gamma \times [1, 1+\epsilon)$? We should be able to work this out, since we have chosen coordinates in the region. We defined Z by $\beta = i(Z)d\beta$. And $i(Z)\beta = 0$, either from the previous equation or since Z directs $\mathcal{F}$, which is the kernel of β. So $L_Z\beta = di(Z)\beta + i(Z)d\beta = d0 + \beta = \beta$, and flowing along $Z = -s(\partial/\partial s)$ dilates β. This gives us a differential equation for what happens to β as you increase s. Note that flowing along Z, in our coordinates β always has $\partial/\partial s$ in its kernel, so $\beta(s) = f(s)\beta(1)$. Now $L_Z\beta = \beta$ becomes

$$-s\frac{df}{ds} = f \quad \text{which implies} \quad f(s) = \frac{f(1)}{s}.$$

So $\beta(s) = \beta(1)/s$. Write γ for $\beta(1)$, and we can consider γ as a 1-form on all of $\Gamma \times (-1-\epsilon, 1+\epsilon)$. For $s \in (-1-\epsilon, -1] \cup [1, 1+\epsilon)$ we have $\beta = \gamma/s$ and our original contact form was $\beta + dt = \gamma/s + dt$. But this gives the same contact structure as $\gamma + s\,dt$, and $\gamma + s\,dt$ gives us a vertically invariant contact structure on all of $\Gamma \times (-1-\epsilon, 1+\epsilon)$. We just need to check it's contact for $s \in (-1, 1)$, which follows upon inspecting the expression $(\gamma + s\,dt) \wedge d(\gamma + s\,dt)$ and noting it's contact for $s = \pm 1$.

A silly little fiddle indeed. ∎

5 The power of characteristic foliations

We now know that if we have a convex surface, we can choose coordinates so a neighbourhood looks like $S \times \mathbb{R}$, with the contact vector field $X = \partial/\partial t$, so that the contact form is $\beta + u\, dt$, with β a 1-form and u a function on S, and with the characteristic foliation $\mathcal{F}$ directed by a Liouville vector field away from $\Gamma = \{u = 0\}$, which changes direction over Γ.

On the other hand, suppose you have all this data on the surface S: a Γ, an $\mathcal{F}$, an X, Liouville forms, Liouville vector fields and all the rest of it. Does it follow that the surface is convex?

We don't know yet: to conclude that the surface is convex, we need to know that the contact structure near S looks like our standard picture. We *could* conclude that, if we knew that the characteristic foliation determines the contact structure nearby.

In other words, to go any further we need to harness the power of the characteristic foliation. We need the following proposition (Giroux's proposition II.1.2(b)).

Proposition 5.1 *Let $\mathcal{F}$ be a singular foliation on a closed surface S. Fix an orientation on $\wedge^2 TS$, and we are interested only in positive contact structures (i.e. with $\alpha \wedge d\alpha > 0$). Two germs of contact structures which induce the same characteristic foliation $\mathcal{F}$ on S are conjugate by a germ of a diffeomorphism which is isotopic to the identity through diffeomorphisms preserving $\mathcal{F}$. (Consequently, the two germs of contact structures are isomorphic!)*

So, we commence a thorough study of what characteristic foliations look like, from the ground up.

5.1 What singularities can occur in characteristic foliations?

A first question is: characteristic foliations can be singular. How bad can those singularities be? We have to worry about the types of singularities in the foliation $\mathcal{F}$.

Nobody else in the history of contact geometry (at least, the recent history) seems to have worried about this issue. Everyone else just says you can assume, generically, that there are only nice singularities. Not Giroux. Well, he is proving the result that everyone will use, so we all owe him a debt of gratitude.

To be sure, we are assuming *something*, namely that our foliation is a *singular foliation*. We are assuming that $\mathcal{F}$ is only degenerate at singularities which are isolated points. You can easily find characteristic surfaces where the singularities are worse (e.g. take a surface tangent to the contact structure along a whole curve). Beyond this, we are making no assumptions on the types of singularities.

A good way to study a singularity is by looking at the *linearisation* there. We have a singular foliation $\mathcal{F}$ on S and a vector field Y directing it. At a singularity x of $\mathcal{F}$ we have $Y = 0$. Now we have a family of diffeomorphisms φ_t, the flow of Y. Clearly $\varphi_t(x) = x$; x isn't going anywhere! However around

x, we can take note of what happens. For a tangent vector in T_xS, we can exponentiate it, see what happens to it under φ_t, and (for small enough t) return to the tangent space. As $t \to 0$ we obtain the linearisation of φ_t, which is a linear map $T_xS \longrightarrow T_xS$. The trace of this linear transformation tells us what φ_t is doing to areas as $t \to 0$; this trace is the divergence of Y at x.

Definition 5.2 *A singularity x of a vector field Y is* isochore *if the divergence of Y at x is zero.*

Perhaps there's an English translation for "isochore", I don't know. I'm just copying the word from the French.

Note this includes hyperbolic fixed points where the attracting and repelling eigenvalues are equally strong — so there is no effect on area. Note also that a degenerate singularity can be okay, and not isochore — a degenerate singularity has non-invertible linearisation, but the trace may still be nonzero.

It turns out that isochore singularities can't occur in characteristic foliations; but this is all. Happily, a singular foliation without isochore singularities can be realised as a characteristic foliation. Thank goodness! This may seem somewhat strange, since it implies that some fairly degenerate singularities can occur in characeristic foliations; but that's not a problem, as we'll see, we can construct the contact structure explicitly nearby anyway.

This is Giroux's proposition II.1.2(a)

Proposition 5.3 *Let $\mathcal{F}$ be a singular foliation on a surface S. We fix an orientation on the manifold $\wedge^2 TS$ and we are interested only in germs of contact structures along S which give this orientation.*

$\mathcal{F}$ is the characteristic foliation induced on S by a germ of contact structures if and only if $\mathcal{F}$ is without isochore singularities.

