Universidad Autónoma de Madrid
Instituto de Ciencias Matemáticas

Engel structures and symplectic foliations

On their global topology: flexibility and rigidity

Author:
Álvaro del Pino Gómez

Advisor:
Francisco Presas Mata

Para mi abuelo. To Dóra

Contents

Agradecimientos

He tenido la buena suerte de llevar a cabo un doctorado que ha sido, en casi todo momento, un disfrute. Esto no hubiera sido posible sin el apoyo de muchas personas, tanto a nivel personal como profesional.

Primero tengo que darle las gracias a Fran, por haber sido un genial director de tesis, un estupendo colaborador, y un gran amigo. He aprendido mucho de tí y creo que en los últimos años hemos aprendido muchas mates el uno del otro también. ¡Espero que esto continúe siendo así! También le debo mucho a Klaus por acogerme con entusiasmo en Budapest, por las muchas cervezas compartidas y por las horas dedicadas a discutir mates. Los dos tenéis actitudes matemáticas muy diferentes y me alegro de haber podido ver ambas.

Quiero mencionar también a todos mis coautores: R. Casals, D. Martínez Torres, D. Peralta y J.L. Pérez. Las matemáticas no las hace uno solo y ha sido un placer compartirlas con vosotros.

I want to thank also the people at the Renyi Institute for tolerating me, even when there was a major lack of space and I kept occupying random desks. My stay there was incredibly pleasant due to the nice mathematical environment that everyone created. Thanks particularly to András, Keon, Stefan, Gilberto, Roman, Baldur, and Antonio.

I am also very grateful to the people at the LMU Munich for having me there for three weeks and for showing an interest in the things I am doing. The questions you have asked me and the discussions we have had have motivated a lot of the things I have been thinking about for the past few months. Thanks to Jonathan, Dieter, Nicola, and, particularly, Thomas: discussing math with you has been very enlightening (and entertaining) and I cannot thank you enough for reviewing this thesis with the depth you did.

Gracias también a la vieja guardia de los estudiantes de Eva (Anna, Romero y Arnau) por tratarme súper bien cada vez que he visitado Barcelona.

I have a trip pending to India to visit Sushmita and Dishant, both of which have been incredibly kind to me. Dishant: I promise we will finish the project we have pending. Sushmita: you have spent an incredible amount of time looking at this thesis and thanks to you other people reading it will have a much easier time; I am (very much) indebted to you for this.

He pasado muchos años compartiendo mi tiempo con mucha gente del ICMAT. Con algunos de vosotros he echado un montón de horas organizando/asistiendo/hablando en seminarios (a veces un poco locos y sin orden ni concierto). Me lo he pasado muy bien (que es lo importante) y he aprendido cosas (o, al menos, me da la sensación de que sé algo más que antes). Aquí va una lista (en orden aleatorio) y espero no olvidarme de nadie: Pablo, Paco, Carlos, Ángel, Juan, Juanjo, Alfaya, José Ángel, Sonja, Víctor, Mariángeles, Manuel, Miguel, Tania. Gracias también a Dani, al que he debido de preguntarle como ocho millones de veces sobre formalidades para presentar la tesis.

He tenido el placer de poder hablar y pensar sobre matemáticas con gente que sabe mucho más que yo y que han estado dispuestos a compartir lo que saben conmigo. Muchas gracias a Vicente (Muñoz), Eva (Miranda) y Luis (Hernández). Thanks also to Gael (Meigniez) for spotting that one fatal mistake.

También quiero acordarme del personal del ICMAT, de la Gerencia del CFTMAT, y de la UAM que han estado ahí todos estos años haciendo papeleo por mí. ¡Muchas gracias!

Finally, I would like to thank the members of my thesis committee for agreeing to partake in such an event: Javier Aramayona, Viktor Ginzburg, Alberto Ibort, Eva Miranda, and Thomas Vogel. It really is an honour for me. Muchas gracias también a Luis Guijarro, mi tutor en la UAM, por hacer la cuestión de los papeleos fácil, y a Gabino González-Díez y Tomás Gómez por estar de apoyo con la defensa de la tesis.

Dejo lo más importante para el final. Muchas gracias a mis padres, a mis abuelos, y a mi hermana. Me veis ir y venir de un lado a otro con esta excusa de las mates y no os quejáis demasiado. Gracias por estar siempre ahí. Vosotros no hace falta que me llaméis doctor de ahora en adelante.

Gracias también a mis amigos: sois demasiados para mencionaros a todos, así que hagamos lo siguiente. Thanks, ...![1]

Lastly... Dóra, you are pretty cool (for someone dating a soon–to–be math PhD).

[1] Write your name here for a personalised experience

Contenidos y conclusiones

Esta tesis es la culminación del trabajo que he llevado a cabo durante mis cuatro años de doctorado. Buena parte de los contenidos de la misma han aparecido con anterioridad en forma de artículo:

[9] R. Casals, A. del Pino, and F Presas. "h-Principle for Contact Foliations". *Int. Math. Res. Not.* 20 (2015), pp. 10176–10207.

[10] R. Casals, J.L. Pérez, A. del Pino, and F Presas. "Existence h-Principle for Engel structures". *arXiv e-prints* (2015). arXiv: 1507.05342.

[14] A. del Pino. "On the classification of prolongations up to Engel homotopy". *Submitted* (2016).

[15] A. del Pino and F. Presas. "Flexibility for tangent and transverse immersions in Engel manifolds". *arXiv e-prints* (2016). arXiv: 1609.09306.

[16] A. del Pino and F. Presas. "The foliated Weinstein conjecture". *To appear in Int. Math. Res. Not.* (2015). arXiv: 1509.05268.

[53] D. Martínez Torres, A. del Pino, and F. Presas. "Transverse geometry of foliations calibrated by non-degenerate closed 2-forms". *To appear in Nagoya Mathematical Journal* (2014). arXiv: 1410.3043.

[68] D. Peralta–Salas, A. del Pino, and F. Presas. "Foliated vector fields without periodic orbits". *To appear in Israel Journal of Mathematics* (2014). arXiv: 1412.0123.

Aunque algunas de las publicaciones han sido reproducidas con mínimos cambios de estructura y estilo, otras han sido expandidas (muy) sustancialmente. Los Capítulos 3 y 4 de la Parte II presentan resultados nuevos cuyo punto de partida son las técnicas introducidas en mi artículo [10]. Se trata, posiblemente, del material más importante de la tesis.

Temáticamente, la tesis se puede dividir en dos bloques claramente diferenciados aunque relacionados. He decidido hacer esta dicotomía evidente estructurando el documento en Parte I y Parte II. Mi idea en este capítulo es introducir al lector a algunos de los conceptos básicos que son comunes a ambas partes, dar una idea general de qué se prueba en esta tesis y, finalmente, discutir qué preguntas quedan abiertas para investigación futura.

La tesis en términos sencillos

Esta sección pretende ser una introducción informal para el lector no matemático (familia y amigos del autor).

El h–principio

Aunque tocaremos muchos otros temas, podría argumentarse que buena parte de los resultados que presentamos tratan sobre h–principio. Así que empecemos, ¿qué es el h–principio?

Supongamos que tenemos una cierta ecuación diferencial $F(f, f') = 0$, donde F no es más que una expresión que combina la función f y su derivada f' de alguna manera (por ejemplo: $F(f, f') = f^2 - f'$). Estamos interesados en encontrar funciones f que precisamente hagan que F nos de cero. Está claro que primero podemos intentar resolver el siguiente problema más sencillo: encontrar funciones g y h que no estén relacionadas de ninguna manera pero que cumplan $F(g, h) = 0$. Lo que hemos hecho es *desacoplar* las variables: lo que antes eran una función y su derivada ahora son dos funciones independientes. Se suele decir que h es la *derivada formal* de g y llamamos *solución formal* de F al par (g, h).

En vez de una ecuación diferencial como la que hemos presentado, podríamos trabajar con un concepto un poco más general. F podrían ser varias expresiones que dependen de más de una función y de derivadas de orden mayor. A la vez, en vez de pedir que F sea cero, podríamos pedir que sea distinta de cero, o mayor que un cierto número o, en general, que tome valores en un cierto conjunto. Entonces decimos que estamos considerando una *relación en derivadas parciales*.

Una cierta relación en derivadas parciales se dice que satiface el h–principio (*principio de homotopía*) si toda solución formal se puede homotopar (digamos, modificar adecuadamente) a una solución de verdad. Si esto ocurre, un geómetra está contento, porque el problema que se había planteado ahora se reduce a resolver la ecuación desacoplada y esto suele ser relativamente sencillo (o, al menos, se convierte en un problema que su amigo que estudia topología algebraica sabe resolver). Probar que el h–principio se satisface depende mucho del problema en cuestión. Existen resultados clásicos que dan criterios generales bajo los cuales esto ocurre. Para explicarlo, necesitamos un par de conceptos adicionales.

Los espacios en los que trabajamos son lo que llamamos *variedades*. Una variedad es por ejemplo la recta real, el plano, o la esfera. Variedades como la recta, donde uno puede alejarse infinitamente, se dicen abiertas. Variedades donde esto no ocurre, como la esfera, se dicen cerradas. Entonces, las funciones con las que trabajamos no dependen de un número, si no del punto en la variedad correspondiente. Esto es muy natural. El lector por ejemplo puede pensar que la temperatura es una función que a cada punto de la superficie de nuestro planeta le asigna un número.

Por otro lado, generalmente no estamos interesados en funciones que nos den números. Estamos interesados en funciones que toman valores en lo que llamamos *fibrados*. Un ejemplo sencillo es el siguiente: en cada punto de la superficie de la Tierra (fuera de los polos) ponemos una flecha apuntando al polo norte. Podemos pensar que esto es una función que a cada punto le está asignando un vector en el correspondiente plano tangente.

M. Gromov probó en su tesis que muchas relaciones en derivadas parciales en variedades abiertas satisfacen el h–principio. La idea (vaga) subyacente a este fenómeno es que somos capaces de empujar nuestros problemas al infinito de la variedad. El problema usual al que nos enfrentamos entonces es al de dilucidar si una relación en una variedad cerrada satisface el h–principio o no.

Distribuciones totalmente no–holónomas

El h–principio se estudia en muchos contextos, pero nosotros vamos a centrarnos en usarlo en el estudio de distribuciones. Una distribución es lo siguiente: es una función que a cada punto de nuestra variedad le asigna un cierto espacio lineal de direcciones (por ejemplo, un plano). La idea es que una partícula en la variedad sólo puede moverse siguiendo direcciones tangentes a dicho espacio. Un ejemplo (más o menos) de este fenómeno es por ejemplo un ascensor. El ascensor sólo puede moverse hacia arriba y hacia abajo, así que el espacio lineal que describe esta restricción es una línea vertical.

Muchos sistemas mecánicos y de teoría de control pueden entenderse en este lenguaje. Una pregunta natural en este contexto es la siguiente: si tenemos una cierta distribución, ¿podemos llegar desde el punto A de nuestra variedad al punto B (siguiendo caminos tangentes a la distribución)? Para ciertas distribuciones esto no es posible (por ejemplo, el ascensor no puede salirse de la vertical en la que está situado). Una distribución se dice *totalmente no-holónoma* si podemos llegar a un punto desde cualquier otro (la definición dice que además esto se ha de cumplir infinitesimalmente, es decir, si tomamos bolitas cada vez más pequeñas podemos llegar desde un punto de la bolita a cualquier otro sin salirnos de ella). Esto no quiere decir que la distribución le permita a uno moverse en todas direcciones; el ejemplo usual que se suele dar es el siguiente: si estamos montados en un coche, sólo podemos hacer dos cosas: girar el volante o movernos hacia adelante o hacia atrás (en particular, no podemos movernos hacia los lados). Sin embargo, todos sabemos que a base de hacer maniobras moviendo el volante y yendo un poco hacia adelante y un poco hacia atrás, sí que podemos realizar un desplazamiento a los lados (esto es lo que hacemos al aparcar). Es decir, la dirección de movimiento que nos falta la obtenemos por combinación de las otras dos.

Volviendo a nuestra discusión anterior, resulta que ser totalmente no–holónomo es una relación en derivadas parciales. Esta tesis está dedicada a probar que existen distribuciones totalmente no–holónomas en variedades de dimensión 4.

El estudio de las distribuciones totalmente no-holónomas es importante en dinámica. El primer ejemplo es el de las distribuciones de rango 2 (esto es, planos) en variedades de dimensión 3, a las que llamamos *estructuras de contacto*. Desde los años setenta han sido estudiadas en profundidad y se ha descubierto que interaccionan fuertemente con las variedades en las que viven (esto es, que podemos deducir propiedades de la variedad a partir de la clasificación de las estructuras de contacto en ella), así como con otras estructuras geométricas.

Lo que es relevante para nuestra discusíon es lo siguiente: existen dos tipos de estructuras de contacto. Las estructuras *overtwisted* y las estructuras *tight*. Una estructura es overtwisted si existe una bola donde la estructura tiene un aspecto especial; a este modelo especial lo llamamos el *disco overtwisted*. Ocurre el siguiente fenómeno: la relación en derivadas parciales de ser de contacto y contener a un disco overtwisted satisface el h–principio. Esto quiere decir que cualquier distribución de planos en dimensión 3 puede ser deformada para ser una estructura de contacto overtwisted. Esto automáticamente dice que las estructuras overtwisted no son muy interesantes (porque las podemos construir fácilmente); de esta forma, durante los últimos treinta años el área de la topología de contacto se ha dedicado a entender las estructuras tight, que son las que tienen propiedades interesantes y no detectables con h–principio.

La Parte II de esta tesis trata del siguiente posible ejemplo de distribución totalmente no–holónoma: las *estructuras Engel*. Éstas son distribuciones de planos pero ahora en variedades de dimensión 4. Esta tesis define lo que es una estructura Engel overtwisted y prueba que la correspondiente relación en derivadas parciales satisface el h–principio. En particular, se muestra cómo construir estructuras Engel en abundancia y cómo clasificarlas si son overtwisted.

La gran pregunta que se deja abierta es: ¿hay estructuras Engel no overtwisted?

Foliaciones

La Parte I de la tesis trata sobre foliaciones, así que empecemos explicando qué son. Si cogemos un taco de folios, podemos pensar en él de dos maneras. Por un lado, podemos tratarlo como un objeto tridimensional indivisible. Por otro, podemos ver que en realidad se trata de una multitud de folios diferentes apilados. Idealmente, podríamos imaginar un taco de folios donde hay infinitos folios infinitamente finos apilados. En esta situación, el taco de folios es nuestra variedad, y cada uno de los folios es lo que llamamos una *hoja*. La división de la variedad en hojas es lo que llamamos la *foliación*. Una foliación es un ejemplo particular de distribución (invito al lector a pensar en ello un rato), pero no es un ejemplo de distribución totalmente no holónoma. Al contrario: la idea es que sólo nos podemos mover de un punto a otro si están en la misma hoja.

Resulta que la teoría de foliaciones en dimensión 3 es increíblemente rica. Su historia es muy similar a la de las estructuras de contacto. Por un lado, se cumple el h–principio: hay una receta general para construir foliaciones en cualquier variedad. Sin embargo, esta receta es muy particular: en el último paso se introduce un cierto modelo local por todos lados. A este modelo se le llama la *componente de Reeb* y juega un papel similar al del disco overtwisted. Lo que queremos entonces es estudiar aquellas foliaciones que no contienen componentes de Reeb. Se definen entonces lo que se llaman las foliaciones *taut*, que más o menos vienen a ser esto mismo.

Hoy en día una de las conjeturas abiertas más importantes en topología de dimensión baja dice que para ciertas variedades de dimensión 3 (las esferas de homología) el poseer o no una foliación taut puede caracterizarse de otras dos maneras que son puramente topológicas. En general, la idea que queremos transmitir es que, si queremos entender "la forma" que tiene una variedad, una de las herramientas que tenemos a nuestra disposición es ver si admite foliaciones taut e intentar clasificarlas.

La definición de foliación taut (que no daremos al lector porque resulta un poco innecesaria para la discusión), puede darse también para variedades de dimensión superior. Sin embargo, la situación es muy diferente de dimensión 4 en adelante, ya que un resultado muy reciente prueba que existe una receta para construirlas. La pregunta que exploramos en la Parte I de la tesis es cómo generalizar esto de ser taut a dimensión superior de forma que las foliaciones que consideremos tengan propiedades interesantes.

Parte I

Una foliación por superficies en una 3–variedad se dice taut si existe una 2–forma global cerrada que es área sobre cada hoja. La generalización usual a dimensión superior consiste en considerar foliaciones de codimensión–1 calibradas por una forma cerrada de codimensión–1. Sin embargo, como explicamos en la sección anterior, un resultado de Meigniez [57] prueba que a partir de dimensión 4 esta clase de foliaciones satisface el h–principio.

En [52], Martínez–Torres propuso una generalización alternativa más natural desde el punto de vista simpléctico: una foliación se dice *fuertemente simpléctica* si existe una 2–forma global cerrada que es no–degenerada sobre cada hoja. Entonces, de la misma manera en que las variedades de contacto aparecen como condición de frontera natural para las variedades simplécticas, uno puede introducir el concepto de *foliación de contacto* en el contexto foliado. Éstas no son más que foliaciones que admiten una distribución de corango–1 que es de contacto en cada hoja.

El h–principio de existencia

El primer resultado que presentamos muestra que, si no introducimos restricciones adicionales, las foliaciones de contacto no son muy interesantes:

Theorem 2.1. *Sea V una 4–variedad y sea $\mathcal{F}$ una foliación de rango 3. La inclusión*

$$\mathfrak{ContFol}(V, \mathcal{F}) \to \mathcal{F}\mathfrak{ContFol}(V, \mathcal{F})$$

del espacio de foliaciones de contacto en el espacio de foliaciones de contacto formales es una sobreyección en todos los grupos de homotopía.

Una foliación de contacto formal no es más que una solución formal de la relación en derivadas parciales que define a las foliaciones de contacto. Es decir, este enunciado simplemente dice que se cumple el h–principio (de existencia).

El resultado no es muy sorprendente de acuerdo a lo que hemos explicado antes. Existe una clase de estructuras de contacto, las overtwisted, que satisfacen el h–principio y el teorema es una manifestación de esto mismo en el contexto foliado. La referencia original es mi artículo [9], pero la prueba que damos aquí ha sido simplificada considerablemente.

Técnicas de Donaldson

A continuación, dejamos las foliaciones de contacto para centrarnos en las foliaciones simplécticas fuertes. El teorema principal en esta dirección, que apareció primero en [53], es una análogo del teorema del hiperplano de Lefschetz:

Theorem 3.5. *Sea M^{2n+1} una variedad cerrada. Sea $(\mathcal{F}, \omega)$ una foliación fuertemente simpléctica en M y sea W un divisor de Donaldson de codimensión 2.*

Para toda hoja $\mathcal{L}$ de $\mathcal{F}$ se cumple:

$$\pi_k(\mathcal{L}, \mathcal{L} \cap W) = \{1\},\ 0 \leq k \leq n-1.$$

Necesitamos introducir un poco de terminología adicional antes de explicar el teorema. Cuando una foliación admite una 2–forma global que es simpléctica en cada hoja, se dice que es *(débilmente) simpléctica*. La diferencia entre fuerte y débil está en que sólo en el caso fuerte la 2–forma define una clase de cohomología, y esto es esencial para adaptar ciertas técnicas de topología simpléctica al contexto foliado.

Cuando la 2–forma no sólo es cerrada si no que además es integral, existe un fibrado hermítico $L \to M$ del cual es la curvatura. Bajo estas hipótesis, uno es capaz de construir secciones de L^k, para k grande, cuyos ceros son subvariedades que heredan la estructura de foliación simpléctica fuerte. A estas subvariedades las llamamos divisores de Donaldson. El teorema dice que (al igual que en el caso proyectivo y en el caso simpléctico), los divisores recuerdan parte de la topología del ambiente.

Una consecuencia particularmente relevante es que toda variedad con una foliación simpléctica fuerte tiene una subvariedad con una foliación taut y sus espacios de hojas son isomorfos.

Dinámica foliada

En los Capítulos 4 y 5 estudiamos dos cuestiones que se alejan un poco del hilo conductor de la tesis (aparentemente al menos), pero que están relacionadas entre sí.

Consideremos la siguiente pregunta natural: ¿existe alguna relación entre la topología del espacio ambiente y la dinámica de los campos de vectores en el mismo? De forma aún más precisa: ¿existe alguna condición topológica que fuerce la aparición de órbitas periódicas para cualquier campo?

La respuesta a esta segunda pregunta es que no. Por trabajos de Wilson [87] y Kuperberg [48] sabemos que cualquier variedad de dimensión al menos 3 posee campos sin órbitas periódicas. El resultado principal del Capítulo 4 generaliza este fenómeno a familias de campos de vectores de cualquier dimensión:

Theorem 4.1. *Sea $(M^{n+m}, \mathcal{F}^n)$, $n \geq 3$, una foliación. Denotamos por $\mathfrak{X}_{ns}(M, \mathcal{F})$ al espacio de campos de vectores sin ceros tangentes a $\mathcal{F}$. Escribimos $\mathfrak{X}_{no}(M, \mathcal{F})$ para denotar al subespacio de aquellos campos que además no tienen órbitas cerradas.*

La siguiente inclusión es una equivalencia de homotopía débil:

$$\iota_n : \mathfrak{X}_{no}(M, \mathcal{F}) \to \mathfrak{X}_{ns}(M, \mathcal{F}).$$

Aunque se trata de una cuestión de dinámica, la respuesta tiene la forma de un h–principio.

Por otro lado, sabemos que ciertas clases de campos de vectores siempre tienen que tener órbitas cerradas. Para justificar esta afirmación necesitamos introducir algunas nociones nuevas. Las variedades de contacto no son más que el lenguaje natural (aunque ligeramente más general) para formalizar el concepto de nivel de energía en mecánica hamiltoniana. Como tal, toda variedad de contacto tiene una serie de flujos asociados (los *flujos de Reeb*) que juegan el papel del flujo hamiltoniano del sistema. Una famosa conjetura de Weinstein que ha guiado buena parte de la investigación en dinámica de contacto dice que todo flujo de Reeb tiene una órbita cerrada. A principios de este milenio, Taubes probó que este era el caso en toda variedad de dimensión 3.

En el Capítulo 5, que se corresponde a mi artículo [16], estudiamos la cuestión análoga en el contexto foliado:

Theorem 5.1. *Sea $(M^{3+m}, \mathcal{F}^3, \xi^2)$ una foliación de contacto en una variedad cerrada M. Sea α una forma de contacto para una extensión de ξ y sea R su campo de Reeb. Sea $\mathcal{L}^3 \hookrightarrow M$ una hoja.*

i. Si $(\mathcal{L}, \xi|_{\mathcal{L}})$ es una variedad de contacto overtwisted, R tiene una órbita cerrada en la clausura de $\mathcal{L}$.

ii. Si $\pi_2(\mathcal{L}) \neq 0$, R tiene una órbita cerrada en la clausura de $\mathcal{L}$.

Las técnicas que usamos para probar el resultado se basan en el uso de curvas pseudoholomorfas, que fueron introducidas por Hofer para probar la conjetura de Weinstein para variedades overtwisted. En particular, parte de nuestro trabajo ha sido el de sentar las bases de la teoría para foliaciones. Es de esperar que estas ideas puedan ser usadas para atacar cuestiones de clasificación de foliaciones de contacto.

Hay que recalcar que en este mismo capítulo se construyen varios ejemplos que demuestran que las hipótesis del teorema no son sólo suficientes, sino necesarias. En particular, probamos que la conjetura de Weinstein no es cierta para foliaciones que tengan sólo hojas tight.

Parte II

Por poner las cosas en contexto, la situación actual en el estudio de espacios de distribuciones es un poco la siguiente. Por un lado, existen sendas industrias dedicadas a entender estructuras de contacto y foliaciones. Al menos en dimensión tres, se ha llegado a un punto donde buena parte de ello es zoología: ¿cuántas estructuras de contacto tight admiten las variedades de la familia tal? Muy recientemente, nuevos resultados [6] han iniciado el estudio de variedades de contacto de dimensión superior; esto es un campo todavía inexplorado.

Fuera de estos casos particulares, el estudio de distribuciones es puramente local. No se entiende si otras clases de distribuciones interaccionan de alguna manera interesante con la topología del espacio ambiente en el que viven o si tienen alguna relación con otras estructuras geométricas.

Las estructuras Engel son el siguiente ejemplo más sencillo de distribución totalmente no–holónoma y, por tanto, son un candidato natural para iniciar un estudio sistemático de otros espacios de distribuciones. La segunda parte de esta tesis pretende ser un texto de referencia para aquellos que quieran iniciarse en este área, explicando los avances recientes llevados a cabo por mis colaboradores y yo mismo.

Capítulo 1

Este capítulo sienta las bases de la teoría, dando definiciones y resultados elementales. Parte del material presentado aquí puede encontrarse también en los artículos [10, 15, 14]. Adicionalmente, es muy recomendable la lectura de [60], que condensa buena parte de la literatura aparecida antes del año 2000.

El resultado más importante en esta parte es la caracterización de la condición Engel para 2–planos que han sido trivializados usando un campo de líneas. Al igual que en topología de contacto se dice que "ser de contacto" es equivalente a girar consistentemente respecto a un campo de líneas legendriano, uno puede caracterizar "ser Engel" en términos de curvas en la 2–esfera. Entender esto es clave para construir y clasificar estructuras Engel.

Por otro lado, y por completitud, damos algunas construcciones (conocidas por algunos expertos) que muestran la rigidez de las estructuras Engel. Esto es manifiesta de dos maneras: existe móduli localmente y tienen muy pocos automorfismos genéricamente.

Adicionalmente, se introduce el concepto de *proyección de Geiges* para estudiar inmersiones de $\mathbb{S}^1$ tangentes a la estructura Engel. Aunque se trata de una idea sencilla, demuestra ser muy útil, pues transforma la manipulación de dichas inmersiones en un problema sobre curvas con cúspides en el plano cumpliendo una condición de área.

Capítulo 2

A partir de una 3–variedad de contacto podemos construir una variedad Engel a la que llamamos la *prolongación de Cartan*. En mi artículo [14] (y en este capítulo) se estudia el tipo de homotopía del espacio de prolongaciones. Es claro que esto no es exactamente un problema de topología Engel, sino un problema acerca de este espacio más restrictivo. En particular, el estudio de las prolongaciones está muy relacionado con el estudio de las estructuras de contacto subyacentes. Esto se manifiesta en el hecho de que es muy sencillo describir el espacio de prolongaciones obtenidas a partir de estructuras de contacto overtwisted. El resultado principal del capítulo dice precisamente esto.

Capítulos 3 y 4

Estos dos capítulos constituyen el grueso de la tesis, tanto en volumen como en relevancia. En ellos nos enfrentamos, finalmente, al problema de construcción y clasificación de estructuras Engel y probamos que hay un h–principio completo para estructuras Engel overtwisted. Buena parte del trabajo es precisamente entender cómo se manifiesta la flexibilidad en topología Engel y usar esto para dar una definición buena de overtwistedness.

Las técnicas usadas aparecieron por vez primera en [10], donde se probaba el h–principio de existencia. El teorema principal del artículo decía:

Theorem 3.1. *Sea M una 4–variedad, no necesariamente orientable. El siguiente mapa es sobreyectivo:*

$$\pi_k(i) : \pi_k(\mathfrak{Engel}(M)) \longrightarrow \pi_k(\mathcal{F}\mathfrak{Engel}(M)),$$

donde $\mathfrak{Engel}(M)$ denota el espacio de estructuras Engel y $\mathcal{F}\mathfrak{Engel}(M)$ es su análogo formal. En particular, toda estructura formal Engel puede ser homotopada a una estructura Engel.

En la tesis (Capítulo 3) ofrecemos diferentes estrategias de prueba para llegar a este resultado. Cada una de ellas proporciona una perspectiva diferente para manipular estructuras Engel. Mi intención es que esto permita al lector enfrentarse a problemas similares teniendo más herramientas a su disposición (además de las de [10]).

Los contenidos del Capítulo 4 no han aparecido todavía en forma de artículo. Basándonos en una de las nuevas estrategias de prueba dadas para Teorema 3.1, somos capaces de identificar qué ha de pedirse para que una estructura Engel sea overtwisted. Esto nos permite probar un h–principio completo (a nivel de componentes conexas) para ellas:

Theorem 4.22. *Sea M una 4–variedad, no necesariamente orientable. Sea K una variedad compacta.*

Cualesquiera dos familias $\mathcal{D}_0, \mathcal{D}_1 : K \to \mathfrak{Engel}(M)$ que sean loose y formalmente homotópas son homotópas a través de familias loose. En particular, el espacio de estructuras loose tiene las mismas componentes conexas que el espacio de estructuras Engel formales.

A estas estructuras las llamamos *loose* y no overtwisted porque resulta que la definición que tenemos depende de la dimensión de la familia que consideremos. Esto es: una familia de estructuras que son individualmente loose no tiene por qué ser loose ella misma. Este problema es puramente técnico y la expectativa es que podamos definir overtwistedness para obtener un h–principio en todos los π_k.

Capítulo 5

La cuestión que nuestro trabajo de clasificación deja abierta es si existen estructuras Engel que no sean loose/overtwisted. Es decir: ¿existen estructuras Engel más allá del h–principio? ¿Es posible que algunas de ellas tengan propiedades geométricas/topológicas interesantes que sólo puedan ser detectadas con técnicas nuevas?

La idea clave en este último capítulo es intentar utilizar subvariedades para tantear esta posibilidad. Para ello, tenemos que restringirnos a subvariedades adaptadas a la estructura Engel (por ejemplo, que sean tangentes o transversas). Como ocurre en topología de contacto, el primer resultado es negativo: si simplemente nos restringimos a variedades inmersas, no hay nada interesante:

Theorem 5.14. *Sea $(M, \mathcal{D})$ una variedad Engel. Sea $\mathcal{HI}^{\text{n.e.t.}}(\mathbb{S}^1, \mathcal{D})$ el espacio de inmersiones tangentes que no son órbitas del núcleo de $\mathcal{D}$. Sea $\mathcal{FHI}(\mathbb{S}^1, \mathcal{D})$ el espacio de inmersiones tangentes formales. La siguiente inclusión es una equivalencia de homotopía débil:*

$$\mathcal{HI}^{\text{n.e.t.}}(\mathbb{S}^1, \mathcal{D}) \to \mathcal{FHI}(\mathbb{S}^1, \mathcal{D}).$$

Theorem 5.22. *Sea $(M, \mathcal{D})$ una variedad Engel y sea V una variedad cualquiera. Mapas e inmersiones $f : V \to M$ transversas a $\mathcal{D}$ satisfacen el h–principio completo y C^0–denso.*

No es importante explicar qué es el núcleo de una estructura Engel. Basta con decir que se trata de un campo de líneas contenido en ella y definido de forma canónica. En particular, dentro del espacio de todas las inmersiones de $\mathbb{S}^1$ tangentes, sus órbitas cerradas son un espacio de codimensión de Hausdorff infinita.

Conclusiones

Aunque esta tesis es la conclusión de cuatro años de esfuerzo, los resultados que contiene son sobre todo un comienzo. Los Teoremas 3.1 y 4.22 responden a una pregunta que llevaba abierta 40 años: ¿satisfacen las estructuras Engel alguna forma de h–principio? (véase por ejemplo [20][Intrigue F2]). La respuesta es que sí, pero en lugar de zanjar por completo la cuestión, desterrándola al olvido matemático, esto nos abre nuevas puertas. La más importante es: ¿podemos construir una estructura Engel no overtwisted?

En el mundo de las estructuras de contacto, el primer ejemplo de tightness vino de la mano de Bennequin [3], que probó que había dos estructuras de contacto que eran homótopas formalmente pero no geométricamente. La forma en que lo hizo fue probando que para una de ellas y no para la otra existía un nudo tangente con ciertas propiedades. Esto motiva el trabajo que hemos presentado en el Capítulo 5. Actualmente estoy intentando entender variedades embebidas para reproducir un resultado de ese estilo.

Uno de los problemas fundamentales en este enfoque es que no existe un análogo de la estabilidad de Gray para estructuras Engel. Esto es, a diferencia de las estructuras de contacto, dos estructuras Engel homótopas no son necesariamente isótopas. Esto hace difícil la construcción de invariantes bajo homotopía Engel. La idea entonces, si queremos usar subvariedades para definir invariantes, es entender qué puede pasar en una homotopía de pares $(\mathcal{D}, N)$ con $\mathcal{D}$ una estructura Engel y N una subvariedad, donde N satisface alguna condición geométricamente relevante como ser tangente o transversa.

Uno de los corolarios del Teorema 3.1 dice que las 3–variedades transversas al núcleo de la estructura Engel no satisfacen la homotopy lifting property (¿propiedad de alzado homotópico?). Esto es, dada una homotopía de estructuras Engel $\mathcal{D}_s$, $s \in [0, 1]$, y una subvariedad N_0 transversa, no podemos en general definir un camino de subvariedades con N_s transversa al núcleo de $\mathcal{D}_s$. Esto implica que estos objetos no son buenos candidatos para definir invariantes. Dos de los proyectos en los que estoy involucrado tienen cómo objetivo entender si los nudos tangentes a la estructura Engel y las superficies transversas a ella dan lugar a invariantes.

Otra dirección de investigación interesante, a la que ya he apuntado antes, es la de estudiar otros espacios de distribuciones. Las estructuras Engel forman parte de una jerarquía de distribuciones en todas dimensiones llamadas *estructuras de Goursat*. No es inmediato reproducir los resultados que hemos presentado aquí en este contexto más general, y por ello se trata de un problema interesante. Una cuestión más profunda es si diferentes estructuras de Goursat interaccionan de alguna manera y, en particular, si interaccionan con las estructuras Engel.

The (hardly rigorous) preamble

Disclaimer: *The aim of this chapter is to serve as an introductory point for those not coming from contact topology. On the one hand, I try to introduce some of the language that I assume later on in the thesis and, on the other, I try to motivate the questions this thesis deals with. For a roadmap of the results proven in the thesis, I invite the reader to read the next chapter.*

Often, one is interested in understanding smooth manifolds endowed with structures that satisfy some geometric condition. The structures that we will care about will be described as sections of some bundle X over a base manifold M and, in local coordinates, the geometric condition that a section $s : M \to X$ must satisfy will be given as a **partial differential relation** (PDR) depending on s and its derivatives $s^{(k)}$.

In the cases of interest for us, the bundle X will have a natural action of $\mathrm{Diff}(M)$, lifting the obvious action on the base. This is the case, for instance, of the tangent bundle, the cotangent bundle, and any other bundle that arises by direct sum, tensor product, and (anti-)symmetrisation from them. We then require for the PDR to be invariant under the action of $\mathrm{Diff}(M)$; this invariance amounts to being able to give an intrinsic definition of the PDR, without resorting to local coordinates.

Many cases of interest fall within this framework: immersions, complex, symplectic, and contact structures, foliations... Once one such a structure is defined, some natural questions arise, for instance: *which manifolds admit structure "X"? What is the topology of the space of structures "X" over a given manifold M?* These are the sort of questions that we aim to answer in this thesis.

The h–principle

A PDR is nothing but some relation that a section and its (suitably defined) derivatives are asked to verify. For instance, when we say a map $f : N \to M$ it is an immersion, it simply means that its associated tangent map $Tf : TN \to TM$ is a monomorphism. Imagine we are given manifolds N and M and we ask the question: *is there any immersion from N to M?* Then, necessarily, we must be able to find a morphism $TN \to TM$ of rank $\dim(N)$ lifting some map $N \to M$. That is, we have decoupled f from its derivative Tf, and we have tried to solve this decoupled problem. If this cannot be achieved then, certainly, our PDR will have no solutions either. Given a PDR, we will say that a solution of the corresponding decoupled problem is a **formal solution**. Going back to our example of immersions, a formal solution is simply a pair $(f : N \to M, F : TN \to f^*TM)$ with F a monomorphism.

In many cases, of course, one is unable to transform a formal solution into an real one. Consider, for instance, the following problem: find a smooth function $f : \mathbb{R} \to \mathbb{R}$ that is strictly increasing and agrees with $x \to x$ in $(-\infty, 0]$ and with $x \to x - 1$ in $[1, \infty)$. Bolzano's theorem readily implies that no solution exists, since $f(0) = f(1) = 0$ and yet, there is a formal solution given by $(f, 1)$, with f any function satisfying the boundary conditions and 1 its constant *"formal"* derivative.

This very simple example might discourage the reader: maybe there are not many meaningful problems in which one is able to obtain useful information from the space of formal solutions. The following theorem of Whitney [86] says otherwise:

Theorem 0.1. *Immersions $\mathbb{S}^1 \to \mathbb{R}^2$ are classified by their rotation number.*

Given an immersion $f : \mathbb{S}^1 \to \mathbb{R}^2$, its rotation is nothing but the degree of its Gauss map $\frac{\dot{f}}{|\dot{f}|} : \mathbb{S}^1 \to \mathbb{S}^1$. The theorem is saying that, in order to show that two immersions are homotopic (through immersions), it suffices to take their derivatives and check that they are homotopic as maps $\mathbb{S}^1 \to \mathbb{R}^2 \setminus \{0\}$ (this is what having the same rotation number means). That is, we just have to check that they are homotopic as formal solutions!

Theorem 0.1 can be phrased by saying that the space of solutions of the PDR describing immersions has the same connected components as the space of formal solutions. Thinking in these terms it can readily be seen that the existence and classification questions for geometric structures that we posed before are essentially questions about the action of the natural inclusion

$$i : \{\text{Solutions of the PDR}\} \to \{\text{Formal solutions of the PDR}\}$$

at the level of homotopy groups. For instance, showing that the map

$$\pi_0(i) : \pi_0\{\text{Solutions of the PDR}\} \to \pi_0\{\text{Formal solutions of the PDR}\}$$

is surjective means that any formal solution is homotopic to an honest one. Similarly, saying that it is injective means that any two solutions that are formally homotopic are homotopic through actual solutions.

Without being too precise about the concepts involved, we are ready to define the essential concept that this thesis deals with:

Definition 0.2. *For a given partial differential relation, consider the inclusion*

$$i : \{\textit{Solutions of the PDR}\} \to \{\textit{Formal solutions of the PDR}\}$$

The PDR is said to satisfy the ***existence*** *h–**principle** if $\pi_0(i)$ is surjective. It satisfies the* ***complete*** *h–principle if i is a (weak) homotopy equivalence.*

The h–principle (where the h stands for *homotopy*) is, in some sense, a *philosophy.* On the one hand, it encompasses many techniques that can be regarded as "standard" and that allow the user to show that, for certain families of structures, it holds. Because of this, the interesting problems in the area are those that lie beyond those standard techniques. At the same time, there are usually a number of big picture strategies that apply in a wide variety of settings even if the specifics change.

We will introduce the techniques we need as we go along, instead of reviewing them here. The point of this is making the text flow slightly better, without making the reader go back to check for a reference.

Distributions

Given a smooth manifold, its tangent space is a vector bundle. That is, at every point the tangent fibre has the structure of a real vector space. Then, we define a (tangent) distribution to be a smooth choice of linear subspace (of constant dimension) at every point. Equivalently, it is just a smooth section of the Grassmann fibration $\mathrm{Gr}(TM,k) \to M$.

The physical motivation of such an object is quite immediate. If we think of our manifold M as the set of possible positions/states, the directions contained in a distribution $\xi : M \to \mathrm{Gr}(TM,k)$ can be thought of as "*allowed directions of motion*" for an object living in M. As such, one is led to consider curves that are tangent to the distribution, since those are the possible trajectories an object in our mechanical system can take. Then, a meaningful question one can pose is: "*given two points, is there an admissible trajectory connecting them?*". A celebrated theorem of Chow [11] says that this is always possible for a certain class of distributions:

Theorem 0.3. *Let (M,ξ) be a manifold endowed with a distribution ξ. Assume that ξ is bracket–generating. Then, for any two points $p,q \in M$, there is a map $\gamma : [0,1] \to M$ with $\gamma(0) = p$, $\gamma(1) = q$, $\gamma'(t) \in \xi_{\gamma(t)}$. That is, there is a trajectory tangent to ξ connecting p with q.* [2]

So, what does **bracket–generating** mean? Some of the readers might know this notion under the name of being **non–holonomic**. Essentially, it boils down to the following: at every point $p \in M$ one must be able to realise any possible direction in T_pM by infinitesimal motions tangent to ξ. More rigorously:

Definition 0.4. *Let (M,ξ) be a manifold endowed with a distribution ξ. Then, the sequence:*

$$\xi_0 = \xi \subset \xi_1 = [\xi_0,\xi_0] \subset \xi_2 = [\xi_1,\xi_0] \subset \cdots \subset \xi_l = [\xi_{l-1},\xi_0]$$

stabilises, i.e. $\xi_{l+j} = \xi_l$ for some l and all $j \geq 0$. If $\xi_l = TM$, we say that ξ is bracket–generating.

Maybe this definition confuses more than it explains. ξ_1 is defined by taking Lie brackets of pairs of vector fields tangent to ξ. Observe that, in general, this is not a distribution anymore; ξ_1 is still given as a family of linear subspaces, but now its rank might vary with the point (we could say that ξ_1 is then a sheaf) [3]. This is, however, not a problem, one can still consider vector fields that are tangent to ξ_1 and take Lie brackets. The bracket–generating condition means that, by iterated Lie bracket, we are able to obtain any direction of motion. Then, Chow's theorem is probably not very surprising! It is in some sense a local to global statement: since infinitesimally we can go wherever we want, the same holds true at large scales.

Let us go back to the h–principle terminology we introduced before. Chow's theorem says something like this: for $p,q \in M$, the space of maps

$$\Omega_{p,q}(\xi) = \{\gamma : [0,1] \to M | \gamma(0) = p, \gamma(1) = q, \gamma'(t) \in \xi_{\gamma(t)}\}$$

is non–empty if ξ is bracket–generating. Of course now the natural question is whether it has several connected components and what the topology of each one of them is. At the same time, we can define a formal counterpart for $\Omega_{p,q}(\xi)$. In view of our example about immersions, it is clear what it should be:

$$\mathcal{F}\Omega_{p,q}(\xi) = \{(\gamma : [0,1] \to M, F : [0,1] \to \xi \setminus \{0\}) | \gamma(0) = p, \gamma(1) = q, F(t) \in \xi_{\gamma(t)}\}.$$

[2] This statement is not an if and only if. Being bracket–generating is a sufficient (but not necessary) condition for the conclusion to hold. A counterexample is given by a transitive confoliation that is a foliation in an open set (we invite the reader to look up the beautiful theory of confoliations [21, 84]).

[3] In this thesis we will be interested in distributions whose iterated Lie brackets always produce distributions and not just sheaves.

That is, we decouple the equation $\gamma'(t) \in \xi_{\gamma(t)}$ by replacing γ' by an independent map F. The topology of $\mathcal{F}\Omega_{p,q}(\xi)$ is particularly simple. Firstly, the maps from p to q can be (non–canonically) identified with the loop space of M. Secondly, since $\gamma^*\xi$ is a trivial bundle, F is just a map of the interval into $\mathbb{R}^{\mathrm{rank}(\xi)} \setminus \{0\}$ whose ends are not fixed; therefore, the space of all the possible Fs is contractible! Note that the homotopy groups of the loop space of M are not trivial to compute in many cases, but at least it is not anymore a question about geometry, but a question about algebraic topology.

Having (somewhat) understood $\mathcal{F}\Omega_{p,q}(\xi)$, we pose the question: if ξ is bracket–generating, what is the homotopy theoretical nature of the inclusion

$$i : \Omega_{p,q}(\xi) \to \mathcal{F}\Omega_{p,q}(\xi)?$$

Chow's theorem actually is stronger than what we stated above; it can be readily seen that it actually produces a path in each homotopy class connecting the two given points: that is, i is surjective at the level of connected components. However, π_0–injectivity and the higher homotopy groups are much harder to address. In fact, it is known [7] that π_0–injectivity fails in general. In Part II, Chapter 5, we will study this question for a specific family of distributions: Engel structures.

Studying spaces of distributions

We have introduced bracket–generating distributions as the natural class of distributions where Chow's theorem holds. Being bracket–generating (up to a given order) is itself an instance of a PDR, since the Lie bracket is a first order linear operator. It is therefore natural to ask whether the h–principle holds for meaningful classes of bracket–generating distributions (as a subset of the space of all distributions of a given rank in a given ambient space). Since line fields can never be bracket–generating (as soon as the ambient is not a 1–dimensional manifold....), let us jump to the first non–trivial example.

Contact structures (in dimension 3)

Considering 2–distributions in 3–manifolds that are bracket–generating in one step yields the following definition:

Definition 0.5. *Let N be a 3–dimensional manifold. A 2–dimensional distribution $\xi \subset TN$ is said to be a* ***contact structure*** *if it is everywhere non–integrable. That is, $[\xi, \xi] = TN$.*

A contact structure ξ can always be given locally as $\ker(\alpha) = \xi$, with α a 1–form. Then, the non–integrability condition reads as $\alpha \wedge d\alpha \neq 0$. If ξ has a global defining equation α, it is said to be *coorientable*.

Contact structures have a rich mathematical history. They appear as natural boundary conditions of *symplectic* manifolds and, as such, both notions are often studied together. More important to this introduction is the role the h–principle plays in their study. The following is a remarkable theorem of Gromov:

Theorem 0.6. *The complete h–principle holds for contact structures if M is open.*

This statement is a particular instance of a much more general phenomenon. Gromov [33] actually proved in his thesis that many PDR's satisfy the h–principle if the ambient manifold is open. Consider the following baby example: suppose we wanted to build a function in M without critical points. Then, what we can do is build a exhaustion $M_0 \subset M_1 \subset \cdots$ of M by compact sets and a Morse function $f : M \to \mathbb{R}$. We can then iteratively push the critical points of f by an isotopy so that, at step i, there are no critical points of f in M_i. Taking the limit yields the desired function. Effectively, we have *"pushed our problems to infinity"*. [4]

What about the case where the manifold is closed? In the 70's, work of Lutz and Martinet showed that any closed 3–manifold admits a contact structure, so the existence h–principle is true. The proof used four ingredients: any closed 3–manifold can be obtained from the 3–sphere by performing surgery on a link, there is a canonical contact structure in the 3–sphere, links can be homotoped to be *transverse* to a contact structure, transverse links have a well understood local model and one can explicitly give a formula to extend the contact structure after the surgery.

A particular instance of this construction is the following: one can start with a transverse knot and perform a surgery that is *topologically* trivial but not *contactly* trivial (a priori). That is, we cut out a torus and we glue in back in exactly the same way, but we make the contact structure describe an extra *turn* radially. This is called a **Lutz twist**; we will review it in detail in Part I, Chapter 2, it is not important how it works for the discussion now.

The punchline is that this provided a method of constructing not only contact structures in every 3–manifold, but also of constructing new contact structures in $\mathbb{S}^3$. Not only that, but a certain version of the Lutz twist (the so–called *full Lutz twist*), preserves the homotopy class of the contact structure as a distribution. This raises the question: *are any of these not homotopic as contact structures?* An answer came in 1983 with work of Bennequin: the standard structure and the one obtained from it by performing full Lutz twist along the unknot are genuinely different!

At that point, one could have hoped that maybe one can keep performing Lutz twists along different unknots to yield new contact structures that are not homotopic to one another. This is not the case:

Theorem 0.7. (Eliashberg, 1989) *The following inclusion is a weak homotopy equivalence:*

$$\mathcal{C}\text{-}\mathcal{S}\mathfrak{trs}_{OT}(M,\Delta) \to \mathcal{P}\mathfrak{lanes}(M,\Delta).$$

What does this mean? Let us setup some notation. Suppose we perform a Lutz twist in a tubular neighbourhood of a knot γ, and we take one of the normal disks of γ in such a neighbourhood. Such a disc, endowed with the germ of contact structure around it, is called an **overtwisted disc**. A solid torus in which we have performed a Lutz twist is then foliated by an $\mathbb{S}^1$–family of overtwisted disks. $\mathcal{C}\text{-}\mathcal{S}\mathfrak{trs}_{OT}(M,\Delta)$ denotes the space of contact structures that have a fixed overtwisted disk Δ. $\mathcal{P}\mathfrak{lanes}(M,\Delta)$ are the 2–distributions that are tangent to Δ at its origin. The theorem says (in particular): if a contact structure ξ contains an overtwisted disk, any other contact structure obtained from it by performing full Lutz twist is homotopic to ξ!

All of these facts brought together explain why contact topology has become a vibrant field. On the one hand, there is a class of contact structures, the ones containing an overtwisted disc, that are flexible and, thus, easy to understand thanks to the h–principle. On the other hand, thanks to Bennequin's result, we know that not all contact structures are of this form! We call the former **overtwisted**, and the latter **tight**. Most of the work in the area for the last 30 years has focused on developing techniques for classifying tight contact structures. We will not explore this angle in this thesis.

[4] I was reminded of this very nice example by a talk of T. Vogel.

Even–contact structures

Having covered dimension 3, we can consider bracket–generating distributions in 4–manifolds. Part II of this thesis actually deals with the very next case: rank–2 distributions in dimension 4 that are bracket–generating in two steps. These are called **Engel structures**. However, it turns out that rank–3 distributions are much easier to understand:

Definition 0.8. *Let M be a 4–dimensional manifold. A 3–dimensional distribution $\mathcal{E} \subset TM$ is said to be an even–contact structure if it is everywhere non–integrable, i.e. if $[\mathcal{E}, \mathcal{E}] = TM$.*

As above, one can consider a (possibly local) defining equation α for $\mathcal{E} = \ker(\alpha)$. The even–contact condition amounts to asking $\alpha \wedge d\alpha \neq 0$. We also ask for $\mathcal{E}$ to be coorientable.

At this point maybe one could hope for the theory of even–contact structures to be as rich as the theory of contact ones. However, we are in for a disappointment:

Theorem 0.9. (McDuff) *Even–contact structures satisfy the complete h–principle.*

This result is actually a consequence of previous work of Gromov about partial differential relations that satisfy a property called *ampleness*: if they do, the complete h–principle holds even in closed manifolds. Roughly speaking, the idea behind ampleness is that the PDR is defined by an inequality that fails in large codimension.

This does not mean that even–contact structures are completely uninteresting. Even if the classification question is completely understood, there are still things to be studied from the point of view of dynamics. Even–contact structures contain a canonical line field called the **kernel**: very little is known about its dynamical properties. Even–contact structures will reappear again in Part I and in Part II.

Remark: Contact structures (respectively, even–contact structures) can be defined in all odd (resp. even) dimensions. A proper definition will be given later on, but it is irrelevant for our purposes now.

Topologically stable distributions

We have already (somewhat) motivated our interest in Engel structures: after contact structures (in dimension 3), they conform the next simplest instance of a bracket–generating distribution. There is one result of Cartan that is often quoted to further support the idea that Engel structures are worth studying:

Theorem 0.10. *There are four classes of* ***topologically stable*** *distributions: line fields, contact structures, even–contact structures, and Engel structures.*

Topological stability means two things. First, we require that the class is an *open subset* of the space of distributions (of a given rank and dimension). Secondly, we require that all of the structures of the class are *locally diffeomorphic* (meaning that, given any two points, we can find a diffeomorphism of their neighbourhoods identifying the distributions).

The second condition means in particular that, for topologically stable distributions, there is nothing to be studied locally: there cannot possibly be local invariants, since any two points look the same. This makes them good subjects of study in *topology*, since we are forced to look for global invariants. As we explained above, contact structures have interesting invariants, but even–contact ones do not. The motivating question for the second part of this thesis is: *what about Engel structures?* (Spoiler: we do not know yet.)

Foliations

Instead of focusing on bracket–generating distributions, we can consider the completely opposite assumption: distributions satisfying $[\xi, \xi] = \xi$. By Frobenius theorem, this is equivalent to the fact that, through every point, we can find a little disc of dimension $\mathrm{rank}(\xi)$ that is everywhere tangent to ξ. Even more: all these discs glue with one another to decompose the ambient manifold in a collection of (possibly open) manifolds of dimension $\mathrm{rank}(\xi)$. This we call a **foliation** and each of the lower dimensional submanifolds is called a **leaf.**

We have defined foliations as a particular class of distributions satisfying a certain PDR. As such, we can pose the existence and classification questions we started this preamble with. It turns out that the existence h–principle holds for foliations in all dimensions:

Theorem 0.11. (Thurston, '73) *Let M be a closed orientable n–manifold. Then, M admits a foliation in any given homotopy class of plane fields.*

In dimension 3, constructing these foliations is achieved by paying a certain price: all of them have many *Reeb components*, that is, solid tori whose boundary is a leaf of the foliation and whose interior is foliated by planes that are asymptotic at infinity to the torus leaf. Reeb components play a similar role to that of Lutz tubes in a contact structure: they introduce flexibility.

Looking at Thurston's result, one can wonder whether foliations without Reeb components behave flexibly as well. However, at that point, it was already known by work of Novikov that **taut** foliations (a notion closely related to having no Reeb components) display rigidity:

Theorem 0.12. (Novikov, '65) *Let M be a closed orientable 3–manifold and let $\mathcal{F}$ be a* ***taut*** *foliation by surfaces. Then:*

- *any map $\gamma : \mathbb{S}^1 \to M$ everywhere transverse to $\mathcal{F}$ represents a non–trivial element of $\pi_1(M)$,*
- *any leaf $\mathcal{L}$ π_1–injects into M,*
- *M is irreducible.*

We will define tautness with precision in Part I. The main question that motivates the work in Part I of this thesis is: *what is a good analogue of tautness in higher dimensions?* In particular, knowing that foliations without further restrictions seem to be very flexible in light of Thurston's result, what is a reasonable condition that guarantees an interesting behaviour that is still topological in nature?

So what is this thesis about?

Simply put: this thesis is about distributions and their global topology.

Guide to the contents of the thesis

Much of what is contained in the thesis can be found in the articles I have written during my PhD. Some of them have been reproduced almost literally, cutting down the fat and streamlining some of the exposition, and some others have been fully reworked and expanded. Many of them shared a certain amount of preliminary material, which I have collected in the first chapters of Parts I and II, respectively.

The idea behind this guide is to have all the main theorems in one place, so that the reader can navigate the thesis a bit better. In agreement with these nice intentions, I have also tried, hopefully successfully, to make this chapter self–contained (assuming the reader is familiar with the contents of the Preamble).

List of articles

The articles this thesis borrows from are the following:

[9] R. Casals, A. del Pino, and F Presas. "h-Principle for Contact Foliations". *Int. Math. Res. Not.* 20 (2015), pp. 10176–10207.

[10] R. Casals, J.L. Pérez, A. del Pino, and F Presas. "Existence h-Principle for Engel structures". *arXiv e-prints* (2015). arXiv: 1507.05342.

[14] A. del Pino. "On the classification of prolongations up to Engel homotopy". *Submitted* (2016).

[15] A. del Pino and F. Presas. "Flexibility for tangent and transverse immersions in Engel manifolds". *arXiv e-prints* (2016). arXiv: 1609.09306.

[16] A. del Pino and F. Presas. "The foliated Weinstein conjecture". *To appear in Int. Math. Res. Not.* (2015). arXiv: 1509.05268.

[53] D. Martínez Torres, A. del Pino, and F. Presas. "Transverse geometry of foliations calibrated by non-degenerate closed 2-forms". *To appear in Nagoya Mathematical Journal* (2014). arXiv: 1410.3043.

[68] D. Peralta–Salas, A. del Pino, and F. Presas. "Foliated vector fields without periodic orbits". *To appear in Israel Journal of Mathematics* (2014). arXiv: 1412.0123.

Part I

The first half of this thesis deals with contact and symplectic foliations. These are simply foliations that possess a leafwise contact or symplectic structure that is varying smoothly. The expectation/hope is that the rich behaviour seen in 3–dimensional foliation theory should emerge in higher dimensions by restricting to a suitable class of foliations. Perhaps because of bias, we propose that contact/symplectic foliations conform a reasonable candidate for this.

Chapter 1

In this chapter we will go over the basic definitions and constructions, many of which appeared already in [9]. In some sense, this reduces to adapting many standard constructions arising in contact and symplectic topology to the foliated setting. Examples include the analogues of the space of contact elements, the symplectisation, the Boothby–Wang construction, or the Lutz twist.

Chapter 2

The second chapter of Part I goes over the main result of [9], although the argument has been rewritten completely in a more streamlined fashion. It reads as follows:

Theorem 2.1. *Let V be a 4–manifold and let $\mathcal{F}$ be a foliation of rank 3. Then, the inclusion:*

$$\mathfrak{ContFol}(V, \mathcal{F}) \to \mathcal{F}\mathfrak{ContFol}(V, \mathcal{F})$$

of the contact foliations into the formal contact foliations induces a surjection in homotopy groups.

Let us elaborate on this. As seen in the Preamble, overtwisted contact structures present a flexible behaviour (meaning that the only obstructions to their existence are of algebraic topological nature) and are thus not well–suited to be used as tools in the topological study of manifolds. This is not the case of *tight* contact structures: their rich theory interacts heavily with the topology of the ambient manifold (there are many such results in dimension 3, but this is still largely unexplored in higher dimensions). Theorem 2.1 says that a similar phenomenon can be observed in the foliated setting: given a foliation, it is easy to endow it with a leafwise contact structure, but it turns out that the method by which this is done makes all leaves overtwisted. There are still no examples showing that the existence of a contact foliation with all leaves tight has interesting topological consequences for the ambient manifold; this is a current subject of research by the author and others.

It is worth pointing out that after [9] was published, it was proven in [6] that a complete h–principle holds for contact foliations of any (odd) rank and any dimension where a system of transverse overtwisted disks has been fixed. In particular, the results in [6] recover Theorem 2.1.

Chapter 3

This chapter corresponds to my article[53]. The main result is the analogue of the Lefschetz hyperplane theorem for (strong) symplectic foliations, namely:

Theorem 3.5. *Let $(M^{2n+1}, \mathcal{F}, \omega)$ be a strong symplectic foliation on a closed manifold. Let W be a Donaldson divisor of dimension $2n-1$. Then, for every leaf $\mathcal{L}$ of $\mathcal{F}$ it holds that:*

$$\pi_k(\mathcal{L}, \mathcal{L} \cap W) = \{1\},\ 0 \leq k \leq n-1.$$

For this statement to be understood, we should briefly explain what strongness means in this context. A foliation is symplectic if there exists a global 2–form ω that is symplectic on each leaf. The observation is that this implies that ω is closed leafwise, but not necessarily globally. If it is closed, we say that the foliation is *strong* symplectic.

From a symplectic perspective, strongness is a natural condition, since many constructions require for ω to define a cohomology class. This is the case for the results presented in this chapter. The idea is as follows: if ω is a closed 2–form defining an integral cohomology class, we can consider the associated complex line bundle $L \to M$. Without going into details, one is able to construct sections of L whose zeroes are submanifolds that inherit a symplectic foliation structure from the ambient space; these are the so–called Donaldson divisors. The result can then be understood by saying that the leaves of these submanifolds recall much of the topology of the original leaves. In particular, the leaf spaces of M and its Donaldson divisors (down to dimension 3) are isomorphic.

One of the key observations is that a 3–dimensional strong symplectic foliation is exactly the same as a taut foliation. The result is then saying that the leaf space of any strong symplectic foliation is homeomorphic to the one arising from a taut foliation (which have been deeply studied in 3–dimensional topology).

Chapters 4 and 5

These chapters correspond to the articles [68, 16], whose contents developed simultaneously. Let us explain.

A natural question one can pose is the following: is there any relation between the topology of a manifold and the dynamics of the non–vanishing vector fields on it? In particular, are there topological conditions forcing all vector fields to have closed orbits? By work of Wilson and K. Kuperberg, it is known that the answer to the latter question is *no*: there are vector fields with no closed orbits in all manifolds of dimension at least 3.

The main theorem in [68], which we explain in Chapter 4, generalises this result to families of vector fields of all dimensions:

Theorem 4.1. *Let $(M^{n+m}, \mathcal{F}^n)$, $n \geq 3$, be a foliated manifold. Denote by $\mathfrak{X}_{ns}(M, \mathcal{F})$ the space of non–singular vector fields tangent to $\mathcal{F}$. Denote by $\mathfrak{X}_{no}(M, \mathcal{F})$ the subspace of those with no closed orbits. The inclusion:*

$$\iota_n : \mathfrak{X}_{no}(M, \mathcal{F}) \to \mathfrak{X}_{ns}(M, \mathcal{F})$$

is a weak homotopy equivalence.

However, it is also known that certain classes of vector fields always have closed orbits. We need some additional explanations before we give a precise statement. A contact manifold can be understood as a generalisation of the notion of energy level in Hamiltonian dynamics. In particular, it has a set of associated flows (the *Reeb flows*), that play the role of the Hamiltonian flow describing the motions in the system. It was conjectured by Weinstein that all Reeb flows possess a closed orbit. Much later, it was proven by Taubes that this is the case in dimension 3; it has also been proven in several higher dimensional instances.

In Chapter 5, corresponding to my article [16], we study this very same question in the foliated setting:

Theorem 5.1. *Let $(M^{3+m}, \mathcal{F}^3, \xi^2)$ be a contact foliation in a closed manifold M. Let α be a defining 1–form for an extension of ξ and let R be its Reeb vector field. Let $\mathcal{L}^3 \hookrightarrow M$ be a leaf.*

i. If $(\mathcal{L}, \xi|_{\mathcal{L}})$ is an overtwisted contact manifold, R possesses a closed orbit in the closure of $\mathcal{L}$.

ii. If $\pi_2(\mathcal{L}) \neq 0$, R possesses a closed orbit in the closure of $\mathcal{L}$.

The techniques used are based on Hofer's approach to prove the Weinstein conjecture for overtwisted contact manifolds. In particular, the main ingredient is the introduction of pseudoholomorphic curves to the foliated setting; this material is in some sense, foundational. Some of the most interesting contents of this chapter are the examples, where we show that the assumptions of Theorem 5.1 are actually sharp. In particular, the Weinstein conjecture is not true for foliations with all leaves tight.

Part II

The second half of this thesis deals with Engel structures. It turns out that, apart from the classes we have discussed (contact, even–contact, and foliations), the global topology of other spaces of distributions has not really been explored that much. We do not know if they meaningfully interact with the topology of the ambient manifold they live in, or if they interact with other geometrical structures.

Engel structures, being the next simplest example of bracket–generating distribution to consider, are therefore a great candidate to start our programme with. What is proven about them in this thesis?

Chapter 1

In this chapter we go over the main definitions and basic results about Engel structures; some of the material can be found in [10, 15, 14] and, for an alternate view, there is the excellent survey [60] by Montgomery.

The emphasis is put on characterising the Engel condition when a plane field has been trivialised by applying the flowbox theorem to a line field contained in it. In 3–dimensional topology we usually say that the contact condition amounts to "turning" with respect to a legendrian line field, and what we do here is analogous. This innocuous result is actually key to all the subsequent manipulations we do to construct and classify Engel structures in Chapters 3 and 4.

We also present results that are quite similar to some of the constructions in [60]. In particular, we show that Engel structures are very rigid. This manifests in two forms: they have local moduli and they have very few automorphisms (generically).

Somewhat separately, we explain the basics in the study of immersions of $\mathbb{S}^1$ tangent to Engel structures. In particular, we define the concept of *Geiges projection*, which transforms the problem into a question about planar curves with cusps satisfying an area constraint.

Chapter 2

There is an operation, called *Cartan prolongation*, that assigns an Engel manifold to each contact 3–manifold. The contents of this chapter reproduce the first half of [14], which is dedicated to the study of this particular family of Engel structures. As such, we are not strictly studying Engel topology, but just the space of Cartan prolongations in itself.

A major difference with more general Engel structures is that the study of prolongations is very intimately related to the study of the underlying contact structures. In particular, just like the overtwisted contact structures present a flexible behaviour, the associated prolongations do too; this is exactly what the main theorem of the chapter says.

Chapter 3

This chapter and the next constitute the main body of the thesis. The results presented in Chapter 3 are based on the article [10], but the material has been almost completely rewritten and largely expanded. The main theorem reads as follows:

Theorem 3.1. *Let M be a smooth 4–manifold, not necessarily orientable. Then, the map*

$$\pi_k(i) : \pi_k(\mathfrak{Engel}(M)) \longrightarrow \pi_k(\mathcal{F}\mathfrak{Engel}(M))$$

is surjective for every $k \geq 0$. In particular, every formal Engel structure is homotopic to the flag of a genuine Engel structure.

The additional content compared to [10] is that several alternate approaches to the proof are provided, each of them with a slightly different flavour. The hope of the author is that this will provide a bigger toolbox for those who want to deal with Engel structures.

Chapter 4

The material in this chapter is original and has not yet appeared in article form. It uses one of the new proofs of Theorem 3.1 to show that the existence h–principle can be improved to a complete h–principle (in π_0) for a suitably defined class of overtwisted Engel structures. The chapter begins introducing several notions of looseness that are then shown to be equivalent. Then, the main theorem is proven:

Theorem 4.22. *Let M be a smooth 4–manifold. Let K be a compact manifold. Let $\mathcal{D}_0, \mathcal{D}_1 : K \to \mathfrak{Engel}(M)$ be two formally homotopic loose families. Then, they are Engel homotopic through loose families.*

In particular, the class of loose Engel structures has the same connected components as the class of formal Engel structures.

In particular, we will say that a loose Engel structure is (0–)overtwisted, because it allows us to prove the h–principle at the level of π_0. Upgrading this to a full h–principle (in all π_k) is work in progress.

Chapter 5

The material in this chapter appeared originally in [15]. Transverse and tangent immersions to Engel manifolds are shown to abide to the h–principle (for immersions, it is shown that a set of finite Hausdorff measure has to be discarded for this to be true). Namely:

Theorem 5.14. *Let $(M, \mathcal{D})$ be an Engel manifold. Let $\mathcal{HI}^{\text{n.e.t.}}(\mathbb{S}^1, \mathcal{D})$ be the space of the horizontal immersions that are not orbits of the kernel. Let $\mathcal{FHI}(\mathbb{S}^1, \mathcal{D})$ be the space of formal horizontal immersions.*

The inclusion

$$\mathcal{HI}^{\text{n.e.t.}}(\mathbb{S}^1, \mathcal{D}) \to \mathcal{FHI}(\mathbb{S}^1, \mathcal{D})$$

is a weak homotopy equivalence.

Theorem 5.22. *Let $(M, \mathcal{D})$ be an Engel manifold and V be any manifold. Maps and immersions $f : V \to M$ transverse to $\mathcal{D}$ satisfy a C^0–close, parametric, relative, and relative to the parameter h–principle.*

Part I

Symplectic and contact foliations

Chapter 1

Basics on symplectic and contact foliations

Disclaimer: we will henceforth assume that our manifolds/distributions are smooth unless we explicitly say otherwise. We will additionally assume that they are orientable. This is done for simplicity.

In Part I of this thesis we will study foliations. In the Preamble we claimed that there is a well–understood dichotomy for foliations by surfaces in dimension 3. On the one hand, we have the foliations with Reeb components produced by Thurston's theorem. These are a paradigm of flexibility. On the other hand, Novikov's result says that taut foliations impose topological constraints on the manifolds they live in; in particular, not every manifold admits one.

We shall not explore this any further in this thesis, but the theory of 3–dimensional taut foliations is possibly one of the better understood instances of interaction between a geometric object (the taut foliation) and the topology of the ambient manifold. Indeed, there is a wealth of results, starting with the work of Gabai [24, 25, 26], relating taut foliations to several flavours/ideas of 3–dimensional topology (knot theory, the geometrisation conjecture...)

We have delayed introducing the definition long enough:

Definition 1.1. *Let M be a closed n–dimensional manifold. A codimension–1 foliation $\mathcal{F}$ is said to be* ***taut*** *if any of the following equivalent properties hold:*

a. for each point $p \in M$, there exists a closed curve $\gamma : \mathbb{S}^1 \to M$ that is transverse to $\mathcal{F}$ and satisfies $\gamma(1) = p$,

b. there exists a metric g in M making all the leaves of $\mathcal{F}$ minimal hypersurfaces,

c. there exists a closed codimension–1 form Ω in M that is volume on the leaves of $\mathcal{F}$.

Equivalence between all these three notions was shown by Rummler and Sullivan [76, 75]. Property (a.) is topological in nature, whereas the other two are more geometrical. It is worth remarking that their equivalence relies on M being closed. For open manifolds, Property (a.) is stronger than the other two.

Since there is such a rich theory lurking in the background in the 3–dimensional case, one could hope for the same to be true in higher dimensions. However, a recent result of Meigniez shows that the 3–dimensional case and the higher dimensional cases are markedly different:

Theorem 1.2. (Meigniez, '12) *Let M be a closed n–manifold, $n > 3$. Then, M admits a minimal (and in particular, taut) foliation in every homotopy class of hyperplane fields and Haefliger structures.*

That is, Thurston's theorem holds in higher dimensions even if we add the minimality hypothesis. This is disappointing, so we would rather find an alternate way of defining tautness. A way that is more topologically flavoured.

In this chapter, we will define two classes of foliations: *(strong) symplectic foliations* (Section 1.1) and *contact foliations* (Section 1.2). In dimension 3, the strong symplectic foliations are precisely the taut ones, and we will wave our hands and argue that, perhaps, this is a more natural generalisation of tautness. Much like contact manifolds arise naturally in symplectic topology, contact foliations relate naturally to strong symplectic foliations. A bold claim (just like the one before) would be that contact foliations with all leaves tight conform a class of foliations whose behaviour is interesting.

The main theme of Part I is that techniques arising in contact and symplectic topology transfer nicely to this foliated setting. Indeed, contact and symplectic foliations can be understood as generalisations of parametric families of contact/symplectic structures, so the phenomena taking place in the latter must occur in the former.

1.1 Symplectic foliations

1.1.1 Definitions

Let us start with the following definition.

Definition 1.3. *Let M^{2n+m} be a smooth manifold. Let $\mathcal{F}$ be a rank $2n$ foliation. Let $\omega \in \Omega^2(\mathcal{F}) = \Lambda^2(T^*\mathcal{F})$ be a 2–form along the foliation. The triple $(M, \mathcal{F}, \omega)$ is called a* ***symplectic foliation*** *if, for every leaf $i_\mathcal{L} : \mathcal{L} \to M$, $(\mathcal{L}, i_\mathcal{L}^*\omega)$ is a symplectic manifold.*

$(M, \mathcal{F}, \omega)$ is said to be ***strong*** *if there exists a global closed 2–form $\tilde{\omega}$ satisfying $\tilde{\omega}|_\mathcal{F} = \omega$. The form $\tilde{\omega}$ is said to be an* ***extension*** *of ω.*

Remark 1.4. Often we will abuse notation and simply say that $(M, \mathcal{F}, \tilde{\omega})$ is a strong symplectic foliation.

We will sometimes say that a symplectic foliation is **weak** to highlight the fact that it is not strong. Weak symplectic foliations are also referred to as *regular Poisson structures* in the literature, see [35], and can be understood as a generalisation of *symplectic fibrations*, see [56, Chapter 6]. Strong symplectic foliations are referred to as 2–**calibrated structures** in [43], where they were introduced first.

A moment of thought shows that, in dimension 3, any foliation by surfaces is symplectic, and tautness amounts to being strong, so we can legitimately call this a generalisation. It is also worth noting that, in higher dimensions, any strong symplectic foliation (of codimension–1) is taut, since $\tilde{\omega}^n$ is a leafwise volume form that is globally closed.

A symplectic foliation being strong is meaningful also in the sense that many constructions from symplectic topology are likely to extend to the foliated setting only under this hypothesis. This is the case of the approximately holomorphic techniques (see Chapter 3): in the strong, codimension–1 case they have been shown to work and they offer a fruitful approach for understanding the leaf space.

Another example will be shown in Chapter 5. Pseudoholomorphic curve techniques can be used in the strong setting because energy bounds retain their cohomological nature, as in the non–foliated case. However, in the weak setting this is not true anymore.

1.1.2 Formal symplectic foliations

A question that we will not tackle in this thesis but that is of clear interest (as our discussion so far has shown), is whether we can construct and classify strong/weak symplectic foliations. It is convenient to phrase this problem in terms of the h–principle.

First, it is worth noting that there are two related but distinct problems one can pose. The first is as follows: we fix an ambient manifold M and a foliation $\mathcal{F}$ and then we ask whether there exists a symplectic foliation of the form $(M, \mathcal{F}, \omega)$ (possibly with a extension $\tilde{\omega}$). From this perspective, we are interested in the foliation itself and maybe we expect the auxiliary symplectic data to shed some light on its properties. As such, the *formal counterpart* of a symplectic foliation should be a 2–form in $\Lambda^2(T^*\mathcal{F})$ that is non–degenerate but not necessarily closed. Note that this is the formal analogue of both weak and strong symplectic foliations.

Another question, which is probably more topological in nature, asks us to only fix the manifold M and then to find a foliation $\mathcal{F}$ and a leafwise 2–form such that $(M, \mathcal{F}, \omega)$ is a strong/weak symplectic foliation. Now the object of interest is the manifold M, and the main motivation is to relate its topology to the existence/classification of the symplectic foliations it admits. This question is precisely the one that is interesting if M is 3–dimensional: to see whether M admits some taut foliation or not.

In this latter case, a *formal symplectic foliation* would be a pair conformed of a distribution and a non–degenerate 2–form over it. If the distribution is of codimension 1, this is the same as a *formal contact structure*. If we were interested in the foliation per se, instead of a distribution we could consider a Haefliger structure.

The problem of constructing and classifying symplectic structures in closed manifolds is extremely hard, whereas in open manifolds, Gromov's h–principle completely solves the problem. Similarly, there has been some work (in the form of h–principles) in the open foliated setting [22, 4] and some explicit constructions for closed foliated manifolds [59, 64].

1.1.3 Basic properties

A reasonable question is whether being strong symplectic is much more restrictive than being just symplectic. Let us introduce some notation. Given a symplectic foliation, consider the space $\mathfrak{E}(\mathcal{F}, \omega)$ of all extensions of ω, which we endow with the C^1–topology.

Lemma 1.5. *The space $\mathfrak{E}(\mathcal{F}, \omega)$ is either empty or has a natural affine space structure.*

Proof. Suppose that $\mathfrak{E}(\mathcal{F}, \omega)$ is non–empty and fix an extension $\tilde{\omega}_0$. We set:

$$\{\tilde{\omega} \in \Omega^2(M) : \tilde{\omega}|_{\mathcal{F}} = \tilde{\omega}_0|_{\mathcal{F}}, d\tilde{\omega} = 0\}$$

is an affine space modelled on the space of closed 2–forms that vanish along the foliation. □

The way in which we check that the space of extensions is not empty is cohomological in nature and can be phrased in terms of obstruction classes, see [12, Example 10] and [81]. Let us elaborate. Suppose we have fixed a foliation $\mathcal{F}$. A first question to tackle is whether it is symplectic, as discussed before. Suppose it is, and we are able to exhibit a particular ω achieving this. Then, for that particular ω, we can do this obstruction theory computation to check whether ω possesses an extension.

Strongness is relevant in regard to the concept of a symplectic connection.

Definition 1.6. *Let $(M, \mathcal{F}, \omega)$ be a symplectic foliation. Given an extension $\tilde{\omega}$, the* ***symplectic connection*** *is the distribution $\mathcal{H}_{\tilde{\omega}} = (T\mathcal{F})^{\perp\tilde{\omega}}$.*

In the codimension–1 case, if a defining form β for the foliation is chosen (i.e. $\ker(\beta) = \mathcal{F}$), a symplectic connection determines a distinguished **transverse vector field** T:

$$\tilde{\omega}(T) = 0, \quad \beta(T) = 1,$$

which is the analogue of the Reeb vector field in contact topology. Just as closedness of the extension is the condition required in symplectic fibrations for the parallel transport to be by symplectomorphisms, see [56, Lemma 6.18], closedness of $\tilde{\omega}$ implies that T preserves $\tilde{\omega}$:

$$\mathcal{L}_T\tilde{\omega} = di_T\tilde{\omega} + i_T d\tilde{\omega} = 0.$$

The vector field T has been known in Poisson geometry for a while [81, Definition 4.8]. For strong symplectic foliations defined by a closed 1–form, it was proven in [34] that T preserves the Poisson structure. From our perspective, this can be rephrased by saying that, since β is closed, T not only preserves $\tilde{\omega}$, but also $\mathcal{F}$. A particularly simple case is when $[\beta] \in H^1(M)$ is rational: then $(M, \mathcal{F}, \omega)$ is actually the mapping torus of a symplectomorphism.

1.1.4 Moser's stability

Classic *Moser's stability* says that homotopies of symplectic structures representing the same cohomology class arise from isotopies. The same statement holds true in the foliated setting. Since we are talking about cohomology, closedness is necessary:

Lemma 1.7 (Moser Stability [37])**.** *Consider a foliation $\mathcal{F}$ on a closed manifold M. Let $\{\omega_t\}_{t\in[0,1]}$ be a smooth family of foliated 2–forms such that $(M, \mathcal{F}, \omega_t)$ is a symplectic foliation for every t. Let $\{\tilde{\omega}_t\}$ be a smooth family of extensions. Suppose that $[\tilde{\omega}_t] \in H^2_{DR}(M)$ is constant.*

Then, there exists a global flow $\{\phi_t\}_t \in \mathrm{Diff}(M)$ tangent to $\mathcal{F}$ and such that $\phi_t^*\omega_t = \omega_0$.

Proof. Let us assume that the flow ϕ_t is tangent to the leaves. Denote by X_t the vector field generating the flow ϕ_t. Differentiating with respect to t, we obtain:

$$[\phi_t^*(\mathcal{L}_{X_t}\tilde{\omega}_t + \frac{d\tilde{\omega}_t}{dt})]|_{\mathcal{F}} = 0. \tag{1.1}$$

Take the push forward by ϕ_t in the Equation (1.1) and expand using Cartan's formula to yield:

$$(di_{X_t}\tilde{\omega}_t + \frac{d\tilde{\omega}_t}{dt})|_{\mathcal{F}} = 0,$$

Applying Hodge theory, we can find a 1–parametric family of 1-forms λ_t satisfying $d\lambda_t = \dfrac{d\tilde{\omega}_t}{dt}$, by our assumption on the cohomology classes of the $\tilde{\omega}_t$. Now the equation reads:

$$(i_{X_t}\tilde{\omega}_t + \lambda_t)|_{\mathcal{F}} = 0.$$

Since $\tilde{\omega}_t$ is non-degenerate when restricted to $\mathcal{F}$, we can solve for X_t uniquely. □

Similarly, one can prove an analogue of Darboux's Theorem, providing a local normal form for a symplectic foliation. We will state it in the codimension–1 case for simplicity.

Example 1.8. Consider the product $M = \mathbb{R} \times \mathbb{C}^n$ with coordinates $(t; x, y)$ with $(x, y) = (x_1, \cdots, x_n, y_1, \cdots, y_n)$. This manifold can be endowed with a strong symplectic foliation:

$$\mathcal{F}_{\mathrm{std}} = \coprod_{t\in\mathbb{R}} \{t\} \times \mathbb{C}^n, \quad \tilde{\omega}_{\mathrm{std}} = \sum_{i=1}^{n} dx_i \wedge dy_i.$$

The standard defining form for $\mathcal{F}_{\mathrm{std}}$ is the 1–form $\beta_{\mathrm{std}} = dt$.

Lemma 1.9 ((Strict) Darboux's Theorem)**.** *Consider a symplectic foliation $(M^{2n+1}, \mathcal{F}^{2n}, \omega)$ and a point $p \in M$. Then there exist a small $\varepsilon > 0$ and an embedding*

$$\phi : \mathbb{D}^{2n}_{\varepsilon} \times (-\varepsilon, \varepsilon) \longrightarrow M$$
$$\phi(0, 0) = p, \quad \phi^*\mathcal{F} = \mathcal{F}_{\mathrm{std}}, \quad \phi^*\omega = \tilde{\omega}_{\mathrm{std}}|_{\mathcal{F}_{\mathrm{std}}}.$$

If $(M, \mathcal{F}, \omega)$ is actually strong symplectic with extension $\tilde{\omega}$, it can further be achieved that $\phi^*\tilde{\omega} = \tilde{\omega}_{\mathrm{std}}$.

1.2 Contact foliations

Contact manifolds appear in symplectic topology as the natural boundary conditions for symplectic manifolds. We will now introduce their foliated analogue. In Section 1.3 we will relate the two foliated notions seen so far through examples and several constructions.

1.2.1 Definition

Definition 1.10. *Let V^{2n+1+v} be a smooth manifold. Let $\mathcal{F}$ be a rank $2n + 1$ foliation in V. Let ξ be a codimension–1 distribution in $T\mathcal{F}$ such that $(\mathcal{L}, i^*_{\mathcal{L}}\xi)$ is a contact manifold for every leaf $i_{\mathcal{L}} : \mathcal{L} \to V$. Then, the triple $(V, \mathcal{F}, \xi)$ is said to be a* ***contact foliation****.*

If V and $\mathcal{F}$ have been fixed, the space of contact foliations is denoted as follows:

$$\mathfrak{ContFol}(V, \mathcal{F}) = \{\xi \subset T\mathcal{F} \mid (V, \mathcal{F}, \xi) \text{ is a contact foliation}\}$$

and we endow it with the C^1–topology.

Let $\Theta \subset TV$ be a codimension–1 distribution that satisfies $\xi = \Theta \cap T\mathcal{F}$. Then, Θ is said to be an **extension** of the contact foliation. The reader should compare this with a *contact fibration* [49]. Consider the space of possible extensions with the C^1–topology:

$$\mathcal{E}(V, \mathcal{F}, \xi) = \{\Theta \mid \Theta \text{ is an extension of } (V, \mathcal{F}, \xi)\}.$$

Just like the space of connection 1–forms in a smooth vector bundle has a natural affine structure, one can prove the following:

Lemma 1.11. *The space $\mathcal{E}(V, \mathcal{F}, \xi)$ has an affine structure.*

Proof. Consider an extension Θ_0 with defining form α_0. The space

$$A = \{\alpha \in \Omega^1(V) : \alpha|_{\mathcal{F}} = \alpha_0|_{\mathcal{F}}\}$$

is an affine space modelled on the space of 1–forms vanishing on $\mathcal{F}$. The map $\Psi : A \longrightarrow \mathcal{E}(V, \mathcal{F}, \xi)$ defined by $\Psi(\alpha) = (V, \mathcal{F}, \ker \alpha)$ is bijective and endows $\mathcal{E}(V, \mathcal{F}, \xi)$ with an affine structure. □

In particular, $\mathcal{E}(V, \mathcal{F}, \xi)$ is contractible and the choice of extension for a foliated contact structure is unique up to homotopy.

The reader might wonder whether there is a sensible notion of *closedness* for contact foliations. The key observation is that the form Θ is exact and hence closed, regardless of the choice of extension. In particular, this will later imply that the *symplectisation* of a contact foliation is strong symplectic, showing the desired relation between both notions.

1.2.2 Formal contact foliations

In Chapter 2, we will address the following question: given a manifold V and a foliation $\mathcal{F}$, can it be endowed with a contact foliation $(V, \mathcal{F}, \xi)$? A big difference with the symplectic case is that we will be able to provide a complete answer: if the formal problem is solvable, so is the geometric one. This is a joint result [9] with C. Casals and F. Presas.

What does this mean? Fix the manifold V and the foliation $\mathcal{F}$. Then, the space of**formal contact foliations** is the following:

$$\begin{aligned} \mathcal{F}\mathfrak{ContFol}(V, \mathcal{F}) = &\{(\xi, [\omega]) \mid \xi \text{ is a codimension–1 distribution in } T\mathcal{F}, \\ &\omega \in \Lambda^2(\xi) \text{ is non–degenerate}\} \end{aligned}$$

and we endow it with the C^0–topology; $[\omega]$ means the conformal class of the 2–form ω. Contact foliations can be regarded naturally as a subset of the formal contact ones by taking ω to be $d\alpha|_\xi$ with α any defining 1–form for an extension Θ of ξ.

Remark 1.12. When $\mathcal{F}$ has rank 3 (and thus ξ has rank 2), the formal data reduces to recalling ξ as a plane field tangent to the foliation.

We could have opted to let the foliation $\mathcal{F}$ be part of the PDR as well; however, even with the foliation fixed, our result reads:

Theorem 2.1. *Let V be a 4–manifold and let $\mathcal{F}$ be a foliation of rank 3. Then, the inclusion:*

$$\mathfrak{ContFol}(V, \mathcal{F}) \to \mathcal{F}\mathfrak{ContFol}(V, \mathcal{F})$$

induces a surjection in homotopy groups.

However, this does not make contact foliations uninteresting. The foliations we produce will have all leaves overtwisted, so the reasonable class of foliations to study is the class of *contact foliations with tight leaves.* Future work should focus on understanding whether a given manifold admits a foliation that can be made contact with tight leaves.

1.2.3 The contact connection

Contact foliations relate to a class of distributions we have seen already:

Lemma 1.13. *Let $(V, \mathcal{F}, \xi)$ be a contact foliation with $\mathcal{F}$ of codimension–1. Then, any extension Θ is an even–contact structure.*

Indeed, Θ restricted to $\ker(\Theta)$ has maximal rank. This naturally ties in with the next definition:

Definition 1.14. *Let Θ be an extension of a contact foliation $(V, \mathcal{F}, \xi)$. Let α be its defining form. The associated **contact connection** is the distribution $\mathcal{H}_\Theta = \xi^{\perp d\alpha} \subset (\Theta, d\alpha)$.*

In the language of even–contact structures (when $\mathcal{F}$ is of codimension–1), this is known as the **kernel**. It is worth remarking that the contact connection does not depend on the choice of α, it only depends on the extended distribution Θ; it is always transverse to $\mathcal{F}$.

If $\mathcal{F}$ is of codimension–1, $\mathcal{H}_\Theta$ is the kernel, a line field. A particular vector field spanning it can be chosen if we pick a defining form α for Θ and a defining form β for $\mathcal{F}$. Then, we can set:

$$\alpha(T) = 0, \quad (i_T d\alpha) \wedge \alpha = 0, \quad \beta(T) = 1$$

We call T the **transverse vector field**. It is then obvious that:

Lemma 1.15. *Let $\mathcal{F}$ be of codimension–1. Let $(V, \mathcal{F}, \xi)$ be a contact foliation. Let Θ be an extension with α its defining 1–forms. Then*

$$\mathcal{L}_T \alpha = d\alpha(T, R)\alpha.$$

In particular the distribution Θ is preserved by the flow of the transverse field T of (β, α).

It does not hold in general that T preserves $\mathcal{F}$ since that would imply $d\beta = 0$.

1.2.4 The Reeb vector field

Just like in standard contact topology, once we fix a defining–1 form for the extension Θ, there is a uniquely defined vector field R that is tangent to each leaf. Namely:

$$\alpha(R) = 1, \quad (i_R d\alpha)|_{\mathcal{F}} = 0, \quad R \subset T\mathcal{F}.$$

In Chapter 5, we will pose the following question: *does R have any closed orbits?* In the closed non–foliated case, the conjecture that says that it always does (for any choice of contact form) is called the *Weinstein conjecture*. We shall see below (Subsection 1.3.2) that this is not the case in the foliated setting, although the result can be salvaged when there are *overtwisted* leaves.

1.2.5 The space of transformations of a contact foliation

A diffeomorphism of V that preserves both $\mathcal{F}$ and ξ will be called a **foliated contactomorphism**. The infinitesimal symmetries are described as follows.

Definition 1.16. *Let $(V, \mathcal{F}, \xi)$ be a contact foliation. The space of* ***contact vector fields*** *of $(V, \mathcal{F}, \xi)$ is defined as the subspace of $\mathfrak{X}(V)$ of those vector fields that preserve $\mathcal{F}$ and ξ.*

If an extension Θ is fixed and a defining 1–form α is given, the space of contact vector fields is equivalently defined as:

$$\begin{aligned}\mathfrak{X}(\xi) = \{X \in \mathfrak{X}(V) : \quad & \mathcal{L}_X\alpha|_{\mathcal{F}} = f\alpha|_{\mathcal{F}} \quad \text{for some } f \in C^\infty(V) \text{ and} \\ & [X, \mathcal{F}] \subset \mathcal{F}\}.\end{aligned}$$

A flow by foliated contactomorphisms is induced by a 1–parametric family of contact vector fields.

Lemma 1.17. *The space $\mathfrak{X}(\xi)$ is a Lie algebra. It contains a distinguished ideal $\mathfrak{X}(\xi) \cap \mathfrak{X}(\mathcal{F})$.*

Proof. The first statement follows since the space of vector fields preserving a distribution is a Lie algebra. The second claim is immediate, since $[X, Y] \in T\mathcal{F}$, for X preserving $\mathcal{F}$ and $Y \in \mathcal{F}$. $\square$

1.2.6 Gray stability

Just like we proved Moser's stability for strong symplectic foliations, Gray's stability has a foliated analogue. Similar Moser–type theorems were proven in [2] in the setting of symplectic and contact pairs.

Proposition 1.18 (Foliated Gray's Stability)**.** *Let $\mathcal{F}$ be a codimension 1–foliation on a closed manifold V. Let $\{\xi_t\}_{t\in[0,1]}$ be a family of codimension–2 distributions such that $(V, \mathcal{F}, \xi_t)$ is a foliated contact structure for every $t \in [0, 1]$.*

Then, there exists a global flow $\{\phi_t\}_{t\in[0,1]} \in \mathrm{Diff}(V)$ tangent to $\mathcal{F}$ such that $\phi_t^*\xi_t = \xi_0$.

Proof. We consider a smooth 1–parametric family of extensions Θ_t and their defining 1–forms α_t. We require a flow ϕ_t tangent to the leaves (and therefore preserving $\mathcal{F}$) such that $\phi_t^*\xi_t = \xi_0$. This reads as

$$\phi_t^*\alpha_t|_{\mathcal{F}} = g_t\alpha_0|_{\mathcal{F}}$$

for a suitable choice of $\{g_t\}_t \in C^\infty(V)$. We now apply the foliated version of Moser's argument.

Denote by X_t the vector field generating the required flow ϕ_t (that is, $X_t \circ \phi_t = \dot{\phi}_t$) and we further suppose that X_t is contained in the contact structures ξ_t. Differentiating the above condition with respect to t we obtain:

$$\phi_t^*(\mathcal{L}_{X_t}\alpha_t + \dot{\alpha}_t)|_{\mathcal{F}} = g_t'\alpha_0|_{\mathcal{F}} = \frac{g_t'}{g_t}(\phi_t^*\alpha_t)|_{\mathcal{F}}. \tag{1.2}$$

Define $\lambda_t = (\phi_t)_*\dfrac{g_t'}{g_t}$. Equation (1.2) implies

$$(\mathcal{L}_{X_t}\alpha_t + \dot{\alpha}_t)|_{\mathcal{F}} = (i_{X_t}d\alpha_t + \dot{\alpha}_t)|_{\mathcal{F}} = \lambda_t\alpha_t|_{\mathcal{F}}. \tag{1.3}$$

This is an equation in 1–forms. In particular, it has to be satisfied by the Reeb vector R_t, thus yielding the condition

$$i_{R_t}\dot{\alpha}_t = \lambda_t.$$

This reads $\phi_t^*(i_{R_t}\dot{\alpha}_t) = (\ln g_t)'$, which is an ODE with an unique solution once the initial condition $g_0 = 1$ is fixed. Now, since $R_t \in \ker(\dot{\alpha}_t - \lambda_t\alpha_t)$, Equation (1.3) can solved uniquely for $X_t \in \xi_t$. $\square$

It is worth pointing out that we are not able to provide an isotopy taking the extensions to one another. The isotopy only deals with the leafwise structures. Indeed, if $\mathcal{F}$ is codimension–1, an extension is just an even–contact structure and it is well–known that Gray stability fails unless the kernel is fixed [31, 60].

1.3 Examples and constructions

In this section we shall produce several examples of contact and symplectic foliations. Essentially, the methods by which we produce them are a straightforward generalisation of their non–foliated analogues.

1.3.1 Contact mapping tori

Let $(\mathcal{L}, \xi_{\mathcal{L}} = \ker \alpha_{\mathcal{L}})$ be a contact manifold. The manifold $\mathcal{L} \times [0,1]$, with coordinates (p,t), has a natural contact foliation structure given by

$$\tilde{\mathcal{F}} = \mathcal{L} \times \{t\}, \quad \tilde{\xi}(p,t) = (\iota_{\mathcal{L}\times\{t\}})_* \xi_{\mathcal{L}}(p).$$

Let ϕ be a contactomorphism of $\xi_{\mathcal{L}}$ and consider the associated mapping torus $M(\phi)$, which is a contact fibration. From the contact foliation viewpoint, it inherits the contact foliation structure $(\mathcal{F}, \xi)$ from $(\mathcal{L} \times [0,1], \tilde{\mathcal{F}}, \tilde{\xi})$ as a quotient. Given a vector field $X \subset T\mathcal{F}$ such that $\phi_* X = X$, an extension is obtained by declaring

$$\Theta = \xi \oplus \langle \partial_t + X \rangle.$$

Denote $H = \alpha_{\mathcal{L}}(X)$. Then the contact connection $\mathcal{H}_\Theta$ is the distribution generated by the vector field $\langle \partial_t + \widetilde{X} \rangle$ satisfying the equations

$$\alpha_{\mathcal{L}}(\widetilde{X}) - H = 0, \quad d\alpha_{\mathcal{L}}(\widetilde{X}, v) + dH(v) = 0 \quad \forall v \in \ker \alpha_{\mathcal{L}}.$$

Hence $\widetilde{X}$ is the Hamiltonian vector field associated to H.

1.3.2 A contact foliation with no closed Reeb orbits

The following example is due to V. Ginzburg.

There are four natural types of foliations on the 4–torus obtained by quotienting the horizontal foliation of a 4–polydisc by 3–polydiscs. In coordinates $\mathbb{T}^4(t,x,y,z)$ their defining equations have the form

$$\beta = (p,q,r,s) \cdot (dx, dy, dz, -dt), \quad (p,q,r,s) \in \mathbb{R}^4.$$

The numbers (p,q,r,s) generate a $\mathbb{Q}$–submodule A of $\mathbb{R}$. The leaves are diffeomorphic to $(\mathbb{S}^1)^{(4-\mathrm{rank}(A))} \times \mathbb{R}^{\mathrm{rank}(A)-1}$.

Let us endow such foliations with a foliated contact structure. Suppose that $s = 1$ and consider the form

$$\alpha = \sin(2\pi z)dx + \cos(2\pi z)dy.$$

This is a well–defined 1–form on $\mathbb{T}^4$. The condition for a foliated contact structure reads

$$\alpha \wedge d\alpha \wedge \beta = (\sin(2\pi z)dx + \cos(2\pi z)dy) \wedge (2\pi \cos(2\pi z)dz \wedge dx - 2\pi \sin(2\pi z)dz \wedge dy) \wedge \beta =$$

$$(2\pi dx \wedge dy \wedge dz) \wedge (p \cdot dx + q \cdot dy + r \cdot dz - dt) = 2\pi \cdot dt \wedge dx \wedge dy \wedge dz > 0$$

Hence, α defines a foliated contact structure for any 1–form β as above. In particular, we obtain a contact foliation with (dense) tight contact $\mathbb{R}^3$ leaves on $\mathbb{T}^4$. Precisely in this case it is immediate that all Reeb orbits are lines. This shows that the *foliated Weinstein conjecture* fails in general.

