Symplectic 4-Dimensional 2-Handles and Contact Surgery along Transverse Knots

by

David Thomas Gay

A.B. (Harvard University) 1991

A dissertation submitted in partial satisfaction of the
requirements for the degree of
Doctor of Philosophy

in

Mathematics

in the

GRADUATE DIVISION
of the
UNIVERSITY of CALIFORNIA at BERKELEY

Committee in charge:

Professor Robion C. Kirby, Chair
Professor Andrew J. Casson
Professor Hiroshi Oguri

Spring 1999

The dissertation of David Thomas Gay is approved:

Chair Date

Date

Date

University of California at Berkeley

Spring 1999

Symplectic 4-Dimensional 2-Handles and Contact Surgery along Transverse Knots

Abstract

Symplectic 4-Dimensional 2-Handles and Contact Surgery along Transverse Knots

by

David Thomas Gay

Doctor of Philosophy in Mathematics

University of California at Berkeley

Professor Robion C. Kirby, Chair

A standard convexity condition on the boundary of a symplectic manifold involves an induced positive contact form (and contact structure) on the boundary; the corresponding concavity condition involves an induced negative contact form. We present two methods of symplectically attaching 2-handles to convex boundaries of symplectic 4-manifolds along links *transverse* to the induced contact structures. One method results in a weaker convexity condition on the new boundary, with an associated positive contact structure constructed by a contact surgery along the transverse link; the weaker convexity is sufficient to imply tightness of the new contact structure. These handles can be attached with any framing *less than* some upper bound. The second method results in *concave* boundaries and depends upon a fibration of the link complement over S^1; here the handles can be attached with any framing *larger than* some lower bound (determined by the fibration). Along the way we use a notion of symplectic 4-manifold boundaries which are partially convex and partially concave, so a theory is developed to make this notion rigorous and useful.

Professor Robion C. Kirby
Dissertation Committee Chair

To Juliette Blanca Benitez

and my parents.

Contents

List of Figures

Acknowledgements

I would like to thank Rob Kirby for teaching me the right way to think about topology, Andy Lewis for listening to my weird ideas and actually understanding them and Aaron Abrams for providing perspective. I also owe a great debt to the numerous people with whom I have had useful mathematical conversations around these topics, especially John Etnyre, Danny Calegari, Alan Weinstein, Ana Cannas da Silva, Yasha Eliashberg and Andrew Casson.

Chapter 1

Introduction

Existence and uniqueness questions for symplectic structures on 4-manifolds and contact structures on 3-manifolds are increasingly important in low-dimensional geometric topology. Here we present new "handle-by-handle" constructions of symplectic 4-manifolds with corresponding surgery constructions of contact 3-manifolds. In this introduction we will first review the basic definitions in symplectic and contact topology. Then we will motivate our results by reviewing some standard constructions. Finally we will outline the main new results in this paper.

We begin with a few conventions: All manifolds and all maps between manifolds are smooth. Unless explicitly stated otherwise, all manifolds and submanifolds are oriented and do not have boundaries. A vector field V in a manifold X which is transverse to a codimension 1 submanifold M is *positively transverse* if, for all $p \in M$, $V(p)$ followed by a positively oriented basis for T_pM is a positively oriented basis for T_pX. If X is a manifold with boundary, then ∂X is oriented so that a vector field which points out of X is positively transverse to ∂X. A codimension 0 submanifold of X always inherits its orientation from X. We use the notation $-X$ to refer to X with its orientation reversed.

1.1 Basic definitions

For a comprehensive introduction to symplectic topology the reader is referred to [10]. For excellent background on contact topology with connections to symplectic topology, especially in the low dimensions, the reader is referred to [5]. Here we present only the skeleton of an introduction.

A *symplectic form* on a $2n$-manifold X is a closed 2-form ω for which ω^n is nowhere zero. The standard example on $\mathbb{R}^{2n}$ is $\omega = dx_1 \wedge dy_1 + \ldots + dx_n \wedge dy_n$. Note that ω^n orients X; we will always assume this agrees with the given orientation on X. A *symplectomorphism* is a diffeomorphism ϕ from one symplectic manifold (X_1, ω_1) to another (X_2, ω_2) for which $\phi^*\omega_2 = \omega_1$. A *contact form* on a $(2n-1)$-manifold M is a 1-form α for which $\alpha \wedge (d\alpha)^{n-1}$ is nowhere zero. The standard example on $\mathbb{R}^{2n-1} = \mathbb{R} \times \mathbb{R}^{2(n-1)}$ is $\alpha = dt + x_1 dy_1 + \ldots + x_{n-1} dy_{n-1}$. A contact form is *positive* or *negative* according to whether the orientation given by $\alpha \wedge (d\alpha)^{n-1}$ agrees with or disagrees with the given orientation on M. The *Reeb vector field* for a contact form α, written R_α, is the unique vector field for which $\alpha(R_\alpha) = 1$ and $\imath_{R_\alpha} d\alpha = 0$.

A (co-orientable) contact *structure* on a $(2n-1)$-manifold is a hyperplane field ξ which is the kernel of a contact form. A contact structure is positive or negative according to the sign of the contact form. A co-oriented contact structure is a contact structure together with a transverse orientation. Given a contact form α, the plane field $\ker \alpha$ together with the transverse orientation given by α will be called the *underlying co-oriented contact structure* for α. Two contact forms α_1 and α_2 have the same underlying co-oriented contact structures if and only if $\alpha_1 = f\alpha_2$ for some positive function f.

A *contactomorphism* is a diffeomorphism ϕ from one contact manifold $(M_1, \xi_1 = \ker \alpha_1)$ to another $(M_2, \xi_2 = \ker \alpha_2)$ for which $\phi^*\xi_2 = \xi_1$ or, equivalently, $\phi^*\alpha_2 = f\alpha_1$ for some function f. A contactomorphism is co-oriented if the contact structures are co-oriented and the diffeomorphism preserves the co-orientation, or equivalently if f is positive.

We will generally suppress mention of co-orientations and assume that all contact structures and all contactomorphisms are co-oriented.

On a 3-manifold a contact structure ξ is a 2-plane field and there are two special classes of contact structures. We quote Eliashberg [4]: "The contact structure ξ is called *tight* if there is no embedded disk D such that its boundary ∂D is tangent to ξ while D is transversal to ξ along the boundary. Dichotomy of contact structures into tight and the complementary class of *overtwisted* plays an important role in 3-dimensional contact geometry."

A vector field V in a symplectic manifold (X, ω) is a *symplectic dilation* (or Liouville vector field) if $\mathcal{L}_V \omega = \omega$ and is a *symplectic contraction* if $\mathcal{L}_V \omega = -\omega$. For example, in $\mathbb{R}^{2n}$ with the standard symplectic form, $V = \frac{1}{2}(x_1 \partial_{x_1} + y_1 \partial_{y_1} + \ldots + x_n \partial_{x_n} + y_n \partial_{y_n})$ is a symplectic dilation. We will say that a boundary component M of a symplectic manifold

(X, ω) is *strongly convex* if there exists a symplectic dilation V defined in a neighborhood of and positively transverse to M. (Here we are deviating a little from standard terminology: Often in this case one says that M is a boundary component "of contact type", see [10].) Thus the boundary S^{2n-1} of $(B^{2n}, \omega) \subset (\mathbb{R}^{2n}, \omega)$ is strongly convex. We say M is *strongly concave* if instead V is a symplectic contraction. A 3-dimensional boundary component M of a symplectic 4-manifold (X, ω) is *weakly convex* if there exists a positive contact structure ξ on M such that $\omega|_\xi$ is nowhere zero. (Again this is not quite standard. In Eliashberg's terminology [4], when $M = \partial X$, we say that (X, ω) is a "symplectic filling" of (M, ξ).) In general, when we say "convex" or "concave" we will mean "strongly convex" or "strongly concave".

A knot K in a contact 3-manifold (M, ξ) is *transverse* if TK is transverse to ξ everywhere and is *Legendrian* if TK is tangent to ξ everywhere. By a *framing* of a submanifold we mean a homotopy class of framings of the normal bundle, and for a knot K this can be specified by a single vector field defined along K and transverse to K. Thus a Legendrian knot K has a canonical framing, known as the *Thurston-Bennequin framing* and written $\mathrm{tb}(K)$, given by any vector field in ξ along K transverse to TK.

When F is a framing of a link L and K is a component of L, we will use the notation F_K to refer to the framing of K induced by F.

If Σ is a surface in a contact 3-manifold (M, ξ) then $T\Sigma \cap \xi$ is generically a line field with isolated singularities. This integrates to give a 1-dimensional singular foliation on Σ called the *characteristic foliation* of Σ.

1.2 Background results

Here we present some standard results which will motivate the new results presented later.

Legendrian and transversal knots become useful because any knot can be made either Legendrian or transverse via a C^0-small isotopy (see [3]). Note that we need a canonical 0-framing to associate a framing of a knot with an integer, but that the difference between two framings of a knot is always associated with an integer. With this in mind, the Thurston-Bennequin framing of any Legendrian knot K can always be decreased in the sense that there exists a C^0-small isotopy taking K to another Legendrian knot K' such that $\mathrm{tb}(K) - \mathrm{tb}(K') = 1$ (see [7]). It is much harder in general to increase $\mathrm{tb}(K)$.

If M is a convex boundary component of (X,ω), then the symplectic dilation V induces a positive contact form $\alpha = \imath_V\omega|_M$ on M. If M is concave then α is a negative contact form. Suppose that M is a convex boundary component of (X,ω_X) with induced positive contact form α_M and that N is a concave boundary component of (Y,ω_Y) with induced negative contact form α_N. If ϕ is a diffeomorphism from $-N$ to M such that $\phi^*(\alpha_M) = \alpha_N$ then (X,ω_X) can be glued along ϕ to (Y,ω_Y) to form a new symplectic manifold. (This will be shown in section 3.2.) Thus one approach to building closed symplectic manifolds is to understand the construction of manifolds with convex and concave boundaries and to understand the induced contact forms.

If M is a (strongly) convex boundary component of a symplectic 4-manifold with induced positive contact form α, then M is also a weakly convex boundary component with respect to the positive contact structure $\xi = \ker\alpha$. If all the boundary components of (X,ω) are weakly convex, then the positive contact structure on each boundary component is tight (see [2]). Thus the construction of symplectic 4-manifolds with (strongly or weakly) convex boundaries is useful as a construction of tight contact structures on 3-manifolds.

In [11], Weinstein shows how to attach $2n$-dimensional symplectic k-handles, as long as $k \leq n$, along certain special $(k-1)$-spheres in a convex boundary of a symplectic $2n$-manifold, to build a new symplectic $2n$-manifold with convex boundary. (For details on smooth constructions using handles, see [9]. For background on handles in the construction of 4-manifolds, see [8]. We will also present some of the details in section 2.2.) In the case where $n = 2$, we can summarize Weinstein's results as follows: Suppose (X,ω) is a symplectic 4-manifold with strongly convex boundary. We consider two cases:

- Case 1: Given a pair of points $K = (p^+, p^-)$ in ∂X (i.e. an embedding of S^0), let $Y \supset X$ be a smooth manifold constructed by attaching a 1-handle along K.
- Case 2: Given a knot K in ∂X (an embedding of S^1) which is *Legendrian* with respect to the induced contact structure on ∂X, let $Y \supset X$ be a smooth manifold constructed by attaching a 2-handle along K with framing $\mathrm{tb}(K) - 1$.

Theorem 1.1 (Weinstein). *In both cases above, Y can be constructed so that ω extends across the handle and so that ∂Y is strongly convex. The handle can be attached in an arbitrarily small neighborhood of K.*

(This will be proved in section 3.3.) In particular, ∂Y carries an induced contact form which is obtained from the contact form on ∂X by a certain "contact surgery" along K. The convexity implies that these surgeries yield tight contact structures. Starting with B^4 with the standard symplectic form, we can then form many symplectic 4-manifolds with convex boundaries, and hence many 3-manifolds with tight contact structures. Gompf has extensively investigated the possibilities using this construction in [7].

A simple variation of Weinstein's construction yields the following result: Suppose (X, ω) is a symplectic 4-manifold with strongly concave boundary and $K \subset \partial X$ is a knot which is Legendrian with respect to the induced (negative) contact structure on ∂X. Let $Y \supset X$ be a smooth manifold constructed by attaching a 2-handle along K with framing $\mathrm{tb}(K) - 1$.

Theorem 1.2. *In this case Y can be constructed so that ω extends across the handle and so that ∂Y is strongly concave. Again the handle can be attached in an arbitrarily small neighborhood of K.*

1.3 New results

The new results presented in this paper require some control on the behaviour of contact forms in neighborhoods of knots. Given a knot K in a 3-manifold M, this control is best stated in terms of "polar coordinates on a neighborhood of K". By this we mean functions (r, μ, λ) describing an orientation-preserving embedding of a neighborhood ν of K into $\mathbb{R}^2 \times S^1$, where (r, μ) map to polar coordinates on $\mathbb{R}^2$ and λ maps to S^1, with $K = \{r = 0\}$. (On a torus $\{r = R > 0\}$, the "μ-axis" is a meridian curve while the "λ-axis" is a longitude.) We will always assume that such a neighborhood ν is mapped onto $\{r < R_\nu\} \subset \mathbb{R}^2 \times S^1$ for some radius $R_\nu > 0$. The coordinates determine a framing of K; in fact the function μ is enough to determine the framing, so we label this framing F_μ.