So this theorem answers two questions. It answers our original question: what singularities can occur in characteristic foliations? But it also answers a more general question: which singular foliations are characteristic foliations? This is a very important answer to know.

To prove this, we'll need to understand how contact structures work in neighbourhood $S \times \mathbb{R}$ — but, more generally than before, we no longer require the structure to be vertically invariant.

5.2 Contact structures on $S \times \mathbb{R}$ and germs

Let us consider what a general contact 1-form α on $S \times \mathbb{R}$ looks like. No matter what the t coordinate, at any point in $S \times \{t\}$, α can be written in the form $\beta_t + u_t \; dt$ where β_t is a 1-form and u_t a function on $S \times \{t\}$. When is this a contact form? We do the usual computation:

$$
\begin{aligned}
(\beta_t + u_t\, dt) \wedge d\,(\beta_t + u_t\, dt) &= (\beta_t + u_t\, dt) \wedge \left(d\beta_t + \frac{\partial \beta_t}{\partial t} \wedge dt + du_t \wedge dt \right) \\
&= \beta_t \wedge \frac{\partial \beta_t}{\partial t} \wedge dt + \beta_t \wedge du_t \wedge dt + u_t\, dt \wedge d\beta_t \\
&= dt \wedge \left(u_t\, d\beta_t + \beta_t \wedge \left(du_t - \frac{\partial \beta_t}{\partial t} \right) \right)
\end{aligned}
$$

Note that when we write $d\beta_t$ we mean taking the differential, of a form on S; so the differential as a form on $S \times \mathbb{R}$ is $d\beta_t + (\partial\beta_t/\partial t)dt$. So $\beta_t \wedge d\beta_t = 0$ as a 3-form on S. The 2-form in brackets is a 2-form on S; so we have a contact form if and only if this is nowhere degenerate. The condition is:

$$
u_t\, d\beta_t + \beta_t \wedge \left(du_t - \frac{\partial \beta_t}{\partial t} \right) \neq 0
$$

In our situation, we are given the foliation on $S \times \{0\}$, so we are given β_0. To define the contact structure nearby we need nearby β_t and u_t and $\partial\beta_t/\partial t$. However, to define the *germ* of a contact structure near $S \times \{0\}$ we need less: we only need, in addition to our β_0, the function u_0 and the partial derivative $(\partial\beta_t/\partial t)|_{t=0}$. These need to satisfy

$$
u_0\, d\beta_0 + \beta_0 \wedge \left(du_0 - \frac{\partial \beta_t}{\partial t}|_{t=0} \right) \neq 0
$$

on $S \times \{0\}$. A pair $(u_0, (\partial\beta_t/\partial t)|_{t=0})$ satisfying this condition is all we need.

One important observation is that he set of all pairs $(u_0, (\partial\beta_t/\partial t)|_{t=0})$ satisfying our condition is basically *convex*: no, not a convex surface in a contact manifold, it's a convex set in a vector space, in the old-fashioned sense. If we require our pair $(u_0, (\partial\beta_t/\partial t)|_{t=0})$ to satisfy the condition above, with a particular sign (so say > 0 rather than $\neq 0$), the space of all such pairs is convex.

For if we have two such pairs $(u_0, (\partial\beta_t/\partial t)|_{t=0})$ and $(u_0', (\partial\beta_t/\partial t)|_{t=0}')$, with

$$
\begin{aligned}
u_0\, d\beta_0 + \beta_0 \wedge \left(du_0 - \frac{\partial \beta_t}{\partial t}|_{t=0} \right) &> 0 \\
u_0'\, d\beta_0 + \beta_0 \wedge \left(du_0' - \frac{\partial \beta_t}{\partial t}|_{t=0}' \right) &> 0
\end{aligned}
$$

then

$$
[(1-s)u_0 + su_0']\ d\beta_0 + \beta_t \wedge \left(d\,[(1-s)u_0 + su_0'] - \left[(1-s)\frac{\partial \beta_t}{\partial t}|_{t=0} + s\frac{\partial \beta_t}{\partial t}|_{t=0}' \right] \right) > 0
$$

also.

End disgusting computation. A disgusting computation indeed, but this is the essential reason why a characteristic foliation determine the germ of contact structure. It means that the "germ data" $(u_0, (\partial\beta_t/\partial t)|_{t=0})$ forms a contractible space; and with a bit of Mosering, it means that all the relevant germs are isomorphic.

5.3 A characterisation of isochore singularities

So far we've defined an isochore singularity as one satisfying a condition about linearisations and divergences. There's a slightly more approachable way of describing them.

As usual, let our foliation $\mathcal{F}$ be defined by a 1-form β. But now we will take a dual. So let ω be an area/symplectic form on S and take the vector field X dual to β with respect to ω. So $\beta = i(X)\omega$ and X directs $\mathcal{F}$. The isochore condition can be written using the divergence with respect to the area form ω, i.e. $L_X\omega$. To be isochore is to have $L_X\omega = 0$ at a singularity. We can expand this and obtain

$$0 = L_X\omega = i(X)d\omega + di(X)\omega = 0 + d\beta.$$

So at a singularity, i.e. when $\beta = X = 0$, we have $d\beta = 0$. The above computation shows that the converse is true also.

Lemma 5.4 *A singularity x of a foliation $\mathcal{F}$ on a surface S defined by a 1-form β (so that $\beta = 0$ at x) is isochore if and only if $d\beta = 0$ there.*

5.4 Which foliations are characteristic foliations?

We can now answer the question: precisely those without isochore singularities. And we can prove it. We can prove proposition 5.3.