1.3.3 Cotangent bundle

Let W be a smooth manifold and let $\mathcal{F}$ be a foliation on W. The total space M of the fibration $\pi : T^*\mathcal{F} \longrightarrow W$ can be endowed with a strong symplectic foliation. The foliation is simply given as a lift of $\mathcal{F}$:

$$\mathcal{F}_M = \coprod_{\mathcal{L} \in \mathcal{F}} \pi^{-1}\mathcal{L}.$$

The leafwise symplectic form ω_M (or rather, its extension $\tilde{\omega}_M$) is exact. A distinguished primitive is constructed as follows. At a point $(p, \lambda) \in T^*\mathcal{F}$, where p is a point in W and λ is a covector of $\mathcal{F}$, the *Liouville 1–form* is defined to be:

$$\lambda_{\text{std}}(p, \lambda) = \lambda \circ d_{(p,\lambda)}\pi,$$

that is, vectors *tangent to* $\mathcal{F}_M$ are projected to the base manifold and then they are evaluated with the covector their basepoint in $M = T^*\mathcal{F}$ represents. This defines a leaf–wise 1–form, and we set $\omega_M = d\lambda_{\text{std}}$. Then, λ_{std} can be extended arbitrarily to a global 1–form and any such choice yields an extension $\tilde{\omega}_M$ (in particular, the extension is non–unique).

Restricted to each leaf, this is the standard construction, and hence the resulting structures are indeed leafwise symplectic.

1.3.4 Foliated contact elements

The contact analogue of what we just did also works. Let W be a smooth manifold and let $\mathcal{F}$ be a foliation on W. Consider the manifold $V = \mathbb{P}(T^*\mathcal{F})$, which we are going to endow with a contact foliation. Using the projection $\pi : V \longrightarrow W$, we pullback the foliation $\mathcal{F}$ on the base; that is, we set

$$\mathcal{F}_V = \coprod_{\mathcal{L} \in \mathcal{F}} \pi^{-1}\mathcal{L}.$$

Now we can tautologically define the following codimension 1–distribution in $T\mathcal{F}$:

$$\xi_V(p, [\lambda]) = d\pi^{-1}(\ker(\lambda))$$

that is, being p a point in W and $[\lambda]$ a conformal class of covectors of $\mathcal{F}$ at p, we set the contact plane to be the preimage of their kernel.

In particular, the foliated contact structure restricted to a leaf coincides with the standard *space of contact elements* of the leaf. In the same vein, the sphere bundle $S = \mathbb{S}(T^*\mathcal{F})$ associated to the cotangent space of the foliation is a foliated contact manifold $(S, \mathcal{F}_S, \xi_S)$ that restricts to the *space of cooriented contact elements* over each leaf. The foliated contact structure $(\mathcal{F}_S, \xi_S)$ can also be obtained via the pullback of $(\mathcal{F}_V, \xi_V)$ through the double–cover $S \longrightarrow V$.

1.3.5 Contactisation

Let (M, λ) be an exact symplectic manifold, the contactisation of (M, λ) is the manifold $M \times \mathbb{R}$ endowed with the contact structure $\xi = \ker(\lambda - dt)$.

Definition 1.19. *Let $(M, \mathcal{F}, \omega)$ be a strong symplectic foliation admitting an exact extension $\tilde{\omega} = d\lambda$. Then,*

$$(M \times \mathbb{R}, \mathcal{F} \times \mathbb{R}, \ker(\lambda - dt))$$

*where t is the coordinate in the factor $\mathbb{R}$, is called the **contactisation** of $(M, \mathcal{F}, \omega)$.*

Of course, the contactisation is an open manifold. However, there exists another construction, called the **prequantisation**, that allows us to construct a closed manifold endowed with a contact foliation from a *strong* symplectic foliation. Assume that $\tilde{\omega}$ is a closed 2–form such that $[\tilde{\omega}/(2\pi)]$ is integral, which can be assumed after a small perturbation and scaling. Consider the principal circle bundle $\pi : L_{\tilde{\omega}} \to M$ associated to $\tilde{\omega}$. Construct a connection 1–form $\alpha \in \Omega^1(L_{\tilde{\omega}})$ with curvature $\tilde{\omega}$.

Definition 1.20. *The foliated contact manifold $(L_{\tilde{\omega}}, \pi^*\mathcal{F}, \ker(\alpha))$ is said to be the **Boothby–Wang contact foliation** associated to the strong symplectic foliation $(M, \mathcal{F}, \omega)$ with extension $\tilde{\omega}$.*

The Boothby–Wang contact foliation over a closed base is taut because the strong symplectic foliation is taut, so a transverse loop in the base can simply be lifted. Observe that this is not necessarily the case for a general contact foliation (as Theorem 2.1 shows).

1.3.6 Symplectisation

Let $(V, \xi = \ker \alpha)$ a contact manifold. The symplectic manifold $(V \times \mathbb{R}, d(e^t\alpha))$ is known as the symplectisation of (V, ξ).

Definition 1.21. *Let $(V, \mathcal{F}, \xi)$ be a contact foliation and let Θ be an extension. Fix a defining 1–form α for Θ. The manifold*

$$(\mathbb{R} \times V, \mathbb{R} \times \mathcal{F}, \omega = d(e^t\alpha)|_{\mathcal{F}}).$$

*is called the **symplectisation** of $(V, \mathcal{F}, \xi)$, where t is the coordinate in the $\mathbb{R}$ factor.*

Notice that $\tilde{\omega} = d(e^t\alpha)$ is an extension of ω. The strong symplectic foliation obtained in the construction does not depend on the particular choice of α and Θ up to foliated symplectomorphism.

1.3.7 The parametric Lutz twist

In the Preamble we discussed briefly how Lutz and Martinet were able to prove that any 3–dimensional manifold admits a contact structure using surgery and a construction that we called the *Lutz twist*. Later on, in Chapter 2, we will need a parametric version of this procedure, so we will introduce it here.

1.3.7.1 The non–foliated case

Let us go over the standard Lutz twist before we get to its parametric/foliated counterpart. A good reference is [27, Section 4.3].

Lemma 1.22. *Fix a contact manifold (M^3, ξ) and let $\gamma : \mathbb{S}^1 \to M$ be a loop that is transverse to the contact structure ξ. Then, there is a diffeomorphism*

$$\phi : \mathcal{O}p(\gamma) \to \mathbb{S}^1 \times \mathbb{D}^2_\varepsilon$$

satisfying $\phi^* \ker(dt - r^2 d\theta) = \xi$, where t is the coordinate in $\mathbb{S}^1$ and (r, θ) are polar coordinates in the disc.

We will not go over its proof. The reader is invited to read [27, Example 2.5.16].

Geometrically, we see that the structure ξ is rotating as we move radially away from γ. A Lutz twist consists of replacing ξ by a structure that agrees with ξ near γ and near the boundary of $\mathcal{O}p(\gamma)$ but that in the middle performs an additional half turn or complete turn. These are called the (half) Lutz twist and the full Lutz twist, respectively.

Namely, find functions $g, h : [0, \varepsilon] \to \mathbb{R}$ satisfying:

- $g(r) = 1$ if $r \in \mathcal{O}p(\{0, \varepsilon\})$,
- $h(r) = -r^2$ if $r \in \mathcal{O}p(\{0, \varepsilon\})$,
- the curve $(g, h)/|(g, h)| : [0, \varepsilon] \to \mathbb{S}^1$ is an immersion describing 1 turn around the sphere.

These functions describe the full Lutz twist. The following lemma collects the relevant properties:

Lemma 1.23. *Fix a contact manifold (M^3, ξ) and let $\gamma : \mathbb{S}^1 \to M$ be a loop that is transverse to the contact structure ξ. Using the model*

$$\phi : \mathcal{O}p(\gamma) \to \mathbb{S}^1 \times \mathbb{D}^2_\varepsilon$$

$$\phi^* \ker(dt - r^2 d\theta) = \xi,$$

define a contact structure $\tilde{\xi}$ that agrees with ξ outside of $\mathcal{O}p(\gamma)$ and is given by $\phi^* \ker(g(r)dt - h(r)d\theta)$ inside.

The structure $\tilde{\xi}$ is said to be obtained from ξ by a ***full Lutz twist****. There exists a family of plane fields ξ_t, $t \in [0, 1]$, with $\xi_0 = \xi$ and $\xi_1 = \tilde{\xi}$, all of them agreeing outside of $\mathcal{O}p(\gamma)$.*

That is, this method produces new contact structures that formally are homotopic to the original one, but not necessarily geometrically. Indeed, as explained in the Preamble, this method produces *flexible* contact structures, as the following definition and proposition show.

Definition 1.24. *Endow $\mathbb{S}^1 \times \mathbb{D}^2_\varepsilon$ with the contact structure $\xi_{\mathrm{Lutz}} = \ker(g(r)dt - h(r)d\theta)$. Let $r_0 > 0$ be the first time where $h(r) = 0$.*

Let (M, ξ) be a contact manifold. A disc $D \subset M$ is said to be ***overtwisted*** *if there is*

$$\phi : \mathcal{O}p(D) \to \mathcal{O}p(\{1\} \times \mathbb{D}^2_{r_0})$$

such that $\phi^* \xi_{\mathrm{Lutz}} = \xi$. If (M, ξ) contains an overtwisted disc, ξ is an ***overtwisted contact structure****.*

Proposition 1.25. (Eliashberg, '89) *Let M be a 3–manifold. The following inclusion is a weak homotopy equivalence:*

$$\mathcal{C}\text{-}\mathcal{S}\mathfrak{trs}_{\mathrm{OT}}(M,D) \to \mathcal{P}\mathfrak{lanes}(M,D).$$

Here $\mathcal{C}\text{-}\mathcal{S}\mathfrak{trs}_{\mathrm{OT}}(M,D)$ denotes the contact structures on M having D as an overtwisted disc and $\mathcal{P}\mathfrak{lanes}(M,D)$ the plane fields that are tangent to D at its origin. This result will play a key role when we try to prove Theorem 2.1, because it provides the flexibility required to manipulate formal contact foliations into honest ones.

1.3.7.2 The foliated case

Let $(V,\mathcal{F},\xi)$ be a contact foliation. The parametric Lutz twist relies on the fact that, given a loop transverse to $\mathcal{F}$, it is possible to find a local model around it in which the leafwise contact structures are standard:

Lemma 1.26. *Let $(V,\mathcal{F},\xi)$ be a contact foliation with* $\mathrm{rank}(\mathcal{F}) = 3$. *Let $\eta \subset V$ be a curve everywhere transverse to $\mathcal{F}$. Then, there is a diffeomorphism*

$$\phi_\eta : \mathcal{O}p(\eta) \to I \times \mathbb{D}^3_\varepsilon$$

$$\phi_\eta^*\xi_{\mathrm{std}} = \xi,$$

where $\xi_{\mathrm{std}} = \ker(dz - xdy) \cap \{t\} \times \mathbb{D}^3_\varepsilon$. *$I$ denotes some 1–dimensional manifold.*

Proof. This is an immediate consequence of Moser's trick. A more geometrical way to do it is to find a vector field $X \subset \xi$ in $\mathcal{O}p(\eta)$ and use its flow to simultaneously find an I–family of flowboxes. Then, the implicit function theorem puts the contact structures in standard position. □

Remark 1.27. Let $K \subset V$ be a manifold of dimension at most $\mathrm{codim}(\mathcal{F})$ that is everywhere transverse to $\mathcal{F}$. Then, if $\xi|_K$ is the trivial bundle, the obvious analogue of Lemma 1.26 holds.

Given some transverse knot $\gamma : \mathbb{S}^1 \to (\mathbb{D}^3_\varepsilon, \xi_{\mathrm{std}})$, we can consider the cylinder $C = \mathbb{S}^1 \times \gamma \subset \mathbb{S}^1 \times \mathbb{D}^3_\varepsilon$. Let $\eta : \mathbb{S}^1 \to V$ be a loop transverse to $\mathcal{F}$ and let ϕ_η be defined as in Lemma 1.26. Then $\phi_\eta^{-1}(C)$ is a 1–parametric family of transverse knots in V.

Any embedded loop of transverse knots can be assumed to have a standard local model given by a parametric application of Lemma 1.22. That is, there is a map $\Psi : \mathcal{O}p(C) \to \mathbb{S}^1 \times \mathbb{S}^1 \times \mathbb{D}^2_\varepsilon$ satisfying:

$$\Psi^*(\ker(ds)) = \mathcal{F}, \qquad \Psi^*[\ker(dt - r^2 d\theta) \cap \ker(ds)] = \xi.$$

We can then define ξ_{pLutz}, a leafwise contact structure that agrees with ξ outside of $\mathcal{O}p(C)$ and in $\mathcal{O}p(C)$ is given by $\Psi^*[\xi_{\mathrm{Lutz}} \cap \ker(ds)]$. We say that it has been obtained from ξ by a **parametric Lutz twist** along the family C.

Now we can prove the following result:

Proposition 1.28. *The inclusion*

$$\mathfrak{ContFol}(V,\mathcal{F}) \to \mathcal{F}\mathfrak{ContFol}(V,\mathcal{F})$$

is not a weak homotopy equivalence. Indeed, it is not injective in π_0.

Proof. Consider the Reeb foliation $(\mathbb{S}^3, \mathcal{F})$ and its associated space of foliated contact elements $(\mathbb{S}^3 \times \mathbb{S}^1, \mathcal{F}_c, \xi_c)$. There exists a homotopy at the formal level that produces a contact foliation ξ_{OT} with tight and overtwisted leaves. Assuming that this is true, Gray's stability (Proposition 1.18) says that ξ_c and ξ_{OT} being homotopic as contact structures implies that they are isotopic. But this is a contradiction with the fact that ξ_c has all leaves tight.

The Reeb foliation in $\mathbb{S}^3$ is given by gluing two Reeb components along their torus boundary suitably. We can find a knot $\eta : \mathbb{S}^1 \to \mathbb{S}^3$ running along the core of one of the solid tori that is everywhere transverse to $\mathcal{F}$. In particular, η does not intersect the torus leaf nor the planar leaves of the other Reeb component. Lift η to a curve $\tilde{\eta}$ in $\mathbb{S}^3 \times \mathbb{S}^1$ and use Lemma 1.26 to trivialise the leafwise contact structures along $\tilde{\eta}$. This allows us to find an embedded family of transverse knots that is C^0–close to $\tilde{\eta}$; we perform a parametric Lutz twist along this family.

The (lift of the) leaves in the interior of the Reeb component containing η are now overtwisted, whereas all other leaves are tight. This proves the claim. The same construction works for any foliation $(M^3, \mathcal{F})$ with a Reeb component. $\square$

1.3.8 Foliated contact divisor connected sum

We will end this chapter carrying out one last construction. It is not really necessary for the rest of the contents of the thesis, but it has some interest nonetheless.

Consider a contact foliation $(V, \mathcal{F}, \xi)$ on a $(2n+2)$–fold V and let S be a $2n$–dimensional submanifold transverse to $\mathcal{F}$ and ξ. Then S inherits a codimension–1 foliation $\mathcal{F}_S$ with foliated contact structure $\xi_S = \xi \cap T\mathcal{F}_S$. We say that $(S, \mathcal{F}_S, \xi_S)$ is a **foliated contact divisor**. Generalizing [27, Theorem 2.5.15], Proposition 1.18 implies that the tubular neighbourhood of S is uniquely determined by the conformal symplectic structure of its normal bundle. Since its normal bundle is a 2–dimensional disc bundle, the foliated contact structure of its tubular neighbourhood depends only on its oriented topological type.

Suppose that the normal bundle of S is trivial. Then, there is a diffeomorphism $\phi : S \times \mathbb{D}^2_{2\varepsilon} \to \mathcal{O}p(S)$, $\varepsilon > 0$, where $\mathcal{O}p(S) \subset V$ denotes a small tubular neighbourhood of S. Let α_S be a defining 1– form for an extension of ξ_S. Then, the pullback of the contact foliation by the embedding ϕ can be chosen to be:

$$(S \times \mathbb{D}^2_{2\varepsilon}, \mathcal{F}_S \times \mathbb{D}^2_{2\varepsilon}, \ker(\alpha_{\mathcal{O}p(S)})),$$

where $\alpha_{\mathcal{O}p(S)} = \phi^*\alpha_S + r^2 d\theta$. This is the local model along the foliated contact divisor, which now we can use to perform a connected sum.

Let $(V_0, \mathcal{F}_0, \xi_0)$ and $(V_1, \mathcal{F}_1, \xi_1)$ be two contact foliations and $f_0 : S \longrightarrow V_0$, $f_1 : S \longrightarrow V_1$ two embeddings of S as a foliated contact divisor with trivial normal bundle. There exist two neighbourhoods $\mathcal{O}p(S, V_0)$ and $\mathcal{O}p(S, V_1)$ and two embeddings

$$f_0 : S \times \mathbb{D}^2_{2\varepsilon} \longrightarrow \mathcal{O}p(S, V_0), \quad f_1 : S \times \mathbb{D}^2_{2\varepsilon} \longrightarrow \mathcal{O}p(S, V_1)$$

conforming to the local model described above (and extending f_0 and f_1).

The gluing region is the open manifold $\mathcal{S} = S \times (-\varepsilon^2, \varepsilon^2) \times \mathbb{S}^1$, endowed with the contact foliation:

$$(\mathcal{S}, \mathcal{F}_S \times (-\varepsilon^2, \varepsilon^2) \times \mathbb{S}^1, \ker(\alpha_{\mathcal{S}})$$

where $\alpha_{\mathcal{S}} = \alpha_S + t d\theta$); note the linearity in the t–coordinate. Define maps

$$F_0 : S \times (0, \varepsilon^2) \times \mathbb{S}^1 \longrightarrow \mathcal{O}p(S, V_0), \quad (p, t, \theta) \longmapsto f_0(p, t^2, \theta)$$
$$F_1 : S \times (-\varepsilon^2, 0) \times \mathbb{S}^1 \longrightarrow \mathcal{O}p(S, V_1), \quad (p, t, \theta) \longmapsto f_1(p, t^2, -\theta)$$

Then the topological connected sum

$$V_0 \#_S V_1 = (V_0 \setminus f_0(S)) \cup_{F_0} \mathcal{S} \cup_{F_1} (V_1 \setminus f_1(S))$$

with the foliated contact models introduced above inherits a foliated contact structure. The related construction for symplectic foliations is discussed in [42], though it does not preserve strongness.

Chapter 2

Existence of contact foliations

As we advanced during the introductory chapter, our aim now is proving that the existence h–principle holds for contact foliations of rank 3 and codimension 1:

Theorem 2.1. *Let V be a 4–manifold and let $\mathcal{F}$ be a foliation of rank 3. Then, the inclusion:*

$$\mathfrak{ContFol}(V, \mathcal{F}) \to \mathcal{F}\mathfrak{ContFol}(V, \mathcal{F})$$

induces a surjection in homotopy groups.

We should point out, first of all, that the restriction on the codimension is not necessary. However, for simplicity, only this case was tackled in [9]. The restriction on the rank was, at the time of writing, essential. The classification of overtwisted structures in all dimensions [6] had not been announced yet and hence, overtwistedness itself had not been defined in dimension greater than 3. However, due to the nature of overtwistedness, it was already clear that a higher rank analogue of Theorem 2.1 would follow from the non–foliated h–principle. Indeed, in [6, Theorem 1.5], the result we present here was proven in full generality.

Theorem 2.1 did not try to solve the classification question, although good guesses of what an overtwisted class has to be follow naturally from its proof. In [6] they defined an analogue of the overtwisted disc for contact foliations, dealing with this loose end.

Remark 2.2. Recall that in Subsection 1.3.7 we already provided an example where it is shown that the inclusion $\mathfrak{ContFol}(V, \mathcal{F}) \to \mathcal{F}\mathfrak{ContFol}(V, \mathcal{F})$ is not a weak homotopy equivalence in general, as one would expect from the non–foliated case. There, we introduced the *parametric Lutz twist*, which will be key for the upcoming arguments.

The proof of Theorem 2.1 goes roughly as follows. We will first triangulate the ambient manifold V in a manner that is adapted to the foliation (Subsection 2.1.3); this relies on Thurston's jiggling [79, Section 5]. Once that is done, we will first homotope our leafwise plane field to an honest contact structure in a neighbourhood of the lower dimensional simplices (Section 2.2); we call this the **reduction** process. Then, we have to solve the **extension problem**: that is, the contact condition has to be achieved in the interior of each top dimensional cell, relative to what we did in the previous steps (Section 2.3). In Section 2.4 we prove the parametric case of the statement, which goes through without major differences.

The two main ingredients needed are Gromov's h–principle for open manifolds (Subsection 2.1.1), and Eliashberg's h–principle for closed 3–manifolds (Subsection 2.1.2, where a more detailed statement of Proposition 1.25 is given).

2.1 Technical ingredients

In this section we state the main lemmas involved in the proof of Theorem 2.1. Subsection 2.1.1 introduces Gromov's h–principle. Subsection 2.1.2 discusses Eliashberg's. Then, in Subsection 2.1.3, we discuss triangulating the manifold and producing suitable neighbourhoods for the simplices: this will allow us to apply the aforementioned h–principle results.

Let us setup some basic notation that we will keep using during this chapter. Consider the standard closed 3–ball $\mathbb{D}^3_r \subset \mathbb{R}^3$ of radius r and fix a sequence of equators $\mathbb{S}^0_r \subset \mathbb{S}^1_r \subset \mathbb{S}^2_r$, respectively bounding closed flat discs $\mathbb{D}^1_r \subset \mathbb{D}^2_r \subset \mathbb{D}^3_r$. When we omit the radius r, we will mean $r = 1$. For any 1–dimensional manifold I, we write

$$\mathcal{F}_{\mathbb{D}^3_r \times \{t\}} = \coprod_{t \in I} \mathbb{D}^3_r \times \{t\}$$

for the product foliation in $\mathbb{D}^3_r \times I$.

2.1.1 h–Principle for open manifolds

Gromov's work [32] provides a parametric (and relative) h–principle for the existence of contact structures on open manifolds. This statement is also proven in [20]. The precise version we will use in the proof of Theorem 2.1 reads as follows:

Proposition 2.3 ([20, p. 7.2.1]). *Let V be a smooth manifold. Let $B \subset A \subset V$ be CW–complexes of positive codimension. Let K be a compact parameter space with $L \subset K$ a closed subset. Consider a continuous family ξ_t, $t \in K$, of formal contact structures (that is, plane fields) on V which are contact in $\mathcal{O}p(B)$ and, if $t \in L$, are contact everywhere.*

Then, there exists a continuous deformation $\xi_{t,s}$, $t \in K$, $s \in [0,1]$, relative to $\mathcal{O}p(B)$ in the domain and to $\mathcal{O}p(L)$ in the parameter, such that $\xi_{t,1}$, $t \in K$ is contact in A.

Observe that the dependence on the parameter can be supposed to be smooth since the condition for a contact structure is open (and thus preserved by small smoothing perturbations). Proposition 2.3 allows us to construct a foliated contact structure in a neighborhood of the 3–skeleton. This will be explained in Section 2.2.

2.1.2 Classification of overtwisted contact structures

For our purposes, we will need a fairly precise version of Proposition 1.25. It reads as follows:

Proposition 2.4 ([19, Theorem 3.1.1]). *Let V be a smooth 3–manifold. Let $A \subset V$ be a CW–complex such that $V \setminus A$ is connected. Let K be a compact space and $L \subset K$ a closed subset.*

Fix a family of discs $\Delta_t \subset V \setminus A$. Consider a continuous family ξ_t, $t \in K$, of almost contact structures on V such that:

- *they are contact in V for $t \in L$,*
- *ξ_t is contact in $\mathcal{O}p(A) \cup \mathcal{O}p(\Delta_t)$,*

- *ξ_t has Δ_t as an overtwisted disc.*

Then, there exists a continuous deformation $\xi_{t,s}$, $t \in K$, $s \in [0,1]$, relative to $\mathcal{O}p(L)$ in the parameter and relative to $\mathcal{O}p(A)$ in the domain, such that $\xi_{t,1}$, $t \in K$, is a family of contact structures on V.

Perhaps the reader is confused about the fact that Δ_t is actually not a fixed disc but a parametric family of them. The point is the following: suppose we are given two K–families of overtwisted contact structures, with corresponding K–families of overtwisted discs. It might happen that the families of structures are formally homotopic but there is an algebraic topology obstruction to connecting the families of overtwisted discs; this is the reason why we fix the overtwisted disc. However, the assumptions of Proposition 2.4 do provide a homotopy between the discs as well.

2.1.3 Triangulating the ambient manifold

2.1.3.1 Construction of the triangulation

Consider a pair $(V^4, \mathcal{F}^3)$. We want to construct a triangulation of V that is nicely adapted to $\mathcal{F}$.

A smooth simplex $\sigma : \Delta \longrightarrow V$ is said to be *linear* with respect to $\mathcal{F}$ if its image is contained in the image of a trivialising foliation chart for $\mathcal{F}$ and $\mathcal{F}$ is transverse to all its faces. In a linear simplex, the height function in the foliation chart yields a function in Δ with one maximum and one minimum in two vertices and no critical points elsewhere. See Figure 2.1.

A triangulation $\mathcal{T}$ of V is **adapted** to the foliation $\mathcal{F}$ if all its simplices are linear with respect to $\mathcal{F}$. This corresponds to a distribution being in general position with respect to a triangulation. There always exists a triangulation on V adapted to $\mathcal{F}$, see [79, Section 5]. The i–skeleton of V with respect to this triangulation $\mathcal{T}$ is denoted by $\mathcal{T}^{(i)}$.

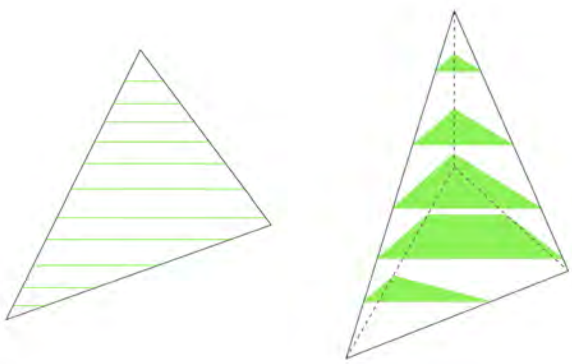

Figure 2.1: A linear 2–simplex (left) and a linear 3–simplex (right) with their induced foliations.

2.1.3.2 Construction of the neighbourhoods

In order to use Propositions 2.3 and 2.4 we require somewhat precise local models and the following lemma provides them:

Lemma 2.5. *Consider a triangulation $\mathcal{T}$ adapted to $(V, \mathcal{F})$. Let σ be a j–simplex, $j = 0, 1, 2, 3, 4$. Let $G \subset \mathcal{T}$ containing $\mathcal{T}^{(j-1)}$ and not containing σ nor any simplex of dimension greater than j.*

Then, there are neighbourhoods $\mathcal{O}p(G)$ and $\mathcal{O}p(\sigma)$, and a diffeomorphism $\phi(\sigma) : \mathcal{O}p(\sigma) \to \mathbb{D}^3 \times [0,1]$ such that the following properties hold:

- $\phi(\sigma)(\sigma) \subset \mathbb{D}^{j-1} \times [0,1]$.
- $\phi(\sigma)^* \mathcal{F} = \mathcal{F}_{\mathbb{D}^3 \times [0,1]}$.
- *For* $j = 0$: $\phi(\sigma)(\mathcal{O}p(G)) = \emptyset$,
- *For* $j = 1$: $\phi(\sigma)(\mathcal{O}p(G)) = \mathbb{D}^3 \times \mathcal{O}p(\{0,1\})$,
- *For* $j = 2, 3, 4$: $\phi(\sigma)(\mathcal{O}p(G)) = (\mathbb{D}^3 \times \mathcal{O}p(\{0,1\})) \cup \mathcal{O}p(\mathbb{S}^{j-2} \times [0,1])$.

See Figure 2.2 for a schematic representation of the statement.

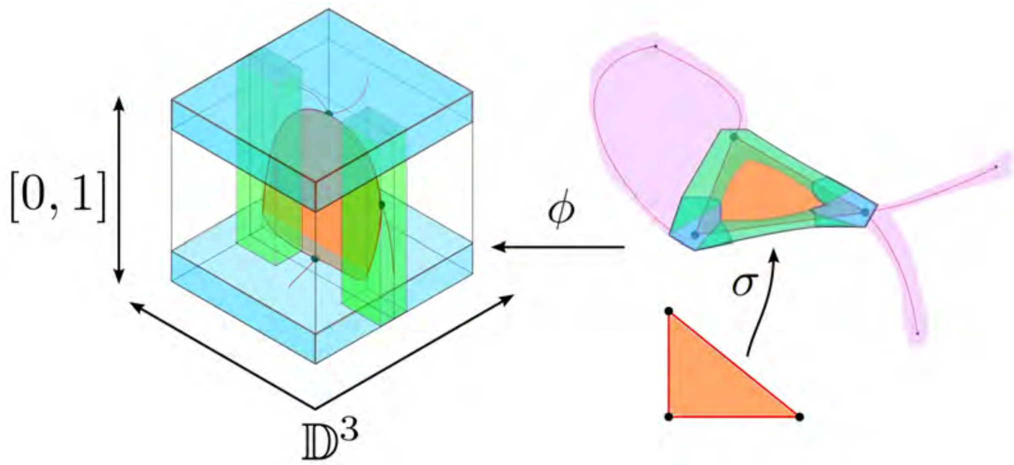

Figure 2.2: Statement of Lemma 2.5 for the case $j = 2$. The figure on the left depicts the local model and the one on the right a neighbourhood of its image in the manifold. The simplex σ is the surface with red boundary. The neighbourhood $\mathcal{O}p(G)$ is colored in pink, and in this example it covers $\tau \in G$, to the left of σ, and two edges connected to the rightmost vertex of σ. The subsets $\mathbb{D}^3 \times \mathcal{O}p(\{0,1\})$ (blue) and $\mathcal{O}p(\mathbb{S}^0 \times [0,1])$ (green) cover the intersection of $\mathcal{O}p(G)$ with the image of the local model.

Proof. In the case $j = 0$, simply take $\phi(\sigma)$ to be a small foliation chart with domain containing σ. G is just a collection of points, so $\mathcal{O}p(\sigma)$ and $\mathcal{O}p(G)$ are disjoint if they are taken to be sufficiently small.

Suppose $j = 1, 2, 3$. By hypothesis, σ is transverse to the foliation. Consider an embedding $i : \mathbb{D}^j \longrightarrow V$ of a closed j–disc extending σ such that $i(\partial \mathbb{D}^j)$ is arbitrarily close to $\partial\sigma$. In particular, we can assume that $\mathcal{O}p(\partial\mathbb{D}^j) = i^{-1}(\mathcal{O}p(G))$ by making the extension small and shrinking $\mathcal{O}p(G)$.

After a small isotopy, we can suppose that $i^*\mathcal{F}$ foliates the disc $\mathbb{D}^j$ horizontally. In order to construct the embedding ϕ we use a normal frame contained in $\mathcal{F}$ along $i(\mathbb{D}^j)$. Then, the exponential map (and rescaling) yields an identification

$$\phi(\sigma) : \mathcal{O}p(\sigma) \to \mathbb{D}^j \times \mathbb{D}^{4-j} \cong \mathbb{D}^3 \times [0,1], \quad \phi^{-1}|_{\mathbb{D}^j \times \{0\}} = i.$$

By construction, this map can be assumed to satisfy $\phi^*\mathcal{F} = \mathcal{F}_{\mathbb{D}^3 \times \{t\}}$.

Since $i(\partial\mathbb{D}^j) \subset \mathcal{O}p(G)$ by construction, the exponential map used to define the chart $\phi(\sigma)$ can be used for a very short time to guarantee

$$(\mathbb{D}^3 \times \mathcal{O}p(\{0,1\})) \cup \mathcal{O}p(\mathbb{S}^{j-2} \times [0,1]) \subset \phi(\sigma)(\mathcal{O}p(G)).$$

To ensure that an strict equality is obtained, $\mathcal{O}p(G)$ can be slightly adjusted.

Lastly, in the case $j = 4$, the triangulation being adapted means that we can find $\phi(\sigma)$ with domain a smoothing of σ. By adjusting $\mathcal{O}p(G)$ the claimed properties follow. □

2.1.4 Constructing a system of transverse segments

Without spoiling much of it, let us motivate the next technical result. The triangulation $\mathcal{T}$ and the charts $\{\phi(\sigma)\}$ that we just constructed provide a nice way of translating our problem about contact foliations into a problem about 1–parametric families of (formal) contact structures in $\mathbb{D}^3$. The special form they have allows us to apply Proposition 2.3 to homotope our starting formal contact foliation to one that is contact along the 3–skeleton.

At that point, in the local model $\phi(\sigma) : \mathcal{O}p(\sigma) \to \mathbb{D}^3 \times [0,1]$ associated to 4–simplex σ, the formal contact foliation has already been homotoped to a contact foliation along the boundary. If we want to apply Eliashberg's theorem (Proposition 2.4) to further homotope it in the interior, we need to produce a family of overtwisted discs (one in each leaf $\mathbb{D}^3 \times \{t\}$). The way in which we do it is by taking a segment that is transverse to each of the leaves of the model and whose ends lie inside some other simplex. We will then deform the formal contact foliation along it, introducing a Lutz twist in all the leaves in $\mathcal{O}p(\sigma)$, and interpolating back to the original structure near its ends.

Lemma 2.6. *Consider a 4–simplex $\sigma \in \mathcal{T}^{(4)}$ and let $\phi(\sigma)$ be the corresponding chart. There exists a map $\gamma(\sigma) : [-1,2] \to V$ transverse to $\mathcal{F}$ such that:*

- $\gamma(\sigma)(t) \in \phi(\sigma)(\mathbb{D}^3 \times \{t\})$ *for* $t \in [0,1]$,
- $\gamma_\sigma(\mathcal{O}p(\{-1,2\}))$ *is disjoint from* $\mathcal{O}p(\mathcal{T}^{(3)})$.

Also, $\gamma(\sigma)$ and $\gamma(\tau)$ are disjoint if $\tau \in \mathcal{T}^{(4)}$ is different from σ.

Proof. Set $\gamma(\sigma)(t) = \phi(0,t)$ if $t \in [0,1]$. Then, extend this embedding to $t \in [1,2]$ by prolonging it into one of the 4–simplices whose minimum vertex is the maximum vertex of ϕ. Proceed analogously to define it in $t \in [-1,0]$.

Perturb the embeddings $\gamma(\sigma)$ so that they do not intersect each other and $\gamma(\sigma)$ intersects $\mathcal{T}^{(3)}$ only in two points close to $t = 0, 1$. □

Lemma 2.7. *Let $\gamma(\sigma) : [-1, 2] \to V$, $\sigma \in \mathcal{T}^{(4)}$ be a collection of transverse segments as provided by Lemma 2.6.*

There exists a collection of diffeomorphisms

$$\kappa(\sigma) : \mathbb{D}^3 \times [-1, 2] \to \mathcal{O}p(\gamma(\sigma)),$$

and a neighbourhood $\mathcal{O}p(\mathcal{T}^{(3)})$ of the 3–skeleton, conforming to the following properties:

- $\kappa(\sigma)^*\mathcal{F} = \mathcal{F}_{\mathbb{D}^3\times[-1,2]}$,
- *all the $\kappa(\sigma)$ are disjoint,*
- $\gamma(\sigma)(t) = \kappa(\sigma)(0, t)$,
- *there exists a union of closed intervals $I(\sigma) \subset [-1, 2]$ such that*

$$\kappa(\sigma)^{-1}(\mathcal{O}p(\mathcal{T}^{(3)})) = \mathbb{D}^3 \times I.$$

Proof. Any small thickening of the $\gamma(\sigma)$ satisfies the properties. □

2.2 The reduction process

The following proposition packages the fact that we are able to achieve contactness close to the 3–skeleton:

Proposition 2.8. *Let ξ be a foliated almost contact structure on $(V, \mathcal{F})$. There exists a homotopy ξ_s, $s \in [0, 1]$, of foliated almost contact structures such that $\xi_0 = \xi$ and ξ_1 is a contact foliation in $\mathcal{O}p(\mathcal{T}^{(3)})$.*

Proof. Order the simplices of $\mathcal{T}$ increasingly in the dimension, up to and including dimension 3. Suppose that a given inductive step the homotopy has been performed in $\mathcal{O}p(G)$, with G some collection of simplices respecting the order. Let σ be the next simplex. Lemma 2.5 provides a foliation chart:

$$\phi(\sigma) : \mathcal{O}p(\sigma) \to \mathbb{D}^3 \times [0, 1].$$

The foliated almost contact structure $\phi(\sigma)_*\xi_0$ can be considered as a 1–parametric family $\widetilde{\xi}_t$, $t \in \mathbb{D}^3$, of almost contact structures in the disc $\mathbb{D}^3$.

If σ is zero dimensional, we can apply Proposition 2.3 to $\widetilde{\xi}_t$ with $K = [0, 1]$, $V = D^3$, $A = \{0\}$, $B = \emptyset$, and $L = \emptyset$. This can be used to deform ξ in $\mathcal{O}p(\sigma)$ to a formal contact foliation that is contact close to σ.

Suppose instead that σ is j–dimensional with $j = 1, 2, 3$. Then Proposition 2.3 applies similarly with $K = [0, 1]$, $V = D^3$, $A = \mathbb{D}^{j-1}$, $B = \mathbb{S}^{j-2}$, and $L = \mathcal{O}p(\{0, 1\})$. This inductive procedure implies the statement. □

2.3 The extension problem

The extension process has three steps: first, we perform a first homotopy so that our formal contact foliation is honestly contact close to the system of transverse segments produced by Lemmas 2.6 and 2.7. This will then allow us to perform a parametric Lutz twist along them. The subtle point in this is that we do not have closed curves, but segments. We will have to define the so–called Lutz twist to deal with this. Finally, we will use Proposition 2.4, using the overtwisted discs we just introduced, to conclude the proof.

2.3.1 Homotopy along the system of transverse segments

Let σ be a top–dimensional simplex. Let $\phi(\sigma) : \mathcal{O}p(\sigma) \to \mathbb{D}^3 \times [0,1]$ be the corresponding flowbox chart. Let $\gamma(\sigma) : [-1,2] \to V$ be the associated segment and let $\kappa(\sigma) : \mathbb{D}^3 \times [-1,2] \to \mathcal{O}p(\gamma(\sigma))$ be its thickening, as in Lemmas 2.6 and 2.7. Our assumption at this point, thanks to Proposition 2.8, is that the formal contact foliation ξ is honestly contact in $\mathcal{O}p(\mathcal{T}^{(3)})$.

Regard $\kappa(\sigma)^*\xi$ as a 1–parametric family of contact structures ξ_t in $\mathbb{D}^3$. We can apply Proposition 2.3 to them with $K = [-1,2]$, $A = \{0\}$, $B = \emptyset$, $L = I(\sigma)$. Recall that $I(\sigma)$ is the interval provided by Lemma 2.7. This allows us to assume that ξ is honestly contact also along the curves $\gamma(\sigma)$.

2.3.2 The vanishing Lutz twist

Recall Lemma 1.26. It allows us to find a new trivialisation

$$\phi = \phi(\gamma(\sigma)) : \mathcal{O}p(\gamma(\sigma)) \to \mathbb{D}^3_\varepsilon \times [-1,2]$$

satisfying $\phi(\gamma(\sigma))^*\xi = \ker(dz - xdy) = \xi_{\text{std}}$. Let $\eta \subset (\mathbb{D}^3_\varepsilon, \xi_{\text{std}})$ be some transverse knot.

Let $\chi : [-1,2] \to [0,1]$ be a bump function satisfying $\chi = 0$ in $\mathcal{O}p(\{-1,2\})$ and 1 slightly away from it. Let $\widetilde{\xi}_s$, $s \in [0,1]$, be a homotopy (as plane fields) between ξ_{std} and the structure ξ_{Lutz} obtained from it by performing Lutz twist along η. Then, we can define a new contact foliation in V that agrees with ξ outside of each $\mathcal{O}p(\gamma(\sigma))$ and inside it is given, leafwise, by $\phi(\gamma(\sigma))^*\widetilde{\xi}_{\chi(t)}$. In particular, let $\Delta \subset \mathbb{D}^3_\varepsilon$ be an overtwisted disc for ξ_{Lutz}. Then, we can define

$$\Delta(\sigma)_t = \phi(\sigma) \circ \phi(\gamma(\sigma))^{-1}(\Delta \times \{t\}), \quad t \in [0,1],$$

which is a family of overtwisted discs seen from the perspective of the chart $\phi(\sigma)$.

We call this the **vanishing Lutz twist**, and it is exactly like a parametric Lutz twist, but we interpolate back to the original structure at the ends. In particular, this destroys the contact condition in $\gamma(\sigma)(\mathcal{O}p(\{-1,2\}))$, but this is okay, since these regions are in the interior of the 4–cells. Let us abuse notation and denote by ξ the formal contact structure we have just produced; it is of course homotopic to the one we started with.

2.3.3 Solving the extension problem

Now we can apply Proposition 2.4 to $\phi(\sigma)_*\xi$, with $K = [0,1]$, $V = \mathbb{D}^3$, $A = \mathbb{S}^2$, and Δ_t as constructed above. This solves the extension problem relative to the boundary and concludes the proof of Theorem 2.1 at the level of π_0. $\square$

2.4 Proof of the existence theorem

In the previous section we concluded the proof of Theorem 2.1 at the level of connected components. We will now address the main points on how to adapt it to the higher homotopy groups. Let P be some compact parameter space and let $\{\xi_p\}_{p\in P}$, be a P–family of formal contact foliations.

I. Since the parameter space P is compact, all the statements regarding the selection of neighbourhoods at each step of the proof go through as in the non–parametric case by taking their intersection. Now the family ξ_p can be understood as a $[0,1]\times P$ family of formal contact structures every time we push them forward using a chart $\phi(\sigma) : \mathcal{O}p(\sigma) \to \mathbb{D}^3 \times [0,1]$.

II. In the zero skeleton we proceed as in Proposition 2.8, by setting $K = [0,1]\times P$, $V = \mathbb{D}^3$, $A = \{0\}$, $B = \emptyset$, and $L = \emptyset$ when we apply Proposition 2.3. Similarly, for the case $j \in \{1,2,3\}$, define $K = I \times P$, $V = \mathbb{D}^3$, $A = D^{j-1}$, $B = \mathbb{S}^{j-2}$, and $L = \mathcal{O}p(\{0,1\}) \times P$.

III. The system of transverse segments does not have to depend on parameters, and we can produce a first homotopy close to them exactly as in the non–parametric case by setting $K = [-1,2]\times P$, $A = \{0\}$, $B = \emptyset$, $L = I(\sigma)$ in Proposition 2.3. Similarly, the vanishing Lutz twist works in the same way.

IV. Lastly, we apply Proposition 2.4 with $K = [0,1]\times P$, $V = \mathbb{D}^3$, $A = \mathbb{S}^2$, and Δ_t now a $[0,1]\times P$– overtwisted discs.

These four steps conclude the proof. □

2.4.1 A corollary regarding even–contact structures

The following corollary of Theorem 2.1 was pointed out to me by J. Bowden:

Corollary 2.9. *Given $(V^4, \mathcal{F}^3)$ admitting a formal contact foliation, there exists an even–contact structure whose kernel is transverse to $\mathcal{F}$.*

Proof. Indeed, the contact foliation produced by the h–principle is exactly of this form. □

As we pointed out in the Preamble, even–contact structures abide to the h–principle, so there are no interesting phenomena to be detected regarding their classification. However, this does not mean that they are uninteresting. Indeed, not much is known about the *dynamical properties of their kernel*. The corollary can be understood as a mild result in this direction.

Chapter 3

Donaldson techniques for strong symplectic foliations

In Chapter 1 we introduced the notion of strong symplectic foliation as a generalisation of 3–dimensional tautness. From the definition of tautness we gave, it is not hard to see that a (smooth) foliation is taut if and only if there is a transverse loop that intersects all the leaves. Is there an analogous notion in the strong symplectic setting?

Remark 3.1. In this chapter a strong symplectic foliation $(M^{2n+1}, \mathcal{F}^{2n}, \omega)$ is a triple where ω is already an extension. We do this because the extension is the central object in all the constructions we carry out.

Consider the following obvious definition:

Definition 3.2. *Let $(M^{2n+1}, \mathcal{F}^{2n}, \omega)$ be a (strong) symplectic foliation. A submanifold $W \hookrightarrow M$ is said to be a* ***(strong) symplectic foliated submanifold*** *if it is everywhere transverse to $\mathcal{F}$ and $\omega|_W$ is a symplectic form in each leaf of $\mathcal{F}|_W$.*

Codimension–2 strong symplectic foliated submanifolds that intersect every leaf of $\mathcal{F}$ can be regarded then as generalisations of the phenomenon above. Then, there are two natural questions: does every strong symplectic foliation admit one such submanifold? And, conversely, suppose that a (weak) symplectic foliation admits a W whose (weak) symplectic foliated structure can be homotoped to be strong: can this be extended to a global homotopy to a strong symplectic foliation?

A positive answer to both questions can be understood as an analogue of the Rummler–Sullivan theorem. In this chapter we will address the first question. The second question is at the moment open and, probably, not even well–posed.

3.1 Statement of the results and some history

Half of the story is already contained in the following theorem of A. Ibort and D. Martínez–Torres:

Proposition 3.3 ([43, Corollary 1.2])**.** *Let $(M^{2n+1}, \mathcal{F}^{2n}; \omega)$ be a strong symplectic foliation on a closed manifold, with ω of integral class. Then, for any integer k large enough, there are strong symplectic foliated submanifolds W_k^{2n-1} representing the Poincaré dual of $[k\omega]$.*

Additionally, the maps

$$i_* : \pi_j(W_k) \to \pi_j(M)$$
$$i_* : H_j(W_k, \mathbb{Z}) \to H_j(M, \mathbb{Z})$$

are isomorphisms for $j < n-1$ and surjections for $j = n-1$.

Using the proposition the following corollary follows:

Corollary 3.4. *Let $(M^{2n+1}, \mathcal{F}^{2n}; \omega)$ be a strong symplectic foliation with M closed. It admits strong symplectic foliated submanifolds of all even codimensions.*

Proof. Take a collection of closed 2–forms whose cohomology classes generate $H^2_{DR}(M)$. By adding a linear combination of them to ω, we produce a closed 2–form ω' of rational class which then can be scaled to be integral; if the perturbation we add is small enough, ω' is still strong symplectic. It can be shown (but it is not immediate from the statement of Proposition 3.3) that the submanifolds W_k^{2n-1} we produce for ω' are strong symplectic for ω as well. This is a consequence of the fact that we have precise estimates on how ω' restricts to each W_k^{2n-1} (we will elaborate more later in this section, once we have introduced the appropriate language). $\square$

The submanifolds W_k, as produced by Proposition 3.3, are called **(Donaldson) divisors**. In [17], S. Donaldson introduced *approximately holomorphic techniques* in symplectic topology. The idea is as follows: in *projective algebraic geometry*, complex codimension–1 subvarieties (the so–called divisors) arise naturally as zeroes of (holomorphic) sections of complex line bundles. In symplectic geometry, one can mimick this process. Naturally, the idea of a holomorphic section does not make sense anymore but, by taking progressively higher powers of a fixed line bundle, one is able to construct sequences of sections whose behaviour is increasingly more complex–linear. In particular, their zeroes eventually become symplectic submanifolds.

A. Ibort, D. Martínez–Torres, and F. Presas [41] extended these techniques to the *contact* setting. Then, they were adapted to strong symplectic foliations in [43], where Proposition 3.3 was proven. The construction goes through because both families of objects can be understood as 1–parametric analogues of symplectic structures and, much like generic paths of projective varieties are smooth, generic paths of symplectic divisors are smooth too.

The second part of Proposition 3.3 is the *foliated Lefschetz hyperplane theorem*. It states that the divisors recall some of the topology of the ambient space. However, this does not answer the question we posed: can we make W_k^{2n-1} intersect every leaf of $\mathcal{F}$? The main result of this chapter says that yes:

Theorem 3.5. *Let $(M^{2n+1}, \mathcal{F}, \omega)$ be a strong symplectic foliation on a closed manifold. Let W be a Donaldson divisor of dimension $2n-1$. Then, for every leaf $\mathcal{L}$ of $\mathcal{F}$ it holds that:*

$$\pi_k(\mathcal{L}, \mathcal{L} \cap W) = \{1\},\ 0 \leq k \leq n-1.$$

In particular, $\pi_0(\mathcal{L}, \mathcal{L} \cap W) = 0$, proving our claim. It should be remarked that precisely the π_0 case had already been proven in [51], where the question was tackled constructing Lefschetz pencils. However, the method of proof we present here is different and yields a simpler proof. It first appeared in my article [53] with D. Martínez–Torres and F. Presas.

If $2n+1 \geq 5$, Theorem 3.5 produces a divisor that is a classical 3–dimensional taut foliation. In this case, we have that $\pi_1(\mathcal{L}, \mathcal{L} \cap W) = 0$, which implies that the map induced on leaf spaces $W/\mathcal{F}_W \to M/\mathcal{F}$ is a homeomorphism, and that both foliations have the same *transverse geometry*.

3.2 Ingredients of the proof

Our proof of Theorem 3.5 follows Donaldson's proof of the Lefschetz hyperplane theorem for approximately holomorphic divisors. His proof followed the Andreotti-Frankel proof in the affine/projective case: in the complement of a divisor, the modulus of its defining approximately holomorphic section can be regarded, after a small perturbation, as a Morse function with critical points of index at least n, which implies that the ambient manifold is obtained from the divisor by attaching handles of index at least n. It readily follows that the relative homology and homotopy groups of degree less than n vanish.

In the foliated case, however, the critical points of this function come in $\mathbb{S}^1$–families and a non–compact leaf will, in general, have infinitely many critical points. Hence, the relative homotopy type of a leaf with respect to the divisor is not readily understood.

We shall first review the essentials of the approximately holomorphic machinery that we will need for the proof of Theorem 3.5. Then, we will discuss some conditions for Morse functions in open manifolds guaranteeing a nice behaviour of their gradient flow.

3.2.1 Approximately holomorphic theory for strong symplectic foliations

Let M^{2n+1} be a closed manifold endowed with a strong symplectic foliation $(\mathcal{F}, \omega)$. After a small perturbation, we may assume without loss of generality that $[\omega]$ is a rational class; by scaling the class, we may also assume that it is integral. We let $L \to M$ be the pre–quantum line bundle associated to ω; this is a Hermitian line bundle with a compatible connection ∇ whose curvature is $-2\pi i\omega$.

We let $\nabla^{\mathcal{F}}$ denote the component of ∇ tangential to $\mathcal{F}$. After choosing an almost complex structure J compatible with ω, the tangential connection can be further decomposed into its complex linear and antilinear parts, yielding $\nabla^{\mathcal{F}} = \partial + \bar{\partial}$.

According to [43, Corollary 1.2], upon choosing the almost complex structure J, it is possible to construct a family $s_k : M \to L^k$ of sections of the k–th tensor powers of L, for k large enough, such that $W_k := s_k^{-1}(0)$ are closed, strong symplectic submanifolds of codimension two.

To state the conditions that are required for the sequence s_k, we fix a metric g on M which is required to satisfy $g = \omega(\cdot, J\cdot)$ leafwise. Further, we define a family of scaled metrics $g_k = kg$.

Definition 3.6.

1. *A sequence of sections $s_k : M \to L^k$ is said to be **approximately holomorphic** if there is a universal constant $C > 0$ such that:*
$$|s_k|_{g_k}, |\nabla s_k|_{g_k} < C; \qquad |\bar{\partial} s_k|_{g_k}, |\nabla\bar{\partial} s_k|_{g_k} < Ck^{-1/2},$$
for k large enough.
2. *A sequence of sections $s_k : M \to \mathcal{L}^k$ is said to be ν–**transverse** to zero along the foliation $\mathcal{F}$ if at any point either $|s_k|_{g_k} \geq \nu$ or $|\nabla^{\mathcal{F}} s_k|_{g_k} \geq \nu$.*

To every approximately–holomorphic, transverse to zero sequence s_k one associates the sequence of functions
$$f_k : M \setminus W_k \to \mathbb{R}$$

$$f_k = \log |s_k|^2.$$

A simple computation [17, 43] shows, just like in the non–foliated setting:

Proposition 3.7. *Let $\mathcal{L}$ be a leaf. The function*

$$f_k = \log |s_k|^2 : \mathcal{L} \setminus (W_k \cap \mathcal{L}) \to \mathbb{R},$$

might not be Morse. However, its critical points have index at least n.

This proposition, when applied to a closed leaf, implies Theorem 3.5 immediately. The case of interest is, of course, when $\mathcal{L}$ is open.

3.2.2 Gradient flows and the topology of open manifolds.

The study of flows which behave well on open manifolds appears in the literature on foliation theory [23]. For the sake of completeness, we review these facts tailored to the applications we have in mind.

Let f be a Morse function on a manifold $\mathcal{L}$. For any $a \in \mathbb{R}$ set $\mathcal{L}_a = \{x \in \mathcal{L} \,|\, f(x) \leq a\}$, and denote by $\mathrm{Crit}_a(f)$ the subset of critical points of f lying in $\mathcal{L}\backslash\mathcal{L}_a$.

Let a be a regular value for f and let $b > a$. Assume for the moment that $\mathcal{L}$ is compact. It is customary to study the relative topology of the pair $(\mathcal{L}_b, \mathcal{L}_a)$ using minus the gradient flow of f with respect to some fixed metric g. The key point is that the following dichotomy holds: for any $x \in \mathcal{L}_b\backslash\mathcal{L}_a$ the trajectory of $-\nabla_g f$ starting at x either enters $\mathcal{L}_a$ in finite time, or converges to one of the finitely many critical points in $\mathrm{Crit}_a(f)$.

If $\mathcal{L}$ is no longer compact but f is proper, then of course the study of the relative topology of the pair $(\mathcal{L}_b, \mathcal{L}_a)$ goes exactly as in the compact case. There might be cases –as in our setting coming from approximately holomorphic geometry– that the natural Morse functions to be used are not proper, and one needs to impose an appropriate form of the above dichotomy for trajectories of $-\nabla_g f$:

Lemma 3.8. *Let f be a Morse function on a manifold $\mathcal{L}$ and let g be a metric on $\mathcal{L}$ so that $\nabla_g f$ is complete. Let a be a regular value, $b > a$, and assume that the following holds:*

I. *For every compact subset $X \subset \mathcal{L}_b$, there exist finitely many critical points $c_1, \ldots, c_{i_X}$ in $\mathrm{Crit}_a(f)$ such that the following dichotomy holds: a trajectory of $-\nabla_g f$ starting at $x \in X$ either reaches $\mathcal{L}_a$ in finite time, or converges to a critical point in $\{c_1, \cdots, c_{i_X}\}$.*

II. *Every $c \in Crit_a(f)$ has index $\geq j$.*

Then, we have that $\pi_k(\mathcal{L}_b, \mathcal{L}_a) = 0$, for $k = 0, \ldots j-1$.

Proof. Let us start by making the following observation: if X is as in assumption (I.) and the collection $\{c_1, \ldots, c_{i_X}\}$ is empty, then we claim that X is taken in finite time to $\mathcal{L}_a$ by the flow ϕ of $-\nabla_g f$. Indeed, for every $x \in X$ there exists a time $t_x > 0$ such that $f(\phi_{t_x}(x)) < a$; further, since for fixed t, ϕ_t is continuous, there is a small ball $B_g(x, \varepsilon_x)$ centered at x such that $\phi_{t_x}(B_g(x, \varepsilon_x)) \subset \mathcal{L}_a$. Then, the result follows by compactness of X.

Now, let N be a compact manifold and $h\colon (N, \partial N) \to (\mathcal{L}_b, \mathcal{L}_a)$ be a smooth map. Let U be a relatively compact neighborhood of $h(N)$. Then assumption (I.) implies that trajectories starting at points in $\bar{U}$ can only enter $\mathcal{L}_a$ in finite time or converge to one of the finitely many critical points $\{c_1, \dots, c_{i_{\bar{U}}}\}$.

Observe that there is a small relatively compact neighborhood V of $h(\partial N)$ such that the flow of $-\nabla_g f$ sends V into $\mathcal{L}_a$: this follows if $V \subset U$ is selected so that $f(V)$ lies below the critical values $\{f(c_1), \dots, f(c_{i_{\bar{U}}})\}$.

We now construct h', an arbitrarily small perturbation of h relative to V. Proceeding inductively over the finite list $\{c_1, \dots, c_{i_{\bar{U}}}\}$, as in [58], we obtain h' that is transverse to the ascending disks of the critical points and that satisfies $h'(N) \subset U$.

If N has dimension at most $j-1$ then, by hypothesis (II.), transversality to the ascending disks means empty intersection. The hypotheses of the claim at the start of the proof are satisfied and it follows that $\pi_k(\mathcal{L}, \mathcal{L}_a) = 0$, for $k = 0, \dots j-1$. □

The following result describes quantitative conditions on the gradient vector field granting the dichotomy in point (I.) of Lemma 3.8.

Proposition 3.9. *Let f be a Morse function, g be a complete metric on $\mathcal{L}$, and $a < b \in \mathbb{R}$. Assume that there exist real constants $D, E > 0$ and open subsets $\mathcal{C}_i \subset \mathcal{L}_b$, $i \in I$, such that:*

a. For any pair $i, i' \in I$, $i \neq i'$, we have $d_g(\mathcal{C}_i, \mathcal{C}_{i'}) > D$.

b. The diameter of the sets $\mathcal{C}_i$ is at most E.

c. There exist real numbers $\delta_1, \delta_2 > 0$, such that

$$\delta_2 \geq |\nabla_g(f)(p)| \geq \delta_1, \forall p \in \mathcal{L}_b \setminus \left(\bigcup_{i \in I} \mathcal{C}_i \right).$$

Then $-\nabla_g f$ is complete and the dichotomy in point (I.) of Lemma 3.8 for $-\nabla_g f$ holds.

Essentially, the proposition states that the critical points of f come in families, indexed by I and contained in the sets $\mathcal{C}_i$, that are far from each other. In order to prove Proposition 3.9, let us introduce some notation and prove an auxiliary lemma. Given any $x \in \mathcal{L}$, we denote by γ_x the positive half of the flow line that contains x. Denote by ϕ_t the flow of f at time t. Let γ_x^t designate the segment of the curve γ_x between x and $\phi_t(x)$. Then:

Lemma 3.10. *Under the assumptions of Proposition 3.9, there is a constant R, independent of $t \in \mathbb{R}$ and $x \in \mathcal{L}_b$, such that $d_g(\phi_t(x), x) > R$ implies $f(\phi_t(x)) < a$.*

Proof. For every curve γ we denote by $\tilde{\gamma}$ the (possibly disconnected) curve:

$$\tilde{\gamma} = \left\{ p \in \gamma : p \notin \bigcup_{i \in I} \mathcal{C}_i \right\},$$

that is, the union of segments of γ that are disjoint from the sets $\mathcal{C}_i$.

Given any curve $\gamma \subset B(x, R)$ starting at x and intersecting the boundary of $B(x, R)$ at y, we can associate to it another curve, which we denote by $\eta = \eta_\gamma$, using the following procedure:

1. list, in order, all the sets $\mathcal{C}_i$ that γ intersects. Remove all the consecutive repetitions of the same $\mathcal{C}_i$, listing just the first one in each series of repetitions. Write $\{\mathcal{C}_{i_j}\}_{j\in[1,..k]}$ for this finite list,

2. mark the entry and exit points e_j and f_j of γ into each $\mathcal{C}_{i_j}$. In the case of consecutive repetitions of the same $\mathcal{C}_i$, just mark the first entry point and the last exit point of the series. For simplicity, denote $f_0 = x$ and $e_{k+1} = y$,

3. call η the piecewise smooth curve formed by connecting these marked points in the order they appear. From e_j to f_j, take the shortest geodesic between the two points. From f_j to e_{j+1}, take the shortest path not intersecting any C_i. Denote these paths by $l(e_j, f_j)$ and $l(f_j, e_{j+1})$ respectively.

Assume $R > E + D$. If $k = 0, 1$, it is immediate that

$$\frac{\text{length}(\tilde{\eta})}{\text{length}(\eta)} \geq \frac{D}{E+D},$$

otherwise, the following estimate holds:

$$\frac{\text{length}(\tilde{\eta})}{\text{length}(\eta)} = \frac{\sum_{j=0}^{k} \text{length}(l(f_j, e_{j+1}))}{\sum_{j=0}^{k} \text{length}(l(f_j, e_{j+1})) + \sum_{j=1}^{k} \text{length}(l(e_j, f_j))} \geq$$

$$\frac{\sum_{j=1}^{k-1} \text{length}(l(f_j, e_{j+1}))}{\sum_{j=1}^{k-1} \text{length}(l(f_j, e_{j+1})) + kE} \geq \frac{(k-1)D}{(k-1)D + kE} \geq \frac{D}{2(E+D)}.$$

For any radius $r > E + D$, denote by τ the time at which the curve γ_x first intersects $\partial B(x, r)$. Denote this intersection point by y. Consider the segment γ_x^τ and its associated curve $\eta = \eta_{\gamma_x^\tau}$. Use the fact that over $\tilde{\gamma}_x^\tau$ we have a lower bound for the gradient $|\nabla_g f| > \delta_1 > 0$:

$$|f(y) - f(x)| \geq \delta_1 \text{length}(\tilde{\gamma}_x^\tau) \geq \delta_1 \text{length}(\tilde{\eta}) \geq \delta_1 \text{length}(\eta) \frac{D}{2(E+D)} \geq r \frac{\delta_1 D}{2(E+D)}$$

which implies that, if r is taken to be large enough, $|f(y) - f(x)| > b - a$, and hence $y \in M_a$. $\square$

Proof of Proposition 3.9. Let $X \subset \mathcal{L}_b$ be a compact set. Let R be the universal constant given by Lemma 3.10. Denote by $X(R)$ the R-neighborhood of X, which is a relatively compact set. Lemma 3.10 implies that any trajectory starting at X either reaches the interior of $\mathcal{L}_a$ – which is equivalent to saying that it reaches $\mathcal{L}_a$ in finite time – or it remains in $X(R)$ for all time.

It must be shown that if a trajectory γ_x remains within $X(R)$ for all times then it must converge to a critical point. Since $X(R)$ is relatively compact and f is a Morse function, there is a finite number k of critical points in its closure. Each of those critical points $\{c_i\}_{i=1}^k$ has an arbitrarily small neighborhood V_i which corresponds to a ball in the standard Morse model around c_i. In particular, a trajectory that intersects V_i must intersect just once, either converging to c_i or escaping from V_i eventually. From this it follows that there is a time $t_0 > 0$ such that $\gamma_x(t) \notin V_i$, for all $t > t_0$ and every i. Since the gradient $|\nabla_g f| > \delta > 0$ is bounded from below in $X(R) \setminus \cup_{i=1..k} V_i$, this shows that $f(\gamma_x(t)) < a$ for t large enough, which is a contradiction. $\square$

3.3 Proof of the Lefschetz hyperplane theorem for leaves

Fix some leaf $\mathcal{L} \in M$. All we need to do now is to check that, for a suitable choice of Morse function and metric on $\mathcal{L}$, the hypotheses of Proposition 3.9 are satisfied for $\mathcal{L}$. Our candidate is the restriction to the leaf of the function $f_k = \log |s_k|^2$, and the restriction to the leaf of any Riemannian metric on M.

We shall prove a couple of preliminary lemmas, for which we need to recall some notation. Given a function f, defined on a manifold endowed with a codimension one foliation $(M, \mathcal{F})$, the tangential differential $d^{\mathcal{F}} f$ is the composition of the differential with the projection $T^*M \to (T\mathcal{F})^*$. The points in which $d^{\mathcal{F}} f$ vanishes are the *tangential critical points of* f, which we denote by $\Sigma^{\mathcal{F}}(f)$. Of course, $\Sigma^{\mathcal{F}}(f)$ are nothing but the critical points of the restriction of f to each leaf of $\mathcal{F}$.

Lemma 3.11. *For every k large enough, the strong symplectic foliated submanifold $W_k \subset M$ has a tubular neighborhood that contains a full regular level set of $f_k = \log |s_k|^2$ and which is also disjoint from $\Sigma^{\mathcal{F}}(f_k)$.*

Proof. It is enough to check that $h_k = ||s_k||^2$ satisfies the Lemma, since log is an increasing monotone function.

We claim that the neighborhood $U = \{x \in M \,|\, ||s_k(x)|| < \nu\}$ of the submanifold W_k does not intersect $\Sigma^{\mathcal{F}}(f_k)$. Assume that $p \in U$. By the ν–transversality along $\mathcal{F}$ of the section s_k, there is a unitary vector field $v \in T_p\mathcal{F}$ such that $||\nabla_v s_k(p)|| \geq \nu$. By asymptotic holomorphicity, for k large, we have that the unitary vector field $Jv \in T_p\mathcal{F}$ satisfies $||\nabla_{Jv} s_k(p) - i\nabla_v s_k(p)|| = O(k^{-1/2})$. Therefore, the map $\nabla^{\mathcal{F}} s_k(p)$ is surjective. We conclude that $p \notin \Sigma^{\mathcal{F}}(f_k)$. □

Lemma 3.12. *Let $\mathcal{L}$, a leaf of $\mathcal{F}$, be fixed. After a perturbation of the sequence s_k, preserving transversality to zero and approximately holomorphicity, it can be assumed that:*

1. *the restrictions of the f_k to $\mathcal{L}$ are Morse functions.*
2. *$\Sigma^{\mathcal{F}}(f_k)$ is a finite union of disjoint circles in general position with respect to $\mathcal{F}$. Their tangency points are turning points, i.e., birth-death type singularities for the restriction of f_k to the corresponding leaf.*

Proof. According to [23], after an arbitrarily small C^r perturbation, $r \geq 2$, the set of tangential critical points $\Sigma^{\mathcal{F}}(f_k)$ can be assumed to fit into a 1-dimensional manifold that is transverse to $\mathcal{F}$ everywhere but at the finite collection of turning points $c_1, \dots, c_d$. Every other point is a non–degenerate critical point for the restriction of f_k to the corresponding leaf. The turning points satisfy the following relevant property: in a small foliated chart, a plaque not containing the turning point intersects $\Sigma^{\mathcal{F}}(f_k)$ either in the empty set or in two tangential critical points.

Assertion (I.) in the Lemma follows by showing that none of the $c_1, \dots, c_d$ belong to the fixed leaf $\mathcal{L}$: if any of them do, a C^r–small isotopy, transverse to $\mathcal{F}$ at the turning point, can be used to move it to a nearby leaf. This is described in detail in [23].

These C^r–perturbations of f_k can be taken to be the result of a C^r–perturbation of s_k. Indeed, let ε_k be a C^r–perturbation of f_k. The function ε_k can be assumed to be identically zero away from an arbitrary small neighborhood of $\Sigma^{\mathcal{F}}(f_k)$ so, by Lemma 3.11, the following expression is well defined:

$$\tilde{s}_k = s_k \sqrt{1 + \varepsilon_k / f_k},$$

since f_k is bounded from below in the support of ε_k. It is clear that

$$||\tilde{s}_k|| = f_k + \varepsilon_k.$$

The asymptotic holomorphicity of the sequence $\tilde{s}_k$ can be readily checked:

$$\nabla \tilde{s}_k = \nabla s_k \sqrt{1 + \varepsilon_k / f_k} + s_k \frac{f_k \nabla \varepsilon_k - \varepsilon_k \nabla f_k}{2 f_k^2 \sqrt{1 + \varepsilon_k / f_k}},$$

where the second term is C^r–small and the first is C^r–close to ∇s_k. A similar computation for the higher order derivatives concludes the claim. □

We can finally address the proof of the theorem.

Proof of Theorem 3.5. Fix a leaf $\mathcal{L}$ and assume that we have all the data needed for developing approximately–holomorphic geometry in M^{2n+1}. The metrics g_k induce complete metrics in $\mathcal{L}$. Given an approximately–holomorphic sequence s_k, with corresponding Donaldson-type submanifolds W_k, an application of Lemma 3.12 yields a new approximately–holomorphic sequence, still denoted by s_k, that induces Morse functions $(f_k)_{|\mathcal{L}}$ in $\mathcal{L} \setminus W_k$.

By Lemma 3.11, W_k has a ε–neighborhood containing a regular level a_k. Lemmata 3.11 and 3.12 together mean that $\Sigma^{\mathcal{F}}(f_k)$ has a small tubular neighborhood of positive radius not intersecting the level a_k.

By Lemma 3.12, the manifold $\Sigma^{\mathcal{F}}(f_k)$ is transverse to $\mathcal{F}$ except in a finite number of turning points $c_1, \dots, c_d$. Fix a closed geodesic arc T_i through each c_i, transverse to the foliation. Let $B^{2n}(0, r) \subset \mathbb{R}^{2n}$ be the closed ball of radius r. For $r > 0$ sufficiently small, the exponential map for the leafwise metric $g_k^{\mathcal{F}}$ yields disjoint foliated charts $\phi_i : U_i \to [0,1] \times B^{2n}(0, r)$ satisfying $\phi_i(T_i) = [0,1] \times \{0\}$. Having fixed r, by taking the T_i sufficiently short – effectively shrinking U_i in the vertical direction – it can be assumed that:

$$\phi_i(\Sigma^{\mathcal{F}}(f_k) \cap U_i) \subset [0,1] \times B^{2n}(0, r/2)$$

Consider the family of open arcs $I_j \cong (0,1) \subset \Sigma^{\mathcal{F}}(f_k)$, $j \in [1, 2, .., l]$, and circles $I_j \cong \mathbb{S}^1 \subset \Sigma^{\mathcal{F}}(f_k)$, $j \in [l+1, 2, .., m]$, comprising $\Sigma^{\mathcal{F}}(f_k) \setminus (\cup_{i=1..d} U_i)$. For sufficiently small $0 < s < r$, the exponential map for the metric $g_k^{\mathcal{F}}$ defines disjoint charts $\psi_j : V_j \to I_j \times B^{2n}(0, s)$. The union of the U_i and the V_j covers $\Sigma^{\mathcal{F}}(f_k)$.

The subsets $\mathcal{C}_i$, as in Proposition 3.9, can be defined and they come in two families:

1. $s/2$–neighborhoods, in the metric $g_k^{\mathcal{F}}$, of the points $x \in I_j \cap \mathcal{L}$, for any j,
2. $r/2$–neighborhoods, in the metric $g_k^{\mathcal{F}}$, of the points $x \in T_i \cap \mathcal{L}$, for any i.

By construction, the $g_k^{\mathcal{F}}$–diameter of the $\mathcal{C}_i$ is bounded above by $r/2$. Further, the $g_k^{\mathcal{F}}$–distance between any two sets $\mathcal{C}_i$ and $\mathcal{C}_{i'}$ is bounded below by s. Therefore Conditions (a.) and (b.) in Proposition 3.9 hold. Condition (c.) follows immediately from the fact that the union of the $\mathcal{C}_i$ is the intersection of a neighborhood of $\Sigma^{\mathcal{F}}(f_k)$ with the leaf $\mathcal{L}$.

An application of Lemma 3.8 shows that the relative homotopy groups $\pi_j(\mathcal{L}, \mathcal{L} \cap W_k)$ vanish for $j < n$ and for k large enough, since we already did the index computation in Proposition 3.7. □

Chapter 4

Dynamics of leafwise vector fields

The last two chapters of Part I (namely, this chapter and the next) deal with dynamics and, as such, they could be thought to be thematically independent from the rest of the thesis. However, this is not the case: let us explain how they fit into the big picture and how they relate to each other.

This chapter is about dynamics, but from an *h–principle perspective*. Our main result states that vector fields with no closed orbits satisfy a complete h–principle. As the reader knows from Chapter 2, the analogous complete h–principle must then hold for vector fields *tangent to foliations* as well.

Let us review the history of the problem. The *Seifert Conjecture* [73] stated that all non–singular vector fields in $\mathbb{S}^3$ have at least one closed orbit. A more general but related question was whether the topology of a manifold imposes some dynamical rigidity forcing closed orbits to appear.

A first result indicating that this was not the case appeared in [87], where Wilson showed that, for manifolds of dimension at least 4, any non–singular vector field can be homotoped to one without closed orbits. He also proved that any non–singular vector field in a 3–manifold can be homotoped to a vector field with only finitely many closed orbits. In both cases, one could assume any degree of differentiability for the vector fields involved.

Inspired by the ideas of Wilson, the Seifert Conjecture was settled in the negative first by Schweitzer [71] in the C^1–category and then by Krystyna Kuperberg [48] in the smooth case. They proved that a non–singular vector field in a 3–dimensional manifold can be homotoped to a vector field with no closed orbits.

These results can be restated, using the language of the h–principle, as follows. Given a manifold M, denote by $\mathfrak{X}_{ns}(M)$ the space of smooth non–singular vector fields on M and by $\mathfrak{X}_{no}(M)$ the space of smooth non–singular vector fields with no closed orbits, both of them endowed with the C^∞–topology. Then Wilson's and Kuperberg's constructions show that the inclusion

$$\iota_n : \mathfrak{X}_{no}(M) \to \mathfrak{X}_{ns}(M)$$

induces a surjection in π_0 as long as $\dim(M) \geq 3$.

Both Wilson's and Kuperberg's constructions are based around the notion of a *plug*, which is a local model for modifying a vector field in a flowbox. *Wilson's plug* traps a non–empty open subset of orbits and, in dimension greater than 3, creates no new closed ones. *Kuperberg's plug* in dimension 3 creates no new closed orbits and a later result by Matsumoto [55] shows that the set of orbits that are trapped in Kuperberg's plug contains a non–empty open subset.

Let $(M^{n+m}, \mathcal{F}^n)$, $n \geq 3$, be a closed smooth $(n+m)$-dimensional manifold endowed with a smooth foliation of codimension m. Denote by $\mathfrak{X}_{ns}(M, \mathcal{F})$ and $\mathfrak{X}_{no}(M, \mathcal{F})$ the subsets of, respectively, $\mathfrak{X}_{ns}(M)$ and $\mathfrak{X}_{no}(M)$ consisting of vector fields tangent to $\mathcal{F}$. The main result of this chapter reads, as promised:

Theorem 4.1. *The inclusion:*

$$\iota_n : \mathfrak{X}_{no}(M, \mathcal{F}) \to \mathfrak{X}_{ns}(M, \mathcal{F})$$

is a weak homotopy equivalence.

In particular, in the case where $\mathcal{F}$ is comprised of a single leaf, the whole of M, this recovers and improves the results of Wilson and Kuperberg, that dealt only with π_0. Also, it shows that the *foliated Seifert conjecture* – every vector field tangent to a foliation has a closed orbit – does not hold for foliations of dimension $n \geq 3$. For contrast, the case where the leaves are 2–dimensional is discussed in Section 4.3: it will be shown that there is an ample class of foliations for which all foliated vector fields must have a closed orbit.

Similarly, as in the classical case, it is possible to find classes of vector fields tangent to foliations that always possess a periodic orbit. This holds true, for instance, for Reeb vector fields in contact foliations with overtwisted leaves. The next chapter will precisely study this problem, the *foliated Weinstein conjecture* (as introduced in Chapter 1).

4.1 Setup and proof of the theorem

For the rest of the section $\mathcal{M}^{n+l}$ will denote a smooth compact manifold, possibly with boundary and corners. Endow $\mathcal{M}$ with a smooth n-dimensional foliation $\mathcal{F}^n_{\mathcal{M}}$ and a smooth non–singular vector field X, tangent to $\mathcal{F}_{\mathcal{M}}$. Homotopies of vector fields will be of particular interest and, unless stated otherwise, they will always be through smooth non–singular vector fields tangent to $\mathcal{F}_{\mathcal{M}}$.

By a **foliated flowbox**, or simply a flowbox, it is meant an embedding

$$\phi : [-2, 2] \times \mathbb{D}^{n-1} \times \mathbb{D}^l \to \mathcal{M}$$

with image $U \subset \mathcal{M}$, a smooth submanifold with corners. In the domain of ϕ there are coordinates $(z; x_2, \ldots, x_n; y_1, \ldots, y_l)$. We require for ϕ to satisfy $\phi^*\mathcal{F}_{\mathcal{M}} = \ker(dy_1) \cap \cdots \cap \ker(dy_l)$ and $\phi^*X = \partial_z$. U^+, U^- and U^v denote the components of ∂U in which X is outgoing, ingoing, and tangent, respectively. If $V \subset U$ is another foliated flowbox such that $V^+ \subset U^+$, $V^- \subset U^-$ and $V^v \subset \mathring{U}$ then the pair (U, V) will be called **nice**. The following proposition is key in the construction, and the proof is standard, as in [87, Theorem A].

Proposition 4.2. *Let $\mathcal{M}$, $\mathcal{F}_{\mathcal{M}}$ and X be as above. Fix $A \subset \mathcal{M}$ an open neighbourhood of $\partial\mathcal{M}$. Then there is a finite number of pairs (U_i, V_i) satisfying:*

- *each (U_i, V_i) is a nice pair of foliated flowboxes,*
- *any orbit of X is either fully contained in A or it intersects one of the V_i,*

- *the U_i are disjoint from $\partial\mathcal{M}$ and disjoint from one another.*

The idea now is homotoping X within the flowboxes in order to "open up" all closed orbits without introducing new ones. Let N^{n-1} be a manifold with boundary, possibly with corners, and denote $\mathbf{N} = [-2,2] \times N \times \mathbb{D}^l$, with coordinates $(z; p; y_1, \dots, y_l)$. Assume that there is an embedding $\psi : \mathbf{N} \to [-2,2] \times \mathbb{D}^{n-1} \times \mathbb{D}^l$ such that:

- ψ is the identity in the y coordinates,
- ψ preserves the vertical direction, i.e. $\psi^*\partial_z = \delta\partial_z$ with δ a positive function.

If $\mathbf{N}$ is endowed with a vector field X_N that agrees with ∂_z close to $\partial\mathbf{N}$ and that has no ∂_{y_i} components for all i, we say that the pair $(\mathbf{N}, X_N)$ is a **parametric plug**.

Denote by $\mathbf{N}^+$, $\mathbf{N}^-$, and $\mathbf{N}^v$ the different components of $\partial\mathbf{N}$, as in the case of flowboxes. A trajectory of X_N intersecting $\mathbf{N}^-$ is said to be **entering** the plug and a trajectory intersecting $\mathbf{N}^+$ is said to be **exiting** the plug. Since these plugs are meant to be embedded in foliated flowboxes in order to replace X by X_N, there are a number of properties that a parametric plug must satisfy:

i. X_N must be homotopic to ∂_z, relative to the boundary, and through non–vanishing vector fields with no ∂_{y_i} components, $i = 1, \dots, l$,

ii. if a trajectory of X_N enters and exits the plug, then it must do so at opposite points $(-2, x_0, y_0)$ and $(2, x_0, y_0)$.

A trajectory entering the plug and remaining there for infinite time is called **trapped**. The first property ensures that if the plug is used within a foliated flowbox, then the homotopy obtained is indeed through non–vanishing vector fields tangent to $\mathcal{F}_\mathcal{M}$. The second property ensures that no new closed orbits are created by connecting two previously different orbits.

Proposition 4.3. *Following with the notation of Proposition 4.2, suppose that there is a parametric plug* $\mathbf{N}$ *that additionally satisfies that:*

iii. X_N has no closed orbits within $\mathbf{N}$*,*

iv. the set of trajectories of X_N trapped by $\mathbf{N}$ *contains a non–empty open set.*

Then there is a homotopy of X, relative to $\partial\mathcal{M}$ and through non–singular vector fields tangent to $\mathcal{F}_\mathcal{M}$, to a vector field X' whose closed orbits are contained in A.

Proof. Consider the open set of trapped trajectories given by Property (iv.). Denote by $T_N^- \subset \mathbf{N}^-$ its intersection with the lower boundary of $\mathbf{N}$ and by $T_N^+ \subset \mathbf{N}^+$ its intersection with the upper boundary.

Consider a finite cover by nice pairs (U_i, V_i) as in Proposition 4.2. Since T_N^- and T_N^+ have non–empty interior, there are embeddings $\psi_i : \mathbf{N} \to U_i$ satisfying $\psi_i^* X = \delta\partial_z$, with δ a positive function, $V_i^- \subset \psi_i(T_N^-)$, and $V_i^+ \subset \psi_i(T_N^+)$.

Within each $\psi_i(\mathbf{N})$, we homotope $X = (\psi_i)_*(\delta\partial_z)$ to $(\psi_i)_*\partial_z$ and then to $(\psi_i)_* X_N$ as in Property (i.). This yields a new vector field X'. Let γ' be some trajectory of X'. If it is not fully contained in A, it shares some segment with a trajectory γ of X that was not fully contained in A. Follow this segment forward or backward in time as part of γ and γ'. If they differ at some point it is because γ entered a plug. If it escaped the plug, it did so at opposing points of the plug, by Property (ii.), and hence it will agree with γ' on the other side. If it does not escape the plug, the orbit becomes trapped and hence cannot be closed. Since every orbit of X intersected some V_i, the latter must happen eventually, and hence γ' has at least one end trapped. Since no new closed trajectories have been introduced in the plugs, by Property (iii.), the claim follows. □

As soon as the existence of such a plug $\mathbf{N}$ is proven for $n \geq 3$, Theorem 4.1 is an easy corollary.

Proof of Theorem 4.1 (assuming the existence of a suitable plug). Using the homotopy exact sequence for inclusions, the theorem is equivalent to showing that

$$\pi_j(\mathfrak{X}_{ns}(M, \mathcal{F}), \mathfrak{X}_{no}(M, \mathcal{F})) = 0 \quad \text{for all } j \in \mathbb{Z}.$$

Let X_t, $t \in \mathbb{D}^j$, be a j–parametric family of non–vanishing vector fields tangent to $\mathcal{F}$, defining an element in $\pi_j(\mathfrak{X}_{ns}(M, \mathcal{F}), \mathfrak{X}_{no}(M, \mathcal{F}))$. What has to be proven now is that this family can be homotoped, leaving those X_t, $t \in \mathbb{S}^{j-1}$, fixed, to a family fully contained in $\mathfrak{X}_{no}(M, \mathcal{F})$.

Consider the manifold $\mathcal{M} = M \times \mathbb{D}^j$ with the foliation $\mathcal{F}_{\mathcal{M}} = \coprod_{t_0 \in \mathbb{D}^j} \mathcal{F} \times \{t_0\}$ of codimension $m + j$. Then X_t can be regarded as a vector field X in $\mathcal{M}$ tangent to $\mathcal{F}_{\mathcal{M}}$. Since X_t is an element in the relative homotopy group $\pi_j(\mathfrak{X}_{ns}(M, \mathcal{F}), \mathfrak{X}_{no}(M, \mathcal{F}))$, we can assume that X has no closed orbits in a neighborhood A of $\partial \mathcal{M} = M \times \mathbb{S}^{j-1}$. Then an application of Proposition 4.3 readily implies that X can be homotoped, relative to $\partial \mathcal{M}$ and through non–vanishing vector fields tangent to $\mathcal{F}_{\mathcal{M}}$, to a vector field X' with no closed orbits.

Equivalently, the family X_t of vector fields can be homotoped, relative to the boundary of $\mathbb{D}^j$, to a family X'_t fully contained in $\mathfrak{X}_{no}(M, \mathcal{F})$, thus proving the claim. □

4.2 Construction of the parametric plugs

In this section we describe the parametric versions of Wilson's plug (which is needed for Theorem 4.1 if $n \geq 4$) and Kuperberg's plug (for the case $n = 3$). Note that Kuperberg's plug could be used also for the higher dimensional case, but Wilson's is easier to describe and paves the way to explain Kuperberg's.

4.2.1 The Wilson Plug in dimensions 4 and higher

Consider the manifold with boundary and corners $\mathbf{W}^{n,l} = [-2, 2] \times \mathbb{T}^2 \times [-2, 2] \times \mathbb{D}^{n-4} \times \mathbb{D}^l$, with coordinates $(z; s, t; r; x_5, \ldots, x_n; y_1, \ldots, y_l)$, $s, t \in [0, \pi)$, embedded in $\mathbb{R}^{n+l}$, $n \geq 4$ as follows:

$$i : \mathbf{W}^{n,l} \to \mathbb{R}^{n+l}$$

$$i(z, s, t, r, x, y) = (z, \cos(s)(6 + (3 + r)\cos(t)), \sin(s)(6 + (3 + r)\cos(t)), (3 + r)\sin(t), x, y).$$

Construct a vector field X_W in $\mathbf{W}^{n,l}$ as follows:

$$X_W = f(z, r, x, y)(\partial_s + b\partial_t) + g(z, r, x, y)\partial_z,$$

with b some irrational number and f, g smooth functions satisfying the following constraints:

1. g is symmetric and f is antisymmetric in the z coordinate,
2. $g(z,r,x,y) = 1$, $f(z,r,x,y) = 0$ close to the boundary of $\mathbf{W}^{n,l}$,
3. $g(z,r,x,y) \geq 0$ everywhere and $g(z,r,x,y) = 0$ only in $\{|z| = 1, |r| \leq 1, |x| \leq 1/2, |y| \leq 1/2\}$,
4. $f(z,r,x,y) = 1$ in $\{z \in [-3/2,-1/2], |r| \leq 1, |x| \leq 1/2, |y| \leq 1/2\}$.

This is the usual construction for **Wilson's plug**, but we have explicitly split the additional coordinates into $(x_i)_{i=5,\ldots,n}$ and $(y_j)_{j=1,\ldots,l}$, so that the y coordinates denote the parameter space. Write $\mathbf{W}^{n,l}_{y_0}$ for the n–dimensional plug one obtains for $y = y_0$ fixed.

Proposition 4.4 (Wilson [87]). *Wilson's plug satisfies all 4 properties required for Proposition 4.3 to hold.*

Proof. Property (i.) follows by interpolating linearly between g and the constant function 1 and then between f and the constant function 0. The symmetry of g and the antisymmetry of f imply Property (ii.). The only possible closed orbits within $\mathbf{W}^{n,l}$ would lie in the zero set of g, and by construction the flow in the zero set consists of invariant tori in which the vector field has irrational slope, so Property (iii.) follows. Finally, the orbits touching $\{z = \pm 2, |r| \leq 1, |x| \leq 1/2, |y| \leq 1/2\}$ are trapped, proving Property (iv.). $\square$

4.2.2 The Wilson plug in dimension 3

It is clear from the construction above that Wilson's method cannot be used in dimension 3. However, a 3–dimensional version can be constructed. This object will be used later on when defining Kuperberg's plug. The treatment here follows very closely the one in [40], where everything is described in more detail.

Consider the manifold $\mathbf{W} = [-2,2] \times \mathbb{S}^1 \times [1,3]$, with coordinates (z,θ,r), embedded in $\mathbb{R}^3$ cylindrically in the obvious fashion. Define a vector field X_S in $\mathbf{W}$ as follows:

$$X_W = f(z,r)\partial_\theta + g(z,r)\partial_z,$$

with the functions f and g satisfying:

- f is antisymmetric and g is symmetric in the z coordinate,
- f is 0 and g is 1 near the boundary of $\mathbf{W}$,
- $g(z,r) \geq 0$ and $g(z,r) = 0$ only in $B = \{|z| = 1, r = 2\}$,
- $f(z,r) \geq 0$ in $\{|z| > 0\}$ and $f(z,r) = 1$ in $\{1/4 \leq z \leq 7/4, 5/4 \leq r \leq 11/4\}$.

This version of Wilson's plug satisfies Properties (i.) and (ii.), as is easily verified. Further, it contains a pair of closed orbits, namely, $\{|z| = 1, r = 2\}$ and a closed set of orbits that get trapped, those touching $\{z = \pm 2, r = 2\}$. Observe that the flow of X_W is tangent to the cylinders with $r = r_0$ fixed.

4.2.3 The Kuperberg Plug

The plug $\mathbf{W}$ described above is the basis for Kuperberg's plug. See [48] for the original article and [40] for a very detailed account of the construction.

The key objects are as follows. We have shown that the two closed orbits γ_i, $i = 1, 2$, of $\mathbf{W}$, lie in the cylinder $[-2, 2] \times \mathbb{S}^1 \times \{2\}$. We are going to construct two cylinders D_i, each one of them intersecting the orbit γ_i in a segment. The D_i will be reinserted into the plug in order to destroy the γ_i while introducing no new periodic orbits.

Construct two disjoint convex discs $L_i \subset \mathbb{S}^1 \times [1, 3]$, $i = 1, 2$. L_i has a piecewise smooth boundary comprised of two closed connected intervals: α_i', whose ends are attached to $\mathbb{S}^1 \times \{3\}$ and whose interior lies in $\mathbb{S}^1 \times (1, 3)$, and $\alpha_i \subset \mathbb{S}^1 \times \{3\}$. We additionally assume that each L_i intersects the curve $\mathbb{S}^1 \times \{2\}$ in a segment. We define $D_i = [-2, 2] \times L_i$.

Let $\mathcal{D}_i$, $i = 1, 2$, be two flowboxes for X_W, disjoint from one another and from the D_i, satisfying:

- $\mathcal{D}_i$ contains an interval $\{((-1)^i, \theta, 2),\ \theta_i^- \leq \theta \leq \theta_i^+\}$ of the closed orbit γ_i,
- each $\mathcal{D}_i$ is diffeomorphic, as a manifold with boundary and corners, to D_i by a map $\sigma_i : D_i \to \mathcal{D}_i$ satisfying $\sigma_i^* X_W = \partial_z$. Denote $\mathcal{L}_i^{\pm} = \sigma_i(\{\pm 2\} \times L_i)$, $i = 1, 2$.
- there is a closed connected arc $\beta_i' \subset \mathbb{S}^1 \times \{1\}$ such that $[-2, 2] \times \beta_i'$ is the region of the boundary of $\mathcal{D}_i$ lying in the vertical boundary of $\mathbf{W}$. We require for σ_i to map $\{z\} \times \alpha_i'$ to $\{z\} \times \beta_i'$, for all $z \in [-2, 2]$.

These properties imply that the identification σ_i can be realised by an immersion with self–intersections of $\mathbf{W}$ into $\mathbb{R}^3$. Further, the flow X_W in $\mathcal{D}_i$ can be replaced by $(\sigma_i)_* X_W$. See Figure 4.1 for a picture of all these elements.

For some θ_i, $i = 1, 2$, we require for the vertical interval $[-2, 2] \times \{\theta_i\} \times \{2\} \subset D_i$ to be the preimage of $\gamma_i \cap \mathcal{D}_i$ under σ_i. Then we further require for the following property to hold:

- **Radius inequality:** "for all $(z, \theta, r) \in D_i$, with image $\sigma_i(z, \theta, r) = (z', \theta', r') \in \mathcal{D}_i$, it holds that $r' < r$ except for the points $(z, \theta_i, 2)$, $z \in [-2, 2]$, where it is actually an equality."

The quotient manifold constructed by identifying in $\mathbf{W}$ the solid cylinders D_i and $\mathcal{D}_i$ using σ_i will be denoted $\mathbf{K}$. We call it **Kuperberg's plug**; see the left hand side of Figure 4.1. The quotient vector field obtained out of X_W by replacing it with $(\sigma_i)_* X_W$ in $\mathcal{D}_i$ will be denoted X_K.

The following theorem of Matsumoto shows that Property (iv.) of plugs is satisfied by Kuperberg's plug.

Theorem 4.5 ([55]). *There is $\delta > 0$ such that every orbit entering the Kuperberg plug at $\{\pm 2\} \times \mathbb{S}^1 \times (2 - \delta, 2)$ is trapped inside.*

The following Lemma will be useful in the next subsection. The right hand side of Figure 4.1 depicts the different intervals in the construction.

Lemma 4.6. *There is a homotopy in $\mathbf{W}$ of non–singular vector fields X_W^t, $t \in [0, 2]$, with $X_W^0 = X_W$ and $X_W^2 = \partial_z$, such that:*

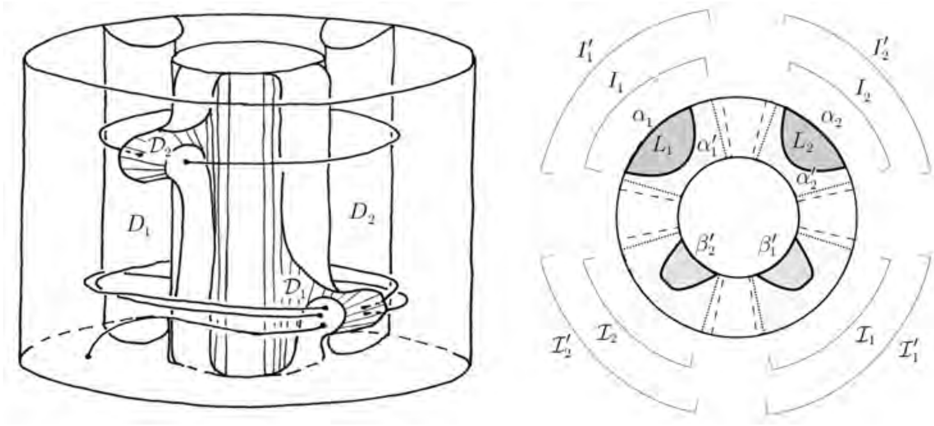

Figure 4.1: On the left hand side, the Kuperberg manifold seen as a quotient of the Wilson cylinder. D_i is identified with $\mathcal{D}_i$, $i = 1, 2$. On the right hand side, a horizontal slice of the Wilson cylinder. These figures originally appeared in [30] and [40].

- *X_W^t agrees with ∂_z in D_i for $t \in [1, 2]$,*
- *X_W^t agrees with X_W in $\mathcal{D}_i$ for $t \in [0, 1]$,*
- *X_W^t defines a plug with no closed nor trapped orbits for $t > 0$.*

Proof. Let f and g be the defining functions for $X_W = f(z, r)\partial_\theta + g(z, r)\partial_z$. Fix disjoint open subintervals of the circle $I_i, \mathcal{I}_i \subset \mathbb{S}^1$, $i = 1, 2$, such that $D_i \subset [-2, 2] \times I_i \times [1, 3]$ and $\mathcal{D}_i \subset [-2, 2] \times \mathcal{I}_i \times [1, 3]$. Fix slightly larger intervals $I_i', \mathcal{I}_i'$, still disjoint, such that $I_i \subset I_i'$ and $\mathcal{I}_i \subset \mathcal{I}_i'$. Construct bump functions

$$\alpha, \beta : \mathbb{S}^1 \to [0, 1]$$

$$\alpha(p) = 1, p \in I_i; \quad \alpha(p) = 0, p \notin I_i'; \quad i = 1, 2,$$

$$\beta(p) = 1, p \in \mathcal{I}_i; \quad \beta(p) = 0, p \notin \mathcal{I}_i'; \quad i = 1, 2.$$

Let $\phi : [0, 2] \to [0, 1]$ be a smooth function that is increasing in $[0, 1]$ and satisfies $\phi(0) = 0$ and $\phi(t) = 1$ for $t \in [1, 2]$. Similarly, let $\psi : [0, 2] \to [0, 1]$ be a smooth function that is increasing in $[1, 2]$ and satisfies $\psi(t) = 0$ for $t \in [0, 1]$ and $\psi(2) = 1$. Now define:

$$f_t(z, \theta, r) = f(z, r)(1 - \phi(t)\alpha(\theta) - \psi(t)(1 - \alpha(\theta)))$$

$$g_t(z, \theta, r) = g(z, r) + (1 - g(z, r))(\phi(t)\alpha(\theta) + \psi(t)(1 - \alpha(\theta)))$$

$$X_W^t = f_t\partial_\theta + g_t\partial_z.$$

It is immediate that X_W^t is non–singular and that the first two claims hold. For the last one, observe that $g_t > C_t > 0$ in I_i for $t > 0$, with C_t some positive constant. $\square$

See Fig. 4.2 for a pictorial representation of this construction.