We begin with results which will be proved in chapter 4, concerning the construction of symplectic 4-manifolds with weakly convex boundaries.

Lemma 1.3. *Given a contact structure ξ on a 3-manifold M and a transverse knot K, there exists a neighborhood ν of K with coordinates (r, μ, λ) and a function $s : (0, R_\nu^2] \to \mathbb{R}$ such that $\xi = \mathrm{span}\{\partial_r, s(r^2)\partial_\mu + \partial_\lambda\}$ on $\nu \setminus K$. If ξ is positive then $s' > 0$ and $\lim_{x\to 0} s(x) = -\infty$. If ξ is negative then $s' < 0$ and $\lim_{x\to 0} s(x) = \infty$.*

This function s will be called the "framing function" associated to α and the coordinates (r, μ, λ) because, when $s(R^2) \in \mathbb{Z}$, the characteristic foliation on the torus $\{r = R\}$ is a foliation by parallel simple closed curves linking K and realizing the framing $F_\mu + s(R^2)$.

Definition 1.4. *In the setup of lemma 1.3, when ξ is positive and we are given a framing F of K, we say that ν is* fat with respect to F *if $s(R_\nu^2) > F - F_\mu$. A transverse link $L = K_1 \cup \ldots K_n$ is fat with respect to a framing F if their exists a family of mutually disjoint neighborhood $\nu_i \supset K_i$ each of which is fat with respect to F_{K_i}.*

(One should think of the "fatness" of a transverse knot as a framing-valued contact injectivity radius.)

Given two framings F and F' of a link L, we say that $F' \leq F$ if, for each component K of L, $F_K - F'_K \geq 0$. So by lemma 1.3 every transverse link L is fat with respect to some framing, and if L is fat with respect to F and $F' \leq F$ then L is fat with respect to F'.

Now let (X, ω) be a symplectic 4-manifold with (strongly) convex boundary and let L be a link in ∂X with some framing F. Let Y be a 4-manifold built by attaching 2-handles along L with framing F.

Theorem 1.5. *In terms of the induced contact structure on ∂X, if L is a transverse link and is fat with respect to F, then Y can be constructed so that ω extends across the 2-handles and so that ∂Y is* weakly *convex.*

When we prove this theorem we will describe explicitly how the contact structure on ∂Y is related to that on ∂X. In this case we cannot ask that the handles be attached in arbitrarily small neighborhoods of the knots; we will actually "use up" each fat neighborhood.

Compare theorem 1.5 to theorem 1.1, case 2. Using Weinstein's symplectic 2-handles and the fact that we can always decrease the Thurston-Bennequin framing of a Legendrian knot, we can construct contact 3-manifolds by surgeries on a (topological) link L with all framings up to some upper bound given by the possible Legendrian realizations of the isotopy class of L. Theorem 1.5 also produces contact 3-manifolds by surgeries on links with framings *less than* a certain upper bound, and the proposition below says that in fact theorem 1.5 can achieve all the framings achievable by Legendrian surgeries.

Proposition 1.6. *Suppose that K is a Legendrian knot in a positive contact 3-manifold with a given neighborhood ν and a framing $F \leq \mathrm{tb}(K) - 1$. Then there exists a* transverse *knot K' inside ν, isotopic to K, which is fat with respect to F.*

(This fact was pointed out to me by John Etnyre and will be proved in section 4.1.) Since weak convexity is sufficient for tightness, both constructions yield tight contact structures.

Next we present results which will be proved in chapter 6, concerning the construction of symplectic 4-manifolds with concave boundaries. We begin with more explicit control on 1-forms near knots.

Definition 1.7. *Suppose (r, μ, λ) are polar coordinates on a neighborhood ν of a knot K. A contact form α on ν is* well-behaved with respect to (r, μ, λ) *if:*

1. *there exist constants A and B, with $B > 0$, such that $R_\alpha = A\partial_\mu + B\partial_\lambda$,*
2. *$\mathcal{L}_{\partial_\mu}\alpha = 0$ and*
3. *$\alpha(\partial_r) = 0$.*

A closed 1-form α^0 on $\nu \setminus K$ is well-behaved with respect to (r, μ, λ) *if $\alpha^0 = Cd\mu + Dd\lambda$ for some constants C and D, with $C > 0$.*

The 1-form in question is said to be well-behaved near K *if there exist coordinates on a neighborhood of K with respect to which it is well-behaved. In this situation we may also say that K is well-behaved with respect to the 1-form.*

In this definition, the first two conditions in fact say that α is invariant under the obvious action of $S^1 \times S^1$ on ν. One consequence of this definition is that if a contact form α is well-behaved near K then K is necessarily a closed orbit of R_α. Also if α (or α^0) is well-behaved near K then it is well-behaved with respect to coordinates realizing any desired framing of K. A closed 1-form α^0 which is well-behaved near K has some bearing on framings of K: We will say that a framing F of K is *positive* with respect to α^0 if, given well-behaved coordinates (r, μ, λ), we have $F - F_\mu > -D/C$. (One can think of α^0 as specifying an "$\mathbb{R}$-valued framing" $F_{\alpha^0} = F_\mu - D/C$.)

Now suppose that (X, ω) is a symplectic 4-manifold with convex boundary and let L be a link in ∂X with framing F. Let α be the induced contact form on ∂X and let Y be a smooth manifold built by attaching 2-handles along L with framing F.

Theorem 1.8. *Suppose there exists a closed 1-form α^0 on $\partial X \setminus L$ with $\alpha^0(R_\alpha) > 0$ and suppose that ∂X is compact. If, for each component K of L, there exist coordinates near*

K with respect to which α and α^0 are both well-behaved and such that F_K is positive *with respect to α^0, then we can construct Y such that ω extends across the 2-handles and so that ∂Y is* concave*. These handles can be attached in arbitrarily small neighborhoods of the knots.*

Again, the induced contact form on ∂Y will be described explicitly when we give the construction. Compare this result to the previous theorems; here we can attach the handles with any framing *larger than* some lower bound.

In our examples, we will have $\alpha^0 = dP$ where $P : \partial X \setminus L \to S^1$ is a fibration with R_α positively transverse to the fibers. For some simple examples where this theorem applies, suppose X is the unit $B^4 \subset \mathbb{R}^4$ with the standard symplectic form and the standard symplectic dilation mentioned earlier, giving $S^3 = \partial B^4$ the standard contact form α. In chapter 7 we will show that when L is the standard unknot in S^3, theorem 1.8 applies for any framing $F \geq 1$ and that when L is the standard Hopf link in S^3, theorem 1.8 applies for any framing $F \geq 0$ (where 0 refers to the 0-framing on each component). In fact in chapter 7 we will present a large class of surgery descriptions of 3-manifolds which can be realized as the concave boundaries of symplectic 4-manifolds using this theorem.

Chapter 6 proves more general results, but these depend on the following ideas, developed in chapter 5. In theorem 1.8, if L has more than one component but only some of the handles are attached, we do not get a concave boundary but rather a boundary which is partially concave and partially convex. This notion is integral to the constructions and is made rigorous as follows:

Let (X, ω) be a symplectic 4-manifold.

Definition 1.9. *Given vector fields V^+ and V^- defined respectively on (possibly empty) open subsets X^+ and X^- of X, we say that (V^+, V^-) is a* dilation-contraction pair *if the following equations hold (on the sets where they make sense):*

$$\mathcal{L}_{V^\pm}(\omega) = \pm\omega, \quad \omega(V^+, V^-) = 0 \,.$$

(The second equation provides the necessary rigidity for these structures.) Suppose that M is a 3-dimensional submanifold of (X, ω) and that (V^+, V^-) is a dilation-contraction pair in X with domains $X^\pm$.

Definition 1.10. *(V^+, V^-)* transversely covers *M if $M \subset X^+ \cup X^-$ and if each vector field is positively transverse, where defined, to M.*

Letting $M^{\pm} = M \cap X^{\pm}$ and $\alpha^{\pm} = \imath_{V^{\pm}}(\omega)|_{M^{\pm}}$, the pair (α^+, α^-) is an instance of the following:

Definition 1.11. *A* contact pair *on M is a pair (α^+, α^-) of 1-forms, defined respectively on (possibly empty) open subsets $M^{\pm}$ of M, such that $M = M^+ \cup M^-$ and such that the following equations hold (on the sets where they make sense):*

$$\pm\alpha^{\pm} \wedge d\alpha^{\pm} > 0, \quad -d\alpha^- = d\alpha^+$$

By a boundary which is partially convex and partially concave, we mean a boundary which is transversely covered by a dilation-contraction pair. We will show in chapter 5 that the induced contact pair on the boundary uniquely determines the germ of the symplectic form along the boundary. If (X, ω_X) and (Y, ω_Y) are two symplectic 4-manifolds with boundary components M and $-N$ respectively, suppose that M and N are both transversely covered by dilation-contraction pairs. If there is a diffeomorphism from N to M preserving the contact pairs we can then symplectically glue (X, ω_X) to (Y, ω_Y), by the uniqueness of the germs. This generalizes the "convex to concave" glueings described in section 1.2.

In general we will use the language of bordisms to describe all these results carefully and we will find it worthwhile to introduce a slight generalization of the familiar notion of "bordism".

Chapter 2

Terminology for Symplectic Constructions

When constructing smooth manifolds by glueing together elementary building blocks such as handles, one has to be careful about terminology because some of the building blocks may not technically be manifolds, or because one may need to "round corners" at various stages. When adding symplectic data to such constructions one should be even more careful. For this reason we develop some abstractions and nonstandard terminology here for some standard ideas; with this terminology the proofs and statements of later results will be simpler and more precise than they would be without it.

2.1 Topological terminology

Definition 2.1. *Given oriented $(n-1)$-manifolds M_1 and M_2, a* patch from M_1 to M_2 *is a triple $P = (X, X_1, X_2)$ with the following properties:*

- *$X = X_1 \cup X_2$ is an n-manifold (without boundary), and X_1 and X_2 are n-dimensional submanifolds of X* with boundary, *with $\partial X_1 = -M_1$ and $\partial X_2 = M_2$.*
- *The* support *of P, written* $\mathrm{supp}(P)$ *and defined to be the closure of the interior of $X_1 \cap X_2$, is compact.*

(Note that a closed manifold is a patch from the empty manifold to itself.)

We introduce the following notation, for $i = 1, 2$:

$$\partial_i P = M_i, \quad \text{and} \quad \mathrm{supp}_i(P) = \mathrm{supp}(P) \cap \partial_i P\ .$$

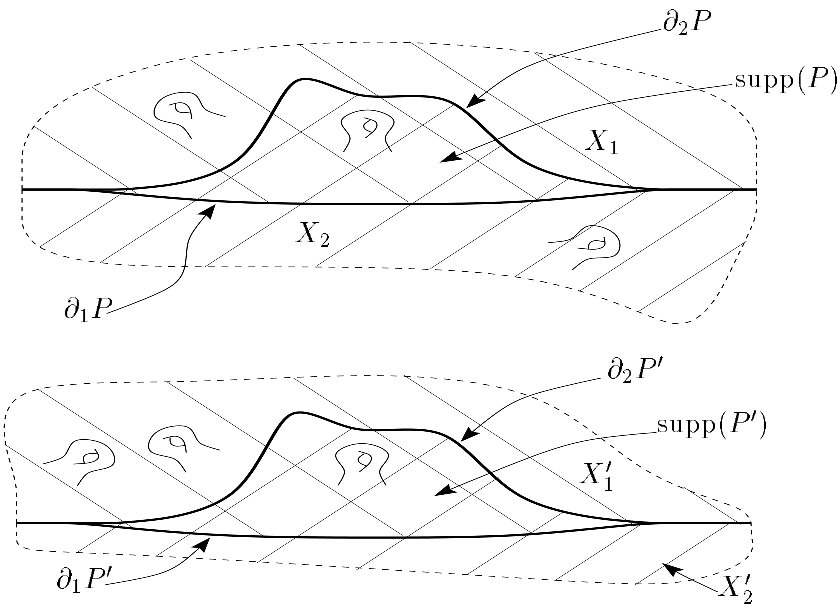

Figure 2.1: Two *equivalent* patches $P = (X, X_1, X_2)$ and $P' = (X', X_1', X_2')$.

Another patch $P' = (X', X_1', X_2')$ is a *restriction* of P if X' is an open neighborhood of $X_1 \cap X_2$ in X with $X_i' = X_i \cap X'$. Two patches are *equivalent* if they share a common restriction. (See figure 2.1.)

If $P = (X, X_1, X_2)$ and $Q = (Y, Y_1, Y_2)$ are patches, then a *patch diffeomorphism* from P to Q is a diffeomorphism of triples

$$\phi : (X', X_1', X_2') \to (Y', Y_1', Y_2')$$

where $P' = (X', X_1', X_2')$ is a restriction of P and $Q' = (Y', Y_1', Y_2')$ is a restriction of Q. If P and Q are equivalent and R and S are equivalent, then a patch diffeomorphism from P to R and a patch diffeomorphism from Q to S are equivalent if they agree on a common restriction of P and Q.