To see why, suppose we have a characteristic foliation $\mathcal{F}$ on S; we will show there are no isochore singularities. Let α be the contact form, and β the induced form on S. At a singularity x of $\mathcal{F}$, α is tangent to S, so for any vector V tangent to S, $\alpha(V) = 0$. Since $\alpha \wedge d\alpha = 0$, $d\alpha$ is nondegenerate on T_xS; hence so is its restriction $d\beta$. So $d\beta \neq 0$, and by the above characterisation, x is not isochore.

In the other direction, given a singular foliation $\mathcal{F}$ without isochore singularities, we want to construct a germ of a contact structure on $S \times \mathbb{R}$. As discussed in the previous section, giving the foliation is equivalent to giving a 1-form β_0 on $S \times \{0\}$, defined up to multiplication by nonzero function. To define the germ of the contact structure near $S \times \{0\}$ we need to add the information of u_0 and $(\partial\beta_t/\partial t)|_{t=0}$, satisfying

$$u_0 \, d\beta_0 + \beta_0 \wedge \left(du_0 - \frac{\partial \beta_t}{\partial t}|_{t=0} \right) \neq 0.$$

So, we just need to find such u_0 and $(\partial\beta_t/\partial t)|_{t=0}$. The isochore condition, we have discovered, means that when $\beta_0 = 0$, then $d\beta_0 \neq 0$.

The condition we must satisfy basically says that a certain 2-form on S must be an area form. So let's take an area form ω and compare everything to it; any 2-form will be some function times ω. So, in particular, $d\beta_0 = f\omega$ for some function f. Then the first term is $u_0 d\beta_0 = u_0 f\omega$. We want to choose u_0 to get things positive, say; so one sneaky thing to do would be to take $u_0 = f$. Then $u_0 d\beta_0 = u_0^2\omega \geq 0$. This is never negative; and in fact, from the isochore condition we know that $u_0 \neq 0$ at singularities, so $u_0^2 > 0$ there.

What about the other term? Well, note that whenever β_0 is nonzero, there is always something you can wedge it with to get ω. Let's call this 1-form γ, so $\beta \wedge \gamma = \omega$. This can't work at singularities; but there does exist γ such that $\beta \wedge \gamma = g\omega$ where $g \geq 0$. It would be very good to have the term in brackets being γ, namely $du_0 - (\partial\beta_t/\partial t)|_{t=0} = \gamma$. We can certainly do this: we have already chosen u_0, and we can define $(\partial\beta_t/\partial t)|_{t=0}$ by this equation. Then we have

$$u_0 \, d\beta_0 + \beta_0 \wedge \left(du_0 - \frac{\partial \beta_t}{\partial t}|_{t=0} \right) = u_0^2 \omega + \beta \wedge \gamma = (u_0^2 + g)\omega.$$

Now we know $u_0^2 + g > 0$. At the singularities of β, thanks to the non-isochore condition we know $u_0 > 0$; and away from the singularities of β from our definition of γ we have $g > 0$. So this is a positive form and we are done.

This proves the proposition. Well, you might want to worry a little about orientations, but any problem with orientations are solvable by taking a double (or quadruple!) cover.

5.5 The characteristic foliation determines the germ of the contact structure... up to isotopy...

Now we have seen that, from the characteristic foliation, you can construct the germ of a contact structure. A characteristic foliation, we know, has no isochore singularities; and from a foliation without isochore singularities we can determine the germ of a contact structure by choosing $(u_0, (d\beta_t/dt)|_{t=0})$ as we just did above.

What we want to know, is that the germ of the contact structure is uniquely determined. This will prove proposition 5.1.

Well, as we mentioned previously, given the foliation $\mathcal{F}$ and hence β_0, the space of all germ data $(u_0, (d\beta_t/dt)|_{t=0})$ for contact structures forms a contractible space, in fact, convex (in the vector space sense!).

That is, given any two germs of contact structures $((u_0^0, (d\beta_t/dt)|_{t=0}^0), (u_0^1, d\beta_t/dt)|_{t=0}^1)$ with the same $\mathcal{F}$, and hence the same β_0, we can linearly homotope their germ data one to the other:

$$\left(u_0^s, \frac{\partial \beta_t}{\partial t}|_{t=0}^s \right) = \left((1-s)u_0^0 + s u_0^1, (1-s)\frac{\partial \beta_t}{\partial t}|_{t=0}^0 + s \frac{\partial \beta_t}{\partial}|_{t=0}^1 \right).$$

These give a germ of a contact form for all $s \in [0, 1]$.

So, to the proof! Take two contact structures ξ^0, ξ^1 with forms α^0, α^1, inducing the same foliation $\mathcal{F}$ on S. We write $\alpha^0 = \beta_t^0 + u_t^0 \, dt$, $\alpha^1 = \beta_t^1 + u_t^1 \, dt$. The fact that they both give the characteristic foliation means that we can take $\beta_0^0 = \beta_0^1$. The two pairs of germ data are above. We let $\alpha^s = (1-s)\alpha^0 + s\alpha^1$. Then the germ data is also linearly interpolated, and is exactly as given above, and from the convexity argument the contact property is always satisfied; so α^s gives a (germ of a) contact structure near S.

This proves that any two contact structures inducing the characteristic foliation $\mathcal{F}$ on S have germs which are isotopic (through germs of contact structures). But we want more.

5.6 ... and up to isomorphism

Given our ξ^1, ξ^2 as above, we not only want their germs to be isotopic. We want them to be related by a "germ of a diffeomorphism which is isotopic to the identity through diffeomorphisms preserving $\mathcal{F}$". That is, we want a one-parameter family of diffeomorphisms — or rather, germs of such diffeomorphisms near S — starting at the identity, ending at a (germ of a) diffeomorphism which takes one contact structure to the other, and which always preserves the surface S and the foliation $\mathcal{F}$. That is, this one-parameter family of (germs of) diffeomorphisms must look like flows along leaves of $\mathcal{F}$. And indeed this is precisely what we do: we will flow along leaves of $\mathcal{F}$, realising the family of (germs of) contact forms α^s.