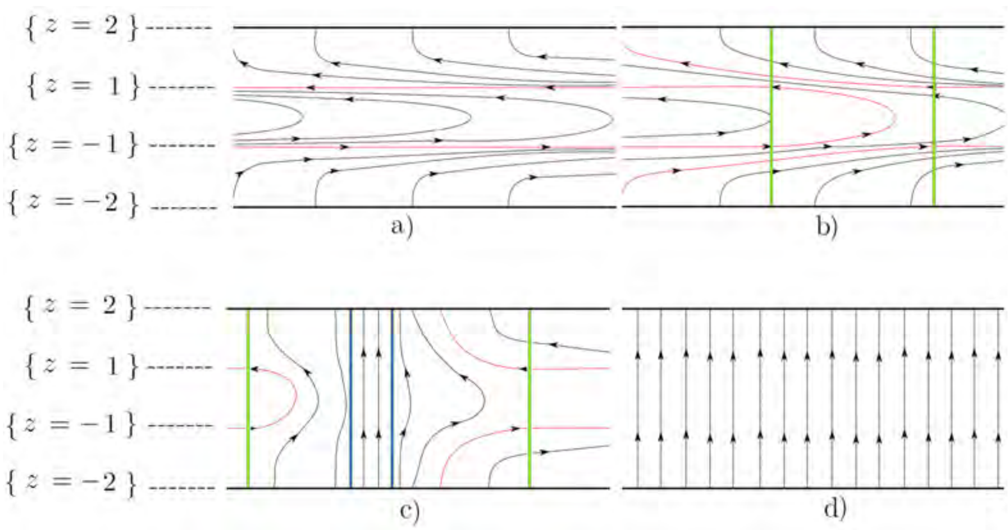

Figure 4.2: The flow of X_W^t at $\{r = 2\}$ in a neighbourhood of D_i. Image a) corresponds to $t = 0$, b) to $t = 1/2$, c) to $t = 1$, and d) to $t = 2$. The thickened dotted lines correspond to the orbit(s) that is(are) tangent to the curves $\{|z| = 1\}$ outside of I_i'.

4.2.3.1 The parametric Kuperberg plug

The *radius inequality* is the key to showing that the Kuperberg plug traps a non–empty open set of orbits and that it contains no closed orbits. Similarly, consider the following property:

- The **strict radius inequality** holds for a diffeomorphism $\phi_i : D_i \to \mathcal{D}_i$ if $r' < r$ for every $(z, \theta, r) \in D_i$ with $\phi_i(z, \theta, r) = (z', \theta', r')$.

In the process of interpolating to a trivial plug, we will need for the intermediate plugs to satisfy this strict radius inequality, since it will guarantee that all orbits enter and exit the plug.

A family of diffeomorphisms

$$\sigma_i^t : D_i \to \mathcal{D}_i, \quad t \in [0, 2], i = 1, 2; \quad \text{satisfying}$$

$$\sigma_i^0 = \sigma_i; \quad (\sigma_i^t)^* X_W = \partial_z$$

and satisfying the strict radius inequality for $t > 0$ can be constructed easily. The diffeomorphisms σ_i can be precomposed with diffeomorphisms of D_i that preserve the z component, that restrict to the identity in $[-2, 2] \times (\partial L_i)$ and that, away from there, take points to points with smaller radius. This produces diffeomorphisms σ_i^t that are C^∞–close to σ_i. The quotients of $\mathbf{W}$ induced by the gluings σ_i^t are all diffeomorphic to the Kuperberg manifold $\mathbf{K}$ and it is possible to fix a smooth t–parametric family of identifications with $\mathbf{K}$, which we henceforth assume.

Recall the explicit homotopy X_W^t constructed in Lemma 4.6. We define a family of vector fields in $\mathbf{W}$ as follows:

- $Y_W^t = X_W^t$ in $(\mathbf{W} \setminus (D_1 \cup D_2 \cup \mathcal{D}_1 \cup \mathcal{D}_2))$,
- $Y_W^t = X_W^t$ in D_i for $t \in [0, 1]$,

- $Y_W^t = (\sigma_i^t)_* X_W^t$ in $\mathcal{D}_i$ for $t \in [0,1]$,
- $Y_W^t = X_W^t$ in $\mathcal{D}_i$ for $t \in [1,2]$,
- $Y_W^t = (\sigma_i^t)^* X_W^t$ in D_i for $t \in [1,2]$.

Note that this vector field does not define a plug in $\mathbf{W}$, since it is not vertical close to the boundary in D_i for $t \in [1,2]$. However, it does descend to the quotient $\mathbf{K}$ and automatically induces a family of plugs $(\mathbf{K}, X_K^t)$, $t \in [0,2]$ interpolating from $X_K = X_K^0$ to $\partial_z = X_K^2$.

Lemma 4.7. *$(\boldsymbol{K}, X_K^t)$ has no closed orbits. Further, for $t > 0$, all orbits enter and exit the plug at opposing points.*

Proof. Since X_K^0 is Kuperberg's plug, it has no closed orbits. Let us set up some notation for the case $t > 0$. There are smooth bijective projections

$$\tau : \mathbf{W} \setminus (\mathcal{D}_1 \cup \mathcal{D}_2) \to \mathbf{K},$$

$$\tau' : \mathbf{W} \setminus (D_1 \cup D_2) \to \mathbf{K}.$$

The discontinuous radius function $\rho : \mathbf{K} \to [1,3]$ at a point p is defined to be the radius of $\tau^{-1}(p)$. Similarly, $r(p)$ will be the radius of the preimage by τ'. [1] Compactness of D_i and the *strict radius inequality*, imply that there is a lower bound

$$\rho - r \geq \epsilon > 0 \tag{4.1}$$

in the points where they disagree (that is, over the points in $\tau(D_i) = \tau'(\mathcal{D}_i)$).

Let us first explain the proof and then provide analytical details. We argue by contradiction by assuming that an orbit is trapped. The way in which we think about $\mathbf{K}$ is as a quotient of $\mathbf{W}$. Therefore, every time an orbit enters the self–insertion of the plug we imagine that the subsequent part of the orbit is entering $\mathbf{W}$. We keep track of the fact that we are inside the self–insertion with a *level function* (that we will define soon); this level function keeps increasing as we keep entering nested self–insertions. [2] Equation (4.1) states that every time we go up a level, the radius increases (by an amount bounded from below); this means that the level function cannot become arbitrarily large (because the radius cannot go up arbitrarily).

In $\mathbf{W}$, no orbits are trapped and, in particular, orbits exit opposite to where they entered from. Using induction, we deduce that if an orbit exits a self–insertion (regardless of how many nested self–insertions it went through inside), it also exits through the opposite point. Then, if an orbit is trapped, the discussion implies a level (in the level function) must be repeated infinitely many times. This implies that there are orbits in $\mathbf{W}$ that enter the self–insertion infinitely many times, which is a contradiction.

The level function. Fix a point $p \in \mathbf{K}$. Let $\Phi_s(p)$ be the flow of X_K^t at time s valued at p. Recall that $\mathcal{L}_i^{\pm} = \sigma_i(\{\pm 2\} \times L_i)$, for $i = 1, 2$, are the boundaries of the self–insertion that are transverse to the flow. We can denote

$$E_i(s_0) = \{\Phi_s(p) \cap \tau'(\mathcal{L}_i^-); s \in (0, s_0)\}$$

$$S_i(s_0) = \{\Phi_s(p) \cap \tau'(\mathcal{L}_i^+); s \in (0, s_0]\}$$

[1] Maybe this naming convention is a bit unfortunate. r here corresponds to r' in the statement of the strict radius inequality, and ρ here corresponds to r there.

[2] I do not know if this helps, but this is very much like the dream within a dream (within a dream...) of Inception. In order not to get lost we keep track of how many levels we have gone down (or, in this case, up).

the sets of points where the forward orbit of p enters and exits, respectively, the self–insertion of the plug. We shall call the points in $E_i = \cup_{s\geq 0} E_i(s)$ entry points and those in $S_i = \cup_{s\geq 0} S_i(s)$ exit points. Define the level function associated to p as follows:

$$\nu_p(s) = (\#E_1(s) + \#E_2(s)) - (\#S_1(s) + \#S_2(s)), \quad s \geq 0.$$

Consider the collection of points $E_1 \cup E_2 \cup S_1 \cup S_2$ and regard it as an ordered list $L = \{x_j = \Phi_{s_j}(p)\}$ in terms of increasing s_j, so the points appear in L as the forward orbit intersects the sets $\mathcal{L}_i^{\pm}$.

Let x_j and x_{j+1} be two consecutive points in L. If they both are entry points, then $\nu_p(s_{j+1}) = \nu_p(s_j)+1$. If they both are exit points, then $\nu_p(s_{j+1}) = \nu_p(s_j)-1$. Otherwise $\nu_p(s_j) = \nu_p(s_{j+1})$.

Obtaining Wilson orbits. Consider two points x_j and x_k, $k > j$, with x_j an entry point. Take the list $\{x_i\}_{i\in\{j,\dots,k\}} \subset L$ of points lying in–between. Recall $x_i = \Phi_{s_i}(p)$. If $\nu_p(s_j) = \nu_p(s_k) \leq \nu_p(s_i)$, $i \in \{j+1,\dots,k-1\}$, then we claim that $(\tau')^{-1}x_j, (\tau')^{-1}x_k \in \mathbf{W}$ lie in the same orbit of X_W^t, which we henceforth call a *Wilson orbit*. We proceed by induction on the size of $\{j,\dots,k\}$.

In the base case $k = j+1$, x_k must be an exit point. Having no other entry or exit points between $\tau^{-1}x_j$ and $\tau^{-1}x_k$, they are joined by a Wilson orbit and hence are opposite to each other in the bottom and top boundaries of $\mathbf{W}$. This implies that $(\tau')^{-1}x_j$ and $(\tau')^{-1}x_k$ are connected by a Wilson orbit.

For the induction step, find the first $l \in \{j+1,\dots,k\}$ satisfying $\nu_p(s_l) = \nu_p(s_j)$. This implies that x_l must be an exit point and all points in-between satisfy $\nu_p(s_i) > \nu_p(s_j) = \nu_p(s_l)$, $i \in \{j+1,\dots,l-1\}$. If $l < k$, the induction hypothesis applies and $(\tau')^{-1}x_j$ and $(\tau')^{-1}x_l$ are connected by a Wilson orbit. Additionally, x_{l+1} must be an entry point and $\nu(s_{l+1}) = \nu(s_l)$, so there is a Wilson orbit connecting $(\tau')^{-1}x_l$ and $(\tau')^{-1}x_{l+1}$. Iterating this process and concatenating the paths proves the induction step in this case.

Assume otherwise that $l = k$. Then we have that $\nu_p(s_i) > \nu_p(s_j) = \nu_p(s_k)$, $i \in \{j+1,\dots,k-1\}$, which in particular means that x_k is an exit point. It is also clear that x_{j+1} must be an entry point and x_{k-1} an exit point. This means that the induction hypothesis applies to the shorter list of points in-between x_{j+1} and x_{k-1}.

We have then that $(\tau')^{-1}x_{j+1}$ and $(\tau')^{-1}x_{k-1}$ are joined by a Wilson orbit. Recall that $\tau^{-1}x_j$ and $(\tau')^{-1}x_{j+1}$ are joined by a Wilson orbit. The same is true for $(\tau')^{-1}x_{k-1}$ and $\tau^{-1}x_k$. Concatenating all these segments yields a Wilson orbit between $\tau^{-1}x_j$ and $\tau^{-1}x_k$, which implies that they lie in opposing points in the lower and upper boundaries of $\mathbf{W}$. In particular, $(\tau')^{-1}x_j$ and $(\tau')^{-1}x_k$ are connected by a Wilson segment and the claim follows.

Concluding the argument. Let x_j be an entry point and let x_k, $k > j$. Assume $\nu_p(s_i) > \nu_p(s_j) = \nu_p(s_k)$, $i \in \{j+1,\dots,k-1\}$. Then $\rho(x_j) = \rho(x_k)$ and $r(x_j) = r(x_k)$ and we can say that x_k is the exit point corresponding to the entry point x_j. Equation 4.1 implies that the radius increases by ε at every entry point and this observation shows that at an exit point the radius goes back to the value it had at the corresponding entry point.

We conclude that, since ρ cannot be arbitrarily large, the elements in the list $N = \{\nu_p(s_j)\}$ have an upper bound. If L is infinite, then there is a minimum number k that gets repeated infinitely many times in N. In particular, we can choose x_j and x_{j+l} with $\nu_p(s_j) = \nu_p(s_{j+l}) = k$, all points in–between with $\nu_p \geq k$, and l arbitrarily large.

This means that we can find Wilson segments that intersect $\mathcal{L}_i^{\pm}$ arbitrarily many times, which is a contradiction with the fact that X_W^t has no trapped orbits for $t > 0$. Therefore, L must be finite and every orbit eventually escapes the plug. A similar analysis for negative time shows that it must enter the plug too. The level analysis above shows that it must do so at opposing points. □

Construct a smooth non–decreasing function $\eta : [0,1] \to [0,2]$, satisfying:

- η is identically 0 in $[0,1/2]$,
- $\eta > 0$ in $(1/2,1]$,
- η is identically 2 close to 1.

Define a family of functions $\eta_s = (1-s)\eta + 2s$, $s \in [0,1]$. Let $\mathbb{D}^l$ be the disk with coordinates $(y_1,\dots,y_l)$. A 1–parametric family of foliated vector fields $\mathcal{X}_K^s$ in $\mathbf{K} \times \mathbb{D}^l$ can be defined by

$$(\mathcal{X}_K^s)|_{\{y=y_0\}} = X_K^{\eta_s(|y_0|)}.$$

Proposition 4.8. *$(\boldsymbol{K}, \mathcal{X}_K^0)$ satisfies all 4 properties required for Proposition 4.3 to hold.*

Proof. $\mathcal{X}_K^s$ is the necessary homotopy between $\mathcal{X}_K^0$ and $\mathcal{X}_K^1 = \partial_z$. That this homotopy is through non–vanishing foliated vector fields follows from the fact that the X_K^t were non–vanishing. Property (i.) holds.

Theorem 4.5 states that Kuperberg's plug traps a non–empty open set of orbits $T_{\mathbf{K}}$. Since $(\mathcal{X}_K^0)_{\{y=y_0\}}$ agrees with the vector field in Kuperberg's plug for $y_0 \in \mathbb{D}^l_{1/2}$, it is immediate that $(\mathbf{K}, \mathcal{X}_K^0)$ traps the open set $T_{\mathbf{K}} \times \mathbb{D}^l_{1/2}$. Property (iv.) follows.

For $|y_0| > 1/2$ it holds that $\eta(|y_0|) > 0$. Hence, applying Lemma 4.7 to the flow $(\mathcal{X}_K^0)|_{\{y=y_0\}} = X_K^{\eta(|y_0|)}$ shows that $\mathcal{X}_K^0$ has no closed orbits in $|y_0| > 1/2$ and all orbits there go through the plug entering and exiting at opposing points. For $|y_0| \le 1/2$, $(\mathbf{K}, (\mathcal{X}_K^0)|_{\{y=y_0\}})$ is the Kuperberg plug. This proves Properties (ii.) and (iii.). □

4.3 Foliations with leaves of dimension 2

In this section M^3 will denote a connected orientable compact smooth 3-manifold, possibly with boundary. It will be endowed with a 2-dimensional foliation $\mathcal{F}^2$, which is assumed to be orientable and tangent to the boundary of M. Further, let X be a non–singular vector field tangent to $\mathcal{F}$.

Lemma 4.9. *Let $(\mathbb{T}^2, \mathcal{F}_T)$ be a smooth foliation by lines in the torus. If $\mathcal{F}_T$ has no Reeb components, then it is equivalent, up to conjugation by a homeomorphism of $\mathbb{T}^2$, to the foliation induced by the suspension of a diffeomorphism of the circle. If $\mathcal{F}_T$ has no closed orbits then the diffeomorphism of the circle is an irrational rotation.*

This is a well known fact. A proof can be found in [36]. The following proposition establishes the existence of at least two periodic orbits for any vector field tangent to the standard Reeb component. The corollary after the proposition is an immediate consequence of Novikov's compact leaf theorem.

Proposition 4.10. *Let $(M^3, \mathcal{F}^2)$ be a standard Reeb component. Let X be a non–singular vector field tangent to $\mathcal{F}$. Then X induces a Reeb component on its boundary torus. In particular, X has at least 2 closed orbits.*

Proof. Denote by $\mathcal{X}$ the oriented foliation by lines induced by X on the boundary torus T of the Reeb component. Assume that $\mathcal{X}$ has a Reeb component in T. Since this foliation is orientable, the Reeb component cannot have as boundary a single leaf $\mathbb{S}^1$, so the vector field X must have at least 2 closed orbits. Let us now assume that $\mathcal{X}$ does not have a Reeb component.

Parametrise $M = \mathbb{D}^2 \times \mathbb{S}^1$ explicitly with coordinates (r, θ, t), $|r| \leq 1$. Consider the one sided neighbourhood $\phi : (0,1] \times \mathbb{T}^2 \to M$, $\phi(r, \theta, t) \to (\frac{1+r}{2}, \theta, t)$ of the boundary torus T. Any curve representing the homology class $m \in H_1(T; \mathbb{Z})$ that vanishes by inclusion into M is called a meridian.

Using Lemma 4.9 in the $\mathbb{T}^2$ coordinates yields a new (maybe topological) embedding $\psi : (0,1] \times \mathbb{T}^2 \to M$ such that $\psi^*\mathcal{X}$ is a suspension of a diffeomorphism of the circle in the torus $\{1\} \times \mathbb{T}^2$.

Suppose that $\psi^*\mathcal{X}$ corresponds to the irrational rotation, then any curve with rational slope makes a constant angle with $\psi^*\mathcal{X}$. Note that, in particular, the homology class $(\psi|_{\{1\}\times\mathbb{T}^2})^* m$ of the meridian under this new parametrisation can be represented by some smooth curve γ with rational slope. Accordingly, $\psi^*\mathcal{X}$ and the tangent vector $\dot{\gamma}$ define, at each point in the image of γ, a positively oriented basis.

Suppose instead that $\psi^*\mathcal{X}$ corresponds to a suspension of a diffeomorphism of $\mathbb{S}^1$ with fixed points. The meridian class $(\psi|_{\{1\}\times\mathbb{T}^2})^* m$ can be represented by a smooth curve $\gamma : \mathbb{S}^1 \to \{1\} \times \mathbb{T}^2$. Denoting this class by (a, b), where the first component stands for the suspension direction, the curve γ can be set to agree with a compact leaf of $\mathcal{X}$ for almost a turns and then to turn b times transversely. Accordingly, the foliation $\psi^*\mathcal{X}$ and the tangent vector $\dot{\gamma}$ are, at each point in the image of γ, either colinear or define a positively oriented basis.

Summarizing, if the foliation $\psi^*\mathcal{X}$ does not have a Reeb component, it admits a smooth curve $\gamma : \mathbb{S}^1 \to \{1\} \times \mathbb{T}^2$ representing the meridian class $(\psi|_{\{1\}\times\mathbb{T}^2})^* m$, such that $\psi^*\mathcal{X}$ and $\dot{\gamma}$ are either colinear or define a positively oriented basis at every point. The degree of $\psi^*\mathcal{X}$ restricted to the image of γ is therefore 0. Since the degree is invariant by homeomorphism, we conclude that $\mathcal{X}$ has degree 0 on the image of the curve $\psi \circ \gamma$.

Now every leaf inside the Reeb component has a family of circles that asymptotically approach the image of $\psi \circ \gamma$. The previous discussion implies that X restricted to any given $\mathbb{R}^2$ leaf in the Reeb component is a non–singular vector field that restricted to some circle has degree 1 (with respect to the standard basis of $\mathbb{R}^2$). Using the Poincaré-Hopf index theorem we get a contradiction, thus implying that X has a Reeb component in the boundary torus T, as we desired to prove. □

Corollary 4.11. *Any non–singular vector field tangent to a codimension one foliation of $\mathbb{S}^3$ has at least 2 closed orbits.*

Proposition 4.10 can be proved in more generality. Following [45] and [72] we introduce the following definition.

Definition 4.12. *A foliation $(M, \mathcal{F})$ is called a generalised Reeb component if M is connected, ∂M is a union of leaves of $\mathcal{F}$, no pair of points on ∂M can be joined by a curve transverse to the foliation, and all the leaves in $\mathcal{F}|_{\mathring{M}}$ are proper and without holonomy.*

In particular, this means that ∂M is a union of tori. The following lemma, which is a straightforward consequence of [44, Corollary 2] and [63, Theorem 1], states that the behaviour near the boundary components is just like the one found in a standard Reeb component:

Lemma 4.13. *Let $(M, \mathcal{F})$ be a generalised Reeb component and let $T \subset \partial M$ be one of the boundary components. Then the one sided holonomy along T is an infinite cyclic group. In particular, there is a basis (α, β) for $H_1(T)$ such that the holonomy along α is contracting and the holonomy along β is the identity.*

We shall see in Theorem 4.15 below that in most generalised Reeb components X must carry closed orbits. First we characterise the exceptions. Consider the annulus $\mathbb{S}^1 \times [0,1]$ and denote by $\mathcal{F}_R$ the 1-dimensional Reeb foliation on the annulus. We will abuse notation and still denote by $\mathcal{F}_R$ its lift as a codimension one foliation to $\mathbb{T}^2 \times [0,1]$.

Lemma 4.14. *Let $(M^3, \mathcal{F}^2)$ be a generalised Reeb component. Suppose one of the leaves F is a cylinder. Then $(M, \mathcal{F})$ is homeomorphic to $(\mathbb{T}^2 \times [0,1], \mathcal{F}_R)$.*

Proof. By [44] it follows that $(\overset{\circ}{M}, \mathcal{F}|_{\overset{\circ}{M}})$ is a fibration over $\mathbb{S}^1$ whose leaves are diffeomorphic to cylinders. Since M is orientable, the fibration $\pi : \overset{\circ}{M} \to \mathbb{S}^1$ is trivial.

Let $\phi_i : (0,1] \times \mathbb{T}^2 \to M$, $i = 1, 2$, be one–sided charts of the 2 boundary components $\phi_i(\{1\} \times \mathbb{T}^2)$, with coordinates (r, s, θ). By Lemma 4.13, it can be assumed that the holonomy is the identity in the s–direction and contracting in the θ–direction. Then these local models can be assumed to agree with that of the standard Reeb component.

Since the leaves are proper, there are numbers r_1, r_2, such that the tori $S_i = \{r = r_i\} \subset \text{Image}(\phi_i)$ intersected with each leaf bound a compact cylinder. Then the ϕ_i can be reparametrised in the θ–direction so that $\pi \circ \phi_i^{-1}(r_i, s, \theta) = \pm\theta$. The sign depends on whether the coorientation of $\mathcal{F}$ agrees with the direction in which the holonomy is contracting. Denote by $B \subset \overset{\circ}{M}$ the manifold bounded by the tori S_i. Since $\pi : B \to \mathbb{S}^1$ is a submersion that is a fibration over each S_i, the Ehresmann fibration theorem implies that B is a trivial $\mathbb{S}^1 \times (-1, 1)$ bundle over $\mathbb{S}^1$.

The boundary torus S_i is endowed with two trivialisations, one coming from B and the other from ϕ_i. They might disagree by a number of Dehn twists in the s–direction. Denote their composition by $\tau : \mathbb{T}^2 \to \mathbb{T}^2$. Since the foliation structure in the chart ϕ_i is invariant under the action of τ on the (s, θ) coordinates, $\psi_i = \phi_i \circ \tau^{-1}$ is a new chart structure that makes the two trivialisations of S_i agree. Therefore, the trivialisation from B glues with the charts ψ_i to yield $\mathbb{T}^2 \times [0,1]$ as a manifold. Further, if $\pi \circ \psi_1^{-1}(r_1, s, \theta) = \pi \circ \psi_2^{-1}(r_2, s, \theta)$, then $\mathcal{F}$ is isomorphic to $\mathcal{F}_R$. Otherwise, that is if the orientations of the boundary components are reversed, $(M^3, \mathcal{F}^2)$ has a transverse path connecting two points of the boundary and is not a generalised Reeb component. $\square$

Now the main result is immediate:

Theorem 4.15. *Let $(M^3, \mathcal{F}^2)$ be a generalised Reeb component. If M is not homeomorphic to $\mathbb{T}^2 \times [0,1]$, then any vector field X tangent to $\mathcal{F}$ has at least 2 closed orbits.*

Proof. Since M is not homeomorphic to $\mathbb{T}^2 \times [0,1]$, none of the non–compact leaves of $\mathcal{F}$ are cylinders. In particular, they must have non–zero Euler characteristic. Assume that X, when restricted to all boundary components of M, induces no Reeb component. Applying Lemma 4.13 and proceeding as in Proposition 4.10 shows that, given some non–compact leaf F, there is a finite collection of closed curves $\gamma_i \subset F$ satisfying:

- $F \setminus \{\gamma_i\}$ is comprised of a compact component G that is a deformation retract of F and a collection of non–compact half–cylinders,

- X is either tangent or defines a positively oriented basis at each point of γ_i (endowed with appropriate orientations).

These properties again yield a contradiction using the Poincaré–Hopf index theorem. □

Chapter 5

The foliated Weinstein conjecture

The **Weinstein conjecture** [85] states that the Reeb vector field associated to a contact form α in a closed $(2n+1)$–manifold M always carries a closed periodic orbit. Hofer proved in [38] that the Weinstein conjecture holds for any 3-dimensional contact manifold (M^3, α) overtwisted or satisfying $\pi_2(M) \neq 0$. Then, it was proven in every 3–manifold by Taubes [77] by localising the Seiberg-Witten equations along Reeb orbits.

This chapter is dedicated to proving an analogue of Hofer's result for contact foliations. Namely:

Theorem 5.1. *Let $(M^{3+m}, \mathcal{F}^3, \xi^2)$ be a contact foliation in a closed manifold M. Let α be a defining 1–form for an extension of ξ and let R be its Reeb vector field. Let $\mathcal{L}^3 \hookrightarrow M$ be a leaf.*

i. If $(\mathcal{L}, \xi|_{\mathcal{L}})$ is an overtwisted contact manifold, R possesses a closed orbit in the closure of $\mathcal{L}$.

ii. If $\pi_2(\mathcal{L}) \neq 0$, R possesses a closed orbit in the closure of $\mathcal{L}$.

Comparing this result with Theorem 4.1, we confirm that Reeb dynamics are distinct to the dynamics of more general vector fields.

The proof of Theorem 5.1, based on Hofer's methods, occupies the last section of the note. Before that, several examples showing the sharpness of Theorem 5.1 are discussed. Some of them are unexpected and show that care is needed to state a Weinstein–type conjecture in full generality in the foliated case:

- *Overtwistedness is a necessary condition:* In Subsection 5.2.2, several examples of foliations with tight leaves are presented. Proposition 5.15 constructs a contact foliation in the 4–torus $\mathbb{T}^4$ that has *all leaves tight* and that has no Reeb orbits. Naturally, in this example all leaves are open. This shows that the **foliated Weinstein conjecture** does not necessarily hold as soon as we drop the assumption on overtwistedness. Then Proposition 5.12 presents a more sophisticated example of a contact foliation in $\mathbb{S}^3 \times \mathbb{S}^1$.
- *Jumps to a nearby leaf are necessary:* In Subsection 5.2.3 we construct a foliation in $\mathbb{S}^2 \times \mathbb{S}^1 \times \mathbb{S}^1$ that has two compact leaves $\mathbb{S}^2 \times \mathbb{S}^1 \times \{0, \pi\}$ on which all others accumulate. We then endow it with a foliated contact structure that makes all leaves overtwisted but that has closed Reeb orbits only in the compact ones. Theorem 5.1 is therefore sharp in the sense that an overtwisted leaf might not possess a Reeb orbit *itself*.

- *Being a leaf is necessary:* In Subsection 5.2.1 we construct Reeb flows with no closed orbits in every *open contact manifold*.

Remark 5.2. Section 5.3 is dedicated to setting up the J–holomorphic machinery in the foliated setting. Although we follow a direct route towards proving Theorem 5.1, some of the results are foundational and have applicability outside of the proof. It is the author's hope that they can be used to prove rigidity statements for contact foliations: this is currently work in progress.

5.1 Reviewing contact structures (again)

We have gone over some of the basic facts regarding contact structures already as we needed them in the thesis. However, some of the material that we will need in this chapter has not appeared yet. Let us go through it briefly.

5.1.1 The basic examples

Since we shall reference them over and over, let us explicitely introduce (again) the standard contact structure and the standard overtwisted structure in $\mathbb{R}^3$.

Example 5.3. Consider $\mathbb{R}^{2n+1}$ with coordinates $(x_1, y_1, \cdots, x_n, y_n, z)$. The $2n$–distribution $\xi_{\text{std}} = \ker(dz - \sum_{i=1..n} x_i dy_i)$ is called the standard **tight** contact structure.

Example 5.4. Consider $\mathbb{R}^3$ with cylindrical coordinates (r, θ, z). The 2–distribution $\xi_{ot} = \ker(\cos(\theta)dz + r\sin(r)d\theta)$ is called the **standard overtwisted** contact structure. The disc $\Delta = \{z = 0, r \leq \pi\}$ is called the **overtwisted disc**.

As explained in the Preamble, Bennequin showed in [3] that the structures $(\mathbb{R}^3, \xi_{\text{std}})$ and $(\mathbb{R}^3, \xi_{ot})$, although homotopic as plane fields, are distinct as contact structures (meaning that they are not *diffeomorphic*). The classification question in $\mathbb{R}^3$ was completely solved later on by Eliashberg:

Proposition 5.5 ([18])**.** *Let ξ be a contact structure in $\mathbb{R}^3$ that is overtwisted in the complement of every compact subset. Then ξ is isotopic to ξ_{ot}.*

Contact structures with the property that they remain overtwisted after removing any compact subset are called **overtwisted at infinity**.

5.1.2 Contact embeddings into overtwisted contact structures

The following lemma will be useful in Subsection 5.2.1.

Lemma 5.6 ([6, Corollary 1.4])**.** *Let (M^{2n+1}, ξ_M) be a connected overtwisted contact manifold and let (N^{2n+1}, ξ_N) be an open contact manifold of the same dimension. Let $f : N \to M$ be a smooth embedding covered by a contact bundle homomorphism $\Phi : TN \to TM$ – that is, $\Phi|_{\xi_M(p)}$ maps into $\xi_N(f(p))$ and preserves the conformal symplectic structure*[1] *– and assume that df and Φ are homotopic as injective bundle homomorphisms $TN \to TM$.*

Then f is isotopic to a contact embedding $\tilde{f} : (N, \xi_N) \to (M, \xi_M)$.

[1]This is the formal data associated to the contact structure. The bundle $\xi = \ker(\alpha)$ has a symplectic structure $(\xi, d\alpha)$ that is uniquely defined up to conformal factor as we rescale α.

5.1.3 Convex surfaces

Let (W^3, ξ^2) be a contact manifold. Let $\Sigma^2 \subset W$ be an immersed surface. The intersection $\xi \cap T\Sigma$ yields a singular foliation by lines on Σ, which is called the **characteristic foliation**. In the generic case, it can be assumed that the singularities – the points where $\xi_p = T_p\Sigma$ – are isolated non–degenerate points, that can then be classified into **elliptic** and **hyperbolic**.

Example 5.7. By our characterisation of overtwistedness, any overtwisted manifold (W, ξ) contains a disc Σ with a single singular point, which is elliptic, and whose boundary is legendrian. All other leaves spiral around the legendrian boundary in one end and converge to the elliptic point in the other. Such a disc appears as a C^∞–small perturbation of the overtwisted disc Δ.

Example 5.8. Consider the unit sphere $\mathbb{S}^2$ in $(\mathbb{R}^3, \xi_{\text{std}})$. Its singular foliation has two critical points located in the poles, which are elliptic. All other leaves are diffeomorphic to $\mathbb{R}$ and they connect the poles.

Theorem 5.9 (Eliashberg, Giroux, Fuchs)**.** *Let $\Sigma = \mathbb{S}^2$ and let (W, ξ) be tight. Then, the characteristic foliation of Σ is conjugate to the one of the unit sphere in $\mathbb{R}^3$ tight.*

5.2 The theorem is sharp

5.2.1 (Non–complete) Reeb vector fields with no closed orbits

It is first reasonable to wonder about the Weinstein conjecture for open manifolds in general. In this direction, not much is known. In [65, 66, 74] it is shown that the Weinstein conjecture holds for non–compact energy surfaces in cotangent bundles as long as one imposes certain topology conditions on the hypersurface and certain growth conditions on the hamiltonian, which is assumed to be of mechanical type.

Proposition 5.10. *Let (N^{2n+1}, ξ) be an open contact manifold. Then there is a contact form α, $\ker(\alpha) = \xi$, whose (possibly non–complete) associated Reeb flow has no periodic orbits.*

Proof. Fix some small ball $U \subset N$. Modify ξ within U to introduce an overtwisted disc Δ in the sense of [6]. By applying the relative h–principle for overtwisted contact structures, there is ξ_{OT} in N that agrees with ξ outside of U and that has Δ as an overtwisted disc. This new contact structure is homotopic to the original one as almost contact structures.

Let $\{N_i\}_{i \in \mathbb{N}}$ be an exhaustion of N by compact sets, $N_i \subset N_{i+1}$. Fix a non–degenerate contact form α_{OT} for the overtwisted structure ξ_{OT}. Its closed Reeb orbits are isolated and countable; moreover, we may assume that no closed orbit is fully contained in Δ. We index them as follows: each compact set N_i is intersected by finitely many closed orbits and hence we write $\{\gamma^i_j\}_{j \in I_i}$ for the collection of closed Reeb orbits intersecting N_i but not N_{i-1}.

Construct a path $\beta : [0, \infty) \to N$, avoiding Δ, that is proper and such that $N \setminus \beta([0, \infty))$ is diffeomorphic to N by a map isotopic to the identity. Then, for each i, and each $j \in I_i$, we can construct paths $\beta^i_j : [0, 1] \to N_i$ such that the β^i_j are all pairwise disjoint, they intersect Image(β) only at $\beta^i_j(0) \in$ Image(β), they satisfy $\beta^i_j(1) \in \gamma^i_j \cap N_i$, and they avoid Δ.

Since the images of β and the β^i_j avoid Δ, we can fix a closed contractible neighbourhood V of Δ disjoint from them as well. Construct a path $\beta_{\text{OT}} : [0, 1] \to N$ with $\beta_{\text{OT}}(0) \in \partial V$, $\beta_{\text{OT}}(1) \in$ Image(β) and otherwise avoiding V and all other paths.

Consider the tree $T = \beta \cup \{\cup_{i\in\mathbb{N}, j\in I_i} \beta^i_j\} \cup \beta_{\mathrm{OT}}$. Denote by $\nu(T)$ a small closed neighbourhood that deformation retracts onto T. We can assume that N is diffeomophic to $N' = N \setminus (\nu(T) \cup V)$ by a diffeomorphism $f : N \to N'$ that is isotopic to the identity.

The embedding $f : (N, \xi) \to (N' \cup V, \xi_{\mathrm{OT}})$ has image N' and is covered by a contact bundle homomorphism. This follows because f is isotopic to the identity in N and ξ and ξ_{OT} are homotopic. Now an application of Lemma 5.6 implies that there is an isocontact embedding $\tilde{f} : (N, \xi) \to (N' \cup V, \xi_{\mathrm{OT}})$. The form α_{OT} has no periodic orbits in $N' \cup V$ by construction and hence the pullback form $\alpha = \tilde{f}^* \alpha_{\mathrm{OT}}$ does not either. □

Remark 5.11. A natural open question is whether it is true that every open contact manifold can be endowed with a contact form inducing a complete Reeb flow with no closed orbits.

5.2.2 The Weinstein conjecture does not hold for contact foliations with all leaves tight

We shall construct first a contact foliation with all leaves tight and with periodic orbits lying in the only compact leaf.

Proposition 5.12. *Let $(\mathbb{S}^3, \mathcal{F}_{Reeb})$ be the Reeb foliation on the 3–sphere and let g be the round metric in $\mathbb{S}^3$. Consider the contact foliation $(\mathbb{S}^3 \times \mathbb{S}^1, \lambda_{can})$ on the unit cotangent bundle of $\mathcal{F}_{Reeb}$. Its only closed Reeb orbits lie in the compact torus leaf.*

The proposition is an easy consequence of the following Lemma.

Lemma 5.13. *Consider the Riemannian manifold $(\mathbb{R}^2, g)$, where g is of the form $dr \otimes dr + f(r)d\theta \otimes d\theta$, with $f(r)$ an increasing function satisfying $f(r) = r^2$ close to the origin. $(\mathbb{R}^2, g)$ has no closed geodesics.*

Proof. Applying the Koszul formula yields the following equations for the Christoffel symbols:

$$g(\nabla_{\partial_r}\partial_\theta, \partial_\theta) = f'/2 = \Gamma^\theta_{r\theta} g(\partial_\theta, \partial_\theta) = \Gamma^\theta_{r\theta} f,$$

$$g(\nabla_{\partial_\theta}\partial_r, \partial_\theta) = f'/2 = \Gamma^\theta_{\theta r} g(\partial_\theta, \partial_\theta) = \Gamma^\theta_{\theta r} f,$$

$$g(\nabla_{\partial_\theta}\partial_\theta, \partial_r) = -f'/2 = \Gamma^r_{\theta\theta} g(\partial_r, \partial_r) = \Gamma^r_{\theta\theta}.$$

And hence the geodesic equations read:

$$\ddot{r} = f'\dot{\theta}^2,$$

$$\ddot{\theta} = -\log(f)'\dot{\theta}\dot{r}.$$

If at any point $\dot{\theta} = 0$, then $\dot{\theta} = 0$ for all times and $\dot{r}$ is a constant. This situation corresponds to radial lines.

All other geodesics have always $\dot{\theta} \neq 0$ and hence $\ddot{r} > 0$. In particular, as soon as a geodesic has $\dot{r} \geq 0$ at some point, it will have $\dot{r} > 0$ for all the points in the forward orbit and hence it will not close up.

For a geodesic to close up we deduce then that it must have $\dot{r} < 0$ for all times, but then it cannot close up either. □

Proof of Proposition 5.12. Consider $\mathbb{S}^3$ lying in $\mathbb{C}^2$, with coordinates $(z_1, z_2) = (r_1, \theta_1, r_2, \theta_2)$. The Reeb foliation can be assumed to have the Clifford torus $|z_1|^2 = |z_2|^2 = 1/2$ as its torus leaf. One of the solid tori, denote it by T, corresponds to $\{|z_1|^2 \leq 1/2, |z_2|^2 = 1 - |z_1|^2\}$ and the other one is given by the symmetric equation. Let us multiple cover the solid torus T with the map $\phi : \mathbb{R} \times \mathbb{D}^2_{1/\sqrt{2}} \to T$ given by $\phi(s, r, \theta) = (r, \theta, \sqrt{1-r^2}, s)$. For all purposes we can work in $\mathbb{R} \times \mathbb{D}^2_{1/\sqrt{2}}$, which is the universal cover of the solid torus, and hence we shall do so.

The restriction of the flat metric of $\mathbb{C}^2$

$$g = \sum_{i=1,2} dr_i \otimes dr_i + r_i^2 d\theta_i \otimes d\theta_i$$

to $\mathbb{S}^3$ is precisely the round metric. In the parametrisation of T given above it reads as:

$$\phi^* g = \frac{1}{1-r^2} dr \otimes dr + r^2 d\theta \otimes d\theta + (1 - r^2) ds \otimes ds.$$

Which in particular readily shows that the metric induced in the Clifford torus is flat.

Consider the embeddings

$$\psi_c : \mathbb{R}^2 \to \mathbb{R} \times \mathbb{D}^2_{1/\sqrt{2}}$$

$$\psi_c(\rho, \theta) = (f_1(\rho) + c, f_2(\rho), \theta),$$

with $f_1 : \mathbb{R}^+ \to \mathbb{R}^+$ a smooth increasing function that agrees with ρ^2 near the origin and with the identity away from it, and $f_2 : \mathbb{R}^+ \to \mathbb{R}^+$ also smooth and increasing, agreeing with the identity near the origin, and converging to $1/\sqrt{2}$ as $\rho \to \infty$. As c is allowed to vary, these embeddings realise the non–compact leaves of the Reeb foliation in T. The pullback metric on each one of them is of the form

$$\psi_c^* \phi^* g = \left[\frac{(f_2')^2}{1 - f_2^2} + (1 - f_2^2)(f_1')^2 \right] d\rho \otimes d\rho + f_2^2 d\theta \otimes d\theta = h_1(\rho) d\rho \otimes d\rho + h_2(\rho) d\theta \otimes d\theta,$$

that is, $h_2(\rho)$ is increasing and converges to $1/2$ as $\rho \to \infty$ and $h_1(\rho)$ is bounded from above and behaves as $O(1)$ near the origin.

Now we claim that reparametrising $\mathbb{R}^2$ suitably yields a metric like the one in Lemma 5.13; this would immediately allow us to conclude the proof. Consider the vector field $X = \sqrt{h_1(\rho)} \partial_\rho$: it is a radial vector field that is of unit length for the metric $\psi_c^* \phi^* g$; as such, it is only defined over $\mathbb{R}^2 \setminus \{0\}$. However, it still allows us to define a diffeomorphism Φ of $\mathbb{R}^2$ to itself: the point (ρ, θ) is taken to the time ρ flow of X, starting at the origin with angle θ. By construction, it must hold:

$$\Phi^* \psi_c^* \phi^* g = d\rho \otimes d\rho + h(\rho) d\theta \otimes d\theta,$$

with $h(\rho)$ increasing and converging to $1/2$ as $\rho \to \infty$ (since h_2 satisfied those properties and we have essentially just reparametrised the radius function). □

Remark 5.14. Taking the universal cover of a leaf yields the standard tight $\mathbb{R}^3$, so all leaves are tight.

One can actually construct a contact foliation with no periodic orbits of the Reeb flow.

Proposition 5.15. *Consider the manifold $\mathbb{T}^3$, endowed with the Euclidean metric g, and the foliation $\mathcal{F}$ by planes given by two rationally independent slopes. The space of foliated cooriented contact elements $\mathbb{S}(T^*\mathcal{F})$ has no closed Reeb orbits.*

Proof. Let $\mathcal{L}$ be any leaf of $\mathcal{F}$. $\mathcal{L}$ is diffeomorphic to $\mathbb{R}^2 \times \mathbb{S}^1$ and its universal cover of is the standard tight $\mathbb{R}^3$. Hence it is a tight contact manifold. Since the restriction of g to $\mathcal{L}$ is Euclidean, there are no closed geodesics on $\mathcal{L}$ and hence no closed Reeb orbits in its sphere cotangent bundle. $\square$

5.2.3 A sharp example. Overtwisted leaves with no closed orbits

Proposition 5.16. *There is a contact foliation on $\mathbb{S}^2 \times \mathbb{S}^1 \times \mathbb{S}^1$ having all leaves overtwisted and such that the only closed Reeb orbits appear in the two compact leaves.*

We shall dedicate the rest of this subsection to the proof of Proposition 5.16. We structure it in three parts.

5.2.3.1 $\mathbb{R}^3$ overtwisted at infinity with no closed orbits

Consider the following 1–form in $\mathbb{R}^3$ in cylindrical coordinates:

$$\alpha = \cos(r)dz + (r\sin(r) + f(z)\phi(r))d\theta$$

If $f(z)\phi(r) = 0$ identically, this is the standard form α_{OT} for the contact structure ξ_{OT} that is overtwisted at infinity. We well henceforth assume that $f(z)\phi(r)$ is C^1–small, and therefore α will be a contact form as well. In particular, by Proposition 5.5, the contact structure it defines is contactomorphic to ξ_{OT}. Let us compute:

$$d\alpha = -\sin(r)dr \wedge dz + [\sin(r) + r\cos(r) + \phi'(r)f(z)]dr \wedge d\theta + f'(z)\phi(r)dz \wedge d\theta$$

whose kernel, away from the origin, is spanned by:

$$X = -f'(z)\phi(r)\partial_r + [\sin(r) + r\cos(r) + \phi'(r)f(z)]\partial_z + \sin(r)\partial_\theta.$$

It is easy to check that $\alpha(X) > 0$ far from the origin, and hence the Reeb vector field is a positive multiple of X.

Assume that $\phi(r)$ is a monotone function that is identically 0 close to 0 and identically 1 in $[\delta, \infty)$, for $\delta > 0$ small. Assume further that f is strictly decreasing, and sufficiently small to guarantee $|\phi'(r)f(z)| << |\sin(r) + r\cos(r)|$ for $r \in [0, \delta]$, this is indeed possible because $\phi'(r)$ can be taken to behave as $O(r)$. Then, the Reeb v.f. is ∂_z in $r = 0$ and has a positive vertical component for $r \in [0, \delta]$. Away from this neighbourhood of the vertical axis, the Reeb flow has a positive radial component, so we conclude that it has no closed orbits.

5.2.3.2 $\mathbb{S}^2 \times \mathbb{R}$ overtwisted at infinity with no closed orbits

Choose coordinates (z, θ), $z \in [0, 2\pi]$, for $\mathbb{S}^2$ using the map

$$(z, \theta) \to (\sqrt{\pi^2 - |z - \pi|^2}\cos(\theta), \sqrt{\pi^2 - |z - \pi|^2}\sin(\theta), z).$$

That is, we consider the sphere of radius π centered at $(0, 0, \pi) \in \mathbb{R}^3$. The z–coordinate is not smooth at the poles, just like the radius is not smooth at the origin of $\mathbb{R}^2$. Take now coordinates $(z, \theta; s)$ in $\mathbb{S}^2 \times \mathbb{R}$, and construct the following 1–form:

$$\lambda_0 = \cos(z)ds + z(z - 2\pi)\sin(z)d\theta.$$

It is easy to see that it is a contact form that defines two families of overtwisted discs sharing a common boundary: $\{z \in [0, \pi], s = s_0\}$ and $\{z \in [\pi, 2\pi], s = s_0\}$. It is therefore overtwisted at infinity.

The form λ_0 defines two cylinders foliated by closed Reeb orbits: $\{z = \pi/2\}$ and $\{z = 3\pi/2\}$. Therefore, proceeding like in the previous example, we will add a small perturbation that gets rid of them. Consider the form:

$$\lambda = \cos(z)ds + [z(z - 2\pi)\sin(z) + f(s)\phi(z)]d\theta.$$

Here we require for $\phi(z)$ to be constant close to the points 0, $\pi/2$, π, $3\pi/2$ and 2π, to satisfy:

$$\phi(0) = \phi(\pi) = \phi(2\pi) = 0, \quad \phi(\pi/2) = \phi(3\pi/2) = 1$$

and to be monotone in the subintervals inbetween. We assume that f is strictly monotone and C^1 small. Computing:

$$d\lambda = -\sin(z)dr\wedge ds+[(z-2\pi)\sin(z)+z\sin(z)+z(z-2\pi)\cos(z)+f(s)\phi'(z)]dz\wedge d\theta+f'(s)\phi(z)ds\wedge d\theta$$

so the Reeb v.f. is a multiple of:

$$X = -f'(s)\phi(z)\partial_z + [(z - 2\pi)\sin(z) + z\sin(z) + z(z - 2\pi)\cos(z) + f(s)\phi'(z)]\partial_s + \sin(z)\partial_\theta,$$

away from $z = 0, \pi, 2\pi$. Near $z = 0, \pi, 2\pi$, the Reeb v.f. is very close to $\pm\partial_s$. Away from those points, it has a non–zero z–component. It follows that it cannot have closed orbits.

5.2.3.3 Constructing the foliation

Consider $\mathbb{S}^2 \times \mathbb{S}^1 \times \mathbb{S}^1$ with coordinates $(z, \theta; s, t)$, $t \in [0, 2]$. It can be endowed with the following 1–form:

$$\tilde{\lambda} = \cos(z)ds + [z(z - 2\pi)\sin(z) + F(t)\phi(z)]d\theta,$$

with F strictly increasing in $(0, 1)$, strictly decreasing in $(1, 2)$, C^1–small and having vanishing derivatives to all orders in $\{0, 1\}$. ϕ is the bump function defined in the previous subsection.

Let $\Phi : \mathbb{S}^1 \to \mathbb{S}^1$ be a diffeomorphism of the circle that fixes $\{0, 1\}$ and no other points, is strictly increasing in $(0, 1)$ as a map $(0, 1) \to (0, 1)$, and is strictly decreasing in $(1, 2)$ as a map $(1, 2) \to (1, 2)$. Φ defines a foliation $\mathcal{F}_\Phi$ on $\mathbb{S}^2 \times \mathbb{S}^1 \times \mathbb{S}^1$ called the **suspension** of Φ.

$\mathcal{F}_\Phi$ can be constructed as follows. Find a family of functions $\Phi_s : \mathbb{S}^1 \to \mathbb{S}^1$, $s \in [0, 1]$, satisfying:

$$\begin{cases} \Phi_0 = \mathrm{Id}, \quad \Phi_1 = \Phi, \\ s \to \Phi_s(t) \text{ is strictly increasing in } (0,1) \text{ and strictly decreasing in } (1,2), \\ \left.\dfrac{\partial}{\partial s}\right|_{s=1} \Phi_s(t) = \left.\dfrac{\partial}{\partial s}\right|_{s=0} \Phi_s(\Phi_1(t)) \quad \text{for all } t. \end{cases} \tag{5.1}$$

Then the curves $\gamma_t(s) = (s, \Phi_s(t))$, induce a foliation in $[0, 1] \times \mathbb{S}^1$ which glues to yield a foliation by curves in the 2–torus. $\mathcal{F}_\Phi$ is the lift of such a foliation.

The leaves of the foliation in the 2–torus are obtained by concatenating the segments γ_t. γ_0 and γ_1 yield closed curves $\tilde{\gamma}_0$ and $\tilde{\gamma}_1$. All other curves are diffeomorphic to $\mathbb{R}$, and we denote them by $\tilde{\gamma}_t(s) = (s, h_t(s))$, $t \in (0, 1) \cup (1, 2)$. By our assumption on Φ_s, the functions h_t are strictly increasing if $t \in (0, 1)$ and strictly decreasing if $t \in (1, 2)$. Observe that the non–compact leaves accumulate onto the two compact ones.

The contact structure in the compact leaves $\mathbb{S}^2 \times \tilde{\gamma}_t$, $t = 0, 1$, is given by

$$\cos(z)ds + z(z - 2\pi)\sin(z)d\theta.$$

In particular, they both have infinitely many closed orbits.

The contact structure in the non compact leaves $\mathbb{S}^2 \times \tilde{\gamma}_t$, $t \in (0, 1) \cup (1, 2)$, reads

$$\cos(z)ds + [z(z - 2\pi)\sin(z) + F(h_t(s))\phi(z)]d\theta$$

Since $F \circ h_t$ is non–zero, strictly monotone and C^1–small, it is of the form described in the previous section. It follows that they have no periodic orbits. □

Remark 5.17. In this example all leaves involved are overtwisted. Further, the non–compact leaves are overtwisted at infinity. It would be interesting to construct an example of a contact foliation where the non–compact leaves are overtwisted, the leaves in their closure are tight and the only periodic orbits appear in the tight leaves.

5.3 Pseudoholomorphic curves in the symplectisation of a contact foliation

In this section we generalise the standard setup for moduli spaces of pseudoholomorphic curves to the foliated setting. The main result is Theorem 5.28, which deals with the removal of singularities. The proof is standard and closely follows that of [38], and indeed the only essential difference lies in the fact that, although the leaves might be open, they live inside a compact ambient manifold, so the Arzelá–Ascoli theorem can still be applied when carrying out the bubbling analysis.

5.3.1 Setup

Consider the contact foliation $(M^{m+2n+1}, \mathcal{F}^{2n+1}, \xi^{2n})$, with extension Θ^{2n+m} given by a 1–form α, and write $(\mathbb{R} \times M, \mathcal{F}_{\mathbb{R}}, \omega)$ for its symplectisation.

5.3.1.1 The space of almost complex structures

The symplectic bundle $(\xi, d\alpha)$ can be endowed with a complex structure compatible with $d\alpha$, which we denote by J_ξ. The space of such choices is non–empty and contractible. J_ξ induces a unique $\mathbb{R}$–invariant leafwise complex structure, $J \in \mathrm{End}(T\mathcal{F}_{\mathbb{R}})$, $J^2 = -\mathrm{Id}$, as follows:

$$J|_\xi = J_\xi$$

$$J(\partial_t) = R$$

Observe that J is **compatible** with ω, and hence they define a metric, which turns each leaf of the symplectisation into a manifold which is not complete. Instead, we shall consider the better behaved $\mathbb{R}$–invariant leafwise riemannian metric g in $\mathbb{R} \times \mathcal{F}$ given by:

$$g = dt \otimes dt + \alpha \otimes \alpha + d\alpha(J_\xi \circ \pi_\xi, \pi_\xi).$$

5.3.1.2 J–holomorphic curves

Let (S, i) be a Riemann surface, possibly with boundary. A map satisfying

$$\begin{cases} F : (S, i) \to (\mathbb{R} \times M, J) \\ dF(TS) \subset T\mathcal{F}_{\mathbb{R}} \\ J \circ dF = dF \circ i \end{cases} \tag{5.2}$$

is called a parametrised **foliated J–holomorphic curve**. The second condition implies that $F(S)$ is contained in a leaf $\mathbb{R} \times \mathcal{L}$ of $\mathcal{F}_{\mathbb{R}}$. Indeed, J is an almost complex structure in the open manifold $\mathbb{R} \times \mathcal{L}$, and F, regarded as a map into $\mathbb{R} \times \mathcal{L}$, is a J–**holomorphic curve** in the standard sense.

By our choice of J, there is an $\mathbb{R}$–action on the space of foliated J–holomorphic curves given by translation on the $\mathbb{R}$ term of $\mathbb{R} \times M$.

5.3.1.3 Foliated J–holomorphic planes and cylinders

A solution of Equation (5.2)

$$F = (a, u) : (\mathbb{C}, i) \to (\mathbb{R} \times M, J)$$

is called a **foliated J–holomorphic plane**. If we write $\mathcal{M}_J^{\mathcal{F}}$ for the space of such maps, it is clear that the space of complex automorphisms of $\mathbb{C}$ acts on it by its action on the domain.

$\mathcal{M}_J^{\mathcal{F}}$ is non–empty. Every Reeb orbit $\gamma : \mathbb{R} \to M$ has an associated foliated J–holomorphic plane given by

$$F(s, t) = (s, \gamma(t)) \qquad \text{where } z = s + it \text{ are the standard complex coordinates in } \mathbb{C}.$$

We call these the *trivial* solutions.

Similarly, a solution of Equation 5.2

$$F = (a, u) : (-\infty, \infty) \times \mathbb{S}^1 \to \mathbb{R} \times M$$

is called a **foliated J–holomorphic cylinder**. We let (s, t) be the coordinates in the cylinder and its complex structure to be given by $i(\partial_s) = \partial_t$. A closed Reeb orbit $\gamma : \mathbb{S}^1 \to M$, gives a *trivial* cylinder $F(s, t) = (s, \gamma(t))$.

Recall that the cylinder $(-\infty, \infty) \times \mathbb{S}^1$ is biholomorphic to $\mathbb{C} \setminus \{0\}$ by the exponential map, and for convenience we will often consider both domains interchangeably. In particular, given some foliated J–holomorphic plane, we could define a foliated J–holomorphic cylinder by introducing a pucture in the domain. Therefore, we say that a foliated J–holomorphic map

$$F = (a, u) : \mathbb{C} \setminus \{0\} \to \mathbb{R} \times M$$

can be **extended** over zero (or ∞) if there is a foliated J–holomorphic map with domain $\mathbb{C}$ (resp. the puctured Riemann sphere $\hat{\mathbb{C}} \setminus \{0\}$) that agrees with F in $\mathbb{C} \setminus \{0\}$.

5.3.1.4 Energy

After introducing the *trivial* foliated J–holomorphic curves, we would like to introduce an *energy constraint* that singles out more interesting solutions of Equation 5.2. This leads us to the following definitions.

Definition 5.18. *Consider the space of functions*

$$\Gamma = \{\phi \in C^\infty(\mathbb{R}, [0,1]) | \quad \phi' \geq 0\}$$

Let $F : S \to \mathbb{R} \times M$ be a foliated J–holomorphic curve.

*Its **energy** is defined by:*

$$E(F) = \sup_{\phi \in \Gamma} \int_S F^* d(\phi\alpha). \tag{5.3}$$

*Its **horizontal energy** is defined by:*

$$E^h(F) = \int_S F^* d\alpha. \tag{5.4}$$

Trivial solutions correspond to the following general phenomenon.

Lemma 5.19. *Let $F = (a,u) : (S,i) \to (\mathbb{R}\times M, J)$ be a foliated J–holomorphic curve. $E^h(F) = 0$ if and only if* Image$(F) \subset \mathbb{R} \times \gamma$*, where γ is a Reeb orbit.*

Proof. Given a ball $U \subset S$ find complex coordinates (s,t). Then:

$$\int_U F^* d\alpha = \int_U d\alpha(u_s, u_t) ds \wedge dt = \int_U d\alpha(u_s, Ju_s) ds \wedge dt =$$

$$\int_U d\alpha(\pi_\xi u_s, \pi_\xi \circ J u_s) ds \wedge dt = \int_U |\pi_\xi u_s|^2 ds \wedge dt$$

and since

$$E^h(F) = \int_S F^* d\alpha = \int_S u^* d\alpha$$

the claim follows. □

The following lemma states that cylinders with finite energy that cannot be extended to planes have to be necessarily trivial and hence imply the existence of a Reeb orbit.

Lemma 5.20. *Let F be a foliated J–holomorphic map*

$$F = (a,u) : \hat{\mathbb{C}} \setminus \{0, \infty\} \to \mathbb{R} \times M$$

satisfying $E(F) < \infty$ and $E^h(F) = 0$. If F cannot be extended over its punctures, then $t \to u(e^{2\pi i t})$, $t \in [0,1]$, is a parametrised closed Reeb orbit.

Proof. By Lemma 5.19, we know that there is some Reeb orbit γ (not necessarily closed) such that Image$(F) \subset \mathbb{R} \times \gamma$. We can identify the universal cover of $\mathbb{R} \times \gamma$ with $\mathbb{C}$ with its standard complex structure. If we let $\Omega \subset \mathbb{C}$ be some region of the complex plane, we can regard it as a J–holomorphic curve using the inclusion $i : \Omega \to \mathbb{C} \to \mathbb{R} \times \gamma$; its energy is just given by:

$$E(i) = \sup_{\phi \in \Gamma} \int_\Omega \phi'(s) ds \wedge dt = \sup_{\phi \in \Gamma} \int_{\partial\Omega} \phi(s) dt$$

where $ds \wedge dt$ is the standard area form in $\mathbb{C}$.

We claim that γ is a closed orbit and that F is a non contractible map into $\mathbb{R} \times \gamma$. Assuming otherwise, regard F as a holomorphic map $f : \hat{\mathbb{C}} \setminus \{0, \infty\} \to \mathbb{C} \subset \hat{\mathbb{C}}$. As such, its punctures are either removable or essential singularities. They cannot be removable singularities with values in $\mathbb{C}$ by assumption.

If f has a removable singularity that is a pole, a neighbourhood of the puncture branch covers a neighbourhood of ∞ in the Riemann sphere. In particular, there is a band $[a, b] \times \mathbb{R} \subset \text{Image}(f) \subset \mathbb{C}$, with $a < b$ large enough. Since the energy of the band is infinite, the energy of F must be too, which is a contradiction.

If f has an essential singularity, then Picard's great theorem states that every point in $\mathbb{C}$, except possibly one, is contained in Image(f). Again, this contradicts the assumption that $E(F)$ was finite.

We deduce that γ is a closed orbit and that F is a non–contractible map into the cylinder $\mathbb{R} \times \gamma$. The exponential is a biholomorphism between the cylinder and $\hat{\mathbb{C}} \setminus \{0, \infty\}$, so now we regard F as a holomorphic map $h : \hat{\mathbb{C}} \setminus \{0, \infty\} \to \hat{\mathbb{C}} \setminus \{0, \infty\}$.

Suppose one of the punctures was an essential singularity for h. Since h has no zeroes or poles, Picard's theorem states that all other points in the Riemann sphere have infinitely many preimages by h. This contradicts $E(F) < \infty$.

Therefore, h can be extended over its punctures to be zero or ∞. h is then a meromorphic function on the Riemann sphere, and hence it is nothing but the quotient of two polynomials. By our assumption that there are no other zeroes or poles this implies that $h(z) = az^k$, for some $k \in \mathbb{Z} \setminus \{0\}$, $a \in \mathbb{C}$. This shows that $t \to u(e^{2\pi i t})$ parametrises the k–fold cover of γ. □

Exactly the same analysis yields the following lemma.

Lemma 5.21. *Let F be a foliated J–holomorphic map*

$$F = (a, u) : \mathbb{C} \to \mathbb{R} \times M$$

satisfying $E^h(F) = 0$. Then either F is the constant map or $E(F) = \infty$.

Proof. Let γ be the Reeb orbit such that Image$(F) \subset \mathbb{R} \times \gamma$. By taking the universal cover of $\mathbb{R} \times \gamma$, regard F as a map $\mathbb{C} \to \mathbb{C}$, as in Lemma 5.20. Now study the extension problem of F to ∞. If it corresponds to a removable singularity with values in $\mathbb{C}$, then F is the constant map. Otherwise, if it is either a pole or a non–removable singularity, it has infinite energy. □

5.3.1.5 Riemannian and symplectic area

In the case of compact symplectic manifolds, there is an interplay between the symplectic area of a J–holomorphic curve and its riemannian area for the metric given by the symplectic form and the compatible almost complex structure.

In our case, g is not of that form. Rather, it is $\mathbb{R}$–invariant, while ω is not: $\mathbb{R}$–translations of the same J–holomorphic curve have different symplectic energy and indeed there are no universal constants relating the ω–area and the g–area.

However, E and E^h are invariant under the $\mathbb{R}$-action. Given F, a foliated J–holomorphic curve, let $\text{area}_g(F)$ be its riemannian area in terms of g, and let $\text{area}_{\omega_\phi}(F)$ be its symplectic area in terms of $\omega_\phi = d(\phi\alpha)$.

Lemma 5.22. *Let $F = (a, u) : (S, i) \to (\mathbb{R} \times M, J)$ be a parametrised foliated J–holomorphic curve. Then, if a is bounded below and above:*

$$\text{area}_g(F) < C\text{area}_{\omega_\phi}(F) < C' \int_{\partial S} \alpha,$$

for some constants C, C' depending only on the upper and lower bounds of a.

Proof. Consider a_0 and a_1 satisfying $a_0 < a < a_1$. Let $\phi(t) = \frac{t-a_0}{3(a_1-a_0)} + 1/3$ in $[a_0, a_1]$ and belonging to Γ. Then ω_ϕ is a symplectic form in $[a_0, a_1] \times M$ and J is ω_ϕ–compatible. Since $0 < D < \phi, \phi' < D' < \infty$, there are universal constants relating the metrics g and $g_\phi = \omega_\phi(-, J-)$ in $[a_0, a_1] \times M$.

Since J is ω_ϕ–compatible, F being J–holomorphic implies that $\text{area}_{g_\phi}(F) = \text{area}_{\omega_\phi}(F)$, and the first inequality follows. The second inequality follows by applying Stokes. □

An immediate consequence of Lemma 5.22 is that there cannot be *closed* foliated J–holomorphic curves in $\mathbb{R} \times M$.

5.3.2 Bubbling

As we shall see in Section 5.4, the way in which we will prove the existence of a periodic orbit of the Reeb vector field will be by constructing a 1–dimensional moduli of pseudoholomorphic discs that necessarily will be open in one of its ends. The following lemma shows that the reason for it to be open must be that the gradient is not uniformly bounded for all discs in the moduli.

Proposition 5.23. *Fix $\mathcal{L}$ a leaf of $\mathcal{F}$. Let $W \subset \mathbb{R} \times \mathcal{L}$ be a totally real compact submanifold, possibly with boundary.*

Let (S, i) be a compact Riemann surface with boundary. Consider the sequence of foliated J–holomorphic maps

$$F_k : (S, \partial S) \to (\mathbb{R} \times \mathcal{L}, W), \quad k \in \mathbb{N}.$$

Suppose that there is a uniform bound $||dF_k|| < C < \infty$. Then there is a subsequence F_{k_i}, $k_i \to \infty$, convergent in the C^∞–topology to a foliated J–holomorphic map

$$F_\infty : (S, \partial S) \to (\mathbb{R} \times \mathcal{L}, W)$$

Proof. Observe that since we have a uniform gradient bound and $F_k(\partial S) \subset W$, for all k, it necessarily follows that the images of all the F_k lie in a compact subset of $\mathbb{R} \times \mathcal{L}$. Then, one can proceed as in the standard case to prove C^∞ bounds from C^1 bounds and then apply the Arzelá-Ascoli theorem to complete the proof. □

Remark 5.24. The same statement holds for surfaces without boundary as long as one imposes for the images of all the F_k to lie in a compact set of the leaf.

Proposition 5.23 suggests that we should study sequences of maps

$$F_k : (S, \partial S) \to (\mathbb{R} \times \mathcal{L}, W), \quad k \in \mathbb{N}$$

in which $||dF_k||$ is not uniformly bounded. We have to consider two separate cases.

5.3.2.1 Plane bubbling

Proposition 5.25. *Consider a sequence of foliated J–holomorphic curves*

$$F_k : (S, \partial S) \to (\mathbb{R} \times \mathcal{L}, W), \quad k \in \mathbb{N}$$

and a corresponding sequence of points q_k in S having $M_k = ||d_{q_k} F_k|| \to \infty$ and converging to a point $q \in S$.

Suppose that there is an uniform bound $E(F_k) < C < \infty$. If $\operatorname{dist}(q_k, \partial S)M_k \to \infty$, there is a foliated J–holomorphic plane

$$F_\infty : \mathbb{C} \to \mathbb{R} \times \mathcal{L}'$$

with $E(F_\infty) < C$, where $\mathcal{L}'$ is a leaf in the closure of $\mathcal{L}$.

Proof. After possibly modifying the q_k slightly, there are charts

$$\phi_k : \mathbb{D}^2(R_k) \to S$$

$$\phi_k(z) = q_k + \frac{z}{M_k}$$

with $R_k < \operatorname{dist}(q_k, \partial S)M_k$, $R_k \to \infty$, $R_k/M_k \to 0$, and $||d(F_k \circ \phi_k)|| < 2$ – this last condition is achieved by the so called Hofer-Viterbo lemma, see [38, Lemma 26] and [38, p. 536. Equation 49].

The maps $F_k \circ \phi_k$ have C^1 bounds by construction, but they have no C^0 bounds. By our construction of J, the vertical translation of a J–holomorphic map is still J–holomorphic and hence we can compose with a vertical translation τ_k guaranteeing that $\tau_k \circ F_k \circ \phi_k$ takes the point 0 to the level $\{0\} \times \mathcal{L}$. Then, for every compact subset $\Omega \subset \mathbb{C}$, the maps $\tau_k \circ F_k \circ \phi_k : \Omega \to \mathbb{R} \times M$ are equicontinuous and bounded – note that this is where we use that $\mathcal{L}$ lies inside the compact manifold M.

Recall that having uniform C^1 bounds implies that we have uniform C^∞ bounds. Hence, an application of the Arzelá–Ascoli theorem shows that a subsequence converges in C^∞_{loc} to a map $F_\infty : \mathbb{C} \to \mathbb{R} \times M$ that must be foliated and J–holomorphic, but not necessarily lying in $\mathbb{R} \times \mathcal{L}$, but maybe in some new leaf $\mathbb{R} \times \mathcal{L}'$.

Note that the energy of the map $\tau_k \circ F_k \circ \phi_k$ is bounded above by that of F_k. Since we have uniform bounds for the energy of the F_k, we have uniform energy bounds for the maps $\tau_k \circ F_k \circ \phi_k$ and hence for their limit F_∞. Note that F_∞ is necessarily non constant, since $||d_0 F_\infty|| = 1$ by construction. In particular, it has non–zero energy. □

Remark 5.26. We say that the map F_∞ as given in the proof is called a **plane bubble**. If the map F_∞ could be extended over the pucture to a map with domain the Riemann sphere $\mathbb{S}^2$, this would yield a contradiction with Lemma 5.22.

5.3.2.2 Disc bubbling

Proposition 5.27. *Consider a sequence of foliated J–holomorphic curves*

$$F_k : (S, \partial S) \to (\mathbb{R} \times \mathcal{L}, W), \quad k \in \mathbb{N}$$

and a corresponding sequence of points q_k in S having $M_k = ||d_{q_k} F_k|| \to \infty$ converging to a point $q \in S$.

Suppose that there is an uniform bound $E(F_k) < C < \infty$. If $\mathrm{dist}(q_k, \partial S)M_k$ *is uniformly bounded from above, there is a foliated J–holomorphic disc*

$$F_\infty : (\mathbb{D}^2, \partial\mathbb{D}^2) \to (\mathbb{R} \times \mathcal{L}, W)$$

with $E(F_\infty) < C$.

Proof. Recall that W is compact. An application of the standard rescaling argument yields a finite energy J–holomorphic map F_∞ of the upper half plane $\mathbb{H}^+$ into $\mathbb{R} \times M$. Since the rescaling is done close to ∂S and $F_k(\partial S) \subset W$, we deduce that F_∞ maps the boundary of $\mathbb{H}^+$ to $W \subset \mathbb{R} \times \mathcal{L}$; from this, it follows that the image of $\mathbb{H}^+$ lies in $\mathbb{R} \times \mathcal{L}$. Now one can apply the removal of singularities theorem from [54, Proposition A.1] to conclude. □

5.3.3 Removal of singularities

The aim of this subsection is to prove the following result, which is one of the key ingredients for proving Theorem 5.1.

Theorem 5.28 (Removal of singularities)**.** *Let $F = (a, u) : \mathbb{D}^2 \setminus \{0\} \to \mathbb{R} \times \mathcal{L} \subset \mathbb{R} \times M$ be a J–holomorphic curve with $0 < E(F) < \infty$, $\mathcal{L}$ a leaf of $\mathcal{F}$.*

Then, either F extends to a J–holomorphic map over $\mathbb{D}^2$ or for every sequence of radii $r_k \to 0$ the curves $\gamma_{r_k}(s) = u(e^{r_k+is})$ converge in C^∞ –possibly after taking a subsequence– to a parametrised closed Reeb orbit lying in the closure of $\mathcal{L}$.

Proof of Theorem 5.28. Let us state the problem in terms of cylinders. Identify $\mathbb{D}^2 \setminus \{0\}$ with $[0, \infty) \times \mathbb{S}^1$ by using the biholomorphism $-\log(z)$, and regard F as a foliated J–holomorphic map $[0, \infty) \times \mathbb{S}^1 \to \mathbb{R} \times M$. Then, the following maps are foliated J–holomorphic:

$$F_k = (a_k, u_k) : [-R_k/2, \infty) \times \mathbb{S}^1 \to \mathbb{R} \times M$$

$$F_k(s, t) = (a(s + R_k, t) - a(R_k, 0), u(s + R_k, t))$$

and by assumption they have a uniform bound $E(F_k) < C < \infty$ and $\lim_{k\to\infty} E^h(F_k) = 0$. Here $R_k = -\log(r_k) \to \infty$.

Suppose that the gradient was not uniformly bounded for the family F_k. We can then find a sequence of points $q_k \in [0, \infty) \times \mathbb{S}^1$ escaping to infinity and satisfying $|d_{q_k}F| \to \infty$. Then we are under the assumptions of Proposition 5.25, and this yields a plane bubble $G : \mathbb{C} \to \mathbb{R} \times M$ with $E^h(G) = 0$, which must lie on top of a Reeb orbit by Lemma 5.19. By our bubbling analysis, it cannot be constant, since its gradient at the origin is 1, which is a contradiction with it having $E(G) < \infty$, by Lemma 5.21.

We conclude that the family F_k has uniform C^1 bounds and hence uniform C^∞ bounds. By construction $a_k(0, 0) \in \{0\} \times M$, which means that we have uniform C^0 bounds on every compact subset of $(-\infty, \infty) \times \mathbb{S}^1$ –here is where we use the compactness of M. The Arzelá-Ascoli theorem implies that –after possibly taking a subsequence– the maps F_k converge in C^∞_{loc} to a map $F_\infty : (-\infty, \infty) \times \mathbb{S}^1 \to \mathbb{R} \times M$ with $E(F_\infty) < \infty$ and $E^h(F_\infty) = 0$, which might of course have image in $\mathbb{R}$ times a different leaf $\mathcal{L}'$.

Observe that

$$\lim_{r\to 0} \int_{\gamma_r} \alpha = \int_{\gamma_1} \alpha - \int_{\mathbb{D}^2\setminus\{0\}} d\alpha.$$

If this limit is zero, then the argument above shows that the γ_r, $r \to 0$, tend to the constant map in the C^∞ sense, and hence F extends to a map over $\mathbb{D}^2$. Assuming otherwise, it is clear that F_∞ cannot be the constant map and hence Lemma 5.20 implies the conclusion. $\square$

5.4 Existence of contractible periodic orbits in the closure of a leaf

After setting up the study of foliated J–holomorphic curves in the previous section and dealing with its compactness issues, we use this machinery to conclude the proof of Theorem 5.1. The setting of the theorem is as follows: $(M^{m+3}, \mathcal{F}^3, \xi^2)$ is a contact foliation with Θ^{2+m} an extension given by a 1–form α. We write $(\mathbb{R} \times M, \mathcal{F}_\mathbb{R}, \omega)$ for its symplectisation. $\mathcal{L}^3$ is a leaf of $\mathcal{F}$.

5.4.1 The Bishop family

The following results have a local nature and hence do not depend on whether $\mathcal{L}$ is compact or not. Their proofs can be found in [38].

5.4.1.1 The Bishop family at an elliptic point

If $(\mathcal{L}, \xi)$ is an overtwisted manifold, let Σ be an overtwisted disc for ξ. Otherwise, if $\pi_2(\mathcal{L}) \neq 0$, let Σ be some sphere realising a non–zero class in π_2. Assume, after a small perturbation, that the characteristic foliations are as described in Subsection 5.1.3 in Exercises 5.7 and 5.8 and Theorem 5.9. Denote by Γ_Σ the set of singular points of the characteristic foliation of Σ.

Let $p \in \Gamma_\Sigma$, a elliptic point. The maps satisfying:

$$\begin{cases} F = (u, a) : (\mathbb{D}^2, \partial\mathbb{D}^2) \to (\mathbb{R} \times \mathcal{L}, \{0\} \times \Sigma) \\ dF \circ i = J \circ dF, \\ \text{wind}(F, p) = \pm 1, \\ \text{index}(F) = 4, \end{cases} \tag{5.5}$$

will be called the **Bishop family**. wind(F, p) refers to the winding number of $F(\partial\mathbb{D}^2)$ around the elliptic point p.

The condition index$(F) = 4$ is implied by the other assumptions. It means that the linearised Cauchy–Riemann operator at F has index 4, and hence, if there is transversality, the solutions of Equation 5.5 close to F form a smooth 4–dimensional manifold. Since the Mobius transformations of the disc have real dimension 3, this implies that the image of F is part of a 1–dimensional family of distinct discs.

The Bishop family is not empty under some integrability assumptions.

Proposition 5.29 ([5], [38, Section 4.2])**.** *For a suitable choice of J_ξ, J is integrable close to p. Then there is a smooth family of maps F_s, $s \in [0, \varepsilon)$, with $F_0(z) = p$ and F_s, $s > 0$, disjoint embeddings satisfying Equation 5.5.*

Additionally, there is a small neighbourhood U of p such that any other disc satisfying Equation 5.5 and interesecting U is a reparametrisation of one of the F_s.

5.4.1.2 Continuation of the Bishop family

The following statement shows that transversality always holds for the linearised Cauchy–Riemann operator for maps belonging to the Bishop family.

Proposition 5.30 ([38, Theorem 17]). *Let F satisfy Equation 5.5. Then there is a smooth family of disjoint embeddings F_s, $s \in (-\varepsilon, \varepsilon)$, satisfying Equation 5.5, such that $F_0 = F$. Additionally, any two such families are related by a reparametrisation of the parameter space and a smooth family of Mobius transformations.*

5.4.1.3 Properties of the Bishop family

Convexity of $\{0\} \times \mathcal{L}$ inside of $\mathbb{R} \times \mathcal{L}$ and an application of the maximum principle yield the following lemma. It will be useful to show that there is no disc bubbling.

Lemma 5.31 ([38, Lemma 19]). *Let $F : (\mathbb{D}^2, \partial\mathbb{D}^2) \to (\mathbb{R} \times \mathcal{L}, \{0\} \times \Sigma)$ be a J–holomorphic map. Then $F(\partial\mathbb{D}^2)$ is transverse to the characteristic foliation of Σ and $F(\mathbb{D}^2)$ is transverse to $\{0\} \times \mathcal{L}$.*

In order to apply Theorem 5.28 we must have energy bounds, which are provided by the following result.

Proposition 5.32 ([38, Lemmas 33 and 35]). *There are uniform energy bounds $0 < C_1 < E(F), E^h(F) < C_2 < \infty$ for every F satisfying Equation 5.5 and having*

$$\mathrm{dist}(\mathrm{Image}(F), \Gamma_\Sigma) > \varepsilon > 0.$$

Proof. By Stokes' theorem:

$$E(F) = \sup_{\phi\in\Gamma} \int_{\mathbb{D}^2} F^* d(\phi\alpha) = \sup_{\phi\in\Gamma} \int_{\partial\mathbb{D}^2} F^*\phi\alpha =$$

$$\sup_{\phi\in\Gamma} \phi(0) \int_{\partial\mathbb{D}^2} F^*\alpha = \int_{F(\partial\mathbb{D}^2)} \alpha.$$

$F(\partial\mathbb{D}^2)$ winds around the critical point exactly once and hence bounds a disc within Σ. The area of such a disc is always bounded above by a universal constant and is bounded below under the assumption that they have radius at least ε. The claim follows.

A similar estimate holds for E^h. □

5.4.2 Proof of Theorem 5.1

Now we tie all the results we have discussed so far.

Lemma 5.33. *Let $\mathcal{L}$ be a leaf of $\mathcal{F}$ and assume that $(\mathcal{L}, \xi)$ is an overtwisted contact manifold. Then there is a finite energy plane contained in $\mathbb{R} \times \mathcal{L}' \in \mathcal{F}_{\mathbb{R}}$, with $\mathcal{L}'$ lying in the closure of $\mathcal{L}$.*

Proof. Denote by $\mathcal{M}$ the set of solutions of Equation 5.5, which is non–empty by Proposition 5.29 and open by Proposition 5.30. Recall that Σ is the overtwisted disc and define a non–negative constant

$$C = \inf\{\mathrm{dist}(u(\partial\mathbb{D}^2), \partial\Sigma) | u \in \mathcal{M}\}.$$

By construction there is a sequence of maps $u_k \in \mathcal{M}$ such that $\lim_{k\to\infty} \operatorname{dist}(u_k(\partial\mathbb{D}^2), \partial\Sigma) = C$. Suppose that the gradient of the sequence is unbounded. Then Propositions 5.25 and 5.27 show that either a plane or a disc bubble appears. In the case of a disc bubble, Lemma 5.31 states that its boundary must wind around the elliptic point once. Lemma 5.31 also implies that, for every k, $u_k(\partial\mathbb{D}^2)$ intersects each leaf of the characteristic foliation exactly once. These two facts show that (after possibly taking a subsequence), the circles $u_k(\partial\mathbb{D}^2)$ converge to the boundary circle of the disc bubble, which is then a solution of Equation 5.5 whose distance to the boundary is exactly C. Proposition 5.30 shows that it must be part of a 1–parametric family of unparametrised discs, contradicting the fact that C was the infimum of the distances. We conclude that necessarily a plane bubble must appear instead.

Otherwise, if the gradient is uniformly bounded for the sequence, Proposition 5.23 shows that the u_k converge to a new J–holomorphic map u_∞. If $C = 0$, $u_\infty(\partial\mathbb{D}^2)$ is not transverse to $\partial\Sigma$, which is a contradiction with Lemma 5.31. If otherwise $C > 0$, then Proposition 5.30 shows that u_∞ is part of a 1–parametric family of unparametrised discs, contradicting the fact that C was the infimum of the distances. □

Lemma 5.34. *Let $\mathcal{L}$ be a leaf of $\mathcal{F}$ and assume that $\pi_2(\mathcal{L}) \neq 0$. Then there is a finite energy plane contained in $\mathbb{R}\times\mathcal{L}' \in \mathcal{F}_{\mathbb{R}}$, with $\mathcal{L}'$ lying in the closure of $\mathcal{L}$.*

Proof. Let us denote by p_- and p_+ the two elliptic points of the convex 2–sphere Σ realising a non trivial element of $\pi_2(\mathcal{L})$. Denote by $\mathcal{M}$ the set of solutions of Equation 5.5. There are two connected components $\mathcal{M}^-, \mathcal{M}^+ \subset \mathcal{M}$, distinguished by the fact that they contain the Bishop families arising from the points p_- and p_+, respectively.

We now prove that actually $\mathcal{M}^- = \mathcal{M}^+$. Define a constant

$$C = \inf\{\operatorname{dist}(u(\partial\mathbb{D}^2), p_+) | u \in \mathcal{M}^-\}.$$

Reasoning as in Lemma 5.33 shows that, unless the gradient explodes and hence a plane bubble appears, we must necessarily have $C = 0$. By Proposition 5.29, the only curves in a neighbourhood of p_+ are those in $\mathcal{M}^+$, and hence $\mathcal{M}^- = \mathcal{M}^+$. The evaluation map

$$\operatorname{ev} : \mathcal{M}^- \times \mathbb{D}^2 \approx [0,1]\times\mathbb{D}^2 \to \mathcal{L}$$

$$\operatorname{ev}(F = (a,u), z) = u(z)$$

satisfies $\operatorname{ev}(\partial(\mathbb{M}^- \times \mathbb{D}^2)) = \Sigma$, which contradicts the fact that Σ was non–trivial in $\pi_2(\mathcal{L})$.

Therefore, the gradient must explode and the claim follows. □

We are now ready to prove Theorem 5.1.

Proof of Theorem 5.1. Let $(\mathcal{L}, \xi)$ be overtwisted. Lemma 5.33 yields a finite energy plane $F : \mathbb{C} \to \mathbb{R}\times\mathcal{L}'$, with $\mathcal{L}'$ a leaf of $\mathcal{F}$ contained in the closure of $\mathcal{L}$. By Lemma 5.22 this plane cannot be completed to a sphere. Now an application of Theorem 5.28 shows that there is a closed Reeb orbit in some leaf $\mathcal{L}''$ lying in the closure of $\mathcal{L}'$. Since $\mathcal{L}''$ is in the closure of $\mathcal{L}$ the claim follows.

Same argument goes through by applying Lemma 5.34 if $\pi_2(\mathcal{L}) \neq 0$. □

Remark 5.35. As we have seen, Lemmas 5.33 and 5.34 yield a finite energy plane in a leaf that might not be the one containing the overtwisted disc or the convex 2–sphere. Then, an application of Theorem 5.28 shows that the plane is asymptotic to a trivial cylinder that might live yet in another leaf.

Our example in Subsection 5.2.3 shows that at least one of these two phenomena must take place. Is it possible for a "*double jump*" to actually happen?

Remark 5.36. Let $(M^{2n+1+m}, \mathcal{F}^{2n+1}, \xi)$ be a contact foliation. Let $\mathcal{L}$ be a leaf of $\mathcal{F}$ and let $(\mathcal{F}, \xi)$ be overtwisted in the sense of [6]. More generally, assume that $(\mathcal{F}, \xi)$ contains a plastikstufe [61]. It is immediate that the Bishop family arising from the plastikstufe can be employed to show that there must be a Reeb orbit, so Theorem 5.1 also holds true for overtwisted manifolds in all dimensions. Similarly, Theorem 5.1 also holds for manifolds containing a Lob, as defined in [62], generalising the case $\pi_2(\mathcal{L}) \neq 0$.

5.5 The non–degenerate case

In this section we show that under non–degeneracy assumptions none of the *jumps* between leaves can happen.

Definition 5.37. *Let $(M^{3+m}, \mathcal{F}^3, \xi)$ be a contact foliation and let α be the defining 1–form for some extension Θ of ξ. The form α is called* ***non–degenerate*** *if the set of closed orbits of its Reeb vector field in any leaf of $\mathcal{F}$ is discrete.*

More precisely we claim that the space of orbits in any given leaf, understood as a topological subspace of the space of loops of that leaf equipped with the C^1–topology, consists of isolated elements.

The reader might wonder why we do not define non–degenerate contact forms to be those having leafwise *non–degenerate* Reeb orbits. The reason is that, for a generic choice of α, the Reeb orbits that appear behave in the same way as the Reeb orbits appearing in an m–dimensional family of (non–foliated) contact forms. In particular, as soon as $m > 0$, degenerate Reeb orbits do appear. However, it is still generic for the orbits to be leafwise isolated.

The statement we want to show is the following. It is a stronger version of the Removal of Singularities (Theorem 5.28) in the non–degenerate case.

Theorem 5.38. *Let $(M, \mathcal{F}, \xi)$ be a contact foliation and let α be the defining 1–form for some extension Θ of ξ. Assume α is non–degenerate.*

Let $F = (a, u) : \mathbb{D}^2 \setminus \{0\} \to \mathbb{R} \times \mathcal{L} \subset \mathbb{R} \times M$ be a J–holomorphic curve with $0 < E(F) < \infty$, $\mathcal{L}$ a leaf of $\mathcal{F}$.

Then, either F extends to a J–holomorphic map over $\mathbb{D}^2$ or the curves $\gamma_r(s) = u(e^{r+is})$ converge in C^∞ to a closed Reeb orbit γ lying in $\mathcal{L}$.

Proof. We proceed by contradiction. Assume that γ, the limit of some γ_{r_i}, $r_i \to \infty$, is contained in some leaf $\mathcal{L}' \neq \mathcal{L}$.

Denote $T = \int_\gamma \alpha$, the period of γ. By our assumption on α, we can find a closed foliation chart $U \subset M$ diffeomorphic to $\mathbb{D}^2 \times \mathbb{S}^1 \times [-1, 1]$ around γ such that the plaque in U containing γ intersects no other orbits of period approximately T. Write $h : U \to [-1, 1]$ for the height function of the chart: we can assume that $h^{-1}(0)$ is the plaque containing γ.

Since the curves γ_{r_i} converge in C^∞ to γ, their images are contained in U for large enough i. Assume, by possibly restricting to a subsequence, that each Image(γ_{r_i}) lies in a different plaque of $\mathcal{F} \cap U$. Then, for each i, there is a smallest radius $r_i < R_i < r_{i+1}$ such that Image(γ_{R_i}) intersects ∂U.

Consider the maps

$$F_i : [r_i - R_i, r_{i+1} - R_i] \times \mathbb{S}^1 \to \mathbb{R} \times M$$
$$F_i(t,s) = (a(e^{t+R_i+is}) - a(e^{R_i}), u(e^{t+R_i+is}))$$

By construction, $F_i(0,0) \in \{0\} \times M$, $F_i(0,s) \cap \{0\} \times (\partial U) \neq \emptyset$, and $\lim_{i\to\infty} h \circ F_i = 0$

By carrying out the bubbling analysis, we can assume that the F_i have bounded gradient. In particular, $r_{i+1} - r_i$ must be uniformly bounded from below by a non–zero constant. The Arcelá–Ascoli theorem states that the F_i converge in C^∞_{loc} –maybe after taking a subsequence– to a map F_∞ with $E^h(F_\infty) = 0$ and therefore lying on top of some Reeb orbit.

By the properties of the F_i, F_∞ must have image contained in $\mathbb{R} \times \mathcal{L}'$ and intersecting $\mathbb{R} \times (h^{-1}(0) \cap \partial U)$. In particular, Image$(F_\infty)$ is not contained in $\mathbb{R} \times \gamma$. If $\lim_{i\to\infty} R_i - r_i < \infty$, the curves $s \to F_i(r_i - R_i, s)$ would converge to γ, which is a contradiction. Similarly we deduce that $\lim_{i\to\infty} r_{i+1} - R_i = \infty$.

Since it has finite energy, $F_\infty : (-\infty, \infty) \times \mathbb{S}^1 \to \mathbb{R} \times \mathcal{L}'$ must yield a periodic orbit of the Reeb flow. It must be a closed orbit different from γ, having period T and intersecting the plaque containing γ, which is a contradiction. Since the only additional assumption we made was that γ was contained in $\mathcal{L}' \neq \mathcal{L}$, we deduce that γ must lie in $\mathcal{L}$ itself. Arguing as above, it is clear that the limit γ does not depend on the chosen sequence r_i. □

Remark 5.39. Theorem 5.38 immediately implies that a finite energy plane is asymptotic to a trivial cylinder lying in the same leaf.

Similarly, it shows that the Bishop family always yields a plane bubble in the original leaf $\mathcal{L}$: outside of a finite set of points, the Bishop family converges to foliated J–holomorphic curve with boundary in the overtwisted disc and possibly many punctures that are asymptotic at $-\infty$ to a number of Reeb orbits necessarily lying in $\mathcal{L}$. In particular, under the non–degeneracy assumption, an overtwisted leaf (or a leaf with $\pi_2(\mathcal{L}) \neq 0$) contains a closed Reeb orbit.

Do note, however, that this does not prove the Weinstein conjecture for non–degenerate overtwisted open manifolds arising as leaves of compact contact foliations. Indeed, the space of non–degenerate contact forms of $\mathcal{L}$ as an abstract manifold is, a priori, larger than the space of non–degenerate contact forms of $\mathcal{L}$ as a leaf.

Part II

Engel structures

Chapter 1

What are Engel structures and why do we care?

We motivated our interest in Engel structures in the Preamble: they constitute one of the first examples of bracket–generating distributions one might consider, and they are one of the four classes of distributions that are topologically stable; this probably qualifies as an answer to the question posed in the title. In this chapter we will review the literature on Engel structures, and we will explain some basic results that will be useful later, the most important one being Proposition 1.12.

The paradigm the reader should have in mind is contact topology. Our motto is that the study of Engel structures should be approached from a topological point of view, trying to find analogues of the landmark results about contact structures. The results in this thesis should be understood as the Engel versions of the classification of overtwisted structures [19], and the flexibility of legendrian/transverse immersions.

Until examples of rigidity are found, it is unclear whether there is an interesting topological theory to be studied. However, the work in this thesis might shed some light on what a reasonable candidate of rigid behaviour might be.

The reader should recall the definitions of contact and even–contact structures from Part I, Chapter 1; we will not review them here.

1.1 Definition and some elementary results

Bracket–generating rank 2 distributions in 4–manifolds receive the following more commercial name:

Definition 1.1. *Let M a 4–dimensional manifold. A 2–dimensional distribution $\mathcal{D} \subset TM$ is said to be an* ***Engel structure*** *if it is everywhere maximally non–integrable, i.e. if $\mathcal{E} = [\mathcal{D}, \mathcal{D}]$ is an even–contact structure.*

Engel structures interact heavily with contact/even–contact structures. This is particularly true locally:

Proposition 1.2. *Let M be a 4–dimensional manifold. Let $\mathcal{D} \subset TM$ be an Engel structure. Let $\mathcal{E} = [\mathcal{D}, \mathcal{D}]$ be the associated even–contact structure. Let $\mathcal{W}$ be the kernel of $\mathcal{E}$. Then:*

- *the line field $\mathcal{W}$ is contained in $\mathcal{D}$,*
- *let $N \subset M$ be a (possibly open) 3–dimensional submanifold of M that is transverse to $\mathcal{W}$. Then, $\xi = TN \cap \mathcal{E}$ is a contact structure in N. Additionally, $\mathcal{X} = TN \cap \mathcal{D} \subset \xi$ is a distinguished legendrian line field.*
- *there are two canonical isomorphisms given by Lie bracket:*

$$\det(\mathcal{E}/\mathcal{W}) \cong TM/\mathcal{E}. \tag{1.1}$$

$$\det(\mathcal{D}) \cong \mathcal{E}/\mathcal{D}. \tag{1.2}$$

Proof. Let p be a point of M where it holds that $\mathcal{W} \not\subset \mathcal{D}$ in $\mathcal{O}p(p) \subset M$. Then, we have the equality

$$TM = [\mathcal{E}, \mathcal{E}] = [\mathcal{D} \oplus \mathcal{W}, \mathcal{D} \oplus \mathcal{W}] \subset \mathcal{E}$$

since, by definition, any flow along $\mathcal{W}$ preserves $\mathcal{E}$, and $[\mathcal{D}, \mathcal{D}] = \mathcal{E}$. This is a contraction, so we deduce that $\mathcal{W} \subset \mathcal{D}$ everywhere.

Checking that $\xi = TN \cap \mathcal{E}$ is a contact structure in N can be checked locally, so let us assume that $\mathcal{E}$ is coorientable with α its defining 1–form. Then, the restriction of $\alpha \wedge d\alpha$ to N has maximal rank, because $\mathcal{W} = \ker(\alpha \wedge d\alpha)$ is transverse to N by hypothesis. The second claim follows.

Lastly, consider the following morphisms:

$$[-,-] : \Gamma(\mathcal{E}) \times \Gamma(\mathcal{E}) \to \Gamma(TM) \to \Gamma(TM/\mathcal{E}),$$

$$[-,-] : \Gamma(\mathcal{D}) \times \Gamma(\mathcal{D}) \to \Gamma(\mathcal{E}) \to \Gamma(\mathcal{E}/\mathcal{D}),$$

where Γ denotes the sheaf of sections. By definition, Lie bracket is antisymmetric. A simple computation shows that these operators (after composing the arrows) are of zero order. This, along with the fact that $\mathcal{D}$ is bracket generating, immediately implies the third claim. ☐

The second statement in Proposition 1.2 has the nice consequence (as we shall see later in Proposition 1.12 and Example 1.17), that locally one can think about Engel structures "slice by slice": we fix a 3–submanifold N and we flow it by a vector field contained in $\mathcal{W}$. As we do so, the contact structure in N remains fixed (because $\mathcal{W}$ preserves $\mathcal{E}$), but the preferred legendrian line field $\mathcal{X}$ moves.