Definition 2.2. *A* generalized bordism $\mathcal{P}$ *from M_1 to M_2 is an equivalence class of patches from M_1 to M_2. A* generalized bordism diffeomorphism *is an equivalence class of patch diffeomorphisms.*

If P is a patch representing $\mathcal{P}$ then $\mathrm{supp}(\mathcal{P}) = \mathrm{supp}(P)$, $\partial_i \mathcal{P} = \partial_i P$ and $\mathrm{supp}_i(\mathcal{P}) = \mathrm{supp}_i(P)$. We will sometimes call $\partial_1 \mathcal{P}$ the *bottom* boundary and $\partial_2 \mathcal{P}$ the *top* boundary. The existence of collar neighborhoods for boundaries and the tubular neighborhood theorem

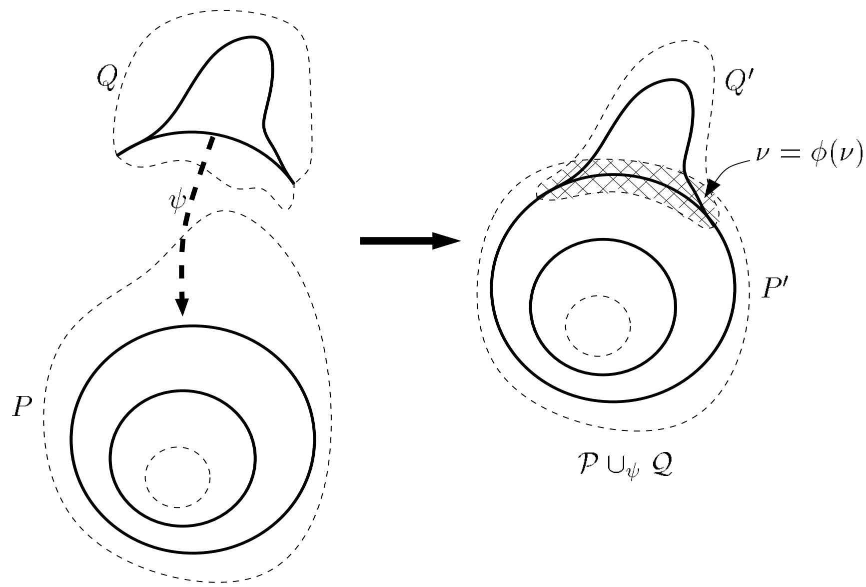

Figure 2.2: Glueing two patches together after trimming away excess material.

(see [9]) together say that an ordinary (compact) bordism B defines a unique generalized bordism whose support is B. Henceforth the word "bordism" will refer to a generalized bordism.

This terminology gives a convenient way to describe a glueing procedure. Given two n-dimensional bordisms $\mathcal{P}$ and $\mathcal{Q}$ represented respectively by patches $P = (X, X_1, X_2)$ and $Q = (Y, Y_1, Y_2)$, suppose that $\psi : \partial_1 \mathcal{Q} \hookrightarrow \partial_2 \mathcal{P}$ is an embedding. Then the tubular neighborhood theorem and the compactness of the supports implies that there exists a bordism $\mathcal{P} \cup_\psi \mathcal{Q}$ represented by a patch constructed as follows: There exist restrictions P' and Q' of P and Q, respectively, and an imbedding ϕ from a neighborhood ν of $\partial_1 \mathcal{Q}$ in Y into X with $\phi|_{\partial_1 \mathcal{Q}} = \psi$ and such that

$$(X' \cup_\phi Y',\ X_1' \cup_{\phi|_{\nu \cap Y_1'}} Y_1',\ X_2' \cup_{\phi|_{\nu \cap Y_2'}} Y_2')$$

is a patch from $\partial_1 \mathcal{P}$ to $(\partial_2 \mathcal{P} \setminus \psi(\mathrm{supp}_1(\mathcal{Q}))) \cup_\psi \partial_2 \mathcal{Q}$. (When glueing two manifolds together along open sets one has only to check that the resulting space is Hausdorff in order to conclude that it is a manifold. In this case we choose the restrictions carefully so that the result of the glueing is Hausdorff.) Then $\mathcal{P} \cup_\psi \mathcal{Q}$ is the equivalence class of this patch and does not depend on the choice of P', Q' or ϕ. (See figure 2.2.)

A *product patch* (see figure 2.3) is a patch $P = (X, X_1, X_2)$ for which:

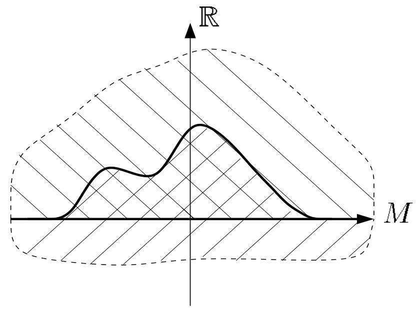

Figure 2.3: A product patch

- $X = \{(t,p) \mid -\epsilon(p) < t < f(p)\} \subset \mathbb{R} \times M$ for a manifold M and positive functions ϵ and f on M,
- $X_1 = \{(t,p) \mid t \geq 0\}$ and
- $X_2 = \{(t,p) \mid t \leq h(p)\}$ for a function $h : M \to [0,\infty)$ with compact support and with $h < f$.

In this case $\partial_1 P = M_0 = \{0\} \times M$ and $\partial_2 P = M_h = \{(h(p),p) \mid p \in M\}$, the graph of h in $\mathbb{R} \times M$; in particular there is a natural diffeomorphism from $\partial_1 P$ to $\partial_2 P$. We will often suppress mention of this diffeomorphism and think of M, M_0 and M_h as the same manifold. We will call h the *height* of P. The only topological significance to the height is that $\text{supp}_1(P) = \{0\} \times \text{supp}(h)$; the height will be more important when we add symplectic data.

A *product bordism* is a bordism represented by a product patch. Two *equivalent* product patches have the same height, so we can speak of the *height* of a product bordism. Given a manifold M and a non-negative function h on M, we in fact have a unique product bordism of height h which we will call "the product bordism of height h constructed on M" and which we will label $\mathcal{B}_h$.

2.2 Handles and how to build them

Here we describe a special class of bordisms carrying nontrivial topology, given two integers n and k with $0 \leq k \leq n$. An *n-dimensional k-handle* is a bordism diffeomorphic to a bordism $\mathcal{H}$ constructed as follows: (We will frequently refer back to this construction and the notation established here.)

When $k=0$, we take $X=X_1=\mathbb{R}^n$, let X_2 be any closed ball centered at 0 and let $\mathcal{H}=[(X,X_1,X_2)]$, so that $\partial_1\mathcal{H}=\emptyset$ and $\partial_2\mathcal{H}\cong S^{n-1}$. When $k=n$ we take $X=X_2=\mathbb{R}^n$ and X_1 to be any ball centered at 0, so that now $\partial_1\mathcal{H}\cong -S^{n-1}$ and $\partial_2\mathcal{H}=\emptyset$. These two cases actually fit into the following more general construction given a few conventions, but it is simplest to see them separately.

When $0<k<n$, we present a construction tailored to unify the various constructions of symplectic handles in later chapters, but this construction is easily seen to be equivalent to more standard constructions, such as that in [9].

Consider $\mathbb{R}^n=\mathbb{R}^k\times\mathbb{R}^{n-k}$ with coordinates $x_1,\dots,x_n$ and the two radius functions:

$$r_1^2=\sum_{i=1}^{k}x_i^2\ ,\quad r_2^2=\sum_{i=k+1}^{n}x_i^2\ .$$

Then $f=-r_1^2+r_2^2$ is the standard Morse function with an index k critical point at 0. Suppose we are given an orientation on $\mathbb{R}^n$ (we will have one example where this is not the standard orientation). By convention we orient non-critical level sets $f^{-1}(t)$ as $\partial(f^{-1}(-\infty,t])$. Let $Z^{desc}=\{r_2=0\}$, the "descending manifold" or "core", and let $Z^{asc}=\{r_1=0\}$, the "ascending manifold" or "co-core". Pick constants $\epsilon_1<0<\epsilon_2$ and consider the submanifolds $f^{-1}(\epsilon_1)$ and $f^{-1}(\epsilon_2)$, with $K_1=Z^{desc}\cap f^{-1}(\epsilon_1)$, the "descending sphere", and $K_2=Z^{asc}\cap f^{-1}(\epsilon_2)$, the "ascending sphere".

Now suppose V is a gradient-like vector field for f (i.e. $df(V)>0$ when $df\neq 0$ and $V=0$ when $df=0$) and suppose that V is tangent to Z^{desc} and Z^{asc} (so that these really are the descending and ascending manifolds for V). (See figure 2.4.) Because V is positively transverse to $f^{-1}(\epsilon_1)$ there exists an open neighborhood $\mathcal{N}$ of $\{0\}\times f^{-1}(\epsilon_1)$ in $\mathbb{R}\times f^{-1}(\epsilon_1)$ with an embedding $\Phi:\mathcal{N}\hookrightarrow\mathbb{R}^n$ given by flow along V starting on $f^{-1}(\epsilon_1)$. Since V is gradient-like for f, $\mathcal{N}$ can be made large enough so that Φ maps $\mathcal{N}$ onto $\mathbb{R}^n\setminus Z^{asc}$. In particular there is some function $T:f^{-1}(\epsilon_1)\setminus K_1\to(0,\infty)$ such that $\Phi(T(p),p)\in f^{-1}(\epsilon_2)$ for all $p\in f^{-1}(\epsilon_1)\setminus K_1$. ($T$ is the time it takes to get from $f^{-1}(\epsilon_1)$ to $f^{-1}(\epsilon_2)$ along V). Thus we also get a diffeomorphism $\phi:f^{-1}(\epsilon_1)\setminus K_1\to f^{-1}(\epsilon_2)\setminus K_2$ defined by $\phi(p)=\Phi(T(p),p)$.

Let C, C' and U be tubular neighborhoods of K_1 in $f^{-1}(\epsilon_1)$, with C and C' compact, U open, $C\subset\text{interior}(C')$ and $C'\subset U$. Choose a function $g:f^{-1}(\epsilon_1)\to[0,1]$

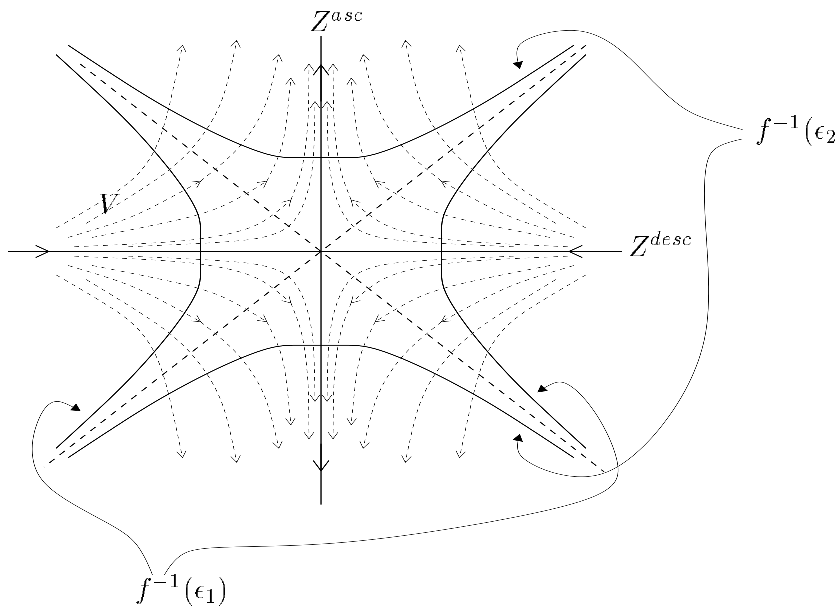

Figure 2.4: Two level sets of a Morse function f, the flow lines of a gradient-like vector field V and the ascending and descending manifolds.

with $\mathrm{supp}(g) = C'$ and $g|_C \equiv 1$ and let $h = gT$. Finally let

$$X = \Phi(\mathcal{N} \cap (\mathbb{R} \times U)) \cup Z^{asc}$$
$$X_1 = \Phi(\mathcal{N} \cap ([0, \infty) \times U)) \cup Z^{asc}$$
$$X_2 = \Phi(\{(t,p) \mid t \leq h(p), p \in U \setminus K_1\}) \cup (Z^{asc} \cap f^{-1}(-\infty, \epsilon_2]) \cup Z^{desc} .$$

Then let $\mathcal{H} = [(X, X_1, X_2)]$; we have $\partial_1 \mathcal{H} = U$ and $\partial_2 \mathcal{H} = \Phi(\{(h(p), p) \mid p \in U\}) \cup K_2$, with $\mathrm{supp}_1(\mathcal{H}) = C' = \mathrm{supp}(g) = \mathrm{supp}(h) \cup K_1$. We will call the function h the *height* of the handle $\mathcal{H}$, which again has little topological significance but which will later help us keep track of symplectic data. Notice that V is positively transverse to both $\partial_1 \mathcal{H}$ and $\partial_2 \mathcal{H}$. (See figure 2.5.)