We are already very close. We have shown that simply linearly interpolating the (germs of) contact forms α^0, α^1 to obtain α^s, we have a family of (germs) of contact forms already. They just need to be realised. If we can find a 1-parameter family of diffeomorphisms φ_s near S, which fix S and flow along leaves of $\mathcal{F}$, and have $\varphi_s^* \alpha^s$ proportional to α^0, we are done. (They don't have to be equal, just proportional, to define the contact structure.)

We will use Moser's method. How does Moser's method work? Instead of looking at the diffeomorphisms, you look at the flow; you look at vectors; and using the properties of the contact form and whatever else, you work out what they must be. A good way of writing the condition we want is:

$$\varphi_s^* \alpha^s \wedge \frac{\partial}{\partial s} \left(\varphi_s^* \alpha^s \right) = 0.$$

This only refers to one particular point in time s, which is useful. If X^s is the vector field with flow φ_s, this translates to

$$L(X^s) \alpha^s = -\frac{\partial \alpha^s}{\partial s}.$$

Expanding this out gives

$$d i(X^s) \alpha^s + i(X^s) d\alpha^s = -\frac{\partial \alpha^s}{\partial s}.$$

Given the situation, it seems natural to require $i(X^s)\alpha^s = 0$. (In particular, $X^s \in \xi^s$.) Then we must have $i(X^s) d\alpha^s = -\partial \alpha^s / \partial s$. But since $X^s \in \xi^s$, and for any contact form $d\alpha^s|_{\xi^s}$ is nondegenerate, we can find such an X^s; and we have our one-parameter family of (germs of) diffeomorphisms!

One more thing to check: we wanted $X^s \in \mathcal{F}$, so that this was a flow along leaves of $\mathcal{F}$. Since all α^s have the same part β_0 which describes the foliation, for any vector V pointing along a leaf, $(\partial \alpha_s / \partial s)(V) = 0$. From the equation $i(X^s) d\alpha^s = -\partial \alpha^s / \partial s$, we then have $d\alpha^s(X^s, V) = 0$. But both X^s, V lie in ξ^s, so $d\alpha^s$ is nondegenerate on them; hence they are scalar multiples. So indeed φ_s is a flow along leaves.

And that is the power of the characteristic foliation.

6 Singularities and their orientations

Recall that the singularities which can occur in characteristic foliations are the non-isochore singularities. We have avoided the issue of orientations up to now; unfortunately now we have to deal with it.

Note that M is naturally oriented by the contact structure ξ (for any contact form α we get an orientation from $\alpha \wedge d\alpha$). If the orientation on S is reversed, this amounts to reversing the $\mathbb{R}$ coordinate on an $S \times \mathbb{R}$ neighbourhood, so $\alpha = \beta + u\ dt$ changes to $\beta - u\ dt$. In effect we replace u with $-u$ while β is unaffected. It's then easy to see that we must replace $\theta = u\ d\beta + \beta \wedge du$ with $-\theta$; recalling that θ is an area form on S, we see that the orientation is reversed as it should be. Then $\beta = i(X)\theta$ implies that we must replace X with $-X$; so the direction on the foliation is reversed. Basically, orienting the surface (and hence foliation) amounts to choosing a sign for u.

According to the choice of orientation, an elliptic singularity becomes a source or a sink. At a hyperbolic singularity x, we have two eigenvalues λ_1, λ_2 associated with the two separatrices, which are the eigenvalues of the linearisation of the flow at x. Here $\lambda > 0$ means that the foliation points outward and $\lambda < 0$ inward; for a hyperbolic singularity obviously there is one eigenvalue of each sign. The non-isochore condition means that $\lambda_1 + \lambda_2 \neq 0$. Reversing the orientation of S and $\mathcal{F}$ will replace the eigenvalues λ_1, λ_2 with the "same eigenvalues, but in the opposite direction" — namely, they become $-\lambda_1, -\lambda_2$.

Now, at a singularity x we have $\beta = 0$. Recall that at a singularity $u \neq 0$. (If $u = 0$ and $\beta = 0$ at x then $\theta = 0$ there too.) So x is not on Γ; and choosing an orientation at x amounts to choosing the sign of u.

Supposing that u is given in advance, of course, we have an orientation; let us call this the *pre-orientation.* But we can also choose a standard orientation at the singularity x. Recall that away from Γ, $\mathcal{F}$ is dilating for some area form. Well, with one direction $\mathcal{F}$ will be dilating, and with the opposite direction it will be contracting. We can define the *positive orientation* of $\mathcal{F}$ at x to be the one that makes it expanding.

Definition 6.1 *Let x be a non-isochore singularity of a singular foliation $\mathcal{F}$. The* positive orientation *of $\mathcal{F}$ (and hence S) at x is a direction of $\mathcal{F}$ near x by a vector field for which the divergence with respect to some area form at x is positive.*

So, at an elliptic singularity, with the positive orientation it is a source; with the negative orientation it is a sink. At a hyperbolic singularity, with eigenvalues $\lambda_1 < 0 < \lambda_2$, we have $\lambda_1 + \lambda_2 > 0$, which means that the "out direction", the "unstable direction", dominates, and $|\lambda_2| > |\lambda_1|$.

Since an orientation at a singularity x amounts to choosing a sign for u, we can ask: if x is positively oriented, what is the sign of u? It seems like it should be positive: and indeed it is. Let us see why.