Remark 1.3. Observe that the isomorphism of Equation (1.1) has nothing to do with the Engel structure. Neither does the fact that N inherits a contact structure. Both facts are true for general even–contact structures $\mathcal{E}$.

1.1.1 Orientations

A recurring issue in this thesis (as often happens in life) will be that we have to keep track of orientations.

Lemma 1.4. *Let $\mathcal{D}$ be an Engel structure. Then $\mathcal{E} = [\mathcal{D}, \mathcal{D}]$ is canonically oriented.*

Proof. Using the isomorphism given in Equation (1.2) we obtain the following bundle identifications:

$$\mathcal{E} \cong \mathcal{D} \oplus \mathcal{E}/\mathcal{D},$$

$$\det(\mathcal{E}) \cong \det(\mathcal{D}) \otimes \det(\mathcal{E}/\mathcal{D}) \cong \underline{\mathbb{R}}.$$

No auxiliary choice is needed for these isomorphisms, so $\det(\mathcal{E})$ receives the orientation of the trivial line bundle $\underline{\mathbb{R}}$. This is alternately seen by finding a frame $\{W, X\}$ for $\mathcal{D}$ and setting $\mathcal{E}$ to be oriented by the frame $\{W, X, [W, X]\}$. Switching X and W or changing their sign leaves the orientation of the frame the same, so this orientation does not depend on how W or X are chosen. The reader can check that the two methods of orienting $\mathcal{E}$ are actually the same. □

It is easy to see that if $\mathcal{E}$ is an arbitrary even–contact structure, it is not canonically oriented; what we did relies on the existence of $\mathcal{D}$. In particular, we shall see that different $\mathcal{D}$ might produce the same even–contact structure with *opposite orientations*.

Now, regarding the orientation of M:

Lemma 1.5. *Let $\mathcal{E}$ be an even–contact structure. Then, an orientation of TM is equivalent to an orientation of the kernel $\mathcal{W}$.*

Proof. This is a consequence of Isomorphism (1.1). Indeed, write:

$$TM \cong \mathcal{W} \oplus \mathcal{E}/\mathcal{W} \oplus TM/\mathcal{E},$$

$$\det(TM) \cong \det(\mathcal{W}) \otimes \det(\mathcal{E}/\mathcal{W}) \otimes \det(TM/\mathcal{E}) \cong \det(\mathcal{W}).$$

Alternately, we can take a frame $\{W \subset \mathcal{W}, X, Y\}$ for $\mathcal{E}$ and consider the orientation of TM given by the frame $\{W, X, Y, [X, Y]\}$. Then, it is immediate that the result does not depend on the choice of X or Y or how they are ordered, but it does depend on the sign of W. □

In particular, a consequence of this discussion is that the orientation of $\mathcal{D}$ (or lack thereof) has nothing to do with the orientation of the manifold. In any case, it is immediate that:

Lemma 1.6. *Let $\mathcal{D}$ be an oriented Engel structure with oriented kernel $\mathcal{W}$. Then, the manifold M has a parallelisation that is unique up to homotopy.*

1.1.2 Formal Engel structures

The Engel condition is a partial differential relation of second order. Namely, the 3–distribution $\mathcal{E}$ is obtained from $\mathcal{D}$ by a first order operator, and requiring for $\mathcal{E}$ to be even–contact is a first order relation itself. In this sense, the PDR we are considering is (a priori) more complicated than the PDRs describing contact structures and even–contact structures.

Remark 1.7. C^0–closeness of two Engel structures means that we have absolutely no control of the corresponding even–contact structures. C^1–closeness implies that the even–contact structures are C^0–close, but their kernels might be quite different.

In any case, we are interested in decoupling the PDR. The statements in Proposition 1.2 imply that the formal analogue of an Engel structure is as follows:

Definition 1.8. *Let M be a 4–manifold. A complete flag $\mathcal{W} \subset \mathcal{D} \subset \mathcal{E} \subset TM$ endowed with bundle isomorphisms as in Equations (1.1) and (1.2) is said to be a* ***formal Engel structure****.*

Often, we will write the formal Engel structure as a triple $(\mathcal{W}, \mathcal{D}, \mathcal{E})$, and we will omit the bundle isomorphisms. We will use terms like **Engel homotopy** and **Engel deformation** to refer to homotopies and deformations through Engel structures. This will contrast with **formal homotopies and deformations**, which are through formal Engel structures.

The question now, of course, is what is the nature of the inclusion of the space of Engel structures $\mathfrak{Engel}(M)$ into the space of formal Engel structures $\mathcal{F}\mathfrak{Engel}(M)$. Before we go any further, we should discuss what topology we take in each space. For $\mathfrak{Engel}(M)$, we are considering smooth Engel structures and, in order to deal with the PDR, a reasonable choice would be the C^k–topology with $k \in \{2, \cdots, \infty\}$. Turns out that all these choices yield homotopically equivalent spaces (see [8] for similar results), so we shall stick to C^2.

Regarding $\mathcal{F}\mathfrak{Engel}(M)$, we will restrict to smooth structures as well, but we will endow it with the C^0–topology. From the point of view of the h–principle, this is the natural choice, since the inclusion

$$\mathfrak{Engel}(M) \to \mathcal{F}\mathfrak{Engel}(M)$$
$$\mathcal{D} \to (\mathcal{W}, \mathcal{D}, \mathcal{E})$$

is then a continuous map. The formal Engel structure $(\mathcal{W}, \mathcal{D}, \mathcal{E})$ associated to $\mathcal{D}$ is called the **Engel flag** of $\mathcal{D}$.

The main theorem of Chapter 3, first proven in [10], reads:

Theorem 3.1. *Let M be a smooth 4–manifold. Then, the map*

$$\pi_k(i) : \pi_k(\mathfrak{Engel}(M)) \longrightarrow \pi_k(\mathcal{F}\mathfrak{Engel}(M))$$

is surjective for every $k \geq 0$. In particular, every formal Engel structure is homotopic to the flag of a genuine Engel structure.

Further, using a different but related approach, we shall prove the following result in Chapter 4:

Theorem 4.22. *There exists a subspace $\mathfrak{Engel}_{\mathrm{OT}}(M)$ of $\mathfrak{Engel}(M)$ such that the inclusion*

$$\mathfrak{Engel}_{\mathrm{OT}}(M) \to \mathcal{F}\mathfrak{Engel}(M)$$

is an isomorphism in π_0.

Additionally, if $\mathfrak{Engel}_{\mathrm{OT}}(M)$ contains an Engel structure, it contains its path connected component.

This result states that an overtwisted class of Engel structures exists (although it might as well be that every Engel structure is overtwisted, we do not know). In Chapter 4 we will discuss what the main issue is with extending this result to a homotopy equivalence.

Remark 1.9. Knowing that the existence h–principle holds, Theorem 4.22 is in some sense a consequence of the axiom of choice. However, the construction leading to the result designates certain connected components as overtwisted based on a very geometric criterion. Additionally, we are able to provide classification results for higher dimensional families (although not quite an h–principle).

1.2 The Engel condition expressed locally

In this section we will introduce a convenient method for describing Engel structures locally. The contact condition in dimension 3 is often stated by saying that the contact planes are turning with respect to a foliation by legendrian lines. Similarly, the Engel condition is easily understood in a flowbox of a line field contained in the Engel structure.

Consider a (possibly open) parallelisable 3–manifold N and take the product manifold $N \times [0,1]$ with coordinates (p,t). We are interested in those 2–distributions $\mathcal{D}$ of the form $\langle \partial_t \rangle \oplus \mathcal{X}$, with $\mathcal{X}$ a line field tangent to the slices $N \times \{t\}$. Consider the bundle isomorphism:

$$d_{(p,t)}\pi : T_{(p,t)}(N \times \{t\}) \to T_pN$$

where $\pi : N \times [0,1] \to N$ is the obvious projection. It is clear that a trivialisation of TN lifts to a trivialisation of the horizontal tangent space $\coprod_{t\in[0,1]} T(N \times \{t\})$. In particular, a trivialisation of TN induces an identification of the projectivised tangent bundle $\mathbb{P}TN$ with $\mathbb{RP}^2$, which lifts to an identification

$$\mathbb{P}T_{(p,t)}(N \times \{t\}) \cong \mathbb{RP}^2$$

which we still denote by $d_{(p,t)}\pi$.

For $p \in N$ fixed, consider the segment $\{p\} \times [0,1]$. The line field $\mathcal{X}$ can be regarded as an N–family of curves

$$\mathcal{X}_p : [0,1] \to \mathbb{RP}^2,$$

$$\mathcal{X}_p(t) = d_{(p,t)}\pi([\mathcal{X}(p,t)]),$$

where the brackets mean taking the line in $T_{(p,t)}N \times \{t\}$ spanned by $\mathcal{X}(p,t)$.

Remark 1.10. The analogous construction for a 3–dimensional contact structure gives a 2–dimensional family of curves that map into $\mathbb{RP}^1$. The contact condition amounts to these curves being immersed (that is, they turn consistently).

Alternatively, instead of the line field $\mathcal{X}$, we could have considered a vector field X spanning it. This would have produced an N–family of curves

$$X_p : [0,1] \to \mathbb{S}^2,$$

$$X_p(t) = d_{(p,t)}\pi([X(p,t)]),$$

where now we are considering the space of *oriented* lines. To keep track of orientations, this approach is often more useful, so we will phrase all further results in terms of X and not $\mathcal{X}$. Additionally, observe that the following identities hold:

$$X_p', X_p'' : [0,1] \to \mathbb{S}^2,$$

$$X_p'(t) = d_{(p,t)}\pi([[\partial_t, X(p,t)]]), \quad X_p''(t) = d_{(p,t)}\pi([[\partial_t, [\partial_t, X(p,t)]]]).$$

Note that the vector fields $X' = [\partial_t, X]$ and $X'' = [\partial_t, [\partial_t, X]]$ are indeed tangent to the slices $N \times \{t\}$.

In Chapter 3 we will review in detail the theory of curves in $\mathbb{S}^2$, but now we need the following concept:

Definition 1.11. *Let $\gamma : [0,1] \to \mathbb{S}^2$ be an immersion. A point $\gamma(t)$, $t \in [0,1]$, is said to be an **inflection point** if any of the following equivalent conditions hold:*

- *the geodesic curvature of γ (i.e. the curvature seen within $\mathbb{S}^2$) at time t vanishes,*
- *$\gamma(t)$ has a higher order tangency in $\gamma(t)$ with the great circle tangent to γ at $\gamma(t)$,*
- *as vectors in $\mathbb{R}^3$, the triple $\{X_p(t), X_p'(t), X_p''(t)\}$ is linearly dependent.*

Without any further delays, we can finally state how Engel structures are locally characterised:

Proposition 1.12 ([10])**.** *Consider the manifold $N \times [0,1]$, with N a parallelisable 3–manifold. Fix a trivialisation of TN. Consider a 2–distribution $\mathcal{D} = \langle \partial_t, X \rangle$ with X tangent to the foliation $\coprod_{t\in[0,1]} N \times \{t\}$ and $X_p : [0,1] \to \mathbb{S}^2$ described as above.*

Then, $\mathcal{D}$ is non–integrable at a point (p,t) if and only if X_p is immersed at time t.

Further, $\mathcal{D}$ is Engel at a point (p,t) if and only if, additionally, one of the following two conditions holds:

A. the curve X_p has no inflection point at time t,

B. the 2–distribution $q \to \langle X(q,t), X'(q,t)\rangle$ is a contact structure in $\mathcal{O}p(p) \times \{t\}$.

Proof. Non–integrability means that $\langle \partial_t, X, [\partial_t, X]\rangle$ is a 3–distribution. Since the last two vectors are horizontal, we have that this is equivalent to $\langle X, [\partial_t, X]\rangle$ being a plane in $T_{(p,t)}N \times \{t\}$. Phrased in terms of X_p, we are requiring for the vectors $X_p(t)$ and $X_p'(t)$ to be linearly independent, so the curve X_p is immersed at time t.

For Engelness the reasoning is the same. Suppose that we have $\mathcal{E} = [\mathcal{D}, \mathcal{D}] = \langle \partial_t, X, [\partial_t, X]\rangle$. Then, if the Engel condition is satisfied, either $[\partial_t, [\partial_t, X]]$ or $[X, [\partial_t, X]]$ are not contained in $\mathcal{E}$. The first case means that $X_p(t)$, $X_p'(t)$, and $X_p''(t)$ are linearly independent, so the curve X_p has no inflection point at time t. The second case means that $q \to \langle X_q, X_q'\rangle$ is not just a plane field, but actually a contact structure in $\mathcal{O}p(\{q\}) \times \{t\}$. □

Remark 1.13. It is immediate from the proof that if $\mathcal{D}$ is Engel, X_p has an inflection point at time t if and only if the vector field ∂_t is precisely the kernel of $\mathcal{D}$.

Property (A.) is easy to check because it only depends on the particular curve X_p. Additionally, since convexity is preserved by transformations in $\mathbb{P}\mathrm{GL}(2)$, a curve X_p will remain convex (or concave) if we change the trivialisation of TN. Condition (B.) is harder to work with: it does not depend on a particular curve, but on how they vary as p moves and t remains fixed. For this reason, we will usually choose framings of TN where condition (B.) is easily understood.

Remark 1.14. A couple of (trivial) observations are in order. The manifold N will often just be the 3–disc $\mathbb{D}^3$. Indeed, if we start with an Engel manifold $(M, \mathcal{D})$, and we have fixed some vector field $\mathcal{Y} \subset \mathcal{D}$, the description we just did applies in any flowbox for $\mathcal{Y}$. Additionally, the classical examples of Engel structures actually arise as $\mathbb{S}^1$–bundles over 3–manifolds, so the model of interest is $N \times \mathbb{S}^1$ (or some other $\mathbb{S}^1$–bundle). Proposition 1.12 works equally well when we take $\mathbb{S}^1$ instead of the interval.

1.2.1 Non–integrable distributions

A 2–distribution $\mathcal{D}$ is (everywhere) **non–integrable** if $[\mathcal{D}, \mathcal{D}]$ is a 3–distribution. We shall show that this condition, which certainly is necessary to be Engel, is actually much weaker. Proposition

1.12 can be used to show that non–integrable 2–distributions satisfy a complete h–principle if the dimension is at least 4.

Let M be a closed manifold of dimension $n > 3$. Denote the space of non–integrable distributions by $\mathfrak{Dist}_{\text{n.i.}}(M)$; its natural topology is the C^1–topology. Its formal analogue are pairs $\mathcal{D} \subset \mathcal{E}$ of distributions of rank 2 and 3 where $\mathcal{E}$ is oriented (in the non formal case, the orientation is given by Lie bracket, as in Lemma 1.4). Denote the space of formal non–integrable 2–distributions by $\mathcal{F}\mathfrak{Dist}_{\text{n.i.}}(M)$; we endow it with the C^0–topology.

If we additionally require for the 2–distribution $\mathcal{D}$ to contain a given line field $\mathcal{Y}$, we can define subspaces $\mathfrak{Dist}_{\text{n.i.}}(M, \mathcal{Y}) \subset \mathfrak{Dist}_{\text{n.i.}}(M)$ and $\mathcal{F}\mathfrak{Dist}_{\text{n.i.}}(M, \mathcal{Y}) \subset \mathcal{F}\mathfrak{Dist}_{\text{n.i.}}(M)$.

Theorem 1.15. *Let M be an n–manifold, $n > 3$. Then, there are full h–principles for the inclusions:*

$$\mathfrak{Dist}_{\text{n.i.}}(M) \to \mathcal{F}\mathfrak{Dist}_{\text{n.i.}}(M),$$
$$\mathfrak{Dist}_{\text{n.i.}}(M, \mathcal{Y}) \to \mathcal{F}\mathfrak{Dist}_{\text{n.i.}}(M, \mathcal{Y}).$$

This theorem follows in more generality from Gromov's convex integration. We are focusing on this particular case because: we shall need it later on, the proof is fairly simple, and some of the ideas relate to the arguments of Chapters 3 and 4.

Proof. The proof relies on the Smale–Hirsch theorem for immersions. Let K be a compact manifold, possibly with boundary. We think of K as the parameter space. Let $\mathcal{D} : K \to \mathcal{F}\mathfrak{Dist}_{\text{n.i.}}(M, \mathcal{Y})$ be a family of formal non-integrable 2–distributions containing $\mathcal{Y}$, such that $\mathcal{D}|_{\partial K}$ maps into $\mathfrak{Dist}_{\text{n.i.}}(M, \mathcal{Y})$. Proving the h–principle amounts to showing that the map $\mathcal{D}$ is homotopic, relative to ∂K, to a map into $\mathfrak{Dist}_{\text{n.i.}}(M, \mathcal{Y})$.

The line field $\mathcal{Y}$ can be lifted to the product manifold $M \times K$, which can then be covered by finitely many flowboxes $\{U_i\}$. Each flowbox is diffeomorphic to $\mathbb{D}^n \times \mathbb{D}^{\dim(K)} \times [0, 1]$, with coordinates (p, x, t); the line field $\mathcal{Y}$ corresponds to the last factor.

Proceed now inductively over the family $\{U_i\}$. In each flowbox, the 2–distribution can be expressed as $\mathcal{D} = \langle \partial_t, X\rangle$, where X is tangent to $\mathbb{D}^n \times \{(x, t)\}$. We deduce that $\mathcal{D}$ is non–integrable if and only if each curve $X_{(x,p)} = X(p, x, -) : [0, 1] \to \mathbb{S}^{n-1}$ is immersed; this follows as in Proposition 1.12.

But now the claim is immediate, because the Smale–Hirsch theorem, being a complete h–principle that is relative in the parameter and the domain, provides a homotopy of the curves $X_{(p,x)}$ to be honest immersions, relative to what we did in previous flowboxes. This shows that the inclusion

$$\mathfrak{Dist}_{\text{n.i.}}(M, \mathcal{Y}) \to \mathcal{F}\mathfrak{Dist}_{\text{n.i.}}(M, \mathcal{Y})$$

is a weak homotopy equivalence.

For the inclusion

$$\mathfrak{Dist}_{\text{n.i.}}(M) \to \mathcal{F}\mathfrak{Dist}_{\text{n.i.}}(M),$$

we have to proceed in the same way, but instead we cover $M \times K$ by flowboxes of *different* line fields (since we cannot assume for $\mathcal{D}$ to have a global section). □

Theorem 1.15 plays an important role later on. The idea is that, much like non–integrability is closely related to immersions of curves, Engelness is related to convex curves (see Property (A.) of Proposition 1.12). Convex curves do *not* satisfy the h–principle, but there is a good understanding of how they manifest flexibility [70]. The proofs of Theorem 3.1 and Theorem 4.22 follow a scheme similar to the one we just presented, and the main complication follows from the fact convex curves are simply not as flexible as immersions.

As a side note, one of the main ideas introduced in the proof of Theorem 1.15 is that of considering the product manifold $M \times K$ as the natural ambient space to work in.

1.3 Examples of Engel structures

1.3.1 The canonical structure in $\mathbb{R}^4$

Given a function $f : \mathbb{R} \to \mathbb{R}$, we can consider its k–th jet:

$$J^k(f) : \mathbb{R} \to \mathbb{R}^{k+2}$$

$$J^k(f)(x) = (x, f(x), f'(x), \cdots, f^{(k)}(x)).$$

That is, it packages the k–th order information of f. Jet spaces play an important role in the h–principle, since the act of decoupling a PDR amounts to considering maps

$$F : \mathbb{R} \to \mathbb{R}^{k+2}$$

whose components are independent functions that we think of as being the "formal derivatives" of one another in the obvious way. Take coordinates $(x, y_1, \cdots, y_{k+1})$ in $\mathbb{R}^{k+2}$. Then, the function $F = (x, F_1, \cdots, F_{k+1}) : \mathbb{R} \to \mathbb{R}^{k+2}$ is the $J^k(f)$ of some $f : \mathbb{R} \to \mathbb{R}$ if and only if $\partial_x F_i = F_{i+1}$. Equivalently, the image of F has to be tangent to the distribution:

$$\xi_k = \bigcap_{i=0,\cdots,k} \ker(dy_i - y_{i+1}dx).$$

If $k = 1$, the plane field ξ_1 is the standard contact structure in $\mathbb{R}^3$. If $k = 2$, $\mathcal{D}_{\text{std}} = \xi_2$ is the **standard Engel structure** in $\mathbb{R}^4$. For higher $k > 2$, these distributions are called the standard *Goursat structures.*

In the case of interest $k = 2$, an easy computation shows that the distribution $\mathcal{D}_{\text{std}}$ is indeed Engel. A convenient way to see this in the language of Proposition 1.12 is as follows. Instead of using the $(x; y_i)$ coordinates, let us take coordinates (x, y, z, w) in $\mathbb{R}^4$. Then:

$$\mathcal{D}_{\text{std}} = \ker(dy - zdx) \cap \ker(dz - wdx) = \langle \partial_w, (\partial_x + z\partial_y) + w\partial_z \rangle.$$

The vector fields $\partial_x + z\partial_y$ and ∂_z span, on each horizontal level $\mathbb{R}^3 \times \{w\}$, the contact structure $\xi = T(\mathbb{R}^3 \times \{w\}) \cap \ker(dy - zdx)$, which is w–invariant. We are thus in a setting where Property (B.) of Proposition 1.12 holds: $(\partial_x + z\partial_y) + w\partial_z$ is a vector field that is rotating within ξ. It is immediate that ∂_w spans the kernel $\mathcal{W}$, and the even contact structure is $\mathcal{E} = \ker(dy - zdx)$.

At this point, it is reasonable to state Engel's theorem. We already mentioned that Engel structures are topologically stable, which in particular means that they have a local model:

Proposition 1.16 (Engel)**.** *Let $(M, \mathcal{D})$ be an Engel 4–manifold and let $p, q \in M$. Then, there are neighbourhoods $\mathcal{O}p(\{p\}), \mathcal{O}p(\{q\}) \subset M$, and a diffeomorphism $\phi : \mathcal{O}p(\{p\}) \to \mathcal{O}p(\{q\})$ such that $\phi^*\mathcal{D} = \mathcal{D}$.*

In particular, any small ball in an Engel manifold can be identified with a neighbourhood of the origin in $(\mathbb{R}^4, \mathcal{D}_{\text{std}})$. In Chapter 5, we will study immersions of $\mathbb{S}^1$ or the interval that are tangent to an Engel structure (as explained in the Preamble), and we will keep coming back to this local model.

1.3.1.1 An alternate parametrisation of the standard structure

Consider now the following 2–distribution in $\mathbb{R}^4$:

$$\mathcal{D}_{\text{Lorentz}} = \ker(dy - wdx) \cap \ker(dz - w^2dx) = \langle \partial_w, X = \partial_x + w\partial_y + w^2\partial_z \rangle.$$

It is clearly an Engel structure, since the vector field $X/|X|$ is describing a convex curve in $\mathbb{S}^2$ as w moves. To see this, take an affine chart whose image is the hemisphere $\{x > 0\}$: the curve is precisely the conic $z^2 = y$.

The even contact structure $\mathcal{E}_{\text{Lorentz}} = [\mathcal{D}_{\text{Lorentz}}, \mathcal{D}_{\text{Lorentz}}]$ is readily seen to be:

$$\mathcal{E}_{\text{Lorentz}} = \langle \partial_w, X, X' = \partial_y + 2w\partial_z \rangle,$$

which readily implies that the kernel is spanned by the vector field X. We can fix the slice $\{x = 0\}$ and use the flow of X to obtain a change of coordinates where the kernel is exactly ∂_x. Indeed, doing this yields a diffeomorphism of $\mathbb{R}^4$:

$$\Psi : \mathbb{R}^4 \to \mathbb{R}^4$$

$$\Psi(x, y, z, w) = (x, y + wx, z + w^2x, w)$$

It is easy to check that it satisfies:

$$\begin{aligned}\Psi^*(\mathcal{D}_{\text{Lorentz}}) &= \ker(dy + xdw) \cap \ker(dz + 2wxdw) = \ker(dy + xdw) \cap \ker(dz - 2wdy) \\ &= \langle \partial_x, \partial_w - x(\partial_y + 2w\partial_z) \rangle.\end{aligned}$$

This is not quite $\mathcal{D}_{\text{std}}$, but it is close: now ∂_x spans the kernel, and the complementary vector is indeed turning inside the contact structure $\ker(dz - 2wdy)$. We can define another diffeomorphism of $\mathbb{R}^4$:

$$\Phi : \mathbb{R}^4 \to \mathbb{R}^4$$

$$\Phi(x, y, z, w) = (w/2, z/2, y - xz, -x)$$

and indeed:

$$(\Psi \circ \Phi)^*\mathcal{D}_{\text{Lorentz}} = \ker(dz - wdx) \cap \ker(dy - zdx) = \mathcal{D}_{\text{std}}.$$

For some purposes (see Lemma 5.2), using $\mathcal{D}_{\text{Lorentz}}$ instead of $\mathcal{D}_{\text{std}}$ is more convenient.

1.3.2 Cartan prolongations

We are now ready to explain two classic constructions of Engel manifolds due to Cartan (although a more modern treatment can be found in [60], which is a very recommendable read). Both of them can be understood within the framework of Proposition 1.12. We start by the better known one.

Example 1.17. Let (N, ξ) be a contact 3–manifold. The total space of the $\mathbb{S}^1$–bundle $\pi : \mathbb{S}(\xi) \to N$ carries an Engel structure given by the universal family construction, called the **(oriented) Cartan prolongation**. Let us elaborate. Points in $\mathbb{S}(\xi)$ are pairs (p, L) with $p \in N$ and L an oriented line in ξ_p, then, we can define:

$$\mathcal{D}_{(p,L)} = d_{(p,L)}\pi^{-1}(L).$$

Consider now a disc $\mathbb{D}^3 \subset N$, so that we can give a more explicit expression in coordinates (sorry). Since $\xi|_{\mathbb{D}^3}$ is the trivial bundle, we can find a legendrian framing $\{Y, Z\}$. Then, the restriction of the bundle $\mathbb{S}(\xi)$ to $\mathbb{D}^3$ is naturally identified with $\mathbb{D}^3 \times \mathbb{S}^1$ using this framing. That is, the point (p, L) gets identified with (p, e^{ti}), where the oriented line L is spanned by $\cos(t)Y + \sin(t)Z$. Then:

$$\mathcal{D}_{(p,L)} = \langle \partial_t, \cos(t)Y + \sin(t)Z \rangle$$

which is exactly as in Property (B.) of Proposition 1.12. It follows that the fibre direction spans the kernel and the even–contact structure is simply the lift of the contact structure in N. □

Of course, this construction can be carried out analogously for the space of unoriented lines $\mathbb{P}(\xi)$. This would define the **non–orientable Cartan prolongation** of (N, ξ). Sometimes we call it the Cartan prolongation of one *projective* turn, as well.

1.3.3 Lorentz prolongations

The following construction, also due to Cartan, is actually quite similar to the previous one, but it received less attention historically, as far as the author knows. In [69] holomorphic Engel structures were described and it was shown that, possibly except for a few exotic cases, all of them fit into one of the two families we are describing.

Example 1.18. Let (N, g) be a lorentzian manifold of signature $(2, 1)$. At each point $p \in N$, one can consider the light–like cone $C_p \subset T_pN$. Let $\pi : \mathbb{S}(g) \to N$ be the total space of the $\mathbb{S}^1$–bundle given by quotienting the cone C by the $(\mathbb{R} \setminus \{0\})$–action of rescaling. Points in $\mathbb{S}(g)$ are pairs (p, L), where L is an unoriented line in the cone C_p. Then, much like before, there is a canonical Engel structure

$$\mathcal{D}(p, L) = d_{(p,L)}\pi^{-1}(L),$$

which we call the **Lorentz prolongation**.

Now, over a disc $\mathbb{D}^3 \subset N$, we can find a orthonormal framing $\{V, Y, Z\}$ for g with Y and Z space–like, and V time–like. Then, in coordinates, we have that $(p, L = [V + \cos(t)Y + \sin(t)Z])$ gets identified with (p, e^{ti}) and the Engel structure is explicitly:

$$\mathcal{D}_{(p,L)} = \langle \partial_t, V + \cos(t)Y + \sin(t)Z \rangle.$$

Now Property (A.) of Proposition 1.12 holds at every point. It is clear that $\mathcal{W}$ is everywhere transverse to $\langle \partial_t \rangle$. □

Up to homotopy, there is a well–defined plane associated to each lorentzian metric. Namely, any plane that is space–like for the metric. Sometimes, that plane can be taken to be a contact structure ξ. Imagine then the following process: we can flatten the light–cone C until it agrees with the plane ξ. We can follow this flattening process by a family of Lorentz prolongations since, at every step except at the limit, we still have a cone. The limit yields precisely the oriented Cartan prolongation of ξ. This manner of thinking, where (oriented) Cartan prolongations are understood as limits of Lorentz prolongations, will be revisited in Chapter 2.

1.3.4 Open Engel manifolds

In the Preamble, when we reviewed the history of contact structures, we pointed out that Gromov, in his thesis, had completely solved the classification problem in open manifolds. His work did not apply just to contact structures, but to a more general class of PDRs: any open Diff–invariant relation automatically satisfies the (complete) h–principle in an open manifold.

We know what open means: formal solutions conform an open subset of the space of all possible sections. We also introduced already the idea of Diff–invariance: the definition of the PDR does not depend on the particular choice of chart, so it is invariant under diffeomorphisms. It is immediate that most of the relations we have considered are of this type and, indeed, Engel structures are a particular example.

For completeness we can just state:

Proposition 1.19 (Gromov). *Let M be an open manifold. Then, Engel structures in M satisfy the complete h–principle.*

This settles the classification problem in open manifolds.

1.3.5 Geiges' construction for mapping tori

Going back to closed manifolds, we can study the case of parallelisable mapping tori. The following construction is due to Geiges [29], and is very closely related to the notion of Lorentz prolongation.

Indeed, take a closed 3–manifold N and let $\phi : N \to N$ be a diffeomorphism whose mapping torus M_ϕ is parallelisable. Find an explicit embedding $N \times [0,1] \to M_\phi$ realising it as the mapping torus.

Then, there exists an almost–quaternionic structure in M_ϕ which allows us to find a parallelisation $\{\partial_t, A, B, C\}$, where t is the coordinate in the $[0,1]$ direction. Consider the following 2–distribution:
$$\mathcal{D}_\varepsilon = \langle \partial_t, A + \varepsilon(\cos(2\pi nt)B + \sin(2\pi nt)C)\rangle.$$
for ε non–negative and small and n some large integer.

From the perspective of the embedded $N \times [0,1]$, the line fields A, B, C are each given by a N–family of curves in $\mathbb{S}^2$, as in Proposition 1.12. Over any sufficiently short interval $I \subset [0,1]$, the line fields are almost constant; therefore, if n is large enough and $\varepsilon > 0$ is fixed, $[A + \varepsilon(\cos(2\pi nt)B + \sin(2\pi nt)C)]|_I$ is a curve obtained from $A|_I$ by performing a convex loop. Thus, for $\varepsilon > 0$ fixed and n sufficiently large, Proposition 1.12 implies that $\mathcal{D}_\varepsilon$ is Engel.

1.3.6 Vogel's theorem. Existence of Engel structures

The closed 4–manifolds admitting a complete flag (which is a necessary condition if they are to admit an Engel structure) conform a much larger class than just the $\mathbb{S}^1$–bundles and the mapping tori. However, until 2004, there were no known constructions in manifolds other than other than these. The following theorem, due to Vogel [82], constitutes a landmark in the history of Engel geometry:

Proposition 1.20 (Vogel). *Let M be a parallelisable manifold. Then, M admits an Engel structure.*

Vogel's result highlights the interplay between Engel and contact structures. We shall not go into detail about its proof, because certainly we wouldn't be able to do it justice. Instead, we encourage the reader to look at [82].

However, to give the reader some idea about what the proof entails, we shall give a rough sketch. First, M, being parallelisable, admits a round handle decomposition (i.e. a decomposition into manifolds of the form $\mathbb{S}^1 \times \mathbb{D}^j \times D^{3-j}$). We then want to proceed handle by handle constructing the Engel structure, until we get to the final handle and we are able to close up. At every step, one ensures that the boundary of the resulting open 4–manifold is transverse to $\mathcal{W}$, so that it is endowed with a contact structure.

At this point, the proof boils down to being able to keep going at every step and, to achieve this, one has to use all the tools available in contact topology to keep control of the contact structure in the boundary and its corresponding line field. Some of the flexibility is provided by making the contact structure be overtwisted.

The approach in Proposition 1.20 is markedly different from the one developed in this thesis. Vogel's method, as we just explained, relies mostly on contact topology and uses a particular topological decomposition of M to explicitly construct the desired Engel structure. The route we take follows the line of a "standard" h–principle and is not very explicit: the resulting structure is very wild everywhere. Papers like [79, 19] are much closer to the methods in this thesis.

1.4 Two remarks about the global topology of Engel structures

1.4.1 (The lack of) Gray stability

We know that Engel structures are *stable* locally because there is a unique local model (Proposition 1.16). One then may pose the question of whether this is true globally: is a small perturbation of an Engel structure diffeomorphic to it? To put this into context, it is worth recalling that this does hold for contact structures; this is called *Gray stability*.

One can see that this is not the case. Indeed, any Engel structure has an associated line field, its kernel $\mathcal{W}$, and we know that global stability does not hold for line fields. Now, it is not true that any perturbation of $\mathcal{W}$ (as a line field) arises as a perturbation of $\mathcal{D}$, but we can still say something if we understand a bit better the space of deformations of $\mathcal{W}$. Most of what will be explained now is better explained in [60], but we shall sketch part of it because it is implicitly assumed later on when we discuss flexibility for prolongations (see Subsubsection 4.2.2.1).

In any case. Let (N, ξ) be a closed contact 3–manifold with ξ trivial as a bundle. Then, the ambient manifold of its Cartan prolongation $(\mathbb{S}(\xi), \mathcal{D})$ is readily identified with $N \times \mathbb{S}^1$ if we choose a framing $\{Y, Z\}$ for ξ. In particular, we can think of $N \times \mathbb{S}^1$ as the mapping torus of the identity. The $\mathbb{S}^1$ direction is spanning the kernel $\mathcal{W}$.

Suppose we now perturb $\mathcal{W}$ slightly, so that its Poincaré return map is some diffeomorphism $\phi : N \to N$. We want to find a perturbation $\mathcal{D}'$ of $\mathcal{D}$ inducing the given perturbation $\mathcal{W}'$ of $\mathcal{W}$, but we run into a problem. Since the even–contact structure $\mathcal{E}' = [\mathcal{D}', \mathcal{D}']$ is a perturbation of $\mathcal{E}$, the contact structure ξ' induced on every level is a perturbation of ξ. This implies that the ϕ must be a contactomorphism of ξ'. In particular, ϕ is constrained to live in the set of diffeomorphisms that are conjugate to a contactomorphism.

Conversely, given some contactomorphism ϕ of ξ, its mapping torus M_ϕ can be endowed with many Engel structures. The model to keep in mind is the one described in Proposition 1.12: we construct a family of legendrian line fields X_t in Nsatisfying $\xi = \langle X_t, X_t' \rangle$ and $\phi^* Y_0 = Y_1$.

Then, the structure $\mathcal{D}' = \langle \partial_t, X_t \rangle$ is an Engel structure. Even though there are many (even non homotopic) choices, if ϕ is chosen C^∞–close to the identity, $\mathcal{D}'$ can be assumed to be C^∞–close to $\mathcal{D}$.

All of this allows us to prove the claim that an Engel analogue of Gray stability cannot possibly hold. The author arrived to this statement and proof thanks to discussions with N. Pia, to whom he is grateful. However, it is probably well–known to experts and, indeed, a related result can be found in [60][Theorem 6]:

Proposition 1.21. *Let (N, ξ) be a closed contact 3–manifold with ξ the trivial bundle. Let $(\mathbb{S}(\xi), \mathcal{D})$ be its (oriented) Cartan prolongation. Then, there are Engel structures C^∞–close to $\mathcal{D}$ that have no closed orbits of $\mathcal{W}$.*

Proof. Find a contact form for ξ such that the Reeb flow is generic. This implies that it has only countably many closed orbits. In particular, the set of periods of the Reeb flow is a countable subset of $\mathbb{R}^+$, and so is the set A of numbers having an integer multiple who is a period. Then, taking ϕ the time T Reeb flow with T not in A, yields, by the discussion above, a perturbation $\mathcal{D}'$ of $\mathcal{D}$ whose kernel has no closed orbits. Since T can be taken arbitrarily close to 0, $\mathcal{D}'$ can be assumed to be C^∞–close to $\mathcal{D}$. The Engel structures produced this way cannot be diffeomorphic to $\mathcal{D}$. □

The lack of Gray stability marks a big difference with contact structures. Its main consequence is that, if we are hoping to understand the space of Engel structures better, we are forced to necessarily construct invariants that remain well–behaved under Engel deformations. Our characterisation of looseness/overtwistedness has this built–in artificially, as we shall see (Definition 4.11).

1.4.2 The group of Engel symmetries

Another difference with contact structures is that the group of symmetries of an Engel structure is generically very small. In [60], this group is studied in the case of Cartan prolongations, where it is shown that each isometry is simply a lift of a contactomorphism of the base in the obvious fashion. This lift is defined as follows. If $\phi : (N, \xi) \to (N, \xi)$ is a contactomorphism, we set:

$$\mathbb{S}(\phi) : (\mathbb{S}(\xi), \mathcal{D}) \to (\mathbb{S}(\xi), \mathcal{D})$$

$$\mathbb{S}(\phi)(p, L) = (\phi(p), \phi_* L).$$

However, prolongations are very special Engel structures. In the same article, Montgomery shows that, after a perturbation akin to the one described in Proposition 1.21, the group of symmetries is much smaller.

In line with these results, we shall show that there are no *compactly supported* symmetries in a Darboux ball:

Proposition 1.22. *Let $\mathbb{R}^4$ be endowed with the standard Engel structure $\mathcal{D} = \ker(dy - zdx) \cap \ker(dz - wdx)$. Then, there are no compactly supported diffeomorphisms of $\mathbb{R}^4$ such that $\phi^*\mathcal{D} = \mathcal{D}$.*

Proof. Assuming that ϕ preserves the Engel structure implies that it preserves $\mathcal{W}$ and $\mathcal{E}$ as well. In particular, since it is compactly supported, it has to take each line $\{(x, y, z)\} \times \mathbb{R}$ to itself (not necessarily pointwise, a priori). We deduce that $\phi(\mathbb{R}^3 \times \{w\})$ must be graphical over the slice $\mathbb{R}^3 \times \{w\}$. We write:

$$\phi(x, y, z, w) = (x, y, z, f(x, y, z, w))$$

and we compute:

$$\phi^*(dz - wdx) = dz - fdx$$

$$\phi^*\mathcal{D} = \langle \partial_w, (\partial_x + z\partial_y) + f\partial_z.$$

We deduce that $f(x, y, z, w) = w$. □

1.5 Horizontal curves in Engel manifolds

Going back to our original motivation for studying bracket–generating distributions, we will now review some of the basics about curves tangent to Engel structures. In the literature these are often called *horizontal curves*, and we will follow this naming convention.

Before we get into the definitions, we should should be honest with the reader and admit that our interest in horizontal curves at this point has nothing to do with the motivation we presented in the Preamble; they are rather seen as a tool to study Engel structures themselves. Ideally, one would be able to show that some property related to the horizontal curves in the manifold persists under Engel deformation. This would then allow to detect whether two Engel structures are not Engel homotopic to one another even if they are formally homotopic.

1.5.1 Definitions

Definition 1.23. *Let $(M, \mathcal{D})$ be an Engel manifold. An immersion $\gamma : \mathbb{S}^1 \to M$ is said to be* ***horizontal*** *if $\gamma'(t) \in \mathcal{D}_{\gamma(t)}$ for all $t \in \mathbb{S}^1$. When γ is an embedding, we say that it is an* ***Engel knot*** *or a* ***horizontal knot****.*

Observe that, since $\mathcal{D}$ is a non–integrable 2–distribution, higher dimensional manifolds cannot be tangent to $\mathcal{D}$.

Following the h–principle philosophy, one can understand an immersion as two separate maps, the map itself and its derivative, that are coupled together. Decoupling this relation yields the following definition:

Definition 1.24. *Let $(M, \mathcal{D})$ be an Engel manifold. A* ***formal horizontal immersion*** *is a pair (γ, F) satisfying:*

1. *$\gamma : \mathbb{S}^1 \to M$ is a smooth map,*
2. *$F : T\mathbb{S}^1 \to \gamma^*\mathcal{D}$ is a monomorphism.*

Now. We shall write $\mathcal{I}(I, M)$ for the space of immersions of the 1–dimensional manifold I into M; its natural topology is the C^1–topology, but all the C^k, $k \geq 1$, are essentially equivalent. Then, the subspace of interest is the space of horizontal immersions of the Engel structure $\mathcal{D}$, which we call $\mathcal{HI}(I, \mathcal{D})$ (the manifold M is omitted from the notation).

Similarly, we will write $\mathcal{FI}(I,M)$ for the formal analogue of $\mathcal{I}(I,M)$, endowed with the C^0–topology. Then, $\mathcal{FHI}(I,\mathcal{D})$, the space of formal horizontal immersions, is a subspace. The following inclusions are continuous maps:

$$\mathcal{I}(I,M) \to \mathcal{FI}(I,M)$$
$$\mathcal{HI}(I,\mathcal{D}) \to \mathcal{FHI}(I,\mathcal{D}).$$

Exercise 1.25. Since we are interested in the nature of the inclusion $\mathcal{HI}(I,\mathcal{D}) \to \mathcal{FHI}(I,\mathcal{D})$, it is convenient to compute what the homotopy groups of the latter space are (particularly because they are fairly easy to obtain). The case where I is the interval is uninteresting, because all the spaces involved are contractible, so assume $I = \mathbb{S}^1$.

We start by noting that there is a locally trivial fibration

$$\mathrm{Mon}(T\mathbb{S}^1, \gamma^*\mathcal{D}) \to \mathcal{FHI}(\mathbb{S}^1,\mathcal{D}) \to \mathcal{M}\mathrm{aps}(\mathbb{S}^1, M)$$

where the projection map is simply the forgetful map that takes the formal immersion and outputs the map into M. Let $\gamma_k \in \mathcal{FHI}(I,\mathcal{D})$, $k \in K$, be a compact family. Then, the bundle $\gamma_k^*\mathcal{D}$, understood as a plane bundle over $K \times \mathbb{S}^1$, is either the trivial bundle or the pullback of the non–orientable plane bundle over $\mathbb{S}^1$. This follows because $\mathcal{D}$ contains the line bundle $\mathcal{W}$. Whether one thing or the other holds depends only on the homology class of the γ_k in $H_1(M,\mathbb{Z}_2)$. From this we deduce that the fibres over a given connected component are all isomorphic and, further, the fibration is just a product.

If $\gamma \in \mathcal{M}\mathrm{aps}(\mathbb{S}^1, M)$ satisfies that $\gamma^*\mathcal{D}$ is trivial, the fibres over the connected component containing γ are all homotopically equivalent to $\mathcal{M}\mathrm{aps}(\mathbb{S}^1,\mathbb{S}^1)$. Then, the homotopy groups of the fibre are:

$$\pi_0(\mathcal{M}\mathrm{aps}(\mathbb{S}^1,\mathbb{S}^1)) = \mathbb{Z}, \quad \pi_1(\mathcal{M}\mathrm{aps}(\mathbb{S}^1,\mathbb{S}^1)) = \mathbb{Z}, \quad \pi_j(\mathcal{M}\mathrm{aps}(\mathbb{S}^1,\mathbb{S}^1)) = 0 \text{ for } j > 1.$$

Otherwise, if $\gamma^*\mathcal{D}$ is the twisted bundle, the fibre has $\mathbb{Z}$ contractible components.

Denote by $\mathcal{M}\mathrm{aps}_t(\mathbb{S}^1, M)$ the subspace of maps such that the restriction of $\mathcal{D}$ is the trivial bundle. Then:

$$\pi_0(\mathcal{FHI}(\mathbb{S}^1,\mathcal{D})) \cong \pi_0(\mathcal{M}\mathrm{aps}(\mathbb{S}^1, M)) \times \mathbb{Z}.$$

For $(\gamma, F) \in \mathcal{FHI}(\mathcal{D})$, the fundamental group based on (γ, F) is either

$$\pi_1(\mathcal{FHI}(\mathbb{S}^1,\mathcal{D}),(\gamma,F)) \cong \pi_1(\mathcal{M}\mathrm{aps}(\mathbb{S}^1, M),[\gamma]) \times \mathbb{Z}, \text{ or}$$
$$\pi_1(\mathcal{FHI}(\mathbb{S}^1,\mathcal{D}),(\gamma,F)) \cong \pi_1(\mathcal{M}\mathrm{aps}(\mathbb{S}^1, M),[\gamma]),$$

depending on the orientability of $\gamma^*\mathcal{D}$. For all the higher homotopy groups:

$$\pi_j(\mathcal{FHI}(\mathbb{S}^1,\mathcal{D}),(\gamma,F)) \cong \pi_j(\mathcal{M}\mathrm{aps}(\mathbb{S}^1, M),[\gamma]).$$

It is convenient to describe a bit more the $\mathbb{Z}$–valued formal invariants that just appeared. Let us focus in the case where the Engel flag is oriented. In this case, $\mathcal{D}$ has a unique framing up to homotopy: it is given by $\mathcal{W}$ and the orientation of $\mathcal{D}$. Then, the formal invariant associated to a formal horizontal immersion (γ, F) simply counts the winding of F with respect to this framing. This number we call the **rotation number**; it does not depend on the parametrisation.

Similarly, if we are given a loop of formal horizontal immersions (γ_s, F_s), $s \in \mathbb{S}^1$, we can instead compute, for fixed $t \in \mathbb{S}^1$, the winding of $s \to F_s(t)$ with respect to the framing of $\mathcal{D}$. This is the formal invariant in π_1, which we call the **looping number**. It is clear that different parametrisations of the curves in the family make the looping number vary; if we consider unparametrised families, it is well–defined modulo the rotation number. □

We shall see that the homotopy type of $\mathcal{HI}(\mathbb{S}^1,\mathcal{D})$ is not quite the same as $\mathcal{FHI}(\mathbb{S}^1,\mathcal{D})$. In general, it contains additional connected components that correspond to *rigid* curves [7]. However, the set of rigid curves, up to parametrisation, will be a space of finite Haussdorf dimension, whereas $\mathcal{HI}(\mathbb{S}^1,\mathcal{D})$ is not; in this sense, the failure of the h–principle is due to a very small collection of curves. All these notions will be explored in Chapter 5.

1.5.2 Projections

When one studies *legendrian* knots in standard contact $(\mathbb{R}^3,\xi_{\text{std}})$, projections make one's life easier (as they usually do when it comes to knots). Of course, due to the presence of the contact structure, certain projections are better than others, and the ones considered are called the *front* projection and the *lagrangian* projection. The reader should refer to [27][Chapter 3].

For Engel manifolds it is quite similar. Let us consider $\mathbb{R}^4$ with coordinates (x,y,z,w) and endowed with the standard Engel structure $\mathcal{D}_{\text{std}} = \ker(dy - zdx) \cap \ker(dz - wdx)$. We described how this structure is the tautological distribution given by identifying $\mathbb{R}^4$ with the jet space $J^2(\mathbb{R},\mathbb{R})$. As such, there are a couple of natural projections to $\mathbb{R}^3$ that are of interest:

$$\pi_{\text{Adachi}}, \pi_{\text{Geiges}} : J^2(\mathbb{R},\mathbb{R}) \to J^1(\mathbb{R},\mathbb{R}),$$

$$\pi_{\text{Adachi}}(x,y,z,w) = (x,y,z), \quad \pi_{\text{Geiges}}(x,y,z,w) = (x,z,w).$$

The first one can be thought as taken the function y, with formal derivatives z and w, and then forgetting the second derivative w, effectively landing us in the jet space $J^1(\mathbb{R},\mathbb{R})$. The projection π_{Geiges} instead forgets about the function y and projects to a jet space where z plays the role of the function and w is its derivative.

Fun fact 1.26. *Adachi [1] and Geiges [28] proved, independently, that the π_0 h–principle for horizontal knots holds in standard $\mathbb{R}^4$ (see Proposition 5.17). The approach of the former relies on understanding the projection π_{Adachi}, whereas the latter relies on using the projection π_{Geiges}.*

The two projections are meaningful maps $J^2(\mathbb{R},\mathbb{R}) \to J^1(\mathbb{R},\mathbb{R})$ and not just maps $\mathbb{R}^4 \to \mathbb{R}^3$. More precisely:

Lemma 1.27. *The map*

$$d\pi_{\text{Geiges}} : T_{(x,y,z,w)}(J^2(\mathbb{R},\mathbb{R})) \to T_{(x,z,w)}(J^1(\mathbb{R},\mathbb{R}))$$

induces a linear isomorphism between $\mathcal{D}_{\text{std}}(x,y,z,w)$ and $\xi_{\text{std}}(x,z,w)$.

The map

$$d\pi_{\text{Adachi}} : T_{(x,y,z,w)}(J^2(\mathbb{R},\mathbb{R})) \to T_{(x,y,z)}(J^1(\mathbb{R},\mathbb{R}))$$

maps $\mathcal{D}_{\text{std}}(x,y,z,w)$ into $\xi_{\text{std}}(x,y,z)$ with kernel $\langle \partial_w \rangle$.

Proof. For completeness recall that, respectively, ξ_{std} is given as $\ker(dz - wdx)$ or $\ker(dy - zdx)$. In both cases it is obvious that the projection must take $\mathcal{D}_{\text{std}}$ into ξ_{std}. The proof is just a computation. $\square$

The reader should think of π_{Geiges} as being the analogue of the *lagrangian* projection and π_{Adachi} as the analogue of the *front* projection.

1.5.2.1 Geiges projection of a horizontal immersion/knot

The projections are handy to deal with horizontal immersions and knots. In particular:

Lemma 1.28. *Let $\gamma : \mathbb{S}^1 \to (\mathbb{R}^4, \mathcal{D}_{\mathrm{std}})$ be a horizontal immersion. Then, $\pi_{\mathrm{Geiges}} \circ \gamma : \mathbb{S}^1 \to (\mathbb{R}^3, \xi_{\mathrm{std}})$ is a legendrian immersion that, additionally, satisfies the area constraint:*

$$\int_\gamma zdx = 0.$$

Proof. We know, by Lemma 1.27, that $d\pi_{\mathrm{Geiges}}$ identifies $\mathcal{D}_{\mathrm{std}}$ with ξ_{std}, so any horizontal immersion projects down to a legendrian immersion. The area constraint comes from the fact that γ is a closed curve with $\gamma' \in \ker(dy - zdx)$: the integral being zero implies that the projection indeed lifts to a closed curve. $\square$

Let $\pi_{\mathrm{front}} : (\mathbb{R}^3, \xi_{\mathrm{std}}) \to \mathbb{R}^2$ be the *front* projection. In the case of interest it is simply $(x, z, w) \to (x, z)$. From the lemma, we deduce that there is a correspondence between horizontal curves in $\mathbb{R}^4$ and curves with cusps in $\mathbb{R}^2$ that bound signed area zero (indeed, note that the integral $\int_\gamma zdx$ is, by Stokes' theorem, exactly the area bounded by the front projection).

Exactly the same thinking yields:

Lemma 1.29. *Let $\gamma : \mathbb{S}^1 \to (\mathbb{R}^4, \mathcal{D}_{\mathrm{std}})$ be a horizontal immersion. Let $t_0, t_1 \in \mathbb{S}^1$. Then, there is a self–intersection $\gamma(t_0) = \gamma(t_1)$ if and only if $\pi_{\mathrm{front}} \circ \pi_{\mathrm{Geiges}} \circ \gamma$ has a self–tangency at t_0, t_1 and:*

$$\int_{t_0}^{t_1} \gamma^*(zdx) = 0.$$

The integral condition is equivalent to the fact that the area bounded by $\pi_{\mathrm{front}} \circ \pi_{\mathrm{Geiges}} \circ \gamma|_{t_0}^{t_1}$ is zero.

Proof. Indeed, a self–intersection means that all the four coordinates (x, y, z, t) agree for $\gamma(t_0)$ and $\gamma(t_1)$. When we project to the plane (x, z), we must then see a self–intersection too. Since w has too match, and w is given as the derivative of z with respect to x, we further deduce that the curve has to be tangent to itself in the self–intersection. The y–coordinates agreeing is precisely the integral criterion. $\square$

This provides a convenient way of dealing with horizontal knots.

Considering the (front of) Adachi projection has its own advantages. Mostly, there are no computations of areas to keep track of (which is often convenient, since drawings of projections are not well suited for it). However, two different types of cusps appear (because the legendrians we obtain projecting the first time are already cuspidal). In any case, we will not make use of it in this thesis.

1.5.3 Closed characteristics of the kernel

In an Engel manifold, there are certain horizontal curves that are very special, since there is a preferred horizontal vector field: the kernel $\mathcal{W}$. Curves tangent to it (in particular, its closed orbits) display rigid behaviours, meaning that the space of possible deformations of such a curve is much smaller than for other horizontal curves. This will be explained in Section 5.1.

In this subsection, we will introduce the concept of developing map (which will play an important role later on when we discuss rigidity). Then, we will discuss the behaviour of $\mathcal{W}$ when $\mathcal{D}$ is assumed to be generic.

1.5.3.1 The developing map

Recall Proposition 1.12 or Example 1.17: the Engel condition, when the trivialising vector field W is tangent to $\mathcal{W}$, is seen as consistent turning of a legendrian vector field X within a contact structure ξ. This implied that, if we projectivise ξ and we follow an orbit of W, we obtain a map into $\mathbb{S}(\xi)$. This is what the *developing map* is. Let us formalise this a bit more.

Let $(M, \mathcal{D})$ be an Engel manifold and fix a point $p \in M$. Denote by ϕ the germ of an embedding of a disc transverse to $\mathcal{W}$ and passing through p at the origin. After suitable reparametrisation, $\phi^*\mathcal{E}$ and $\phi^*\mathcal{D}$ can be assumed to be the germs $\xi_0 = \ker(dy - zdx)$ and $\langle \partial_z \rangle$, respectively.

We can choose some orientation for γ, the orbit of $\mathcal{W}$ passing through p, and find, locally, a vector field $W \subset \mathcal{W}$ that is compatible with this orientation. This readily extends ϕ to a (germ of) immersion $\Psi_p : \mathcal{O}p(\{0\}) \times \mathbb{R} \to M$ using the equation

$$\Psi_p(x, y, z, t) = \varphi_{\phi(x,y,z)}(t),$$

where $\varphi_q(t)$ is the time t flow of W starting at the point q. Since flows tangent to $\mathcal{W}$ preserve $\mathcal{E}$, we have a projection $(\mathcal{O}p(\{0\}) \times \mathbb{R}, \mathcal{E}) \to \mathcal{O}p(\{0\})$ that maps $\mathcal{E}$ onto ξ_0 with kernel the vertical direction. Then:

Definition 1.30. *The map*

$$\begin{array}{rcl} \mathbb{P}(W,p) : \mathbb{R} & \to & \mathbb{P}(\xi_0) \\ t & \to & [X(\Psi_p(0,t))] \end{array}$$

is called the ***developing map*** *at p (with respect to W).*

Observe that the choice of ϕ plays no role in this definition but, for a particular identification of $\mathbb{P}(\xi_0)$ with $\mathbb{RP}^1$, we do need to choose a framing for ξ. The collection of times in which $[X(\Psi_p(0,t))] = [X(\Psi_p(0,0))]$, i.e. the times where a **projective turn** is completed, is well–defined regardless of this choice.

If a curve $\gamma : [0,1] \to M$ tangent to $\mathcal{W}$ is given, and a particular framing for the contact structure has been fixed, one can assign to γ its developing map

$$\mathbb{P}_\gamma : [0,1] \to \mathbb{RP}^1$$

given by the restriction of the developing map of $\gamma(0)$:

$$\mathbb{P}_\gamma(s) = \mathbb{P}(\gamma', \gamma(0))(\varphi^{-1}_{\gamma(0)}(\gamma(s))).$$

One can then compute the number of projective turns γ describes. If γ is a closed orbit, it is worth remarking that γ might not exactly describe an integer number of turns, since its Poincaré return map might not be the identity.

Exercise 1.31. Let us revisit the examples of Engel structures that we already know. It is immediate that all the $\mathcal{W}$-orbits in both $(\mathbb{R}^4, \mathcal{D}_{\mathrm{std}})$ and $(\mathbb{R}^4, \mathcal{D}_{\mathrm{Lorentz}})$ describe less that one projective turn. If $(\mathbb{S}(\xi), \mathcal{D})$ is an oriented Cartan prolongation, the non–multiply–covered orbits describe 2 projective turns; if the Cartan prolongation is non–orientable, they describe 1 turn. In general, it is hard to determine what $\mathcal{W}$ is doing.

Please refer to Chapter 5 for a more in–depth discussion of these ideas.

1.5.3.2 A Kupka–Smale–type theorem for $\mathcal{W}$

We finish our discussion about $\mathcal{W}$–orbits by studying how they behave if $\mathcal{D}$ is chosen generically. As we saw when discussing prolongations, the Poincaré return map of a $\mathcal{W}$–orbit is necessarily a contactomorphism. At a linear level, contactomorphisms fixing the origin but no other points, and having conformal factor different from 1 are generic. Exactly the same should hold for Poincaré maps of kernels of even–contact structures.

Proposition 1.32. *All the $\mathcal{W}$–orbits of a C^∞–generic even–contact structure are isolated and have Poincaré map that is not a strict contactomorphism. The same holds for a C^∞–generic Engel structure.*

Before we prove the proposition, we will need the following auxiliary lemma.

Lemma 1.33. *Let $(\mathbb{D}^3, \xi = \ker(\alpha))$ be the standard contact Darboux ball. Consider the space of maps from $\mathbb{D}^3$ to itself that are contact (but not necessarily bijective). Then, the subset of those having non–degenerate fixed points is open and dense.*

Proof. Consider the manifold $V = \mathbb{R}^3 \times \mathbb{R}^3 \times \mathbb{R}$, and let π_1 and π_2 be the projections onto its first and second factors, respectively. V can be endowed with the contact structure $\ker(\lambda = \pi_1^*\alpha - e^t\pi_2^*\alpha)$; any contact map $\phi : \mathbb{D}^3 \to \mathbb{D}^3$ lifts to a legendrian $\Gamma_\phi(x) = (x, \phi(x), \log[\alpha/(\phi^*\alpha)])$, where the last term accounts for the conformal factor of ϕ. We need for Γ_ϕ to be transverse to $\Delta \times \mathbb{R}$, with $\Delta \subset \mathbb{R}^3 \times \mathbb{R}^3$ the diagonal.

This follows from Thom's transversality (see, for instance, [20, p. 17, 2.3.2]). Indeed, let $p \in \Gamma_\phi \cap (\Delta \times \mathbb{R})$. Then, there is a neighbourhood $U \ni p$ contactomorphic to $J^1(\mathbb{R}^3, \mathbb{R})$ with $\Gamma_\phi \cap U$ being taken to the zero section. Then, Thom's transversality states that a generic C^∞–small deformation of $\Gamma_\phi \cap U$ (which is given as the graph of a function) is transverse to the submanifold $(\Delta \times \mathbb{R}) \cap U$. Proceeding chart by chart, capping the deformations off, and using progressively smaller deformations allows us to conclude. Since the deformations are C^∞-small, they are graphical over the first factor of V, and hence give rise to a contact map. $\square$

Now we can complete our proof:

Proof of Proposition 1.32. Fix a metric on M. For simplicity, focus on even–contact structures having orientable and oriented kernel. Any such $\mathcal{E}$ has an associated unitary vector field W spanning $\mathcal{W}$ positively. Otherwise, observe that the argument that follows can be applied by taking a double cover and proceeding in a $\mathbb{Z}_2$–equivariant fashion.

Consider the subset of even–contact structures such that the W–orbits of length at most $T > 0$ are non degenerate. We claim that it is open and dense. We claim that it is still open and dense if we further require for the Poincaré return maps of the orbits to be non–strict contactomorphisms. Assuming these statements, the subset of even–contact structures such that this is true for orbits of all periods is a countable intersection of open and dense sets.

Our claims readily follow from arguments of Peixoto [67][p. 219-220], which we briefly sketch. Take $(M, \mathcal{E})$. Given any W–orbit γ of period $\tau < T$, Lemma 1.33 produces a C^∞–small deformation of W such that the Poincaré return has only isolated fixed points. However, this might produce new orbits of period $2\tau - \varepsilon \geq N\tau < T$ for some integer N. Therefore, one starts deforming orbits that are close to the minimal period and introduces progressively smaller deformations as the period goes up to T. If we additionally want the orbits not to have return map a strict contactomorphism, we take the isolated orbits we have produced and we replace their Poincaré return maps by their linearised version, which we then make generic.

This concludes the proof for the statement regarding even–contact structures. We then note that any C^∞–perturbation of $\mathcal{E} = [\mathcal{D}, \mathcal{D}]$ can be realised by a C^∞–perturbation of $\mathcal{D}$, so we conclude that the same holds for a C^∞–generic Engel structure. □

Proposition 1.32 is consistent with Proposition 1.21, where we showed that there are perturbations of prolongations with no closed orbits.

1.5.3.3 An Engel structure with short open orbits

We will now construct an Engel structure whose kernel has an interesting behaviour. Let us consider the manifold T^4 with coordinates (x, y, z, w) (which we think as ranging from 0 to 1). Endow it with the following 2–distribution:

$$\mathcal{D}(x, y, z, w) = \langle \partial_w, W = \partial_x + \varepsilon(\cos(2\pi w)\partial_y + \sin(2\pi w)\partial_z)\rangle.$$

It is indeed Engel, because it is nothing but the Lorentz prolongation of the flat metric in T^3. What can be said about its $\mathcal{W}$–orbits?

A computation shows that the kernel is spanned by the vector field W. Consider the immersion $\mathbb{R} \times T^3 \to T^4$ given by:

$$\Psi(x, y, z, w) = (x, y + \varepsilon \cos(2\pi w)x, z + \varepsilon \sin(2\pi w)x, w)$$

which is given by fixing the slice $\{x = 0\}$ and then flowing using W. This implies that the developing map will be easy to read once we pull back $\mathcal{D}$ by Ψ. Indeed:

$$\Psi^*\mathcal{D} = \langle \partial_x, \partial_w + 2\pi\varepsilon x(\sin(2\pi w)\partial_y - \cos(2\pi w)\partial_z)\rangle$$

which, on every $\{x\} \times T^3$, induces the standard contact structure

$$\xi = \ker(\cos(2\pi w)dy + \sin(2\pi w)dz) = \langle \partial_w, L = \sin(2\pi w)\partial_y - \cos(2\pi w)\partial_z\rangle.$$

In particular, we immediately see that the developing map of every orbit describes almost one projective turn. Even though all orbits are open, their developing map is, in some sense, *short*.

A way of understanding this is as follows. Using this immersion, we see that the structure in T^4 can be understood as a mapping torus using the contactomorphism:

$$\psi(y, z, w) = (y - \varepsilon \cos(2\pi w), z - \varepsilon \sin(2\pi w), w).$$

This contactomorphism not only fixes ξ, but it fixes the vector field L. Effectively, the linearisation of ψ is performing a shear parallel to the line $\langle L\rangle$; as such, this line serves as a limit at infinity of the legendrian line field corresponding to the Engel structure. In [47] a similar setting is considered, where instead of a shear the map considered is hyperbolic.

Chapter 2

The spaces of Cartan and Lorentz prolongations

In this chapter we aim to describe the homotopy type of the spaces of Cartan and Lorentz prolongations. This is not a question about *Engel topology* as such, but rather a question about these more restrictive classes of structures. We shall see that the study of Cartan prolongations boils down to the study of contact structures and that flexibility phenomena in contact topology translate to flexibility in this setting. We will also see that the space of Lorentz prolongations can be described in simple algebraic topology terms.

Before we get more technical, let us make some preliminary observations to motivate what we want to do. In Subsections 1.3.2 and 1.3.3 we introduced Cartan and Lorentz prolongations. In the Cartan case we distinguished two families: the orientable and the non–orientable ones. Since they correspond, respectively, to considering oriented and unoriented lines, it is clear that the former case is obtained as a double cover of the latter.

It is clear then that one can obtain other Engel structures in $\mathbb{S}^1$–bundles by taking higher order covering maps. Additionally, one can make the gauge group act on these manifolds, providing further examples. In the Cartan case, the structures constructed using these two methods, which we still call prolongations, satisfy that the orbits of the kernel are the fibres of the bundle. In this chapter we will show that, conversely, any Engel structure satisfying this can be obtained using these two operations and then we will classify them up to homotopy.

Results similar to this (we will explain in which sense later) can be found in [46].

2.1 Spaces of prolongations and formal prolongations

For the rest of the chapter fix a closed, orientable 3–manifold N. Each oriented $\mathbb{S}^1$–bundle over N is given by its Euler class c; denote its total space by $N(c)$ and write $\pi : N(c) \to N$ for the projection. We shall focus on *oriented* Cartan prolongations of *oriented* contact structures. The orientable (but not oriented) case and the non–orientable case are left as exercises for the reader.

2.1.1 Cartan prolongations

Denote by $\mathcal{C}\text{-}\mathcal{S}\mathfrak{trs}$ the space of oriented contact structures on N. It naturally decomposes into several components $\mathcal{C}\text{-}\mathcal{S}\mathfrak{trs}(c)$ corresponding to contact structures having a particular Euler class $c \in H^2(N, \mathbb{Z})$. We can further denote $\mathcal{C}\text{-}\mathcal{S}\mathfrak{trs}(\xi)$ for the connected component containing the contact structure $\xi \in \mathcal{C}\text{-}\mathcal{S}\mathfrak{trs}$.

Suggestively, denote $\mathfrak{Cartan}(c)$ for the space of all oriented Engel structures on $N(c)$ having the fibre direction as their kernel. Any Engel structure $\mathcal{D} \in \mathfrak{Cartan}(c)$ defines a contact structure $\xi = d\pi(\mathcal{E})$ on N, since $\mathcal{W} = \ker(d\pi)$. Orient the line field $\mathcal{W}$ using the orientation of the fibre. Then, ξ inherits an orientation from $\mathcal{E}/\mathcal{W}$. Hence, there is a projection:

$$\mathfrak{Cartan}(c) \to \mathcal{C}\text{-}\mathcal{S}\mathfrak{trs}\,.$$

It is immediate that the Euler class of ξ must be of the form kc, with $k > 0$. This integer is called the **turning number** and is computed as follows. Take any $\mathbb{S}^1$–fibre of $N(c)$. Find some $\mathbb{S}^1$–invariant, positively oriented framing of $\mathcal{E}/\mathcal{W}$. Compute the degree of $\mathcal{D}/\mathcal{W}$ with respect to this framing. The resulting number k does not depend on the choices involved and is necessarily positive.

Denote by $\mathfrak{Cartan}(c, k) \subset \mathfrak{Cartan}(c)$ the space of Cartan prolongations having turning number k. Write $\mathfrak{Cartan}(c, k, \xi) \subset \mathfrak{Cartan}(c, k)$ for the subspace of those that additionally project down to $\xi \in \mathcal{C}\text{-}\mathcal{S}\mathfrak{trs}$. Observe that a path of Cartan prolongations projects down to a path of contact structures; write $\mathfrak{Cartan}(c, k, [\xi])$ for the subspace of those prolongations that lift contact structures homotopic to ξ.

Denote by $\mathcal{C}\mathrm{over}(c, k)$ the space of k–fold covers from $N(c)$ to $N(kc)$; i.e. positively oriented fibrewise submersions with k–sheets lifting the identity on N. Once we fix a bundle isomorphism between the sphere bundle of ξ and $N(kc)$, we can construct the following homeomorphism:

$$f : \mathfrak{Cartan}(c, k, \xi) \to \mathcal{C}\mathrm{over}(c, k)$$

$$f(\mathcal{D})(p, t) = (p, d_{(p,t)}\pi(\mathcal{D}(p, t))),$$

where we regard $d_{(p,t)}\pi(\mathcal{D}(p, t))$ as an oriented line (using the orientations of the fibre and $\mathcal{D}$). Note that $f(\mathcal{D})$ pulls back the canonical Cartan prolongation in $N(kc) \cong \mathbb{S}(\xi)$ to $\mathcal{D}$.

All the contact structures in a neighbourhood of ξ can be identified with ξ itself using a projection along a complementary line field. This implies that the corresponding sphere bundles can consistently be identified with $N(kc)$. This readily implies that

$$\mathfrak{Cartan}(c, k) \to \mathcal{C}\text{-}\mathcal{S}\mathfrak{trs}(kc) \quad \text{and} \quad \mathfrak{Cartan}(c, k, [\xi]) \to \mathcal{C}\text{-}\mathcal{S}\mathfrak{trs}(\xi)$$

are locally trivial fibrations with fibre $\mathcal{C}\mathrm{over}(c, k)$.

Our aim in this chapter is to understand the homotopy type of the spaces $\mathfrak{Cartan}(c, k, [\xi])$ using the fibration structure we have just presented.

2.1.2 Formal Cartan prolongations (of plane fields)

Just like Cartan prolongations are obtained from contact structures by projectivisation, we will now define projectivisations of arbitrary plane fields. The resulting structures will be non–integrable 2–distributions, and we will naturally regard the Cartan prolongations as a subclass. They conform a convenient framework to address the questions we have posed.

Denote by $\mathcal{P}\mathfrak{lanes}$ the space of oriented plane fields in N. Write $\mathcal{P}\mathfrak{lanes}(c)$ for those of Euler class $c \in H^2(N, \mathbb{Z})$ and $\mathcal{P}\mathfrak{lanes}(\xi)$ for the connected component containing the plane field $\xi \in \mathcal{P}\mathfrak{lanes}$. By fixing a parallelisation of N, $\mathcal{P}\mathfrak{lanes}$ can be readily identified with $\mathcal{M}\text{aps}(N, \mathbb{S}^2)$.

We write $\mathcal{F}\mathfrak{Cartan}(c)$ for the space of oriented 2–distributions in $N(c)$ that contain the fibre direction, are everywhere non–integrable (but not necessarily maximally), and whose induced 3–distribution obtained by Lie bracket is preserved by flows along the fibre. The elements in $\mathcal{F}\mathfrak{Cartan}(c)$ are called **formal Cartan prolongations**.

Let $\mathcal{D} \in \mathcal{F}\mathfrak{Cartan}(c)$. Then, by definition, $\xi = d\pi(\mathcal{E} = [\mathcal{D}, \mathcal{D}])$ is a plane field in N; ξ being contact amounts to $\mathcal{D}$ being an element in $\mathfrak{Cartan}(c)$. The orientation of $\mathcal{E}/\mathcal{W}$ orients ξ, just like in the case of Cartan prolongations. The turning number k can also be defined; we write $\mathcal{F}\mathfrak{Cartan}(c, k) \subset \mathcal{F}\mathfrak{Cartan}(c)$ for the subspace of those formal Cartan prolongations with turning number k: they necessarily project down to plane fields of Euler class kc. Similarly, write $\mathcal{F}\mathfrak{Cartan}(c, k, [\xi])$ for those lifting plane fields homotopic to ξ.

There are locally trivial fibrations:

$$\mathcal{C}\text{over}(c, k) \longrightarrow \mathcal{F}\mathfrak{Cartan}(c, k) \longrightarrow \mathcal{P}\mathfrak{lanes}(kc),$$

$$\mathcal{C}\text{over}(c, k) \longrightarrow \mathcal{F}\mathfrak{Cartan}(c, k, [\xi]) \longrightarrow \mathcal{P}\mathfrak{lanes}(\xi),$$

where $\mathcal{C}\text{over}(c, k)$ is defined as before.

2.1.3 Prolongations of rank–2 bundles

To take it one step further, we will now consider projectivisations of rank–2 bundles. The reason why we introduce them is that the fibration structure we discussed in the two previous cases simplifies in this setting (as we shall see in the next section), and serves as a useful intermediate step.

There is a natural inclusion of the space of oriented plane fields into the space of oriented rank 2 bundles:

$$\mathcal{M}\text{aps}(N, \mathbb{S}^2) \cong \mathcal{P}\mathfrak{lanes} \to \mathcal{B}\mathfrak{undles} \cong \mathcal{M}\text{aps}(N, \text{Gr}(2, \infty)),$$

where $\text{Gr}(2, \infty)$ is the infinite Grassmanian of oriented 2–planes, which is the Eilenberg–Maclane space $K(2, \mathbb{Z})$. We write $\mathcal{B}\mathfrak{undles}(c)$ for the subspace of bundles having Euler class $c \in H^2(N, \mathbb{Z})$. Let $\mathcal{V}(2, \infty)$ be the Stiefel manifold of ordered pairs of orthonormal vectors in $\mathbb{R}^\infty$; recall that there is a tautological fibration

$$\mathbb{S}^1 \to \mathcal{V}(2, \infty) \to \text{Gr}(2, \infty).$$

We will now explain what a prolongation is in this setting. We define $\mathcal{F}\mathfrak{Cartan}^\infty(c)$ to be the space of maps of $N(c)$ into $\mathcal{V}(2, \infty)$ which are lifts of maps $N \to \text{Gr}(2, \infty)$ and are fibrewise submersions respecting the orientation. This space has several components $\mathcal{F}\mathfrak{Cartan}^\infty(c, k)$ distinguished by the Euler class kc of the underlying 2–plane bundle, $k > 0$.

Let ξ be a oriented plane field of Euler class kc. The following diagram commutes:

$$\begin{array}{ccccc}
\mathcal{C}\mathrm{over}(c,k) & \longrightarrow & \mathfrak{Cartan}(c,k,[\xi]) & \longrightarrow & \mathcal{C}\text{-}\mathcal{S}\mathfrak{trs}(\xi) \\
\cong\downarrow & & \downarrow & & \downarrow \\
\mathcal{C}\mathrm{over}(c,k) & \longrightarrow & \mathcal{F}\mathfrak{Cartan}(c,k,[\xi]) & \longrightarrow & \mathcal{P}\mathfrak{lanes}(\xi) \\
\cong\downarrow & & \downarrow & & \downarrow \\
\mathcal{C}\mathrm{over}(c,k) & \longrightarrow & \mathcal{F}\mathfrak{Cartan}^\infty(c,k) & \longrightarrow & \mathcal{B}\mathfrak{undles}(kc)
\end{array}$$

where each row is a fibration.

2.1.4 Lorentz prolongations

Now we finally get to Lorentz prolongations. We shall see that their study boils down to the study of formal Cartan prolongations.

Each Lorentz prolongation arises from the light–like cone of a lorentzian metric. We will denote by $\mathfrak{Cones}$ the space of such cones (note that different metrics might yield the same cone; this is true, for instance, if they are related by a conformal factor). We will additionally assume that the space–like plane fields of an element in $\mathfrak{Cones}$ are orientable and we shall fix a preferred orientation (equivalently, we fix an orientation of the projectivised circle bundle of the cone).

We can write $\mathfrak{Lorentz}$ for the space of non–integrable, oriented 2–distributions in $N(c)$ that project down to a cone using the tangent map $d\pi : TN(c) \to TN$. Being non–integrable, the fact that they project to a cone immediately implies that they are Engel (as expected). The non–integrability orients the cones (or, identically, the space–like plane fields), and hence projecting yields elements in $\mathfrak{Cones}$. The space–like plane fields must have Euler class kc with k positive (and, again, we call this number the turning number).

Elements in $\mathcal{P}\mathfrak{lanes}$ can be slightly pushed to yield elements in $\mathfrak{Cones}$. This provides a homotopy inclusion of $\mathcal{P}\mathfrak{lanes}$ into $\mathfrak{Cones}$ that is a homotopy equivalence. This lifts to a homotopy equivalence between (each of the two connected components of) $\mathfrak{Lorentz}$ and $\mathcal{F}\mathfrak{Cartan}$. Let us formalise this.

Fix a metric g and an orientation in N, this allows us to regard the sphere bundle $\mathbb{S}TN$ as a submanifold of TN. Each element $\mathcal{D}$ in either $\mathcal{F}\mathfrak{Cartan}$ or $\mathfrak{Lorentz}$ has an associated tautological map

$$f : N(c) \to \mathbb{S}TN$$

$$f(p,t) = (p, d_{(p,t)}\pi(\mathcal{D}))$$

where we are regarding the right hand side as an oriented line. For those $\mathcal{D} \in \mathcal{F}\mathfrak{Cartan}$, f has image in the circle bundle $\mathbb{S}(\xi)$ of the underlying plane field ξ. Let ν be unique unitary vector field orthogonal to ξ such that $\xi \oplus \langle\nu\rangle = TN$ as oriented bundles. Let g_L be the lorentzian metric such that:

$$g_L|_\xi = g|_\xi, \qquad g_L(\nu,\nu) = -1, \qquad g_L(\nu,-)|_\xi = 0.$$

We can then define an inclusion

$$i : \mathcal{P}\mathfrak{lanes} \to \mathfrak{Cones}$$

by assigning to ξ the light–cone of g_L (ξ will then be space–like and oriented). Similarly, let $\tilde{\mathcal{D}} \in \mathfrak{Lorentz}$ be the the 2–distribution whose tautological map is

$$\tilde{f}(p,t) = \frac{f(p,t) + \nu(p)}{|f(p,t) + \nu(p)|}.$$

This defines an inclusion

$$\tilde{i} : \mathcal{F}\mathfrak{Cartan} \to \mathfrak{Lorentz}$$

that assigns $\tilde{\mathcal{D}}$ to $\mathcal{D}$. The inclusion $\tilde{i}$ is equivalently obtained by composing the tautological map of any formal Cartan prolongation with (the tautological map associated to) i.

Having fixed an orientation of N, we have that the curves $t \to \tilde{f}(p,t) \in \mathbb{S}T_pN \cong \mathbb{S}^2$ are *convex.* If we had chosen $-\nu$ to push ξ (or chosen the opposite orientation of N), we would have obtained *concave* curves. We write $\mathfrak{Lorentz}^+$ for the subspace of $\mathfrak{Lorentz}$ where the structures are given by convex curves. We denote by $\mathfrak{Lorentz}^-$ the subspace where they are concave. The inclusion we have defined is of the form $\mathcal{F}\mathfrak{Cartan} \to \mathfrak{Lorentz}^+$.

Proposition 2.1. *The inclusions*

$$i : \mathcal{P}\mathfrak{lanes} \to \mathfrak{Cones}$$

$$\tilde{i} : \mathcal{F}\mathfrak{Cartan} \to \mathfrak{Lorentz}^+$$

are homotopy equivalences.

Proof. Using the metric g, we are able to recover a plane field uniquely from an element in $\mathfrak{Cones}$. Its orientation is obtained from the orientation of the cone. This defines a map $j : \mathfrak{Cones} \to \mathcal{P}\mathfrak{lanes}$ that is the homotopy inverse of the map i. The tautological map associated to j yields a projection $\tilde{j} : \mathfrak{Lorentz}^+ \to \mathcal{F}\mathfrak{Cartan}$ that is a homotopy inverse of $\tilde{i}$. $\square$

In the sections that remain we will study the homotopy type of $\mathcal{F}\mathfrak{Cartan}$. This will not only give us the homotopy type of $\mathfrak{Lorentz}$, but also will help us compute the homotopy groups of $\mathfrak{Cartan}$.

2.2 Cartan prolongations over a fixed contact structure

First we will describe the homotopy groups of the space $\mathcal{C}\mathrm{over}(c,k)$, the fibre in all the fibrations we have introduced. The π_0 case was already described in [46].

Observe that, by fixing some element $\tau \in \mathcal{C}\mathrm{over}(c,k)$, one readily obtains an inclusion $\mathcal{C}\mathrm{over}(kc,1) \subset \mathcal{C}\mathrm{over}(c,k)$ by making τ act by pullback. Since $\mathcal{C}\mathrm{over}(kc,1)$ is a group that contains the gauge transformations $\mathcal{G}(kc)$ of $N(kc)$ as an abelian subgroup, we can regard $\mathcal{G}(kc)$ as a subspace of $\mathcal{C}\mathrm{over}(c,k)$ as well.

Lemma 2.2. *For any $\tau \in \mathcal{C}\mathrm{over}(c,k)$, the inclusions*

$$\mathcal{G}(kc) \to \mathcal{C}\mathrm{over}(kc,1) \to \mathcal{C}\mathrm{over}(c,k)$$

are weak homotopy equivalences.

Proof. Let $\phi : (\mathbb{D}^j, \partial\mathbb{D}^j, 1) \to (\mathcal{C}\mathrm{over}(c,k), \mathcal{G}(kc), \tau)$ be a continuous function representing an element in $\pi_j(\mathcal{C}\mathrm{over}(c,k), \mathcal{G}(kc), \tau)$. It is sufficient to show that it retracts to $\mathcal{G}(kc)$.

Restricted to an i–simplex Δ_i of N, the bundles $N(c)$ and $N(kc)$ are trivial. There, ϕ can be thought as a $(\Delta_i \times \mathbb{D}^j)$ –family of positively oriented submersions of $\mathbb{S}^1$ onto itself with k–sheets. The SO(2)–bundle structure on $N(c)$ can be taken to be the one induced from $N(kc)$ by using τ, and hence τ can be assumed to be the map z^k on each fibre; the elements of $\mathcal{G}(kc) \subset \mathcal{C}\mathrm{over}(c,k)$ are those of the form $\phi \circ z^k$ with ϕ a $(\Delta_i \times \mathbb{D}^j)$–family of rotations.

Assume that a suitable homotopy has already been found in the $(i-1)$th skeleton of N. The $(\Delta_i \times \mathbb{D}^j)$–family of submersions of $\mathbb{S}^1$ onto itself can be lifted to define a family in $\mathrm{Diff}^+(\mathbb{S}^1)$ such that its boundary lies in $\mathbb{S}^1$, the rotations. Recalling that $\mathbb{S}^1 \to \mathrm{Diff}^+(\mathbb{S}^1)$ is a weak homotopy equivalence concludes the inductive step. $\square$

Lemma 2.3. *The homotopy groups of $\mathcal{G}(kc)$, and hence of $\mathcal{C}\mathrm{over}(c,k)$, are given by:*

$$\pi_0 = H^1(N,\mathbb{Z}), \qquad \pi_1 = \mathbb{Z}, \qquad \pi_j = 0, \text{ for } j > 1.$$

Proof. Recall that $\mathbb{S}^1$ is the classifying space for the discrete group $\mathbb{Z}$. Then:

$$\pi_0(\mathcal{G}(N)) = \pi_0(\mathcal{M}\mathrm{aps}(N,\mathbb{S}^1)) = H^1(N,\mathbb{Z}).$$

In general, it is a result of Thom [78] that $\pi_j(\mathcal{M}\mathrm{aps}(N,K(G,n))) = H^{n-j}(N,G)$. $\square$

Remark 2.4. Lemma 2.3 can be proved using obstruction theory as in Lemma 2.2. This is useful to provide a geometrical interpretation of the result. Let us outline the argument, which is similar to the one presented in [46]. We need to fix a basepoint $\tau \in \mathcal{C}\mathrm{over}(c,k)$.

An explicit identification between $\pi_0(\mathcal{C}\mathrm{over}(c,k))$ and $H^1(N,\mathbb{Z})$ can be given as follows. Take an element $\nu \in \mathcal{C}\mathrm{over}(c,k)$. Over each loop $\gamma \subset N$, the bundles $N(c)$ and $N(kc)$ trivialise. Given any section $s \in \Gamma(N(c)|_\gamma)$, one can compute the degree of $\nu(s)$ with respect to $\tau(s)$. This gives a homomorphism $H_1(N,\mathbb{Z}) \to \mathbb{Z}$ and thus an element in $H^1(N,\mathbb{Z})$. This element only depends on the connected component of ν; we call it the **horizontal distance** between τ and ν.

Similarly, let $\nu_t \in \mathcal{C}\mathrm{over}(c,k)$, $t \in \mathbb{S}^1$, be a loop with $\nu_1 = \tau$. Take a point $p \in N$ and lift it to a point $P \in N(c)$. We say that the degree of $t \to \nu_t(P)$, $t \in \mathbb{S}^1$, as a loop in the fibre of $N(kc)$ over p, is the **looping number**. This identifies $\pi_1(\mathcal{C}\mathrm{over}(c,k))$ with $\mathbb{Z}$.

2.3 Statement and proof of the results

2.3.1 Formal Cartan prolongations

Let ξ be a plane field with Euler class kc. Recall that our objective is to understand the homotopy type of $\mathcal{F}\mathfrak{Cartan}(c,k,[\xi])$. From Lemmas 2.2 and 2.3, and the homotopy long exact sequence for the fibration, it easily follows that:

$$\pi_j(\mathcal{F}\mathfrak{Cartan}(c,k,[\xi])) = \pi_j(\mathcal{P}\mathfrak{lanes}(\xi)), \qquad \text{for } j > 2.$$

However, the cases of π_0, π_1, and π_2 are more subtle. The key is understanding the connecting morphism

$$\pi_j(\mathcal{P}\mathfrak{lanes}(\xi)) \to \pi_{j-1}(\mathcal{C}\mathrm{over}(c,k)) \qquad j = 1,2,$$

which is not zero in general. However, according to the commutative diagram above, this connecting morphism factors through $\pi_j(\mathcal{B}\mathfrak{undles}(kc))$. This motivates us to consider the subgroup:

$$\pi_j^{\mathrm{trivial}}(\mathcal{P}\mathfrak{lanes}(\xi)) = \ker(\pi_j(\mathcal{P}\mathfrak{lanes}(\xi)) \to \pi_j(\mathcal{B}\mathfrak{undles}(kc))).$$

The key observation is that these groups can be computed explicitly using obstruction classes (which is something we will not do).

Then, the desired statement is:

Theorem 2.5. *Let ξ be a plane field of Euler class kc. Then:*

$$\begin{aligned}
&\pi_0(\mathcal{F}\mathfrak{Cartan}(c,k,[\xi])) = \pi_0(\mathcal{P}\mathfrak{lanes}(\xi)) \times H^1(N,\mathbb{Z}_2)\\
&\pi_1(\mathcal{F}\mathfrak{Cartan}(c,k,[\xi])) = \pi_1^{\text{trivial}}(\mathcal{P}\mathfrak{lanes}(\xi)) \times \mathbb{Z}_2\\
&\pi_2(\mathcal{F}\mathfrak{Cartan}(c,k,[\xi])) = \pi_2^{\text{trivial}}(\mathcal{P}\mathfrak{lanes}(\xi))\\
&\pi_j(\mathcal{F}\mathfrak{Cartan}(c,k,[\xi])) = \pi_j(\mathcal{P}\mathfrak{lanes}(\xi)) \qquad \textit{if } j > 2.
\end{aligned}$$

The term $H^1(N,\mathbb{Z}_2)$ is the mod 2 reduction of the horizontal distance. Similarly, the term $\mathbb{Z}_2$ is the parity of the looping number.

In particular, this computes the homotopy groups of the space $\mathfrak{Lorentz}$ as well.

2.3.2 Cartan prolongations

If ξ is a contact structure with Euler class kc, the reasoning is analogous. We are interested in the subgroups:

$$\pi_j^{\text{trivial}}(\mathcal{C}\text{-}\mathcal{S}\mathfrak{trs}(\xi)) = \ker(\pi_j(\mathcal{C}\text{-}\mathcal{S}\mathfrak{trs}(\xi)) \to \pi_j(\mathcal{B}\mathfrak{undles}(kc))).$$

The main result about Cartan prolongations states:

Theorem 2.6. *Let ξ be an overtwisted contact structure of Euler class kc. Then:*

$$\begin{aligned}
&\pi_0(\mathfrak{Cartan}(c,k,[\xi])) = \pi_0(\mathcal{C}\text{-}\mathcal{S}\mathfrak{trs}(\xi)) \times H^1(N,\mathbb{Z}_2)\\
&\pi_1(\mathfrak{Cartan}(c,k,[\xi])) = \pi_1^{\text{trivial}}(\mathcal{C}\text{-}\mathcal{S}\mathfrak{trs}(\xi)) \times \mathbb{Z}_2\\
&\pi_2(\mathfrak{Cartan}(c,k,[\xi])) = \pi_2^{\text{trivial}}(\mathcal{C}\text{-}\mathcal{S}\mathfrak{trs}(\xi))\\
&\pi_j(\mathfrak{Cartan}(c,k,[\xi])) = \pi_j(\mathcal{C}\text{-}\mathcal{S}\mathfrak{trs}(\xi)) \qquad \textit{if } j > 2.
\end{aligned}$$

The term $H^1(N,\mathbb{Z}_2)$ is the mod 2 reduction of the horizontal distance. Similarly, the term $\mathbb{Z}_2$ is the parity of the looping number.

The proof relies on understanding the inclusion $\pi_j(\mathcal{C}\text{-}\mathcal{S}\mathfrak{trs}(\xi)) \to \pi_j(\mathcal{B}\mathfrak{undles}(kc))$. For an arbitrary contact structure ξ this is very difficult. However, if ξ is assumed to be overtwisted the problem simplifies considerably, due to Eliashberg's theorem [19] (which we already stated in the Preamble). Let us recall it here in a form that is useful for us:

Lemma 2.7. *Let ξ be an overtwisted contact structure. The inclusion*

$$\mathcal{C}\text{-}\mathcal{S}\mathfrak{trs}(\xi) \to \mathcal{P}\mathfrak{lanes}(\xi)$$

is surjective in all homotopy groups, where ξ is assumed to be the basepoint. Additionally, this map is a bijection at the level of connected components.

Eliashberg's theorem additionally states that the inclusion is a weak homotopy equivalence if we restrict to contact structures/plane fields with a fixed overtwisted disc. What is relevant about this is that the groups $\pi_j^{\text{trivial}}(\mathcal{C}\text{-}\mathcal{S}\mathfrak{trs}(\xi))$ are not trivial to compute because they do not correspond to $\pi_j^{\text{trivial}}(\mathcal{P}\mathfrak{lanes}(\xi))$ in general.

2.3.3 Technical ingredients of the proof(s)

The proofs of Theorems 2.5 and 2.6 are nearly identical. Both of them rely on understanding the connecting morphisms for the corresponding fibrations.

2.3.3.1 The connecting morphism for bundles

The next two lemmas show that the connecting morphism is a bijection in the case of bundles.

Lemma 2.8. *The connecting morphism $\pi_j(\mathcal{B}\mathfrak{undles}(kc)) \to \pi_{j-1}(\mathcal{C}\mathrm{over}(c,k))$ is injective.*

Proof. Take any element in $\pi_j(\mathcal{F}\mathfrak{Cartan}^\infty(c,k))$, and find a representative $K \subset \mathcal{F}\mathfrak{Cartan}^\infty(c,k)$. Its image $\xi \subset \mathcal{B}\mathfrak{undles}(kc)$ can be understood as a 2–plane bundle over $N \times \mathbb{S}^j$ with vanishing Euler class when restricted to $\{x\} \times \mathbb{S}^j$; this follows from the fact that it is k–covered by $(N(c) \times \mathbb{S}^j)|_{\{x\}\times\mathbb{S}^j}$. Therefore, the family ξ is trivial. □

Similarly:

Lemma 2.9. *The connecting morphism $\pi_j(\mathcal{B}\mathfrak{undles}(kc)) \to \pi_{j-1}(\mathcal{C}\mathrm{over}(c,k))$ is surjective.*

Proof. Take ξ to be the basepoint in $\mathcal{B}\mathfrak{undles}(kc)$ and fix a lift $\tau \in \mathcal{F}\mathfrak{Cartan}^\infty(c,k)$. Take a class $G \in \pi_{j-1}(\mathcal{C}\mathrm{over}(c,k))$, which we can think of as a homotopy class in the gauge transformations of ξ, by Lemma 2.2.

Lemma 2.3 implies that G is given by a cohomology class $g \in H^{2-j}(N,\mathbb{Z})$. Similarly, ξ is given, up to bundle isomorphism, by its Euler class $e \in H^2(N,\mathbb{Z})$. Recall that the Künneth formula yields an isomorphism

$$H^{2-j}(N,\mathbb{Z}) \oplus H^2(N,\mathbb{Z}) \xrightarrow{(\alpha,\beta)} H^2(N \times \mathbb{S}^j,\mathbb{Z}).$$

Consider the unique, up to homotopy, j–sphere K of bundles based on ξ and having $\alpha(g) + \beta(e)$ as its Euler class when regarded as a plane bundle over $N \times \mathbb{S}^j$. We claim that the connecting morphism maps $[K]$ to G.

Take $j = 1$. Write $P : N \times \mathbb{S}^1 \to N$. Write $Q : N \times [0,1] \to N \times \mathbb{S}^1$ for the obvious quotient map. There is a unique, up to homotopy, isomorphism between Q^*K and $Q^*P^*\xi$ extending the identification $(Q^*K)|_{N\times\{0\}} = \xi$. The identification of $(Q^*K)|_{N\times\{0\}}$ with $(Q^*K)|_{N\times\{1\}}$ yields a gauge transformation ϕ of ξ.

We claim that ϕ is a representative of G. Recall that the Euler class of a 2–plane bundle over the torus can be computed as follows: find a section over the complement of the meridian γ and compare the degrees of the two resulting sections over γ. Let now γ be some embedded loop in N, and let T be the corresponding torus on $N \times \mathbb{S}^1$. By construction, $K|_T$ has Euler class $\alpha(g)|_T$, which implies that $\phi|_\gamma$ is described by $g|_\gamma$. Since gauge transformations are characterised by their action over loops, the claim follows.

The case $j = 2$ is similar. In that case, we have to study what happens over a single point $x \in N$ and the corresponding sphere $\{x\} \times \mathbb{S}^2$. □

2.3.3.2 Non–trivial families of plane fields

The following proposition shows that there are many families of plane fields which are non–trivial as families of vector bundles.

Proposition 2.10. *Let $d_j = 2v_j \in H^2(N \times \mathbb{S}^j,\mathbb{Z})$, $j = 1,2$. Fix $\xi \in \mathcal{P}\mathfrak{lanes}$ with Euler class $d_j|_N$. Then, there is a sphere K_j in $\mathcal{P}\mathfrak{lanes}$ based at ξ whose Euler class as a 2–plane bundle over $N \times \mathbb{S}^j$ is d_j.*

Proof. Assume $j = 1$. Take a CW–decomposition of N with only one top cell. Take the CW–decomposition of $\mathbb{S}^1$ with a single 1–cell and x the unique 0–simplex. Denote by $\mathcal{T}$ the product CW–decomposition in $N \times \mathbb{S}^1$. Write $\mathcal{T}^* \subset \mathcal{T}$ for the collection of cells not contained in $N \times \{x\}$. Deform ξ to be constant (as a map into the Grassmannian) over the 1–skeleton of N. We define $(K_1)|_{N\times\{x\}} = \xi$ and we aim to extend it to $\mathcal{T}^*$.

Over the 1–cells, K_1 can be defined to be constant, like ξ. Over a 2–cell Δ_2, we define it to be a map into $\mathbb{S}^2$ of degree $\phi_1(\Delta_2)$, where $[\phi_1] = v_1$. This provides the desired Euler class. Over the 3–skeleton of $\mathcal{T}^*$, the obstruction for extending is given by $d\phi_1$ by construction, which evaluates zero over the cells of $\mathcal{T}^*$. This leaves the single 4–cell Δ_4 to fill. The obstruction is the degree of the map $K_1 : \partial\Delta_4 \to \mathbb{S}^2$. We can trace back our steps and modify K_1 over some 3–cell to make sure this degree is zero.