The choices involved in this construction are the constants ϵ_1 and ϵ_2, the vector field V, the tubular neighborhoods C, C' and U and the function g. It is easy to see that different choices yield diffeomorphic patches, but it will be important later to make these choices carefully.

Via the splitting $\mathbb{R}^n = \mathbb{R}^k \times \mathbb{R}^{n-k}$, we get diffeomorphisms $f^{-1}(\epsilon_1) \cong S^{k-1} \times \mathbb{R}^{n-k}$ and $f^{-1}(\epsilon_2) \cong \mathbb{R}^k \times S^{n-k-1}$. So $K_1 \cong S^{k-1}$ and $K_2 \cong S^{n-k-1}$ and both spheres come with canonical framings from the splitting $\mathbb{R}^n = \mathbb{R}^k \times \mathbb{R}^{n-k}$. We will refer to these framings as

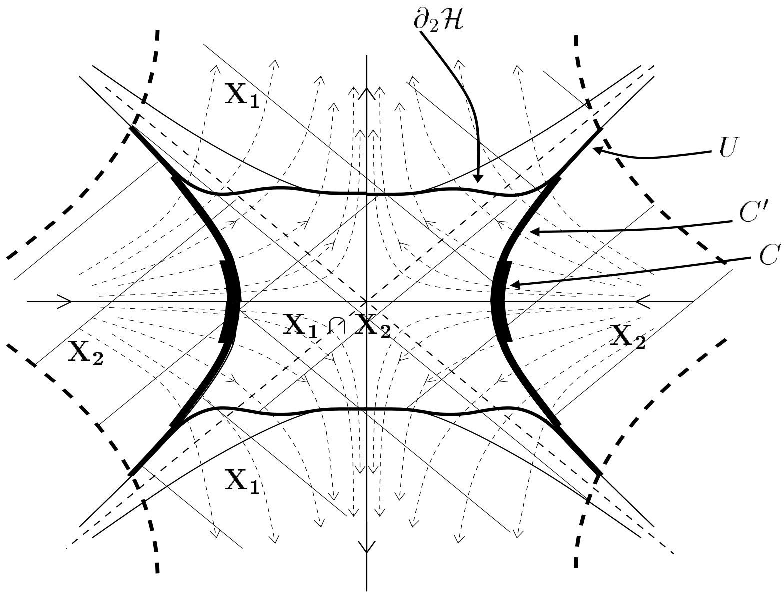

Figure 2.5: Constructing the handle $\mathcal{H}$ beginning with sets $C \subset C' \subset U \subset f^{-1}(\epsilon_1)$ and flowing along V.

the "handle-framings" of K_1 and K_2. (Depending on n and k one may need to be careful about orientations.) Thus, attaching an n-dimensional k-handle $\mathcal{H}$ to another bordism $\mathcal{B}$ along $\psi : \partial_1\mathcal{H} \hookrightarrow \partial_2\mathcal{B}$ involves choosing an embedding of a neighborhood of $S^{k-1} \times \{0\}$ in $S^{k-1} \times \mathbb{R}^{n-k}$ into $\partial_2\mathcal{B}$. In particular this involves a choice of an embedding of S^{k-1} into $\partial_2\mathcal{B}$ and a framing F of the image submanifold K. One of the basic results in handlebody theory (see [9]) is that, up to a diffeomorphism, $\mathcal{B} \cup_\psi \mathcal{H}$ depends only on K and F, so that we speak of attaching a handle "along K with framing F".

Note that, unless $k = 0$ or n, the support of $\mathcal{H}$ will not be a manifold. The significance of these building blocks, however, is that every closed n-manifold can be built by glueing handles together, starting with 0-handles and ending with n-handles (see [9]). In fact every compact bordism, in the usual sense of the word "bordism", can be built by starting with a product bordism with positive height and attaching handles.

A 1-handle is attached along a 0-sphere or, in other words, a pair of points. When everything is oriented there is only one framing of a pair of points, so that we do not need to specify a framing when attaching a 1-handle.

We will be particularly interested in 4-dimensional 2-handles, which are attached along framed knots in 3-manifolds. We develop some more explicit notation for this case, paying attention to our orientation convention for level sets. When $n = 4$ and $k = 2$, the submanifolds $f^{-1}(\epsilon_1)$, $\partial_1\mathcal{H}$, $f^{-1}(\epsilon_2)$ and $\partial_2\mathcal{H}$ are all diffeomorphic to $\mathbb{R}^2 \times S^1$ and we have the ascending and descending *circles* K_1 and K_2. We will write down explicit oriented polar coordinates on $f^{-1}(\epsilon_1)$ and $f^{-1}(\epsilon_2)$ as tubular neighborhoods of K_1 and K_2. We may then use coordinates on $\partial_1\mathcal{H}$ coming from those on $f^{-1}(\epsilon_1)$ and we may either use coordinates on $\partial_2\mathcal{H}$ coming from those on $f^{-1}(\epsilon_2)$ via backward flow along V, or we may use coordinates on $\partial_2\mathcal{H} \setminus K_2$ coming from those on $f^{-1}(\epsilon_1)$ via forward flow along V.

In the specific symplectic constructions in later chapters, we will work with positively oriented coordinates labelled (x_1, y_1, x_2, y_2) on $\mathbb{R}^4$. In the new constructions (chapters 4 and 6) these will line up with the coordinates in the general construction above in the sense that we will have $r_1^2 = x_1^2 + y_1^2$ and $r_2^2 = x_2^2 + y_2^2$. If we let (r_i, θ_i) be polar coordinates in the (x_i, y_i) plane, then correctly oriented polar coordinates on $f^{-1}(\epsilon_1)$ are given by:

$$(r = r_2, \mu = \theta_2, \lambda = -\theta_1)$$

For future reference note that $r_1^2 = r^2 - \epsilon_1$ on $f^{-1}(\epsilon_1)$. Correctly oriented polar coordinates on $f^{-1}(\epsilon_2)$ are given by:

$$(r = r_1, \mu = \theta_1, \lambda = \theta_2)$$

Now note that $r_2^2 = r^2 + \epsilon_2$ on $f^{-1}(\epsilon_2)$.

When we present Weinstein's symplectic 2-handles in section 3.3, we will instead have $r_1^2 = r_x^2 = x_1^2 + x_2^2$ and $r_2^2 = r_y^2 = y_1^2 + y_2^2$. Then the coordinate system which lines up with the construction above is (x_1, x_2, y_1, y_2), which is negatively oriented. Now let (r_x, θ_x) and (r_y, θ_y) be polar coordinates in the (x_1, x_2) and (y_1, y_2) planes, respectively. Then an orientation-preserving diffeomorphism from $f^{-1}(\epsilon_1)$ to $\mathbb{R}^2 \times S^1$ realizing the handle-framing is described by

$$(x = y_1, y = y_2, \lambda = \theta_x),$$

where (x, y) is the usual rectangular coordinate system on $\mathbb{R}^2$.

After establishing oriented coordinates (r, μ, λ) on $f^{-1}(\epsilon_1)$, attaching a 4-dimensional 2-handle along a knot K with framing F now simply involves finding a neighborhood ν of K with coordinates which we also label (r, μ, λ) such that $F_\mu = F$ and picking three positive radii $R_1 < R_2 < R_3$ with $R_3 < R_\nu$, the radius of ν. Then we construct $\mathcal{H}$ with $C = \{r \leq R_1\}$, $C' = \{r \leq R_2\}$ and $U = \{r < R_3\}$. Identifying the coordinates on

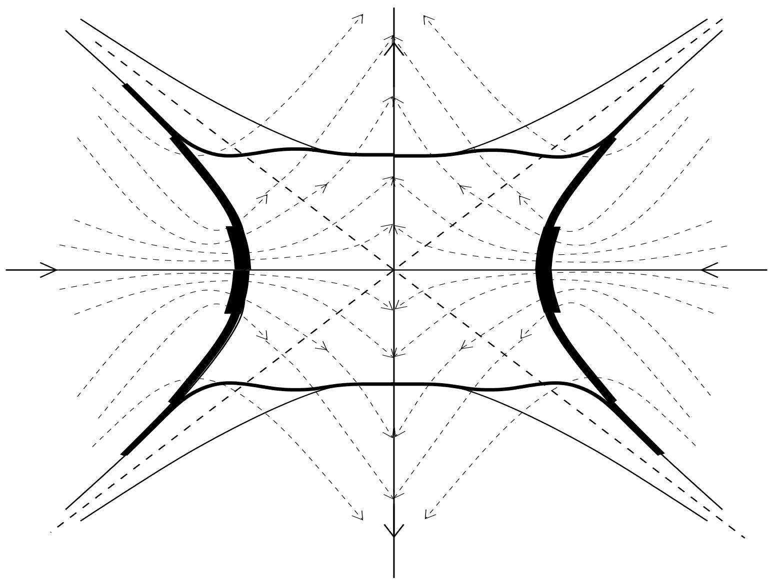

Figure 2.6: Why we don't worry when V fails to extend across Z^{asc} or to be gradient-like outside a neighborhood of $f^{-1}[\epsilon_1, \epsilon_2] \setminus Z^{asc}$.

$U \subset f^{-1}(\epsilon_1)$ with those on ν gives an embedding of $\partial_1\mathcal{H}$ into ν sending the handle-framing of K_1 to the framing F of K.

A useful feature of the examples we construct will be that the "angle coordinates" θ_1 and θ_2 (or θ_x and θ_y) are invariant under flow along the gradient-like vector fields V that we use. Given this, the diffeomorphism $\phi : f^{-1}(\epsilon_1) \setminus K_1 \to f^{-1}(\epsilon_2) \setminus K_2$ induced by flow along V will satisfy the following equations:

$$r \circ \phi = g(r) \quad \text{for some increasing function } g$$
$$\mu \circ \phi = -\lambda, \quad \lambda \circ \phi = \mu$$

In fact, in the new 2-handles constructed in chapters 4 and 6, we will slightly complicate the construction outlined above. The vector field V which we use in each case will preserve the angle coordinates but will not be defined on Z^{asc} and will fail to be gradient-like outside an open neighborhood of $f^{-1}[\epsilon_1, \epsilon_2] \setminus Z^{asc}$. The last point is obviously minor if we can deal with the first point since the entire construction can be carried out inside a neighborhood of $f^{-1}[\epsilon_1, \epsilon_2]$. To deal with the fact that V does not extend across Z^{asc}, we will show in each case that the map induced by flow along V is defined as a diffeomorphism

$$\phi : f^{-1}(\epsilon_1) \setminus \{r \leq R\} \to f^{-1}(\epsilon_2) \setminus K_2$$

Then the construction can be easily modified to work with such a vector field as long as we choose $R_1 > R$ (see figure 2.6). The height of the handle with respect to such a vector field will then be a function $h : f^{-1}(\epsilon_1) \setminus \{r \leq R\} \to [0, \infty)$. We will always assume that $h = h(r^2)$ is a function of r^2.

2.3 Symplectic terminology

Now we want to add symplectic data to the topological constructions above.

Definition 2.3. *If $k < 2n$, a* germ $\mathcal{G}$ *of a $2n$-dimensional symplectic form along a k-manifold M is an equivalence class of symplectic $2n$-manifolds containing M as a submanifold under the following equivalence relation: Two symplectic manifolds $(X, \omega_X) \supset M$ and $(Y, \omega_Y) \supset M$ are equivalent if there exist neighborhoods $\mathcal{N}_X$ and $\mathcal{N}_Y$ of M in X and Y, respectively, with a symplectomorphism $\phi : \mathcal{N}_X \to \mathcal{N}_Y$ which is the identity along M.*

If ϕ is an embedding of another k-manifold N into M and $\imath$ is the inclusion of M into a symplectic manifold (X, ω) representing a germ $\mathcal{G}$ along M, then we can pull $\mathcal{G}$ back to a germ $\phi^*\mathcal{G}$ along N by pulling back ω via the exponential map to the normal bundle of $\imath \circ \phi : N \hookrightarrow X$. Given any non-zero constant c and a germ $\mathcal{G}$ along M represented by (X, ω), we define the germ $c\mathcal{G}$ to be the germ represented by $(X, c\omega)$. (Actually, if n is odd and $c < 0$ then we must reverse the orientation on X, but we are not concerned with that case in this paper.) Of course a germ along M can also be thought of as a germ along $-M$.

Definition 2.4. *Given $(2n-1)$-manifolds M_1 and M_2 with $2n$-dimensional symplectic germs $\mathcal{G}_1$ and $\mathcal{G}_2$, a* symplectic patch from $(M_1, \mathcal{G}_1)$ to $(M_2, \mathcal{G}_2)$ *is a patch $P = (X, X_1, X_2)$ from M_1 to M_2 with a symplectic form ω such that (X, ω) represents both $\mathcal{G}_1$ and $\mathcal{G}_2$.*

A *symplectic restriction* is a patch restriction with the inherited symplectic form. Two symplectic patches are equivalent if they have a common symplectic restriction and a *generalized symplectic bordism* is an equivalence class of symplectic patches. In general we will say "symplectic bordism" when we mean "generalized symplectic bordims". When $\mathcal{P}$ is a symplectic bordism from $(M_1, \mathcal{G}_1)$ to $(M_2, \mathcal{G}_2)$, we may also write $\partial_i\mathcal{P} = (M_i, \mathcal{G}_i)$ or

$\partial_i \mathcal{P} = M_i$ and $\mathcal{G}(\partial_i \mathcal{P}) = \mathcal{G}_i$. (For "$\mathcal{G}(\partial_1 \mathcal{P})$" read "the germ along the bottom boundary of $\mathcal{P}$" and for "$\mathcal{G}(\partial_2 \mathcal{P})$" read "the germ along the top boundary of $\mathcal{P}$".)