Taking our area form θ on S and the vector field Y directing $\mathcal{F}$ given by

$\beta = i(Y)\theta$, the divergence is given by

$$\mathrm{Div}_{\theta} Y = \frac{L_Y \theta}{\theta} = \frac{i(Y)\ d\theta + d\ i(Y)\theta}{\theta} = \frac{d\beta}{\theta}.$$

So to say that the singularity x is oriented positively amounts to u being chosen so that $d\beta$ and θ have the same sign. At the singularity, we have $\beta = 0$ so $\theta = u\ d\beta + \beta \wedge du = u\ d\beta$. Thus, the positive orientation at x puts $u > 0$.

Using this, we can express the orientation condition in another way. We give a more geometric interpretation to the area form $\theta = u\ d\beta + \beta \wedge du$. In the first term, the $d\beta$, as we have seen, is essentially the divergence of Y. In the second term, since β is dual to Y, the wedge product is essentially $du(Y)$. Indeed, expanding out $i(Y)(\theta \wedge du) = 0$ (this is a 3-form on S) we obtain $i(Y)\theta \wedge du = -(i(Y)du)\theta$ so $\beta \wedge du = -(i(Y)du)\theta$. This gives

$$\theta = u\ d\beta + \beta \wedge du = u\theta\ \mathrm{Div}_{\theta} Y - (i(Y)du)\,\theta = (u\mathrm{Div}_{\theta} Y - i(Y)du)\,\theta$$

That θ is nondegenerate amounts to $u\mathrm{Div}_{\theta} Y - i(Y)\ du \neq 0$. The positive choice of orientation (and hence choice of sign of u and θ) gives us, at the singularity x, that $u > 0$, and since at x we have $Y = 0$, we obtain $u\mathrm{Div}_{\theta} Y - i(Y)du > 0$. Since this expression is never allowed to equal 0, it must remain positive over the whole of S.

We record all these conclusions.

Lemma 6.2 *Let S be a convex surface in a contact 3-manifold M. Let $\mathcal{F}$ be the characteristic foliation on S and x a singularity. Take a neighbourhood $S \times \mathbb{R}$ of $S = S \times \{0\}$ with contact form $\alpha = \beta + u\ dt$, where β is a 1-form and u a function on S, so that $\theta = u\ d\beta + \beta \wedge du$ is an area form on S. The positive orientation of $\mathcal{F}$ and S at x gives:*

(i) $d\beta$ and θ differ by a positive factor at x;

(ii) $u > 0$ at x;

(iii) $u\ \mathrm{Div}_{\theta} Y - i(Y)du > 0$ on all of S.

If we start with a pre-determined u, how do the pre-determined orientation and the positive orientation at a singularity differ? Clearly they agree at singularities where $u > 0$, and they disagree at singularities where $u < 0$. So all sources have positive orientation; all sinks have negative orientation; "outward-dominated" hyperbolic singularities have positive orientation; "inward-dominated" hyperbolic singularities have negative orientation. If we cut S along the dividing set $\Gamma = \{u = 0\}$, into the pieces $S_+ = \{u > 0\}$ and $S_- = \{u < 0\}$, then the singularities with positive orientation lie in S_+ and the singularities with negative orientation lie in S_-.

7 What a convex surface looks like II

Let us now return to a question we had before: we have a surface S in (M, ξ) with a characteristic foliation $\mathcal{F}$. Is S convex?

7.1 It has expanding vector fields directing the characteristic foliation, and dividing sets...

We know from previously that if S is convex, then we can choose coordinates so that S is $S \times \{0\}$ in $S \times \mathbb{R}$, that $\alpha = \beta + u\ dt$, and away from $\Gamma = \{u = 0\}$ we have Liouville vector fields and all the rest of it. But *now* we can go in the other direction too. If we have a characteristic foliation $\mathcal{F}$ with this Γ and dilating Liouville vector fields, then we have a nice contact structure. On the other hand though, because of the nature of the nice foliation, by proposition 4.1 there is a vertically invariant contact structure on $S \times \mathbb{R}$ inducing the foliation $\mathcal{F}$. By the power of the characteristic foliation, namely proposition 5.1, these the two germs of these contact structures are equivalent. So they are locally isomorphic; and there is a contact vector field on S which is transverse, because there was for the vertically invariant version. So S is convex. Thus we obtain Giroux's proposition II.2.1.

Proposition 7.1 *Let (M, ξ) be a contact 3-manifold, S an embedded closed orientable surface and $\mathcal{F}$ its characeristic foliation. The following are equivalent:*

(i) *S is convex.*

(ii) *There exists on S a curve Γ transverse to $\mathcal{F}$, possibly disconnected, which decomposes S into subsurfaces where $\mathcal{F}$ can be directed by a dilating vector field for a certain area form, and exiting through the boundary.*

7.2 ...if it's closed, it can't have a characteristic foliation defined by a closed form...

Suppose we have, as usual, a convex surface S, with an $S \times \mathbb{R}$ neighbourhood and contact form $\beta + u\ dt$. Recall that the characteristic foliation is defined on S by the 1-form β; and recall that the condition for β, u to define a contact structure is that the 2-form $\theta = u\ d\beta + \beta \wedge du$ on $S \times \{0\}$ is nondegenerate.

If β is closed, then $\theta = \beta \wedge du$ must be nondegenerate. But on a closed surface, the function u must have critical points, where $du = 0$, and so this requirement of nondegeneracy is impossible to satisfy.

This is something of a strange condition: if β is closed, then taking a nonzero function f on S, $f\beta$ defines the same foliation as β but $f\beta$ may not be closed. But the same argument should still work upon doing some Leibnitz rule and rearranging the terms.

Lemma 7.2 *Let S be a closed surface in a contact manifold M with characteristic foliation $\mathcal{F}$ defined by a 1-form β on S. If β is closed (as a form on S) then S is not convex.*

7.3 ...or a closed leaf with return map tangent to the identity...