Assume $j = 2$. Then, the isomorphism:

$$H^2(\mathbb{S}^2, \mathbb{Z}) \oplus H^2(N, \mathbb{Z}) \xrightarrow{(\alpha,\beta)} H^2(N \times \mathbb{S}^2, \mathbb{Z}),$$

indicates that we can simply compute the Euler class of any plane bundle by evaluating separately on $N \times \{x\}$ or $\{p\} \times \mathbb{S}^2$. Take the manifold $N \times \mathbb{D}^2$: over it, we have (the pullback of) the bundle TN which is trivialised as the trivial $\mathbb{R}^3$–bundle; as a 2–distribution inside, we define K_2, which is $\mathbb{D}^2$–invariant and equal to ξ on every $N \times \{x\}$. We aim two glue two copies of $N \times \mathbb{D}^2$ so that the glued copies of K_2 have the desired Euler class when restricted to each $\{p\} \times \mathbb{S}^2$.

Consider the loop $\mathbb{S}^1 \to \mathrm{SO}(2)$ realising the Euler class $\alpha^{-1}(d_2) \in H^2(\mathbb{S}^2, \mathbb{Z})$ through the clutching construction and denote by $\phi : \mathbb{S}^1 \to \mathrm{SO}(3)$ its inclusion into $\mathrm{SO}(3)$. Observe that, since $d_2 = 2v_2$, ϕ is contractible in $\mathrm{SO}(3)$. We can define then another map $\Phi : N \times \mathbb{S}^1 \to \mathrm{SO}(3)$ so that:

- $\Phi|_{\{p\}\times\mathbb{S}^1} = \phi$, up to a $\mathrm{SO}(3)$–transformation that only depends on p,
- Φ fixes (not pointwise) the plane ξ.

What we are essentially saying is that ϕ was a family of rotations of the XY–plane that was lifted to $\mathbb{R}^3$, and Φ is a p–dependent family of rotations that looks the same but, instead, the plane that $\Phi|_{\{p\}\times\mathbb{S}^1}$ rotates is ξ_p. Since ϕ was contractible, so is Φ, so the resulting $\mathbb{R}^3$–bundle is trivial. However, in each $\{p\} \times \mathbb{S}^2$ the restriction of ξ has been twisted to have Euler class $\alpha^{-1}(d_2)$, proving the claim. □

2.3.4 Concluding the proof(s)

Proof of Theorem 2.6. Let ξ be an overtwisted contact structure and fix $\tau \in \mathfrak{Cartan}(c, k, \xi)$ a basepoint in the fibre over ξ. The existence of τ identifies the connected components of $\mathfrak{Cartan}(c, k, \xi) \cong \mathcal{C}\mathrm{over}(c, k)$ with $H^1(N, \mathbb{Z})$, as in Lemma 2.3.

Let us study the connecting morphism

$$\pi_1(\mathcal{C}\text{-}\mathcal{S}\mathfrak{trs}(\xi)) \to \pi_0(\mathcal{C}\mathrm{over}(c, k)).$$

Using Lemma 2.8 we deduce that its kernel is the space $\pi_1^{\mathrm{trivial}}(\mathcal{C}\text{-}\mathcal{S}\mathfrak{trs}(\xi))$. Let us compute its image. Let $g \in \mathcal{C}\mathrm{over}(c, k)$ and denote by ν the corresponding prolongation in $\mathfrak{Cartan}(c, k, \xi)$. By Lemma 2.9 there is a loop of vector bundles, all of them of Euler class kc, producing the class of g through the connecting morphism $\pi_1(\mathcal{B}\mathfrak{undles}(kc)) \to \pi_0(\mathcal{C}\mathrm{over}(c, k))$. By Proposition 2.10, this loop can be realised by a loop of plane fields based on ξ if and only if $[g] \in H^1(N, \mathbb{Z})$ is even. Then, Lemma 2.7 allows us to turn this into a loop of contact structures ξ_t, $t \in \mathbb{S}^1$ based on $\xi_1 = \xi$.

The case $\pi_2(\mathcal{C}\text{-}\mathcal{S}\mathfrak{trs}(\xi)) \to \pi_1(\mathcal{C}\text{over}(c, k))$ is analogous. □

The proof of Theorem 2.5 is exactly the same but do not have to apply Lemma 2.7.

Chapter 3

Existence of Engel structures

This chapter contains one of the main results of the thesis: any manifold satisfying the obvious topological constraints (say, parallelisability up to a 4-fold cover), admits an Engel structure in each possible homotopy class. The general structure of the proof is standard for an h–principle: one triangulates the ambient manifold and, proceeding inductively over the skeleton, constructs the Engel structure. This is easy for the lower dimensional cells, but for the 4–dimensional ones a more elaborate model is necessary: in the end, the whole problem reduces to understanding how germs of Engel structures in $\partial\mathbb{D}^4$ extend to the interior.

Following this general scheme, there are actually several variations on how to tackle the proof. We shall explain two different methods for dealing with the lower dimensional cells and three different methods to solve the extension problem in $\mathbb{D}^4$. Each approach offers a distinct way of understanding the nature of Engel structures.

The first proof of this result appeared in my article [10] with R. Casals, J.L. Pérez, and F. Presas. The route taken there somehow highlighted the interplay between contact topology and convexity: the combination of both paradigms, exemplified by the *four–leaf clover* construction, was essential for the argument. The approach in [10] is explained in Sections 3.5.4 and 3.6.1. Alternatively, instead of the extension process described in Section 3.6.1, one can use the one from Section 3.6.2; this was actually the original approach in the article.

One of the main themes of this thesis is that, although contact structures interact heavily with Engel structures locally (meaning, when we study local models), globally they do not (except for very particular examples of Engel structures, like prolongations). Following this line of thought, the methods in Sections 3.5.5 and 3.6.3 are able to prove the existence theorem "completely avoiding" contact structures, they just rely on our characterisation of the Engel condition in terms of convex curves. This manner of proceeding is slightly more powerful than the others: it allows for the Engel structure we produce to contain any given line field.

Knowing that a wealth of Engel structures can be produced leads naturally to the question of whether we can classify them. We have chosen to defer this issue to the next chapter.

3.1 Statement of the main theorem and outline of the proof(s)

3.1.1 The existence theorem

As promised, here is the statement of the main result of this chapter:

Theorem 3.1. *Let M be a smooth 4–manifold, not necessarily orientable. Then, the map*

$$\pi_k(i) : \pi_k(\mathfrak{Engel}(M)) \longrightarrow \pi_k(\mathcal{F}\mathfrak{Engel}(M))$$

is surjective for every $k \geq 0$. In particular, every formal Engel structure is homotopic to the flag of a genuine Engel structure.

To put things into perspective, we should recall that T. Vogel had already showed in [82] that any parallelisable 4–manifold admits an Engel structure (Subsection 1.3.6). Equivalently, his result says that the set $\pi_0(\mathfrak{Engel}(M))$ is not empty as soon as M admits a completely oriented full flag.

As explained in the introduction, the main idea (in the π_0 case) is to triangulate M suitably and then construct the Engel structure (by deforming the formal one we are given) cell by cell. The process of deforming it in a neighbourhood of the 3–skeleton we call the **reduction process**. Then, the final step of extending the Engel structure to the interior of the 4–cells (relatively to their boundary) we call **solving the extension problem**.

The two reduction methods are explained, respectively, in Subsections 3.5.4 and 3.5.5. Both of them rely on the triangulation being adapted to the (formal) Engel structure and in having constructed some particularly nice neighbourhoods of the cells; this is explained in Section 3.2. This way of topologically decomposing the manifold will possibly remind the reader of Thurston's papers on the construction of foliations [79, 80] or Eliashberg's paper on 3–dimensional overtwisted contact structures [19]. Indeed, in general lines our approach is similar to theirs, but in our case the parametric case works out almost automatically. It is worth remarking that the results of Section 3.2 will be used again when we discuss immersions tangent to Engel structures in Chapter 5, since they are very useful for h–principles where a line field is involved.

The three extension methods can be found in Section 3.6. They prove that certain germs of Engel structure in $\partial\mathbb{D}^4$ extend to the interior of $\mathbb{D}^4$. The precise form of the germs we are able to deal with is described in Section 3.4; its contents are also necessary for the understanding of Section 3.5, since the final output of the reduction process must be a germ well suited for one of the extension procedures. Additionally, the contents of Subsection 3.6.3 rely on understanding the space of convex curves in $\mathbb{S}^2$ (in light of Proposition 1.12, it is clear that convexity plays a role in the manipulation of Engel structures). To this end, we devote Section 3.3 to a detailed discussion of the concepts and results we shall need; we will use heavily the work of Little [50].

In Section 3.7 we will assemble all the pieces together to prove Theorem 3.1.

It is worth remarking that much of the language we will introduce along the way (*Engel shells* and such) is reminiscent of the one used by Borman, Eliashberg, and Murphy in [6]. However, it is fair to argue that the similarities between both results end there. From a technical point of view, as we already mentioned, the main inspirations behind the reduction process are [79, 80, 19].

3.1.2 First corollary: Engel cobordisms between contact manifolds

The methods used in the proof of Theorem 3.1 are not completely *relative with respect to the domain*: that is, if we are given a formal Engel structure that is Engel in a closed set A, we are, in general, not able to make it Engel everywhere while keeping it the same close to A. This is due to the fact that we do not actually solve the extension problem in $\mathbb{D}^4$ in full generality: we solve it for some germs along $\partial\mathbb{D}^4$ that are *flexible* and, luckily for us, the reduction methods we have are able to produce germs that are precisely flexible.

However, in some cases, it is possible to prove relative statements:

Corollary 3.2. *Let $(M, \partial M)$ be a 4–manifold with boundary and let $(\mathcal{W}, \mathcal{D}, \mathcal{E})$ be a formal Engel structure such that:*

- *$(\mathcal{W}, \mathcal{D}, \mathcal{E})$ is Engel in $\mathcal{O}p(\partial M)$,*
- *$\mathcal{W}$ is transverse to ∂M.*

Then, $(\mathcal{W}, \mathcal{D}, \mathcal{E})$ is homotopic to an Engel structure relative to ∂M.

From Corollary 3.2, it readily follows:

Corollary 3.3. *Let $(M, \partial M)$ be a 4–manifold with boundary and let $(\mathcal{W}, \mathcal{D}, \mathcal{E})$ be a formal Engel structure with $\mathcal{W}$ transverse to ∂M. Let ξ be a contact structure in ∂M that is homotopic to the plane field $\mathcal{E} \cap T(\partial M)$.*

Then, $(\mathcal{W}, \mathcal{D}, \mathcal{E})$ is homotopic to an Engel structure whose restriction to ∂M is ξ.

These statements are proven in Subsection 3.7.2. Corollary 3.3 implies that the notion of Engel cobordism or Engel fillability is not particularly relevant in order to distinguish contact manifolds: Given a smooth 3–manifold V and two homotopic contact structures ξ_0 and ξ_1 with vanishing first Chern class, there exists an Engel structure on the trivial cobordism $V \times [0, 1]$ inducing the contact manifolds $(V \times \{0\}, \xi_0)$ and $(V \times \{1\}, \xi_1)$ on the boundary components. This result is an instance of something we already mentioned at the beginning of the chapter: global Engel topology (seemingly) does not interact meaningfully with global contact topology. For a less pessimistic take on this, refer to the Closing Remarks (the only appendix of the thesis).

Remark 3.4. It is worth pointing out that T. Vogel [82, Remark 4.8] had already constructed an Engel cobordism between the overtwisted and tight contact structures in $\mathbb{S}^3$ that live in the same formal class, solving a question posed by J. Adachi. Corollary 3.3 is a more general version of his statement.

3.1.3 Second corollary: existence of foliated Engel structures

As we saw see in the first part of this thesis, it is often fruitful to endow foliations with additional structures (either leafwise or transversely). In our case, the proof of Theorem 3.1 readily adapts to prove an existence theorem for leafwise Engel structures. The point is that the parametric case π_k, $k > 0$, of Theorem 3.1 and the foliated case are pretty much analogous. As such, they are proven simultaneously in Subsection 3.7.1.

Let us set up the notation. Let M be a smooth manifold endowed with a rank 4 foliation $\mathcal{F}$. A 2–distribution $\mathcal{D} \subset T\mathcal{F}$ is called a **foliated Engel structure** if it is an Engel structure when restricted to each leaf. A flag $\mathcal{W}^1 \subset \mathcal{D}^2 \subset \mathcal{E}^3 \subset T\mathcal{F}^4$ satisfying the bundle isomorphisms:

$$\det(\mathcal{D}) = \mathcal{E}/\mathcal{W} \tag{3.1}$$

$$\det(\mathcal{E}/\mathcal{W}) = T\mathcal{F}/\mathcal{E} \tag{3.2}$$

is said to be a **formal foliated Engel structure**. Denote the space of formal foliated Engel structures, endowed with the C^0–topology, by $\mathcal{F}\mathfrak{Engel}(\mathcal{F})$. Write $\mathfrak{Engel}(\mathcal{F})$ for the space of foliated Engel structures with the C^2–topology. Then:

Corollary 3.5. *The inclusion*

$$\pi_k(i) : \pi_k(\mathfrak{Engel}(\mathcal{F})) \longrightarrow \pi_k(\mathcal{F}\mathfrak{Engel}(\mathcal{F}))$$

is surjective for every $k \geq 0$.

This result can be brought together with Corollary 3.2 to show:

Corollary 3.6. *Let M be a smooth manifold with boundary endowed with a 4–dimensional foliation $\mathcal{F}$ transverse to ∂M. Let $(\mathcal{W}, \mathcal{D}, \mathcal{E})$ be a formal foliated Engel structure with $\mathcal{W}$ is transverse to ∂M. Assume further that $\mathcal{E}$ restricts to a foliated contact structure in $(\partial M, \mathcal{F}|_{\partial M})$.*

Then, $(\mathcal{W}, \mathcal{D}, \mathcal{E})$ is homotopic to a foliated Engel structure inducing the same contact foliation in ∂M.

3.1.4 Third corollary: line fields tangent to Engel structures

A very intriguing question (for the author at least) is the following: how far is a given (non–vanishing) vector field from being legendrian/being the kernel of an even–contact structure/being horizontal for a distribution of a given class? Apart from the obvious restrictions on the linearised return map of a closed orbit in each case, there do not seem to be interesting statements of this flavour (to the author's knowledge). In this direction, Engel structures present a lot of flexibility and one is actually able to prove:

Corollary 3.7. *Let M be a 4–manifold and let $\mathcal{Y}$ be a line field. Then, $\mathcal{Y}$ is tangent to an Engel structure if and only if it is the line field of a formal Engel structure. The same holds parametrically and in the foliated setting.*

The proof relies on the reduction and extension methods based on convexity that are presented in Subsections 3.5.5 and 3.6.3, respectively. Note that we do not claim for $\mathcal{Y}$ to be *kernel* of the resulting Engel structure. In fact, due to the nature of the proof, $\mathcal{Y}$ will be everywhere complementary to the kernel.

3.2 Parametric triangulations adapted to line fields

As pointed out before, our reduction process relies on having a particular decomposition of the ambient manifold M: first we triangulate M and then we find a nice collection of neighbourhoods for the cells. The fundamental property of the triangulation is that it is in general position with respect to the line field $\mathcal{W}$ (meaning that the lower dimensional cells are transverse to it). This is the problem that we concern ourselves with in this section.

One of the key ideas for dealing with the parametric (or foliated) case of Theorem 3.1 is the following: If we are given a family of (formal) Engel structures parametrised by a manifold K, we can just triangulate $M \times K$ and adapt the reduction and extension processes to this higher dimensional manifold. As such, all the results in this section will be stated for manifolds of general dimension.

3.2.1 The triangulation theorem. Statement and proof

Definition 3.8. *Let M be an n–dimensional manifold, $n \geq 2$, endowed with a line field $\mathcal{W}$. Let $\mathcal{T}$ be triangulation of M, which we understand as a finite collection of simplices $\{\sigma\}$. Let $\{\mathcal{U}(\sigma)\}_{\sigma\in\mathcal{T}}$ be a finite collection $\{\mathcal{U}(\sigma)\}_{\sigma\in\mathcal{T}}$ of closed n–discs such that:*

a. Each simplex σ is contained in the union $\cup_{\tau\subseteq\sigma}\mathcal{U}(\tau)$.

b. For each pair of simplices σ, σ', neither of them containing the other, we have $\mathcal{U}(\sigma)\cap\mathcal{U}(\sigma') = \emptyset$.

1. Fix coordinates (p,t) in $\mathbb{D}^{n-1}_{1+\varepsilon} \times [-\varepsilon, 1+\varepsilon]$. For each simplex σ, there exists a map

$$\phi(\sigma) : \mathcal{O}p(\mathcal{U}(\sigma)) \longrightarrow \mathbb{D}^{n-1}_{1+\varepsilon} \times [-\varepsilon, 1+\varepsilon]$$

satisfying $\phi(\sigma)_*\mathcal{W} = \langle\partial_t\rangle$ and $\phi(\sigma)(\mathcal{U}(\sigma)) = \mathbb{D}^{n-1} \times [0,1]$.

2. For each simplex $\sigma \in \mathcal{T}^{(j)}$, $j < n$, any orbit of the line field $\mathcal{W}$ in the disc $\mathcal{U}(\sigma)$ either avoids the set $\cup_{\tau\subsetneq\sigma}\mathcal{U}(\tau)$ or is entirely contained on it. The same is true for $\mathcal{O}p(\mathcal{U}(\sigma))$.

*Then, we say that the triangulation $\mathcal{T}$ and the cover $\{\mathcal{U}(\sigma)\}$ are **adapted** to $\mathcal{W}$.*

Let us make a couple of observations. First, the set $\mathcal{U}(\sigma)$ is not a neighbourhood of σ unless σ is zero dimensional. Rather, the boundary of σ is covered by all the $\mathcal{U}(\tau)$, where $\tau \subset \sigma$ is a subsimplex, and $\mathcal{U}(\sigma)$ just covers the interior of σ.

Secondly, the point of having $\phi(\sigma)$ defined in a thickening $\mathcal{O}p(\mathcal{U}(\sigma))$ of $\mathcal{U}(\sigma)$ is that if we deform a formal Engel structure in $\mathcal{U}(\sigma)$, we need a little region to interpolate between the new structure and the original one that is defined in the remainder of the manifold.

Also, let us introduce some notation. We say that the region $\phi(\sigma)^{-1}(\mathbb{S}^{n-1}\times[0,1])$ is the **vertical boundary** of the flowbox $\mathcal{U}(\sigma)$; in the vertical boundary, $\mathcal{W}$ is tangent. Similarly, we say that $\phi(\sigma)^{-1}(\mathbb{D}^{n-1}\times\{0,1\})$ is the **horizontal boundary**; in the horizontal boundary, $\mathcal{W}$ is transverse.

Theorem 3.9. *Let M be an n–dimensional manifold, $n \geq 2$, endowed with a line field $\mathcal{W}$. Then, there exist a triangulation $\mathcal{T}$ and a cover $\{\mathcal{U}(\sigma)\}$ that are adapted to $\mathcal{W}$. $\mathcal{T}$ can be taken to be arbitrarily fine.*

Let us briefly sketch the idea of the proof. The way in which the $\mathcal{U}(\sigma)$ are constructed is inductively in the dimension of the simplices. Each $\mathcal{U}(\sigma)$ is "tall": it is a flowbox for $\mathcal{W}$ that is very long in the $\mathcal{W}$ direction with respect to the size of its base. This ensures that any higher dimensional simplex τ meets $\mathcal{U}(\sigma)$ in its vertical boundary; in particular, if $\mathcal{U}(\tau)$ is chosen "shorter" than $\mathcal{U}(\sigma)$, Property (2.) holds. Property (b.) holds if all the $\mathcal{U}(\sigma)$ are sufficiently small. All the other properties are automatic. All of this is depicted in Figures 3.1 and 3.2 in the two and three dimensional cases, respectively.

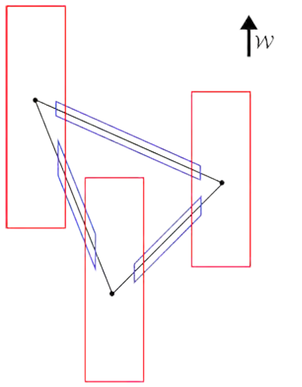

Figure 3.1: Case $n = 2$. *Red*: closed discs for the 0-simplices. *Blue*: closed discs for the 1-simplices.

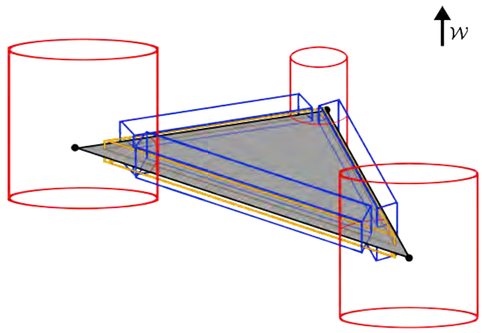

Figure 3.2: Case $n = 3$. *Red*: closed discs for the 0-simplices. *Blue*: closed discs for the 1-simplices. *Yellow*: closed discs for the 2-simplices.

Proof. Fix a Riemannian metric g on the n–manifold M. Apply Thurston's Jiggling Lemma [79, Section 5] to find a triangulation $\mathcal{T}$ whose simplices are all transverse to the line field $\mathcal{W}$. In particular, for every simplex which is not top–dimensional, this means that the angle between the line field and the simplex is strictly positive. We can assume that each simplex is contained in a flowbox of $\mathcal{W}$; over each such flowbox we trivialise $\mathcal{W}$ by a vector field W.

To each j–simplex $\sigma \in \mathcal{T}^{(j)}$, $j < n$, associate a triple of positive real numbers $(r_0, r_1, r_2) \in \mathbb{R}^+ \times \mathbb{R}^+ \times \mathbb{R}^+$. They measure the size of $\mathcal{U}(\sigma) \subset M$, and we shall fix them later on in the construction. Consider a j–dimensional disc $\widetilde{\sigma} \subset \sigma$ satisfying $r_0 < d_g(\partial\widetilde{\sigma}, \partial\sigma) \leq 2r_0$, i.e. it covers almost all of σ. Use the time–r_1 exponential map on an orthonormal basis of $(T\sigma \oplus \mathcal{W})^\perp \subset TM$ and then the time–r_2 flow of W to construct the set:

$$\mathcal{U}(\sigma) \cong \widetilde{\sigma} \times \mathbb{D}^{n-1-j}(r_1) \times [-r_2, r_2].$$

Its boundary can be decomposed in two parts. The region $\partial(\tilde{\sigma} \times \mathbb{D}^{n-1-j}(r_1)) \times [-r_2, r_2]$ will be called the lateral boundary of $\mathcal{U}(\sigma)$.

Let us prove that suitable choices of (r_0, r_1, r_2) (recall that they depend on the particular σ) create a collection $\{\mathcal{U}(\sigma)\}_{\sigma\in\mathcal{T}}$ satisfying all the properties required in the statement for $j < n$. Additionally, the sets $\mathcal{U}(\sigma)$ will be chosen to satisfy the following properties:

3a. For any simplex σ, and any simplex $\tau \supsetneq \sigma$, τ intersects the boundary of the set $\mathcal{U}(\sigma)$ in its lateral region.

3b. For any simplex σ, and any simplex $\tau \subsetneq \sigma$, $\mathcal{U}(\sigma)$ intersects the boundary of $\mathcal{U}(\tau)$ in its lateral region.

These properties imply Property (2.) in the statement; this is because the vertical boundaries are tangent to $\mathcal{W}$. Let us prove that all the desired properties hold by induction on the dimension of the simplices.

For $j = 0$, the first radius r_0 is not defined; but observe that the five properties in the statement are satisfied by choosing $r_1, r_2 > 0$ small enough. Further, Property (3a.) can be satisfied by choosing $r_2 \to 0$ and $r_1/r_2 \to 0$. Indeed, if for each sequence of pairs (r_1, r_2) satisfying $r_2 \to 0$

and $r_1/r_2 \to 0$ Property (3a.) does not hold, then the angle between $\mathcal{W}$ and one of the simplices τ containing the point σ would be zero, and this is impossible. Property (3b.) is vacuous.

Let us explain the inductive step: we suppose that all the properties hold for the k–simplices, $k = 0, \dots, j-1$, and we consider a j–dimensional simplex σ. Choose the first two radii (r_0, r_1) small enough so that

$$\partial(\tilde{\sigma} \times \mathbb{D}^{n-1-j}(r_1)) \subset \cup_{\tau \subsetneq \sigma} \mathcal{U}(\tau),$$

and shrink (r_1, r_2) to guarantee Property (b.). Property (3a.) is achieved by choosing the quotient r_2/r_1 to be large enough and then Property (3b.) is guaranteed if r_2 is chosen small enough.

It remains to consider the n–dimensional simplices $\sigma \in \mathcal{T}^{(n)}$. For each such σ, consider the PL–smooth disc given by the union of the faces of σ where W is inward pointing. By shrinking it slightly and smoothing, we obtain a smooth $(n-1)$–dimensional disc D_-. There is a function $h : D_- \to \mathbb{R}+$ such that $h(p)$ is the first time in the flow of W when $p \in D_-$ is taken to a face of σ where W is outward pointing. Taking a smooth function $H : D_- \to \mathbb{R}+$ arbitrarily close to h allows us to produce the flowbox $\mathcal{U}(\sigma)$ by time–H flowing D_-. By construction, $\partial\mathcal{U}(\sigma)$ is C^0–close to $\partial\sigma$, and therefore $\partial\mathcal{U}(\sigma) \subset \cup_{\tau \subsetneq \sigma} \mathcal{U}(\tau)$. □

3.2.2 Relative triangulations and covers

The natural setting whenever we are trying to prove relative statements is a manifold M with boundary. In this case, the triangulation of interest is actually an extension of a triangulation in ∂M. The two settings of interest are when the line field $\mathcal{W}$ is everywhere transverse or tangent to ∂M.

3.2.2.1 Transverse case

Definition 3.10. *Let M be a compact n–dimensional manifold, $n \geq 2$ with boundary. Let $\mathcal{W}$ be a line field in M that is transverse to ∂M. Let $\mathcal{T} = \{\sigma\}$ be triangulation of M extending a triangulation $\mathcal{T}_\partial$ of ∂M. Let $\{\mathcal{U}(\sigma)\}_{\sigma \in \mathcal{T}}$ be a finite collection of closed n–discs in a thickening $\mathcal{O}p(M)$.*

*We say that $(\mathcal{T}, \{\mathcal{U}(\sigma)\})$ is **adapted** to $(\mathcal{W}, \partial M)$ if the triangulation $\mathcal{T}$ and the covering $\{\mathcal{U}(\sigma)\}_{\sigma \in \mathcal{T}}$ are adapted to $\mathcal{W}$ in the standard sense.*

What we have done is the following. Essentially, we want to use Definition 3.8, but imposing for the triangulation to be relative to ∂M. This is problematic for the simplices contained in the boundary, because their model has to be slightly different from the standard one. An easy fix is simply thickening M by adding a small collar close to ∂M. Using this collar, we can cover M by flowboxes $\mathcal{U}(\sigma) \subset \mathcal{O}p(M)$ as described in Definition 3.8 and proceed in the same way.

It is therefore clear that the proof of Theorem 3.9 readily adapts to show:

Corollary 3.11. *Let M be an n–dimensional manifold, $n \geq 2$, possibly with boundary. Let $\mathcal{W}$ be a line field in M that is transverse to ∂M. Then, there exist a triangulation $\mathcal{T}$ and a cover $\{\mathcal{U}(\sigma)\}$ that are adapted to $(\mathcal{W}, \partial M)$. Additionally, $\mathcal{T}$ can be taken to be arbitrarily fine.*

3.2.2.2 Tangent case

The tangent case is a bit more subtle:

Definition 3.12. *Let M be a compact n–dimensional manifold, $n \geq 2$ with boundary. Let $\mathcal{W}$ be a line field in M that is tangent to ∂M. Let $\mathcal{T} = \{\sigma\}$ be triangulation of M extending a triangulation $\mathcal{T}_\partial$ of ∂M. Let $\{\mathcal{U}(\sigma)\}_{\sigma \in \mathcal{T}}$ be a finite collection of closed n–discs in a thickening $\mathcal{O}p(M)$.*

*We say that $(\mathcal{T}, \{\mathcal{U}(\sigma)\})$ is **adapted** to $(\mathcal{W}, \partial M)$ if the following hold:*

- *the triangulation $\mathcal{T}_\partial$ and the covering $\{\mathcal{U}(\sigma) \cap \partial M\}_{\sigma \in \mathcal{T}_\partial}$ are adapted to $\mathcal{W}|_{\partial M}$,*
- *the triangulation $\mathcal{T} \setminus \mathcal{T}_\partial$ is in general position with respect to $\mathcal{W}$,*
- *the covering $\{\mathcal{U}(\sigma)\}_{\sigma \in \mathcal{T}}$ satisfies the properties of Definition 3.12.*

It any case that the proof of Theorem 3.9 readily yields:

Corollary 3.13. *Let M be an n–dimensional manifold, $n \geq 2$, possibly with boundary. Let $\mathcal{W}$ be a line field in M that is tangent to ∂M. Then, there exist a triangulation $\mathcal{T}$ and a cover $\{\mathcal{U}(\sigma)\}$ that are adapted to $(\mathcal{W}, \partial M)$. Additionally, $\mathcal{T}$ can be taken to be arbitrarily fine.*

3.3 Convex curves in the plane and the sphere

As stated at the beginning of the chapter, we will need a better understanding of the space of convex curves in the sphere. The main reference for this topic is the very excellent paper [70] by N. Saldanha. A lot of what we explain can be found in much greater detail in his work, so we will content ourselves with reviewing the results that are relevant to us. When dealing with matters of local convexity, it is often be fruitful to pass to $\mathbb{R}^2$ using an affine mapping to simplify computations. For this reason, we also include a discussion of convex curves in the plane.

This section, much like the previous one, can be regarded as an appendix containing technical results. However, the discussion is so central to the arguments of latter sections that we very much encourage the reader not to skip it. At the same time, it is completely safe not to go through the proofs of the more technical lemmas as long as their geometrical meaning is understood.

3.3.1 Immersed curves

Before we focus on convex curves, it is worth recalling some essential facts about immersions. Given any smooth manifold N, and any smooth 1–manifold I, we write $\mathcal{I}(I, N)$ for the space of immersions of I into N. A reasonable topology for it is the C^k–topology, with $k \in \{1, 2, 3, \ldots, \infty\}$. One can show that all such choices yield homotopically equivalent topological spaces [8]. For simplicity, we will just take the C^∞–topology. Often, we are interested in spaces of immersions that are fixed at the ends. If $\gamma : I \to N$ is an immersion with I an interval, we shall write $\mathcal{I}(I, N; \gamma)$ for the space of immersions that agree with γ at the ends (as germs).

Given an immersion γ, we can associate to it its **Frenet map** $\Gamma_\gamma = (\gamma, \gamma') \in \mathcal{FI}(I, N)$, the formal immersion corresponding to it. If $\Gamma \in \mathcal{FI}(I, N)$ is some formal immersion, we write

$\mathcal{FI}(I, N; \Gamma)$ for the space of formal immersions agreeing with Γ at their ends. We will often abuse notation and say that $\Gamma \in \mathcal{I}(I, N)$ if $\Gamma = (\gamma, \gamma')$ for some $\gamma \in \mathcal{I}(I, N)$.

Let us state (a particular instance of) the Smale–Hirsch theorem:

Proposition 3.14. *Let N be a smooth manifold of dimension at least 2. Let K be a compact CW–complex, possibly with boundary. Suppose that we are given a K–family of formal immersions $\Gamma_k = (\gamma_k, F_k) \in \mathcal{FI}(I, N)$ satisfying $F_k(t) = \gamma_k'(t)$ if $k \in \partial K$ or $t \in \mathcal{O}p(\{0,1\})$.*

Then, there is a homotopy $\Gamma_{k,s} = (\gamma_{k,s}, F_{k,s}) \in \mathcal{FI}(I, N)$, $k \in K$, $s \in [0,1]$, of formal immersions satisfying:

- *$\Gamma_{k,1} \in \mathcal{I}(I, N)$,*
- *$\Gamma_{k,s}(t) = \Gamma_k(t)$ if $k \in \partial K$, $t \in \mathcal{O}p(\{0,1\})$, or $s = 0$,*
- *$\gamma_{k,s}(t)$ is C^0–close to $\gamma_k(t)$.*

That is, immersions satisfy an h–principle that is relative with respect to the domain and the parameter and also C^0–close. The main consequence of the proposition is that the Frenet map

$$\Gamma : \mathcal{I}(I, N; \gamma) \to \mathcal{FI}(I, N; \Gamma_\gamma),$$

$$\eta \to \Gamma_\eta,$$

induces a *(weak) homotopy equivalence*. This is a much more general version of the Whitney–Graunstein theorem that we stated in the Preamble.

3.3.1.1 Immersions into the sphere and the plane

For the rest of the section, we will particularise $N = \mathbb{R}^2, \mathbb{S}^2$ and we will try to understand the homotopy type of the spaces of *convex* curves. If $N = \mathbb{R}^2$, it is immediate that the space of formal immersions is homotopy equivalent to $\mathcal{M}\mathrm{aps}(I, \mathbb{S}^1)$: i.e. all the homotopy information is contained in the **Gauss map**

$$\mathcal{G}_\gamma(t) =\to \frac{\gamma'(t)}{|\gamma'(t)|}.$$

We will sometimes think of the Frenet map as the pair $(\gamma, \mathcal{G}_\gamma)$.

Similarly, the space of formal immersions if $N = \mathbb{S}^2$ is readily seen to be homotopy equivalent to $\mathcal{M}\mathrm{aps}(I, \mathrm{SO}(3))$. Here the correspondence is given by normalisation too:

$$(\gamma, \gamma') \to (\gamma, \dot\gamma/|\dot\gamma|, \mathfrak{n}),$$

with $\mathfrak{n} : I \to \mathbb{S}^2$ the only vector that completes the basis to a matrix in SO(3). For this reason, we will usually think of the Frenet map as the triple of maps:

$$t \to (\gamma(t), \dot\gamma(t)/|\dot\gamma(t)|, \mathfrak{n}(t)).$$

3.3.2 Convex curves in the plane. Basic notions

For completeness, let us recall the definition: an immersion $\gamma : [0,1] \to \mathbb{R}^2$ is said to be **convex** if, for all t, $\{\gamma'(t), \gamma''(t)\}$ is a positive basis for $T_{\gamma(t)}\mathbb{R}^2$. This is, of course, equivalent to the fact that the line determined by $\gamma(t_0)$ and $\gamma'(t_0)$ lies, locally and as t increases from t_0, to the right of $\gamma(t)$. If $\{\gamma'(t), \gamma''(t)\}$ is instead a negative basis, we say that the curve is **concave**. A time $t \in [0,1]$ or its image $\gamma(t)$ are said to be an **inflection point** if γ is neither convex nor concave at t.

Disclaimer 3.15. What we call "convex" is often called in the literature "locally convex". Convexity is instead a global property amounting to the fact that no geodesic tangent to the curve at a point intersects it at some other point. To avoid repeating "locally" over and over, we have opted to just do without it.

Write $\mathcal{L}([0,1], \mathbb{R}^2; \gamma)$ for the space of convex curves $[0,1] \to \mathbb{R}^2$ that agree with some fixed convex curve $\gamma : [0,1] \to \mathbb{R}^2$ at their ends. Endowing it with the C^∞–topology, it can be regarded as a subspace of $\mathcal{I}([0,1], \mathbb{R}^2; \gamma)$. The Frenet/Gauss functional can be restricted to $\mathcal{L}([0,1], \mathbb{R}^2; \gamma)$. Convexity implies that the Gauss map is necessarily an immersion preserving the standard orientations. This has the following consequence:

Proposition 3.16 (Pohl)**.** *The space $\mathcal{L}([0,1], \mathbb{R}^2; \gamma)$ is homotopy equivalent to the natural numbers $\mathbb{N}$.*

Proof. Consider the covering map

$$\begin{aligned} \mathbb{R} &\to \mathbb{S}^1 \\ s &\to e^{2\pi s i}. \end{aligned}$$

Given any $\Gamma = (\eta, F) \in \mathcal{FI}([0,1], \mathbb{R}^2; \gamma)$, the lift of $F(1)$ minus the lift $F(0)$ only depends on $[\Gamma] \in \pi_0$. This defines a map $\pi_0(\mathcal{FI}([0,1], \mathbb{R}^2; \gamma)) \to \mathbb{R}$ whose image is $c + \mathbb{Z}$. Let $c \in \mathbb{R}$ be the smallest positive number satisfying this. Subtracting c provides an isomorphism

$$\mathrm{wind} : \pi_0(\mathcal{FI}([0,1], \mathbb{R}^2; \gamma)) \to \mathbb{Z}.$$

If γ is a convex curve, we have $\mathrm{wind}([\gamma]) \geq 0$, because the Gauss map $\mathcal{G}_\gamma$ is immersed and orientation preserving. Then, only those components of $\mathcal{FI}([0,1], \mathbb{R}^2; \gamma)$ identified with $\mathbb{N}$ can contain convex curves.

Any curve η in $\mathcal{L}([0,1], \mathbb{R}^2; \gamma)$ can be parametrised by the angle that $\eta'(t)$ makes with the positive x–axis, because this is an strictly increasing function. Since $\mathrm{Diff}^+([0,1])$ is contractible, we can first deformation retract $\mathcal{L}([0,1], \mathbb{R}^2; \gamma)$ to the subspace $\mathcal{L}$ where all curves are indeed parametrised by angle. Let $A \subset \mathcal{L}$ be a connected component.

Let $\eta_0, \eta_1 \in A \subset \mathcal{L}$: they are convex, parametrised by angle, and formally homotopic. The curves $\eta_s = (1-s)\eta_0 + s\eta_1 \in A$ are also parametrised by angle and are therefore convex. This provides a deformation retraction of A onto $\{\eta_0\}$. We conclude that

$$\mathrm{wind} : \mathcal{L}([0,1], \mathbb{R}^2; \gamma) \to \mathbb{N}$$

is a homotopy equivalence. □

Variations of Proposition 3.16 will be useful whenever we want to deform families of convex curves; this is particularly relevant because affine charts map convex curves in the sphere to convex curves in the plane (see Lemma 3.22).

3.3.3 Convex curves in the plane. Families with varying boundary conditions

In Proposition 3.16 we have shown that, having fixed some boundary conditions, the space of convex curves $[0,1] \to \mathbb{R}^2$ satisfying said conditions is a countable collection of contractible components. However, the families of convex curves we will usually encounter will have varying boundary conditions that, nonetheless, we want to preserve under deformations. Our aim in this subsection is to introduce some technical results to deal with such families.

Consider the following problem: we are given two points $p_0, p_1 \in \mathbb{R}^2$ and two directions $v_0, v_1 \in \mathbb{R}^2 \setminus \{0\}$. Can we connect p_0 with p_1 with a convex curve $\gamma : [0,1] \to \mathbb{R}^2$ such that $\gamma'(0) = v_0$ and $\gamma'(1) = v_1$? We invite the readers to convince themselves that the problem is always solvable (even parametrically) if we allow the Gauss map of the solution to describe arbitrarily many turns. However, if we impose bounds on the turning number:

Proposition 3.17. *Fix two compact families of germs of convex curves*

$$\gamma_s : \mathcal{O}p(\{0\}) \subset [0,1] \to \mathbb{R}^2,$$

$$\eta_s : \mathcal{O}p(\{1\}) \subset [0,1] \to \mathbb{R}^2.$$

with $s \in K$ a compact manifold. Assume that $\{\mathcal{G}_{\gamma_s}(0), \mathcal{G}_{\eta_s}(1)\}$ is a positive basis for $\mathbb{R}^2$.

Then, there is a family of convex curves $\phi_s : [0,1] \to \mathbb{R}^2$ satisfying:

- *$\phi_s = \gamma_s$ in $\mathcal{O}p(\{0\})$,*
- *$\phi_s = \eta_s$ in $\mathcal{O}p(\{1\})$,*
- *$\mathcal{G}_{\phi_s} : [0,1] \to \mathbb{S}^1$ is embedded for each s,*

if and only if

$$\eta_s(1) \text{ lies in the interior of the (open) cone } \gamma_s(0) + (\mathcal{G}_{\gamma_s}(0), \mathcal{G}_{\eta_s}(1))\mathbb{R}^+.$$

We should clarify the notation: $(\mathcal{G}_{\gamma_s}(0), \mathcal{G}_{\eta_s}(1)) \subset \mathbb{S}^1$ denotes the interval obtained by going counterclockwise (i.e. positively) from $\mathcal{G}_{\gamma_s}(0)$ to $\mathcal{G}_{\eta_s}(1)$. The proposition says that, if $(\mathcal{G}_{\gamma_s}(0), \mathcal{G}_{\eta_s}(1))$ is smaller than a hemisphere, the problem is solvable if and only if $\eta_s(1)$ can be reached from $\gamma_s(0)$ by following directions contained in it. See Figure 3.3.

Proof. Let us prove the if direction since the only if one is clear.

Let us start by proving a simplified graphical case. Fix a constant $C > 0$. Suppose that we have two families of convex functions $f_s : [-C-1, -C] \to \mathbb{R}$, $g_s : [0,1] \to \mathbb{R}$ such that $f_s'(-C) = 0 < g_s'(0)$ and $f_s(-C) = 0$. If $g_s(0)$ is sufficiently close to 0, we claim that there is $h_s : [-C-1, 1] \to \mathbb{R}$ convex satisfying $h_s|_{[-C-1,-C]} = f_s$ and $h_s|_{[0,1]} = g_s$.

Indeed, assume that $g_s(0)/C < g_s'(0)$, this is exactly the cone condition stated for functions. Define $H_s' : [-C-1, 1] \to \mathbb{R}$ to be strictly increasing and satisfying $H_s'|_{[-C-1,-C]} = f_s'$, $H_s'|_{[0,1]} = g_s'$. Its integral H_s, which we define to agree with f_s in $[-C-1, C]$, will not in general agree with g_s in $[0,1]$, they will differ by some t–independent constant. Consider, however, a reparametrisation $\phi_s : [-C-1, 1] \to [-C-1, 1]$ which is the identity in $[-C-1, C] \cup [0,1]$. We will construct $h_s = H_s \circ \phi_s$ by choosing ϕ_s suitably. If we take $\delta > 0$ small, $H_s'(-C+\delta)$ and $H_s'(-\delta)$ are close

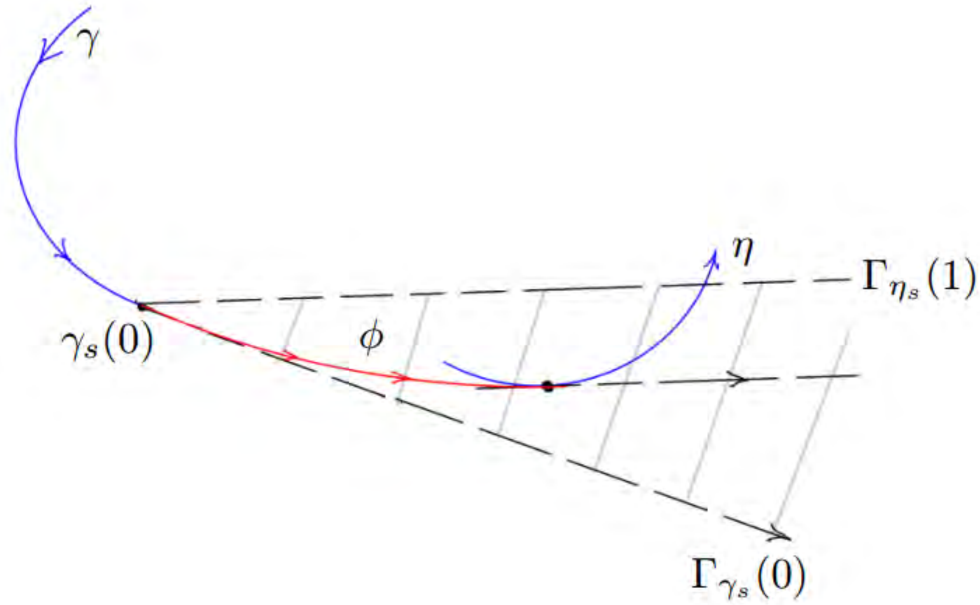

Figure 3.3: The cone condition in Proposition 3.17.

to $f_s'(-C) = 0$ and $g_s'(0)$, respectively. This implies that if we make $\phi_s^{-1}([-C, -C+\delta])$ large, $H_s|_{[0,1]} < g_s$ and if we make $\phi_s^{-1}([-\delta, 0])$ large, $H_s|_{[0,1]} > g_s$. Then, there is some middle-ground choice of ϕ_s guaranteeing $H_s|_{[0,1]} = g_s$.

Since $\{\mathcal{G}_{\gamma_s}(0), \mathcal{G}_{\eta_s}(1)\}$ is a positive basis of $\mathbb{R}^2$, we can find an s–depending family of affine transformations that takes us to the graphical setting and the claim follows. $\square$

If, instead, one has to describe between half and a full turn to reach $\mathcal{G}_{\eta_s}(1)$ from $\mathcal{G}_{\gamma_s}(0)$, the problem we posed is always solvable:

Proposition 3.18. *Fix two families of germs of convex curves, $s \in K$, with K a compact manifold:*

$$\gamma_s : \mathcal{O}p(0) \subset [0,1] \to \mathbb{R}^2,$$
$$\eta_s : \mathcal{O}p(1) \subset [0,1] \to \mathbb{R}^2.$$

Assume that $\{\mathcal{G}_{\gamma_s}(0), \mathcal{G}_{\eta_s}(1)\}$ is a negative basis for $\mathbb{R}^2$.

Then, there is a family of convex curves $\phi_s : [0,1] \to \mathbb{R}^2$ satisfying:

- *$\phi_s = \gamma_s$ in $\mathcal{O}p(\{0\})$,*
- *$\phi_s = \eta_s$ in $\mathcal{O}p(\{1\})$,*
- *$\mathcal{G}_{\phi_s} : [0,1] \to \mathbb{S}^1$ is embedded for each s.*

For a pictorial depiction, refer to Figure 3.4.

Proof. Observe that the points $\eta_s(1)$, $s \in K$, lie in a compact set A of $\mathbb{R}^2$. This implies that, for $C > 0$ large enough, the half–space determined by the point $C \cdot \gamma_s'(0) + \gamma_s(0)$ and the directions $(-\mathcal{G}_{\eta_s}(1), \mathcal{G}_{\eta_s}(1)) \subset \mathbb{S}^1$, contains A.

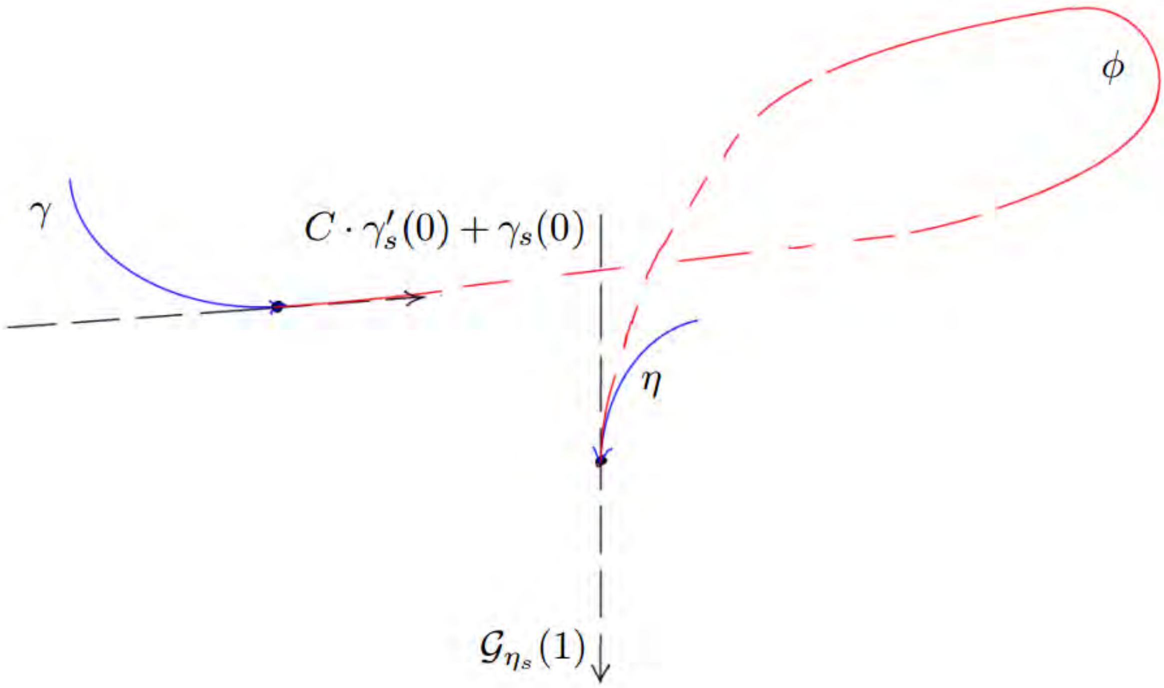

Figure 3.4: The elements involved in the proof of Proposition 3.18.

This allows us to do the following: we trace a convex arc agreeing with γ_s in $\mathcal{O}p(\{0\})$ and having endpoint arbitrarily close to $C \cdot \gamma_s'(0) + \gamma_s(0)$. Then, we turn to so that the tangent vector points in the direction of $-\mathcal{G}_{\eta_s}(1)$. By Proposition 3.17 it is possible to connect this curve with η_s. $\square$

3.3.4 Convex curves in the plane. Glueing, flattening, and stretching

In Subsections 3.6.1 and 3.6.2, we will solve the extension problem by constructing Engel structures in terms of curves that pass from being convex to being tangent to a maximal circle that describes a contact structure. As such, we need three technical results. The first allows us to smoothly transition from being tangent to a maximal circle to being convex; this is achieved by a "glueing" procedure. The second will allow us to "flatten" a point where a convex curve is tangent to a maximal circle into an ∞–order point of contact. The third "stretches" this ∞–order point of contact into an arbitrarily large segment along which both curves coincide. We phrase the lemmas for plane curves, since it is simpler to do it in this setting.

3.3.4.1 The glueing lemma

Lemma 3.19. *Let K be a compact manifold. Let $L \subset \mathbb{R}^2$ be the x–axis. Consider a smooth family of curves $\gamma_s : [0,1] \to \mathbb{R}^2$, $s \in K$, which are either convex or reparametrisations of a segment of L. Assume further that $\gamma_s(0) = 0$ is a tangency with L.*

Then, for any $\varepsilon \in \mathbb{R}^+$ small enough, there is a smooth family of immersions $\eta_s : [0,1] \to \mathbb{R}^2$, $s \in K$, satisfying:

I. $\|\eta_s - \gamma_s\|_{C^1} \leq \varepsilon$ *and* $\eta_s|_{[\varepsilon,1]} = \gamma_s|_{[\varepsilon,1]}$,

II. $\eta_s(0) = (-\varepsilon, 0)$ *is a tangency with* L,

IIIa. *If the curve* γ_s *is convex, the curve* η_s *is convex for* $t \in (0,1]$ *and* $\eta_s(0)$ *is an* ∞*–order tangency with* L.

IIIb. *If the curve* γ_s *is a reparametrisation of a segment of* L, *so is the curve* η_s.

That is, by slightly pushing the tangency, we are able to glue the family with the equator at $\gamma_s(0)$. This lemma readily follows from Proposition 3.17, but it can be easily proven independently:

Proof. There is an interval $[0, 2\varepsilon]$ in which the functions γ_s are graphical over L. Denote by f_s the functions depending on x that describe them.

Construct an increasing cut–off function $\chi : [-\varepsilon, 2\varepsilon] \to [0, 2\varepsilon]$ satisfying:

$$\chi^{(k)}(-\varepsilon) = 0 \text{ for } k \in \mathbb{N}, \quad \chi''|_{[-\varepsilon,\varepsilon)} > 0, \text{ and } \chi(t)|_{[\varepsilon,2\varepsilon]} = t.$$

Since the composition of increasing functions is convex as soon as one of them is strictly convex and the other is (non–strictly) convex, the family $f_s \circ \chi(t)$ defines a family of convex curves

$$\eta_s : [-\varepsilon, 2\varepsilon] \to \mathbb{R}$$

that glues with $\gamma_s|_{[\varepsilon,1]}$. We can then reparametrise to obtain the family claimed in the statement. □

3.3.4.2 The flattening lemma

Proceeding pretty much like in the previous subsubsection, we can smoothly flatten a given point in a convex curve to achieve an ∞–order of contact with respect to an equator. This is the content of the following lemma:

Lemma 3.20. *Let* K *be a compact manifold. Let* $L \subset \mathbb{R}^2$ *be the* x*–axis. Consider a smooth family of curves* $\gamma_s : [-1,1] \to \mathbb{R}^2$, $s \in K$, *which are convex and have* $\gamma_s(0) = 0$ *as a tangency with* L.

For any $\varepsilon \in \mathbb{R}^+$ *small enough, the family* γ_s *can be extended to a* $K \times [0,1]$*–family parametrised by* (s,l) *satisfying:*

I. $\gamma_{(s,0)} = \gamma_s$.

II. $|\gamma_{(s,l)} - \gamma_s|_{C^1} \leq \varepsilon$ *and* $\gamma_{(s,l)}|_{[-1,-\varepsilon]\cup[\varepsilon,1]} = \gamma_s|_{[-1,-\varepsilon]\cup[\varepsilon,1]}$,

III. $\gamma_{(s,l)}(0) = 0$ *is a tangency with* L,

IV. *for* $l \in [0,1)$, *the curves* $\gamma_{(s,l)}$ *are convex.* $\gamma_{(s,1)}$ *is convex for* $t \in [-1,0) \cup (0,1]$ *and has an* ∞*–order of contact with* L *at* $t = 0$.

Proof. By assumption, there is $\varepsilon > 0$ small such that γ_s is graphical over L if $t \in [-\varepsilon, \varepsilon]$. We shall abuse notation and regard it as a function of x with domain $[-\varepsilon, \varepsilon]$. Then, since γ_s is convex and vanishing at at the origin, there are constants $c_0, c_1 \in \mathbb{R}^+$ such that

$$0 < c_0 \leq \gamma_s''(t), \quad 0 \leq \|\gamma_s'(t)\| \leq c_1 \quad \forall t \in [-\varepsilon, \varepsilon].$$

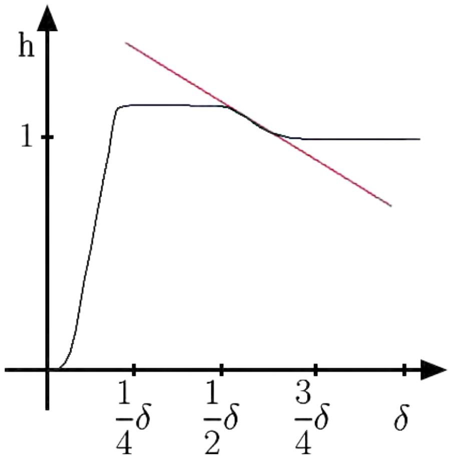

Figure 3.5: Graph of h.

Given a smooth function $g : [-\varepsilon, \varepsilon] \to [-\varepsilon, \varepsilon]$, a condition for $F = \gamma_s \circ g$ to be convex is the differential inequality

$$F'' = (\gamma_s'' \circ g)(g')^2 + (\gamma_s' \circ g)g'' > 0.$$

Having bounds as the one above in terms of constants c_0, c_1, it is sufficient that g satisfies the inequality

$$c_0(g')^2 - c_1|g''| > 0. \tag{3.3}$$

Consider a function $h : [-\varepsilon, \varepsilon] \to [0, 1]$ such that:

a. $h(-t) = h(t)$ and $h^{(k)}(0) = 0$ for $k \in \mathbb{N}$,

b. $\int_0^\varepsilon h(t)dt = \delta$ and $h|_{[3\varepsilon/4,1]} = 1$,

c. $h'|_{(0,\varepsilon/4)} > 0$ and $h'|_{[\varepsilon/4,\varepsilon/2]} = 0$,

d. $c_0 > c_1|h'| \geq 0$ in $[\varepsilon/2, 3\varepsilon/4]$.

See Figure 3.5 for a pictorial description.

The desired family is constructed by linear interpolation:

$$g_l(t) = \int_0^t [(1-l) + lh(t)]dt, \quad l \in [0, 1], \quad t \in (-\varepsilon, \varepsilon),$$

$$\gamma_{(s,l)} = \gamma_s \circ g_l.$$

Property (b.) allows for the function $\gamma_{(s,l)}$ to be defined in $t \in [-1, -\varepsilon) \cup (\varepsilon, 1]$ to be simply γ_s. Property (c.) implies that the resulting curves are convex in $(0, \varepsilon/2)$. Property (c.) implies that the inequality 3.3 holds in $[\varepsilon/2, 3\varepsilon/4]$. Property (a.) ensures convexity in $[-1, 0)$ by symmetry and gives the ∞–order tangency in 0. □

3.3.4.3 The stretching lemma

The following lemma concerns the stretching of a flattened point into a segment, the details of the proof are left to the reader.

Lemma 3.21. *Let K be a compact manifold. Let $L \subset \mathbb{R}^2$ be the x–axis. Consider a smooth family of curves $\gamma_s : [-1,1] \to \mathbb{R}^2$, $s \in K$, which are convex and have $\gamma_s(0) = 0$ as an ∞–order tangency with L.*

Given any number $C > 0$, the family γ_s can be extended to a $K \times [0, C]$–family of curves parametrised by (s,l) with domains $[-1, 1+l]$ that satisfy:

- $\gamma_{(s,l)}(t) = \gamma_s(t)$ *if* $t \in [-1, 0]$,
- $\gamma_{(s,l)}(t) = \gamma_s(t-l) + (l, 0)$ *if* $t \in [l, 1+l]$,
- $\gamma_{(s,l)}$ *is a reparametrisation of the segment* $[0,l] \times \{0\} \subset L$ *if* $t \in [0,l]$.

3.3.5 Convex curves in the sphere. Basic notions

Let $\gamma : \mathbb{S}^1 \to \mathbb{S}^2$ be a curve in the 2–sphere. We say that $\gamma(t)$ is an **inflection point** for γ if the maximal circle tangent to γ at time t locally divides γ in two parts. A curve $\gamma : \mathbb{S}^1 \to \mathbb{S}^2$ having no inflection points has an associated Frenet map $\Gamma_\gamma : \mathbb{S}^1 \to \mathrm{O}(3)$ given at t by the matrix

$$(\gamma(t), \dot{\gamma}(t)/|\dot{\gamma}(t)|, \mathfrak{n}(t)),$$

with $\mathfrak{n} : \mathbb{S}^1 \to \mathbb{S}^2$ satisfying $\langle \ddot{\gamma}(t), \mathfrak{n}(t) \rangle > 0$. The curve γ is said to be **convex** if Γ_γ has positive determinant and thus it lives in $\mathrm{SO}(3)$; then the definition agrees with the one given for immersions. We say it is **concave** if the Frenet map has negative determinant.

The following lemma gives a correspondence between convex curves in the plane and the sphere, allowing us to use the machinery built in the previous subsections:

Lemma 3.22. *The affine chart*

$$\Phi : H^2 = \{(x,y,z) \in \mathbb{S}^2 : x > 0\} \longrightarrow \mathbb{R}^2$$

$$\Phi(x,y,z) = (y/x, z/x)$$

maps geodesics to geodesics, inflection points to inflection points, and convex curves to convex curves. So does any other affine chart.

Proof. The map Φ is readily seen to preserve geodesics from the correspondence between geodesics and planes passing through the origin. Convex curves are also preserved because convexity can be defined in terms of the order of contact with the corresponding geodesics (i.e. through inflection points). $\square$

The main question to tackle in what remains of this section is what the homotopy type of the space of convex curves in $\mathbb{S}^2$ is. We shall phrase it (as expected) in the language of the h–principle, much like we did in the case of immersions. To avoid cluttering the notation, let us make some simplifications.

Let $\mathcal{I}^f$ be the space of smooth immersions of $\mathbb{S}^1$ into $\mathbb{S}^2$ (here, f stands for "free"). We denote by $\mathcal{L}^f \subset \mathcal{I}^f$ the subspace of convex curves. $\mathcal{I}^f$ is a subspace itself of the space of formal immersions $\mathcal{FI}^f$ (we think of them as maps $\mathbb{S}^1 \to \mathrm{SO}(3)$).

Write $\mathcal{I}$, $\mathcal{L}$, and $\mathcal{FI}$, respectively, for the subspaces of curves additionally satisfying $\Gamma_\gamma(1) = \mathrm{Id}$. Taking Frenet maps at $t = 1$ gives a projection $\mathcal{L}^f \to \mathrm{SO}(3)$ whose fibre over the identity matrix is $\mathcal{L}$. The action of $\mathrm{SO}(3)$ on $\mathcal{L}^f$ identifies all the fibres with $\mathcal{L}$ and shows that $\mathcal{L}^f = \mathrm{SO}(3) \times \mathcal{L}$. An analogous statement holds for $\mathcal{I}^f$ and $\mathcal{FI}^f$. Having this simple product structure means that, henceforth, we shall focus on the topology of $\mathcal{FI}$, $\mathcal{I}$, and $\mathcal{L}$.

3.3.5.1 Revisiting immersions in $\mathbb{S}^2$

The Hirsch–Smale theorem says that the inclusions $\mathcal{I}^f \to \mathcal{FI}^f$ and $\mathcal{I} \to \mathcal{FI}$ are weak homotopy equivalences. At the same time, $\mathcal{FI}^f$ can be understood to be the space of formal convex curves as well. Indeed, if we decouple the position, velocity, and acceleration of a convex curve, which is all the relevant formal data, we obtain, up to homotopy, a map into $\mathrm{SO}(3)$.

What are the homotopy groups of $\mathcal{FI}$? Let us compute:

$$\pi_k(\mathcal{FI}) = \pi_k(\Omega(\mathrm{SO}(3))) = \pi_{k+1}(\mathrm{SO}(3)) = \pi_{k+1}(\mathbb{S}^3) \text{ if } k > 0,$$

$$\pi_0(\mathcal{FI}) = \pi_1(\mathrm{SO}(3)) = \mathbb{Z}_2.$$

Here $\Omega(\mathrm{SO}(3))$ denotes the (based) loop space of $\mathrm{SO}(3)$.

Let $\alpha_j : \mathbb{S}^1 \to \mathbb{S}^2$, $j \in \{1, 2, \dots\}$, be the convex curve $e^{it} \to \Phi^{-1}(e^{itj} + i)$, where Φ is the affine chart described in Lemma 3.22. It is a curve in $\mathcal{L}$ that describes j little loops. The computation that we just did shows that α_j is homotopic to $\alpha_{j\pm 2}$ as an element in $\mathcal{I}$.

3.3.5.2 Adding loops

Let us explain analytically the following straightforward geometric operation: given a family of curves, we want to cut them at a point and add there a little convex *wiggle*.

Let $\alpha_j \in \mathcal{L}$, $j \in \{1, 2, \dots\}$, be a convex curve describing j loops in an affine chart. Let K be a compact parameter space (possibly a manifold). Let $\gamma_k : [0, 1] \to \mathbb{S}^2$, $k \in K$, be a family of *immersions*. Let $t_0 \in [0, 1]$. Denote $A_k = \Gamma_{\gamma_k}(t_0) \in \mathrm{SO}(3)$. Since $\Gamma_{A_k \circ \alpha_j}(1) = A_k = \Gamma_{\gamma_k}(t_0)$, we can concatenate the curves $\gamma_k|_{[0,t_0]}$, $A_k \circ \alpha_j$, and $\gamma_k|_{[t_0,1]}$, smoothing them at the concatenation points. From Proposition 3.16, it is obvious that any two families produced like this are homotopic (where the homotopy leaves the curves fixed away from the cutting point), so we denote any such family by $\gamma_k^{[j\#t_0]}$. [1]

It is clear that the same construction can be done if γ_k has domain any other 1–dimensional manifold I or if its target is $\mathbb{R}^2$ instead of $\mathbb{S}^2$. In the plane, the resulting curves lie in different connected components as we add new loops. In the sphere, this operation, if done once, changes the connected component of the family; if done twice, the family is homotopic to the original one:

[1] This notation was originally introduced by Saldanha in [70], but he writes the point and the number of loops in the opposite order. It seems more natural to me to write "j loops *at* the point t_0", and hence I have chosen to follow this convention instead. I apologise for this.

Proposition 3.23. *Let $\gamma_k : I \to \mathbb{S}^2$ be a family of immersions. The curves γ_k and $\gamma_k^{[j\#t_0]}$ are homotopic as families of immersions if and only if j is even.*

For a pictorial description of the homotopy, see Figure 3.6. Since the process can be iterated over different points $\{t_0, \dots, t_n\}$, we will write $\gamma_k^{[j_0\#t_0,\dots,j_n\#t_n]}$ if we do so.

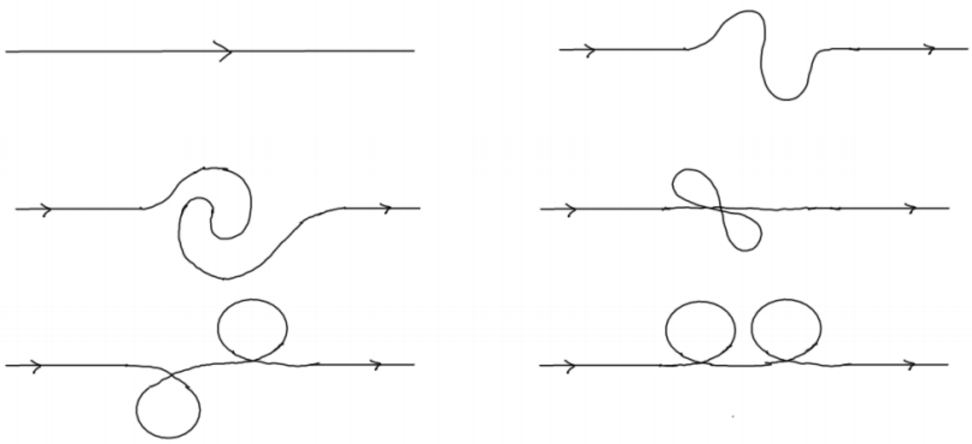

Figure 3.6: A homotopy through immersions between a curve and the same curve with two wiggles added. In the last step, one of the wiggles has to be moved around the sphere to make it appear in the other side.

3.3.6 Convex curves in the sphere. The theorems of Little and Saldanha

The following classic result, due to Little [50], describes the connected components of $\mathcal{L}$:

Proposition 3.24. *Let $\alpha_j \in \mathcal{L}$, $j > 0$, be a little convex curve in $\mathbb{S}^2$ that, in an affine chart, has Gauss map of degree j. The curve α_j is homotopic, through curves in $\mathcal{L}$, to α_{j+2} if and only if $j > 1$. In particular, $\mathcal{L}$ has three connected components, represented by the curves α_j, $j = 1, 2, 3$.*

Write $\mathcal{L}_j \subset \mathcal{L}$ for the connected component containing α_j. See Figure 3.7 for a explicit homotopy connecting α_2 with α_4 through convex curves.

Proposition 3.24 states that the inclusion $\mathcal{L} \to \mathcal{I}$ is not a weak homotopy equivalence: this already fails at the level of π_0. However, this failure for the h–principle to hold for the inclusion $\mathcal{L} \to \mathcal{FI}$ happens in higher homotopy groups as well. This is worked out by Saldanha in [70], where he computes the homotopy groups of all the $\mathcal{L}_j$. Let us review some of the notions introduced in [70].

Within $\mathcal{L}$, there are submanifolds $\mathcal{M}_m$ of codimension $2m$, $m \geq 1$, defined as follows: a convex curve γ belongs to $\mathcal{M}_m$ if there are points $t_1, ..., t_m = 1 \in \mathbb{S}^1$ such that $\Gamma_\gamma(t_i) = \mathrm{Id}$ and the arcs between the t_i belong to $\mathcal{L}_1$. This implies that $\mathcal{M}_1 = \mathcal{L}_1$, $\mathcal{M}_m \subset \mathcal{L}_2$ if m is even, and $\mathcal{M}_m \subset \mathcal{L}_2$ if m is odd. Then, Saldanha's result can be phrased as follows:

Proposition 3.25 ([70][Lemma 10.1)**.** *] There are subspaces $\mathcal{Y}_j \subset \mathcal{L}_j$, $j = 2, 3$, satisfying:*

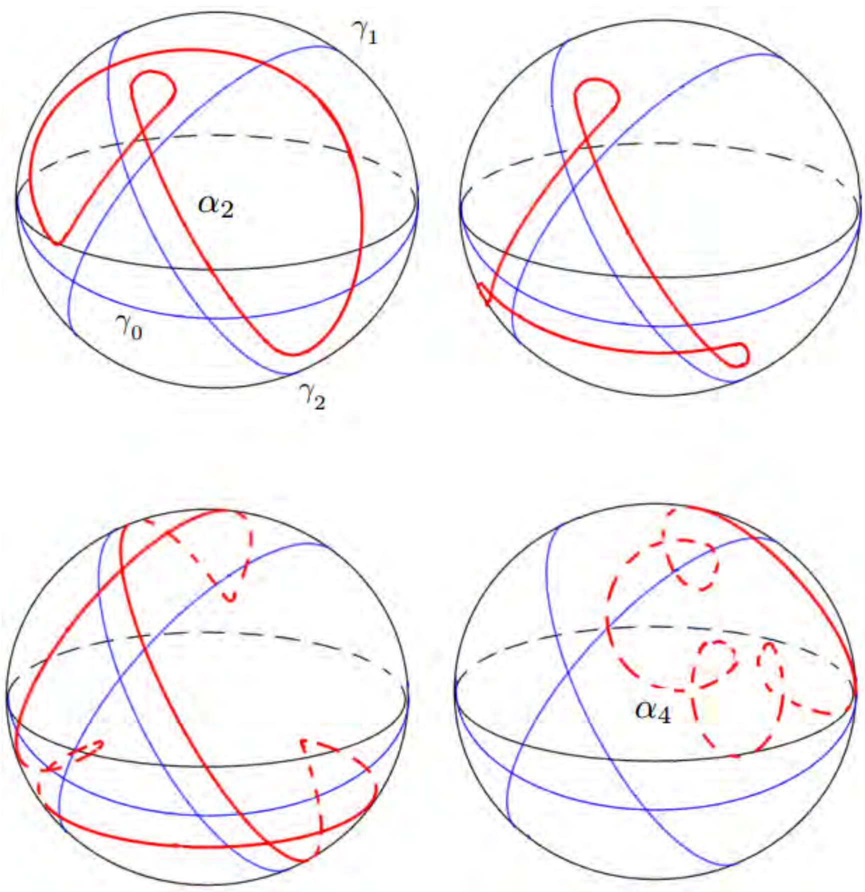

Figure 3.7: The blue lines correspond to maximal circles. The first figure is a convex curve that, in the frontal hemisphere, has winding number two. By pushing the upper strand down, it can be taken to the second figure. It is convex because it is comprised of three segments that are slight push-offs of the blue equators whose corners have been rounded to preserve convexity. The same is true for the third step, which is obtained from the second by following the equators for a longer time. Then, we push everything to the opposite hemisphere, yielding a curve that in said hemisphere has winding number four.

- *$\mathcal{Y}_2 \cup \mathcal{Y}_3 \to \mathcal{I}$ is a deformation retract,*
- *a deformation retract of $\mathcal{L}_2$ is obtained from $\mathcal{Y}_2$ by attaching discs $\mathbb{D}^{4i+2}$, $i = 0, 1, 2, ..,$ along contractible spheres,*
- *a deformation retract of $\mathcal{L}_3$ is obtained from $\mathcal{Y}_3$ by attaching discs $\mathbb{D}^{4i}$, $i = 1, 2, ..,$ along contractible spheres.*

The subspaces $\mathcal{Y}_j$ can be set to be $\mathcal{Y}_j = \mathcal{L}_j \setminus (\cup_m \mathcal{M}_m)$.

It was known classically that $\mathcal{L}_1$ is contractible.

3.3.7 Convex curves in the sphere. Loose maps

Proposition 3.25 should be understood as an h–principle statement. It says that a certain subclass $\mathcal{Y}_2 \cup \mathcal{Y}_3 \subset \mathcal{L}$ of convex curves satisfies a complete h–principle (possibly one would be tempted to say that this is the *overtwisted* class). Even though this class is only meaningfully defined up to homotopy (although we do explicitly describe one representative), one can pose the question of whether a given map into $\mathcal{L}$ is homotopic to a map with image in this flexible class. Such a map will be called **loose**. Proposition 3.25 can be reworded by saying:

Corollary 3.26 (Saldanha)**.** *Let K be a compact set. A continuous map $f : K \to \mathcal{L}$ that is disjoint from each of the $\mathcal{M}_m$ is loose.*

That is, any two maps $f_0, f_1 : K \to \mathcal{L} \setminus (\cup_m \mathcal{M}_m)$ are homotopic as maps into $\mathcal{L} \setminus (\cup_m \mathcal{M}_m)$ if and only if they are homotopic as maps into $\mathcal{I}$ or $\mathcal{FI}$.

Loose maps can also be characterised in a more geometric fashion. Given a map $f : K \to \mathcal{L}$ and a point $t \in \mathbb{S}^1$, recall that we can define the map $f^{[1\#t]} : K \to \mathcal{I}$. It is easy to check that it can be done so that $f^{[1\#t]} : K \to \mathcal{L}$ is still a family of convex curves. Write $f^{[i_0\#t_0,\dots,i_n\#t_n]}$ for the result of adding $i_j > 0$ loops at the point t_j.

It is immediate that the choice of point in which a loop is added is not important: $f^{[i_0\#t_0,\dots,i_n\#t_n]}$ is homotopic, through convex curves, to $f^{[i_0+\cdots+i_n\#1]}$. This is true because we can simply "slide" the basepoints t_i to the point $t = 1$ and this yields the desired homotopy, effectively moving the newly added loops along the curve. We then claim that the following lemma holds, which is a slight improvement of Proposition 3.24:

Lemma 3.27. *Let $f : K \to \mathcal{L}$. Then, $f^{[i_0\#t_0,\dots,i_n\#t_n]}$ is homotopic through convex curves to either $f^{[1\#1]}$ or $f^{[2\#1]}$ depending on whether $i_0 + \dots i_n > 0$ is odd or even.*

Proof. By the previous discussion, $f^{[i_0\#t_0,\dots,i_n\#t_n]}$ is homotopic to $f^{[i_0+\cdots+i_n\#1]}$. The latter family is obtained from f by cutting it at the point $1 \in \mathbb{S}^1$ and concatenating with the curve $\alpha_{i_0+\cdots+i_n}$ (a curve describing $i_0 + \cdots + i_n$ wiggles in an affine chart).

If $i_0 + \cdots + i_n > 3$, we can apply Little's homotopy (Proposition 3.24) to $\alpha_{i_0+\cdots+i_n}$ relative to the ends, and show that $f^{[i_0+\cdots+i_n\#1]}$ is homotopic to $f^{[i_0+\cdots+i_n-2\#1]}$. If $i_0 + \cdots + i_n = 3$, then it is left as an exercise to the reader to prove that the process of Proposition 3.24 (or, equivalently, Figure 3.7), can be done as well by using a little portion of the original curve. $\square$

The geometric intuition now is that, as soon as some extra convexity is added (by introducing additional loops), all possible homotopies through immersions can be approximated by homotopies through convex curves:

Lemma 3.28 (Saldanha)**.** *Let $f : K \to \mathcal{L}$ be homotopic to $f^{[2\#1]}$. Then, f is loose.*

3.3.7.1 Adding wiggles to achieve convexity

To finish the discussion, we shall prove an statement that is somewhat analogous to (our phrasing of) the Smale–Hirsch theorem (Proposition 3.14). The reader can take it as a first step towards understanding Saldanha's work. It will be used later when we solve the Engel extension problem. For a picture, see Figure 3.8.

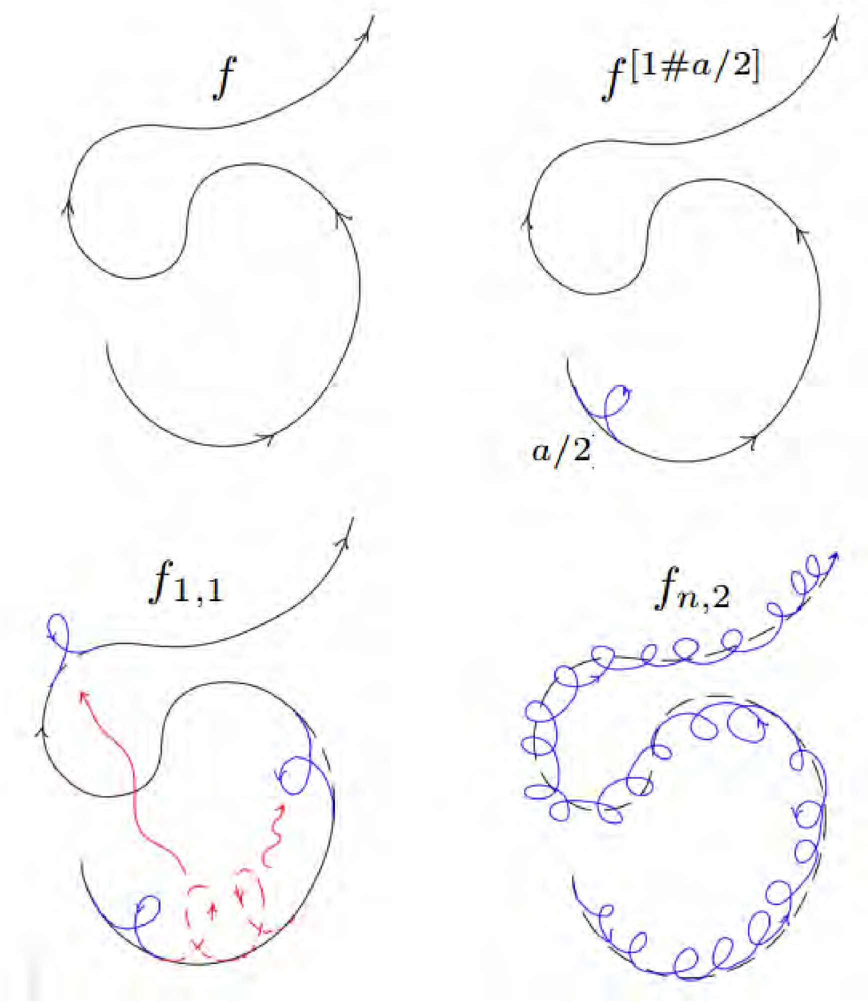

Figure 3.8: A homotopy where, starting from one wiggle, we use Little's homotopy to create many more that we then distribute along the curve to achieve convexity everywhere. We use the notation from Proposition 3.29.

Proposition 3.29. *Let K be a manifold, possibly non–compact and with boundary. Fix $A \subset K$ some closed submanifold. Let $a \in (0,1)$. Let $f : K \to \mathcal{I}([0,1], \mathbb{S}^2)$ be a family of immersions satisfying:*

- *for all k, $f(k)|_{[0,a]\cup\{1\}}$ is convex,*
- *$f(A) \subset \mathcal{L}([0,1], \mathbb{S}^2)$.*

Let $I = [0,1] \setminus \mathcal{O}p(\{0,1\})$. Then, there is a family $f : K \times [0,\infty) \to \mathcal{I}([0,1], \mathbb{S}^2)$ such that:

- *$f(k,0) = f^{[1\#a/2]}(k)$,*
- *$f(k,s)(t)$ is convex if $k \in \mathcal{O}p(A)$,*
- *$f(k,s)(t) = f^{[1\#a/2]}(t)$ if $t \in \mathcal{O}p(\{0,1\})$,*

- *for s large enough and independent of k, $f(k,s)$ is everywhere convex,*
- *the length of $f(k,s)$ is bounded independently of k and s; its curvature in I goes to infinity, independently of k, as $s \to \infty$.*

Proof. Let $f_n : K \to \mathcal{I}([0,1], \mathbb{S}^2)$, $n \in [0,\infty)$, be a family of immersions with $f_n = f^{[1+2n\#a/2]}$ if n is an integer. That is, Little's homotopy between $f^{[1+2n\#a/2]}$ and $f^{[3+2n\#a/2]}$ is performed in the interval $(n, n+1)$ (see Proposition 3.24 and Lemma 3.27). We have $f_0 = f$.

Now we further homotope the family f_n to distribute the wiggles that we are creating uniformly in I. Define $f_{n,r} : K \to \mathcal{I}([0,1], \mathbb{S}^2)$, $n \in [0,\infty)$, $r \in [0,1]$, satisfying:

- $f_{n,0} = f_n$,
- if $n \in [0,1]$, $f_{n,r} = f_n$,
- let $m > 0$ be an integer. Then, for $n \in [m, m+1]$, we have
$$f_{n,1} = f_{n-m}^{[1\#t_0,\dots,1\#t_{2m-1}]}$$
where $t_j : [m, m+1] \to I$ are continuous functions of n whose images are (approximately) evenly spaced.

That such a homotopy exists is a consequence of the discussion prior to Lemma 3.27: we simply slide the new wiggles as they appear, so they are distributed uniformly over the domain.

Fix $k \in K$. Then, there exists N large enough such that, for $n > N$ and all $j \in \{0, \cdots, 2\lfloor n-1 \rfloor\}$, the segment $f_{n,1}|_{[t_j,t_{j+1}]}$ can be made convex by borrowing some of the convexity of the wiggles at its endpoints. This can be packaged as a smooth function $N : K \to [0,\infty)$. Define $f_{n,r} : K \to \mathcal{I}([0,1], \mathbb{S}^2)$, $r \in [1,2]$, $n \geq N(k)$, to be a homotopy with $f_{n,2} : K \to \mathcal{L}([0,1], \mathbb{S}^2)$ doing precisely this.

We now set for the wiggles we introduce to have diameter $O(1/n)$. Then, the curvature behaves like $O(n^2)$ and the length like $O(1)$. □

3.4 Engel shells

Consider coordinates $(x,y,z;t)$ in the cartesian product $\mathbb{D}^3 \times [0,1]$. We endow it with the *standard orientation* and the *Euclidean metric* as a subspace of $\mathbb{R}^4$. In Section 1.2 we discussed Engel structures in $\mathbb{D}^3 \times [0,1]$ containing the vertical direction $\langle \partial_t \rangle$. We shall now study formal Engel structures in $\mathbb{D}^3 \times [0,1]$ containing $\langle \partial_t \rangle$ that, additionally, are Engel in $\partial(\mathbb{D}^3 \times [0,1])$. For ease of notation, we will call this notion an *Engel shell*; they conform the natural framework in which we will phrase the three extension methods.

Definition 3.30. *An* ***Engel shell*** *is a 2–plane distribution $\mathcal{D}$ on the 4–cell $\mathbb{D}^3 \times [0,1]$ conforming to the following properties:*

1. *$\mathcal{D} = \langle \partial_t, X \rangle$, where X is a unitary vector field tangent to the level sets $\mathbb{D}^3 \times \{t\}$,*
2. *in a neighbourhood $\mathcal{O}p(\partial(\mathbb{D}^3 \times [0,1]))$ of the boundary, $\mathcal{D}$ is an Engel structure.*

*If $\mathcal{D}$ is Engel everywhere, the Engel shell is said to be **solid**.*

*Let M be a 4–manifold with $\mathcal{D}$ a 2–plane distribution. A set $\mathcal{U} \subset M$ is said to be an Engel shell if there is a map $\phi : \mathcal{U} \to \mathbb{D}^3 \times [0,1]$ so that $\phi_*\mathcal{D}$ is an Engel shell. The map ϕ is called a **trivialising chart**.*

Remark 3.31. Any *homotopy of Engel shells* is assumed to be *relative to the boundary* unless explicitly stated otherwise.

It is possibly unclear a priori how the formal data in the interior is being encoded. The following lemma states that, given any Engel shell, there is a unique extension (up to homotopy) of the formal data in the boundary (given by the fact that $\mathcal{D}$ is Engel there) to the whole of $\mathbb{D}^3 \times [0,1]$.

Lemma 3.32. *Let $\mathcal{D}$ be an Engel shell. Then, there is a formal Engel structure $(\mathcal{W}, \mathcal{D}, \mathcal{E})$ that in $\mathcal{O}p(\partial(\mathbb{D}^3 \times [0,1]))$ satisfies $\mathcal{E} = [\mathcal{D}, \mathcal{D}]$ and $\mathcal{W}$ is its kernel. This formal Engel structure is unique, up to homotopy relative to the boundary.*

Proof. The (orientable) line field $\mathcal{W}$ yields a section of $\mathcal{D}$ in $\mathcal{O}p(\partial(\mathbb{D}^3 \times [0,1]))$. Since $\mathcal{D}$ is the trivial 2–dimensional real bundle, this section extends to the interior and this extension is unique up to homotopy. Similarly, $\mathcal{E}/\mathcal{D}$ can be thought of as a section of $T(\mathbb{D}^3 \times [0,1])/\mathcal{D}$ and one can reason analogously.

Consider the bundle isomorphism $\det(\mathcal{D}) \cong \mathcal{E}/\mathcal{D}$. Any line field over $\mathbb{D}^3 \times [0,1]$ that is orientable over the boundary, is globally orientable. Over the boundary, the isomorphism is realised because $\mathcal{D}$ is Engel, so we obtain a global isomorphism. The same applies to $\det(\mathcal{E}/\mathcal{W}) \cong TM/\mathcal{E}$. $\square$

In particular, Lemma 3.32 says that, when we tackle the extension problem, we only have to worry about homotoping the 2–distribution so that it becomes Engel, because the resulting formal Engel structure will have the correct homotopy type regardless of any choice we make.

In the next two subsections we shall define two particularly relevant examples of Engel shells.

3.4.0.1 X as a family of curves in $\mathbb{S}^2$

We recall one last piece of notation: as explained in Section 1.2, the vector field X can be understood as a $\mathbb{D}^3$–family of curves. Indeed, for each $p \in \mathbb{D}^3$, the unit tangent bundle of $\mathbb{D}^3 \times \{t\}$ at every point (p,t) can be identified with $\mathbb{S}^2$. For this, a *t–invariant metric* has to be chosen in $\mathbb{D}^3 \times [0,1]$ and, for simplicity and unless stated otherwise, we will take the Euclidean one, as stated above.

A particularly important consequence of this is that it is not immediate to show that certain statements about the curves X_p still hold in a different but overlapping Engel shell (living in the same ambient 4–manifold). Later on, in Subsections 3.5.1 and 3.5.2, we will deal with this issue.

In any case, for a given Engel shell we have the family of curves:

$$X_p : [0,1] \to \mathbb{S}^2,$$

$$X_p(t) = X(p,t) \in \mathbb{S}T_{p,t}(\mathbb{D}^3 \times \{t\}) \cong \mathbb{S}^2.$$

The one result that it is important to have in mind before reading the rest of this section is Proposition 1.12 about how Engelness can be characterised in terms of these curves. Both the statement and its proof are key to our arguments.

3.4.1 Engel shells of contact type

Definition 3.33. *Let $\mathcal{D}$ be an Engel shell. $\mathcal{D}$ is said to be of* ***contact–type*** *if the following properties hold:*

a. *$\mathcal{D}$ is contained in the 3–distribution $\xi \oplus \langle \partial_t \rangle$, with ξ a t–invariant contact structure on the level sets $\mathbb{D}^3 \times \{t\}$,*

b. *in $\mathcal{O}p(\partial(\mathbb{D}^3 \times [0,1]))$, the even-contact structure $[\mathcal{D}, \mathcal{D}]$ is $\xi \oplus \langle \partial_t \rangle$ and $\langle \partial_t \rangle$ is its kernel.*

In particular, the underlying contact structure ξ is part of the data of the contact–type shell.

The reduction process of Subsection 3.5.4 will create Engel shells of contact-type in each 4–ball. The extension procedures of Subsections 3.6.1 and 3.6.2 use Engel shells of contact-type as their starting data.

Remark: Our discussion about how the formal structure extends to the interior means that the extension can be taken to be $\mathcal{E} = \xi \oplus \langle \partial_t \rangle$ and $\mathcal{W} = \langle \partial_t \rangle$.

In an Engel shell of contact-type, $\mathcal{D}$ is constrained to live in $\xi \oplus \langle \partial_t \rangle$. In particular, this means that the line field X is tangent to the contact structure ξ. As such, the functions X_p, that normally take values in $\mathbb{S}^2$, instead take values in $\mathbb{S}(\xi) \cong \mathbb{S}^1$. This allows us to describe such a shell in terms of a $\mathbb{D}^3$–family of real–valued functions that describe how X is turning with respect to a framing of ξ.

Let us be more precise. Consider the euclidean metric in $\mathbb{D}^3 \times [0,1]$. The distribution $\xi \oplus \langle \partial_t \rangle$ agrees with $\mathcal{E}$ in $\mathcal{O}p(\partial(\mathbb{D}^3 \times [0,1]))$; since the latter is oriented, so is the first (not only at the boundary, but everywhere). We can now fix a orthonormal Legendrian frame $\{Y, Z\}$ for the contact structure $(\mathbb{D}^3, \xi)$ so that $\{\partial_t, Y, Z\}$ is a positive frame for $\xi \oplus \langle \partial_t \rangle$. Then, the following formula

$$X(p,t) = \cos(c(p,t))Y + \sin(c(p,t))Z, \tag{3.4}$$

assigns to each contact–type shell a real–valued function $c : \mathbb{D}^3 \times [0,1] \longrightarrow \mathbb{R}$ that is uniquely defined up to shifting by 2π. The function $c = c(\mathcal{D})$ defined by Equation 3.4 is called the **angular function**

Remark 3.34. Whenever $\mathcal{D}$ is Engel, the angular function is the lift to $\mathbb{R}$ of the *developing map* of the $\mathcal{W}$–orbit $\{p\} \times [0,1]$. See 1.5.3.1.

The orientation conventions imply the following fact:

Lemma 3.35. *The contact–type Engel shell $\mathcal{D}$ is an Engel structure at the point (p,t) if and only if $\partial_t c(\mathcal{D})(p,t) \neq 0$.*

Additionally, $\xi \oplus \langle \partial_t \rangle = [\mathcal{D}, \mathcal{D}]$ at (p,t) as oriented bundles if and only if $\partial_t c(\mathcal{D})(p,t) > 0$.

In particular, the differential inequality $\partial_t c(\mathcal{D}) > 0$ always holds on a neighbourhood $\mathcal{O}p(\partial(\mathbb{D}^3 \times [0,1]))$. Conversely, suppose that a contact structure ξ in $\mathbb{D}^3$ and a frame $\{Y, Z\}$ for it have been fixed; then, any function $c : \mathbb{D}^3 \times [0,1] \longrightarrow \mathbb{R}$ satisfying $\partial_t c(\mathcal{D}) > 0$ on a neighbourhood $\mathcal{O}p(\partial(\mathbb{D}^3 \times [0,1]))$ is the angular function of some uniquely defined contact–type Engel shell $\mathcal{D}(c)$. We conclude that there is a bijective correspondence between angular functions up to shifting by 2π and contact–type shells (once a *framed contact structure* has been fixed).

Contractibility of the space of real functions relative to the boundary implies that:

Lemma 3.36. *Fix a contact structure ξ in $\mathbb{D}^3$ with a legendrian frame $\{Y, Z\}$. Let $c_1, c_2 : \mathbb{D}^3 \times [0,1] \longrightarrow \mathbb{R}$ be two angular functions agreeing on $\mathcal{O}p(\partial(\mathbb{D}^3 \times [0,1]))$. Then, the corresponding contact–type shells $\mathcal{D}(c_1)$ and $\mathcal{D}(c_2)$ are homotopic relative to the boundary as contact–type shells.*

One relevant instance of this is when one of the angular functions is strictly increasing, and thus defines a solid Engel shell:

Lemma 3.37. *Let $\mathcal{D}$ be a contact–type Engel shell with c its angular function. $\mathcal{D}$ is homotopic, relative to the boundary, to a solid contact–type Engel shell if and only if $c(p,1) > c(p,0)$ for all $p \in \mathbb{D}^3$.*

Proof. If $c(p,1) > c(p,0)$ holds everywhere, then c is homotopic, relative to the boundary, to a function $\tilde{c} : \mathbb{D}^3 \times [0,1] \to \mathbb{R}$ satisfying $\partial_t \tilde{c} > 0$; then Lemma 3.36 yields the claim. Otherwise, Bolzano's theorem says that we cannot find a function $\tilde{c}$ that is everywhere increasing and agrees with c at the boundary, and therefore the problem is not solvable within the category of contact–type shells. ☐

We are therefore forced to solve the extension problem using Engel shells that are not necessarily of contact-type. Let us define one last concept that will be relevant:

Definition 3.38. *Let $\mathcal{D}$ be a contact–type Engel shell. The **height** of $\mathcal{D}$ is the largest non–negative integer* height$(\mathcal{D})$ *satisfying:*

$$\min_{p \in \partial \mathbb{D}^3} \frac{c(p,1) - c(p,0)}{\pi} > \mathrm{height}(\mathcal{D}).$$

Do note the strict inequality in the definition. Geometrically, the height of a contact–type shell measures how many projective turns the legendrian vector field $X(p,1)$ has rotated with respect to $X(p,0)$ (for those p in the *boundary* of $\mathbb{D}^3$). The extension processes of Subsections 3.6.1 and 3.6.2 require as input contact–type shells of sufficiently large height.

Proposition 3.39. *Let $\mathcal{D}$ be a contact–type shell. Its height is invariant under reparametrisation of $\mathbb{D}^3$ and modification of the framing $\{Y, Z\}$.*

3.4.2 Engel shells of convex type

The following type of shell is the starting point for the extension method explained in Subsection 3.6.3.

Definition 3.40. *Let $\mathcal{D}$ be an Engel shell. $\mathcal{D}$ is said to be of **convex-type** if the following properties hold:*

a. the curves X_p are immersions for all $p \in \mathbb{D}^3$,

b. the curves X_p are convex at time t whenever $(p,t) \in \mathcal{O}p(\partial(\mathbb{D}^3 \times [0,1]))$.

We shall see that a convex–type shell can (under suitable assumptions) be made solid through convex–type shells. This can be found in Subsection 3.6.3. Before that, we should introduce a notion that is somehow analogous to the *height* we defined for contact–type ones.

3.4.2.1 Winding

The winding of a convex–type shell measures whether the curves X_p are very *wiggly*. A practical way of measuring this is to pass to an affine chart and measure the how many turns the Gauss map is describing there. Let us formalise this idea.

Let $\Psi_p : H_p \to \mathbb{R}^2$, $p \in \partial\mathbb{D}^3$, be a smooth family of affine charts such that the hemisphere H_p contains the vector $X_p(0)$. Then, one can consider the segment γ_p of $\Psi_p \circ X_p$ that contains the initial point $\Psi_p \circ X_p(0)$.