A patch symplectomorphism is a patch diffeomorphism which respects the symplectic forms. Likewise we can define equivalent patch symplectomorphisms and thus bordism symplectomorphisms.

Now by the definition of a symplectic germ and the argument in the topological case we can see that symplectic bordisms can be glued together just like topological bordisms: If $\mathcal{P}$ and $\mathcal{Q}$ are symplectic bordisms and if $\psi : \partial_1 \mathcal{Q} \hookrightarrow \partial_2 \mathcal{P}$ is an embedding for which $\psi^* \mathcal{G}(\partial_2 \mathcal{P}) = \mathcal{G}(\partial_1 \mathcal{Q})$, then we can glue $\mathcal{Q}$ to $\mathcal{P}$ along ψ to form the symplectic bordism $\mathcal{P} \cup_\psi \mathcal{Q}$.

Given a non-zero constant c and bordism $\mathcal{P}$ represented by (X, X_1, X_2, ω), we can always form a bordism $c\mathcal{P}$ by replacing ω with $c\omega$. Then we have $\mathcal{G}(\partial_i c\mathcal{P}) = c\mathcal{G}(\partial_i \mathcal{P})$.

Considering product bordisms in this context gives rise to the following question: Given a $(2n-1)$-manifold M and two $2n$-dimensional symplectic germs $\mathcal{G}_1$ and $\mathcal{G}_2$ along M, does there exist a product bordism from $(M, \mathcal{G}_1)$ to $(M, \mathcal{G}_2)$? For certain special classes of germs we will construct such product bordisms and use them extensively in our constructions.

We will construct symplectic handles by following the construction in section 2.2 and imposing some symplectic form on $\mathbb{R}^n$. By choosing the vector fields V in the construction carefully we will able to use contact forms to record the symplectic germs along the top and bottom boundaries.

Chapter 3

Background Results

3.1 Standard local results in contact and symplectic topology

A variety of results describing local behaviours of symplectic forms and contact forms can be deduced from an argument attributed to Moser and described carefully by McDuff and Salamon in [10]. The simplest version is Darboux's theorem, which states that any two symplectic forms on a neighborhood of a point are, after perhaps shrinking the neighborhoods, symplectomorphic. Here we will state only the versions we need.

The symplectic results are very standard so we begin by quoting a lemma verbatim and without proof:

Lemma 3.1 (Lemma 3.14 in [10], page 94). *Let M be a $2n$-dimensional smooth manifold and $Q \subset M$ be a compact submanifold. Suppose that $\omega_0, \omega_1 \in \Omega^2(M)$ are closed 2-forms such that at each point q of Q the forms ω_0 and ω_1 are equal and nondegenerate on T_qM. Then there exist open neighborhoods $\mathcal{N}_0$ and $\mathcal{N}_1$ of Q and a diffeomorphism $\psi : \mathcal{N}_0 \to \mathcal{N}_1$ such that*

$$\psi|_Q = \mathrm{id}, \quad \psi^*\omega_1 = \omega_0 \ .$$

The immediate corollary most useful for us is:

Corollary 3.2. *For $i = 1, 2$, suppose that (X_i, ω_i) are symplectic $2n$-manifolds with compact submanifolds M_i. If $\phi : T_{M_1}X_1 \to T_{M_2}X_2$ is a bundle isomorphism for which $\phi^*\omega_2 = \omega_1|_{T_{M_1}X_1}$ then there exists a* symplectomorphism *Ψ from a neighborhood of M_1 to a neighborhood of M_2 with $\Psi|_{M_1} = \phi|_{M_1}$.*

Proof. Extend ϕ to a diffeomorphism Φ from a neighborhood of M_1 to a neighborhood of M_2. Let $\omega_0 = \Phi^*\omega_2$ and note that ω_0 and ω_1 are equal and nondegenerate on $T_{M_1}X_1$, so we can apply the lemma quoted above to get a map ψ which is the identity along M_1 and such that $\psi^*\omega_1 = \omega_0$. Then let $\Psi = \Phi \circ \psi^{-1}$. □

We will develop analogous results for contact structures beginning with a lemma which is proved (but not stated as a lemma) in [10].

Lemma 3.3. *For any family α_t of contact forms on a manifold M there exists a family X_t of vector fields, generating a flow ψ_t, such that $\psi_t^*\alpha_t = f_t\alpha_0$ (for some family f_t of functions on M). Furthermore, wherever $\frac{d}{dt}\alpha_t = 0$, we will have $X_t = 0$.*

Proof. (This proof is cobbled together from a proof in [10] and a proof in [5].) Let $\alpha_t' = \frac{d}{dt}\alpha_t$ and let $h_t = \alpha_t'(R_{\alpha_t})$, where R_{α_t} is the Reeb vector field for α_t. We claim that there exists a unique vector field $X_t \in \ker\alpha_t$ such that $h_t\alpha_t - \alpha_t' = \imath_{X_t}d\alpha_t$. Given this claim, and since $\imath_{X_t}\alpha_t = 0$, we get that $\frac{d}{dt}\alpha_t + \mathcal{L}_{X_t}\alpha_t = h_t\alpha_t$.

To prove this claim, note that the map $X \mapsto \imath_X d\alpha$, for a contact form α, is a linear isomorphism from $\ker\alpha$ to the space of 1-forms $\{\beta \mid \beta(R_\alpha) = 0\}$. But $\beta_t = h_t\alpha_t - \alpha_t'$ is in this subspace for each α_t.

Now let ψ_t be the flow generated by X_t starting from $\psi_0 = \mathrm{id}$. We want to show that $\psi_t^*\alpha_t = f_t\alpha_0$ for some positive functions f_t. Certainly this is true when $t = 0$. Our next claim is that $\frac{d}{dt}(\psi_t^*\alpha_t) = g_t\psi_t^*\alpha_t$, where $g_t = h_t \circ \psi_t$. Given this claim, we get an initial value problem at each point in the manifold which can be solved for $\psi_t^*\alpha_t$ to get:

$$\psi_t^*\alpha_t = e^{\int_0^t g_s ds}\psi_0^*\alpha_0 = f_t\alpha_0$$

and $f_t > 0$.

To prove this last claim, we use the "Lie derivative formula for time-dependent vector fields" (see [1] theorem 5.4.5) which says that $\frac{d}{dt}(\psi_t^*\alpha_t) = \psi_t^*(\frac{d}{dt}\alpha_t + \mathcal{L}_{X_t}\alpha_t)$ which, by the first paragraph, is equal to $\psi_t^*(h_t\alpha_t) = g_t\psi_t^*\alpha_t$. □

(Lemma 3.1 is proved by a very similar method.) Note that we make no assertions about how long the flows persist. However if M is closed then they exist for all time and we get:

Corollary 3.4 (Gray's theorem). *Suppose that α_t is a 1-parameter family of contact forms on a closed manifold M. Then there exists an isotopy $\psi_t : M \to M$ and a family of functions $f_t : M \to (0, \infty)$ such that $\alpha_t = \psi_t^*(f_t\alpha_0)$.*

In other words, any isotopy of contact structures on a closed manifold is the pull-back via an ambient isotopy of the initial contact structure.

We also use lemma 3.3 to prove the following:

Corollary 3.5 (Darboux's theorem for contact structures). *Let M_1 and M_2 be $(2n-1)$-manifolds with contact forms α_i and compact submanifolds K_i, suppose that $\phi : T_{K_1}M_1 \to T_{K_2}M_2$ is a bundle isomorphism such that $\phi^*\alpha_2 = f\alpha_1|_{T_{K_1}M_1}$ for some positive function f on K_1. Then there exists a* contactomorphism *Ψ from a neighborhood of K_1 to a neighborhood of K_2 with $\Psi|_{K_1} = \phi|_{K_1}$*

Note that, even if $f = 1$ we are not asserting that $\Psi^*\alpha_2 = \alpha_1$ but simply that $\Psi^*\alpha_2 = F\alpha_1$ for some positive function F.

Proof. Extend ϕ to a diffeomorphism Φ from a neighborhood of K_1 to a neighborhood of K_2 and let $\alpha_0 = \Phi^*\alpha_2$. On a neighborhood of K_1 let $\alpha_t = (1-t)\alpha_0 + t\alpha_1$. Since $\alpha_0 = f\alpha_1$ along K_1 and K_1 is compact, α_t is contact for all $t \in [0,1]$ in a (possibly smaller) neighborhood of K_1. Now apply lemma 3.3 to this neighborhood to get the vector field X_t and flow ψ_t. Since $X_t = 0$ along K_1, ψ_t is the identity on K_1. Since K_1 is compact, ψ_t is defined for t between 0 and 1, after again shrinking the neighborhood if necessary. So $\psi_1^*\alpha_1 = f_1\alpha_0$ and we let $\Psi = \Phi \circ \psi_1^{-1}$. Ψ preserves the co-orientation because it does so along K_1 (assuming each neighborhood of each component of K_1 is connected). □

We mainly use this result in the following form:

Corollary 3.6. *Let (M_1, α_1) and (M_2, α_2) be 3-manifolds with contact forms, either both positive or both negative. We consider three cases, for $i = 1, 2$:*

1. *K_i is a single point in M_i.*

2. *K_i is a Legendrian knot in M_i.*

3. *K_i is a transverse knot in M_i.*

Then in each case there exist neighborhoods ν_i of K_i in M_i and a diffeomorphism $\phi : \nu_1 \to \nu_2$ such that $\phi(K_1) = K_2$ and $\phi^*(\alpha_2) = f\alpha_1$ for some positive function f on ν_1. In case 2, ψ necessarily takes the framing $\mathrm{tb}(K_1)$ *to the framing* $\mathrm{tb}(K_2)$.

Proof. In each case we choose a diffeomorphism from K_1 to K_2 and extend it to a bundle isomorphism from $T_{K_1}M_1$ to $T_{K_2}M_2$ taking the underlying co-oriented contact structure for α_1 to that for α_2. Then apply corollary 3.5. In case 2, ψ takes $\text{tb}(K_1)$ to $\text{tb}(K_2)$ simply because the Thurston-Bennequin faming is canonically determined by the contact structure. □

(We will see later that in case 3 we can take any chosen framing of K_1 to any chosen framing of K_2.)

3.2 Contact forms, contact structures and symplectic germs

If M is a $(2n-1)$-dimensional manifold with contact form α then $\mathbb{R} \times M$ with the form $\omega = d(e^t\alpha)$ (where t is the $\mathbb{R}$-coordinate) is a symplectic $2n$-manifold called the *symplectification* of (M, α). Note that ∂_t is a symplectic dilation and that $\alpha = \imath_{\partial_t}\omega|_{\{0\}\times M}$. More generally, if M is a $(2n-1)$-dimensional submanifold of a symplectic $2n$-manifold (X, ω) with a symplectic dilation V defined in a neighborhood of and positively transverse to M, then $\alpha = \imath_V\omega|_M$ is a positive contact form on M. This is exactly because $d\alpha = \omega|_M$ so $\alpha \wedge (d\alpha)^{n-1} = \imath_V\omega^n|_M > 0$. In this situation, flow along V starting on M gives an embedding of a neighborhood of $\{0\} \times M$ in $\mathbb{R} \times M$ into X which, by the uniqueness of solutions to ordinary differential equations, is in fact a symplectic embedding. This is an instance of the following:

Lemma 3.7. *A positive contact form α on a 3-manifold M defines a unique 4-dimensional symplectic germ $\mathcal{G}(\alpha)$ in the following sense: There exists a symplectic 4-manifold (X, ω) containing M with a symplectic dilation positively transverse to M inducing α, and any other symplectic 4-manifold (X', ω') containing M with the property that $\omega'|_M = d\alpha$ represents the same germ along M.*

Proof. We can take (X, ω) to be the symplectification of (M, α), identifying M with $\{0\} \times M$. Now given (X', ω') with the indicated property, we will construct a bundle isomorphism ψ from T_MX' to T_MX covering the identity and preserving the symplectic forms along M. Then we will apply corollary 3.2. To get this bundle isomorphism we need to construct a vector field V' *along M and transverse to M in X'* (i.e. a section of T_MX' transverse to TM) such that $\imath_{V'}\omega'|_{TM} = \alpha$. Then we construct ψ by sending V' to V. The fact that $\psi^*\omega' = \omega$ follows from the facts that $\omega|_{TM} = \omega'|_{TM}$, that $\imath_V\omega|_{TM} = \imath_{V'}\omega'|_{TM}$ and that V

and V' are transverse to M. To construct V', extend α to a maximal rank 1-form on $T_M X'$; then there exists a unique V' such that $\imath_{V'}\omega' = \alpha$ (this follows from the nondegeneracy of ω). That V' is positively transverse to M comes from the fact that α is a positive contact form with $d\alpha = \omega'|_{TM}$. □

Germs defined by contact forms behave well under pull-backs in the sense that if ϕ is an embedding then $\phi^*\mathcal{G}(\alpha) = \mathcal{G}(\phi^*\alpha)$. Also, for any non-zero c, $\mathcal{G}(c\alpha) = c\mathcal{G}(\alpha)$.