Nor can there be a closed leaf F in the characteristic foliation $\mathcal{F}$ of a certain type. Every closed leaf F defines a return map, by taking a transverse interval which is sufficiently small; we then get a return map $f : [-1, 1] \longrightarrow [-1, 1]$, say, where the point 0 on the interval corresponds to the closed leaf. Then we have $f(0) = 0$. It turns out that $f'(0) = 0$ can't happen; this contradicts the convexity of the surface.

Why is this? Along the leaves of the foliation we have $\beta = 0$. The return map is determined by $d\beta|_F$; if we choose a product neighbourhood of F and a good coordinate system, we can easily see that we can take β such that $d\beta|_F = 0$ identically. But then along F, $\theta = u\ d\beta + \beta \wedge du = \beta \wedge du$... and u must have a maximum on the leaf F, at which $\beta \wedge du = 0$, contradicting the requirement $\theta \neq 0$.

Lemma 7.3 *Let S be a surface in a contact manifold M with characteristic foliation $\mathcal{F}$. If $\mathcal{F}$ contains a closed leaf F with return map tangent to the identity along F, then S is not convex.*

7.4 ... the dividing set avoids certain separatrices...

If S is convex, we can also consider where the dividing set Γ is in relation to our foliation $\mathcal{F}$. Of course we know about expanding area forms and so on, but we can pin down more: it turns out the dividing set can't go in certain places.

The idea is the following: taking a product neighbourhood and nice coordinate system $S \times \mathbb{R}$ we may take $\alpha = \beta + u\ dt$ as we have previously. The dividing set Γ is defined by $u = 0$. The characteristic foliation is defined by β, so the singular points of the foliation are where $\beta = 0$. If we know something about how u interacts with the characteristic foliation, we can say more. We will find that on certain leaves of the foliation — namely, leaves connecting hyperbolic singularities in a particular way — u cannot equal zero along them; and so Γ cannot intersect them.

Take a hyperbolic singularity x. Considering u as given a priori, and defining a pre-orientation, S has either a positive or negative orientation at x according to the sign of u. Suppose S has a positive orientation at x. As discussed previously, we may positively orient S and $\mathcal{F}$ at x. This means that $u > 0$ at x; and it means that $u \operatorname{Div}_\theta Y - i(Y)du > 0$ on all of S.

Now, there is one crucial point, which we discovered back in section 4.3. (Recall that there we characterised vertically invariant contact structures: but now we know that the germ of the contact structure on any convex surface can be described in this way.) The crucial point is: across Γ, Y points in the direction of decreasing u. So from a positively oriented hyperbolic singularity, where $u > 0$, Γ cannot intersect the stable (inwards) separatrix: the stable separatrix is oriented towards x, and u would have to decrease as we cross Γ, where $u = 0$, towards x; but $u > 0$ at x; this is a contradiction. Similarly, from a negatively oriented hyperbolic singularity x, Γ cannot intersect the unstable (outwards) separatrix: for $u < 0$ at x, and at an intersection with Γ on the unstable separatrix we would have $u = 0$ and increasing towards x, another contradiction.

Now, if such the stable separatrix from a positively oriented hyperbolic fixed point *intersects* the unstable separatrix from a negatively oriented hyperbolic fixed point, then we are in trouble! For u is positive at one fixed point and negative at the other, hence must be zero somewhere in between; but for the above reasons, this is impossible. This means that there cannot be such a leaf connecting the two hyperbolic fixed points; and it means that the two separatrices cannot intersect at a separate singularity.

Let us record these conclusions.

Lemma 7.4 *Let S be a convex surface in a contact 3-manifold M. Let the contact form be $\beta + u\ dt$ as before and let x be a hyperbolic singularity of the characteristic foliation of S.*

(i) If x is positively (resp. negatively) oriented then Γ does not intersect the stable (resp. unstable) manifold/separatrix of x.

(ii) The stable separatrix of a positively oriented hyperbolic singularity and the unstable separatrix of a negatively oriented hyperbolic singularity do not meet.

7.5 ...it's sufficient that the characteristic foliation is almost Morse-Smale...

Now we return to the situation where we have a surface S in a contact manifold, and want to know whether it is convex, by looking at its characteristic foliation. We already have a sufficient condition in terms of expanding area forms and dividing sets, but we can do even better.

In the previous section we noted that in the characteristic foliation of a convex surface, then certain separatrices do not meet. We can actually formulate a chain of possible properties of our foliation, where each clearly implies the next.

(i) (Morse–Smale) For any two separatrices of hyperbolic fixed points (possibly the same point) they do not intersect. (Except possibly in emanating from the same singularity.)

(ii) (Almost-Morse–Smale) Positively orient all hyperbolic fixed points and taking their stable separatrices. For any two such separatrices, they do not meet. (Except possibly in emanating from the same singularity.)

(iii) With respect to a given pre-orientation, we consider each hyperbolic fixed point to be positively or negatively oriented. For any given pair of hyperbolic fixed points, one positive and one negative, the stable separatrix of the positively oriented hyperbolic singularity and the unstable separatrix of the negatively oriented hyperbolic singularity do not meet.

If the foliation is otherwise nice, then the first property is the Morse–Smale property; the second property is called by Giroux the almost-Morse–Smale property. And the third is not called anything in particular, being such a mess, but is a property of characteristic foliations on convex surfaces, by the previous section. They are all strict implications and not equivalences. (It's easy to find an almost Morse–Smale foliation which is not Morse–Smale. And to find a foliation satisfying property (iii) but not (ii), take two positively oriented hyperbolic fixed points whose stable separatrices both connect to an elliptic fixed point, i.e. a source.) So convex implies property (iii). And we will shortly see that property (ii) implies convex. So in fact (i) (Morse–Smale) implies (ii) (almost Morse–Smale) implies convex implies (iii). In fact having (iii) but not (ii) implies we have connected elliptic and hyperbolic fixed points of the same orientation and can apply the elimination lemma, which we will come to shortly. But let us not get ahead of ourselves!