Let $\rho_p \in (0,1)$, $p \in \partial\mathbb{D}^3$, be a smooth family of constants depending on p such that the domain of γ_p contains $[0,\rho_p]$. Given the convex curves:

$$\gamma_p : [0,\rho_p] \to \mathbb{R}^2$$

we can consider their Gauss map

$$\mathcal{G}_{\gamma_p} = \frac{\gamma_p'}{||\gamma_p'||} : [0,\rho_p] \to \mathbb{S}^1,$$

$$\mathcal{G}_{\gamma_p}(s) = e^{w_{\gamma_p}(s)i}.$$

Definition 3.41. *Let $\gamma : [0,\rho] \to \mathbb{R}^2$ be a convex curve. Then, the* ***winding*** *of γ is the non–negative real number:*

$$\mathrm{wind}(\gamma) = \frac{w_\gamma(\rho) - w_\gamma(0)}{\pi}.$$

Let $\mathcal{D}$ be a convex–type shell. The ***winding*** *of $\mathcal{D}$ is the largest non–negative integer* $\mathrm{wind}(\mathcal{D})$ *satisfying:*

$$\sup_{\Psi_p,\rho_p} \min_{p \in \partial\mathbb{D}^3} \frac{w_{\gamma_p}(\rho_p) - w_{\gamma_p}(0)}{\pi} > \mathrm{wind}(\mathcal{D}).$$

where the supremum is taken over all the possible choices of Ψ_p and ρ_p.

Remark 3.42. The winding is simply measuring the minimum amount of half–turns, over all $p \in \partial\mathbb{D}^3$, that γ_p' makes. The curves γ_p themselves depend on which hemisphere we choose, and for this reason we take supremum over all possible choices.

For all practical uses, we will check winding with Φ_p satisfying $\gamma_p(0) = (0,0)$ and $\mathcal{G}_{\gamma_p}(0) = (1,0)$. Such a Φ_p is not uniquely defined but, since we are assuming that $\mathbb{D}^3 \times [0,1]$ is endowed with the Euclidean metric, we could set for Φ_p to preserve orthogonality at the point $X_p(0)$.

We saw already in Section 3.3 that convex curves become flexible/loose if a wiggle is added to them. We shall see (Subsection 3.6.3) that having large winding will immediately imply that the curves X_p, $p \in \partial\mathbb{D}^3$, are loose, which will allow us to solve the extension problem.

3.4.3 The hierarchy of shells

Definition 3.43. *Let $\mathcal{D}_0$ and $\mathcal{D}_1$ be two shells. We say that $\mathcal{D}_1$* ***dominates*** *$\mathcal{D}_0$ if there is an embedding $\phi : \mathbb{D}^3 \times [0,1] \to \mathbb{D}^3 \times [0,1]$ satisfying:*

- *there is a deformation of $\mathcal{D}_1$, relative to the boundary, such that $\phi^*\mathcal{D}_1 = \mathcal{D}_0$,*

- *this deformation can be chosen to be Engel in* $\mathbb{D}^3 \times [0,1] \setminus \mathrm{Image}(\phi)$.

That is, $\mathcal{D}_1$ dominates $\mathcal{D}_0$ if solving the extension problem for $\mathcal{D}_0$ automatically solves the extension problem for $\mathcal{D}_1$. The idea of this subsection is to show that, when we study the extension problem, we can restrict to shells that have some additional structure that make them easy to work with.

3.4.3.1 Domination between contact–type shells

In the particular case of contact–type shells, we have the following straightforward proposition:

Proposition 3.44. *Let ξ be a contact structure in $\mathbb{D}^3$ whose frame $\{Y, Z\}$ we fix. Let $\mathcal{D}(c_0)$ and $\mathcal{D}(c_1)$ be two contact–type shells such that, for all $p \in \mathbb{D}^3$, satisfy:*

$$c_0(p,0) < c_1(p,0) \text{ and } c_0(p,1) > c_1(p,1).$$

*Then, $\mathcal{D}(c_0)$ **dominates** $\mathcal{D}(c_1)$.*

Proof. First we can find a contact embedding $\phi : (\mathbb{D}^3, \xi) \to (\mathbb{D}^3, \xi)$ that slightly shrinks $\mathbb{D}^3$, i.e. $\phi^{-1}(\partial\mathbb{D}^3) = \emptyset$. It can be assumed that ϕ satisfies $c_0(\phi(p), 0) < c_1(p, 0)$ and $c_0(\phi(p), 1) > c_1(p, 1)$.

Using Lemma 3.36, we can deform c_0 to be strictly increasing in the intervals

$$\mathcal{O}p([0, h_0(p)]) \cup \mathcal{O}p([h_1(p), 1]),$$

where $h_0, h_1 : \phi(\mathbb{D}^3) \to [0,1]$ are smooth functions satisfying $c_0(\phi(p), h_i(p)) = c_1(p, i)$, $i = 0, 1$. This allows us, if we adjust c_0 suitably, to ensure that the embedding

$$\Phi : \mathcal{O}p(\partial\mathcal{D}(c_1)) \to \mathcal{D}(c_0)$$

$$\Phi(p, i) = (\phi(p), h_i(p))$$

satisfies $\Phi^* c_0 = c_1$. Extending Φ to the interior of $\mathcal{D}(c_1)$ arbitrarily, $\mathcal{D}(c_0)$ can be further homotoped within $\Phi(\mathcal{D}(c_1))$, relative to its boundary, to achieve $\Phi^* c_0 = c_1$ everywhere. The claim follows. ☐

Remark 3.45. One of the key assumptions in Proposition 3.44 is that the two shells we are comparing have the same underlying framed contact structure. This raises the following question: what if we allow for reparametrisations of $\mathbb{D}^3$ that preserve the contact structure? For instance: consider an Engel shell $\mathcal{D}(c)$ with $c(p,1) < c(p,0)$, for some p. This implies that the extension problem cannot be solved by applying Lemma 3.37; geometrically, the legendrian vector field X is turning negatively as we move in the interval $p \times [0,1]$. However, there might be a contactomorphism ϕ of $\mathbb{D}^3$ such that $(\phi^* X)(p, 1)$ turns positively with respect to $X(p, 0)$ for all p, allowing us to apply Lemma 3.37 after this reparametrisation. The underlying question is: to what extent do (compactly supported) contactomorphisms of $\mathbb{D}^3$ preserve the *partial order* in the space of legendrian vector fields (that are fixed at the boundary)? It seems very hard to give an interesting answer.

We can now introduce a particular type of contact–type shell that is well suited for performing the extension processes of Subsections 3.6.1 and 3.6.2. Roughly speaking, it is a shell where there is a region $t \in [0, \rho]$, $\rho > 0$ small, in which the angular function is strictly increasing and does *not* depend on p. As such, the regions $\mathbb{D}^3 \times [0, \rho]$ and $\mathbb{D}^3 \times [\rho, 1]$ are shells themselves, and the first one is actually solid. Our aim (in Section 3.6) is to modify the former (through solid shells that are not necessarily of contact–type), modifying in particular the germ along its upper boundary, so that the angular function of the latter shell fits in the hypothesis of Lemma 3.37.

Definition 3.46. *Let C be a positive integer. A function $c : \mathbb{D}^3 \times [0,1] \longrightarrow \mathbb{R}$ is said to be a C–**angular function** if there is $\rho \in (0,1)$ such that c is strictly increasing in $[0,\rho]$ and*

$$c(p,\rho) - c(p,0) > C\pi \text{ for all } p.$$

*The contact–type shell $\mathcal{D}(c)$ associated to a C–angular function c is said to be a C–**contact–type shell**.*

The most relevant use of Proposition 3.44 is the following corollary.

Corollary 3.47. *The contact–type shell $\mathcal{D}$ dominates a* height$(\mathcal{D})$*–contact–type shell.*

Let us spell out what this is saying: given a shell of height height$(\mathcal{D})$, which is a quantity that measures the turning of the legendrian frame X along the boundary of $\mathbb{D}^3$, it dominates a second shell where every curve X_p describes height$(\mathcal{D})$ turns concentrated in the band $\mathbb{D}^3 \times [0,\rho]$.

3.4.3.2 Domination between convex–type shells

Our aim now is proving an statement that is analogous to Corollary 3.47: given a convex–type shell of winding wind$(\mathcal{D})$, it dominates a convex–type shell where the curves $X_p|_{[0,\rho]}$ have controlled winding for all $p \in \mathbb{D}^3$ and not just those p in the boundary.

Definition 3.48. *Let C be a positive integer. A convex–type shell $\mathcal{D}$ is said to be of C–**convex–type** if there is a constant $\rho \in (0,1)$, a family of convex curves $f_p : [0,\rho] \to \mathbb{S}^2$, such that $X_p|_{[0,\rho]} = f_p^{[C\#\rho/2]}$.*

Recall that $f_p^{[C\#\rho/2]}$ is a curve obtained from f_p by cutting it at time $\rho/2$ and adding C little loops; see Subsubsection 3.3.7.

Remark 3.49. Observe that the curves X_p in a C–convex–type shell describe C little wiggles. This implies that their Gauss maps, in an affine chart centered at $X_p(t_0)$, describe $2C$ *projective turns*. Thus, a C–convex–type shell has winding at least $2C$. We apologise for this potentially confusing notation.

The following is the main result about domination among convex–type shells:

Proposition 3.50. *A convex–type shell $\mathcal{D}$ with* wind$(\mathcal{D}) \geq 2C+1 \geq 3$ *dominates a C–convex–type shell.*

Proof. By hypothesis, there are $\mathbb{D}^3$–families of affine charts $\Phi_p : H_p \to \mathbb{R}^2$, constants $\rho_p \in (0,1)$, and curves $\gamma_p : [0,\rho_p] \to \mathbb{R}^2$ satisfying:

- $\gamma_p(t) = \Phi_p \circ X_p(t)$ if $t \in [0,\rho_p]$,
- the curves γ_p are convex for t small,
- wind$(\gamma_p) > 2C+1$ if $p \in \mathcal{O}p(\partial\mathbb{D}^3)$,

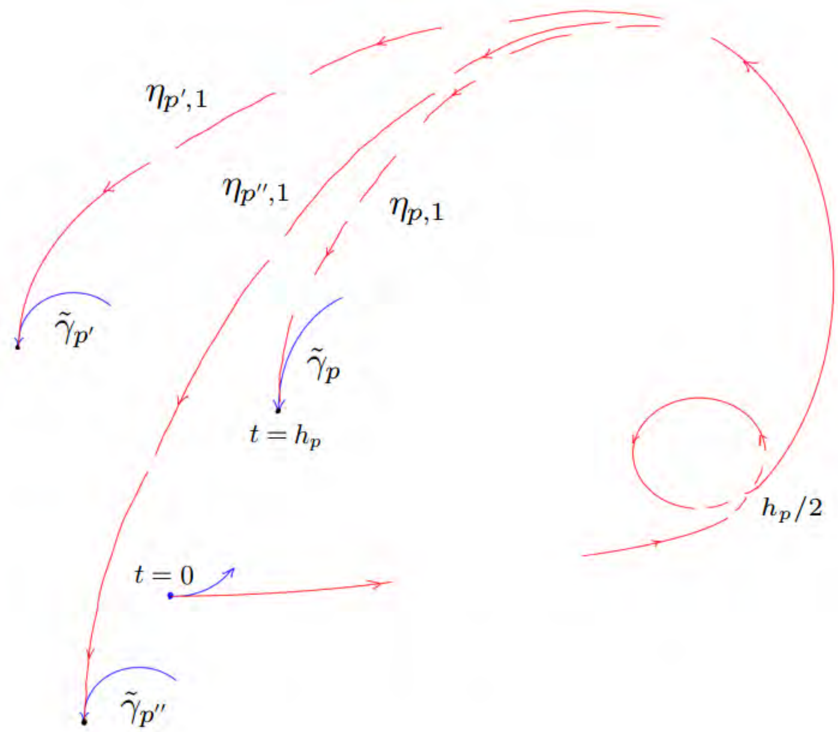

Figure 3.9: Given a family of planar curves all of which have winding at least 3, it is possible to homotope all of them to have a little wiggle.

We will homotope them, relative to $(p, t) \in \mathcal{O}p(\partial(\mathbb{D}^3 \times [0, 1]))$, so that *all of them* have winding greater than $2C + 1$ in the band $A = \{(p, t) \mid t \in [0, \rho_p]\}$. Then, an application of Proposition 3.18 and a few remarks will yield the result.

Consider the tangent vectors γ'_p. In the regions where γ_p is convex, γ'_p turns counterclockwise around the origin; in particular, for $p \in \mathcal{O}p(\partial\mathbb{D}^3)$, it describes (at least) $2C+1$ projective turns. Regarding γ'_p as a formal derivative, homotope it to a $\mathbb{D}^3$–family of curves $\tilde{\gamma}'_p : [0, \rho_p] \to \mathbb{R}^2 \setminus \{0\}$ such that $\tilde{\gamma}'_p/|\tilde{\gamma}'_p|$ is an immersion that describes exactly $2C + 1 + \varepsilon$ turns in the segment $[0, h_p]$, with $h_p \in (0, \rho_p)$ some smooth family, and $\varepsilon > 0$ some very small constant. This can be done relative to ∂A. Denote $B = \{(p, t) \mid t \in [0, h_p]\}$.

We can integrate $\tilde{\gamma}'_p$ to yield a map $\tilde{\gamma}_p : [0, \rho_p] \to \mathbb{R}^2$ where all the curves have winding $2C + 1$ at least. We can set $\tilde{\gamma}_p(t) = \gamma_p(t)$ if $t \in \mathcal{O}p(\{0\})$ but, in general, they do not agree in the region $t \in \mathcal{O}p(\{\rho_p\})$. Although this can be adjusted by hand, there is the following high-tech solution. Define a formal immersion that agrees with γ_p in $\mathcal{O}p(\{\rho_p\})$ and with $\tilde{\gamma}_p$ in $[0, h_p]$. This is possible by construction. Then the Smale–Hirsch theorem provides an immersion satisfying this as well. Since $\tilde{\gamma}_p$ and γ_p agree by construction in $p \in \mathcal{O}p(\partial\mathbb{D}^3)$, all the process is relative to the boundary region ∂A.

Replacing γ_p by $\tilde{\gamma}_p$ yields a first deformation $\tilde{\mathcal{D}}$, rel. boundary, of $\mathcal{D}$. Now our aim it to further homotopy $\tilde{\mathcal{D}}$ so that it is of C–convex–type (see Figure 3.9). Applying Proposition 3.18 to the family $\tilde{\gamma}_p|_{[0,h_p]}$, $p \in \mathbb{D}^3$, we deduce that there is a homotopy $\eta_{p,l} : [0, h_p] \to \mathbb{R}^2$, $l \in [0, 1]$, and a family of convex curves $F_p : [0, h_p] \to \mathbb{R}^2$, such that:

- $\eta_{p,0} = \tilde{\gamma}_p$,
- $\eta_{p,l}(t) = \tilde{\gamma}_p(t)$ if $t \in \mathcal{O}p(\{0, h_p\})$ or $p \in \mathcal{O}p(\partial\mathbb{D}^3)$,
- $\eta_{p,1} = F_p^{[C\#h_p/2]}$.

Then, consider a bump function $\chi : B \to [0,1]$ such that $\chi = 0$ in $\mathcal{O}p(\partial B)$ and $\chi = 1$ outside of a slightly larger neighbourhood. Then, the family $\eta_{p,\chi(p)}$ can be used to replace the family $\tilde{\gamma}_p$ in the band B. The resulting shell is homotopic to $\mathcal{D}$, for any $\delta > 0$ small $(\mathbb{D}^3_{1-\delta}, \mathcal{D})$ is of C–convex–type, and outside of this region it is a honest Engel structure. $\square$

3.4.4 Flowboxes and orientations

A shell is the central object of the extension problem: it describes a germ of Engel structure along $\partial\mathbb{D}^3$ and a formal extension to the interior (in the case of contact–type and convex–type shells, this formal information has already been "polished" to have a certain particular form). However, during the reduction process, what we deal with are formal Engel structures in $\mathbb{D}^3 \times [0,1]$ that we have to make Engel in some parts of the interior, "pushing the problems" to the boundary of the model. The following definitions encapsulate this idea.

Definition 3.51. *A **contact–type flowbox** is a formal Engel structure $(\mathcal{W}, \mathcal{D}, \mathcal{E})$ on the 4–cell $M = \mathbb{D}^3 \times [0,1]$ conforming to the following properties:*

- *$\mathcal{E}$ is even–contact with $\mathcal{E} = \xi \oplus \langle\partial_t\rangle$, where ξ a t–invariant levelwise contact structure,*
- *$\mathcal{W} = \langle\partial_t\rangle$ is the kernel of the even–contact structure $\mathcal{E}$,*
- *the isomorphism $\det(\mathcal{E}/\mathcal{W}) \cong TM/\mathcal{E}$ of the formal Engel structure is realised geometrically through Lie bracket by $\mathcal{E}$ being even–contact.*

The last property simply requires for the bundle isomorphism of Equation (1.2) to be compatible with the fact that $\mathcal{E}$ is even–contact.

Note that the orientation of $\mathcal{E} \cong \xi \oplus \langle\partial_t\rangle$ (which arises from the formal data) naturally orients ξ. Then, $\mathcal{D}$ being Engel and inducing the correct orientation on $\mathcal{E}$ amounts to the line field X turning positively in ξ. Ensuring this takes care of the isomorphism $\det(\mathcal{D}) \cong \mathcal{E}/\mathcal{D}$ (Equation (1.1)).

A contact–type flowbox is readily described by the oriented contact structure ξ, an oriented frame on it, and an *angular function*, exactly like a contact–type shell. Then, the Engel condition (with the adequate orientation) boils down to the angular function being strictly increasing.

Analogously:

Definition 3.52. *A **convex–type flowbox** is a formal Engel structure $(\mathcal{W}, \mathcal{D}, \mathcal{E})$ on the 4–cell $\mathbb{D}^3 \times [0,1]$ conforming to the following properties:*

- *$\mathcal{D} = \langle\partial_t, X\rangle$, where the curves $X_p : [0,1] \to \mathbb{S}^2$ are immersions,*
- *$\mathcal{W}$ is transverse to the line field ∂_t,*
- *$\mathcal{E} = [\mathcal{D}, \mathcal{D}]$ as oriented bundles. Equivalently, the isomorphism $\det(\mathcal{D}) \cong \mathcal{E}/\mathcal{D}$ of the formal Engel structure is geometrically realised through Lie bracket by $\mathcal{D}$ being non–integrable.*

The second assumption states that an orientation of $\langle X\rangle$ yields an orientation of $\mathcal{W}$. Therefore, choosing X (as opposed to $-X$) orients $\mathbb{D}^3 \times [0,1]$. Suppose $\mathcal{D}$ is Engel. This yields a second orientation for $\mathbb{D}^3 \times [0,1]$ by stating that the frame $\langle X, \partial_t, X', X''\rangle$ should be positive. Indeed, this is the natural orientation of M from the perspective of the Engel flag $\mathcal{W} \subset \mathcal{D} \subset \mathcal{E} \subset TM$, since $\mathcal{W}$ is *transverse* to ∂_t.

Changing X by $-X$ changes both orientations too, so it is meaningful to require for these two orientations to agree. If $\mathcal{D}$ is Engel and they agree, the bundle isomorphism of Equation (1.1) holds.

Orient $[0,1]$ with ∂_t; this uniquely defines an orientation of $\mathbb{D}^3$ if we want $\mathbb{D}^3 \times [0,1]$ to have the product orientation. Then, $\langle X, X', X''\rangle$ must be a *negative* basis for $\mathbb{D}^3$. This is unfortunate, because it means that the curves X_p have to be *concave* instead of convex. Since we do not like this very much, in a convex–type flowbox we will always orient $\mathbb{D}^3$ and hence $\mathbb{D}^3 \times [0,1]$ in the manner *opposite* to the one we should. Then, being Engel and inducing the correct formal Engel structure will mean that the curves X_p have to be *convex*.

Definition 3.53. *Let M be a 4–manifold endowed with a formal Engel structure $(\mathcal{W}, \mathcal{D}, \mathcal{E})$. A subset $\mathcal{U} \subset M$ is a contact/convex–type flowbox if there is a diffeomorphism $\phi : \mathcal{U} \to \mathbb{D}^3 \times [0,1]$ such that $\phi_*(\mathcal{W}, \mathcal{D}, \mathcal{E})$ is a contact/convex–type flowbox.*

We say that ϕ is a ***trivialising chart****.*

3.5 Reduction methods

The proof of Theorem 3.1 consists of a reduction process and an extension problem. In this section we present two reduction methods. The first one can be stated as follows:

Theorem 3.54. *Let C be a non–negative integer. Let $(\mathcal{W}_0, \mathcal{D}_0, \mathcal{E}_0)$ be a formal Engel structure on a closed 4–manifold M. Then, there exists a homotopy of formal Engel structures $(\mathcal{W}_s, \mathcal{D}_s, \mathcal{E}_s)$, $s \in [0,1]$, and a collection of 4–discs $B_1, \ldots, B_p \subset M$ such that:*

a. $(\mathcal{W}_1, \mathcal{D}_1, \mathcal{E}_1)$ is a genuine Engel structure in $M \setminus \bigcup_{i=1}^p B_i$, the complement of the 4–balls.

b. For each $i \in \{1, \ldots, p\}$, the restriction of the formal Engel structure $(\mathcal{W}_1, \mathcal{D}_1, \mathcal{E}_1)$ to each 4–ball B_i is a contact–type shell of height greater than C.

The proof can be found in Subsection 3.5.4. Similarly, the other reduction method, whose proof is given in Subsection 3.5.5, reads:

Theorem 3.55. *Let C be a non–negative integer. Let $(\mathcal{W}_0, \mathcal{D}_0, \mathcal{E}_0)$ be a formal Engel structure on a closed 4–manifold M. Then, there exists a homotopy of formal Engel structures $(\mathcal{W}_s, \mathcal{D}_s, \mathcal{E}_s)$, $s \in [0,1]$, and a collection of 4–discs $B_1, \ldots, B_p \subset M$ such that:*

a. $(\mathcal{W}_1, \mathcal{D}_1, \mathcal{E}_1)$ is a genuine Engel structure in $M \setminus \bigcup_{i=1}^p B_i$, the complement of the 4–balls.

b. For each $i \in \{1, \ldots, p\}$, the restriction of the formal Engel structure $(\mathcal{W}_1, \mathcal{D}_1, \mathcal{E}_1)$ to each 4–disc B_i is a convex–type shell of winding greater than C.

The proofs rely on performing local deformations that ensure that the 4–balls that we obtain have large height/winding. We will define two quantities, the contact–type energy and the convex–type energy, that are globally defined but that, up to constants, compute the derivative of the height or winding. This will provide us with a tool to relate winding/height across different flowboxes/shells. We dedicate Subsections 3.5.1, 3.5.2, and 3.5.3 to this.

One important matter we will have to pay attention to is that of *orientations*. A formal Engel structure is not just a flag, it also packages the isomorphisms $TM/\mathcal{E} \cong \det(\mathcal{E}/\mathcal{W})$ and $\mathcal{E}/\mathcal{D} \cong \det(\mathcal{D})$. They play a role in deciding what *positive* energy is. This is very similar to the discussion we had regarding flowboxes.

3.5.1 Contact–type energy

Contact–type energy is a measure of the Engelness of a formal Engel structure whose $\mathcal{E}$ is already even–contact: it is actually a parametrisation–independent way of controlling the derivative of the *developing map* along each $\mathcal{W}$–curve. The point of why we introduce it is this: it is easier to check that this derivative is large pointwise, than to check that the developing map has performed some fixed amount of turns along some $\mathcal{W}$–segment. Whereas geometrically we are interested in the latter, the former is actually more convenient to work with.

Remark 3.56. This notion was defined in [10] under the name of *Engel energy*. Since this quantity is only meaningful for contact–type shells (and we have an alternate definition of energy for convex–type ones), we have opted for this new naming convention.

Definition 3.57. *Let M be a 4–manifold. Let g be a Riemannian metric. Let $(\mathcal{W}, \mathcal{D}, \mathcal{E})$ be a formal Engel structure with $\mathcal{E}$ even–contact and $\mathcal{W}$ its kernel.*

Given any point $p \in M$, we can uniquely construct a orthonormal, positively–oriented frame $\{W \in \mathcal{W}, X \in \mathcal{D}, Y\}$ for $\mathcal{E}$ in $\mathcal{O}p(p)$. The ***contact–type energy*** *of the 2–distribution $\mathcal{D}$ at the point $p \in M$ is*

$$\mathcal{H}_{cont}(\mathcal{D})(p) = \langle [W, X], Y \rangle.$$

The definition analytically captures the geometric intuition that in order for $(\mathcal{W}, \mathcal{D}, \mathcal{E})$ to define an Engel structure, the Legendrian vector field X should rotate towards Y when we flow along the line field $\mathcal{W}$. The following lemma is immediate:

Lemma 3.58. *Let $(\mathcal{W}, \mathcal{D}, \mathcal{E})$ be a formal Engel structure with $(\mathcal{E}, \mathcal{W})$ even–contact. Then $\mathcal{H}_{cont}(\mathcal{D})(p) > 0$ if and only if $\mathcal{D}$ is Engel at p and induces the correct orientation on $\mathcal{E}$.*

By correct orientation we mean that the orientation of $\mathcal{E}$ as $[\mathcal{D}, \mathcal{D}]$ agrees with the one arising from the additional data in the formal Engel structure.

The following result simply states that the contact–type energy is, up to constants, the derivative of the developing map/angular function.

Lemma 3.59. *Let M be a 4–manifold endowed with a metric g. Let $\mathcal{E}$ be an even–contact structure with $\mathcal{W}$ its kernel.*

Then, for each $\mathcal{W}$–flowbox $U \subset M$ with trivialising chart $\phi : U \to \mathbb{D}^3 \times [0, 1]$, there exists a positive constant $H_\phi > 0$ satisfying:

$$\frac{1}{H_\phi} \mathcal{H}_{cont}(\mathcal{D})(\phi^{-1}(p, t)) < \partial_t c(\phi_* \mathcal{D})(p, t) < H_\phi \cdot \mathcal{H}_{cont}(\mathcal{D})(\phi^{-1}(p, t)).$$

The inequalities hold for all choices of $\mathcal{D}$ making $(\mathcal{W}, \mathcal{D}, \mathcal{E})$ a formal Engel structure.

Proof. We can assume that we are working in $\mathbb{D}^3 \times [0,1]$, where some metric g has been fixed. Let W, X, and Y be as in the definition of the contact–type energy. Let $\{Z_1, Z_2\}$ be the frame for ξ used to compute c; it defines a metric g_0 by letting $\{\partial_t, Z_1, Z_2\}$ be a unitary framing. Then, let $\{\tilde{X}, \tilde{Y}\}$ be a positively oriented, g_0–unitary framing for ξ with $\mathcal{D} = \langle \partial_t, \tilde{X} \rangle$.

There are functions $f_1, f_2, f_3, g_2, g_3, h_3 : \mathbb{D}^3 \times [0,1] \to \mathbb{R}$ such that $W = f_1\partial_t$, $X = f_2\partial_t + g_2\tilde{X}$, $Y = f_3\partial_t + g_3\tilde{X} + h_3\tilde{Y}$. Observe that, in terms of the metric g, ∂_t and $\tilde{X}$ are both orthogonal to Y. Then, a computation shows:

$$\langle [W,X], Y\rangle = \langle f_1 g_2[\partial_t, \tilde{X}] + [df_2(W) - df_1(X)]\partial_t + dg_2(W)\tilde{X}, Y\rangle =$$

$$f_1 g_2 \langle [\partial_t, \tilde{X}], Y\rangle = (\partial_t c) f_1 g_2 \langle \tilde{Y}, Y\rangle = (\partial_t c)\frac{f_1 g_2}{h_3}.$$

There are universal bounds, depending only on the metrics g and g_0, controlling the functions f_1, g_2, and h_3 from above and below. The claim follows. □

That is, once we fix $\mathcal{E}$ even–contact, we have an universal estimate to go from the derivative of the angular function of any $\mathcal{D}$ sandwiched between $\mathcal{W}$ and $\mathcal{E}$ to the contact–type energy of $\mathcal{D}$.

3.5.2 Convex–type energy

Much like contact–type energy measures the derivative of the angular function, we want *convex–type energy* to measure the curvature of the convex curves describing a convex–type flowbox. We formalise this intuition in the following definition:

Definition 3.60. *Let M be a 4–manifold. Let g be a metric. Let $(\mathcal{W}, \mathcal{D}, \mathcal{E})$ be a formal Engel structure satisfying $\mathcal{E} = [\mathcal{D}, \mathcal{D}]$ as oriented distributions. Let $\mathcal{Y} \subset \mathcal{D}$ be a line field complementary to $\mathcal{W}$.*

Given any point $p \in M$, there is a unique orthonormal, positively oriented frame $\{X \in \mathcal{D}, Y \in \mathcal{Y}, Z\}$ of $\mathcal{E}$ in $\mathcal{O}p(p)$. Note that the choice of X orients $\mathcal{O}p(p)$. The ***convex–type energy*** *of the 2–distribution $\mathcal{D}$ at p is*

$$\mathcal{H}_{convex}(\mathcal{D})(p) = \det(X, Y, Z, [Y, Z])(p).$$

Regarding orientations. First, observe that $[\mathcal{D}, \mathcal{D}]$ is naturally oriented as a 3–distribution, and what we require is that the orientation of $\mathcal{E}$ (which is part of the data of the formal Engel structure) agrees with it. Secondly, our choice of X orients $\mathcal{W}$ and therefore $\mathcal{O}p(\{p\})$, so the determinant can be computed. It is easy to see that this choice is auxiliary and does not affect the sign of the convex energy. See Subsection 3.4.4 for a similar discussion in terms of flowboxes.

Remark: the convex–type energy is information of second order on $\mathcal{D}$. Indeed, Z is obtained from $[X, Y]$ by the Gram–Schmidt process (where the first two vectors are X and Y). In particular, $\mathcal{H}_{\text{convex}}(\mathcal{D})(p) = 0$ if and only if X, Y, $[Y, X]$, and $[Y, [Y, X]]$ are linearly dependent.

The fundamental claim is that, up to some universal constant, the convex–type energy computes the curvature of the curves X_p of any convex–type flowbox we consider.

Lemma 3.61. *Let M be a 4–manifold endowed with $\mathcal{Y}$, a line field. Let $\mathcal{U} \subset M$ be a $\mathcal{Y}$–flowbox with trivialising chart $\phi : \mathcal{U} \to \mathbb{D}^3 \times [0,1]$. Let $\Phi_p : D_p \subset \mathbb{S}^2 \to \mathbb{D}^2 \subset \mathbb{R}^2$, $p \in \mathbb{D}^3$, be a family of affine charts whose domain we have restricted to the preimage D_p of the unit ball.*

Then, there exists a positive constant H_{ϕ,Φ_p} satisfying:

$$\frac{1}{H_{\phi,\Phi_p}}\mathcal{H}_{convex}(\mathcal{D})(\phi^{-1}(p,t)) < \partial_t \mathcal{G}_{\Phi_p\circ X_p}(t) < H_{\phi,\Phi_p}\cdot \mathcal{H}_{convex}(\mathcal{D})(\phi^{-1}(p,t))$$

for any choice of convex–type flowbox $\mathcal{D}\supset\mathcal{Y}$ described by curves X_p, whenever $\Phi\circ X_p$ is defined.

Indeed, $\Phi_p\circ X_p$ is not be defined if $X_p(t)$ does not map to the domain of Φ_p. The reason why we introduce Φ_p is that in the plane it is immediate to estimate the number of turns the Gauss map is performing from its derivative. Making sense of a similar statement in the sphere requires more work. In any case, the winding of a convex–type shell was computed in a affine chart, so this works well for our purposes.

Notice as well that we did not consider a full affine chart, but its restriction to the unit ball. The reason for doing this is that $\mathbb{D}^2$ is compact, and therefore the distortion that the map Φ_p introduces is controlled by compactness. This wouldn't be the case if we took a full open hemisphere to be the domain.

Proof. We assume that our ambient manifold is $\mathbb{D}^3\times[0,1]$, which we endow with a metric g. Let X, Y, Z be as in the definition of convex–type energy, using g. Let $\tilde{X}$, ∂_t, $\tilde{Z}$ be their analogues with respect to the euclidean metric g_0. Write X_p for the curves corresponding to $\tilde{X}$.

Since the volume forms of g and g_0 are related by a constant, there is a positive constant bounding the following determinants from above and below in terms of each other:

$$\mathcal{H}_{\text{convex}}(\mathcal{D})(p,t)=\det_g(X,Y,Z,[Y,Z])(p,t) \quad\text{and}\quad -\det_{g_0}(X,Y,Z,[Y,Z])(p,t).$$

The difference in sign follows from our convention regarding convex–type flowboxes (Subsection 3.4.4): $\det_g$ is computed with respect to the natural orientation arising from the formal Engel structure (as in the definition of convex–type energy), but $\det_{g_0}$ is computed with respect to the *opposite* orientation. This way, positivity means that the curves X_p are convex and not concave (as we shall see in the remainder of the proof).

Arguing as in Lemma 3.59, it is easy to see that

$$-\det_{g_0}(X,Y,Z,[Y,Z])(p,t) \quad\text{and}\quad -\det_{g_0}(\tilde{X},\partial_t,\tilde{Z},[\partial_t,\tilde{Z}])(p,t),$$

are related by a universal positive constant, only depending on g_0 and g. Thus, we are left with showing that the right hand side is, up to scaling, the derivative of the Gauss map $\mathcal{G}_{\Phi_p\circ X_p}$.

Observe that

$$-\det_{g_0}(\tilde{X},\partial_t,\tilde{Z},[\partial_t,\tilde{Z}])(p,t)=\det_{g_0,\mathbb{D}^3}(\tilde{X},\tilde{Z},[\partial_t,\tilde{Z}])(p,t);$$

that is, since we are dealing with the Euclidean metric, the first determinant is equivalent to computing the second one (which is given by the euclidean metric in $\mathbb{D}^3$ and the orientation of $\mathbb{D}^3$ that is also opposite to the one induced by the formal Engel structure).

Now, Φ_p is not an isometry but, being a map between compact manifolds (since we restricted its domain), there are universal bounds relating their corresponding volume forms. By construction, Z is the normalised velocity vector of X. Bringing both facts together, we deduce that, up to universal positive constants,

$$\det_{g_0,\mathbb{D}^3}(\tilde{X},\tilde{Z},[\partial_t,\tilde{Z}])(p,t) \quad\text{and}\quad \det_{g_0,\mathbb{D}^2}(\mathcal{G}_{\Phi_p\circ X_p}(t),\partial_t\mathcal{G}_{\Phi_p\circ X_p}(t))$$

are the same. This concludes the proof. □

3.5.3 Non–integrability energy

Let us understand convex–type energy a bit better. Consider the following scenario: fix a convex curve $f : \mathbb{S}^1 \to \mathbb{S}^2$ that describes a non–maximal circle in $\mathbb{S}^2$ that is very close to being an equator. Consider now a family of convex–type shells $\mathcal{D}_C$ where the curves X_p^C describing them are of the form $t \to f(e^{iCt})$. It is clear that, since f is convex, the convex–type energy is increasing, linearly, with C. However, this is not because the curves X_p^C get more wiggly as C increases; rather, the curvature remains the same and the curves simply get longer and longer.

For performing certain operations (for instance, in the next chapter, where we define what an overtwisted Engel structure is), we will need to deform an Engel structure so that many wiggles appear. The way in which this is done is by increasing the convex–type energy while keeping the length of the curves X_p bounded. The following notion captures this idea of length in a flowbox/shell–independent fashion:

Definition 3.62. *Let M be a 4–manifold. Let g be a metric. Let $(\mathcal{W}, \mathcal{D}, \mathcal{E})$ be a formal Engel structure satisfying $\mathcal{E} = [\mathcal{D}, \mathcal{D}]$ as oriented distributions. Let $\mathcal{Y} \subset \mathcal{D}$ be a line field complementary to $\mathcal{W}$.*

*At point $p \in M$, find a orthonormal frame $\{Y \in \mathcal{Y}, X \in \mathcal{D}, Z\}$ for $\mathcal{E}$ in $\mathcal{O}p(\{p\})$. Then, the **non–integrability energy** of $\mathcal{D}$ is defined to be:*

$$\mathcal{H}_{n.i.}(\mathcal{D}) = \det_{g,\mathcal{E}}(Y, X, [Y, X]).$$

The determinant in the left hand side is measured with respect to the natural orientation of $\mathcal{E}$. The non–integrability energy measures the *speed* of the curves X_p of any given convex–type flowbox along $\mathcal{Y}$:

Lemma 3.63. *Let M be a 4–manifold endowed with $\mathcal{Y}$, a line field. Let $\mathcal{U} \subset M$ be a $\mathcal{Y}$–flowbox with trivialising chart $\phi : \mathcal{U} \to \mathbb{D}^3 \times [0, 1]$.*

Then, there exists a positive constant H_ϕ satisfying:

$$\frac{1}{H_\phi}\mathcal{H}_{n.i.}(\mathcal{D})(\phi^{-1}(p, t)) < |X_p'(t)| < H_\phi \cdot \mathcal{H}_{n.i.}(\mathcal{D})(\phi^{-1}(p, t))$$

for any choice of convex–type flowbox $\mathcal{D} \supset \mathcal{Y}$ described by curves X_p.

Proof. We assume that our ambient manifold is $\mathbb{D}^3 \times [0, 1]$, which we endow with a metric g. Let $\tilde{X}$, $\tilde{Y}$ play the role of X and Y in the definition of non–integrability energy. Let $\{\partial_t, X\}$ be the unitary frame of $\mathcal{D}$ with respect to the Euclidean metric g_0.

Then, there is a positive universal constant relating

$$\det_{g,\mathcal{E}}(\tilde{Y}, \tilde{X}, [\tilde{Y}, \tilde{X}])(p, t) \quad \text{and} \quad \det_{g_0,\mathcal{E}}(\tilde{Y}, \tilde{X}, [\tilde{Y}, \tilde{X}])(p, t),$$

and yet another one relating

$$\det_{g_0,\mathcal{E}}(\tilde{Y}, \tilde{X}, [\tilde{Y}, \tilde{X}])(p, t) \quad \text{and} \quad \det_{g_0,\mathcal{E}}(\partial_t, X, X')(p, t),$$

both of which depend only on g and g_0. But then, the right hand side is exactly $|X_p'(t)|$. ☐

3.5.4 Reduction to contact–type shells

Let us sketch what goes into the proof of Theorem 3.54. A first homotopy allows us to assume that $\mathcal{E}$ is an even–contact structure and $\mathcal{W}$ is its kernel. We then use the triangulation and covering introduced in Section 3.2 to proceed over (neighbourhoods of) the lower dimensional cells adjusting $\mathcal{D}$, while keeping $\mathcal{W}$ and $\mathcal{E}$ fixed. If $\mathcal{E}$ and $\mathcal{W}$ are of that form, and $\mathcal{D}$ is sandwiched between them, we are dealing with contact–type flowboxes and shells and $\mathcal{D}$ is purely described by angular functions. Having removed some 4–balls, the resulting manifold is open, which allows us to construct angular functions with no critical points (and therefore, everywhere increasing and actually with arbitrarily large derivative).

3.5.4.1 Setting up the proof

Consider a formal Engel structure $(\mathcal{W}_0, \mathcal{D}_0, \mathcal{E}_0)$ on a closed 4–manifold M and $C_0 \in \mathbb{R}^+$ a constant. By applying Theorem 0.9, McDuff's h–principle for even–contact structures, we produce a first homotopy $(\mathcal{W}_s, \mathcal{D}_s, \mathcal{E}_s)$, $s \in [0, 1/2]$, so that $\mathcal{E}_{1/2}$ is even–contact and $\mathcal{W}_{1/2}$ is its kernel. As abstract bundles $\mathcal{W}_s \subset \mathcal{E}_s$ do not depend on s, so it is indeed possible to extend $\mathcal{D}_0$ along the homotopy to a family $\mathcal{D}_s$, $s \in [0, 1/2]$.

Theorem 3.9 provides a triangulation $\mathcal{T} = \{\sigma\}$ and a covering of M by closed discs $\{\mathcal{U}(\sigma)\}$ that is well–suited to the induction process that we want to perform. Each disc $\mathcal{U}(\sigma)$ is actually a contact–type flowbox that comes with a trivialising chart $\phi(\sigma) : \mathcal{O}p(\mathcal{U}(\sigma)) \to \mathbb{D}^3_{1+\varepsilon} \times [-\varepsilon, 1+\varepsilon]$. We shall deform $\mathcal{D}$ flowbox by flowbox obtaining lower bounds on the derivative of the angular functions; using the notion of contact–type energy, we will be able to translate these lower bounds between different flowboxes. In the end, we will be able to conclude that the 4–balls are actually contact–type shells having height as large as we want.

3.5.4.2 Contact–type energy in the lower skeleta.

Fix a positive constant $C_0 \in \mathbb{R}^+$. Fix a metric g on M. Let us construct a deformation $\mathcal{D}_1$ of $\mathcal{D}_{1/2}$ satisfying $\mathcal{H}(\mathcal{D}')|_{\mathcal{U}_3} > C_0$, where we denote

$$\mathcal{U}_j := \bigcup_{\sigma \in \mathcal{T}^{(j)}} \mathcal{U}(\sigma).$$

This is achieved by induction over the dimension j of the simplices. The induction hypothesis is that

$$\mathcal{H}_{\text{cont}}(\mathcal{D})|_{\mathcal{U}_j} > C_0 \cdot H^{2(3-j)}$$

where the constant H is the maximum among all the constants $H_{\phi(\sigma)}$ arising from Lemma 3.59. Do recall that the number of cells is finite.

Suppose that $\mathcal{D}_{1/2}$ has already been deformed on $\mathcal{U}_{j-1}$ suitably to produce some $\mathcal{D}$. Consider the image through $\phi(\sigma)$ of the finite union $\cup_{\tau \subsetneq \sigma} \mathcal{U}(\tau)$. Property (2.) in Theorem 3.9 implies that this closed set is of the form $A \times [-\varepsilon, 1+\varepsilon]$, for some closed set A, if the thickening is small enough. The inductive hypothesis $\mathcal{H}_{\text{cont}}(\mathcal{D})|_{\mathcal{U}_{j-1}} > C_0 \cdot H^{2(4-j)}$ translates into the inequality

$$\partial_t c(\phi(\sigma)_* \mathcal{D})|_{A \times [-\varepsilon, 1+\varepsilon]} > C_0 \cdot H^{2(4-j)-1}|_{A \times [-\varepsilon, 1+\varepsilon]}.$$

Consider a function $f : \mathbb{D}^3_{1+\varepsilon} \times [-\varepsilon, 1+\varepsilon] \longrightarrow \mathbb{R}$ such that

$$f|_{A \times [-\varepsilon, 1+\varepsilon]} = \partial_t c(\phi(\sigma)_* \mathcal{D})|_{A \times [-\varepsilon, 1+\varepsilon]}, \qquad \text{and } f > C_0 \cdot H^{2(4-j)-1}.$$

This function f is the derivative of an angular function for a contact–type shell with contact–type energy greater than $C_0 \cdot H^{2(3-j)}$, and it agrees with the function $\partial_t c(\phi(\sigma)_*\mathcal{D})$ on $\mathcal{U}_{j-1}$, where the contact–type energy of the 2–distribution $\mathcal{D}$ is already greater than $C_0 \cdot H^{2(4-j)} > C_0 \cdot H^{2(3-j)}$.

The linear interpolation serves now as the required deformation of $\mathcal{D}_{1/2}$. In detail, consider a cut–off function $\beta : \mathbb{D}^3_{1+\varepsilon} \times [-\varepsilon, 1+\varepsilon] \to [0,1]$ such that

$$\beta|_{\mathbb{D}^3 \times [0,1]} \equiv 1, \quad \beta|_{\mathcal{O}p(\partial(\mathbb{D}^3_{1+\varepsilon} \times [-\varepsilon, 1+\varepsilon]))} \equiv 0,$$

and the angular function $d : \mathbb{D}^3_{1+\varepsilon} \times [-\varepsilon, 1+\varepsilon] \to \mathbb{R}$ defined as the linear interpolation

$$d(p,t) = (1-\beta(p,t))c(p,t) + \beta(p,t)\left(c(p,0) + \int_0^t f(p,t)dt\right).$$

Then the two angular functions c and d are isotopic relative to the boundary and $A \times [-\varepsilon, 1+\varepsilon]$. Hence, d induces a deformation $\mathcal{D}$ of $\mathcal{D}_{1/2}$ through 2–distributions contained in $\mathcal{E}$ and containing $\mathcal{W}$, relative to $\mathcal{U}_{j-1}$. By applying this deformation to each j–simplex $\sigma \in \mathcal{T}^{(j)}$ and the inductive character of the argument, we obtain a deformation $\mathcal{D}_1$ such that $\mathcal{H}(\mathcal{D}_1)|_{\mathcal{U}_3} > C_0$. □

This provides a deformation satisfying Property (a.) in the statement of Theorem 3.54. The second step in the proof is thus to prove that having large contact–type energy in the neighbourhood $\mathcal{U}_3$ of the 3–skeleton implies that all the 4–cells are actually contact–type shells of height C. Note that the constant C is given by the statement, whereas the constant $C_0 \in \mathbb{R}^+$ in the previous argument can be chosen arbitrarily.

3.5.4.3 Contact–type energy vs. height in the 4–cells

Consider a 4–simplex $\sigma \in \mathcal{T}^{(4)}$. Property (a.) of the triangulation/covering provided by Theorem 3.9 ensures that $\partial\sigma \subset \cup_{\tau \subsetneq \sigma} \mathcal{U}(\tau)$, this implies, first of all, that $\mathcal{U}(\sigma)$ is indeed a contact–type shell (and not just a contact–type flowbox) that additionally satisfies the bound:

$$\partial_t c(\phi(\sigma)_*\mathcal{D}) \mid_{\partial\mathbb{D}^3} > C_0/H_{\phi(\sigma)},$$

by applying the comparison lemma about contact–type energy, Lemma 3.59.

Choose the constant $C_0 \in \mathbb{R}^+$ so that the inequality $C_0/H > C$ is satisfied, where H is the universal constant defined in the previous subsubsection satisfying $H > H_{\phi(\sigma)}$. This implies the inequality $\text{height}(\phi(\sigma)_*\mathcal{D})|_{\partial\mathbb{D}^3} > C$ for the height of each contact–type shell $\mathcal{U}(\sigma)$, concluding the proof. □

3.5.5 Reduction to convex–type shells

The proof of Theorem 3.55 pretty much follows the same lines of the proof of Theorem 3.54. We start by deforming the formal Engel structure to guarantee $[\mathcal{D}, \mathcal{D}] = \mathcal{E}$. Then, any flowbox of $\mathcal{W}$ is automatically a convex–type flowbox. We proceed to apply Theorem 3.9 to find a suitable triangulation/covering of M. We then deform $\mathcal{D}$ inductively over the neighbourhoods of the lower dimensional cells, ensuring that at every step the convex–type energy is large. This guarantees that the 4–cells are actually convex–type shells.

3.5.5.1 Setting up the proof

We start with M a 4–manifold, endowed with a metric g, and a formal Engel structure $(\mathcal{W}_0, \mathcal{D}_0, \mathcal{E}_0)$. We let $\mathcal{Y} \subset \mathcal{D}_0$ be a line field transverse to $\mathcal{W}_0$ (which, up to homotopy, is unique). An application of Proposition 1.15 with respect to the line field $\mathcal{Y}$, allows us to homotope $(\mathcal{W}_0, \mathcal{D}_0, \mathcal{E}_0)$ to a new formal Engel structure $(\mathcal{W}_{1/2}, \mathcal{D}_{1/2}, \mathcal{E}_{1/2})$ such that:

- $\mathcal{E}_{1/2} = [\mathcal{D}_{1/2}, \mathcal{D}_{1/2}]$ as oriented distributions,
- $\mathcal{D}_{1/2}$ and all the intermediate 2–distributions still contain $\mathcal{Y}$.

$\mathcal{E}_{1/2}$ is of course not necessarily even–contact, or we would be done. The first condition means that any $\mathcal{Y}$–flowbox is a convex–type flowbox for $\mathcal{D}_{1/2}$.

3.5.5.2 Introducing convex–type energy in the lower skeleta

We apply Theorem 3.9 to the line field $\mathcal{Y}$; from this, we obtain a triangulation $\mathcal{T}$ of M and a cover $\{\mathcal{U}(\sigma)\}$ by (convex–type) flowboxes of $\mathcal{Y}$, each of which has a trivialising chart $\phi(\sigma) : \mathcal{O}p(\mathcal{U}(\sigma)) \to \mathbb{D}^3_{1+\varepsilon} \times [-\varepsilon, 1+\varepsilon]$, with $\varepsilon > 0$ small.

Let $\sigma \in \mathcal{T}$. Write $X_p^{1/2}$ for the curves describing $\mathcal{D}_{1/2}$ in Image$(\phi(\sigma))$. If $\mathcal{T}$ is fine enough, we can assume that the curves $X_p^{1/2}$ have image in a very small disc in $\mathbb{S}^2$. In particular, we can take a family of (restrictions of) affine charts $\Phi_p(\sigma) : D_p \subset \mathbb{S}^2 \to \mathbb{D}^2 \subset \mathbb{R}^2$ whose domains cover $X_p^{1/2}$ for all p.

Let $C_0 > 0$ be a constant that we shall later fix. It will serve as a lower bound for the convex–type energy in the domains:

$$\mathcal{U}_j := \bigcup_{\sigma \in \mathcal{T}^{(j)}} \mathcal{U}(\sigma), \quad j = 0, 1, 2, 3.$$

If $j > 0$, suppose that $\mathcal{D}_{1/2}$ has already been deformed in $\mathcal{U}_{j-1}$ to produce some $\tilde{\mathcal{D}}$ satisfying

$$\mathcal{H}_{\text{convex}}(\tilde{\mathcal{D}})|_{\mathcal{U}_{j-1}} > C_0 \cdot H^{2(4-j)}.$$

Here H is a constant that satisfies $H > H_{\phi(\sigma), \Phi_p(\sigma)}$ for all $\sigma \in \mathcal{T}$. Suppose additionally that $\tilde{\mathcal{D}}$ is C^0–close to $\mathcal{D}_{1/2}$.

Let $\sigma \in \mathcal{T}^{(j)}$. Write X_p for the curves describing $\tilde{\mathcal{D}}$, they are C^0–close to the curves $X_p^{1/2}$ by hypothesis. In particular, X_p is contained in the domain of Φ_p. Theorem 3.9 implies that $\phi(\sigma)(\mathcal{U}_{j-1})$ is a closed set of the form $A \times [-\varepsilon, 1+\varepsilon]$ and Lemma 3.61 states that the curves X_p, $p \in A$, satisfy:

$$\partial_t \mathcal{G}_{\Phi_p(\sigma) \circ X_p}(t) > C_0 \cdot H^{2(4-j)-1}.$$

Now we proceed as in Subsection 3.5.4. We can construct a family of immersed curves

$$Y_p : [-\varepsilon, 1+\varepsilon] \to \mathbb{D}^2, \quad p \in D^3_{1+\varepsilon},$$

satisfying:

- $\partial_t \mathcal{G}_{Y_p} > C_0 \cdot H^{2(4-j)-1}$ if $(p, t) \in \mathbb{D}^3 \times [0, 1]$,

- $Y_p = \Phi_p(\sigma) \circ X_p$ if $(p,t) \in \mathcal{O}p(\partial(\mathbb{D}^3_{1+\varepsilon} \times [-\varepsilon, 1+\varepsilon]))$ or $p \in \mathcal{O}p(A)$,
- Y_p is in the same formal class as $\Phi_p(\sigma) \circ X_p$, relative to the boundary of the shell.

This is achieved by first defining Y_p in $(p,t) \in \mathbb{D}^3 \times [0,1]$, extending it as a family of formal immersions, and then using the Hirsch–Smale theorem. If we replace $\Phi_p(\sigma) \circ X_p$ by Y_p to homotope $\tilde{\mathcal{D}}$, Lemma 3.61 implies that the convex–type energy in $\mathcal{U}(\sigma)$ is now greater than $C_0 \cdot H^{2(3-j)}$.

Proceeding cell by cell we obtain that this is true over the whole of $\mathcal{U}_j$. After iterating $j = 0, 1, 2, 3$, we deduce that we can introduce C_0 convex–type energy in $\mathcal{U}_3$. The resulting formal Engel structure we denote by $(\mathcal{W}_1, \mathcal{D}_1, \mathcal{E}_1 = [\mathcal{D}_1, \mathcal{D}_1])$.

3.5.5.3 From convex–type energy to convex–type shells

Now we want to show that having C_0 convex–type energy in $\mathcal{U}_3$ (with C_0 sufficiently large depending on C) implies that all the $\mathcal{U}(\sigma)$, $\sigma \in \mathcal{T}^{(4)}$, are actually convex–type shells of winding at least C, proving Theorem 3.55.

Theorem 3.9 states that the boundary of each $\mathcal{U}(\sigma)$, with $\sigma \in \mathcal{T}^{(4)}$, is covered by $\mathcal{U}_3$. In particular, $\mathbb{D}^3_{1+\varepsilon} \times [-\varepsilon, \rho]$ is contained in $\phi(\sigma)(\mathcal{U}_3)$ if $\rho \in (0,1)$ is sufficiently small. Since there are finitely many such σ, ρ can be chosen uniformly for all of them.

Fix $\sigma \in \mathcal{T}^{(4)}$. Write X_p^1 for the curves describing $\mathcal{D}_1$ in $\mathcal{U}(\sigma)$. Our bound on the convex–type energy along the boundary implies:

$$\partial_t \mathcal{G}_{\Phi_p \circ X_p^1}(t) > \frac{C_0}{H_{\phi,\Phi_p}}, \quad t \in [0, \rho].$$

Thus, if we want $\Phi_p \circ X_p^1$ to have winding C in $t \in [0, \rho]$, we need to impose

$$C_0 > \frac{2\pi H_{\phi,\Phi_p}}{\rho} C, \text{ for all } p.$$

Since there are finitely many 4–cells, any C_0 large enough will satisfy this inequality for all the 4–cells simultaneously. □

Remark 3.64. The reader might wonder what happened with the non–integrability energy and why we did not make use of it. Go back to the construction of the curves Y_p. What we did was the following: we deformed $\mathcal{D}_{1/2}$ so that the resulting 2–distribution $\mathcal{D}_1$ was C^0–close to it. In particular, we allowed for its non–integrability energy to grow, but only by letting the curves X_p^1 move very fast in very small regions of $\mathbb{S}^2$. Since this is the case, the curves X_p^1 have many wiggles regardless.

We could have been a bit more careful. The curves Y_p can be obtained from $\Phi_p \circ X_p$ by setting:

$$Y_p(t) = \Phi_p \circ X_p^{[B_1 \# t_1(p), \dots, B_n \# t_n(p)]}(t), \quad t \in [0,1],$$

with $t_i : \mathbb{D}^3 \to (0,1)$ some functions satisfying $t_1 < t_2 < \cdots < t_n$ and $B_i > 0$ some integers. This has to be done carefully to ensure the relative character of the construction.

It is easy to see that the resulting curves $Y_p|_{[0,1]}$ have length arbitrarily close to $\Phi_p \circ X_p|_{[0,1]}$ by taking the wiggles arbitrarily small. Outside of $\mathbb{D}^3 \times [0,1]$, instead of using Hirsch–Smale, we can introduce wiggles going in the opposite direction to preserve the formal class, and these can be taken to be small too. We can thus assume that $\mathcal{H}_{\text{n.i.}}(\mathcal{D}_1) < (1+\delta)\mathcal{H}_{\text{n.i.}}(\mathcal{D}_{1/2})$, for any $\delta > 0$.

We will revisit this in Section 4.3, so we will not give more details for now. □

3.6 Solving the extension problem

After explaining two methods of performing the reduction process, we tackle the central part of the chapter: solving the extension problem. As we already advanced in the first section, we will provide three extension approaches. Two of them, in Subsections 3.6.1 and 3.6.2, start with contact–type shells (although they will not be contact–type anymore in the interior as we deform them), and the other, presented in Subsection 3.6.3, starts and works with convex type–shells.

As the reader might envisage from our discussions about energy, we do not know *whether the extension problem is always solvable*. Indeed: we need to impose energy conditions to be able to do it (using any of the three methods). Whether this is truly necessary remains one of the most intriguing questions in the area.

Without further ado, let us state the main results of this section:

Theorem 3.65. *A C–contact–type shell with $C \geq 6$ is homotopic, through Engel shells, to a solid Engel shell.*

Theorem 3.66. *A C–contact–type shell with $C \geq 2$ is homotopic, through Engel shells, to a solid Engel shell.*

The reader might wonder why we state the first result if the second one is stronger. The point we want to make is that using the method from Subsection 3.6.1, the *four-leaf clover*, we are able to obtain the first result. Using the method from Subsection 3.6.2, the *turning model*, we are able to obtain the second one, which provides a better bound.

Separately, and proven in Subsection 3.6.3:

Theorem 3.67. *A C–convex–type shell with $C \geq 1$ is homotopic, through convex–type shells, to a solid Engel shell.*

One immediate corollary from these theorems using the discussion about domination between shells (see Subsection 3.4.3) is the following:

Corollary 3.68. *A convex–type shell of winding $2C + 1 \geq 3$ is homotopic, through convex–type shells, to a solid Engel shell.*

A contact–type shell of height $C \geq 2$ is homotopic, through Engel shells, to a solid Engel shell.

Proof. We apply Proposition 3.50 first to pass from convex–type shells of winding $2C + 1$ to a C–convex type shell. Then Theorem 3.67 concludes the claim.

Similarly, we apply Corollary 3.47 and then Theorem 3.66 in the contact–type case. □

Observe that, regardless of the extension method, in all the statements the resulting Engel structure still contains the line field $\langle \partial_t \rangle$.

3.6.1 Filling the 4–ball using the four-leaf clover

This extension method appeared first in [10]. The idea is simple. We want to homotope a C–contact–type shell $\mathcal{D}$, $C \geq 3$, to a solid one. By Bolzano's theorem, having $c(p, 1) < c(p, 0)$

implies that this cannot be done through contact–type shells. In particular, to achieve our goal under this assumption, we need for the curves X_p to be convex somewhere along the homotopy. We claim that this can be done letting the curves X_p be either tangent to the underlying contact structure ξ or convex.

We start by describing two families of curves: the kink (or wiggle) and the four–leaf clover. The former family helps with the interpolation to the latter. The latter allows us to modify the angular function of the shell so that we can apply Lemma 3.37 to conclude.

3.6.1.1 A family of wiggles

Remark 3.69. The curves that we shall now describe were called *kink curves* in [10]. In the rest of the text we have called them *wiggles* and hence we will stick to this new naming convention. At any rate, they are simply convex loops that we add to a preexisting curve. Their role now is to interpolate, relative to the boundary, between a segment going around the equator once and a short convex segment strictly contained in an hemisphere and describing one turn, see Figure 3.10. Let us describe them analytically.

For each $\theta \in [0, \pi/2]$, consider the plane given by the equation $\{\sin(\theta)(x-1) + \cos(\theta)z = 0\}$. For $\theta = 0$ this describes the plane $\{z = 0\}$ and for $\theta = \pi/2$ the vertical plane $\{x = 1\}$. Considering $\theta \in [0, \pi/2)$, the intersection of these planes with the 2–sphere $\mathbb{S}^2$ yields the following parametrised curves

$$\beta_\theta(t) = (\sin^2(\theta) + \cos^2(\theta)\cos(t), \cos(\theta)\sin(t), \sin(\theta)\cos(\theta)(1-\cos(t))), \quad t \in [0, 2\pi].$$

The curve β_0 parametrises the equator with constant angular speed, and $\beta_{\pi/2}$ is a constant map with image the point $(1, 0, 0)$.

Lemma 3.70. *Each non–maximal circle in $\mathbb{S}^2$ is given as the intersection of the sphere with a (uniquely defined) plane that does not contain the origin. They all are convex curves (if oriented suitably).*

Proof. We shall prove that they are convex curves, the rest is pretty much by definition. Let γ be the non–maximal circle generated by the plane H. At every point of γ we can take the plane G that is tangent to γ at that point and that passes through the origin. $G \cap H \subset H$ lies to one side of γ (as curves in H), and this implies the claim. □

From the lemma we deduce that the curves β_θ, with $\theta \neq 0, \pi/2$ are convex. Observe that all of them have Frenet frame $\Gamma_{\beta_\theta}(0) = \mathrm{Id}$ at the origin.

3.6.1.2 The four-leaf clover

Let us analytically describe it in an affine chart. It is given by the parametrised plane curve $f(t) = (\cos(t)\sin(2t), \sin(t)\sin(2t))$, $t \in [0, 2\pi]$; the reader can check from the expression that it is indeed a convex curve in the plane. By Lemma 3.22, we deduce that if we use any affine chart $\Phi : H^2 \subset \mathbb{S}^2 \to \mathbb{R}^2$ to map it into the sphere, the curve $\Phi^{-1} \circ f$ is convex too. See Figure 3.11.

Consider the following curves in the 2–sphere

$$\tau_1, \beta : [0, 6\pi] \to \mathbb{S}^2$$

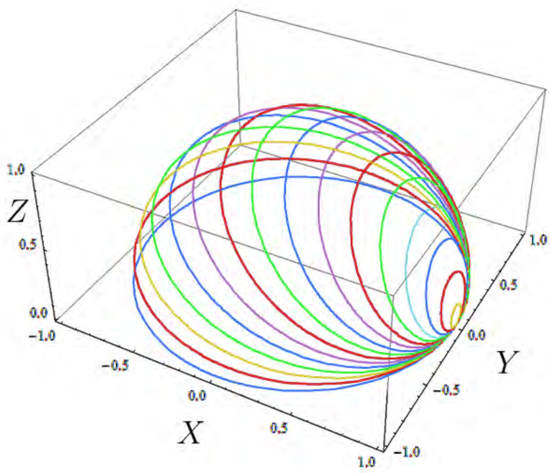

Figure 3.10: The curves β_θ for different values of the parameter θ.

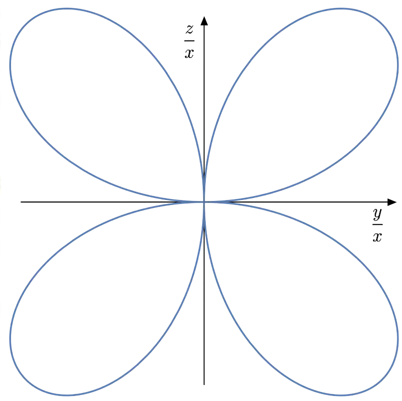

Figure 3.11: The curve $\pi \circ \kappa$.

$$\tau_1(t) = \Phi^{-1} \circ f(t/3)$$
$$\beta(t) = \beta_\theta(t)$$

where the chart Φ is chosen so that $\Gamma_{\tau_1}(0) = \mathrm{Id}$, and $\theta \in (0, \pi/2)$ is fixed but otherwise arbitrary. Then, the following holds:

Lemma 3.71. *There is a smooth family of convex curves:*

$$\tau_s : [0, 6\pi] \to \mathbb{S}^2, \quad s \in [0,1]$$

connecting $\tau_0 = \beta$ *with* τ_1 *and having Frenet frames* $\Gamma_{\tau_s}(0) = \mathrm{Id}, \forall s \in [0,1]$.

Proof. By Proposition 3.24, both curves β and τ_1 lie in the same connected component of the space of convex curves having the identity as their Frenet frame at $t = 0$. A particular interpolation can be given using an affine chart containing the image of both and applying Proposition 3.16 since, in the chart, they are convex plane curves with the same winding number (they both describe 6 projective turns). □

3.6.1.3 Outline of the proof

Let $\mathcal{D}_0(p,t) = \langle \partial_t, X^0_p(t) \rangle$ be a C–contact–type shell, $C \geq 6$, with angular function $c_0 : \mathbb{D}^3 \times [0,1] \to \mathbb{R}$. By assumption, $c_0(p, \rho) > c_0(p, 0) + 6\pi$, for some $\rho \in [0,1]$. We have $\mathcal{E}_0 = \xi \oplus \langle \partial_t \rangle$, where ξ is a contact structure in $\mathbb{D}^3$ with framing $\{Y, Z\}$, and $\mathcal{W}_0 = \langle \partial_t \rangle$.

Suppose $(\mathbb{D}^3 \times [0,1], \mathcal{D}_s)$, $s \in [0,1]$, is a homotopy of $\mathcal{D}_0$ given by curves $X^s_p : [0,1] \to \mathbb{S}^2$ and satisfying:

- the homotopy $(\mathbb{D}^3 \times [0, \rho], \mathcal{D}_s)$ is through solid shells (not of contact–type),
- the homotopy $(\mathbb{D}^3 \times [\rho, 1], \mathcal{D}_s)$ is through contact–type shells with angular function $c_s : \mathbb{D}^3 \times [\rho, 1] \to \mathbb{R}$,

- $c_s(p,t) = c_0(p,t)$ if $t \notin \mathcal{O}p([0,\rho])$ or $p \in \mathcal{O}p(\partial\mathbb{D}^3)$,
- $c_1(p,\rho) < c_1(p,1) = c_0(p,1)$.

Then, an application of Lemma 3.37 in $(\mathbb{D}^3 \times [\rho,1], \mathcal{D}_1)$ would conclude the proof of Theorem 3.65.

Constructing the homotopy $\mathcal{D}_s$, $s \in [0,1]$, is done in three steps. We define it in $s \in [0,1/3]$ so that $X_p^{1/3}$ is a four–leaf clover if $|p| < 1-\delta/3$. Then, we further homotope so that $X_p^{2/3}$ is a four–leaf clover whose midpoint has been flattened to be tangent to the (contact) equator if $|p| < 1-2\delta/3$. Lastly, we define the homotopy in $s \in [2/3,1]$ so that X_p^1, $|p| < 1-\delta$, is a four–leaf clover whose left–most petals have been pulled negatively, along the equator, as much as needed. This makes $c_1(p,\rho)$ arbitrarily small in $|p| < 1-\delta$, and in particular smaller than $c_1(p,1) = c_0(p,1)$. If $\delta > 0$ is taken to be sufficiently small, we have that $c_1(p,1) = c_0(p,1) > c_0(p,\rho) \geq c_1(p,\rho)$ in $|p| > 1-\delta$ too, and we are done.

Figure 3.12 shows in pictures how this is done. The rest of the subsection is dedicated to explaining the picture with formulas.

3.6.1.4 Some additional setup

For each point $(p,t) \in \mathbb{D}^3 \times [0,1]$, consider the t–invariant, orientation–preserving, linear isometry $\varphi_{(p,t)} : T_{(p,t)}(\mathbb{D}^3 \times \{t\}) \longrightarrow \mathbb{R}^3$ defined by the conditions

$$\varphi_{(p,t)}(Y) = (1,0,0), \quad \varphi_{(p,t)}(Z) = (0,1,0),$$

where $\{Y,Z\}$ is the t–invariant frame we chose for the contact structure $(\mathbb{D}^3 \times \{t\}, \xi)$. The isometries $\varphi_{(p,t)}$ identify the unit sphere $\mathbb{S}T_{(p,t)}(\mathbb{D}^3 \times \{t\}) \cap \xi$ of the contact plane ξ with the horizontal equator $\mathbb{S}^2 \cap \{z=0\}$.

We write

$$\mathrm{Rot}(\theta) = \begin{pmatrix} \cos(\theta) & -\sin(\theta) & 0 \\ \sin(\theta) & \cos(\theta) & 0 \\ 0 & 0 & 1 \end{pmatrix}$$

for the rotation of angle θ around the z–axis.

Fix $\varepsilon > 0$ arbitrarily small. Since $c_0(p,\rho) > c_0(p,0) + C\pi \geq c_0(p,0) + 6\pi$, there exists a (unique) family of constants $h_p \in (0,1)$, $p \in \mathbb{D}^3$, such that $c_0(p,h_p) = c_0(p,0) + 6\pi + 2\varepsilon$. Then, there exists a (unique) family of orientation–preserving diffeomorphisms $H_p : [0, 6\pi + 4\varepsilon] \to [0,h_p]$ such that:

$$f_0 : [-2\varepsilon, 6\pi + 2\varepsilon] \to \mathbb{S}^2$$

$$f_0(t) = (\cos(t), \sin(t), 0) = \mathrm{Rot}(-c(p,0) - 2\varepsilon) \circ \varphi_{(p,t)} \circ X_p^0 \circ H_p(t).$$

The function H_p reparametrises X_p^0 using the value of $c_0(p,t)$. The transformations $\varphi_{(p,t)}$ and Rot take the resulting family of curves to a p–independent family of curves parametrised by arclength that turn along the equator $\{z=0\}$. Essentially, we want to manipulate the curves X_p^0, and it is easier to do it if we put them in this standard position.

3.6.1.5 Introducing the clover

Recall the family $\tau_u : \mathbb{S}^1 \to \mathbb{S}^2$, $u \in [0,1]$ introduced in Lemma 3.71. Regard them as maps $[0,1] \to \mathbb{S}^2$ satisfying $\Gamma_{\tau_u}(0) = \Gamma_{\tau_u}(1) = \mathrm{Id}$ instead. By applying the glueing technical Lemma

3.19 at both their ends, we are able to produce a family of curves $f_u : [-2\varepsilon, 6\pi + 2\varepsilon] \to \mathbb{S}^2$, $u \in [0,1]$, satisfying:

- $f_u(t) = \tau_u(t)$, for $t \in [\varepsilon, 6\pi - \varepsilon]$,
- $f_u(t) = f_0(t)$ for $t \in [-2\varepsilon, -\varepsilon] \cup [6\pi + \varepsilon, 6\pi + 2\varepsilon]$,
- $f_u(t)$ is convex if $(u,t) \in (0,1] \times (-\varepsilon, 6\pi + \varepsilon)$,
- the Frenet frame in the midpoint of the four–leaf clover is

$$\Gamma_{f_1}(3\pi)) = \begin{pmatrix} 1 & 0 & 0 \\ 0 & -1 & 0 \\ 0 & 0 & 1 \end{pmatrix}.$$

3.6.1.6 Flattening at the midpoint

At time $t = 3\pi$ the four-leaf clover f_1 has turned and is pointing in the opposite direction. In order to clockwise pull the two left–most leaves of the four–leaf clover, we first need to flatten this point so that it has an ∞–order tangency with the equator: this is done by applying the flattening technical Lemma 3.20. This allows (after slightly reparametrising the parameter in $u \in \mathcal{O}p(\{1\})$) to smoothly extend the family of curves $f_u : [-2\varepsilon, 6\pi + 2\varepsilon] \to \mathbb{S}^2$ to the domain $u \in [0,2]$; satisfying:

- $f_u(t) = f_1(t)$ for $t \in [-2\varepsilon, 3\pi - \varepsilon] \cup [3\pi + \varepsilon, 6\pi + 2\varepsilon]$, $u \in [1,2]$,
- the curves $f_u(t)$, $u \in [1,2)$, are convex if and only if $t \in (-\varepsilon, 6\pi + \varepsilon)$,
- the curve $f_2(t)$ is convex if and only if $t \in (-\varepsilon, 3\pi) \cup (3\pi, 6\pi + \varepsilon)$,
- $\Gamma_{f_u}(3\pi) = \Gamma_{f_1}(3\pi)$, for $u \in [1,2]$.

3.6.1.7 Pulling the left–most petals

The ∞–order tangency point that we have introduced at $t = 3\pi$ allows us to stretch the point into an arbitrarily large interval, and hence clockwise pull the two left–most leaves. This deformation will occur for those values of the parameter $s \in [2,3]$. Consider an arbitrary constant $C_0 < 0$ to be chosen later which captures the amount of stretching and clockwise pulling.

By applying the stretching technical Lemma 3.21 to the flattened four–leaf clover $f_2 : [0,1] \longrightarrow \mathbb{S}^2$, we extend the family $f_u : [-2\varepsilon, 6\pi + 2\varepsilon] \to \mathbb{S}^2$ to the domain $u \in [0,3]$. For this, a slight reparametrisation in u around $u = 2$ might be necessary. At any rate, the resulting curves satisfy:

- $f_u(t) = f_2(t)$ for $t \in [-2\varepsilon, 3\pi - \varepsilon]$, $u \in [2,3]$,
- $f_u(t) = \mathrm{Rot}(C_0(u-2))f_2(t)$ for $t \in [3\pi + \varepsilon, 6\pi + 2\varepsilon]$, $u \in [2,3]$,
- $f_u(t)$ negatively winds around the horizontal equator $\{z = 0\}$ in the interval $t \in [3\pi - \varepsilon, 3\pi + \varepsilon]$ with non–vanishing speed.

The first two conditions say that, away from $\mathcal{O}p(\{t = 3\pi\})$, $f_3(t)$ is obtained from $f_2(t)$ by taking its second half and rotating it an angle of $C_0 < 0$. The third condition says that in $\mathcal{O}p(\{t = 3\pi\})$, the curve $f_3(t)$ is precisely performing this rotation.

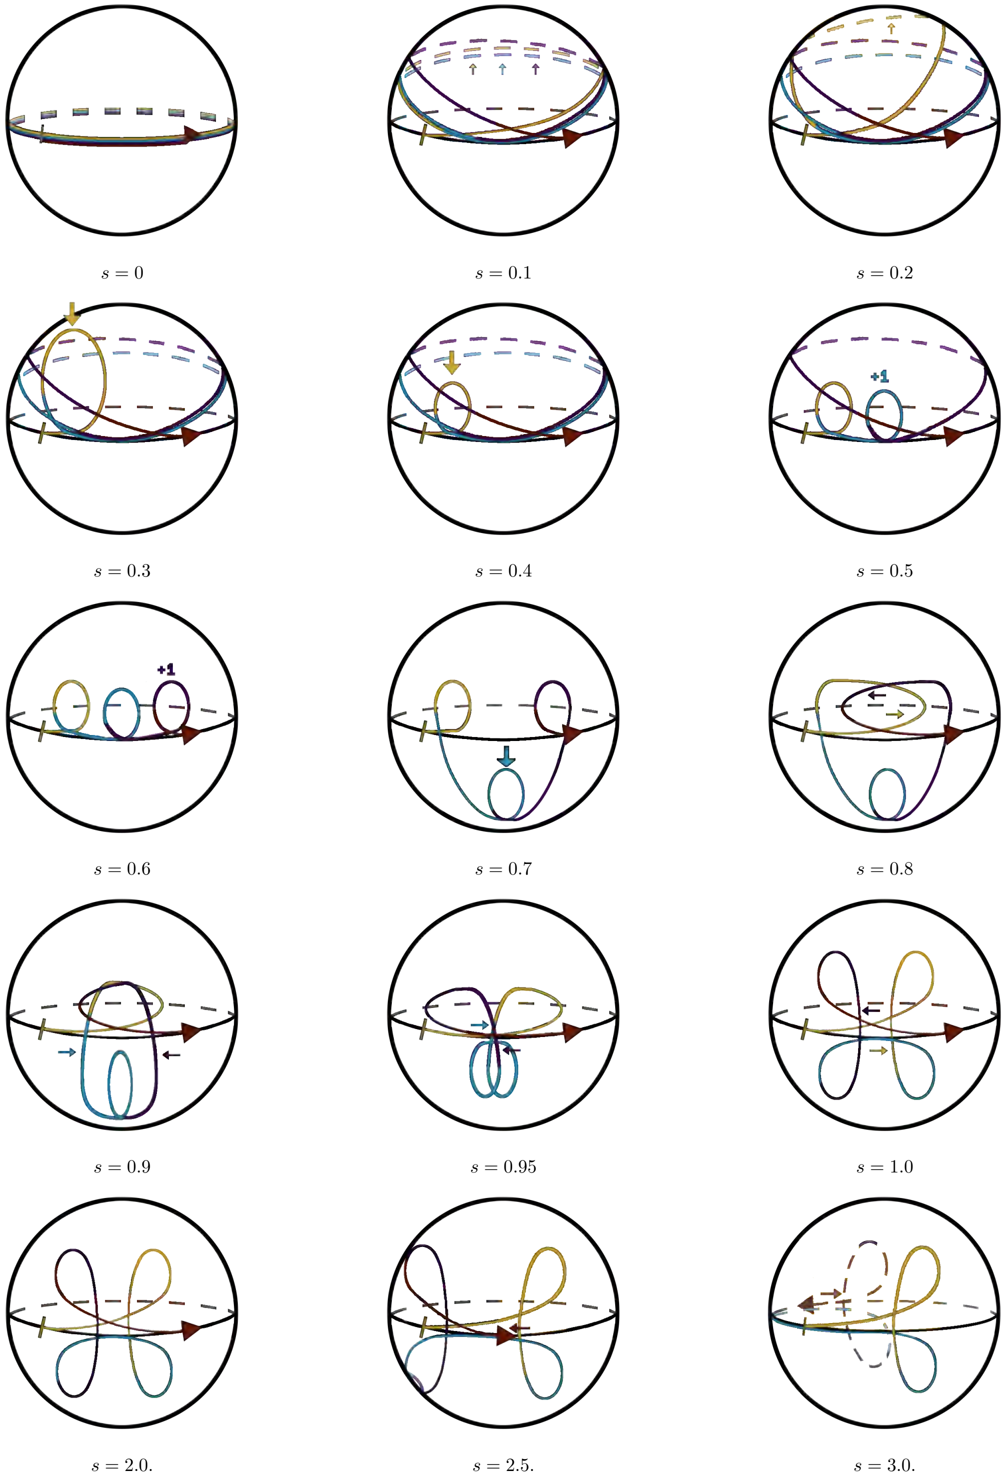

Figure 3.12: The family of curves f_u from the proof of Theorem 3.65. For dramatic effect (and clarity) we have spaced in time the evolution of the three wiggle families. In the formulas, this is done simultaneously.

3.6.1.8 Defining the deformation and checking the Engel condition

The 1–parametric family of curves f_u has been constructed to exactly fit with the desired deformation of X_p^s as $|p|$ varies. Find a bump function

$$\chi : [0,1] \times [0,1] \to [0,3]$$
$$\chi(r,s) = 3 \text{ if } (r,s) \in [0, 1-\delta] \times \mathcal{O}p(\{1\}),$$
$$\chi(r,s) = 0 \text{ if } (r,s) \in [0,1] \times \mathcal{O}p(\{0\}) \cup \mathcal{O}p(\{1\}) \times [0,1],$$
$$\partial_r \chi \leq 0.$$

Which allows us to define, using the correspondence explained in Subsubsection 3.6.1.4, the curves:

$$X_p^s : [0, h_p] \to \mathbb{S}^2$$
$$f_{\chi(|p|,s)}(t) = \mathrm{Rot}(-c_0(p,0) - \varepsilon) \circ \varphi_{(p,t)} \circ X_p^s \circ H_p(t).$$

The curve X_p^s is tangent to ξ in $t = \mathcal{O}p(\{h_p\})$, by construction. It is clear that $X_p^s(h_p) \neq X_p^0(h_p)$; rather, for those $p \notin \mathcal{O}p(\partial \mathbb{D}^3)$, the former is obtained from the latter by rotating negatively around the equator $\mathbb{S}(\xi)$ for some time. This implies that we can find a family of angular functions

$$c_s : \{t \in \mathcal{O}p(\{h_p\}) \subset [0, h_p]\} \subset \mathbb{D}^3 \times [0,1] \to \mathbb{R},$$
$$X_p^s = \cos(c_s(p,t))Y + \sin(c_s(p,t))Z, \qquad t \in [h_p, 1], \quad s \in [0,1]$$

with $c_s(p, h_p) \leq c_0(p, h_p)$. The domain of definition of c_s can be extended to be $\{t \in [h_p, 1]\} \subset \mathbb{D}^3 \times [0,1]$ by setting $c_s(p,t) = c_0(p,t)$ whenever $p \in \mathcal{O}p(\partial \mathbb{D}^3)$ or $t \in \mathcal{O}p(\{1\})$; the extension is unique up to homotopy. This defines a family of 2–distributions $\mathcal{D}_s$ in $\mathbb{D}^3 \times [0,1]$ arising from the curves:

$$X_p^s = \cos(c_s(p,t))Y + \sin(c_s(p,t))Z, \quad \text{if } t \in [h_p, 1].$$

By construction, $c_1(p, h_p) = c_0(p, h_p) - C_0$ if $|p| \leq 1 - \delta$. In particular, if C_0 is chosen large enough, we have $c_1(p, h_p) < c_0(p, 1) = c_1(p, 1)$ everywhere. This implies that Lemma 3.37 can be applied to deform c_1 to a function with $\partial_t c_1 > 0$ everywhere. In particular, the 2–distribution $\mathcal{D}_1$ is Engel in $\{t \in [h_p, 1]\} \subset \mathbb{D}^3 \times [0,1]$.

We claim that this solves the extension problem, concluding the proof of Theorem 3.65. We have to check whether the Engel condition holds in $\{t \in [0, h_p]\}$. By construction, the curves f_u are either convex or C^∞–tangent to the equator $\{z = 0\}$. Whenever $f_{\chi(|p|,s)}$ is convex, so is X_p^s, and the Engel condition holds. Otherwise, if $f_{\chi(|p|,s)}$ is C^∞–tangent to $\{z = 0\}$, X_p^s is C^∞–tangent to the maximal circle $\mathbb{S}(\xi)$. By construction, whenever $f_u(t)$ is C^∞–tangent to the equator, there exists an interval I_u (which is either $[0, u]$ or $[u, 3]$) in which $f_{u'}(t)$, $u' \in I_u$, is also C^∞–tangent to the equator. Equivalently, $X_p^s(t)$ being C^∞–tangent to the equator implies that there there exists an interval I_p containing $|p|$ so that $X_{p'}^s(t)$, $|p'| \in I_p$, is also a C^∞–tangency. We deduce that, for s and t fixed, $\langle X_{p'}^s, (X_{p'}^s)' \rangle$ describes a 2–distribution that agrees with ξ to ∞–order at p, and is thus contact at p. This implies that the Engel condition holds there as well. We have shown that $\mathcal{D}_1$ is a solid shell. □

Remark 3.72. Let us observe that there is a little variation one can do in the proof. In $u \in [1,2]$ the family f_u sees its midpoint flattened to become C^∞–tangent to the equator. Instead of doing this, we could push the midpoint down, away from the equator, and make a little segment around it be tangent to a non–maximal circle in $\mathbb{S}^2$. Then, $u \in [2,3]$, the two left–most petals of f_2 could be pushed as we just did, but now this rotation would be along this non–maximal circle.

It is clear that doing this solves the extension problem as well, because the curves f_u, $u > 0$, are convex everywhere away from their ends, where they agree with the equator.

3.6.2 Filling the 4–ball by turning around

In this subsection we focus on proving Theorem 3.66, i.e. the extension problem for C–contact–type shells is solvable if $C \geq 2$. Our starting point is exactly the same as in the previous subsection (except for the bound on C). We shall use the very same notation, which we briefly recall.

Let $\mathcal{D}_0$ be a C–contact–type shell. Let X_p^0 be the corresponding curves. Let $c_0 : \mathbb{D}^3 \times [0,1] \to \mathbb{R}$ be its angular function. By assumption, $c_0(p,\rho) > c_0(p,0) + C\pi$ for some $\rho \in (0,1)$. We want to produce a homotopy $\mathcal{D}_s$, $s \in [0,1]$, such that the angular function at the end of the homotopy, c_1, satisfies $c_1(p,\rho) < c_1(p,1) = c_0(p,1)$. The idea is similar (but somewhat simpler) than in the previous subsection. We want to follow the equator $\mathbb{S}(\xi)$ negatively for some time; to achieve it, we simply turn around, we follow a curve κ that is C^∞–close to $\mathbb{S}(\xi)$ (but is oriented in the opposite direction), and then we turn to become tangent to $\mathbb{S}(\xi)$ once again.

The curve will remain convex while it turns, ensuring that the Engel condition holds in that region. Thus, it is sufficient to check that it defines an Engel structure as well when it is following κ. Here is where C^∞–closeness is essential: once κ is sufficiently close to $\mathbb{S}(\xi)$ its tangencies define plane fields that are C^∞–close to ξ. Since ξ is a contact structure, this concludes the argument. The whole process can be seen in Figure 3.13.

3.6.2.1 Reminder of notation

In the rest of the subsection we will provide the necessary analytical details of the construction. We recommend the reader to go back and reread the Subsubsections 3.6.1.1 and 3.6.1.4: wiggles and some other notions will reappear in our constructions.

The bare minimum is this. We will construct curves $X_p^s : [0,1] \to \mathbb{S}^2$ describing the desired homotopy $\mathcal{D}_s$, $s \in [0,1]$, with $\mathcal{D}_1$ solid. By hypothesis, we can find a function $h_p : \mathbb{D}^3 \to (0,1)$ such that $c_0(p,h_p) = c_0(p,0) + 2\pi + 2\varepsilon$. Let

$$H_p : [-\varepsilon, 2\pi + \varepsilon] \to [0, h_p]$$

be a family of reparametrisations satisfying:

$$f_0 : [-\varepsilon, 2\pi + \varepsilon] \to \mathbb{S}^2$$

$$f_0(t) = (\cos(t), \sin(t), 0) = \mathrm{Rot}(-c_0(p,0) - \varepsilon) \circ \varphi_{(p,t)} \circ X_p^0 \circ H_p(t).$$

We will focus on defining a homotopy for f_0 that we will then use to produce the homotopy X_p^s of X_p^0 in the interval $[0, h_p]$. The shell $\{t \in [h_p, 1]\}$ will remain a contact–type shell along the homotopy. Showing that the angular function $c_1 : [h_p, 1] \to \mathbb{R}$, corresponding to $\mathcal{D}_1$, satisfies $c_1(h_p) < c_1(1)$ will conclude the proof, just like in the case of the four–leaf clover.

3.6.2.2 Squashing against the equator

Consider the following matrices

$$A_R = \begin{pmatrix} R & 0 & 0 \\ 0 & R & 0 \\ 0 & 0 & 1 \end{pmatrix}$$

which can be used to define the transformation:

$$\mathbb{S}^2 \to \mathbb{S}^2,$$
$$(x,y,z) \to \frac{A_R(x,y,z)}{|A_R(x,y,z)|}.$$

As $R \to \infty$, all the points except the poles are pushed against the equator, uniformly over any compact set disjoint from the poles. Since it is a normalisation of a linear transformation, it preserves geodesics and thus convexity but, clearly, not the metric.

Let $\gamma_p : [0,1] \to \mathbb{S}^2$, $p \in \mathbb{D}^3$ be family of immersions. Let us define sets $U_t, V_t \subset \mathbb{D}^3$, $t \in [0,1]$, and a constant $\tau > 0$, satisfying the following properties:

I. $\{U_t, V_t\}$ is covering of $\mathbb{D}^3$, for all t,

II. if $p \in U_t$, γ_p is graphical over the equator $\{z = 0\}$ at time t and the angle it makes with the meridian passing through $\gamma_p(t)$ is at least τ,

III. if $p \in V_t$, γ_p has no inflection point at time t.

The second property is a quantitative assessment of how graphic γ_p is over the equator at time t: for all p in U_t, the curve γ_p has uniformly bounded slope at time t, when seen as a function over the equator. This property implies:

Lemma 3.73. *Be γ_p, U_t, V_t as above. For each $(p,t) \in \mathbb{D}^3 \times [0,1]$, at least one of the following hold:*

a. the curve $A_R(\gamma_p)$ is convex/concave at time t, for all R,

b. the planes $\langle A_R(\gamma_q(t)), A_R(\gamma_q'(t))\rangle$ C^∞–converge uniformly to the equator as $R \to \infty$, for $q \in \mathcal{O}p(\{p\})$.

Replacing the equator $\{z = 0\}$ by $\mathbb{S}(\xi)$, we want the conclusions of Lemma 3.73 to hold for the family X_p^1, since they imply the Engel condition (after applying A_R with R large).

In a more geometric fashion, the lemma can be rephrased as follows:

Lemma 3.74. *Let $\gamma_p : [0,1] \to \mathbb{S}^2$, $p \in \mathbb{D}^3$. Suppose that, whenever γ_p is tangent to a meridian at time t, there is a neighbourhood $\mathcal{U} \ni p$ such that γ_q is convex/concave at time t if $q \in \mathcal{U}$.*

Then, there are U_t, V_t satisfying properties (I.), (II.), and (III.).

Proof. Indeed, let V_t be the set of $p \in \mathbb{D}^3$ such that γ_p is convex at time t. Let $p \notin V_t$. Then, γ_p is not tangent to a meridian at time t (and, in particular, $\gamma_p(t)$ is not one of the poles). Let $\tau(p,t)$ be the angle that γ_p makes with the corresponding meridian. This angle is a continuous function in the complement of V_t, which is a compact set. Thus, taking U_t a thickening of the complement yields the claim. $\square$

3.6.2.3 Defining the homotopy for the model curve

We will proceed as in Subsection 3.6.1. Recall the family of wiggles:

$$\beta_\theta(t) = (\sin^2(\theta) + \cos^2(\theta)\cos(t), \cos(\theta)\sin(t), \sin(\theta)\cos(\theta)(1-\cos(t))), \quad t \in [0, 2\pi]$$

with $\theta \in [0, \pi/2)$. By Proposition 3.16. Fix $\theta_0 \in (0, \pi/2)$.

Applying the glueing technical Lemma 3.19 to the family β_θ, we construct a family of curves $f_u : [-\varepsilon, 2\pi + \varepsilon] \to \mathbb{S}^2$, $u \in [0, 1]$. Namely:

- $f_0(t) = (\cos(t), \sin(t), 0)$, as above,
- f_u is convex if $t \in (-\varepsilon, 2\pi + \varepsilon)$,
- f_u is C^∞ tangent to $\{z =\}$ in $t = -\varepsilon, 2\pi + \varepsilon$,
- $f_u(t) = \beta_{u\theta_0}(t)$ if $t \in [\varepsilon, 2\pi - \varepsilon]$.

Now we use the stretching technical Lemma 3.21 to pull the leftmost branch of f_1 negatively around the equator. More precisely. Let $C_0 > 0$ be a (sufficiently large) constant to be fixed later on. We can extend our family of immersions (after a slight reparametrisation in the parameter $u \in \mathcal{O}p(\{1\})$) to $f_u : [-\varepsilon, 2\pi + \varepsilon] \to \mathbb{S}^2$, $u \in [0, 2]$, satisfying:

- $f_u(t) = f_1(t)$ for $t \in [-\varepsilon, \pi - \varepsilon]$, $u \in [1, 2]$,
- $f_u(t) = \mathrm{Rot}(-C_0(u - 1))f_1(t)$ for $t \in [\pi + \varepsilon, 2\pi + \varepsilon]$, $u \in [1, 2]$,
- $f_u(t)$ is transverse to all meridians if $t \in [\pi - \varepsilon, \pi + \varepsilon]$, $u \in [1, 2]$.

3.6.2.4 Defining the Engel homotopy

Let $\delta > 0$ be sufficiently small. Construct a bump function

$$\chi : [0, 1] \times [0, 1] \to [0, 2]$$

such that $\chi(r, s)$ is identically 2 in $[0, 1-\delta] \times \mathcal{O}p(\{1\})$, identically 0 in $[0, 1] \times \mathcal{O}p(\{0\}) \cup \mathcal{O}p(\{1\}) \times [0, 1]$, and non–increasing in every interval $[0, 1] \times \{s\}$. Set:

$$\gamma_p^s : [0, h_p] \to \mathbb{S}^2, \quad p \in \mathbb{D}^3, s \in [0, 1],$$

$$\gamma_p^s(t) = f_{\chi(|p|,s)}(t).$$

By construction, the family γ_p^1 satisfies the hypothesis of Lemma 3.74. Then, setting X_p^s to satisfy:

$$A_R(\gamma_p^s(t)) = \mathrm{Rot}(-c_0(p, 0) - \varepsilon) \circ \varphi_{(p,t)} \circ X_p^s \circ H_p(t),$$

with R large enough provides the desired homotopy. Since C_0 can be taken arbitrarily large, the inequality $c_1(p, h_p) < c_1(p, 1)$ can be assumed to hold, concluding the proof. □

3.6.3 Filling the 4–ball using a phone wire

Now we focus on giving a proof for Theorem 3.67. That is, given a C–convex–type shell $\mathcal{D}_0$ with $C \geq 1$, we shall provide a homotopy $\mathcal{D}_s$, $s \in [0, 1]$, relative to the boundary and through convex–type shells, that makes $\mathcal{D}_1$ solid. Write X_p^s for the curves describing $\mathcal{D}_s$ and recall that, by hypothesis:

$$X_p^0(t) = f_p^{[1\#\rho/2]}(t), \quad t \in [0, \rho]$$

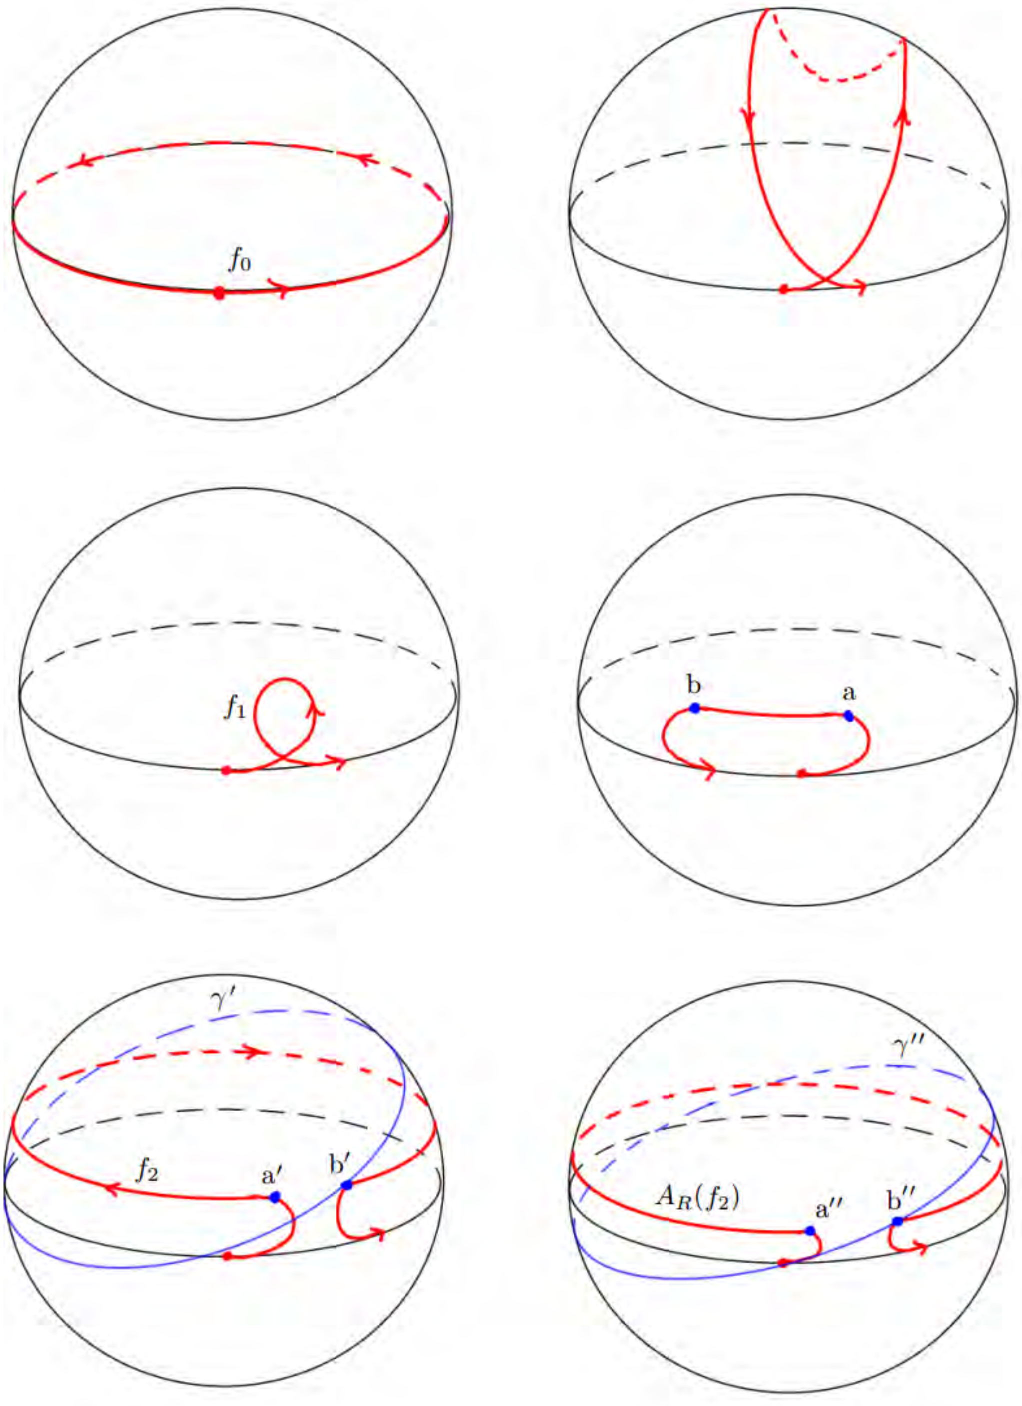

Figure 3.13: The family of curves f_u from Subsection 3.6.2. The points a and b (and their versions with apostrophes) are inflection points of the corresponding curve. We have depicted only two of them, which is the minimum that will appear, but it might happen that there are more and, indeed, in the proof we let that happen. The last figure is obtained from the previous one by squashing everything against the equator. In particular, we see that the maximal circle γ' that is tangent to the curve becomes γ'', which approaches the equator.

with $f_p : [0, \rho] \to \mathbb{S}^2$ a family of convex curves. Define $f_p : [\rho, 1] \to \mathbb{S}^2$ by setting $f_p(t) = X^0_p(t)$.