Definition 3.8. *A symplectic bordism from $(M_1, \mathcal{G}_1)$ to $(M_2, \mathcal{G}_2)$ is* convex *if $\mathcal{G}_i = \mathcal{G}(\alpha_i)$ for positive contact forms α_i on M_i, with $\alpha_1 = \alpha_2$ on $M_1 \cap M_2$.*

Notice that the condition that α_1 and α_2 agree on $M_1 \cap M_2$ and lemma 3.7 together guarantee that the result of glueing one convex bordism onto another is again convex. The easiest way to show that a bordism is convex is to demonstrate a single symplectic dilation defined on a whole patch representing the bordism and then to show that it is positively transverse to both boundaries. Another important point is the following: When glueing two convex manifold $\mathcal{P}$ and $\mathcal{Q}$ along an embedding $\psi : \partial_1\mathcal{Q} \hookrightarrow \partial_2\mathcal{P}$, suppose that we have chosen symplectic dilations positively transverse to $\partial_1\mathcal{Q}$ and $\partial_2\mathcal{P}$ and suppose that ψ respects the induced contact forms. Then, after glueing, the two dilations patch together to give a symplectic dilation positively transverse to $\partial_2\mathcal{P} \subset \mathcal{P} \cup_\psi \mathcal{Q}$. (This is by the comment that a symplectic dilation transverse to a codimension 1 submanifold gives a unique embedding of the symplectification of the submanifold with its induced contact form.)

A symplectic manifold (X, ω) with (strongly) convex boundary as described in the introduction, with induced contact form α on ∂X, is thus the support of a convex bordism from the empty manifold to $(\partial X, \mathcal{G}(\alpha))$. If instead ∂X is concave then X is the support of a convex bordism from $(-\partial X, \mathcal{G}(-\alpha))$ to the empty manifold. (Multiply the symplectic contraction by -1 to get a dilation pointing into X.) The convex to concave glueing mentioned in section 1.2 then follows from the fact that symplectic bordisms can be glued symplectically and the fact that contact forms determine unique symplectic germs.

Definition 3.9. *A* product convex bordism *is a product symplectic bordism $\mathcal{P}$ represented by a symplectic patch $(X = \mathbb{R} \times M, X_1, X_2, \omega)$ for which ∂_t is a symplectic dilation (where t is the $\mathbb{R}$-coordinate).*

Given such a bordism, $\mathcal{G}(\partial_i\mathcal{P}) = \mathcal{G}(\alpha_i)$ where $\alpha_i = \imath_{\partial_t}\omega|_{\partial_i\mathcal{P}}$. Identifying M with $M_0 = \partial_1\mathcal{P}$ and letting $\alpha = \alpha_1$, we can see that $\omega = d(e^t\alpha)$ and that $\alpha_2 = e^h\alpha$. Thus we can conclude:

Proposition 3.10. *Given a $(2n-1)$-manifold M with a positive contact form α and a function $h : M \to [0,\infty)$, there exists a unique product convex bordism $\mathcal{B}_h$ of height h constructed on (M,α) with $\mathcal{G}(\partial_1\mathcal{B}_h) = \mathcal{G}(\alpha)$, $\mathcal{G}(\partial_2\mathcal{B}_h) = \mathcal{G}(\alpha_h)$, where $\alpha_h = e^h\alpha$.*

Corollary 3.11. *Suppose we are given (M,α) and subsets $C \subset U \subset M$, with C compact and U open, and a contact form α_1 on U with the same underlying contact structure as α. Then there exists another contact form α_2 on M with the same underlying contact structure such that $\alpha_2 = k\alpha_1$ on C for some constant $k > 0$ and $\alpha_2 = \alpha$ on $M \setminus U$, and there exists a product convex bordism $\mathcal{B}_h$ from $(M,\mathcal{G}(\alpha))$ to $(M,\mathcal{G}(\alpha_2 = \alpha_h))$.*

Proof. Let f be a non-negative function equal to 1 on C and supported inside U. Note that $\alpha_1 = g\alpha$ for some positive function g on U. Choose k so that $kg > 1$ on C, let $h = f\log(kg)$ and build $\mathcal{B}_h$ as in proposition 3.10. □

These product convex bordisms are most useful for the following construction: Suppose that $\mathcal{P}$ is a convex bordism from $(M_1,\mathcal{G}(\alpha_1))$ to $(M_2,\mathcal{G}(\alpha_2))$ and that $\mathcal{Q}$ is a convex bordism from $(N_1,\mathcal{G}(\beta_1))$ to $(N_2,\mathcal{G}(\beta_2))$. We know that we can glue $\mathcal{Q}$ to $\mathcal{P}$ along an imbedding $\psi : N_1 \hookrightarrow M_2$ as long as $\psi^*\mathcal{G}(\alpha_2) = \mathcal{G}(\beta_1)$ and in particular as long as $\psi^*\alpha_2 = \beta_1$. Suppose instead that ψ is just a contactomorphism. Then $\beta_1 = \psi^*(g\alpha_2)$ for some positive function g on $\psi(N_1)$. As long as g extends to an open neighborhood U of a compact set C containing $\psi(N_1)$ then there is still a sense in which we can glue $\mathcal{Q}$ to $\mathcal{P}$ along ψ. Actually we apply corollary 3.11 to get a constant $k > 0$ and a product convex bordism $\mathcal{B}_h$ from $(M_2,\mathcal{G}(\alpha_2))$ to $(M_2,\mathcal{G}(\alpha_2'))$ where $\alpha_2' = kg\alpha_2$ on $\psi(N_1)$, $kg = e^h$ and h is supported on a C and identically 0 outside U. Then glue $\mathcal{B}_h$ to $\mathcal{P}$ along M_2 and glue $k\mathcal{Q}$ to $\mathcal{B}_h$ along ψ.

The result of this construction is topologically the same as glueing $\mathcal{Q}$ to $\mathcal{P}$ but to achieve the glueing symplectically we need to rescale the form on $\mathcal{Q}$ and to "fill the gap" with $\mathcal{B}_h$. We will call this a "sloppy glueing" of $\mathcal{Q}$ onto $\mathcal{P}$. (See figure 3.1.)

3.3 Weinstein's symplectic handles

The handles that Weinstein constructs in [11] are examples of convex symplectic bordisms and can be described concisely using the construction of handles in section 2.2.

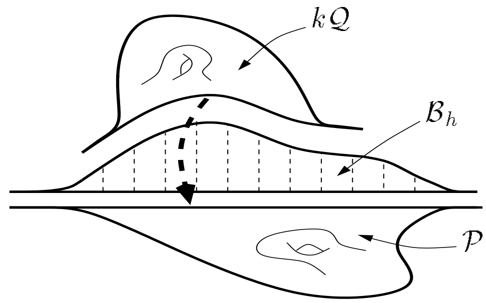

Figure 3.1: Sloppy glueing: filling the gap with a product bordism.

The vector fields V used to construct the handles are symplectic dilations so that the level sets $f^{-1}(\epsilon_1)$ and $f^{-1}(\epsilon_2)$ and the boundaries $\partial_1\mathcal{H}$ and $\partial_2\mathcal{H}$ carry induced contact forms. To understand how to attach these handles to other convex bordisms we need to understand these contact forms.

To build a convex $2n$-dimensional k-handle, start with $\mathbb{R}^{2n}$ with coordinates

$$(x_1, \ldots, x_n, y_1, \ldots, y_n)$$

and the standard symplectic form

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

Note that when $n = 2$, a positively oriented coordinate system is (x_1, y_1, x_2, y_2), making the coordinate system (x_1, x_2, y_1, y_2) negatively oriented. Consider the two radius functions

$$r_1^2 = \sum_{i=1}^{k} x_i^2 \ , \quad r_2^2 = \sum_{i=k+1}^{n} x_i^2 + \sum_{i=1}^{n} y_i^2$$

with the Morse function $f = -r_1^2 + r_2^2$ and the vector field

$$V = \sum_{i=1}^{k}(-x_i\partial_{x_i} + 2y_i\partial_{y_i}) + \sum_{i=k+1}^{n} \frac{1}{2}(x_i\partial_{x_i} + y_i\partial_{y_i}) \ .$$

Check that V is a symplectic dilation and is gradient-like for f. Then, following the notation in section 2.2, after we choose the constants ϵ_1 and ϵ_2, the sets C, C' and U and the function g, we get a convex symplectic k-handle $\mathcal{H}$, since V is positively transverse to both boundaries. Let α_i be the contact forms induced by V on $\partial_i\mathcal{H}$, so that $\mathcal{G}(\partial_i\mathcal{H}) = \mathcal{G}(\alpha_i)$.

It is instructive to consider the case $n = 2$. We have a particular result which will be useful in chapter 7, which could probably be proved by other arguments, but which follows nicely from the existence of these handles. Suppose that Σ is an oriented, n-punctured surface with a proper Morse function $f : \Sigma \to [0, \infty)$ with no critical points of index 2. In other words, there is some $t_0 > 0$ $f^{-1}(t_0, \infty)$ has n components $U_1, \ldots U_n$, each diffeomorphic via a map Φ_i to $\mathbb{R}^2 \setminus \{r^2 \leq t_0\}$ such that $r^2 \circ \Phi_i = f|_{U_i}$. Let $\omega_{\mathbb{R}^2} = dx \wedge dy = r dr \wedge d\theta$, the standard symplectic form on $\mathbb{R}^2$ and let $V_{\mathbb{R}^2} = \frac{1}{2} r \partial_r$, the standard symplectic dilation.

Lemma 3.12. *In this situation, there exists a symplectic form ω on Σ and a symplectic dilation V which is gradient-like for f such that, on each U_i, we have $\Phi_i^* \omega_{\mathbb{R}^2} = \omega$ and $D\Phi_i V = V_{\mathbb{R}^2}$.*

Proof. Build Σ using convex handles with a handle decomposition corresponding to the give Morse function f. We can always attach the 1-handles because a contact form on a 1-manifold is just a nowhere zero 1-form, and we can always make these match up locally. After attaching all the 1-handles we have a bordism from the empty set to n circles, with a symplectic dilation gradient-like for a certain Morse function inducing a contact form on each circle. After reparametrizing this Morse function, we can arrange that the support of this bordism is diffeomorphic to $f^{-1}(-\infty, t_0]$ via a diffeomorphism matching up the Morse functions. We can pick a coordinate θ_i on each boundary component C_i so that the contact form on C_i is $\frac{1}{2} R_i^2 d\theta_i$ for some $R_i^2 > 0$. After rescaling the symplectic form we can arrange that $R_i^2 < t_0$ for each i. Reparametrize the Morse function again so that $f(C_i) = R_i^2$. Now notice that we can attach the annuli $\mathbb{R}^2 \setminus \{r^2 < R_i^2\}$ to the boundary components in such a way that the symplectic forms, Morse functions and symplectic dilations all match up (using the same convex glueing.) ∎

Weinstein describes exactly how much data is needed to determine when one of these convex handle can be attached to another convex bordism in arbitrary dimensions. Rather than present the general argument we will look at the two cases of most interest to us.

Proposition 3.13. *Suppose we are given a 3-manifold M with a positive contact form α and a submanifold K with a compact neighborhood C inside an open neighborhood U.*

1. *If K is a pair of points then there exists a positive function g on U, a convex 4-dimensional 1-handle $\mathcal{H}$ with $\mathcal{G}(\partial_1 \mathcal{H}) = \mathcal{G}(\alpha_1)$ and an embedding $\psi : \partial_1 \mathcal{H} \hookrightarrow C$ such*

that $\psi(K_1) = K$ *and* $\psi^* g\alpha = \alpha_1$.

2. *If* K *is a Legendrian knot then there exists a positive function* g *on* U, *a convex 4-dimensional 2-handle* $\mathcal{H}$ *with* $\mathcal{G}(\partial_1\mathcal{H}) = \mathcal{G}(\alpha_1)$ *and an embedding* $\psi : \partial_1\mathcal{H} \hookrightarrow C$ *such that* $\psi(K_1) = K$ *and* $\psi^* g\alpha = \alpha_1$ *and such that* ψ *takes the handle framing of* K_1 *to the framing* $\mathrm{tb}(K) - 1$.

In both cases the gradient-like vector field V *used in the construction of* $\mathcal{H}$ *is a symplectic dilation inducing* α_1 *on* $\partial_1\mathcal{H}$.