To be more precise about these definitions:

Definition 7.5 *A singular foliation $\mathcal{F}$ on a closed surface S is called* Morse–Smale *if it satisfies the following conditions:*

(i) the singularities and closed leaves of $\mathcal{F}$ are hyperbolic (in the dynamical systems sense!);

(ii) the limit set of each half-leaf is a singularity or a closed leaf;

(iii) For any two separatrices of hyperbolic fixed points (possibly the same point) they do not intersect. (Except possibly in emanating from the same singularity.) As Giroux puts it, there are "no connections between saddles".

Definition 7.6 *A singular foliation $\mathcal{F}$ on a closed surface S is called* almost-Morse–Smale *if it satisfies the first two conditions above, and instead of (iii), satisfies the following: when we orient $\mathcal{F}$ positively near hyperbolic fixed points, the associated stable manifolds do not meet each other.*

Now let us see why almost Morse–Smale implies convex. Given a closed surface S with characeristic foliation $\mathcal{F}$, we explicitly find regions on which a vector field directing $\mathcal{F}$ dilates (or contracts) an area form. Around each elliptic point we take a small disc. Around each closed leaf we take a small annulus. At each hyperbolic singularity, we orient it positively and consider the stable separatrix; we take a small band around it. We take the union of these annuli, discs and bands, which forms a (possibly disconnected) surface S_0. Note that from previous discussion that S_0 will not intersect Γ. In fact the idea is that S_0 is homeomorphic to $S \backslash \Gamma$. Note that the orientations on components of S_0 may not agree with a pre-orientation on S; but on each individual component of S_0, the positive orientations on singularities agree. (This doesn't (and can't!) rely on any convex-specific assumptions — it's clear from the topology.)

By our construction, we can find an area form ω on S_0 and a vector field Y directing $\mathcal{F}$ such that $\mathrm{Div}_\omega Y > 0$. (We can do this because, positively orienting everything, the foliation points out of the discs and bands; and we can easily do the same around closed leaves.) We are then basically done. For what does the rest of the surface look like? It has a foliation without any singularities or closed leaves, and it is part of the connected oriented closed surface S. Hence it must be a union of annuli, and the foliation must be the standard product foliation of intervals from one boundary component to the other. Each one of these annuli will contain a curve of Γ in its core. Taking Γ as such, making some components of S_0 positive and others negative, and gluing together, we have all the requirements of a characteristic foliation required to be convex. Now we have Giroux's proposition II.2.6.

Proposition 7.7 *Let S be an orientable closed surface embedded in a contact 3-manifold. If the characteristic foliation of S is almost Morse–Smale, then S is convex.*

Giroux quotes a theorem of M. Peixoto that a vector field on a closed manifold is C^∞-generically Morse–Smale. Since a C^∞ perturbation of a surface C^∞-perturbs the characteristic foliation, it follows that convex surfaces are generic.

Proposition 7.8 *Let S be a closed orientable surface in a contact 3-manifold. If ξ is transversely orientable, then the characteristic foliation of S is directed by a vector field that we can make Morse–Smale by a C^∞-small isotopy of S in V; and hence generically S is convex.*

8 What Giroux says about eliminating singularities

Giroux has what I believe is the first proof of "the elimination lemma", allowing you to remove elliptic and hyperbolic fixed points which are connected by a leaf and which have the same sign. Giroux's (1991) version is significantly weaker than the version quoted by Eliashberg (1992) in "Contact 3-manifolds 20 years since J. Martinet's work". Eliashberg there says the improvement is due to Fuchs (but cites no paper of Fuchs), and refers to another (1991) Eliashberg paper, "Legendrian and transversal knots in tight contact manifolds", for a proof.

We can compare the two versions of the statement.

Proposition 8.1 (Giroux) *Let S be a closed orientable surface embedded in a contact 3-manifold with a Mose–Smale characteristic foliation $\mathcal{F}$. (Hence S is convex; let Γ be the dividing set.) Let x_0 be an elliptic and x_1 a hyperbolic fixed point such that when we positively orient the foliation near x_1, one or both of the stable separatrices comes from x_0. Then:*

(i) There exists in S an annulus A disjoint from Γ and satisfying the following:

- *the only singularities of $\mathcal{F}$ on A are x_0 and x_1;*
- *$\mathcal{F}|_A$ has no closed leaf;*

- *$\mathcal{F}$ is transverse to the boundary of A/*

(ii) There exists an $S \times \mathbb{R}$ neighbourhood of $S = S \times \{0\}$ and a function $k : A \longrightarrow (-\infty, 0]$ with support in the interior of A such that the characteristic foliation of the graph of k (which is an arbitrarily small perturbation of S) has no singularities.

On the other hand, the stronger version, cited by Eliashberg, is as follows.

Proposition 8.2 (Giroux, Fuchs) *Let S be a surface with (possibly empty) Legendrian boundary and characteristic foliation $\mathcal{F}$ in a contact 3-manifold (M, ξ). Let C be a trajectory of $\mathcal{F}$ whose closure contains an elliptic point x_0 and a hyperbolic point x_1 with the same orientation. Let U be a neighbourhood of C in M which contains no other singular points of $\mathcal{F}$. Then there exists a C^0-small isotopy of S in M which is supported in U, fixed at C and such that the new surface S' has no singular points of the characteristic foliation $\mathcal{F}$ inside U. If x_0 and x_1 belong to the Legendrian boundary of S then one can kill them leaving ∂S fixed.*

I won't detail Giroux's proof, because it's probably better to refer to the later paper. Finding the annulus in part (i) is easy, one just has to fiddle a little to deal with the various possibilities. Then finding the function k one finds a function which depends purely on the "radial" coordinate of the annulus, supported in its interior and which is sufficiently large that its "Hamiltonian" vector field Y_k — defined by $i(Y_k)\omega = dk$ for some area form ω — is very large away from the boundary of the annulus. Hence considering $Y + Y_k$ is nonzero, and the dual form $\beta + dk$ has no singularities nearby; and $\beta + dk$ describes the characteristic foliation on the perturbed surface which is the graph of k.