Now the claim is an immediate application of Proposition 3.29: this result states that, since we have a wiggle at time $\rho/2$, we can apply Little's homotopy to produce arbitrarily many new wiggles that we then distribute evenly along the domain $[0, 1]$, achieving convexity. Indeed, we obtain a family of curves:

$$f_{p,u} : [0, 1] \to \mathbb{S}^2, \quad u \in [0, u_0]$$

satisfying

- $f_{p,0} = f_p^{[1\#\rho/2]} = X^0_p$,
- $f_{p,u}(t) = X^0_p(t)$ if $t \in \mathcal{O}p(\{0, 1\})$,
- $f_{p,u}$ is everywhere convex if $|p| > 1 - \delta$, for $\delta > 0$ small,
- f_{p,u_0} is everywhere convex.

Then, set $X^s_p = f_{p,\chi(p,s)}$ with $\chi : \mathbb{D}^3 \times [0, 1] \to [0, u_0]$ satisfying

$$\chi(p, s) = 0 \quad \text{if } p \in \mathcal{O}p(\partial\mathbb{D}^3) \text{ or } s = 0,$$

$$\chi(p, s) = u_0 \quad \text{if } p \in D^3_{1-\delta}, s = 1.$$

Then, $\mathcal{D}_1$ is an Engel structure, since the curves X^1_p are convex. $\square$

3.7 Proving the Engel existence theorem and its corollaries

The π_0–statement of Theorem 3.1, that is, *every formal Engel structure can be deformed through formal Engel structures to an Engel structure*, is a consequence of putting together the reduction process from Theorem 3.54 with the filling methods of Theorems 3.65 or 3.66 or, alternatively, putting together Theorems 3.55 and 3.67. In both cases we need the results on domination of shells as an intermediate step: Propositions 3.44 and 3.50, respectively. In order to prove the statement for higher π_k, it is convenient to understand it as a particular version of the *foliated* existence theorem. For this reason, we shall prove Theorem 3.1 and Corollary 3.5 simultaneously.

3.7.1 Proof of Theorem 3.1 and Theorem 3.5

Consider a $\mathbb{S}^k$–family of formal foliated Engel structures $(\mathcal{W}_x, \mathcal{D}_x, \mathcal{E}_x)$, $x \in \mathbb{S}^k$, in a smooth foliated manifold $(M^{m+4}, \mathcal{F}^4)$. The product manifold $N = M \times \mathbb{S}^k$ is endowed with the product foliation $\mathcal{F}_N = \coprod_{x\in\mathbb{S}^k} \mathcal{F} \times \{x\}$ and then the family $\{(\mathcal{W}_x, \mathcal{D}_x, \mathcal{E}_x)\}_{x\in\mathbb{S}^k}$ can be understood as a formal foliated Engel structure $(\mathcal{W}, \mathcal{D}, \mathcal{E})$ in the foliated manifold $(N^{4+m+k}, \mathcal{F}^4_N)$. Homotoping this formal foliated Engel flag to a genuine foliated Engel flag amounts to deforming the original family of formal foliated Engel structures to a family of genuine foliated Engel structures.

In consequence, the π_0–surjectivity of Corollary 3.5 applied to the formal foliated Engel structure $(N^{m+4+k}, \mathcal{F}^4_W, \mathcal{W}, \mathcal{D}, \mathcal{E})$ implies the π_k–surjectivity for the formal foliated Engel structure $(M^{m+4}, \mathcal{F}^4, \mathcal{W}, \mathcal{D}, \mathcal{E})$. Since Theorem 3.1 is a particular case of Corollary 3.5 ($m = 0$) it suffices to discuss the proof of the latter.

The two central ingredients in the proof of Corollary 3.5 are the parametric counterparts of the reduction and extension methods. We shall detail them in the next two subsubsections.

3.7.1.1 The reduction process in the foliated case

First we need to define Engel shells in the foliated setting:

Definition 3.75. *A formal foliated Engel structure*

$$(\mathbb{D}^3 \times [0,1] \times \mathbb{D}^m, \coprod_{x \in \mathbb{D}^m} \mathbb{D}^3 \times [0,1] \times \{x\}; \mathcal{W}, \mathcal{D}, \mathcal{E})$$

is said to be a ***foliated Engel*** (***contact–type,*** C***–contact–type, convex–type, or*** C***–convex–type***) ***shell*** *if:*

a. $(\mathbb{D}^3 \times [0,1] \times \{x\}, \mathcal{W}, \mathcal{D}, \mathcal{E})$ *is an Engel (contact–type, C–contact–type, convex–type, or C–convex–type) shell for all $x \in \mathbb{D}^m$,*

b. $(\mathbb{D}^3 \times [0,1] \times \{x\}, \mathcal{W}, \mathcal{D}, \mathcal{E})$ *is solid for $x \in \mathcal{O}p(\partial\mathbb{D}^m)$.*

A foliated Engel shell is ***solid*** *if its formal foliated Engel structure is a foliated Engel structure.*

Note that the role of the parameter space in these definitions is being played by the m–disc $\mathbb{D}^m$. Similarly, we can define flowboxes:

Definition 3.76. *A formal foliated Engel structure*

$$(\mathbb{D}^3 \times [0,1] \times \mathbb{D}^m, \coprod_{x \in \mathbb{D}^m} \mathbb{D}^3 \times [0,1] \times \{x\}; \mathcal{W}, \mathcal{D}, \mathcal{E})$$

is said to be a ***foliated*** (***contact–type, convex–type***) ***flowbox*** *if the restriction of $(\mathcal{W}, \mathcal{D}, \mathcal{E})$ to $(\mathbb{D}^3 \times [0,1] \times \{x\}$is a (contact–type, convex–type) flowbox for all $x \in \mathbb{D}^m$.*

Theorems 3.54 and 3.55 both generalise. They read as follows:

Proposition 3.77. *Let $(N^{4+m}, \mathcal{F}^4; \mathcal{W}_0, \mathcal{D}_0, \mathcal{E}_0)$ be a formal foliated Engel structure and $C \in \mathbb{R}^+$ a constant. Then, there exists a homotopy of formal foliated Engel structures $(\mathcal{W}_s, \mathcal{D}_s, \mathcal{E}_s)$, $s \in [0,1]$, and a collection of $(4+m)$–discs $B_1, \dots, B_p \subseteq W$ such that:*

1. $(\mathcal{W}_1, \mathcal{D}_1, \mathcal{E}_1)$ *is a foliated Engel structure in the complement of $W \setminus \bigcup_{i=1}^p B_i$.*

2. *For each $i \in \{1, \dots, p\}$, $(B_i, \mathcal{F}|_{B_i}; \mathcal{W}_1, \mathcal{D}_1, \mathcal{E}_1)$ is a foliated C–contact–type shell.*

Proof. We shall outline the general argument, indicating any differences with the non–parametric and non–foliated case.

First, do note that McDuff's theorem proving the complete h–principle for even–contact structures implies that a complete h–principle holds for *foliated even–contact* ones as well. From our previous discussion on how the parametric and foliated cases are related, this is clear.

Secondly, observe that Theorem 3.9, the theorem that provides a triangulation and a cover adapted to $\mathcal{W}$, was proven in all dimensions. In particular, it can be applied to N. We claim that we the obtain a triangulation $\mathcal{T}$ and an associated cover by sets $\{\mathcal{U}(\sigma)\}_{\sigma \in \mathcal{T}}$ such that the neighbourhoods $\mathcal{U}(\sigma) \cong \mathbb{D}^3 \times [0,1] \times \mathbb{D}^m$ are, at the same time, flowboxes for the line field $\mathcal{W}$ and foliated charts for the foliation $\mathcal{F}$. This additional condition can be achieved by requiring in its proof that we first follow the exponential flow in the leaf and then in the ambient manifold. Each one of them is then a foliated contact–type flowbox.

Thirdly, we try to deform the angular functions in each foliated contact–type flowbox to ensure that the height of the resulting shells is at least C. Angular functions behave in exactly the same way in this foliated setting; we have:

$$c(p,t,x) : \mathbb{D}^3 \times [0,1] \times \mathbb{D}^m \to \mathbb{R}$$

and they can be made to have arbitrarily large derivative, proving the claim.

This shows that, for each top dimensional shell σ, the corresponding foliated Engel shell $\mathcal{U}(\sigma) \cong \mathbb{D}^3 \times \mathbb{D}^m \times [0,1]$ is of contact–type and each Engel shell $\mathbb{D}^3 \times \{x\} \times [0,1]$ has height at least C. Now we have to prove that this foliated shell dominates a shell that is actually of C–contact type. The point is that Lemma 3.36, that says that the angular function is determined, up to homotopy, by its value at the ends $t \in \mathcal{O}p(\{0,1\})$, holds as well in this setting. Proceeding like we did when we discussed domination between convex–type shells yields the claim. This concludes the proof. $\square$

Proposition 3.78. *Let $(N^{4+m}, \mathcal{F}^4; \mathcal{W}_0, \mathcal{D}_0, \mathcal{E}_0)$ be a formal foliated Engel structure and $C \in \mathbb{R}^+$ a constant. Then, there exists a homotopy of formal foliated Engel structures $(\mathcal{W}_s, \mathcal{D}_s, \mathcal{E}_s)$, $s \in [0,1]$, and a collection of $(4+m)$–discs $B_1, \dots, B_p \subseteq W$ such that:*

1. *$(\mathcal{W}_1, \mathcal{D}_1, \mathcal{E}_1)$ is a foliated Engel structure in the complement of $W \setminus \bigcup_{i=1}^p B_i$.*
2. *For each $i \in \{1, \dots, p\}$, $(B_i, \mathcal{F}|_{B_i}; \mathcal{W}_1, \mathcal{D}_1, \mathcal{E}_1)$ is a foliated C–convex–type shell.*

Proof. Fix a line field $\mathcal{Y} \subset \mathcal{D}_0$ transverse to $\mathcal{W}_0$. Recall Theorem 1.15, it states that 2–distributions that are not–integrable (but possibly non–maximally) satisfy a full h–principle. In particular, we deduce that a foliated analogue holds. We can apply it to $\mathcal{D}_0$ with respect to $\mathcal{Y}$. Then, Theorem 3.9 provides a triangulation $\mathcal{T}$ and a cover by foliated convex–type flowboxes adapted to the line field $\mathcal{Y}$.

Now we have to deform the curves to ensure that we obtain convex–type shells having arbitrarily large winding. This is done proceeding flowbox by flowbox adding arbitrarily many wiggles to $X_{p,x}$ and then interpolating back to whatever we had in the boundary of the flowbox, and relative to the areas we have already fixed, using again Theorem 1.15 (which in terms of curves is simply the Smale–Hirsch theorem). No differences with the the proof of Theorem 3.55.

The last step is passing from foliated convex–type shells having large winding to C–convex–type shells. The key ingredient was the technical Proposition 3.18; it allows to deform a family of convex curves into a family where all the curves agree with $f^{[C\#\rho/2]}$ away from their ends, where f is some convex curve, and C is computed from the winding of the family as $\mathrm{wind}(\mathcal{D}) = 2C+1$. This works equally well in this new setting, and allows us to prove the analogue of the domination Proposition 3.50. This concludes the proof. $\square$

3.7.1.2 The foliated extension problem

Having explained the reduction process, we can turn to filling the last balls:

Proposition 3.79. *A foliated C–contact–type shell, $C \geq 2$, is homotopic through foliated Engel shells to a solid foliated Engel shell.*

We shall prove it using the turning model of Subsection 3.6.2. It will be clear from the proof that it can be done with 3.6.1 as well.

Proof. Denote by $(\mathcal{W}, \mathcal{D}, \mathcal{E})$ the shell we start with. Let $X_{p,x}$, $p \in \mathbb{D}^3$, be the curves that describe the contact–type shell $(\mathcal{W}, \mathcal{D}, \mathcal{E})|_{\mathbb{D}^3 \times [0,1] \times \{x\}}$.

Requiring for the shell $(\mathcal{W}, \mathcal{D}, \mathcal{E})$ to be of C–contact–type means that the curves $X_{p,x}$ describe at least C projective turns in the band $t \in [0, \rho]$, for some ρ. We can then find $h : \mathbb{D}^3 \times \mathbb{D}^m \to (0, 1)$ a smooth function such that $X_{p,x}|_{[0,h(p,x)]}$ turns $2\pi + 2\varepsilon$, exactly. In Subsection 3.6.2 we defined a family of curves f_u, $u \in [0, 2]$, that interpolate between $X_{p,x}|_{[0,\rho]}$ (or rather, its normalised version f_0) and a curve that turns, follows a non–maximal circle, and then turns again. We can use the homotopy f_u to define the corresponding deformation of $X_{p,x}$. □

The counterpart of Theorem 3.67 reads:

Proposition 3.80. *A foliated C–convex–type shell, $C \geq 1$, is homotopic through foliated convex–type shells to a solid foliated Engel shell.*

Proof. Denote by $(\mathcal{W}, \mathcal{D}, \mathcal{E})$ the shell we start with. Let $X_{p,x}$, $p \in \mathbb{D}^3$, be the curves that describe the contact–type shell $(\mathcal{W}, \mathcal{D}, \mathcal{E})|_{\mathbb{D}^3 \times [0,1] \times \{x\}}$.

Saying that the shell $(\mathcal{W}, \mathcal{D}, \mathcal{E})$ is of C–contact–type means that the curves $X_{p,x}|_{[0,\rho]}$ are actually of the form $f_{p,x}^{[C\#\rho/2]}$, with $f_{p,x}$, $(p, x) \in \mathbb{D}^3 \times \mathbb{D}^m$, a family of convex curves. Applying Proposition 3.29 works exactly as in the non–foliated setting. □

3.7.1.3 The proof

It is now immediate that Propositions 3.77 and 3.79 can be brought together to prove the π_0 version of Corollary 3.5. By our discussion above, this proves also the parametric version, and the parametric version of Theorem 3.1.

3.7.2 The proof of Corollaries 3.2 and 3.6

As we pointed out in the introduction, Corollary 3.2 should be understood as a relative existence statement. The key ingredient we need to show is that the reduction process can be adapted to have a relative character (with respect to the contact boundaries). Once this is done, the extension theorems imply our claim.

We shall prove that Theorem 3.54 (the reduction theorem that provides contact–type shells of large height) holds in this relative sense. We leave it as an exercise to the reader to show that Theorem 3.55 (the one that provides convex–type shells) can also be adapted to this setting.

Consider a collar neighbourhood $\mathcal{O}p(\partial M) \cong \partial M \times [0, 1)$ of the boundary, and thicken the filling M to

$$\overline{M} := M \cup_{\partial M \times \{0\}} \partial M \times [-\varepsilon, 0];$$

we do this for technical reasons: this way we can think of $\mathcal{O}p(M)$ as an open manifold without boundary that is contained in $\overline{M}$, and we do not have to redo some of the arguments to deal with the boundary.

By hypothesis, we can assume that $\mathcal{E}$ is even–contact in the collar neighbourhood we have selected, and further assume that the lines $\{p\} \times [-\varepsilon, 1)$ are tangent to the kernel $\mathcal{W}$. An application of McDuff's h–principle for even–contact structures allows us to extend $\mathcal{E}$ to the interior of M as a genuine even–contact structure.

Corollary 3.11, provides a triangulation $\mathcal{T}$ and a cover $\{\mathcal{U}(\sigma) \subset \overline{M}\}$ of M that are adapted to $\mathcal{W}$ and are relative with respect to ∂M. The simplices contained in ∂M are no different than any other, so we can simply apply the rest of the reduction argument of Theorem 3.54 to the flowboxes $\mathcal{U}(\sigma)$. This concludes the proof of the corollary.

It is also immediate that Corollary 3.6, the foliated analogue of Corollary 3.2, can be obtained from Corollary 3.5. □

3.7.3 The proof of Corollary 3.7

Let us go back to the proof of Theorem 3.1, and focus on the proof through convex–type shells, step by step. First we homotope $\mathcal{D}$ so that it is non–integrable: this relies on the Hirsch–Smale theorem and thus leaves the vector field $\mathcal{Y}$ fixed. Then, we go flowbox by flowbox deforming the vector field X_p describing the 2–distribution, it is thus clear that the vector field ∂_t of the flowbox (which spans $\mathcal{Y}$) is always part of $\mathcal{D}$. This is also true when we show the domination statement to pass from a convex–type shell to a C–convex–type one and when we finally solve the extension problem. Thus, we conclude that any $\mathcal{Y}$ that we start with is actually tangent to the resulting Engel structure. Since the structure is described by convex curves when seen through the lense of a $\mathcal{Y}$–flowbox, we deduce that $\mathcal{Y}$ is actually *transverse* to the kernel. This is the content of Corollary 3.7.

Observe that, if M admits a parallelisation, *any* vector field can be made to be tangent to an Engel structure. Indeed, the existence of a parallelisation immediately induces and almost–quaternionic structure on M which can be used to extend $\mathcal{Y}$ to a complete flag.

Chapter 4

Classification results for Engel structures

In [10] it was stated, without proof, that it is possible to define a subclass of Engel structures whose inclusion into the formal ones is a weak homotopy equivalence. This is the question that we want to tackle in this chapter.

Let us motivate what we are doing by first going through the contact topology case. Suppose we only cared about the π_0 case. Then, by Gray stability, any two contact structures having an overtwisted disc and being formally homotopic are actually isotopic. This singles out certain connected components as overtwisted and some others as tight. This dichotomy, has to be stressed, is a pure π_0 phenomenon. To obtain a weak homotopy equivalence into the formal contact structures, one needs to actually restrict to those overtwisted contact structures/ formal contact structures whose overtwisted disc has been fixed! This implies that the homotopy groups of an overtwisted component are not necessarily those of the corresponding component of formal contact structures. Some interesting results in this direction can be found in [83].

A class of Engel/contact structures can be reasonably called overtwisted if it satisfies that its inclusion into the formal ones is a weak homotopy equivalence. Consider the following alternate way of thinking about flexibility: a compact family of structures that is homologically trivial formally is called *loose* if it is *homologically trivial geometrically*; otherwise, the family is said to be *tight*. In the contact case we know many examples of tight families: any pair of contact structures lying in the same formal class, one being tight (in the standard sense) and the other overtwisted. In the Engel case, constructing a single instance of tightness remains the most important open question.

Our intention in this chapter is to understand looseness better. Our hope is that this will help to find potential candidates of *tight* families. The injectivity h–principle holds in π_0 for knots tangent to an Engel structure (see Corollary 5.21): *is the same true for Engel structures?* Maybe a tight family must necessarily be higher dimensional.

The chapter is structured as follows. In Section 4.2 we will give several definitions of looseness. Not very surprisingly, they rely on imposing for the Engel structures in question to have large contact/convex–type energy. In Section 4.3, we will show that all the definitions of looseness are essentially equivalent. We will then prove, in Section 4.4, that two loose families that are homotopic formally are homotopic geometrically (through loose families).

A subtle point is that a family of loose structures will not be loose itself, necessarily. This is a problem towards providing a complete h–principle for loose structures (holding for all π_k). [1]. However, the π_0 h–principle still goes through. We will say that an Engel structure is (0–)overtwisted if it is loose and we will show that **the space of overtwisted Engel structures has the same connected components as the space of formal Engel structures**; this is the content of Corollary 4.25.

Is it possible that Engel structures satisfy a complete h–principle and, therefore, the best possible choice for the overtwisted class is the whole space of Engel structures? Of course we do not know the answer to this question. However, it is worth pointing out that all of our characterisations of looseness rely on a *global property*: they depend on the behaviour of the Engel structure with respect to a vector field contained in it. This contrasts heavily with the contact case, in which the condition to check is of semi–local nature (i.e. we just need to find an overtwisted disc). One goal to work towards would be to define Engel overtwistedness/looseness by saying that some local model can be found in the manifold.

As a little warm–up to all of this, in Section 4.1, we shall tackle the classification of prolongations through Engel homotopies (as explored in my article [14]). Using the machinery presented in the previous chapter, these results will be obtained almost for free.

4.1 Prolongations and Engel homotopies

In this section we will explain a couple of examples of Engel homotopies that can be performed to display flexibility for prolongations. The first example deals with Cartan prolongations and shows that homotopies through *formal* Cartan prolongations can be approximated by homotopies through Lorentz prolongations. The second example shows how Saldanha's work can be applied to produce homotopies when the underlying contact plane (or space–like plane) is trivial as a bundle.

We shall see later that these statements do not tell the full story: Corollary 4.9 states that families of Cartan prolongations are *completely* classified by their formal type. Proving this complete statement is more complicated and we believe this simpler statements give some intuition of what we are doing.

Let us fix some notation. In this section N will be an oriented 3–manifold with a fixed parallelisation, which yields a particular choice of metric. We let $\mathcal{T}$ be a triangulation of N. Fix $c \in H^2(N, \mathbb{Z})$. Denote by $N(c)$ the total space of the $\mathbb{S}^1$–bundle of Chern class c over N; the orientation of N naturally orients $N(c)$. We still invite the reader to refer to Chapter 2 if some of the objects are not familiar.

4.1.1 Example I. Homotopies through Lorentz prolongations

We want to prove the following statement:

Theorem 4.1. *Let K a topological space. Let $\phi_0, \phi_1 : K \to \mathfrak{Cartan}(c)$ be two continuous maps and let $\phi_s : K \to \mathcal{F}\mathfrak{Cartan}(c)$, $s \in [0,1]$, be a homotopy between them. Then ϕ_s can be C^∞–approximated, relative to its ends, by a homotopy with image in $\mathfrak{Engel}(N(c))$.*

[1]We expect a complete h–principle to hold for loose structures. This should follow from a better understanding of the geometry of convex curves in $\mathbb{S}^2$

For a proof by picture, refer to Figure 4.1.

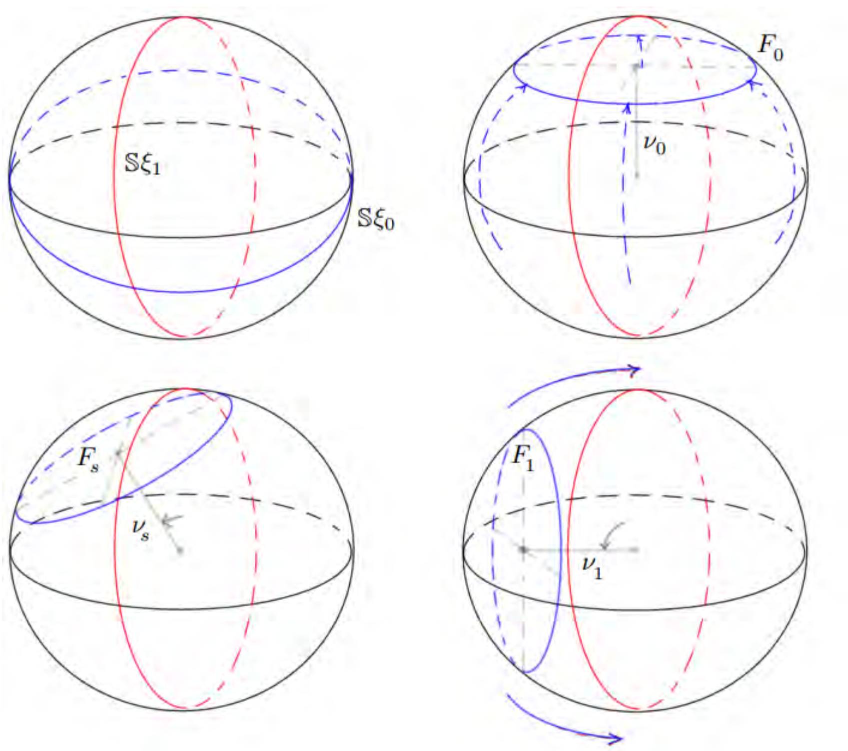

Figure 4.1: In the first figure, the red and blue equators are the sphere bundles of some contact structures ξ_0 and ξ_1. Then, the blue family is pushed upwards to be given by a convex curve describing a Lorentz prolongation (step II). Then, it can be readily moved through Lorentz prolongations to become a convex push-off of the red equator (steps III and IV), which we then flatten (blue arrows in step IV).

Proof. Consider the maps ϕ_s. There is a corresponding family $K \times [0,1]$ of maps

$$f_{x,s} : N(c) \to \mathbb{S}TN$$

$$f_{x,s}(p,L) = d\pi_{p,L}([\phi_s(x)])$$

which are simply the tautological maps associated to each formal Cartan prolongation. Write $\xi_{x,s}$ for the oriented contact plane associated to $\phi_s(x)$.

Write $\nu_{x,s}$ for a family of unit vectors in N such that $\nu_{x,s}$ is orthogonal to $\xi_{x,s}$ for each $(x,s) \in K \times [0,1]$; there is only one choice making $(\xi_{x,s}, \nu_{x,s})$ positively oriented. Define a function $h : [0,1] \to \mathbb{R}$ vanishing to all orders on 0 and 1 and otherwise satisfying $h(s) > 0$, $s \in (0,1)$. Consider the following deformation of f:

$$F_{x,s} = \frac{f_{x,s} + h(s)\nu_{x,s}}{|f_{x,s} + h(s)\nu_{x,s}|}.$$

The tautological distributions associated to $F_{x,s}$ provide a family $\psi_s : K \to \mathfrak{Engel}(N(c))$, $s \in [0,1]$; the Engel structures ψ_s are Lorentz prolongations if and only if $s \in (0,1)$. Making $h(s)$ approach zero, ψ_s becomes arbitrarily close to ϕ_s. $\square$

Let us spell out what the theorem is stating. Take a family $\phi_0 : K \to \mathfrak{Cartan}(c,k)$. From it, we obtain a family of contact structures $\pi \circ \phi_0 : K \to \mathfrak{Cartan}(c,k) \to \mathcal{C}\text{-}\mathcal{S}\mathfrak{trs}(kc)$. Then, any homotopy, as plane fields, of the family $\pi \circ \phi_0$ lifts, by the homotopy lifting property, to a homotopy of ϕ_0 in $\mathcal{F}\mathfrak{Cartan}(c,k)$. If the resulting family of plane fields is actually contact (but possibly not homotopic to the original family in the contact category), the theorem implies that the corresponding Cartan prolongations are homotopic through Lorentz prolongations (and therefore, homotopic as Engel families).

We therefore have that, as long as we stay in the class of Cartan prolongations, they do recover all the contact topology information they arose from. However, as soon as we allow more general Engel homotopies (in this case, still fairly restrictive ones: homotopies through Lorentz prolongations), this is not possible anymore.

We can revisit the results of Chapter 2 using this result. Proposition 2.10 showed that two formal Cartan prolongations, whose only difference was some even horizontal distance, were homotopic as formal Cartan prolongations. From this, an application of Eliashberg's theorem helped us prove the analogous result for Cartan prolongations of overtwisted contact structures. A consequence of Theorem 4.1 is the following corollary:

Corollary 4.2. *Any two Cartan prolongations in $\mathfrak{Cartan}(c,k)$ are homotopic through Engel structures if and only if they lift contact structures homotopic as plane fields and their horizontal distance vanishes mod 2.*

Any two loops of Cartan prolongations in $\mathfrak{Cartan}(c,k)$ are homotopic through Engel structures if and only if they lift loops of contact structures that are homotopic as loops of plane fields and additionally the parity of the difference of their looping numbers is zero.

As we remarked at the beginning of the chapter, the horizontal distance and the looping are well–defined mod 2 even between two elements/loops that are different.

4.1.2 Example II. Prolongations of trivial bundles

Theorem 4.3. *Let K a CW–complex of dimension at most m. Let $\phi_0, \phi_1 : K \to \mathfrak{Engel}(N(0))$ be two continuous maps, with image either in the orientable Cartan prolongations or in the Lorentz prolongations.*

Assume their turning numbers are at least 2. Then, they are Engel homotopic if and only if they are formally homotopic.

The assumption $c = 0 \in H^2(N, \mathbb{Z})$ is made for simplicity. We will later show that prolongations are flexible without assumptions on the turning number (Corollary 4.9), so there is no point in dealing with the general case in this example.

Proof. We break down the proof in three steps, the first of which is some simple setup.

Deformation into Lorentz prolongations. Assume first that ϕ_i are Cartan prolongations. Then, $\phi_i(x)$, $x \in K$, $i = 0,1$, defines an oriented contact plane $\xi_{x,i}$ and a tautological map $f_{x,i} : N(0) \to \mathbb{S}TN$. We can take $\nu_{x,i}$ to be the unique vector field such that $(\xi_{x,i}, \nu_{x,i})$ is positively oriented. Then $(f_{x,i} + \nu_{x,i})/|f_{x,i} + \nu_{x,i}|$ defines a family ψ_i of Lorentz prolongations. This family is Engel homotopic to ϕ_i. Furthermore, the choice of $\nu_{x,i}$ (as opposed to its negative), implies that when we follow the fibres of $N(0)$ positively, the curves that describe $\psi_i(x)$ are convex (i.e. having positive curvature).

Assume otherwise that ϕ_i is a Lorentz prolongation. If the curves describing it are convex, we are done. Otherwise, consider the tautological map $f_{x,i} : N(c) \to \mathbb{S}TN$. There are a plane field $\xi_{x,i}$ and a vector field $\nu_{x,i}$ transverse to it so that ϕ_i is precisely given by $(f_{x,i}+\nu_{x,i})/|f_{x,i}+\nu_{x,i}|$. We can then find homotopies $\xi_{x,i,s}$ and $\nu_{x,i,s}$, $s \in [0,1]$, so that:

- $\xi_{x,i,0} = \xi_{x,i}$ and $\nu_{x,i,s} = \nu_{x,i}$,
- $\nu_{x,i,s}$ is transverse to $\xi_{x,i,s}$,
- $\xi_{x,i,1}$ is contact and overtwisted.

Set $\phi_{i,s}$ to be the Lorentz prolongation obtained by pushing the formal Cartan prolongation of $\xi_{x,i,s}$ with $\nu_{x,i,s}$. This provides a homotopy of ϕ_0 through Lorentz prolongations. Now, $\phi_{i,1}$ is clearly Engel homotopic to the Cartan prolongation of $\xi_{x,i,1}$ and we can apply the previous discussion. Effectively, we pass through Cartan prolongations to go from *concave* curves to *convex*.

Let us henceforth assume that we are dealing with Lorentz prolongations ψ_i, $i \in \{0,1\}$, described by convex curves.

The formal data. As a map into $\mathcal{F}\mathfrak{Engel}(N(c))$, ψ_i, $i \in \{0,1\}$, can be thought of as a map $K \times N(0) \to \mathrm{SO}(4)$, by lifting the parallelisation of N to the obvious parallelisation of $N(0) \cong N \times \mathbb{S}^1$. Being a bit more precise, the map into SO(4) can be given by applying the Gram–Schmidt process to the (local) frames $\{\partial_t, X, X', X''\}$ with $\psi_i(x) = \langle \partial_t, X \rangle$. Note that this differs from our usual choice of Engel parallelisation: we have reversed the first two vectors because it is a bit more suited to this Lorentzian setting.

As a manifold, SO(4) is simply the product $\mathbb{S}^3 \times \mathrm{SO}(3)$; this can be shown using a quaternionic structure. Then, the induced map $K \times N(0) \to \mathbb{S}^3$ into the first factor can be assumed to be constant for both ψ_0 and ψ_1 because it corresponds to the fibre direction. If we assume that the ψ_i are formally homotopic, this discussion implies that there is a formal homotopy $\psi_s : K \to \mathcal{F}\mathfrak{Engel}(N(c))$, $s \in [0,1]$, such that all the $\psi_s(x)$ contain the fibre direction as well.

In particular, each ψ_s, $s \in [0,1]$, is described by a map $N \times K \to \mathcal{F}\mathcal{I}^f$. By applying the Hirsch–Smale theorem, which gives a full (in particular, relative) h–principle, we can further assume that ψ_s maps into $\mathcal{I}^f$. This means that the 2–distributions given by ψ_s are non–integrable, but not necessarily maximally. By applying a family of SO(3)– transformations, instead of $\mathcal{I}^f$ we can work with $\mathcal{I}$.

Deformation as families of convex curves. Now the claim is immediate. The maps ψ_0 and ψ_1 are each described by a $N \times K$ family of curves in $\mathcal{L}$. If the turning numbers are at least 2, these families are immediately loose. ψ_s provides a homotopy as a family of immersions, which can be turned into a homotopy through families of convex curves by Proposition 3.25. This concludes the argument. □

4.2 Loose families

We start by characterising looseness in three different ways. They all reflect the idea that large height/winding provides flexibility. The first one will be the most practical one: it will rely on triangulations and will be adequate for proving that two loose families are homotopic if they lie in the same formal class. The second one will be more geometrically meaningful: it relies on less auxiliary data and will allow us to completely classify prolongations. The third one is invariant under Engel deformations and, therefore, has better properties than the other two. These definitions are given in Subsections 4.2.1, 4.2.2, and 4.2.3, respectively.

The proof that all three definitions are equivalent, up to homotopy, is given in the next section, in Theorem 4.15. With this done, it will be almost immediate that loose families having the same underlying formal data are homotopic (Theorem 4.22).

4.2.1 First definition of looseness

The simplest (or only?) way towards defining looseness is by simply giving the appropriate conditions for the existence theorem to work in a relative fashion. The key point in Theorem 3.1 was that convex–type shells of large winding allow us to solve the extension problem. If we expect to solve the extension problem simultaneously for a 1–parametric family of shells, relative to the ends, the ends better have large winding themselves. This motivates the following definition:

Definition 4.4. *Let M be a 4–manifold. Let K be a compact manifold.*

*A family $\mathcal{D} : K \to \mathfrak{Engel}(M)$ is said to be t–**loose** if there exists a family of line fields $\mathcal{Y}(k) \subset \mathcal{D}(k)$, and a triangulation $\mathcal{T}$ and a cover $\{\mathcal{U}(\sigma)\}_{\sigma\in\mathcal{T}}$ of $M \times K$, satisfying:*

- *$\mathcal{Y}(k) \subset \mathcal{D}(k)$ is transverse to the kernel $\mathcal{W}(k)$,*
- *regarding $\mathcal{Y}$ as line field in $M \times K$, $\mathcal{T}$ and $\{\mathcal{U}(\sigma)\}$ are adapted to it (as in Definition 3.8),*
- *for each $(4 + \dim(K))$–simplex σ,*

$$\phi(\sigma) : \mathcal{U}(\sigma) \to \mathbb{D}^3 \times \mathbb{D}^{\dim(K)} \times [0, 1]$$

is a (solid and foliated) convex–type shell described by curves

$$X_{p,x} : [0, 1] \to \mathbb{S}^2, \qquad (p, x) \in \mathbb{D}^3 \times \mathbb{D}^{\dim(K)}$$

such that every $X_{p,x}$ has winding greater than 3.

*The triple $(\mathcal{Y}, \mathcal{T}, \{\mathcal{U}(\sigma)\})$ is called the **certificate** of $\mathcal{D}$.*

Remark 4.5. The presence of the constant 3 probably comes as no surprise to the reader, since it is the winding we needed to prove the extension Theorem 3.67. An important observation is that the condition on the winding is destroyed as we perform Little's homotopy, so two t–loose families cannot be (naively) connected by a t–loose homotopy.

4.2.2 Second definition of looseness

Definition 4.6. *Let M be a 4–manifold. Let K be a compact manifold. Let C_{Greco} be a universal constant depending only on* $\dim(K)$.

A family $\mathcal{D} : K \to \mathfrak{Engel}(M)$ is said to be Greco*–**loose** if there exists a smooth family of line fields $\mathcal{Y}(k)$, $k \in K$, such that:*

- *$\mathcal{Y}(k) \subset \mathcal{D}(k)$ is transverse to the kernel $\mathcal{W}(k)$,*
- *each orbit $\gamma : \mathbb{R} \to M$ of $\mathcal{Y}(k)$ has an embedded segment of winding greater than C_{Greco}.*

*We say that $\mathcal{Y}$ is a **certificate** for $\mathcal{D}$.*

The interested reader can jump to Fun Fact 4.21 for an explanation of why "Greco".

Before we get to the next definition, let us discuss Definition 4.6 a bit more. The idea we want to explore is that one can pass from having a contact–type picture to having a convex–type one by taking the line field $\mathcal{W}$ and slightly "tilting" it. This will have some interesting consequences.

4.2.2.1 Tilting

The following definition is analogous to Greco–looseness but using the kernel as the trivialising vector field:

Definition 4.7. *Let M be a 4–manifold. Let K be a compact manifold. Let $C_{\mathcal{W}}$ be a universal constant depending only on* $\dim(K)$*.*

*Let $\mathcal{D} : K \to \mathfrak{Engel}(M)$. Assume that for each Engel structure $\mathcal{D}(k)$, $k \in K$, each orbit $\gamma : \mathbb{R} \to M$ of the kernel $\mathcal{W}(k)$ has an embedded segment where the developing map makes at least $C_{\mathcal{W}}$ projective turns. Then, the family is said to be $\mathcal{W}$–**loose**.*

Indeed, both definitions are closely related:

Proposition 4.8. *Let M be a 4–manifold. Let K be a compact manifold. Any family $\mathcal{D} : K \to \mathfrak{Engel}(M)$ of $\mathcal{W}$–loose structures is* Greco*–loose.*

Proof. Let $\mathcal{Y}_s(k) \subset \mathcal{D}(K)$, $k \in K$, $s \in [0,1]$, be a family of line fields such that $\mathcal{Y}_0(k) = \mathcal{W}(k)$ and $\mathcal{Y}_s(k)$, $s \in (0,1]$, is transverse to the kernel $\mathcal{W}(k)$. By our characterisation of the Engel condition, $\mathcal{D}(k)$ is described by convex curves $X_{p,k,s}$ in any $\mathcal{Y}_s(k)$–flowbox if $s \in (0,1]$; the curves $X_{p,k,0}$ are tangent to a contact equator. We think of $\mathcal{Y}_s$, for s fixed, as a line field in the product $M \times K$.

Let $\gamma_{p,k,s}$ be the orbit of $\mathcal{Y}_s$ through the point (p,k) in $M \times K$; fix a vector field $Y_s \subset \mathcal{Y}_s$ to parametrise it starting from (p,k). Observe that $\gamma_{p,k,0}$ possesses a segment $\gamma_{p,k,0}|_{[a,b]}$ where its developing map makes more than $C_{\mathcal{W}}$ projective turns, by assumption. Then, any convex curve that is C^∞–close to $X_{p,k,0}|_{[a,b]}$ must have winding greater than $C_{\mathcal{W}} - 1$ necessarily. This is the case of $X_{p',k',s}|_{[a,b]}$ for any $s > 0$ sufficiently small, and any (p',k') sufficiently close to (p,k).

Fix $s > 0$ sufficiently small. Write $U_{p,k,a,b} \subset M \times K$ for a small flowbox around $\gamma_{p,k,s}|_{[a,b]}$. We can cover $M \times K$ with finitely many of these regions. By compactness, taking $C_{\mathcal{W}} > C_{\text{Greco}}$ concludes the proof. ☐

That is, in agreement with the idea that large height provided flexibility, we have proven that Engel structures having orbits with long developing map are loose. In particular, recall Theorem 4.3, where we proved that some sort of h–principle applied to prolongations having large turning number. Now we can do better:

Corollary 4.9. *Let K be a compact manifold. Let $\mathcal{D} : K \to \mathfrak{Engel}(M)$ be a family of Lorentz or orientable Cartan prolongations. Then, $\mathcal{D}$ is homotopic to a* Greco*–loose family.*

Proof. By assumption, the ambient manifold M is an $\mathbb{S}^1$–bundle $\pi : N(c) \to N$ over some 3–manifold N. There exists a family of plane fields $\xi(k)$ in N such that $\mathcal{E}(k)$ either projects down to the contact plane $\xi(k)$ (if $\mathcal{D}(k)$ is a Cartan prolongation) or $\mathcal{D}(k)$ is a convex push-off of a formal Cartan prolongation over $\xi(k)$ (if $\mathcal{D}(k)$ is a Lorentz prolongation).

Let us start with the case where $\mathcal{D}(k)$ is a Lorentz prolongation for all k. After an Engel homotopy, we can assume that $d\pi(\mathcal{D}(k))$ is arbitrarily close to $\nu(k)$, with $\nu(k) : N \to TN$, $k \in K$, a vector field transverse to $\xi(k)$. Consider a family of line fields $\mathcal{Y}_s(k) \subset \mathcal{D}(k)$, $s \in [0,1]$, with $\mathcal{Y}_0(k)$ the fibre direction and otherwise transverse to it. Let $Y_s(k)$ be a family of line fields spanning them. Over any 3–disc in N, the $Y_s(k)$ provide a family of return maps $\phi_{k,s}$, with $\phi_{k,0}$ the identity. We claim that, for any fixed $n \in \mathbb{N}$, $\phi_{k,s}^{(n)}$ has no fixed points if $s \neq 0$ is close enough to 0.

Indeed, since we have pushed the prolongations to be very convex, $\phi_{k,s}^{(n)}$ becomes a map that is as close as we want to the flow of $\nu(k)$ for an arbitrarily short time. In particular, since there is an upper bound for the length of the orbits of $\nu(k)$, the map cannot have fixed points.

Now the claim follows. We take n larger than the universal constant C_{Greco}, and we find some s small such that the maps $\phi_{k,s}^{(n)}$ have no fixed points. The rotation of the prolongation is at least 1 projective turn, which implies that any orbit of $Y_s(k)$, as it winds n times around, has winding greater than C_{Greco}. This proves Proposition 4.8 in the Lorentzian case. It also proves the statement for orientable Cartan prolongations, since those are Engel homotopic to Lorentzian ones by a push-off. $\square$

Once we prove that loose families are actually flexible, we will very much improve Theorem 4.3, showing that no counterexample to a full Engel h–principle can arise from studying prolongations.

Remark 4.10. The proof of Corollary 4.9 contains an idea that probably can be exploited more. The core of the argument was showing that, after a deformation, the trivialising line field $\mathcal{Y}$ could be assumed to have arbitrarily large period, and large winding appears. Is the same true for more general Engel structures?

Generically, we know that closed $\mathcal{W}$–orbits are isolated (see Proposition 1.32), but there is no reason to presume that they might have long developing map. Even for open orbits, the developing map can be short (see [47] or Subsubsection 1.5.3.3).

4.2.3 Third definition of looseness

Definition 4.11. *Let M be a 4–manifold. Let K be a compact manifold.*

A family $\mathcal{D} : K \to \mathfrak{Engel}(M)$ is said to be ***asymptotically loose*** [2] *if there exists a smooth family of line fields $\mathcal{Y}(k)$, $k \in K$, and a smooth family of Engel structures $\mathcal{D}_C : K \to \mathfrak{Engel}(M)$, $C \in [0, \infty)$, such that:*

- $\mathcal{D}_0 = \mathcal{D}$,
- $\mathcal{Y}(k) \subset \mathcal{D}_C(k)$ *for all C large enough,*
- $\lim_{C\to\infty} \mathcal{H}_{convex}(\mathcal{D}_C(k)) = \infty$, *and* $\mathcal{H}_{n.i.}(\mathcal{D}_C(k)) = O(1)$.

[2]This definition is very much inspired by the definition of loose map in the context of convex curves, as introduced by Saldanha [70]. Hence the naming.

Remark 4.12. We fix the line fields $\mathcal{Y}(k)$ for C large enough so that the bounds on convex–type energy and the non–integrability energy actually mean something. Indeed, imagine that we had a family of line fields $\mathcal{Y}_C(k)$ that vary with C and get progressively more complicated (say, their C^2 norm diverges). Then, we would have no uniform lower bounds on the size of closed orbits of $\mathcal{Y}_C(k)$, making it impossible to say that the $\mathcal{D}_C$ eventually become loose according to the other definitions.

Remark 4.13. We could have asked for the ratio of convex–type energy over non–integrability energy to go to ∞ as C does. This would have been certainly sufficient. It is unclear whether this has any application.

Remark 4.14. Observe that if $\mathcal{D}$ is asymptotically–loose, any family Engel homotopic to $\mathcal{D}$ is asymptotically–loose itself. This is a slightly gimmicky definition of looseness, but it will allow us to say that the space of loose *structures* has the same connected components as the space of *formal structures*. Essentially, the definition singles out the connected components of the space of Engel structures that contain a loose structure. Equivalently, it singles out connected components where there is a sequence of Engel structures whose convex–type energy goes to infinity (with respect to the non–integrability one).

In any case, this definition will prove useful to relate the other two characterisations. Not only that, but will play a key role in showing that loose families with the same formal data are homotopic. Let us briefly explain why that is by sketching how the proof goes in the π_0 case.

4.2.3.1 A little teaser of the results to come

Suppose we are given two t–loose Engel structures $\mathcal{D}_0, \mathcal{D}_1$, and a formal homotopy $(\mathcal{W}_s, \mathcal{D}_s, \mathcal{E}_s) \in \mathcal{F}\mathfrak{Engel}(M)$, $s \in [0,1]$, between them. We regard the homotopy as a foliated formal Engel structure in $M \times [0,1]$. Each of the structures $\mathcal{D}_0$ and $\mathcal{D}_1$ comes with a a line field $\mathcal{Y}_0, \mathcal{Y}_1$, a triangulation $\mathcal{T}_0, \mathcal{T}_1$, and a cover which certify that they are t–loose.

Let $\mathcal{Y}_s$, $s \in [0,1]$, be the (unique up to homotopy) line field contained in $\mathcal{D}_s$ and connecting $\mathcal{Y}_0$ with $\mathcal{Y}_1$. We want to extend the triangulations $\mathcal{T}_0$ and $\mathcal{T}_1$ to a triangulation $\mathcal{T}$ of $M \times [0,1]$ that is transverse to $\mathcal{Y}_s$. This is (in general) impossible without subdividing the triangulation, but it is clear that the nice properties of $\mathcal{D}_0$ and $\mathcal{D}_1$ with respect to $\mathcal{T}_0$ and $\mathcal{T}_1$ are lost under subdivision. Therefore, the first step of the proof is showing that a t–loose Engel structure with certificate $(\mathcal{Y}, \mathcal{T}, \{\mathcal{U}(\sigma)\})$ is homotopic to a t–loose Engel structure with certificate $(\mathcal{Y}, \mathcal{T}', \{\mathcal{U}'(\sigma)\})$, where $\mathcal{T}'$, $\{\mathcal{U}'(\sigma)\}$ are any triangulation and cover adapted to $\mathcal{Y}$ that we want. Knowing that $\mathcal{D}_0$ and $\mathcal{D}_1$ are asymptotically–loose (almost) automatically yields this claim.

Once that is achieved, the second step of the proof is immediate. Now we can assume that there is indeed a triangulation of $M \times [0,1]$ adapted to $\mathcal{Y}_s$ and extending the triangulations that certificate the looseness of $\mathcal{D}_0$ and $\mathcal{D}_1$. Then, following the proof of (the parametric version of) Theorem 3.1 allows us to conclude that $\mathcal{D}_0$ and $\mathcal{D}_1$ are Engel homotopic.

4.3 All definitions of looseness are equivalent

Having given the three definitions and having outlined the main ideas that come into play, we are ready to state one of the main results:

Theorem 4.15. *Let M be a smooth 4–manifold. Let K be a compact manifold. Let $\mathcal{D} : K \to \mathfrak{Engel}(M)$ be a family of Engel structures. Then, the following statements are equivalent:*

I. the family $\mathcal{D}$ is Engel homotopic to a t–loose family,

II. the family $\mathcal{D}$ is Engel homotopic to a Greco*–loose family,*

III. the family $\mathcal{D}$ is asymptotically–loose.

The proof of Theorem 4.15 breaks down in three parts: showing that (III.) implies (I.) and (II.), showing that (I.) implies (III.), and showing that (II.) implies (III.). We will do them in this order, from simpler to more complicated.

Remark 4.16. It is worth pointing out that we could have defined t–looseness by saying that each $\mathcal{D}(k)$ possesses a triangulation $\mathcal{T}(k)$ and cover $\{\mathcal{U}(k)(\sigma)\}$ where the winding is large, without requiring for the pair $(\mathcal{T}(k), \{\mathcal{U}(k)(\sigma)\})$ to vary smoothly (or even continuously) with k. We claim that, in that case, it would have been sufficient to ask for the winding in each 4–cell to be larger than C_{Greco}. We invite the reader to follow the proof of the implication (II.) $\to$ (III.) and check that exactly the same argument proves this claim.

4.3.1 Large energy implies large winding everywhere

4.3.1.1 Asymptotic looseness implies t–looseness

Proposition 4.17. *Let M be a smooth 4–manifold. Let K be a compact manifold. Let $\mathcal{D} : K \to \mathfrak{Engel}(M)$ be a family of asymptotically loose Engel structures. Let $\mathcal{D}_C$, $C \in [0, \infty)$, and $\mathcal{Y}(k) \subset \mathcal{D}_C(k)$, for C large, be given by hypothesis.*

Let $\mathcal{T}$ and $\{\mathcal{U}(\sigma)\}$ be a triangulation and cover of $M \times K$ that is adapted to $\mathcal{Y}$. Then, $\mathcal{D}_C$ is t–loose with certificate $(\mathcal{Y}, \mathcal{T}, \{\mathcal{U}(\sigma)\})$, for C large.

Proof. Fix a (foliated) flowbox $\mathcal{U}(\sigma)$, with $\sigma \in \mathcal{T}$ a $(4+\dim(K))$–cell. Denote by $X^C_{p,k}$ the family of convex curves describing $\phi(\sigma)_*\mathcal{D}_C(k)$ in $\mathcal{U}(\sigma)$. Here $\phi(\sigma)$ is the trivialising chart of $\mathcal{U}(\sigma)$.

Denote by $\Phi^C_{p,k} : D_{p,k} \subset \mathbb{S}^2 \to \mathbb{D}^2 \subset \mathbb{R}^2$ an affine chart centered at $X^C_{p,k}(0)$. We know that, up to constants that do not depend on the particular Engel structure, the quantities $\mathcal{H}_{\text{convex}}(\phi(\sigma)_*\mathcal{D}_C(k))$ and $\partial_t \mathcal{G}_{\Phi^C_{p,k} \circ X^C_{p,k}}$ are comparable.

Recall that the non–integrability energy $\mathcal{H}_{\text{n.i.}}(\mathcal{D}_C(k))$ is bounded above, independently of C. This implies that there exists a constant $\rho > 0$, independent of C, p, and k, such that $X^C_{p,k}|_{[0,\rho]} \subset \mathbb{D}_{p,k}$. By taking C sufficiently large, we can assume that

$$\partial_t \mathcal{G}_{\Phi^C_{p,k} \circ X^C_{p,k}} > \frac{3\pi}{\rho}$$

holds for all p and k, showing that all the curves in the shell have winding 3. Since the number of top dimensional cells is finite, this implies that $(\mathcal{Y}, \mathcal{T}, \{\mathcal{U}(\sigma)\})$ is a certificate for $\mathcal{D}_C$ for C large. ∎

4.3.1.2 Asymptotic looseness implies Greco–looseness

Proposition 4.18. *Let M be a smooth 4–manifold. Let K be a compact manifold. Let $\mathcal{D} : K \to \mathfrak{Engel}(M)$ be a family of asymptotically loose Engel structures. Let $\mathcal{D}_C$, $C \in [0, \infty)$, and $\mathcal{Y}(k) \subset \mathcal{D}_C(k)$, for C large, be given by hypothesis.*

Let $\{\mathcal{U}_i\}$ be a collection of $\mathcal{Y}$–flowboxes in $M \times K$ covering the leaf space of $\mathcal{Y}$. Then, $\mathcal{D}_C$, for C large, is Greco*–loose with $\mathcal{Y}$ as certificate. Further, for every $\mathcal{U}_i$, the curves describing $\mathcal{D}_C$ all have winding greater than C_{Greco}.*

That is, we are showing that, for each $\mathcal{Y}$–orbit γ, the segment where the winding is large can be taken to be as short as we want (and further, that we can take as many segments where this happens as we desire)

Proof. Proceed as in Proposition 4.17, knowing that the non–integrability energy of the curves in each of the $\mathcal{U}_i$ is uniformly bounded, while the convex–type one is increasing with C. □

4.3.2 t–Looseness implies asymptotic looseness

The statement we want to show is the following:

Proposition 4.19. *Let M be a smooth 4–manifold. Let K be a compact manifold. Any t–loose family of Engel structures $\mathcal{D} : K \to \mathfrak{Engel}(M)$ is asymptotically loose.*

Let $(\mathcal{Y}, \mathcal{T}, \{\mathcal{U}(\sigma)\})$ be the certificate of $\mathcal{D}$. Define:

$$\mathcal{U}_j = \bigcup_{\tau \in \mathcal{T}^{(j)}} \mathcal{U}(\tau) \subset M \times K.$$

The proof can be broken down in several steps, mimicking the strategy used for the existence theorem. We want to produce a family $\mathcal{D}_C : K \to \mathfrak{Engel}(M)$, $\mathcal{D}_0 = \mathcal{D}$, satisfying the quantitative conditions of asymptotic looseness.

We will modify $\mathcal{D}_0$ in each $\mathcal{U}(\sigma)$, iterating over the simplices $\sigma \in \mathcal{T}$ inductively in the dimension. We will first explain how this is done if σ is a 0–simplex, then we shall do it up to the $(3+\dim(K))$–skeleton, and then we shall finish the argument by modifying over the top dimensional cells. At every step, the convex–type energy of $\mathcal{D}_C$ over the flowboxes $\mathcal{U}(\sigma)$ that have already been processed will go to infinity with C. Outside of those, the structure will not even be Engel until the completion of the proof. Indeed, the main idea is that we make things very Engel in the neighbourhood $\mathcal{U}_{3+\dim(K)}$ of the $(3+\dim(K))$–skeleton by destroying the Engel condition in the $(4+\dim(K))$–cells and then we use Little's theorem to fix this. This must be done relative to $\mathcal{D}_0$.

Recall that if $\sigma \in \mathcal{T}$ is a simplex, $\mathcal{U}(\sigma)$ is the corresponding flowbox in the cover. $\mathcal{U}(\sigma)$ has an associated map $\phi(\sigma) : \mathcal{O}p(\mathcal{U}(\sigma)) \to \mathbb{D}^{3+\dim(K)}_{1+\varepsilon} \times [-\varepsilon, 1+\varepsilon]$ that takes $\mathcal{U}(\sigma)$ to $\mathbb{D}^{3+\dim(K)} \times [0,1]$. In the target we take coordinates (p,t). We do not distinguish the coordinates on the manifold and on the parameter space because it actually plays no role in the discussion.

4.3.2.1 Adapting $\mathcal{D}$ to the cover

Before we start the inductive process, we have to perform some additional setup. Let $\sigma \in \mathcal{T}^{(4+\dim(K))}$ be a top dimensional cell. Then, we know that the curves X^0_p describing $\mathcal{D}_0$ in $\mathcal{U}(\sigma)$ have winding 3. However, this does not imply that this winding is concentrated in a band contained in $\mathcal{U}(\sigma) \cap \mathcal{U}_{3+\dim(K)}$. We need to perform a first homotopy that achieves precisely this, because the Engel condition will be destroyed outside of $\mathcal{U}_{3+\dim(K)}$.

However, this is almost immediate. In particular, by construction, the lower horizontal boundary $\phi(\sigma)^{-1}(\mathbb{D}^{3+\dim(K)} \times \{0\})$ and the vertical boundary of $\mathcal{U}(\sigma)$ are covered by $\mathcal{U}_{3+\dim(K)}$. Find a family of reparametrisations $H_p : [0,1] \to [0,1]$ such that:

- $H_p(t) = t$ if $(p,t) \in \mathcal{O}p(\partial(\mathbb{D}^{3+\dim(K)} \times [0,1]))$,
- the curves $X_p^0 \circ H_p|_{[0,\rho]}$ have winding 3 if $|p| < 1 - \delta$.

Any family of curves $X_p^0 \circ H_p$ yields an Engel structure, all of which are Engel homotopic. If $\rho > 0$ and $\delta > 0$ are chosen sufficiently small, the winding takes place in $\mathcal{U}(\sigma) \cap \mathcal{U}_{3+\dim(K)}$. This yields an Engel homotopy to an Engel structure satisfying the desired property; we still write $\mathcal{D}_0$ for such an structure.

4.3.2.2 Twisting at the ends

Given σ some lower dimensional simplex, we want to take the curves X_p^0 describing $\mathcal{D}_0$ in $\mathcal{U}(\sigma)$ and start turning their ends to add convexity. This will add new wiggles that we will then distribute along the curves. As we keep doing this, the convex–type energy will go to infinity in $\mathcal{U}(\sigma)$, while the non–integrability energy will stay bounded if we make the wiggles have progressively smaller radius.

Let us write some explicit formulae for the twisting at the endpoints. Given a family of curves $f : K \to \mathcal{L}([a,b],\mathbb{S}^2)$, we are going to produce a homotopy $f_s : K \to \mathcal{L}([a,b],\mathbb{S}^2)$, $s \in [0,1]$, such that $f_0 = f$ and $f_1 = f^{[1\#a;1\#b]}$. What we mean by this notation is that a small convex wiggle has been added at the endpoints a and b, but f and f_1 otherwise agree.

Let $\alpha \in \mathcal{L}_1$, i.e. $\alpha : \mathbb{S}^1 \to \mathbb{S}^2$, is a convex curve with $\Gamma_\alpha(1) = \mathrm{Id}$ describing one turn in an affine chart. Define:

$$\alpha_s, \beta_s : [0,1] \to \mathbb{S}^2, \quad s \in [0,1],$$
$$\alpha_s(t) = \alpha(e^{2\pi sti}),$$
$$\beta_s(t) = \alpha(e^{2\pi s(t-1)i}).$$

Then, $f_s(k)$ is defined as a smoothing of the concatenation of $\Gamma_f(a)(\beta_s)$, f, and $\Gamma_f(b)(\alpha_s)$. The matrices $\Gamma_f(0)$ and $\Gamma_f(1)$ simply make the Gauss maps at the concatenation points agree. Observe that f_s verifies what we claimed.

We should be slightly careful with the explicit parametrisation of f_s, because it is geometrically relevant for the Engel structures. However, the space of possible parametrisations of the interval is contractible, so we have complete freedom.

Let us introduce some additional notation that packages the parameters we are interested in. Given any family $f : K \to \mathcal{L}([a,b],\mathbb{S}^2)$, write $f^{\{\tau,s\}}$ for the family f_s produced as we just explained, and additionally satisfying:

- the intervals parametrising $\Gamma_f(a)(\beta_s)$ and $\Gamma_f(b)(\alpha_s)$ have width $2\pi\tau s$,
- the wiggles added have radius τ,
- $f^{\{\tau,s\}}|_{[a+s\tau,b-s\tau]}$ is a linear reparametrisation of $f|_{[a,b]}$.

4.3.2.3 Increasing the convex–type energy in the 0–skeleton

Let σ be some 0–simplex. We shall define a family of 2–distributions $\mathcal{D}_C$, $C \in [0,\infty)$, in $\mathcal{O}p(\mathcal{U}(\sigma))$, given by curves $X_p^C : [-\varepsilon, 1+\varepsilon] \to \mathbb{S}^2$. The structure $\mathcal{D}_C$ will be Engel in $\mathcal{U}(\sigma)$, but the Engel condition will fail in the region $\mathbb{D}_{1+\varepsilon}^{3+\dim(K)} \times ([-\varepsilon, 0] \cup [1, 1+\varepsilon])$.

Let $D : [0,\infty) \to [0,\infty)$ be an increasing function of C satisfying $D(C) = O(C^2)$ with $D(0)$ sufficiently large; we will fix it along the proof. Let N be an integer. We will first construct curves $Z_p^C : [-\varepsilon, 1+\varepsilon] \to \mathbb{S}^2$, $C \in [0,\infty)$.

Set $Z_p^0 = X_p^0$ and define:

$$Z_p^C = (Z_p^{2N})^{\{1/D(N), C-2N\}}, \qquad C \in [2N, 2N+1].$$

That is, in the interval $[2N, 2N+1]$ we add a new wiggle at each end $t = -\varepsilon, 1+\varepsilon$. We will then use the interval $[2N+1, 2N+2]$ to evenly distribute them. Indeed, define functions

$$t_i^{2N} : [2N+1, 2N+2] \to [-\varepsilon, 1+\varepsilon], \qquad i \in \{0, \ldots, 2N+1\} \quad \text{satisfying:}$$

- $t_0^{2N}(2N+1) = -\varepsilon$, and $t_{2N+1}^{2N}(2N+1) = 1+\varepsilon$,
- $t_i^{2N}(2N+1) = t_{i-1}^{2N-2}(2N)$, $i \in \{1, \ldots, 2N\}$, are evenly spaced in $(-\varepsilon, 1+\varepsilon)$,
- $t_i^{2N} < t_j^{2N}$ if and only if $i < j$,
- t_i^{2N} goes linearly from $t_i^{2N}(2N+1)$ to $t_i^{2N}(2N+2)$.

Then, we can set, for $C \in [2N+1, 2N+2]$:

$$Z_p^C = (Z_p^0)^{[1\#t_0^{2N}(C), \ldots, 1\#t_{2N+1}^{2N}(C)]}.$$

Finally, smooth the family Z_p^C in the parameter C (at the moment, it is only continuous).

These curves behave exactly how we want in $\mathbb{D}^{3+\dim(K)} \times [0,1]$, but we need to interpolate back to X_p^0 in the boundary of the bigger flowbox $\mathcal{O}p(\mathcal{U}(\sigma))$. There are two things to be done. First, we replace Z_p^C in the region $t \in [-\varepsilon, 0] \cup [1, 1+\varepsilon]$ by a family of immersions that glue with X_p^0 at $t = -\varepsilon, 1+\varepsilon$. This can be achieved by finding a formal interpolation and using the Smale–Hirsch theorem. Geometrically, it amounts to undoing the winding we add by winding in the opposite direction. Secondly, find a bump function

$$\chi : \mathbb{D}_{1+\varepsilon}^{3+\dim(K)} \to [0,1]$$

that is identically 1 in $\mathbb{D}^{3+\dim(K)}$ and identically 0 in the boundary of $\mathbb{D}_{1+\varepsilon}^{3+\dim(K)}$. Set:

$$X_p^C = Z_p^{C\cdot\chi(p)}.$$

Since the curves X_p^C glue nicely with X_p^0, the corresponding 2–distributions $\mathcal{D}_C$ can be assumed to be globally defined, but they are not Engel anymore close to the horizontal boundary of $\mathcal{U}(\sigma)$. However, they are still Engel in every $\mathcal{U}(\tau)$ with τ not top dimensional, due to the properties of the triangulation $\mathcal{T}$.

Finally, note that in $\mathcal{U}(\sigma)$ the properties we claimed hold. Indeed, as we keep adding wiggles the curves get increasingly convex everywhere (in–between the wiggles we do some adjusting to go from one to the next). Additionally, since they get progressively smaller quadratically, the length they add is bounded. See Figure 4.2 for a pictorial depiction of this process.

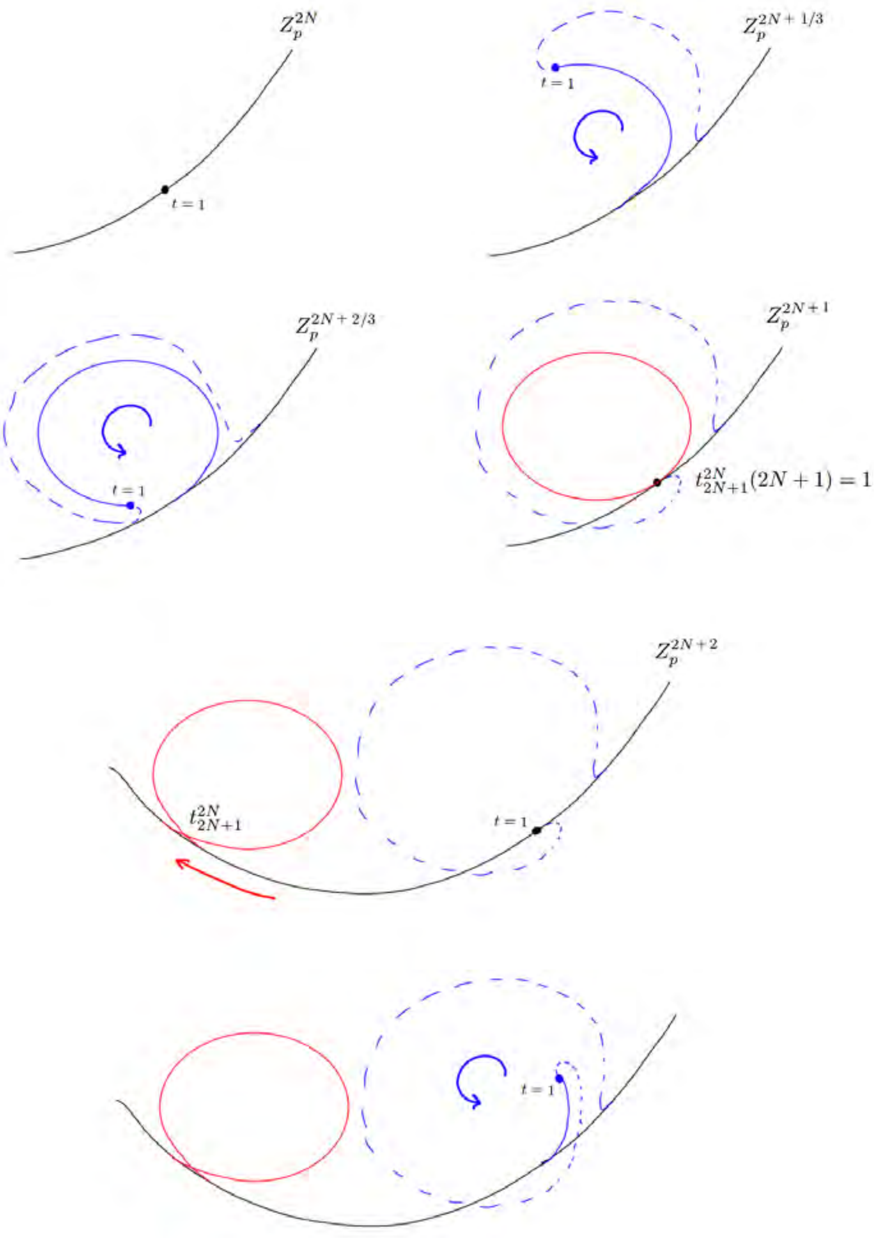

Figure 4.2: In this sequence of figures we display how the twisting at the ends is done. We have included how the "untwisting" is performed, so it is slightly different from the proof given. Instead of twisting at $t = 1 + \varepsilon$ (and then using Hirsch–Smale), the twist is added at $t = 1$ and then we undo it by turning in the opposite direction (in discontinuous blue). When a whole loop has been introduced (in red), we slide it down. It is clear that if the loop is sufficiently small, the curve is becoming more convex.

4.3.2.4 Increasing the convex–type energy in the $(3 + \dim(K))$–skeleton

Suppose σ is now a j–dimensional cell, $j = 1, 2, \dots, 3 + \dim(K)$. We will follow exactly the same method as before, but now it has to be done relative to the areas already arranged in previous steps. By induction hypothesis, we have defined $\mathcal{D}_C$, $C \in [0, \infty]$, and it satisfies the properties we want (regarding convex and non–integrability energy) in $\mathcal{O}p(\mathcal{U}_{j-1})$. Denote by $\tilde{X}_p^C : [-\varepsilon, 1+\varepsilon] \to \mathbb{S}^2$ the corresponding curves.

From the properties of the cover, it holds that:

$$\phi(\sigma)(\mathcal{U}_{j-1} \cup \mathcal{U}(\sigma)) = A \times [-\varepsilon, 1+\varepsilon]$$

$$\phi(\sigma)(\mathcal{O}p(\mathcal{U}_{j-1}) \cup \mathcal{U}(\sigma)) = \mathcal{O}p(A) \times [-\varepsilon, 1+\varepsilon]$$

with A some closed set. Our construction will be relative to A.

Set $Z^0_{p,c} = \tilde{X}^c_p$, $c \in [0, \infty)$, and define:

$$Z^C_{p,c} = (Z^{2N}_{p,c})^{\{1/D(N), C-2N\}}, \qquad C \in [2N, 2N+1],$$

$$Z^C_{p,c} = (Z^0_{p,c})^{[1\# t_0^{2N}(C), \dots, 1\# t^{2N}_{2N+1}(C)]}, \qquad C \in [2N+1, 2N+2],$$

where the functions t_i^{2N} are exactly as before. Effectively, we are doing the same, but we use the functions $\tilde{X}^c_p$ as starting point; this will allow us to glue everything together.

Let $\mathcal{O}p(A)$ be the same neighbourhood we mentioned before in which the convex–type energy was already growing with C. Let $\mathcal{O}p'(A) \subset \mathcal{O}p(A)$ be a much smaller neighbourhood. The, set:

$$\chi : \mathbb{D}^{3+\dim(K)}_{1+\varepsilon} \to [0,1]$$

$$\chi|_{\mathcal{O}p'(A) \cup \mathcal{O}p(\mathbb{S}^{2+\dim(K)}_{1+\varepsilon})} = 0, \qquad \chi|_{\mathbb{D}^{3+\dim(K)} \setminus \mathcal{O}p(A)} = 1,$$

$$X^C_p = Z^{C\cdot\chi(p)}_{p, C\cdot(1-\chi(p))}.$$

These curves do not glue nicely at their ends with the original family $\tilde{X}^C_p$, but they do for those $p \in \mathcal{O}p'(A)$ and those $p \notin \mathcal{O}p(A)$. We can then apply the Smale–Hirsch theorem, relative in the parameter, to modify the curves X^C_p in $t \in [-\varepsilon, 0] \cup [1, 1+\varepsilon]$ so that they do agree with $\tilde{X}^C_p$ in $t = -\varepsilon, 1+\varepsilon$.

We now redefine the family $\mathcal{D}_C$, $C \in [0, \infty)$, in $\mathcal{O}p(\mathcal{U}(\sigma))$ so that it is given by X^C_p. By construction, $\mathcal{D}_0$ is still our original Engel structure. Similarly, $\mathcal{D}_C$ is Engel in $\mathcal{U}_{3+\dim(K)}$. It is obvious that its non–integrability energy remains bounded. We claim that its convex–energy in $\mathcal{U}_j$ tends to ∞. It is enough to check that this is the case in each $\mathcal{U}(\sigma)$ with σ a j–simplex.

In the regions $A \times [-\varepsilon, 1+\varepsilon]$ and $(\mathbb{D}^{3+\dim(K)} \setminus \mathcal{O}p(A)) \times [0,1]$ it is clear. In the interpolation region $(\mathcal{O}p(A) \setminus A) \times [0,1]$, the observation to make is that the curve $Z^{C\cdot\chi(p)}_{p, C\cdot(1-\chi(p))}$ has approximately $2C_0 C \cdot (1 - \chi(p)) + 2C \cdot \chi(p)$ wiggles, where C_0 is a constant relating the "height" of $\mathcal{U}(\sigma)$ with the height of the surrounding $\mathcal{U}(\tau)$. All such constants are controlled from below by a universal positive constant depending only on the cover and hence the claim follows.

4.3.2.5 Extension to the $(4 + \dim(K))$–balls as an Engel structure

We have constructed a family $\mathcal{D}_C$, $C \in [0, \infty)$, of 2–distributions that are not integrable everywhere and Engel in $\mathcal{U}_{3+\dim(K)}$, where they additionally have convex energy going to ∞. Let us henceforth assume that we are in a fixed $(4 + \dim(K))$–ball $\mathcal{U}(\sigma)$ where the structures $\mathcal{D}_C$ are all convex–type shells given by curves X_p^C. The setup we did guarantees that the curve X_p^0 has winding 3 in the segment $[0, h_p]$, where $h_p \in (0, 1)$ is a smooth family of constants, and

$$\{(p, t) \mid t \in [0, h_p]\} \subset \phi(\sigma)(\mathcal{U}_{3+\dim(K)}).$$

This winding is measured with respect to a smooth family of affine charts Φ_p.

We claim that the curves X_p^C also have winding as least 3 in $[0, h_p]$, with respect to Φ_p. At this point we fix the constant $D(0)$ to be large enough to that all the wiggles we have been adding are more convex than the curves X_p^0. Fixing p and varying C, we see that new wiggles progressively appear at the ends of $X_p^C|_{[0,h_p]}$. Once the domain of the wiggle is fully contained in $[0, h_p]$, it is clear that the winding has increased by 1 and the curve is otherwise the same as before the wiggle appeared. While the wiggle is appearing, since it is more convex than the original curve, the winding is increasing as well. This proves the claim.

This shows that the extension problem is solvable simultaneously for all C and relative to $C = 0$. Let us be a bit more precise. Those $\mathcal{D}_C$ with C sufficiently close to 0 can be assumed to be Engel everywhere. Then, homotope all the shells so that $\mathcal{D}_C$, $C \in [\delta, \infty)$, is a 1–convex–shell and those $\mathcal{D}_C$, $C \in [0, 2\delta]$, are solid. This is done as in Proposition 3.50. Then, an application of Proposition 3.29 solves the extension problem parametrically for $C \in [\delta, \infty)$, relative to δ. Additionally, the proposition states that the non–integrability–energy remains bounded and the convex–energy goes to infinity away from the vertical boundary of the shell.

We leave as an exercise to the reader to check that, since the vertical boundary is contained in $\mathcal{U}_{3+\dim(K)}$, one can evenly distribute the convex–energy there while the extension problem is being solved. □

4.3.3 Greco–looseness implies asymptotic looseness

The proof relies on the following result:

Proposition 4.20. *Let M be a smooth compact 4–manifold. Let K be a compact manifold. Let $\mathcal{D} : K \to \mathfrak{Engel}(M)$ be a* Greco*–loose family.*

Then, there exists a cover $\{\mathcal{U}_i\}_{i=1,\dots,N}$ of $M \times K$ by convex–type flowboxes with trivialising charts $\phi_i : \mathcal{U}_i \to \mathbb{D}^{3+\dim(K)} \times [0, 1]$ such that:

- *there is a smaller flowbox $V_i = \phi_i^{-1}(\mathbb{D}^{3+\dim(K)} \times [a_i, b_i]) \subset \mathcal{U}_i$, such that curves X_p^i describing $\mathcal{D}$ in the flowbox V_i have winding 3,*
- *the flowbox V_i does not intersect the vertical boundaries $\phi_i^{-1}(\mathbb{S}^{2+\dim(K)} \times [a_j, b_j])$, $j < i$, of the previous flowboxes V_j.*

During its proof, we will determine the value of the constant C_{Greco}. Assuming the proposition to be true, showing that the family $\mathcal{D}$ is then asymptotically loose is similar to what we explained in the previous subsection, but technically simpler:

Greco*–looseness implies asymptotic looseness.* Consider the flowbox $\mathcal{U}_0$; let X_p be the curves describing $\mathcal{D}$. We want to apply Proposition 3.29 to increase the convex energy. Following its proof yields a family of curves X_p^C, $C \in [0, \infty)$, such that, for each integer n:

- $X_p^0 = X_p$,
- $X_p^n = X_p^{[1\#t_1;\dots,1\#t_{2n}]}$ in the region $\{|p| < 1-2\delta\}$, where

$$t_j, a, b : \mathbb{D}^{3+\dim(K)} \to [0,1]$$

$$t_j(p) = (a(p) - b(p))\frac{j}{n} + b(p)$$

$$a(p) = a_0, \quad b(p) = b_0, \qquad \text{if } |p| > 1-3\delta$$

$$a(p) = 0, \quad b(p) = 1, \qquad \text{if } |p| < 1-4\delta,$$

- $X_p^C = X_p$ in the region $\{t \notin [a(p), b(p)]\}$,
- X_p^n is defined in the band $\{1-\delta > |p| > 1-2\delta\}$ by performing Little's homotopy sequentially n times,
- in the interval $C \in (n, n+1)$, Little's homotopy between X_p^n and X_p^{n+1} is performed, the spacing of the wiggles is modified, and the radius of the wiggles is decreased so that it behaves as $O(1/C^2)$.

These curves define a family $\mathcal{D}_C^0$ whose convex–type energy goes to infinity in $\{|p|, t < 1-4\delta\}$, whose non–integrability energy remains bounded, and satisfying $\mathcal{D}_0^0 = \mathcal{D}$. See Figure 4.3 for a pictorial depiction of all the regions involved.

At every step i, a new family $\mathcal{D}_C^i$ is produced by deforming, over the flowbox $\mathcal{U}_i$, the family $\mathcal{D}_C^{i-1}$ constructed in the previous step. This is done by taking the curves $\tilde{X}_p^C$ describing $\mathcal{D}_C^{i-1}$ and applying to them the recipe we just explained. Let us give a bit more detail. The main geometrical ingredient is that the flowbox V_i intersects all the previous flowboxes V_j, $j < i$, away from their vertical boundary; in particular, V_i intersects $\mathcal{U}_j$ in a region where the $\mathcal{D}_C^{i-1}$ either agrees with $\mathcal{D}$ or has been obtained from it by adding wiggles.

For $C > C_{i-1} > 0$, we can assume that $\mathcal{D}_C^{i-1}$ has winding greater than 3 in the flowbox V_i. Note that for small C this might not be true: while Little's homotopy is being performed in a flowbox overlapping with V_i, some of the curves in V_i are moving around $\mathbb{S}^2$ and the winding goes down; once the homotopy has been completed, the curves are exactly as in $\mathcal{D}_0$, but with an additional loop. As C increases, Little's homotopy is performed over and over, and eventually enough loops appear to guarantee having winding 3.

Then, we construct a family of curves $X_{p,c}^C : [0,1] \to \mathbb{S}^2$, for $c \in [C_{i-1}, \infty)$ and $C \in [0, \infty)$. For any integer n, we set:

- $X_{p,c}^0 = \tilde{X}_p^c$,
- $X_{p,c}^n = (\tilde{X}_p^c)^{[1\#t_1;\dots,1\#t_{2n}]}$ in the region $\{|p| < 1-2\delta\}$, where

$$t_j, a, b : \mathbb{D}^{3+\dim(K)} \to [0,1]$$

$$t_j = (a(p) - b(p))\frac{j}{n} + b(p),$$

$$a(p) = a_i, \quad b(p) = b_i, \qquad \text{if } |p| > 1-3\delta,$$

$$a(p) = 0, \quad b(p) = 1, \qquad \text{if } |p| < 1-4\delta,$$

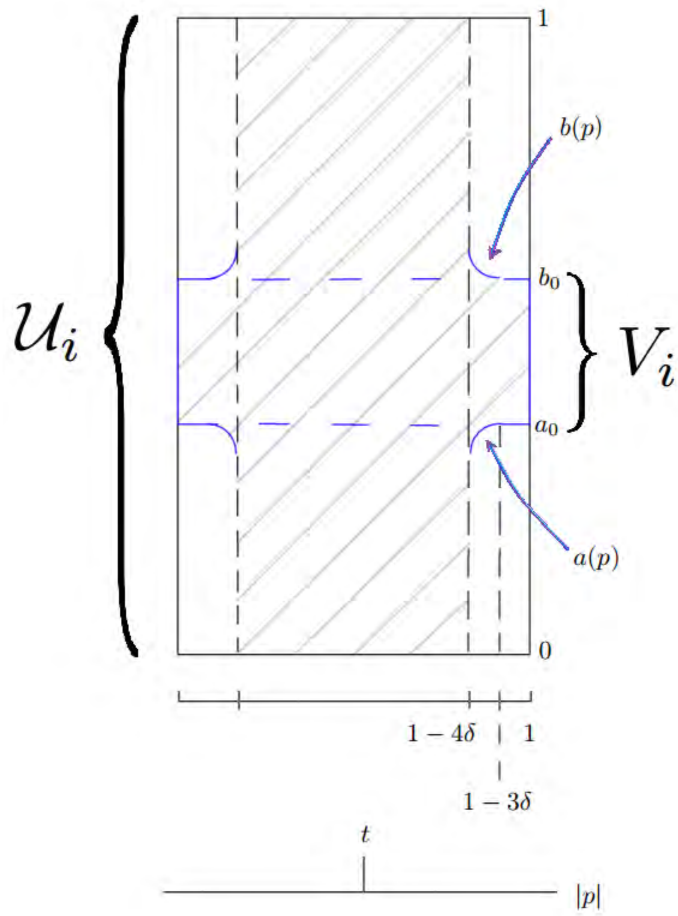

Figure 4.3: A pair of flowboxes $(\mathcal{U}_i, V_i)$ as in the statement of Proposition 4.20. The labels mark the different areas involved in the proof of asymptotic looseness from Greco–looseness.

- $X^C_{p,c} = \tilde{X}^c_p$ in the region $\{t \notin [a(p), b(p)]\}$,
- $X^n_{p,c}$ is defined in the band $\{1-\delta > |p| > 1-2\delta\}$ by performing Little's homotopy sequentially n times,
- in the interval $C \in (n, n+1)$, Little's homotopy between $X^n_{p,c}$ and $X^{n+1}_{p,c}$ is performed, the spacing of the wiggles is modified, and the radius of the wiggles is decreased so that it behaves as $O(1/c^2)$.

Construct a function $\chi : [0,\infty) \to [0,\infty)$ such that $\chi|_{[0,2C_{i-1}]} = 0$ and $\chi(C) = C$ if $C > 3C_{i-1}$. Define the family $\mathcal{D}^i_C$ by:

- $\mathcal{D}^i_C = \mathcal{D}^{i-1}_C$ outside of $\mathcal{U}_i$,
- in $\mathcal{U}_i$, $\mathcal{D}^i_C$ is described by the curves $\tilde{X}^C_p$ if $C \leq C_{i-1}$
- in $\mathcal{U}_i$, $\mathcal{D}^i_C$ is described by the curves $X^{\chi(C)}_{p,C}$ if $C > C_{i-1}$.

Iterating this process over all the $\mathcal{U}_i$, we deduce that $\mathcal{D}_C = \mathcal{D}_C^N$ has the desired properties. $\square$

The rest of the subsection is dedicated to the proof of Proposition 4.20. We break it down in several parts.

4.3.3.1 Constructing the first covering

Let $q \in M\times K$, and let γ be the $\mathcal{Y}$–orbit passing through q; find a parametrisation $\gamma : \mathbb{R} \to M\times K$ for the orbit. By hypothesis, there exists an interval $I_q \subset \mathbb{R}$ and a thickening $\mathcal{V}_q$ of $\gamma|_{I_q}$ such that: $\phi_q : \mathcal{V}_q \to \mathbb{D}^{3+\dim(K)}\times[0,1]$ is an embedded convex–type $\mathcal{Y}$–flowbox whose curves $X_p^q : [0,1] \to \mathbb{S}^2$ have winding greater than the universal constant C_{Greco}. Fix $\delta > 0$ arbitrarily small. Then, moreover, it can be assumed that if X_0^q has winding exactly a over some interval $I \subset [0,1]$, all the curves $X_p^q|_I$ have winding greater than $a-\delta$. Both statements follow by taking the thickening sufficiently small.

The constant C_{Greco} will be determined as we go along in the proof.

Since $M \times K$ is compact, we can find a finite cover by flowboxes $\{\mathcal{V}_{q_i}\}_{i=0,\dots,L}$ like the ones we just described; we denote their trivialising charts by $\phi_{q_i} : \mathcal{V}_{q_i} \to \mathbb{D}^{3+\dim(K)} \times [0,1]$.

Fun fact 4.21. *We have given the artistic name of* ***Greco chart*** *to the flowboxes $\mathcal{V}_{q_i}$. The picture to have in mind is the following: they will usually be incredibly tall, going around the manifold describing a very long and thin tube. Doménikos Theotokópoulos, more known as* El Greco*, was a Greek Renaissance painter established in Toledo, Spain. His style is very distinctive: the characters in his works often display elongated faces and hands, as well as a very personal grace and posture. Our flowboxes, even if not particularly graceful, are indeed elongated.*

4.3.3.2 First flowbox of the refinement

We will obtain the covering claimed in the statement by *refining* $\{\mathcal{V}_{q_j}\}_{j=0,\dots,L}$. Let us start with $\mathcal{V}_{q_0}$. Now we impose for C_{Greco} to be greater than 3. Since all the curves $X_p^{q_0}$ have winding $C > 3$, we can find some interval $[a_0, b_0] \subset [0,1]$ such that the curves $X_p^{q_0}|_{[a_0,b_0]}$ have winding in $(3, 3+\delta)$, with $\delta > 0$ some very small constant. Then, we set

$$\mathcal{U}_0 = \mathcal{V}_{q_0}$$

$$V_0 = \phi_{q_0}^{-1}(\mathbb{D}^{3+\dim(K)} \times [a_0, b_0]).$$

4.3.3.3 General position for flowboxes

Suppose now that we have constructed flowboxes $\{\mathcal{U}_i\}_{i=0,\dots,m}$ covering all the $\mathcal{V}_{q_j}$, $j < j_0$. Denote their trivialising charts by $\phi_i : \mathcal{U}_i \to \mathbb{D}^{3+\dim(K)} \times [0,1]$, and write X_p^i for the convex curves describing $\mathcal{D}$ in each of them. The induction hypothesis is that, just like in the statement:

I. the curves X_p^i all have winding in $(3, 3+\delta)$ in $[a_i, b_i]$,

II. each $V_i = \phi_i^{-1}(\{t \in [a_i, b_i]\})$ does not meet the vertical boundaries of the previous ones.

Consider the flowbox $\mathcal{V}_{q_{j_0}}$. We will construct new flowboxes $\{\mathcal{U}_i\}_{i=m+1,\dots,m'}$ so that the induction hypothesis holds for j_0 too. As a first step, we claim that the following additional properties hold:

III. Let $\pi : \mathbb{D}^{3+\dim(K)} \times [0,1] \to \mathbb{D}^{3+\dim(K)}$ be the obvious projection. The (probably incomplete) cylinders

$$D_{j_0,i} = \{\phi_{j_0} \circ \phi_i^{-1}(\mathbb{S}^{2+\dim(K)} \times [a_i, b_i])\}, \quad i = 0, \dots, m,$$

and the (probably incomplete) piecewise smooth (2+dim(K))–spheres to which they project

$$S_{j_0,i} = \pi \circ \phi_{q_{j_0}} \circ \phi_i^{-1}(\mathbb{S}^{2+\dim(K)} \times [a_i, b_i]), \quad i = 0, \dots, m$$

are in general position. In particular, any collection of $3 + \dim(K)$ spheres intersects in a finite number of points, and any collection of $4 + \dim(K)$ spheres has empty intersection.

IV. For each cylinder $D_{j_0,i}$, there exists an interval $I_i \subset [0,1]$ such that

$$S_{j_0,i} \times I_i \subset D_{j_0,i} \subset S_{j_0,i} \times \mathcal{O}p(I_i),$$

and, for all $p \in S_{j_0,i}$, the curves $X_p^{j_0}|_{I_i}$ have less that 4 winding.

The first statement follows by applying standard transversality, since the number of flowboxes is finite. The nice consequence of this is that, given any point p, the segment $\{p\} \times [0,1]$ only intersects $3 + \dim(K)$ of the cylinders $D_{j_0,i}$, so only a controlled amount of convex–type energy is being lost for the curve $X_p^{j_0}$. We will detail this later.

The second statement follows by construction. Just like in the 0–dimensional case, we will choose the flowboxes along the proof to have winding lying in the interval $(3, 3+\delta)$ and having horizontal boundaries that are almost flat (so that the interval I_i does not depend on the particular point $p \in S_{j_0,i}$). The main point of subtlety here is that winding does depend on how the base $\mathbb{D}^{3+\dim(K)}$ is trivialised, so it might be unclear whether statements about winding transfer between different Greco charts. However, 3 projective turns seen in an affine chart will be always produce less that 4 in any other chart.

4.3.3.4 Triangulating $\mathbb{D}^{3+\dim(K)}$ with respect to the projected spheres

Consider the manifold $\mathbb{D}^{3+\dim(K)}$ and the (possibly disconnected and with boundary) $(2 + \dim(K))$–manifolds $S_{j_0,i}$, $i = 0, \dots, m$. Then, there exists a triangulation $\mathcal{T}$ of $\mathbb{D}^{3+\dim(K)}$ such that:

- $\mathcal{T}$ extends a triangulation $\mathcal{T}_\partial$ of the boundary $\partial(\mathbb{D}^{3+\dim(K)})$,
- the $(3 + \dim(K) - n)$–discs given as intersections of n elements of the collection

$$\{S_{j_0,i}\} \cup \{\mathbb{S}^{2+\dim(K)}\}$$

are covered by simplices of $\mathcal{T}$ of dimension at most $(3 + \dim(K) - n)$.

This is straightforward by starting with the union of the spheres and subdividing sufficiently.

To each simplex $\sigma \in \mathcal{T}$ we assign a disc $U(\sigma)$ so that the collection of discs $\{U(\sigma)\}_{\sigma\in\mathcal{T}}$ is a covering of $\mathbb{D}^{3+\dim(K)}$. Similarly to Theorem 3.9, it can be assumed that $U(\sigma)$ and $U(\tau)$

only intersect if one of the simplices contains the other. Further, it can be assumed that $U(\sigma)$ intersects at most $3+\dim(K)$ of the spheres in the collection $\{S_{j_0,i}\}$. Let D be the sum of $3+\dim(K)$ and the number of subsimplices a simplex of dimension $3+\dim(K)$ possesses; it is a universal constant that depends only on $\dim(K)$.

Refer to Figure 4.4 for a pictorial depiction of the manifolds $S_{j_0,i}$, the corresponding triangulation $\mathcal{T}$, and the subsequent covering $\{U(\sigma)\}_{\sigma\in\mathcal{T}}$.

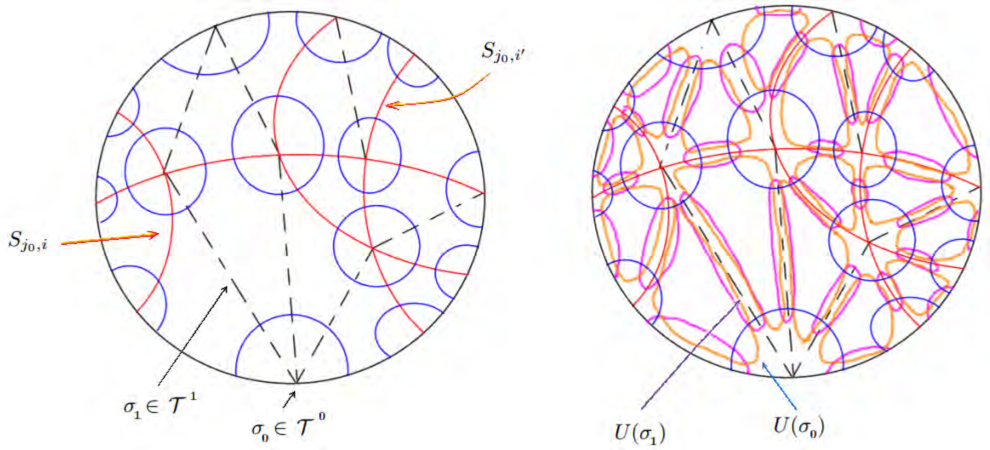

Figure 4.4: Two depictions of $\mathbb{D}^{3+\dim(K)}$: for simplicity we show it as a $\mathbb{D}^2$. On the left hand side, we show in red the spheres $S_{j_0,i}$. In (discontinuous) black, the rest of the triangulation $\mathcal{T}$. The covering $\{U(\sigma)\}_{\sigma\in\mathcal{T}}$ is shown in 3–colours. In blue we depict the (boundary of the) neighbourhoods of the zero simplices. In pink, (in the right hand side of the picture) the neighbourhoods of the edges. In orange, the neighbourhoods of the faces.

4.3.3.5 Refining the Greco chart at step j_0

Let σ be a l–simplex. Assume that, for each simplex τ of dimension lower than l, we have defined flowboxes

$$\mathcal{U}(\tau) = \phi_{j_0}^{-1}(U(\tau)\times[0,1]),$$
$$V(\tau) = \phi_{j_0}^{-1}(U(\tau)\times[a_\tau,b_\tau]),$$

such that, along with the previous $\mathcal{U}_i$, V_i, $i=1,\dots,m$, they satisfy the induction hypothesis. Let us construct $V(\sigma)$.