With this we can give the

Proof of theorem 1.1. Realize X in the statement of the theorem as the support of a convex bordism $\mathcal{X}$ with $\partial_1\mathcal{X} = \emptyset$ and $\partial_2\mathcal{X} = \partial X$. By the convexity we have a positive contact form α such that $\mathcal{G}(\alpha) = \mathcal{G}(\partial_2\mathcal{X})$. Y is then the support of a convex bordism $\mathcal{Y}$ constructed by performing a "sloppy glueing" of a convex handle $\mathcal{H}$ constructed by proposition 3.13 onto $\mathcal{X}$. □

Proof of proposition 3.13. We will use the general construction of convex handles outlined above, with $n = 2$ and $k = 1$ or 2. Let α_1 be the contact form induced on $f^{-1}(\epsilon_1)$ by V. In case $k = 1$, K_1 is a pair of points in $f^{-1}(\epsilon_1)$. In case $k = 2$, we will show that K_1 is a Legendrian knot and that the handle framing is $\mathrm{tb}(K_1) - 1$. Then we use corollary 3.6 to get a contactomorphism ψ from a neighborhood ν_1 of K_1 in $f^{-1}(\epsilon_1)$ onto a neighborhood $\nu_0 \subset U$ of K, which takes the handle framing to $\mathrm{tb}(K) - 1$ when $k = 2$. Then build $\mathcal{H}$ so that $\partial_1\mathcal{H}$ is an open tubular neighborhood of K_1 inside $\psi^{-1}(C)$.

In the case $k = 2$, we show that K_1 is Legendrian by direct calculation of α_1 along K_1. Here we have $f = -x_1^2 - x_2^2 + y_1^2 + y_2^2$ and $V = -x_1\partial_{x_1} + 2y_1\partial_{y_1} - x_2\partial_{x_2} + 2y_2\partial_{y_2}$. Also $K_1 = \{y_1 = y_2 = 0\} \cap f^{-1}(\epsilon_1) = \{x_1^2 + x_2^2 = -\epsilon_1, y_1 = y_2 = 0\}$. Since

$$\imath_V\omega = -x_1 dy_1 - 2y_1 dx_1 - x_2 dy_2 - 2y_2 dx_2$$

we have $\alpha_1|_{K_1} = -x_1 dy_1 - x_2 dy_2$. Converting to the coordinates on $f^{-1}(\epsilon_1)$ described in section 2.2, this becomes $\alpha_1|_{K_1} = c(\cos\theta_x dy_1 + \sin\theta_x dy_2) = c(\cos\lambda dx + \sin\lambda dy)$ for some unimportant constant c. Thus the kernel of α_1 makes exactly one full right-handed twist around K_1, so the handle-framing is $\mathrm{tb}(K_1) - 1$. □

One way to think about attaching such a handle is to use corollary 3.11 to get k and h and to glue together the handle $k\mathcal{H}$ and the product bordism $\mathcal{B}_h$ of height h built on $(U, \mathcal{G}(\alpha|_U))$ to create an enlarged handle $\mathcal{H}' = \mathcal{B}_h \cup_\psi k\mathcal{H}$. This handle can then be glued directly onto a bordism $\mathcal{P}$ with $\partial_2\mathcal{P} = (M, \mathcal{G}(\alpha))$. The point is that the symplectic dilations on $\mathcal{B}_h$ and $k\mathcal{H}$ patch together under the glueing to give a symplectic dilation V on $\mathcal{H}'$ such that flow along V gives a diffeomorphism $\phi : U \setminus K \to \partial_2\mathcal{H}' \setminus K_2$. Since ϕ comes from flow along the dilation that induces the contact forms, we know that $\phi^*\alpha_2 = e^H\alpha$ for some function $H : U \setminus K \to [0, \infty)$. Since V is gradient-like inside $\mathcal{H}$ we also know that $\lim_{p\to K} H(p) = \infty$.

This shows us how to understand the "contact surgery" that results from attaching one of these handles to a convex bordism $\mathcal{P}$. The submanifold K is removed from $\partial_2\mathcal{P}$ and a new submanifold K_2 with its neighborhood $\partial_2\mathcal{H}'$ is glued in via an embedding $\Psi = \phi^{-1} : \partial_2\mathcal{H}' \hookrightarrow \partial_2\mathcal{P}\setminus K$. Then the contact form α on $\partial_2\mathcal{P}$ is replaced by $e^H\alpha$, where H is supported inside the image of Ψ. The crucial property of H is that it blows up along K and in fact blows up in exactly the right way so that $\Psi^*(e^H\alpha)$ extends across K_2.

Now we consider briefly the concave version of this construction.

Proof of theorem 1.2. We first need a definition of a concave bordism, and of the germ $\mathcal{G}(\alpha)$ determined by a negative contact form α. This is given by the negative symplectification $(\mathbb{R} \times M, \omega = -d(e^{-t}\alpha))$ and the proof of uniqueness is essentially identical to that for positive germs. There is an obvious sloppy glueing construction in the concave situation. To construct the concave 2-handles in dimension 4 we look at $\mathbb{R}^4$ with the same symplectic form and the same Morse function as in the convex case, but with the gradient-like symplectic *contraction* $V = -2x_1\partial_{x_1} + y_1\partial_{y_1} - 2x_2\partial_{x_2} + y_2\partial_{y_2}$. Then we calculate that $\alpha_1|_{K_1} = \imath_V\omega|_{K_1} = -2x_1dy_1 - 2x_2dy_2$, the kernel of which again makes one full right handed twist around K_1, so that the handle framing is $\mathrm{tb}(K_1) - 1$. □

Chapter 4

Fat Transverse 2-Handles

In this chapter we will prove theorem 1.5 and give an explicit description of the contact surgery involved. First we prove some preliminary lemmas.

4.1 Framings and neighborhoods of transverse knots

Proof of lemma 1.3. Let $\alpha = d\lambda \pm \frac{1}{2}r^2 d\mu$ on $\mathbb{R}^2 \times S^1$, a positive contact form when $\pm = +$ and a negative contact form when $\pm = -$. Note that $\{0\} \times S^1$ is transverse to $\ker\alpha$. Given a contact structure ξ on a 3-manifold M and a transverse knot K, we apply corollary 3.6 to K and $\{0\} \times S^1$ to get a neighborhood ν of K, a radius R and a contactomorphism

$$\psi : (\nu, \xi) \to (\{r < R\} \times S^1, \ker\alpha) \ .$$

This induces coordinates on ν with $R_\nu = R$. Now just note that $\ker\alpha = \operatorname{span}\{\partial_r, \mp\frac{2}{r^2}\partial_\mu + \partial_\lambda\}$ away from $\{r = 0\}$, so we have $s(r^2) = \mp\frac{2}{r^2}$. □

Lemma 4.1. *Given any function $s : (0, \infty) \to \mathbb{R}$ such that $s' > 0$ and $\frac{1}{s} = t$ where t extends smoothly across 0 with $t(0) = 0$ and $t'(0) < 0$, there exists a unique positive contact structure ξ on $\mathbb{R}^2 \times S^1$ with $\xi = \operatorname{span}\{\partial_r, s(r^2)\partial_\mu + \partial_\lambda\}$ away from $\{r = 0\}$. Given two such functions s_1 and s_2 describing contact structures ξ_1 and ξ_2, suppose that R_1 and R_2 are radii such that $s_2(R_2^2) = s_1(R_1^2) + k$, for some integer k. Then there exist positive constants δ_1 and δ_2 and a function $g : [0, R_1 + \delta_1) \to [0, R_2 + \delta_2)$ such that $s_2(g(r)^2) = s_1(r^2) + k$ giving a contactomorphism $\phi : (r, \mu, \lambda) \mapsto (g(r), \mu + k\lambda, \lambda)$ from $(\{r < R_1 + \delta_1\}, \xi_1)$ to $(\{r < R_2 + \delta_2\}, \xi_2)$.*

Proof. To see that ξ is a well-defined contact structure given s and t simply requires verifying that ξ extends smoothly across $\{r = 0\}$ as a positive contact form. But this is true because, near $r = 0$, ξ is the kernel of $d\lambda - t(r^2)d\mu$ which is a positive contact form because $t'(0) < 0$.

For the other half of the lemma, define g by $g(r)^2 = \sqrt{s_2^{-1}(s_1(r^2) + k)}$ and note that s_2^{-1} is defined on an open neighborhood of $(0, s_1(R_1^2) + k]$. Verify that ϕ is a contactomorphism by direct computation. □

Notice that ϕ takes the framing F_μ to the framing $F_\mu + k$. In particular, if a neighborhood ν of a knot K is fat with respect to a framing F then we can always choose coordinates (r, μ, λ) on ν with framing function s such that $F = F_\mu$ and $s(R_\nu^2) > 0$.

To prove proposition 1.6 we will need the following:

Lemma 4.2. *If β_1 and β_2 are 1-forms on a disk D such that $d\beta_i$ is a symplectic form (positive area form) on D then $\alpha_i = d\lambda + \beta_i$ are positive contact forms on $D \times S^1$. If $\phi : (D, d\beta_1) \to (D, d\beta_2)$ is a symplectomorphism then there exists a function $h : D \to \mathbb{R}$ such that the diffeomorphism $\Phi : D \times S^1 \to D \times S^1$ given by $\Phi(p, \lambda) = (p, \lambda + h(p))$ satisfies $\Phi^*\alpha_2 = \alpha_1$.*

Proof. That each α_i is a positive contact form is a straightforward calculation. Since $\phi^* d\beta_2 = d\beta_1$, the 1-form $\beta_1 - \phi^*\beta_2$ is closed and therefore exact. Choose h so that $\beta_1 = \phi^*\beta_2 + dh$. Then $\Phi^*\alpha_2 = d(\Phi^*\lambda) + \phi^*\beta_2 = d\lambda + dh + \phi^*\beta_2 = \alpha_1$. □

Proof of proposition 1.6. Without loss of generality, by corollary 3.6, we may assume that ν has the form $(\nu = D_\epsilon \times S^1, \xi = \ker\alpha)$ where $\alpha = dy - x d\lambda$, x and y are coordinates on the disk $D_\epsilon = \{x^2 + y^2 < \epsilon^2\}$ of radius ϵ, and λ is the S^1-coordinate. This is because $K = \{0\} \times S^1$ is Legendrian. Note that $\mathrm{tb}(K)$ is the "zero-framing" coming from the product structure on ν. We will measure all framings relative to this product framing, so that $\mathrm{tb}(K) - 1 = -1$. On $\{x > 0\}$, $\xi = \ker\alpha'$, where $\alpha' = d\lambda + \beta$ and $\beta = -\frac{1}{x}dy$. The 2-form $d\beta = \frac{1}{x^2}dx \wedge dy$ is a positive symplectic form on $\{x > 0\}$. We will construct a symplectomorphism ϕ from the disk D_2 of radius 2 in $\mathbb{R}^2$ with the standard symplectic form $dx \wedge dy = r dr \wedge d\mu$ onto a region $D \subset \{x > 0, x^2 + y^2 < \epsilon\}$ with the symplectic form $d\beta$. On D_2 let $\beta_2 = \frac{1}{2}r^2 d\mu$ and note that $d\beta_2 = r dr \wedge d\mu$. Thus the lemma above gives a contactomorphism Φ from $(D_2 \times S^1, d\lambda + \beta)$ onto $(D \times S^1, \alpha')$. taking the zero framing to the zero framing. Then simply note that, for the contact form $d\lambda + \beta$, $D_2 \times S^1$ is fat with respect to the framing $-1 = \mathrm{tb}(K) - 1$.

We construct ϕ directly. Choose two positive constants c_1 and c_2, define ϕ by:

$$x \circ \phi = \frac{c_2}{c_1 - x}, \; y \circ \phi = c_2 y$$

and verify that $\phi^* \frac{1}{x^2} dx \wedge dy = dx \wedge dy$. The map is only defined when $x < c_1$, but as long as we choose $c_1 > \sqrt{2}$, ϕ will be defined on D_2. By choosing c_2 small enough we can guarantee that $\phi(D_2) \subset D_\epsilon$. □

4.2 Fat transverse surgeries and 2-handles

Here we describe a certain surgery on a fat neighborhood of a transverse knot and show that it can be realized by attaching symplectic handles.

Suppose ν is a neighborhood of K with coordinates (r, μ, λ) and contact structure $\xi = \mathrm{span}\{\partial_r, s(r^2)\partial_\mu + \partial_\lambda\}$ as in lemma 1.3 such that $F_\mu = F$ and $s(R_\nu^2) > 0$. Let $R < R_\nu$ be the radius for which $s(R^2) = 0$ and choose some positive $\delta < R$. Remove the set $\{r \leq R - \delta\}$ from ν. Let T be a solid torus with coordinates (r, μ, λ) and radius $R_T = R_\nu - R + \delta$. Glue T into $\nu \setminus \{r \leq R - \delta\}$ via a map Ψ defined on $\{r > 0\}$ as follows:

$$\Psi : (r, \mu, \lambda) \mapsto (r + R - \delta, \lambda, -\mu)$$

Letting $\alpha = d\mu - s(r^2)d\lambda$ (so that $\xi = \ker \alpha$), we get that $\Psi^*\alpha = s((r + R - \delta)^2)d\mu + d\lambda$. Let $t(r^2)$ be a function of r^2 which agrees with $s((r + R - \delta)^2)$ outside $\{r \leq \delta + \delta_1\}$ for some positive $\delta_1 < R_\nu - R$, goes to 0 smoothly as r^2 goes to 0 and has positive derivative. Then the 1-form $t(r^2)d\mu + d\lambda$ is a positive contact form which extends smoothly across $\{r = 0\} \subset T$ and agrees with $\Psi^*\alpha$ outside a small neighborhood of $\{r = 0\} \subset T$.