Giroux also shows how to simplify the characteristic foliation to reduce an overtwisted disc to a standard form; ; and how to eliminate elliptic singularities — much as Eliashberg does in his "20 years on" paper, in greater generality. So I won't detail it here.

9 The dividing set is all powerful

So far we have seen that a convex surface is closely related to the dividing set Γ. A convex surface certainly has a dividing set. And a surface's convexity can be detected by finding a dividing set Γ, on the complement of which the characteristic foliation is directed by a vector field which dilates an area form. If S has an almost Morse–Smale foliation, then we showed that S is convex by constructing regions on which S dilated area forms, and the remainder was just a tubular neighbourhood of Γ; so Γ was quite canonical, topologically.

But the dividing set is even more powerful than this. Suppose you have a closed orientable surface S in a contact 3-manifold M, which is convex. So you have $\mathcal{F}$ and Γ. But suppose you only have the information of Γ; you don't know what the characteristic foliation is. Is $\mathcal{F}$ determined? Well obviously not: there are potentially many foliations dilating area forms away from Γ; if you take one, you can perturb it. But, it turns out that all possible characteristic foliations $\mathcal{F}$ compatible with Γ can be achieved by perturbing S a little. In effect, Γ captures precisely the important information about $\mathcal{F}$; once you have Γ, provided you don't mind wiggling the surface around a tiny bit, you can have whatever (compatible) $\mathcal{F}$ you like.

This is quite amazing. Humans have a hard time visualizing contact planes fluttering around, even on a surface, let alone in space. But it turns out, as we have seen, that the (germ of a) contact structure near a surface is determined by its characteristic foliation. And now, the characteristic foliation, up to a bit of surface-jiggling, is determined by the dividing set. So we need not visualize fluttering planes; we need not even visualize foliations and trajectories; we only need visualize a few curves on a surface.

Let us make this precise. When we say that $\mathcal{F}$ is "compatible" with Γ, we mean something about area forms and stuff. What we actually mean is:

Definition 9.1 *Let S be a closed orientable surface and Γ a 1-manifold in S. Let S_Γ denote the compact surface with boundary obtained by cutting S along Γ. A singular foliation $\mathcal{F}$ on S is* adapted *to Γ on S if $\mathcal{F}$ is directed by a vector field which dilates an area form on S_Γ, and which exits transversely through the boundary ∂S_Γ.*

Given Γ, we are going to wiggle the convex surface S. This is means an isotopy $\delta_s : S \longrightarrow M$ for $s \in [0,1]$ where δ_0 is the identity. Thinking of an $S \times \mathbb{R}$ neighbourhood where the transverse contact vector field is vertical, we see that S will remain convex as long as it is the graph of a function; and the dividing set will always be the intersection of $\Gamma \times \mathbb{R}$ with the jiggled surface. We have a characteristic foliation to begin with, $\mathcal{F}_0$ say, and a foliation we would like to get as a characteristic foliation, $\mathcal{F}_1$; both are adapted to Γ.

So, taking an area form such as θ on S, we see that our contact form α can be written $\beta + u\ dt$. Away from Γ, we can write our contact form as $\beta + dt$; it's a contactization. We can direct $\mathcal{F}_0$ by Y_0, defined by $i(Y_0)\theta = \beta$. We can direct $\mathcal{F}_1$ by Y_1, where Y_1 dilates a (possibly different) area form on S_Γ. But $\text{Div}_{\pm e^g \theta} Y = e^{-g} \text{Div}_\theta (e^g Y)$, so by adjusting Y_1 by a nonzero function, we can arrange that Y_1 dilates θ also. The idea is simply to set $Y_s = (1-s)Y_0 + sY_1$, hence $\beta_s = i(Y_s)\theta$, and once we obtain a u_s we will have a family of contact forms α_s and we can Moser away to get our isotopy. Note the vector fields Y_s can certainly involve new singularities; these are the singularities moving around, being created and destroyed on the shifting foliations $\mathcal{F}_s$

There are some things to think about, though. Near Γ we have to fiddle a little. We take vector fields there to agree with $\pm Y_0, \pm Y_1$ on the boundary of a tubular neighbourhood, and glue together and linearly interpolate. So we can extend Y_s over all of S. We must also take a bunch of functions u_s (they only matter near Γ, though, since elsewhere it's a contactization) such that the contact condition is satisfied: $u_s \text{Div}_\omega (Y_s) - Y_s \cdot u_s > 0$. This requires a fiddle also. But then we have a family of contact structures ξ_s, and can Moser to obtain an isotopy realising them. Since our $S \times \mathbb{R}$ neighbourhood can be arbitrarily small, we obtain the following result.

Proposition 9.2 *Let S be a convex closed orientable surface in a contact 3-manifold M, with transverse contact vector field X and dividing set Γ. Let $\mathcal{F}$ be a singular foliation on S adapted to Γ. Then there exists an isotopy $\delta_s : S \longrightarrow V$ for $s \in [0,1]$ such that:*

(i) δ_0 *is the identity;*

(ii) each surface $\delta_s S$ is transverse to X, hence convex;

(iii) the dividing set of $\delta_s S$ (with respect to X) is $\delta_s \Gamma$.

So indeed; by a small perturbation, we can get whatever characteristic foliation we like. This is the power of convex surfaces. Note, in particular, that it implies all the elimination lemmas and results about standard forms of characteristic foliations, immediately. It subsumes them all.

References

[1] Emmanuel Giroux, "Convexité en topologie de contact", Comment. Math. Helvetici, 66 (1991) 637-677.