For each sphere S in the collection $\{S_{j_0,i}\}$ and intersecting $U(\sigma)$, there is an interval $I_S\subset[0,1]$ that is forbidden (meaning that we cannot take $[a_\sigma,b_\sigma]$ to overlap with it if we want to satisfy the induction hypothesis). Similarly, for each subsimplex $\tau\subsetneq\sigma$, the interval $[a_\tau,b_\tau]$ is forbidden as well. There are at most D forbidden intervals, by our reasoning before. In the interval $[0,1]$, the curves $X_p^{j_0}$ have winding C_{Greco}, and in each forbidden interval, at most 4 winding is lost. We need to find a gap in-between the intervals of winding 3. Since there are at most $D+1$ gaps, the pigeonhole principle states that setting $C_{\text{Greco}}>4D+3(D+1)$ guarantees that there is a suitable gap in which to place $[a_\sigma,b_\sigma]$. Define:

$$\mathcal{U}(\sigma) = \phi_{j_0}^{-1}(U(\sigma)\times[0,1]),$$

$$V(\sigma) = \phi_{j_0}^{-1}(U(\sigma)\times[a_\sigma,b_\sigma]).$$

The collection of all $\mathcal{U}(\sigma)$ (with their respective $V(\sigma)$) is the desired covering of $\mathcal{U}_{j_0}$. This proves the inductive step and concludes the proof of Proposition 4.20. $\square$

4.4 Loose families are flexible

Our main results are almost immediate after setting up all the essential definitions and facts:

Theorem 4.22. *Let M be a smooth 4–manifold. Let K be a compact manifold. Let $\mathcal{D}_0, \mathcal{D}_1 : K \to \mathfrak{Engel}(M)$ be two formally homotopic loose families. Then, they are Engel homotopic through loose families.*

Proof. Let $\mathcal{D}$ be a foliated formal Engel structure in $M \times K \times [0,1]$ connecting $\mathcal{D}_0$ and $\mathcal{D}_1$ as families of formal Engel structures. Since the families $\mathcal{D}_i$, $i = 0,1$, are asymptotically loose, we can extend $\mathcal{D}$ to foliated formal Engel structure in $M \times K \times \mathbb{R}$ such that $\lim_{C \to \pm\infty} \mathcal{H}_{\text{convex}}(\mathcal{D}(-,C)) = \infty$ and all the structures $\mathcal{D}(-,C)$ are Engel for $C \notin [0,1]$. By assumption, there are line fields $\mathcal{Y}_+$ and $\mathcal{Y}_-$ trivialising the structures $\mathcal{D}(-,C)$ with C large and positive or negative, respectively. Let $\mathcal{Y}$ be a trivialising line field for $\mathcal{D}$ that agrees with $\mathcal{Y}_+$ if $C >> 0$ and with $\mathcal{Y}_-$ if $C << 0$.

Then, there exists a triangulation $\mathcal{T}$ of $M \times K \times [-C_0, C_0]$, relative to the boundary $M \times K \times \{\pm C_0\}$, and a cover $\{\mathcal{U}(\sigma)\}$ such that their restrictions to the ends

$$\mathcal{T}_\pm = \{\sigma \in \mathcal{T} \mid \sigma \subset M \times K \times \{\pm C_0\}\}$$

$$\{\mathcal{U}_\pm(\sigma) = \mathcal{U}(\sigma) \cap M \times K \times \{\pm C_0\}\}_{\sigma \in \mathcal{T}_\pm},$$

along with the line fields $\mathcal{T}_\pm$, serve as certificates of t–looseness for $\mathcal{D}_\pm = \mathcal{D}(-, \pm C_0)$.

Now the proof boils down to performing the reduction process in $M \times K \times [-C_0, C_0]$ relative to the ends. We refer the reader to the proof of Theorem 3.55 (Subsection 3.5.5). We go over the main details.

Consider

$$\mathcal{U}_j = \left(\bigcup_{\tau \in \mathcal{T}_\pm} \mathcal{U}(\tau) \right) \cup \left(\bigcup_{\tau \in \mathcal{T}^{(j)}} \mathcal{U}(\tau) \right).$$

Suppose we are performing the reduction process in some $\mathcal{U}(\sigma)$, with $\sigma \notin \mathcal{T}_\pm$ a j–simplex. Then, $\phi(\sigma)(\mathcal{U}_{j-1})$ is of the form $A \times [-\varepsilon, 1+\varepsilon]$, with A some closed set, by Corollary 3.13. By taking C_0 larger (and using the induction hypothesis) we can assume that the convex–type energy in any such A is arbitrarily large. Proceeding as in Theorem 3.55 (or rather, its parametric counterpart Proposition 3.78), we can make the convex–type energy large in $\mathcal{U}(\sigma)$, relative to $A \times [-\varepsilon, 1+\varepsilon]$.

We are left with solving the extension problem in each top dimensional cell. This is an immediate application of Proposition 3.80: since the extension problem is relative to the boundary of the foliated shell, it automatically solves the problem relative to $\mathcal{D}_\pm$. This concludes the proof. $\square$

Remark 4.23. By Theorem 4.15, we have that t–loose families and Greco–loose families are particular examples of asymptotically loose families. However, an asymptotically loose family might not be Greco/t–loose (although it is homotopic to a family that is).

Let us particularise to the case of 0–dimensional families. Consider the following definition:

Definition 4.24. *An Engel structure is said to be* $(0--)$*overtwisted if it is asymptotically loose. Denote the space of overtwisted Engel structures by* $\mathfrak{Engel}_{\mathrm{OT}}(M)$*.*

Then, an immediate corollary of the theorem is the following (the π_0 h–principle):

Corollary 4.25. *The inclusion*

$$\mathfrak{Engel}_{\mathrm{OT}}(M) \to \mathcal{F}\mathfrak{Engel}(M)$$

of the overtwisted Engel structures into the formal Engel structures induces an isomorphism in π_0*.*

Chapter 5

Tangent and transverse submanifolds

In Section 1.5, we discussed some basic material on horizontal immersions. The two most relevant notions were the *Geiges projection* and the *developing map* of an orbit of the kernel. Both of them are important for the contents of this chapter.

Three topics will be addressed: the space of deformations of an horizontal curve, particularly of those tangent to $\mathcal{W}$ (Section 5.1), the h–principle for horizontal immersions (Section 5.2), and the h–principle for transverse immersions (Section 5.3).

In the presence of an Engel structure (or any other distribution), two geometrical conditions are meaningful for a submanifold to satisfy: being tangent or transverse. The punchline of the results proven here is that immersions satisfying either one of these conditions are (almost) completely flexible, just like in the contact case. Future research has to focus on figuring out whether *embedded* submanifolds display rigid behaviour.

5.1 Deformations of horizontal curves

In this section, we will explain, finally, the concept of rigidity for horizontal curves. Given a manifold M and an Engel structure $\mathcal{D}$, we denoted by $\mathcal{HI}(M, \mathcal{D})$ the space of horizontal curves; we endow it with the C^1–topology. Ideally, one would want for this space to be an infinite dimensional manifold, since this would allow us to define interesting operators in it and apply calculus of variations. However, this is far from being the case [7, 39]: whereas curves everywhere transverse to $\mathcal{W}$ have a large space of deformations, segments that are everywhere tangent to it and have "short" developing map (in a sense that we will now define), are actually isolated.

The presence of these curves is a problem for a complete h–principle to hold for horizontal immersions: rigidity is a purely geometrical phenomenon. However, the main point, as we will explain in Section 5.2, is that rigidity is essentially the only problem. Once rigid curves are discarded (and they are not very many), horizontal immersions present a flexible behaviour (see Theorem 5.14).

5.1.1 Local models

Let $(M, \mathcal{D})$ be an Engel manifold with kernel $\mathcal{W}$. Let $\gamma : [-1, 1] \to M$ be an horizontal immersion. We start by explaining the local model around γ in two particular cases: γ is everywhere transverse to the kernel $\mathcal{W}$, or γ is tangent to $\mathcal{W}$ and its developing map is shorter than one projective turn. In particular we will be able to describe the deformations of γ, relative to its ends, as an horizontal curve.

It will be convenient to state these results in parametric form for later use. Let K be some compact manifold that we use as a parameter space.

Lemma 5.1. *Let $\gamma_k : [-1, 1] \to M$, $k \in K$, be a family of horizontal curves that are everywhere transverse to $\mathcal{W}$.*

Then, there are a constant $\varepsilon > 0$ and a family of immersions

$$\phi_k : ([-1, 1] \times \mathbb{D}^2_\varepsilon \times [-\varepsilon, \varepsilon] \subset \mathbb{R}^4, \mathcal{D}_{\text{std}}) \to (M, \mathcal{D}), \qquad k \in K,$$

satisfying $\phi_k(x, 0, 0, 0) = \gamma_k(x)$ and $\phi^*\mathcal{D} = \mathcal{D}_{\text{std}}$.

In the model, all C^1–perturbations of γ_k are of the form $(x, y_k(x), y'_k(x), y''_k(x))$, where

$$y_k : [-1, 1] \to \mathbb{R}, \qquad k \in K.$$

is a family of functions. In particular, the space of compactly supported deformations corresponds to the space of compactly supported functions on the interval $[-1, 1]$.

Proof. What we are effectively doing is saying that any horizontal segment transverse to $\mathcal{D}$ locally looks like the zero section in $J^2(\mathbb{R}, \mathbb{R})$.

Depending on k, we can take two linearly independent vector fields Y_k, Z_k along γ_k such that $\langle \gamma'_k, Y_k, Z_k \rangle$ is a 3–distribution complementary to $\mathcal{W}$. Using the exponential map on $\langle Y_k, Z_k \rangle$ along γ_k, we find an immersion $\psi_k : [-1, 1] \times \mathbb{D}^2_\varepsilon \to M$. Since $\psi_k^*\mathcal{E}$ is a contact structure on this 3–dimensional slice, with $\psi_k^{-1} \circ \gamma_k$ a legendrian curve, we can take a small tubular neighbourhood of γ_k within the slice that is contactomorphic to $([-1, 1] \times \mathbb{D}^2_\varepsilon, \xi = \ker(dy - z dx))$ with γ_k being taken to the central legendrian curve $\tilde{\gamma}_k(x) = (x, 0, 0)$. We assume that ψ_k satisfies exactly this.

Now, on the slice $\text{Image}(\psi_k)$ the Engel structure $\mathcal{D}$ is imprinting a (legendrian) line field. In terms of the model, this line field $\mathcal{X}_k$ is contained in ξ. Along $\tilde{\gamma}_k$, $\mathcal{X}_k$ agrees with $\langle \partial_x + z\partial_y \rangle$, but not anywhere else, necessarily. We need to modify the embedding ψ_k for this to hold everywhere.

Find a vector field W spanning $\mathcal{W}$. As we flow the slice $\text{Image}(\psi_k)$ by W, the preferred legendrian vector field moves. In terms of the local model, $\mathcal{X}_k(x, y, z)$ swipes a cone of lines $C_k(x, y, z) \subset \xi(x, y, z)$. For those points in $\tilde{\gamma}_k$, we have that the cone $C_k(x, 0, 0)$ contains $\langle \partial_x + z\partial_y \rangle = \mathcal{X}_k(x, 0, 0)$, so the same is true in a neighbourhood of $\tilde{\gamma}_k$. This uniquely defines a way of modifying ψ_k graphically to achieve the condition $\psi_k^*\mathcal{D} = \langle \partial_x + z\partial_y \rangle$. We possibly need to take $\varepsilon > 0$ smaller.

Now we take W again to flow the slice, yielding an immersion $\phi_k : [-1, 1] \times \mathbb{D}^2_\varepsilon \times [-\varepsilon, \varepsilon] \to M$ extending ψ_k. The implicit function theorem says that adjusting the length of W appropriately ensures $\phi_k^*\mathcal{D} = \mathcal{D}_{\text{std}}$. The claim about deformations is immediate from the model. □

However, curves tangent to $\mathcal{W}$ are much more rigid, as we pointed out already.

Lemma 5.2. *Let $\gamma_k : [-1,1] \to M$, $k \in K$, be a family of horizontal curves that are everywhere tangent to $\mathcal{W}$ and whose developing maps are shorter than one (projective) turn.*

Then, there are:

- *a family of embeddings $x_k : [-1,1] \to \mathbb{R} \subset \mathbb{R}^4$ satisfying $x(0) = 0$,*
- *a family of immersions :*

$$\phi_k : (\mathcal{O}p(x_k) \subset \mathbb{R}^4, \mathcal{D}_{\mathrm{Lorentz}}) \to (M, \mathcal{D}), \qquad k \in K,$$

satisfying $\phi_k(x_k(s), 0, 0, 0) = \gamma_k(s)$ and $\phi_k^*\mathcal{D} = \mathcal{D}_{\mathrm{Lorentz}} = \ker(dy - tdx) \cap \ker(dz - t^2dx)$.

Any C^1–perturbation of γ is given, in the model, as $(x_k(s), y(x_k(s)), z(x_k(s)), t_k(x_k(s)))$ with

$$y_k(x) = y_k(0) + \int_0^x t_k(u)du$$

$$z_k(x) = z_k(0) + \int_0^x t_k^2(u)du$$

In particular, the expression for z_k implies that there are no deformations relative to the ends.

We are saying that, up to reparametrising γ_k, the model nearby is essentially unique.

Proof. Find some disc D_k transverse to $\mathcal{W}$ and passing through $\gamma_k(0)$. Parametrise it so that the contact structure induced by $\mathcal{E}$ is standard $\ker(dy - zdx)$. In D_k, $\mathcal{D}$ induces a preferred line field and, after suitable reparametrisation of D_k, it can be assumed that the line field is $\langle \partial_x + z\partial_y \rangle$. By thickening γ_k, we obtain an immersion ψ_k of $\mathbb{D}^3_\varepsilon \times [-1,1]$ into M, extending the embedding D_k of $\mathbb{D}^3_\varepsilon \times \{0\}$.

Since the developing map of γ_k is shorter than 1 projective turn, a reparametrisation of ψ_k allows us to assume that $\psi_k^*\mathcal{D} = \mathcal{D}_{\mathrm{std}}$. Effectively, we are just adjusting the speed with which the legendrian vector field is turning by using the implicit function theorem.

Now we recall Subsubsection 1.3.1.1: there is a diffeomorphism $\Psi : \mathbb{R}^4 \to \mathbb{R}^4$ taking $\mathcal{D}_{\mathrm{Lorentz}}$ to $\mathcal{D}_{\mathrm{std}}$. Precomposing ψ_k with Ψ and taking a tubular neighbourhood around $(\psi_k \circ \Psi)^{-1}(\gamma_k)$ yields the result. The claim regarding deformations is immediate from the model. □

Remark 5.3. It is clear that, if the curves γ_k are embedded, the models can be taken to be diffeomorphisms with their images.

5.1.2 A result of Bryant and Hsu

We say that a horizontal curve having no compactly supported C^1–perturbations is **rigid** or **isolated**. According to Lemma 5.2, sufficiently short integral curves of $\mathcal{W}$ are isolated. What about curves having longer developing map? Bryant and Hsu provide an answer to this question:

Proposition 5.4 (Proposition 3.1 in [7]). *Let $(M, \mathcal{D})$ be an Engel manifold. Then, an horizontal immersion $\gamma : [0,1] \to M$ is isolated if and only if it is everywhere tangent to $\mathcal{W}$ and its associated developing map is injective away from its endpoints; i.e. if and only if it makes one projective turn or less.*

In this chapter we are not interested in embeddings of the interval, but embeddings of $\mathbb{S}^1$, since our final aim is understanding the topology of the space $\mathcal{HI}(\mathbb{S}^1, M)$. As such, we will not reprove their result. However, observe that we have already covered the case of those curves with developing map *shorter* than one turn. We will now use Lemma 5.2 to show that, in the case of non–orientable prolongations, the space of simple $\mathcal{W}$–orbits conforms (after quotienting by orientation–preserving reparametrisations) two copies of the base manifold (one for each orientation). Further, these curves possess no other deformations. In this way, we have been able to completely identify two whole connected components in $\mathcal{HI}(\mathbb{S}^1, M)$ that cannot be accounted for formally:

Proposition 5.5. *In a Cartan prolongation* $(\mathbb{P}(\xi), \mathcal{D}(\xi))$ *of* 1 *projective turn, the embedded curves tangent to* $\mathcal{W}(\xi)$ *are isolated, in the* C^1*–topology, from all other curves in* $\mathcal{HI}(\mathbb{S}^1, \mathcal{D}(\xi))$*. They conform two connected components that deformation retract to* $\mathbb{P}(\xi)$*.*

Proof. Everything is reduced to showing that the only C^1–perturbations of such a curve γ are the nearby curves tangent to $\mathcal{W}(\xi)$. Take some curve η that is C^1–close to γ. If it is tangent to a $\mathcal{W}(\xi)$–orbit $\tilde{\gamma}$ in an open interval $I \subset \mathbb{S}^1$, then $\eta|_{\mathbb{S}^1 \setminus I}$ is a compactly supported deformation of $\tilde{\gamma}|_{\mathbb{S}^1 \setminus I}$, which makes less than one projective turn. Applying Lemma 5.2 implies that η is a reparametrisation of $\tilde{\gamma}$.

Otherwise, take some time t_0 where η is transverse to $\mathcal{W}(\xi)$. Lemma 5.1 yields a neighbourhood U of $\eta(t_0)$. The vector field $\eta'(t)|_{[t_0-\delta, t_0+\delta]}$, for δ small, can be extended in U to a vector field X whose flowlines are C^1–close to η (and thus graphical over γ). Now, in $\mathcal{O}p(\eta(t_0)) \subset U$, we perturb X to make it tangent to $\mathcal{W}(\xi)$. A suitable choice yields flowlines that are C^1–close to η in $\mathcal{O}p(\partial U)$, that are still graphical over γ, and that are tangent to $\mathcal{W}$ in an interval. Lemma 5.1 implies that we can take one such flowline and interpolate back to η in $\mathcal{O}p(\partial U)$ to yield a curve $\tilde{\eta}$ that is a C^1–small deformation of γ. The proof concludes by applying to $\tilde{\eta}$ the reasoning in the first paragraph. □

In more general Engel manifolds it is not always true that sufficiently short curves tangent to $\mathcal{W}$ are isolated as tangent immersions. This is explored in Section 5.1.

5.1.3 The Geiges projection for mapping tori

We have just shown, in Proposition 5.5, that the h–principle for horizontal immersions does fail, in general, in the presence of closed orbits of the kernel of the Engel structure. We would like to explore the phenomenon of rigidity in more detail. The main idea is that an Engel structure along a $\mathcal{W}$–orbit is necessarily given by a mapping torus construction, and one is able to define an analogue of the Geiges' projection in this setting. This will allow us to understand perturbations of the orbit in a pictorial way.

Let us explain our setup in detail. Take the standard $(\mathbb{R}^3, \xi = \ker(dy - zdx))$ and let $\phi : (\mathbb{R}^3, \xi) \to (\mathbb{R}^3, \xi)$ be a contactomorphism that fixes the origin. Think about the mapping torus M_ϕ as the quotient $\mathbb{R}^3 \times [0,1]/\phi$ with coordinates (x, y, z, t). M_ϕ can be endowed with a natural even contact structure: the pull–back of ξ, whose kernel is spanned by ∂_t. An Engel structure $\mathcal{D} = \langle \partial_t, L \rangle$ can be defined on M_ϕ, where $L \subset \xi$ is some t-dependent Legendrian vector field rotating positively in the t–direction and satisfying $\langle \phi^*(L(0)) \rangle = \langle L(1) \rangle$.

Fix a framing $\langle X = \partial_x + z\partial_y, Z = \partial_z \rangle$ of ξ. Fix $L(0) = X$. We write $L(1)$ as $\cos(F(x,y,z))X + \sin(F(x,y,z))Z$, where $F : \mathbb{R}^3 \to \mathbb{R}$ is the smallest possible such function that is still positive. F

can be extended to the whole mapping torus to define a possible L:

$$f : \mathbb{R}^3 \times [0,1] \to \mathbb{R}$$

$$f|_{\mathbb{R}^3\times\{1\}} = F, \qquad f|_{\mathbb{R}^3\times\{0\}} = 0, \qquad \partial_t f > 0,$$
$$L = \cos(f(x,y,z,t))X + \sin(f(x,y,z,t))Z.$$

Therefore, the structure equations for the Engel structure $\mathcal{D}$ read as:

$$\alpha = dy - zdx, \qquad \beta = \cos(f(x,y,z,t))dz - \sin(f(x,y,t,z))dx.$$

Observe that there are many possible definitions of L for a given F, but they all yield diffeomorphic Engel structures.

Consider the $\mathcal{W}$–integral curve $\gamma(\theta) = (0,0,0,\theta)$. Any C^1–small deformation of γ is of the form $\eta(\theta) = (x(\theta), y(\theta), z(\theta), \theta)$, and satisfies the equations:

$$\begin{aligned} \tan(f(x,y,z,t)) &= \frac{z'}{x'}, \\ y(t) - y(0) &= \int_0^t zx' ds, \\ \phi(\eta(1)) &= \eta(0). \end{aligned}$$

We say that the plane curve $\pi \circ \eta(\theta) = (x(\theta), z(\theta))$ is the front (of the Geiges projection) of η. These formulas in particular describe how to recover η from its front.

5.1.4 Revisiting Proposition 5.5

We can use the language we have just introduce to reprove Proposition 5.5 in a more geometric way. This will convey the types of arguments that can be carried out using the Geiges projection (or rather, its front).

Alternate proof of Proposition 5.5. Given a C^1–perturbation η of a $\mathcal{W}$–tangent curve γ, we want to show that η is tangent to $\mathcal{W}$ as well. Suppose otherwise; by Theorem 5.12 (below), we can assume that η is in general position with respect to $\mathcal{W}$. We can find a neighbourhood of γ that is a mapping torus M_ϕ with ϕ the identity and $L(1) = -L(0)$; we are in the setup above, with $f(x,y,z,t) = \pi t$. Due to it being in general position, the front $\pi \circ \eta$ is a closed plane curve with cusps.

The front must possess at least one cusp and, choosing our neighbourhood suitably, we assume that $\pi \circ \eta$ has, at $\pi \circ \eta(0) = 0$, a cusp pointing to the left. The first equation above states that the slope of η rotates clockwise π degrees, and thus the curve is piecewise convex. The second one says that the signed area bounded by η must be zero.

Observe that the number of cusps must be odd since the oriented slope approaching $t = \pi$ must be horizontal and pointing to the left and at every cusp the orientation changes sign. Denote the values of the parameter for which the curve has a cusp by $\{t_0 = 0 = \pi, t_1, \dots, t_{2n}\}$. Since the slope is only horizontal at the endpoints, the cusps are alternating; that is, at t_{2i-1} the curve leaves the horizontal line $\{z = z(t_{2i-1})\}$ going downwards and at t_{2i} it leaves it going upwards. In other words, the function $z(t)$ is strictly increasing in the intervals (t_{2i}, t_{2i+1}) and strictly decreasing otherwise. We deform the curve by *enlarging the cusps*: Except for the one at the origin, we add a straight segment to the end of each of the cusps and then we make it convex by a slight deformation (like glueing a thickened needle to its end). This procedure allows us, without changing the total area, to push upwards the odd cusps and push downwards the even ones. Therefore, for $i > 0$:

$$\begin{aligned} z(t_{2i-1}) &= z(t_1) > 0, \\ z(t_{2i}) &= z(t_2) < 0. \end{aligned}$$

Consider the segments $\pi \circ \eta|_{(t_{2i-1},t_{2i})}$ and $\pi \circ \eta|_{(t_{2i},t_{2i+1})}$, and reverse the parametrisation of the former. Then, both of them are segments starting from $\pi \circ \eta(t_{2i})$ and finishing in the same z–coordinate, but the latter has greater slope. This reasoning readily implies that:

$$x(t_1) > x(t_3) > \cdots > x(t_{2n-1}),$$

$$x(t_2) < x(t_4) < \cdots < x(t_{2n}).$$

In particular, the segments $\pi \circ \eta|_{(t_{2i-1},t_{2i})}$ and $\pi \circ \eta|_{(t_{2i+1},t_{2i+2})}$ intersect at a point s_i. This means that in-between t_{2i-1} and t_{2i+2} a *Reidemeister I move* configuration appears, bounding some positive area. Refer to Figure 5.1.

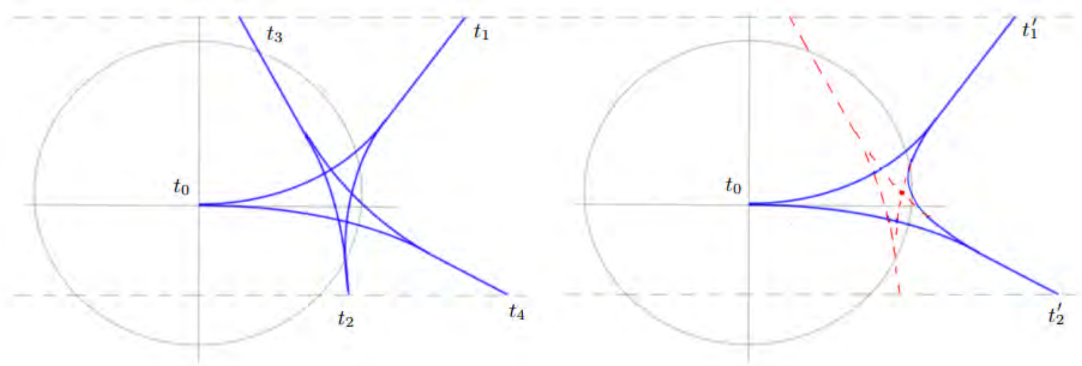

Figure 5.1: On the left hand side, a possible projection for a deformation η with five cusps. On the right hand side, we outline in red the area of the Reidemeister I loop, that has been removed, yielding a curve with only three cusps. Note that the cusps have been made longer so that they would reach the horizontal gray lines.

Now we conclude by induction on $2n+1$, the number of cusps. Our induction hypothesis is that a front conforming to the properties above must bound positive area. This is straightforward for $2n+1 = 3$. For the induction step, the reasoning on the previous paragraph shows that, for $2n+1 > 3$, a Reidemeister I move appears. By removing it (along with the points t_{2i} and t_{2i+1}) and smoothing the curve at s_i, the points t_{2i-1} to t_{2i+2} are now connected by a segment with no cusps. Since the area under this operation decreases and now the number of cusps is $2n-1$, the induction hypothesis concludes the proof. □

5.1.5 Short $\mathcal{W}$–orbits admitting deformations

Using the mapping torus construction of Subsection 5.1.3, it is clear that one can construct $\mathcal{W}$–orbits whose developing map is shorter than one projective turn. We will now go through some examples of this and we will show that, in many cases, these orbits do admit deformations (that are everywhere not tangent to $\mathcal{W}$ themselves).

Example 5.6 (Curves making one projective turn.)**.** Take the mapping torus M_ϕ with $\phi(x, y, z) = (x, y/2, z/2)$. Fix $L(0) = -L(1) = \partial_x + z\partial_y$. Let η be the desired deformation of $(0, 0, 0, \theta)$, which we assume is in general position with respect to $\mathcal{W}$. Its front $\pi \circ \eta(\theta) = (x(\theta), z(\theta))$ satisfies $(x(0), z(0)) = (x(1), 2z(1)) = (0, 0)$ and encloses an area of $y(1)/2$. On the left hand side of Figure 5.2, such a curve is presented; it is clear that the area it bounds can be adjusted to be exactly $y(1)/2$. $\square$

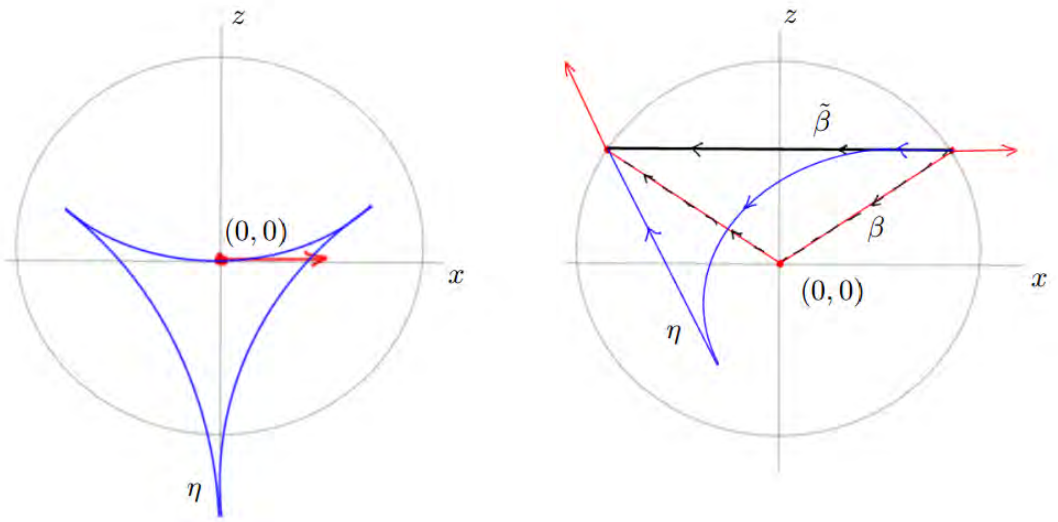

Figure 5.2: On the left hand side, a possible deformation for a curve making one projective turn. On the right, a deformation for a curve with short developing map. The curves are depicted in blue. The tangent vectors at $t = 0, 1$ are shown in red.

Example 5.7 (Curves having an arbitrarily short developing map.)**.** Fix some angle $\alpha \in (0, \pi)$. The following contactomorphism is the lift of the turn of angle $-\alpha$ in the plane (x, z):

$$\psi(x, y, z) = (\cos(\alpha)x + \sin(\alpha)z, y - \sin^2(\alpha)zx + \frac{1}{2}\cos(\alpha)\sin(\alpha)(z^2 - x^2), \cos(\alpha)z - \sin(\alpha)x).$$

We consider the mapping torus of ψ.

Take a deformation η ending at $(x(1), y(1), z(1))$. The projection $\pi \circ \eta$ must bound a signed area of

$$y(1) - y(0) = \sin^2(\alpha)z(1)x(1) - \frac{1}{2}\cos(\alpha)\sin(\alpha)[z(1)^2 - x(1)^2].$$

The right hand side is precisely the integral of zdx over the curve β given by going from $(x(0), z(0))$ to the origin and then to $(x(1), z(1))$ following straight lines, as a computation shows.

Consider $(x(1), z(1))$ lying in the first quadrant and making an angle of $\alpha/2$ with the vertical axis. Let $\tilde{\beta}$ be the straight horizontal segment connecting $(x(0), z(0))$ and $(x(1), z(1))$. In particular, it lies above β and thus $\int_\beta zdx > \int_{\tilde{\beta}} zdx$. Now it is straightforward to create a curve η such that $\int_\eta zdx = \int_\beta zdx = y(1) - y(0)$ by adding some (positive) area to $\tilde{\beta}$ and adjusting it to ensure that it consistently turns clockwise. Refer to the right hand side of Figure 5.2. $\square$

Slightly generalising the first example, it is not hard to show that:

Proposition 5.8. *Let ϕ be a contactomorphism of $(\mathbb{R}^3, \xi = \ker(dy - zdx))$ fixing the origin and with conformal factor different from 1. Let M_ϕ be the corresponding mapping torus with coordinates (x, y, z, t) and endowed with the Engel structure with smallest turning. Then, the $\mathcal{W}$–curve $\gamma(\theta) = (0, 0, 0, \theta)$ admits deformations somewhere not tangent to $\mathcal{W}$.*

Proof. Take $d_0\phi$, the linearisation at the origin. $d_0\phi|_\xi$ is a linear map in $\mathbb{R}^2$ that can be lifted to a contactomorphism $\tilde{\phi}$. By zooming in with the contactomorphism $(x, y, z) \to (\lambda x, \lambda^2 y, \lambda z)$, ϕ becomes C^∞ close to $\tilde{\phi}$, and therefore it is enough to prove the statement for ϕ linear.

If the conformal factor at the origin is different from 1, there is a dilation in the y–coordinate. Then we construct a deformation starting and finishing at the origin and bounding an area $y(1) - y(0) > 0$, which is possible if we select $y(0)$ small enough and with the adequate sign. □

It is clear that we can start with a non–orientable Cartan prolongation and perturb its return map to insert any of these models.

5.2 The h–principle for horizontal immersions

5.2.1 Generic horizontal immersions

Having somewhat understood the extent to which the h–principle fails, we can finally state our main result about horizontal immersions. Before we do so, we have to restrict our attention to those horizontal curves that are not rigid. To be on the safe side of things, we will just discard all curves that are $\mathcal{W}$–orbits. This yields the following definition:

Definition 5.9. *Let I be a 1–dimensional manifold. We denote by $\mathcal{HI}^{\text{n.e.t.}}(I, \mathcal{D}) \subset \mathcal{HI}(I, \mathcal{D})$ the open subspace of those horizontal curves that are not everywhere tangent to $\mathcal{W}$.*

Of course, we are interested in the case $I = \mathbb{S}^1$.

How large is the set of orbits we are discarding? In the most symmetric case, that of Cartan prolongations, the space of closed $\mathcal{W}$–orbits is a countable collection of copies of the base manifold. Two of these copies, those corresponding to simple curves, are independent connected components in the space of horizontal curves (Proposition 5.5). The other copies (corresponding to multiply covered orbits) are not isolated, since the curves with longer developing maps do admit deformations (Proposition 5.4).

Under the opposite assumption, $\mathcal{D}$ being generic, Proposition 1.32 states that we are discarding countably many orbits. However, only those with short developing map are isolated.

Remark 5.10. Recall that there are Engel structures C^∞–close to prolongations that have no closed orbits, so we immediately deduce that the homotopy type of $\mathcal{HI}(\mathbb{S}^1, \mathcal{D})$ is not invariant under deformations.

Consider this other definition:

Definition 5.11. *Let I be a 1–dimensional manifold. We denote by $\mathcal{HI}^{\text{gen}}(I, \mathcal{D}) \subset \mathcal{HI}^{\text{n.e.t.}}(I, \mathcal{D})$ the open subspace of those horizontal curves whose tangencies with $\mathcal{W}$ are isolated.*

$\mathcal{HI}^{\text{gen}}(I, \mathcal{D})$ is a nicer space than $\mathcal{HI}^{\text{n.e.t.}}(I, \mathcal{D})$ in the sense that the curves it contains are defined by a condition at the level of germs, as opposed to a global condition. It is therefore easier to work with $\mathcal{HI}^{\text{gen}}(I, \mathcal{D})$ in certain situations (this will be the case in Section 5.3).

5.2.2 Statement of the results

Let us state the two main theorems of this section. The first one is a genericity result: given a curve that is somewhere not tangent to $\mathcal{W}$ (even if it is tangent to it over an open set), it can be deformed to a curve with isolated $\mathcal{W}$–tangencies. We state it in parametric form:

Theorem 5.12. *Let M be a 4–manifold. Let K be a closed manifold and fix a map $\mathcal{D} : K \to \mathfrak{Engel}(M)$. Denote by $\mathcal{W}(k)$ the kernel of $\mathcal{D}(k)$, $k \in K$.*

Consider $\gamma : K \to \mathcal{I}(M)$ satisfying $\gamma(k) \in \mathcal{HI}^{\text{n.e.t.}}(\mathbb{S}^1, \mathcal{D}(k))$. Then, after a C^∞–perturbation, it can be assumed that the set

$$\{(k,s) \in K \times \mathbb{S}^1 \ ; \gamma(k)'(s) \in \mathcal{W}(k)_{\gamma(k)(s)}\}$$

is a submanifold of codimension–1 in $K \times \mathbb{S}^1$ in generic position with respect to the foliation $\coprod_{k\in K}\{k\} \times \mathbb{S}^1$.

An immediate corollary is the following:

Corollary 5.13. *Let $(M, \mathcal{D})$ be an Engel manifold. The inclusion:*

$$\mathcal{HI}^{\text{gen}}(\mathbb{S}^1, \mathcal{D}) \to \mathcal{HI}^{\text{n.e.t.}}(\mathbb{S}^1, \mathcal{D})$$

is a weak homotopy equivalence.

Theorem 5.12 and Corollary 5.13 are technical results that allow us to pass to a setting where manipulating horizontal curves is easier. They are important in the proof of our main result about horizontal immersions:

Theorem 5.14. *Let $(M, \mathcal{D})$ be an Engel manifold. The inclusion*

$$\mathcal{HI}^{\text{n.e.t.}}(\mathbb{S}^1, \mathcal{D}) \to \mathcal{FHI}(\mathbb{S}^1, \mathcal{D})$$

is a weak homotopy equivalence.

This result pretty much settles the classification of horizontal immersions. However, admittedly, we still need a better understanding of the deformations of the $\mathcal{W}$–orbits to really determine the homotopy type of $\mathcal{HI}(\mathbb{S}^1, \mathcal{D})$.

Remark 5.15. Instead of Theorem 5.14, we will prove the following slightly stronger result: Let K be some compact m–dimensional manifold, possibly with boundary, and fix $\mathcal{D} : K \to \mathfrak{Engel}(M)$. Let $\phi : K \to \mathcal{FI}(\mathbb{S}^1, M)$ be a map satisfying:

- $\phi(k) \in \mathcal{HI}^{\text{n.e.t.}}(\mathbb{S}^1, \mathcal{D}(k))$ for $k \in \partial K$,
- $\phi(k) \in \mathcal{FHI}(\mathbb{S}^1, \mathcal{D}(k))$ for all $k \in K$.

Then, ϕ is homotopic, relative to ∂K, to a map $\tilde{\phi} : K \to \mathcal{I}(\mathbb{S}^1, M)$ with $\tilde{\phi}(k) \in \mathcal{HI}^{\text{n.e.t.}}(\mathbb{S}^1, \mathcal{D}(k))$ for all k.

5.2.3 The genericity result

Theorem 5.12 states that, once one restricts to the subspace $\mathcal{HI}^{\text{n.e.t.}}(\mathbb{S}^1, \mathcal{D})$, there are enough deformations to guarantee a "*generic*" picture. We shall dedicate the rest of this subsection to its proof. It relies on local C^∞–small deformations using Lemmas 5.1 and 5.2.

Setup. Let us construct an adequate cover of the space $K \times \mathbb{S}^1$. Locally, for every $(k, s) \in K \times \mathbb{S}^1$, we can find vector fields $W(k', s)$ spanning $\mathcal{W}(k')$. Denote by $\eta(k', s)$ the integral curve of $W(k', s)$ with domain $[s - \varepsilon, s + \varepsilon]$ and satisfying $\eta(k', s)(s) = \gamma(k')(s)$.

We can apply Lemma 5.2 parametrically to the curves $\eta(k', s)$. If $\gamma(k)'(s) \in \mathcal{W}(k)$, there is a product neighbourhood $\mathbb{D}_\varepsilon(k) \times [s - \varepsilon, s + \varepsilon] \ni (k, s)$ in which every curve $\gamma(k')$, $k' \in \mathbb{D}_\varepsilon(k)$, is graphical over $\eta(k', s)$ in the model. We say that this neighbourhood is of type I.

Otherwise, if (k, s) is such that $\gamma(k)'(s)$ is transverse to $\mathcal{W}(k)$, so are the nearby curves. We use Lemma 5.1 parametrically to yield a product neighbourhood of (k, s) in which the curves $\gamma(k')$ look like the zero section in $J^2(\mathbb{R}, \mathbb{R})$. We call this a neighbourhood of type II.

Then, by compactness of $K \times \mathbb{S}^1$, we can find a finite cover $\{U_{i,j}\}$ comprised of neighbourhoods like the ones we just described. We assume that it is the product of a covering $\{W_i\}$ in K and a covering $\{V_j = \mathcal{O}p([\frac{j}{N}, \frac{j+1}{N}])\}$, $j = 0, .., N - 1$, in $\mathbb{S}^1$. We order the neighbourhoods $\{U_{i,j} = W_i \times V_j\}$ as follows: for any fixed W_i, we find some $j_i \in \{0, ..N - 1\}$ such that $W_i \times V_{j_i}$ is of type II and we order the $W_i \times V_j$ cyclically increasing from $j = j_i + 1$ to $j = j_i$. The order in which we consider each W_i is not important and hence we just proceed as we numbered them. See Figure 5.3.

The idea now is to modify γ over the neighbourhoods $U_{i,j}$ inductively using the order we just constructed. Over those of type I we will deform to achieve the desired transversality. Over those of type II we have more flexibility, so we shall use them to ensure that the deformation $\tilde{\gamma}$ does close up.

Take a neighbourhood $U_{i,j}$. Denote by $\tilde{U}_{i,j}$ the union of the neighbourhoods over which a C^∞–close deformation $\tilde{\gamma}$ of γ has been defined already.

Type I neighbourhoods. Assume that $U_{i,j}$ is of type I. Applying Lemma 5.2, we have a family of curves

$$\gamma(k) : V_j \to (\mathbb{R}^4, \ker(dy - tdx) \cap \ker(dz - t^2 dx)), \qquad k \in W_i$$

that are graphical over the x axis and thus given by functions:

$$\gamma(k)(s) = (x_k(s), y_k(x_k), z_k(x_k), t_k(x_k))$$

with $x_k(s)$ a diffeomorphism with its image, $t_k(x)$ some arbitrary function, and

$$\begin{cases} y_k(x_k(s)) = y_k(x_k(-1)) + \displaystyle\int_{x_k(-1)}^{x_k(s)} t_k(x)dx \\ z_k(x_k(s)) = z_k(x_k(-1)) + \displaystyle\int_{x_k(-1)}^{x_k(s)} t_k^2(x)dx \end{cases} \tag{5.1}$$

where the dependence on k is smooth. Analogously, $\tilde{\gamma}$ is defined by functions $(\tilde{x}_k, \tilde{y}_k, \tilde{z}_k, \tilde{t}_k)$ which are only defined over $U_{i,j} \cap \tilde{U}_{i,j}$.

Tangencies with $\mathcal{W}$ are given by $t'_k, \tilde{t}'_k = 0$. We extend $\tilde{t}_k$ from $U_{i,j} \cap \tilde{U}_{i,j}$ to the whole of $U_{i,j}$ arbitrarily, ensuring that it remains C^∞–close to t_k and that it has generic critical points (for a

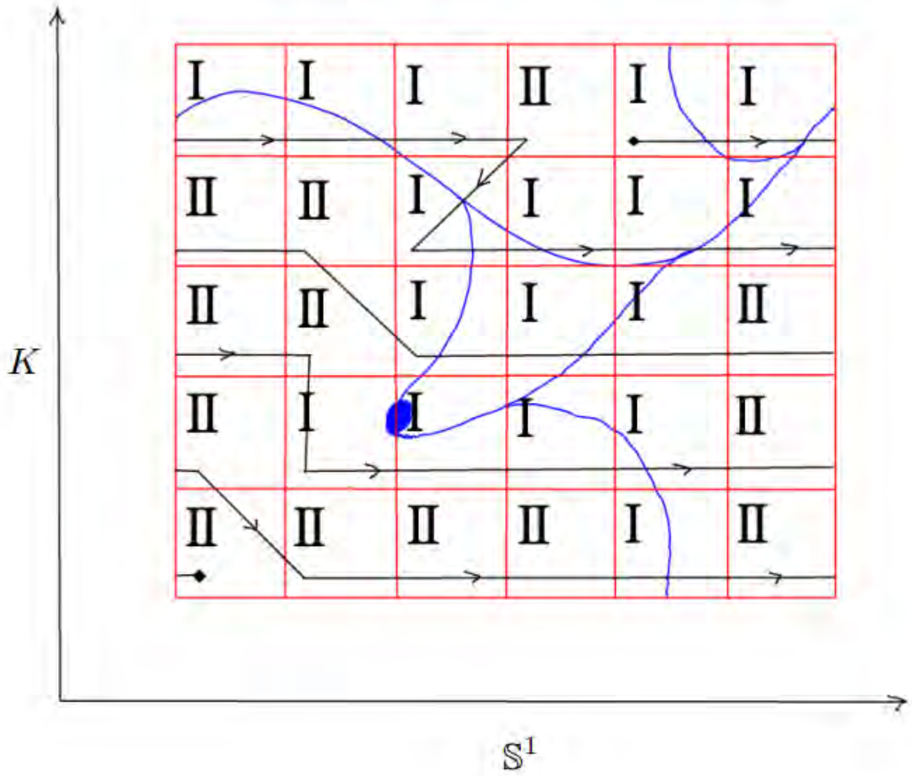

Figure 5.3: In red we depict the manifold $K \times \mathbb{S}^1$. The tangencies with $\mathcal{W}$ are the points in blue. The sets $U_{i,j}$ correspond to neighbourhoods of the smaller red squares; they have been numbered as I or II depending on whether they are of the first or second type. The black line with arrows indicates the order in which we proceed for the induction.

family of dimension $\dim(K)$). We can extend $\tilde{y}_k$ and $\tilde{z}_k$ to the whole of $U_{i,j}$ using the integral expressions (5.1) with initial values those in $U_{i,j} \cap \tilde{U}_{i,j}$. The order that we chose for the induction means that $U_{i,j} \cap \tilde{U}_{i,j} \cap (\{k\} \times \mathbb{S}^1)$ is connected at every such step, so in particular we are defining $\tilde{y}_k$ and $\tilde{z}_k$ as integrals with boundary conditions given only at one end of the interval V_j. Note that this construction is indeed relative to $\tilde{U}_{i,j}$.

Type II neighbourhoods. Assume that $U_{i,j}$ is of type II. In its local model, given by Lemma 5.1, the perturbations $\tilde{\gamma}(k)$ (which are defined only over $\tilde{U}_{i,j}$) can be assumed to be graphical over $\gamma(k)$, which are seen as intervals contained in the zero section of $J^2(\mathbb{R}, \mathbb{R})$. $\tilde{\gamma}$ is thus described by a family of functions $\tilde{y}_k$ and their first and second derivatives $\tilde{z}_k$ and $\tilde{t}_k$, respectively. Extend $\tilde{y}_k$ arbitrarily to $U_{i,j}$ while keeping it C^∞ close to y_k; take $\tilde{z}_k$ and $\tilde{t}_k$ to be the corresponding derivatives of $\tilde{y}_k$. This can be done regardless of the boundary conditions, and this is the reason why we left a neighbourhood of type II for last. No additional tangencies with $\mathcal{W}$ are introduced doing this. $\square$

Remark 5.16. In type II neighbourhoods, after extending $\tilde{y}_k$, one could construct a bump function $\psi : U_{i,j} \to \mathbb{R}$ that is identically 1 near $\partial U_{i,j}$ and identically zero in a slightly smaller ball and then take $\psi\tilde{y}_k$ and its derivatives as the desired extensions to the whole of $U_{i,j}$. In this manner, by taking the cover to be fine enough, one can strengthen Theorem 5.12 saying that the deformation $\tilde{\gamma}$ can be taken to agree with γ in an arbitrarily large closed set disjoint from the tangencies.

5.2.4 The theorem of Adachi and Geiges

To showcase some of the ideas that we will need for Theorem 5.14, we will first prove the following theorem of Adachi and Geiges [1, 28]:

Proposition 5.17. *Horizontal knots in* $(\mathbb{R}^4, \mathcal{D}_{\text{std}})$ *are classified, up to homotopy, by their rotation number with respect to* $\mathcal{W}$.

One of the main ingredients of the proof is the h–principle for legendrian immersions:

Proposition 5.18. *A* C^0*–dense, parametric, relative, and relative to the parameter* h*–principle holds for legendrian immersions in contact manifolds.*

Proof. First of all, observe that in dimension 4 there is no smooth knotting of $\mathbb{S}^1$. Computations similar to those carried out in Exercise 1.25 show that the only other formal invariant in π_0 (apart from the smooth type), is the rotation number, which is defined exactly as in the case of immersions.

The proof that we shall outline is Geiges', although our naming convention for the variables is different from his.

From Lemma 1.27, we deduce that the image of an horizontal knot $\gamma : \mathbb{S}^1 \to \mathbb{R}^4$ under the Geiges projection is a legendrian immersion that satisfies the additional area constraint $\int_\gamma zdx = 0$. The rotation number of γ agrees precisely with the rotation (as a legendrian immersion) of its Geiges projection.

Given two horizontal immersions, we can use the h–principle for legendrian immersions to find a homotopy between their Geiges projections. The area constraint to lift to a homotopy through horizontal curves does not hold a priori, but this can be readily remedied by adding some bumps in the front projection to add or substract area.

Now the issue is ensuring that the curves remain embedded along the homotopy, but this follows by genericity. First, note that there are only finitely many curves in the homotopy that have a self–tangency, generically. By Lemma 1.29, these self–tangencies will lift to a self–intersection only if one of the branches starting and ending there bound area zero. This is a codimension–1 condition and the self–tangencies are isolated, so generically this can be avoided. $\square$

We shall see that there are two main points that we have to be address when trying to adapt Geiges' argument to prove our h–principle. First, Engel Darboux balls are, generally, "smaller" than a full $\mathbb{R}^4$, and Geiges' argument uses the whole of $\mathbb{R}^4$ to guarantee that the area constraint is satisfied. Secondly, and more subtly, in general Engel manifolds one has to account for the presence of rigid curves, as we have often observed out (of which there are none in $\mathbb{R}^4$).

It is worth remarking that the genericity argument used in Proposition 5.17 does not work for 1–parametric families of embeddings anymore: self–intersections generically appear conforming a codimension–2 stratum in a family. A fundamental open problem in Engel topology is whether there exist two loops of horizontal knots that are homotopic formally but not geometrically.

5.2.5 The h–principle for horizontal immersions

This subsection is dedicated to proving Theorem 5.14. Most of the work needed for the theorem is contained in the following proposition, which states that a parametric, relative (with respect to some some subset B of the interval), relative in the parameter (with respect to some subset A of the parameter space K), C^0–close h–principle holds for horizontal immersions of the interval.

Proposition 5.19. *Let $(M = \mathbb{R}^3 \times (-\varepsilon, \varepsilon), \mathcal{D} = \mathcal{D}_{\text{std}})$. Let $A \subset \partial\mathbb{D}^m$ be some closed CW–complex. Let $B \subset [0,1]$ be either $\{0,1\}$, $\{0\}$ or the emptyset. Fix a map $\psi_k \in \mathcal{FHI}(\mathcal{O}p([0,1]), \mathcal{D})$, $k \in \mathcal{O}p(\mathbb{D}^m)$, conforming to the following properties:*

- *$\psi_k \in \mathcal{HI}^{\text{n.e.t.}}(\mathcal{O}p([0,1]), \mathcal{D})$ for $k \in \mathcal{O}p(A)$.*
- *ψ_k is horizontal with respect to $\mathcal{D}$ for $s \in \mathcal{O}p(B)$.*

Then, there is a homotopy $\psi_k^\delta \in \mathcal{FHI}(\mathcal{O}p([0,1]), \mathcal{D})$, $\delta \in [0,1]$, starting at $\psi_k^0 = \psi_k$, such that:

- *$\psi_k^1|_{[0,1]} \in \mathcal{HI}^{\text{n.e.t.}}([0,1], \mathcal{D})$ for $k \in \mathbb{D}^m$,*
- *$\psi_k^\delta = \psi_k$ for $k \in \mathcal{O}p(A)$ or $s \in \mathcal{O}p(B)$,*
- *$\psi_k^\delta = \psi_k$ away from $\mathbb{D}^m \times [0,1]$,*
- *writing $(\gamma_k^\delta, F_k^\delta)$ for the two components of ψ_k^δ, γ_k^δ is C^0–close to γ_k^0.*

Remark: We define the domain to be $\mathcal{O}p(\mathbb{D}^m) \times \mathcal{O}p([0,1])$ so that ψ_k^δ glues with ψ_k away from $\mathbb{D}^m \times [0,1]$. This is important when proving a relative statement. In this direction, by $\mathcal{O}p(A)$ we mean an arbitrarily small neighbourhood of the radial projection of A to $\partial\mathcal{O}p(\mathbb{D}^m)$ that still contains A. The definition of $\mathcal{O}p(B)$ is analogous.

Let us explain how to deduce our main theorem using Proposition 5.19.

Proof of Theorem 5.14. We will prove Remark 5.15 instead. Denote by K' the complement of a small collar neighbourhood of ∂K such that $\phi(k) \in \mathcal{HI}^{\text{n.e.t.}}(\mathbb{S}^1, \mathcal{D}(k))$ for all $k \notin K'$. After applying Theorem 5.12, we can assume that the curves $\phi(k)$, $k \in \mathcal{O}p(\partial K')$, are in general position with respect to $\mathcal{W}(k)$.

Recall now Theorem 3.9 (or rather, Corollary 3.13). It explained how to construct a triangulation adapted to a line field. Our setting is exactly like that: we shall take triangulate $K' \times \mathbb{S}^1$ with respect to the foliation $\coprod_{k \in K} \{k\} \times \mathbb{S}^1$ and with respect to the boundary $(\partial K') \times \mathbb{S}^1$. Denote by $\pi_{K'}$ and $\pi_{\mathbb{S}^1}$ for the obvious projections of $K' \times \mathbb{S}^1$ to its factors.

Denote by $\mathcal{T}$ the triangulation produced and by $\{\mathcal{U}(\sigma)\}$ the associated cover. Each element of the cover has a trivialising map $F(\sigma) : \mathcal{U}(\sigma) \to \mathbb{D}^m \times [0,1]$. For a pictorial aid, look at Figures 3.1 and 3.2.

We shall deform ϕ over the neighbourhoods $\mathcal{U}(\sigma) \subset K' \times \mathbb{S}^1$, for those σ not contained in $\partial K'$. We proceed inductively on the dimension of σ. Observe that, since $\mathcal{T}$ was a jiggling of an arbitrary triangulation, we can assume that $\mathcal{T}$ is fine enough so that $s \to \phi(k)(s)$, for $(k,s) \in \mathcal{U}(\sigma)$, maps into a Darboux ball B_k for $\mathcal{D}(k)$. If the triangulation was fine enough, we can parametrically identify the balls B_k with the k–independent Darboux ball $(M = \mathbb{R}^3 \times (-\varepsilon, \varepsilon), \mathcal{D} = \mathcal{D}_{\text{std}})$.

Now we apply Proposition 5.19 to the map $\psi_k(s) = \phi(\pi_{K'} \circ F^{-1}(k,s))(\pi_{\mathbb{S}^1} \circ F^{-1}(k,s))$. If $\dim(\sigma) < m+1$, we take $B = \emptyset$ and A as in the enumeration above. If σ is top dimensional, we take $A = \partial\mathbb{D}^m$ and $B = \{0,1\}$. $\square$

Now, the strategy of proof for Proposition 5.19 is quite similar in spirit to the proof of Proposition 5.17. We first want to use the h–principle for legendrian immersions to do some work for us (Proposition 5.18). This auxiliary result will be enough for the lower dimensional cells, but for the top ones we will need some adjustments. We break down the proof of Proposition 5.19 in several steps.

Step I. The image of $M = \mathbb{R}^3 \times (-\varepsilon, \varepsilon)$ under the Geiges projection π_{Geiges} is $V = \mathbb{R}^2 \times (-\varepsilon, \varepsilon)$ with coordinates (x, z, t). Horizontal immersions descend to legendrian immersions for the standard contact structure $\xi = \ker(dz - tdx)$. In particular, tangencies with $\mathcal{W}$ upstairs are in correspondence with tangencies downstairs with $\langle \partial_t \rangle$. From this, it follows that, whenever ψ_k is horizontal and generic ($k \in \mathcal{O}p(A)$ or $s \in \mathcal{O}p(B)$), $\pi_{\text{Geiges}} \circ \psi_k$ is in general position with respect to the legendrian foliation given by $\langle \partial_t \rangle$, and thus the singularities of its front are generic. Do note that, since we work with higher dimensional families, singularities more complicated than cusps do appear.

Let us denote $\mathcal{L}eg(V, \xi)$ for the legendrian immersions of the interval $\mathcal{O}p([0,1])$ into (V, ξ) and $\mathcal{FL}eg(V, \xi)$ for its formal counterpart. Much like in the case of horizontal immersions, a formal legendrian immersion is a pair comprised of a map into V and a monomorphism into ξ (in this case, both with domain the interval).

Since $d\pi_{\text{Geiges}}$ maps $\mathcal{D}$ isomorphically onto ξ, the Geiges projection yields a family

$$\begin{aligned} \Psi_k^0 = \pi_{\text{Geiges}} \circ \psi_k \in \quad & \mathcal{FL}eg(V, \xi), && k \in \mathcal{O}p(\mathbb{D}^m), \\ \Psi_k^0 \in \quad & \mathcal{L}eg(V, \xi), && k \in \mathcal{O}p(A), \end{aligned}$$

which is already legendrian for $s \in \mathcal{O}p(B)$. By Proposition 5.18, Ψ_k^0 is homotopic, relative to A and B, to a family $\Psi_k^{1/2} \in \mathcal{L}eg(V, \xi)$ for all k. We can further assume that the front of $\Psi_k^{1/2}$ has generic singularities as well. Denote by $\Psi_k^\delta = (\eta_k^\delta, G_k^\delta)$, $\delta \in [0, 1/2]$, the homotopy as formal legendrians.

Step II. Let us construct a lift $\psi_k^\delta = (\gamma_k^\delta, F_k^\delta)$ of Ψ_k^δ. Since $\mathcal{D}$ projects to ξ under the Geiges projection, we define F_k^δ to be the unique lift of G_k^δ. For γ_k^δ, let us focus first on the case where B is empty or $\{0\}$. Define its y–coordinate $y_k^\delta(s)$, for $(k, s) \in (\mathbb{D}^m \times [0,1]) \cup (\mathcal{O}p(A) \times \mathcal{O}p([0,1]))$, to be given by:

$$y_k^\delta(s) = y_k^0(0) + \int_{\eta_k^\delta|_{[0,s]}} z dx.$$

In the complement, we extend γ_k^δ by interpolating back to γ_k^0. This construction guarantees $\gamma_k^\delta = \gamma_k^0$ for $k \in \mathcal{O}p(A)$.

Step III. If $B = \{0,1\}$, defining y_k^δ by integration means that the y–coordinate of γ_k^δ will not necessarily agree with that of γ_k^0 at $s = 1$, as it should. The idea is to deform $\eta_k^{1/2}$ to yield a new Geiges projection η_k^1 having this integral adjusted. Note that we cannot do wild deformations: for a legendrian not to escape the local model $V = \mathbb{R}^2 \times (-\varepsilon, \varepsilon)$, its front must have a slope bounded in terms of ε. Instead, we introduce type I Reidemeister moves to add or subtract area.

Recall that the front of $\eta_k^{1/2}$ has generic singularities. In particular, given any point $k \in \mathbb{D}^m$, there is s_k such that the curve $\eta_k^{1/2}$ is not tangent to $\langle \partial_t \rangle$ at time s_k. It follows that we can find a small disc $\mathcal{U}_k \subset \mathbb{D}^m$ containing k and an interval $I_k \subset [0,1]$ containing s_k such that the curves $s \to \psi_{k'}^{1/2}(s)$, $(k', s) \in \mathcal{U}_k \times I_k$, are transverse to $\langle \partial_t \rangle$ and therefore their front projection is an interval without cusps. By compactness, a finite number of open subsets $\mathcal{U}_k$ disjoint from A cover $\mathbb{D}^m \setminus \mathcal{O}p(A)$.

Given any even integer N, find an ordered sequence of times $s_k^1, .., s_k^N \in I_k$ and a width $\epsilon > 0$ such that the segments $[s_k^j - \epsilon, s_k^j + \epsilon] \subset I_k$ do not overlap. We construct η_k^1 as follows. Replace the front of the curves $\eta_k^{1/2}|_{[s_k^j - \epsilon, s_k^j + \epsilon]}$, for $k \in \mathcal{U}_k$, by adding a "Reidemeister I" loop such that the sign of the area it encloses is given by the parity of j. Modify the fronts of $\eta_k^{1/2}$, for $k \in \mathcal{O}p(\mathcal{U}_k) \setminus \mathcal{U}_k$, so that they transition, through Reidemeister I moves, from agreeing with those of $\eta_k^{1/2}$ in $\partial \mathcal{O}p(\mathcal{U}_k)$ to agreeing with those of η_k^1 in $\mathcal{U}_k$. Denote by η_k^δ, $\delta \in [1/2, 1]$, the corresponding legendrian homotopy.

A remark is in order. The slopes of the fronts of η_k^δ, $\delta \in [1/2, 1]$, can be assumed to remain arbitrarily close to those of $\eta_k^{1/2}$; in particular, the deformation does not escape the Darboux ball M. In particular, we can find a bound, independent of N but depending on how much we want to C^0–approximate $\eta_k^{1/2}$, for how large the areas enclosed by the Reidemeister I loops can be. This implies that we can adjust N and the size of the loops to modify the area to be exactly the amount we require.

Since η_k^δ is legendrian for $\delta \in [1/2, 1]$, its tangent map extends G_k^δ to the whole of $\delta \in [0, 1]$. G_k^δ lifts to F_k^δ as above. We define ψ_k^1 (or, rather, its y–coordinate) by integrating zdx over η_k^1. Since the $\mathcal{U}_k$ cover $\mathbb{D}^m$, we have that for all $k \in \mathbb{D}^m$ this integral can be adjusted to ensure $\psi_k^1(1) = \psi_k^0(1)$. We define the y–coordinate of ψ_k^δ, $\delta \in (0, 1)$, by lifting it arbitrarily relative to $s = 0, 1$ and $\delta = 0, 1$. □

Bringing together Theorem 5.14 and Proposition 1.32, we obtain the following corollary:

Corollary 5.20. *Let $\mathcal{D}$ be a C^∞–generic Engel structure. Then, the inclusion*

$$\pi_0(\mathcal{HI}(\mathbb{S}^1, \mathcal{D})) \to \pi_0(\mathcal{FHI}(\mathbb{S}^1, \mathcal{D}))$$

is a bijection.

Corollary 5.20 should still be true for higher π_k. This would require carefully analysing families of curves and ensuring that the model from Figure 5.2 can be introduced parametrically.

Similarly, we are able to extend the Adachi–Geiges result to all Engel manifolds:

Corollary 5.21. *Let $(M, \mathcal{D})$ be an Engel manifold and let $\gamma_1, \gamma_2 \in \mathcal{HI}^{\text{n.e.t.}}(\mathcal{D})$ be two horizontal loops. Then, they are isotopic as horizontal loops if and only if they are homotopic as maps and they have the same rotation number.*

Proof. Apply Theorem 5.14 to obtain a connecting family of immersions. One can then proceed in a cover by Darboux charts, much like in Theorem 5.12, in which intersection points, under the Geiges projection, appear as self–tangencies satisfying an area condition. Generically, curves with self–tangencies can be assumed to be isolated in a 1–parametric family. By adding or subtracting area around said points, they can be assumed not to lift to intersections. □

5.3 The h–principle for transverse maps and immersions

Having proven our results on horizontal immersions, we can study the other condition that is geometrically meaningful for a map to satisfy in the presence of a distribution: that of being transverse. We shall review Gromov's strategy for proving flexibility. This was already worked out in detail by Y. Eliashberg and N. Mishachev in [20, p. 136] for the contact case, and indeed the proof goes through without any major differences.

Theorem 5.22. *Let $(M, \mathcal{D})$ be an Engel manifold. The following statements hold:*

a. Let V be a manifold. Maps $f : V \to M$ with $df : TV \to TM \to TM/\mathcal{D}$ surjective satisfy a C^0–close, parametric, relative, and relative to the parameter h–principle.

b. Let V be a manifold with $\dim(V) < \dim(M)$. Then, immersions of V into M transverse to $\mathcal{D}$ satisfy a C^0–close, parametric, relative, and relative to the parameter h–principle.

Remark 5.23. If V has subcritical dimension 1, it was already proven in [20][Prop. 8.3.2] that transverse immersions $V \to M$ satisfy the h–principle. For this, no assumptions on the distribution are needed. However, in dimensions 2 and 3 the assumption on Engelness is geometrically essential. More precisely, *Gromov's conjecture* states that the same statement should hold whenever the distribution is bracket–generating (and it clearly does not if the distribution is, for instance, a foliation).

Remark 5.24. Assume that the Engel flag $\mathcal{W} \subset \mathcal{D} \subset \mathcal{E} \subset TM$ is orientable. Then, if V is an immersed closed transverse 2–dimensional manifold, it must be a torus with trivial normal bundle. If we drop the orientability assumption, V can be a Klein bottle as well.

Remark 5.25. The two PDRs we are considering have the following formal analogues. In the first case, the formal data is a pair (f, F) where f is a mapping of V into M and $F : TV \to TM$ is a bundle morphism that is surjective onto the quotient $TM/\mathcal{D}$. In the second case we have the same, but we additionally ask for F to be a monomorphism.

5.3.1 The h–principle for Diff–invariant, microflexible and locally integrable relations

Let us explain the main ingredients needed to prove Theorem 5.22. The interested reader might want to refer to [20][Chap. 13].

Fix two manifolds W and M and let $\pi : J^r(W, M) \to W$ be the space of r–jets from W to M. r can be any integer, but can also take the value ∞ or the value g, by which we mean germs of maps. A subset $\mathcal{R} \subset J^r(W, M)$, as seen in the Preamble, is a partial differential relation (PDR). Given some function $f : W \to M$, we write $J^r \circ f : W \to J^r(W, M)$ for its r–jet expansion.

Definition 5.26. *A differential relation $\mathcal{R}$ is* ***locally integrable*** *if, for any m, and for any two maps*

$$h : [0, 1]^m \to \mathcal{R} \subset J^r(W, M)$$
$$g_p : \mathcal{O}p(\pi \circ h(p)) \to M,\ p \in \mathcal{O}p(\partial[0, 1]^m)$$

satisfying $(J^r \circ g_p)(\pi \circ h(p)) = h(p)$ and $J^r \circ g_p \subset \mathcal{R}$, there is

$$f_p : \mathcal{O}p(\pi \circ h(p)) \to M,\ p \in [0, 1]^m$$

satisfying $(J^r \circ f_p)(\pi \circ h(p)) = h(p)$ for all p, $f_p = g_p$ for all $p \in \mathcal{O}p(\partial[0, 1]^m)$, and $J^r \circ f_p \subset \mathcal{R}$.

What the definition states is that $\mathcal{R}$ is locally integrable if any r–jet at a point satisfying $\mathcal{R}$ (in this case, the family h) can be locally extended to a solution (the family f). We introduce the parameter space $[0, 1]^m$ to state that this local solvability holds parametrically and relatively as you vary the pointwise condition.

Let us denote $\theta_l = (A = [-1, 1]^n, B = \partial([-1, 1]^n) \cup ([-1, 1]^l \times \{0\}))$.

Definition 5.27. *A relation $\mathcal{R}$ is **microflexible** if, for any small ball $U \subset W$, any m and l, and any maps*

$$h_p : \theta_l \to U,\ p \in [0,1]^m,\ \textit{embeddings},$$

$$F_p : \mathcal{O}p(h_p(A)) \to \mathcal{R}\ \textit{holonomic},$$

$\tilde{F}_p^t : \mathcal{O}p(h_p(B)) \to \mathcal{R}$, $t \in [0,1]$, *holonomic and satisfying* $\tilde{F}_p^t = F_p$ *for* $p \in \mathcal{O}p(\partial[0,1]^m)$ *or* $t = 0$, *there is, for small t, a holonomic family $F_p^t : \mathcal{O}p(h_p(A)) \to \mathcal{R}$ satisfying $F_p^t = F_p$ if $p \in \mathcal{O}p(\partial[0,1]^m)$ or $t = 0$, and satisfying $F_p^t = \tilde{F}_p^t$ in $\mathcal{O}p(B)$. If the F_p^t exists for all $t \in [0,1]$, we say that $\mathcal{R}$ is **flexible**.*

That is, being microflexible amounts to proving that local deformations of a solution of the differential relation can be extended to global solutions, as least for small times. Relations that are open are immediately microflexible and locally integrable.

The following proposition [20, p. 13.5.3] holds:

Proposition 5.28 (Gromov)**.** *Let $\mathcal{R} \subset J^r(V \times \mathbb{R}, M)$ be a locally integrable and microflexible relation that is invariant with respect to diffeomorphisms that leafwise preserve the foliation $\coprod\{v\} \times \mathbb{R}$. Then, a C^0–close, parametric, relative, and relative to the parameter h–principle holds in $\mathcal{O}p(V \times \{0\})$.*

Saying that the h–principle holds means that the space of holonomic sections (sections such that the formal derivatives are the actual derivatives of the zeroth order map) is weak homotopy equivalent –by the inclusion– to the space of all sections into $\mathcal{R}$. Note that, by C^0–close it is meant that the zeroth order components are C^0–close, not its derivatives.

5.3.2 The proof of Theorem 5.22 (a.)

Let $(M, \mathcal{D})$ be an Engel manifold. We claim that the relation $\mathcal{R}_1$ in $\pi : J^1(\mathbb{R}, M) \to \mathbb{R}$ of being tangent to $\mathcal{D}$ but transverse to $\mathcal{W}$ is locally integrable. Suppose we are given maps

$$h : [0,1]^m \to (\mathcal{D} \setminus \mathcal{W}) \subset TM$$

$$g_p : \mathcal{O}p(0) \subset \mathbb{R} \to M,\ p \in \mathcal{O}p(\partial[0,1]^m)$$

where the g_p are horizontal curves transverse to $\mathcal{W}$ satisfying $dg_p(0) = h(p)$. For all $p \in [0,1]^m$ and depending smoothly on p, we can extend the vector $h(p)$ to a vector field H_p in $\mathcal{O}p(\pi \circ h(p))$. We can assume that the maps g_p are embeddings by shrinking the domain. Therefore, for those $p \in \mathcal{O}p(\partial[0,1]^m)$, H_p can be assumed to be an extension of the tangent vector g_p'. Following the flow of H_p for short times gives the desired local extension of g_p.

We claim that $\mathcal{R}_1$ is microflexible as well. Observe that we only have to consider the case θ_0, which can be phrased as follows. Let $F_p^0 : [0,1] \to \mathcal{R}_1$, $p \in [0,1]^m$, be a family of holonomic maps. Let $F_p^t : \mathcal{O}p(\{0,1\}) \to \mathcal{R}_1$, $t \in [0,1]$, be a family of deformations defined around the endpoints of the interval. Let $\psi : [0,1] \to \mathbb{R}$ be a bump function which is identically 1 around $\{0,1\}$ and zero in an arbitrarily large interval in the interior of $[0,1]$. According to Lemma 5.1, the curves F_p^0 possess a local model in which they correspond to the zero section in $J^2(\mathbb{R}, \mathbb{R})$; this implies that, for small t, F_p^t is graphical over F_p^0 and therefore given by a function y_p^0. The extension is given by ψy_p^0 and its derivatives.

Let V be some manifold. Let $\mathcal{R}_2 \subset J^1(V, M)$ be the open relation of having the formal derivative be surjective onto $TM/\mathcal{D}$. The relation $\mathcal{R}_3 \subset J^1(V \times \mathbb{R}, M)$ consists of those maps with formal derivative surjective onto $TM/\mathcal{D}$ that, along the fibres $\{v\} \times \mathbb{R}$, are tangent to $\mathcal{D}$ but transverse to $\mathcal{W}$. Local integrability for $\mathcal{R}_3$ follows by mimicking the argument for $\mathcal{R}_1$.

We claim that $\mathcal{R}_3$ is also microflexible. Take $\theta_j = (A, B)$. Suppose we are given a holonomic family F_p^0 on A and a corresponding deformation F_p^t over $\mathcal{O}p(B)$. Find neighbourhoods $\mathcal{O}p_1(B) \subset \mathcal{O}p_2(B) \subset \mathcal{O}p(B)$ and build a bump function ψ that is 1 in $\mathcal{O}p_1(B)$ and 0 outside of $\mathcal{O}p_2(B)$. Since F_p^t is fibrewise graphical over F_p^0 for small t, we use ψ to interpolate back to F_p^0, as above; this can be achieved even if B is embedded wildly with respect to the foliation $\coprod\{v\} \times \mathbb{R}$. For small times the resulting deformation is C^∞–close to F_p^0, so in particular it is still surjective onto $TM/\mathcal{D}$ in the transverse direction.

By construction, $\mathcal{R}_3$ is invariant under diffeomorphisms preserving the foliation $\coprod\{v\} \times \mathbb{R}$ leaf-wise. Then, Proposition 5.28 allows us to conclude that in $\mathcal{O}p(V \times \{0\})$ a complete h–principle holds, so in particular a complete h–principle holds in V for the relation $\mathcal{R}_2$. □

Remark: Observe that we did not need the h–principle for tangent immersions of Theorem 5.14, instead we just checked the much more simple properties of being microflexible and locally integrable for the relation $\mathcal{R}_1$. However, we will need it in the next subsection.

5.3.3 The proof of Theorem 5.22 (b.)

For statement (b.), note that the case $\dim(V) = 2$ has already been covered by part (a.). This leaves only the case $\dim(V) = 3$. Inspecting the proof presented in the previous subsection, it is clear that it cannot possibly go through, since immersions $V^3 \times \mathbb{R} \to M$ cannot avoid the $\mathcal{W}$–direction, which was a key ingredient in the 2–dimensional case to obtain microflexibility.

Here is where the space of generic horizontal immersions $\mathcal{HI}^{\text{gen}}(\mathbb{S}^1, \mathcal{D})$ that we introduced in Definition 5.11 comes into play. Define the following differential relation $\mathcal{R} \subset J^g(V \times \mathbb{R}, M)$: germs that are transverse to $\mathcal{D}$ along $V \times \{s\}$ and lie in $\mathcal{HI}^{\text{gen}}(\mathbb{R}, \mathcal{D})$ along $\{v\} \times \mathbb{R}$. There exists an obvious projection $J^g(V \times \mathbb{R}, M) \to J^1(V \times \mathbb{R}, M)$ and the image of $\mathcal{R}$ is the relation $\mathcal{R}^1$: maps with formal differential transverse to $\mathcal{D}$ along $V \times \{s\}$ and tangent to $\mathcal{D}$ along $\{v\} \times \mathbb{R}$. $\mathcal{R} \to \mathcal{R}^1$ is a Serre fibration with contractible fibre.

The proof amounts to showing that $\mathcal{R}$ is microflexible and locally integrable: assuming this Proposition 5.28 allows us to conclude. The full h–principle for $\mathcal{HI}^{\text{gen}}(\mathcal{D})$ and the openness of the transverse immersion condition in codimension–1 imply microflexibility. Local integrability is tautological. The claim follows. □

Closing remarks

This chapter/appendix is the end of the thesis. Thank you for making it this far. Now I will go back to certain parts of the thesis that deserve further comment and simultaneously I will discuss directions of future/current research.

Novikov's theorem for symplectic foliations

This is in regard to Part I.

Novikov's remarkable result is the starting point of many of the results on the rigid behaviour of taut foliations. As such, we can pose the question of what a suitable generalisation in higher dimensions is (in the context of strong symplectic foliations). Naively, we could simply ask for π_1 of the leaves to inject in π_1 of the ambient space. If we do so, it is easy to see that the dimensional reduction provided by Theorem 3.5 (the Lefschetz hyperplane theorem for leaves) reduces the proof of this statement to proving the 5–dimensional case. Critically, the theorem does not allow us to go down to the 3–dimensional case (where we know it holds).

However, it might happen that this is not the right generalisation. An alternate way of thinking about Novikov's theorem is as a statement about the non–existence of lagrangian spheres that are vanishing cycles. Is this true in higher dimensions?

Flexibility for Engel structures

All the remaining sections refer to Part II.

As we have pointed out more than once, the main open question regarding Engel structures is whether they display complete flexibility. That is to say, whether the class of all Engel structures satisfies a complete h–principle. I would not dare conjecture such a thing: it might happen that this result is within our reach improving on the techniques we currently have, or it might happen that it is actually not even true.

Even if the latter holds, some work (that is currently in progress) is needed to understand better the nature of looseness. How can we improve on the results of Chapter 4? One of the first issues to deal with is that looseness still depends on the dimension of the family under consideration. This is a consequence of the fact that, every time we deform the Engel structure in a flowbox with winding 3, we have to avoid it in further homotopies (because the argument relies on applying Little's homotopy over and over). It is conceivable that replacing this step by a more complicated homotopy (Saldanha style) will avoid this problem. In general, the question we should understand is the following: can Saldanha's h–principle be extended to an h–principle *relative in the domain* (for a suitable subclass of convex curves)?

Once that is settled, there is a more complicated question to be understood. Can the overtwisted disc be localised? What we mean by this is the following: in the contact case, overtwistedness is identified by displaying a ball with an overtwisted disc. In the Engel case looseness has to be checked orbit by orbit (of some line field contained in the Engel structure). It would certainly be desirable if this was replaced by a more semi–local condition. At the same time, maybe the condition of being loose with respect to $\mathcal{W}$ (Corollary 4.9) can always be achieved by a deformation. We know that generically $\mathcal{W}$ has only isolated orbits. Is it possible to make them have long developing map? Can it be achieved for the open orbits to have long developing map as well?

The particular case of Engel cobordisms

Corollary 3.2 stated that two (framed) contact structures that are formally Engel cobordant are Engel cobordant. We pointed out that this indicates that there might not be a deep interaction between contact and Engel topology as soon as we try to make global statements. However, all the cobordisms we produce are necessarily "flexible". It is possible that if one restricts to non–loose cobordisms (whatever they are) then the Engel cobordism relation is actually meaningful. It is worth pointing out that two homotopic contact structures are isotopic by Gray stability and hence there are trivial Engel cobordisms with arbitrarily short developing map connecting them.

Flexibility/rigidity for tangent/transverse submanifolds

In Chapter 5 we left a little question open: is the inclusion $\mathcal{HI}(\mathbb{S}^1, \mathcal{D}) \to \mathcal{FHI}(\mathbb{S}^1, \mathcal{D})$ a weak homotopy equivalence if $\mathcal{D}$ is generic? We showed that the answer was yes in π_0. We believe that we have provided all the tools needed to answer this, and simply some additional effort to do the technical work is needed. It would be nice to settle this for completeness.

A more pressing question (or series of questions) is to start researching *embedded* submanifolds. This is largely unexplored as of now. Let us start with a remark: given $\mathcal{D}$ loose and N a 3–submanifold, $\mathcal{D}$ is Engel homotopic to another structure that has N as a submanifold transverse to the kernel and endowed with any contact structure we want (as long as there are no formal obstructions). This is a consequence of our h–principle for loose structures and Corollary 3.2. In particular, whether an Engel manifold contains a given contact manifold as a hypersurface is not an Engel invariant.

This construction showcases the usual problem: since homotopic Engel structures are not isotopic, it seems hard to construct invariants under Engel homotopy. We are therefore interested in homotopies of pairs: the structure and the submanifold. For 3–manifolds transverse to the kernel there is no homotopy lifting property (given a path of structures and a initial submanifold, we cannot extend it to a path of pairs in general). Is this the same true for surfaces and 3–manifolds transverse to $\mathcal{D}$? What about (higher dimensional families of) horizontal knots?

A different problem is the study of submanifolds inside a *fixed* Engel structure. Some work in progress of R. Casals, F. Presas, and myself deals with loops of horizontal knots in standard $\mathbb{R}^4$. At the moment, we are working on a Thurston–Bennequin type invariant which we conjecture recovers the formal type completely. This has to be understood, but what is truly interesting is whether a non–formal invariant exists. Among experts, there is a certain expectation/hope that a higher order analogue of a generating function might provide the first instance of a non–formal invariant. At the moment it is unclear how such an analogue should be defined.

Other classes of distributions

I justified our interest in Engel structures arguing that they conform the simplest class of bracket–generating distributions after 3–dimensional contact structures. Having done that, we may consider other classes of distributions. At the moment very little is known in this direction, and we have no expectations (in either direction) regarding whether most classes of distributions outside of the scope of convex integration are hiding an interesting topological theory.

After Engel structures, the next interesting example are the $(2,3,5)$–**distributions** introduced by Cartan. The (fairly descriptive) name means the following: they are 2–distributions in dimension 5 that first generate a 3–distribution by Lie bracket and then the whole tangent space. They are given by an open Diff–invariant relation and hence the h–principle holds for them in open manifolds. Whether the same holds in closed manifolds is unknown. A recent paper of Dave and Haller [13] discusses some constructions of $(2,3,5)$–distributions.

Another class of distributions that is more closely related to Engel structures (probably?) are the **Goursat structures**. They are 2–distributions that generate a complete flag by iterated Lie bracket. As such, from dimension 5 onwards they are given by a partial differential relation that is actually *closed*. It is easy to write a characterisation lemma for these structures (much like in Proposition 1.12). It is possible that techniques similar to the ones presented in this thesis provide an h–principle for them. However, this is certainly not immediate and constitutes an interesting direction of work.

Horizontal curves in bracket–generating distributions

In all these cases one can consider the problem of determining the homotopy type of the space of horizontal curves. In particular, does the h–principle hold once one removes the rigid ones? [1] This question can be considered for any type of bracket–generating distribution and is quite deep. Essentially, it amounts to proving a form of 1–*dimensional convex integration for bracket–generating relations* (relations that might be closed!).

The original proof of Theorem 5.14 actually relied on a simplified form of this more general (but still conjectural) result. The key insight is using the Lorentz model of Subsubsection 1.3.1.1, where being horizontal but transverse to $\mathcal{W}$ is seen as a bracket–generating relation in a particularly convenient fashion. Proving the more general statement is a very interesting open problem.

[1] Somehow separately, what can we say about the subspace of rigid curves?

Bibliography

[1] J. Adachi. “Classification of horizontal loops in standard Engel space”. *Int. Math. Res. Not.* (2007), pp. 29–40.

[2] G. Bande, P. Ghiggini, and D. Kotschick. “Stability theorems for symplectic and contact pairs”. *Int. Math. Res. Not.* 68 (2004), pp. 3673–3688.

[3] D. Bennequin. “Entrelacements et équations de Pfaff”. *Asterisque* 107–108 (1983), pp. 87–161.

[4] M. Bertelson. “A h–principle for open relations invariant under foliated isotopies”. *J. Symplectic Geom.* 2.1 (2002), pp. 369–425.

[5] E. Bishop. “Differentiable manifolds in complex euclidean space”. *Duke Math. J.* 32.1 (1965), pp. 1–21.

[6] M.S. Borman, Y. Eliashberg, and E. Murphy. “Existence and classification of overtwisted contact structures in all dimensions”. *Acta Math.* 215.2 (2015), pp. 281–361.

[7] R.L. Bryant and L. Hsu. “Rigidity of integral curves of rank 2 distributions”. *Invent. Math.* 114.2 (1993), pp. 435–461.

[8] D. Burghelea, N. Saldanha, and C. Tomei. “Results on infinite dimensional topology and applications to the structure of the critical set of non–linear Sturm-Liouville operators”. *J. Diff. Eq.* 188 (2003), pp. 569–590.

[9] R. Casals, A. del Pino, and F Presas. “h-Principle for Contact Foliations”. *Int. Math. Res. Not.* 20 (2015), pp. 10176–10207.

[10] R. Casals, J.L. Pérez, A. del Pino, and F Presas. “Existence h-Principle for Engel structures”. *arXiv e-prints* (2015). arXiv: 1507.05342.

[11] W.L. Chow. “Über Systeme von linearen partiellen Differentialgleichungen erster Ordnung”. *Math. Ann.* 117 (1939), pp. 98–105.

[12] M. Crainic and R.L. Fernandes. “Integrability of Poisson brackets”. *J. Diff. Geom.* 66 (2004), pp. 71–137.

[13] S. Dave and S. Haller. “On 5–manifolds admiting rank–2 distributions of Cartan type”. *arXiv e-prints* (2016). arXiv: 1603.09700v2.

[14] A. del Pino. “On the classification of prolongations up to Engel homotopy”. *Submitted* (2016).

[15] A. del Pino and F. Presas. “Flexibility for tangent and transverse immersions in Engel manifolds”. *arXiv e-prints* (2016). arXiv: 1609.09306.

[16] A. del Pino and F. Presas. “The foliated Weinstein conjecture”. *To appear in Int. Math. Res. Not.* (2015). arXiv: 1509.05268.

[17] S.K. Donaldson. “Symplectic submanifolds and almost-complex geometry”. *J. Differential Geom.* 44.4 (1996), pp. 666–705.

[18] Y. Eliashberg. “Classification of contact structures in $\mathbb{R}^3$”. *Int. Mat. Res. Not.* 3 (1993), pp. 87–91.

[19] Y. Eliashberg. “Classification of overtwisted contact structures on 3-manifolds”. *Invent. Math.* 98.3 (1989), pp. 623–637.

[20] Y. Eliashberg and N. Mishachev. *Introduction to the h-principle*. Vol. 48. Graduate Studies in Mathematics. American Mathematical Society, Providence, RI, 2002, pp. xviii+206.

[21] Y. Eliashberg and D. Thurston. *Confoliations*. Vol. 13. University lecture series. Providence, RI: American Mathematical Society, 1991, pp. xiii+66.

[22] R.L. Fernandes and P. Frejlich. “An h-Principle for Symplectic Foliations”. *Int. Math. Res. Not.* 7 (2012), pp. 1505–1518.

[23] S.C. Ferry and A.G Wasserman. “Morse theory for codimension–one foliations”. *Trans. Amer. Math. Soc.* 298.1 (1986), pp. 227–240.

[24] D. Gabai. “Foliations and the topology of 3-manifolds”. *J. Diff. Geom.* 18 (1983), pp. 445–503.

[25] D. Gabai. “Foliations and the topology of 3-manifolds II”. *J. Diff. Geom.* 26 (1987), pp. 461–478.

[26] D. Gabai. “Foliations and the topology of 3-manifolds II”. *J. Diff. Geom.* 26 (1987), pp. 479–536.

[27] H. Geiges. *An Introduction to Contact Topology*. Vol. 109. Cambr. Studies in Adv. Math. Cambr. Univ. Press, 2008.

[28] H. Geiges. “Horizontal loops in Engel space”. *Math. Ann.* 342.2 (2008), pp. 291–296.

[29] H. Geiges. “Review of [60]” (2001). MathSciNet: `2001h:53127`.

[30] E. Ghys. “Construction de champs de vecteurs sans orbite périodique”. *Séminaire Bourbaki* 36 (1993–1994), pp. 283–307.

[31] A. Golubev. “On the global stability of maximally nonholonomic two–plane fields in four dimensions”. *Int. Math. Res. Not.* 11 (1997), pp. 523–529.

[32] M. Gromov. *Partial differential relations*. Vol. 9. Ergebnisse der Mathematik und ihrer Grenzgebiete (3) [Results in Mathematics and Related Areas (3)]. Springer-Verlag, Berlin, 1986, pp. x+363.

[33] M. Gromov. “Pseudoholomorphic curves in symplectic manifolds”. *Invent. Math.* 82.2 (1985), pp. 307–347.

[34] V. Guillemin, E. Miranda, and A.R. Pires. “Codimension one symplectic foliations and regular Poisson structures”. *Bull. Braz. Math. Soc.* 42.4 (2011), pp. 607–623.

[35] G. Hector. “Une nouvelle obstruction à l’intégrabilité des variétés de Poisson régulières”. *Hokkaido Math. J.* 21.1 (1992), pp. 159–185.

[36] G. Hector and U. Hirsch. *Introduction to the Geometry of Foliations, Part A*. Friedrich Vieweg & Sohn, 1986.

[37] G. Hector, E. Macías, and M. Saralegi. “Lemme de Moser feuilleté et classification des variétés de Poisson régulières”. *Publ. Mat.* 33.3 (1989), pp. 423–430.

[38] H. Hofer. “Pseudoholomorphic curves in symplectizations with applications to the Weinstein conjecture in dimension three”. *Invent. math.* 114 (1993), pp. 515–563.

[39] L. Hsu. “Calculus of variations via the Griffiths formalism”. *J. Diff. Geom.* 36.3 (1992), pp. 551–589.

[40] S. Hurder and A. Rechtman. “The dynamics of generic Kuperberg flows”. *arXiv e-prints* (2013). arXiv: `1306.5025`.

[41] A. Ibort, D. Martínez–Torres, and F. Presas. "On the construction of contact submanifolds with prescribed topology". *J. Diff. Geom.* 56.2 (2000), pp. 235–283.

[42] A. Ibort and D. Martínez Torres. "A new construction of Poisson manifolds". *J. Symp. Geom.* 2.1 (2003), pp. 083–107.

[43] A. Ibort and D. Martínez Torres. "Approximately holomorphic geometry and estimated transversality on 2-calibrated manifolds". *C. R. Math. Acad. Sci. Paris* 338.9 (2004), pp. 709–712.

[44] H. Imanishi. "Structure of codimension one foliations which are almost without holonomy". *J. Math. Kyoto Univ.* 16 (1976), pp. 93–99.

[45] H. Imanishi and K. Yagi. "On Reeb components". *J. Math. Kyoto Univ.* 16 (1976), pp. 313–324.

[46] M. Klukas and B. Sahamie. "On prolongations of contact manifolds". *Proc. Amer. Math. Soc.* 141.9 (2013), pp. 3257–3263.

[47] D. Kotschick and T. Vogel. "Engel structures and weakly hyperbolic flows on four-manifolds". *arXiv e–prints* (2016). arXiv: 1610.04001.

[48] K. Kuperberg. "A smooth counterexample to the Seifert conjecture". *Ann. of Math.* 140 (1994), pp. 723–732.

[49] E. Lerman. "Contact fiber bundles". *J. Geom. Phys.* 49.1 (2004), pp. 52–66.

[50] J.A. Little. "Nondegenerate homotopies of curves on the unit 2–sphere". *J. Diff. Geom.* 4.3 (1970), pp. 339–348.

[51] D. Martínez Torres. "A higher dimensional generalization of taut foliations". *arXiv e-prints* (2006). arXiv: 0602576v1.

[52] D. Martínez Torres. "The geometry of 2-calibrated manifolds". *Port. Math.* 66.4 (2009), pp. 427–512.

[53] D. Martínez Torres, A. del Pino, and F. Presas. "Transverse geometry of foliations calibrated by non-degenerate closed 2-forms". *To appear in Nagoya Mathematical Journal* (2014). arXiv: 1410.3043.

[54] P. Massot and K. Niederkrüger. "Examples of non-trivial contact mapping classes in all dimensions". *Int. Math. Res. Not.* 15 (2016), pp. 4784–4806.

[55] S. Matsumoto. "Kuperberg's C^∞ counterexample to the Seifert conjecture". *Sugaku* 47 (1995), pp. 38–45.

[56] D. McDuff and D. Salamon. *Introduction to symplectic topology*. Second. Oxford Mathematical Monographs. The Clarendon Press, Oxford University Press, New York, 1998, pp. x+486.

[57] G. Meigniez. "Regularization and minimization of Haefliger structures of codimension one". *arXiv e-prints* (2009). arXiv: 0904.2912v5.

[58] J. Milnor. *Morse theory*. Based on lecture notes by M. Spivak and R. Wells. Annals of Mathematics Studies, No. 51. Princeton University Press, Princeton, N.J., 1963, pp. vi+153.

[59] Y. Mitsumatsu. "Leafwise Symplectic Structures on Lawson's Foliation". *arXiv e-prints* (2011). arXiv: 1101.2319v4.

[60] R. Montgomery. "Engel deformations and contact structures". Amer. Math. Soc. Transl. Ser. 196 (1999), pp. 103–117.

[61] K. Niederkrüger. "The plastikstufe: A generalization of the overtwisted disk to higher dimensions". *Algebr. Geom. Topol.* 6 (2006), pp. 2473–2508.

[62] K. Niederkrüger and A. Rechtman. "The Weinstein conjecture in the presence of submanifolds having a Legendrian foliation". *J. Topol. Anal.* 3.4 (2011), pp. 405–421.

[63] T. Nishimori. "Compact leaves with abelian holonomy". *Tohoku Math. J.* 27 (1975), pp. 259–272.

[64] B. Osorno. *Codimension–one symplectic foliations: constructions and examples.* PhD Thesis Utrecht University. 2015, pp. xiii+139.

[65] F. Pasquotto, J.B. van den Berg, and R.C. Vandervorst. "Closed characteristics on non–compact energy surfaces in $\mathbb{R}^{2n}$". *Math. Ann.* 343.2 (2009), pp. 247–284.

[66] F. Pasquotto, T. Rot, J.B. van den Berg, and R.C. Vandervorst. "Closed characteristics on non-compact mechanical contact manifolds". *To appear in J. Symplectic Geom.* (2013). arXiv: 1303.6461.

[67] M. M. Peixoto. "On an Approximation Theorem of Kupka and Smale". *J. Diff. Eq.* 3 (1966), pp. 214–227.

[68] D. Peralta–Salas, A. del Pino, and F. Presas. "Foliated vector fields without periodic orbits". *To appear in Israel Journal of Mathematics* (2014). arXiv: 1412.0123.

[69] F. Presas and L.E. Solá Conde. "Holomorphic Engel structures". *Rev. Mat. Complut.* 27.1 (2014), pp. 327–344.

[70] N. Saldanha. "The homotopy type of spaces of locally convex curves in the sphere". *Geom. Topol.* 19 (2015), pp. 1155–1203.

[71] P.A. Schweitzer. "Counterexamples to the Seifert conjecture and opening closed leaves of foliations". *Ann. of Math.* 100 (1974), pp. 386–400.

[72] P.A. Schweitzer. "Riemannian manifolds not quasi-isometric to leaves in codimension one foliations". *Ann. Inst. Fourier* 61 (2011), pp. 1599–1631.

[73] H. Seifert. "Closed integral curves in 3–space and isotopic two-dimensional deformations". *Proc. Amer. Math. Soc.* 1 (1950), pp. 287–302.

[74] S. Suhr and K. Zehmisch. "Linking and closed orbits." *Abh. Math. Semin. Univ. Hambg.* 86.1 (2016), pp. 133–150.

[75] D. Sullivan. "A homological characterization of foliations consisting of minimal surfaces". *Comment. Math. Helv.* 54 (1979), pp. 218–223.

[76] D. Sullivan. "Cycles for the dynamical study of foliated manifolds and complex manifolds". *Invent. Math.* 36 (1976), pp. 225–255.

[77] C.H. Taubes. "The Seiberg–Witten equations and the Weinstein conjecture". *Geometry & Topology* 11 (2007), pp. 2117–2202.

[78] R. Thom. "L'homologie des espaces fonctionnels". *Colloque de topologie algébrique, Louvain* (1956), pp. 29–39.

[79] W.P. Thurston. "Existence of codimension-one foliations". *Ann. Math.* 104.2 (1976), pp. 249–268.

[80] W.P. Thurston. "The theory of foliations of codimension greater than one". *Comm. Math. Helv.* 87 (1974), pp. 214–231.

[81] I. Vaisman. *Lectures on the geometry of Poisson manifolds.* Vol. 118. Progress in Mathematics. Birkhäuser Verlag, 1994.

[82] T. Vogel. "Existence of Engel structures". *Ann. Math.* 169.1 (2009), pp. 79–137.

[83] T. Vogel. "Non-loose unknots, overtwisted discs and the contact mapping class group of $\mathbb{S}^3$". *arXiv e–prints* (2016). arXiv: 1612.06557.

[84] T. Vogel. "On the uniqueness of the contact structure approximating a foliation". *To appear in Geom. Topol.* (2016).

[85] A. Weinstein. "On the hypotheses of Rabinowitz' periodic orbit theorems". *J. Diff. Eq.* 33.3 (1979), pp. 353–358.

[86] H. Whitney. "On Regular Closed Curves in the Plane". *Compos. Math.* 4 (1937), pp. 276–284.

[87] F.W. Wilson Jr. "On the minimal sets of non-singular vector fields". *Ann. of Math.* 84 (1966), pp. 529–536.