A similar transverse surgery has also been described, in a different context and without 4-dimensional symplectic constructions, in [6].

Recall the definition of weak convexity. We will say that a symplectic bordism $\mathcal{P}$ with symplectic form ω is a "convex-to-weakly-convex" bordism if $\mathcal{G}(\partial_1\mathcal{P}) = \mathcal{G}(\alpha_1)$ for some positive contact form α_1 on $\partial_1\mathcal{P}$ and if there is a positive contact structure ξ_2 on $\partial_2\mathcal{P}$ which agrees with $\ker\alpha_1$ on $\partial_1\mathcal{P} \cap \partial_2\mathcal{P}$ and such that $\omega|_{\xi_2} > 0$. A convex-to-weakly-convex bordism can be glued onto a convex bordism, but in general we do not know how to glue onto a convex-to-weakly-convex bordism.

Proposition 4.3. *Given a 3-manifold M with contact form α, suppose K is a transverse knot with a neighborhood ν with coordinates (r, μ, λ) fat with respect to a given framing F.*

Let R be the radius for which $s(R^2) = F - F_\mu$ and let $A = \{r \leq R\} \subset \nu$. Then there exists a positive function g on ν and a convex-to-weakly-convex 2-handle $\mathcal{H}$, with $\mathcal{G}(\partial_1\mathcal{H}) = \mathcal{G}(\alpha_1)$ for a positive contact form α_1 on $\partial_1\mathcal{H}$ and with an embedding $\psi : \partial_1\mathcal{H} \hookrightarrow \nu$ such that $\psi^* g\alpha = \alpha_1$ and $\psi(K_1) = K$ and such that ψ takes the handle framing to the framing F. We can arrange that $\psi(\partial_1\mathcal{H})$ lie in an arbitrarily small neighborhood of A (but not in an arbitrarily small neighborhood of K). Lastly, the surgery achieved by attaching the handle and extending the contact structure $\ker\alpha_1$ *across the new boundary using ξ_2 is the surgery described above.*

Now we are ready for the

Proof of theorem 1.5. As in the proof of theorem 1.1, let $\mathcal{X}$ be the convex bordism with $\mathrm{supp}(\mathcal{X}) = X$. Then attach the convex-to-weakly-convex handles from proposition 4.3 along L using the "sloppy glueing" construction of section 3.2. □

Proof of proposition 4.3. We construct $\mathcal{H}$ according to the model construction in section 2.2. Recall the usage there of the various radii $R < R_1 < R_2 < R_3$. As advertised, we will use a vector field V for which the angle coordinates θ_1 and θ_2 are invariant under flow along V but which fails to extend across Z^{asc} and is only gradient-like on an open neighborhood of $f^{-1}[\epsilon_1, \epsilon_2] \setminus Z^{asc}$. We will calculate a radius R and show that the map ϕ induced by flow along V is defined as a diffeomorphism

$$\phi : f^{-1}(\epsilon_1) \setminus \{r \leq R\} \to f^{-1}(\epsilon_2) \setminus K_2\ .$$

We will impose the standard symplectic form on $\mathbb{R}^4$ and V will be a symplectic dilation and therefore will induce a contact form α_1 on $f^{-1}(\epsilon_1)$ and a contact form α_2 on $f^{-1}(\epsilon_2) \setminus K_2$. We will show that K_1 is a transverse knot, that $f^{-1}(\epsilon_1)$ is fat with respect to the handle framing, and that given any neighborhood ν of K_1 in $f^{-1}(\epsilon_1)$ which is fat with respect to the handle framing we can construct $\mathcal{H}$ so that $\partial_1\mathcal{H} \subset \nu$. Then we will show that $\partial_2\mathcal{H}$ supports the promised contact structure ξ_2, obtained by modifying $\ker\alpha_2$ so that it extends across K_2, and show that this realizes the surgery described. The rest of the proposition follows from lemma 4.1.

On $\mathbb{R}^4$, using polar coordinates $(r_1, \theta_1, r_2, \theta_2)$, the standard symplectic form becomes $\omega = r_1 dr_1 \wedge d\theta_1 + r_2 dr_2 \wedge d\theta_2$; the Morse function is $f = -r_1^2 + r_2^2$. The symplectic dilation is:

$$V = \frac{1}{2}[(r_1 - \frac{1}{r_1})\partial_{r_1} + r_2\partial_{r_2}]$$

which blows up along $\{r_1 = 0\} = Z^{asc}$. Note that $df(V) = f + 1$ so that V is gradient-like on $f^{-1}(-1,\infty) \setminus Z^{asc}$. Thus we make sure to choose $\epsilon_1 > -1$.

Using the coordinates (r, μ, λ) on $f^{-1}(\epsilon_1)$ and $f^{-1}(\epsilon_2)$ from section 2.2, we calculate that the flow Φ along V starting on $f^{-1}(\epsilon_1)$, as a map from (an open subset of) $\mathbb{R} \times f^{-1}(\epsilon_1)$ into $\mathbb{R}^4$, is given by

$$r_1^2 \circ \Phi = (r^2 - \epsilon_1 - 1)e^t + 1, \quad \theta_1 \circ \Phi^+ = -\lambda$$
$$r_2^2 \circ \Phi^+ = r^2 e^t, \quad \theta_2 \circ \Phi^+ = \mu\ .$$

Note that the flow lines meet $f^{-1}(\epsilon_2)$ when $e^t = \frac{1+\epsilon_2}{1+\epsilon_1} > 1$. Plugging this back in and letting $R = \sqrt{\epsilon_2(\frac{1+\epsilon_1}{1+\epsilon_2})}$, we see that forward flow defines the diffeomorphism $\phi^+ : f^{-1}(\epsilon_1) \setminus \{r \leq R\} \to f^{-1}(\epsilon_2) \setminus K_2$ given by:

$$r^2 \circ \phi = r^2 \frac{1+\epsilon_2}{1+\epsilon_1} - \epsilon_2$$
$$\mu \circ \phi = -\lambda, \quad \lambda \circ \phi = \mu$$

Thus we need to construct $\mathcal{H}$ with $R_3 > R_2 > R_1 > R = \sqrt{\epsilon_2 \frac{1+\epsilon_1}{1+\epsilon_2}}$.

The contact forms induced by V^+ on $f^{-1}(\epsilon_1)$ and $f^{-1}(\epsilon_2) \setminus K_2$ are restrictions of the 1-form $\tilde{\alpha} = \imath_{V^+}(\omega) = \frac{1}{2}[(r_1^2 - 1)d\theta_1 + r_2^2 d\theta_2]$ Restricting to $f^{-1}(\epsilon_1)$ we get $\alpha_1 = \frac{1}{2}[r^2 d\mu - (r^2 - \epsilon_1 - 1)d\lambda]$ which has framing function $s(r^2) = 1 - \frac{1+\epsilon_1}{r^2}$. Thus we see that K_1 is a transverse knot and that a neighborhood $U = \{r < R_3\} \subset f^{-1}(\epsilon_1)$ is fat with respect to the handle framing as long as $R_3 > \sqrt{1+\epsilon_1}$. This and lemma 4.1 tell us that to attach the handle onto a fat neighborhood of the given knot K in ∂X will simply require choosing $R_3 = \sqrt{1+\epsilon_1} + \delta_3$ for some small positive δ_3.

Restricting to $f^{-1}(\epsilon_2)$ we get $\alpha_2 = \frac{1}{2}[(r^2 - 1)d\mu + (r^2 + \epsilon_2)d\lambda]$ which has the same kernel as $\alpha_2' = d\lambda + \frac{r^2-1}{r^2+\epsilon_2}d\mu$. For any small positive δ we claim that we can modify α_2' on $\{r^2 \leq 1 + \delta\}$ to get a new positive contact form α_2'' which extends across K_2, agrees with α_2' on $\{r^2 \geq 1 + \delta\}$, and satisfies $\alpha_2'' \wedge d\alpha_2 > 0$: We let $\alpha_2'' = d\lambda + t(r^2)d\mu$ where $t(r^2)$ goes smoothly to 0 as r^2 goes to zero, has positive derivative, and agrees with $\frac{r^2-1}{r^2+\epsilon_2}$ on $\{r^2 \geq 1 + \delta\}$. Then $t(r^2) < 1$ for all r and so $\alpha_2'' \wedge d\alpha_2 = (1 - t(r^2))rdr \wedge d\mu \wedge d\lambda > 0$. Letting $\xi_2 = \ker \alpha_2''$, the condition $\alpha_2'' \wedge d\alpha_2 > 0$ is the same as the condtion $\omega|_{\xi_2} > 0$, so ξ_2 is the desired contact structure on $\partial_2\mathcal{H}$. Since $\phi(\{r^2 = 1+\epsilon_1\}) = \{r^2 = 1\}$, we find that we just need to pick $R_1 = \sqrt{1+\epsilon_1} + \delta_1$ for some small positive δ_1 to guarantee that this modification of α_2' happens inside $\phi(C = \{r \leq R_1\}) \cup K_2$ and therefore inside $f^{-1}(\epsilon_2) \cap \partial_2\mathcal{H}$. This is compatible with the earlier restriction that $R_1 > R$ since, for any $\epsilon_2 > 0$, $\sqrt{\epsilon_2 \frac{1+\epsilon_1}{1+\epsilon_2}} < \sqrt{1+\epsilon_1}$.

Attaching such a handle realizes the described surgery if we take $T = \partial_2\mathcal{H}$ and $\Psi = \phi^{-1}$. $\square$

Chapter 5

Contact Pairs and Symplectic Bordisms

In the introduction we mentioned a notion of symplectic boundaries which are partially convex and partially concave. In this chapter we express that notion in terms of symplectic bordisms and investigate some important properties and examples.

5.1 Contact pairs

The notion of a contact pair (α^+, α^-) on a 3-manifold M has been defined in the introduction (definition 1.11). Consequences of the definition are that α^+ is a positive contact form while α^- is a negative contact form and that the 2-forms $d\alpha^+$ and $-d\alpha^-$ patch together to define a nowhere-zero 2-form $\gamma = \pm d\alpha^\pm$ on all of M. The first condition in the definition then becomes: $\alpha^\pm \wedge \gamma > 0$.

Henceforth the notation $(M, (\alpha^+, \alpha^-))$, with or without subscripts, will refer to a 3-manifold equipped with a contact pair. We will always refer to the domains of the forms as $M^\pm$ (possibly with subscripts) and we will always take M^0 to be $M^+ \cap M^-$ with $\alpha^0 = \alpha^+ + \alpha^-$ on M^0.

Given $(M, (\alpha^+, \alpha^-))$, we will say that (α^+, α^-) is *purely positive* if $M^+ = M$ and $M^- = \emptyset$; we may also write $\alpha^- = 0$ to indicate this condition. In other words, a positive contact form defined on all of M is the same as a purely positive contact pair on M. The pair is *purely negative* if $M^+ = \emptyset$ while $M^- = M$, in which case we may write $\alpha^+ = 0$. The pair is *pure* if it is purely positive or purely negative. We will say that the pair is *essentially* *positive* if $M^+ = M$, *essentially negative* if $M^- = M$ and *essentially pure* if it is either essentially positive or essentially negative. If (α^+, α^-) is pure or essentially pure we will typically let α refer to whichever of the two 1-forms is defined on all of M. The notation (M, α) will refer to an oriented 3-manifold equipped with a pure contact pair, with either $\alpha^+ = \alpha$ if α is positive or $\alpha^- = \alpha$ if α is negative.

To *extend* a contact pair (α^+, α^-) will mean to enlarge the domain of either 1-form maintaining the contact pair properties. To restrict a contact pair simply involves restricting the domains so that they still cover M. The following lemma tells us one way to extend contact pairs:

Lemma 5.1. *Given $(M, (\alpha^+, \alpha^-))$, the 1-form $\alpha^0 = \alpha^+ + \alpha^-$ on M^0 is closed and satisfies $\alpha^0(R_{\alpha^\pm}) > 1$. Given a* closed *extension α_1^0 of α^0 from M^0 to a larger open set M_1^0 with the property that $\alpha_1^0(R_{\alpha^\pm}) > 1$, there exists a unique contact pair (α_1^+, α_1^-) extending (α^+, α^-) such that $\alpha_1^0 = \alpha_1^+ + \alpha_1^-$.*

Proof. α^0 is closed because $d\alpha^+ = -d\alpha^-$. Let $\gamma = \pm d\alpha^\pm$. Note that for any 1-form δ, $\delta(R_{\alpha^\pm}) = \frac{\delta\wedge\gamma}{\alpha^\pm\wedge\gamma}$. Thus $\alpha^0(R_{\alpha^\pm}) = \frac{\alpha^+\wedge\gamma+\alpha^-\wedge\gamma}{\alpha^\pm\wedge\gamma} > 1$ because $\alpha^\pm \wedge \gamma > 0$.

Given the extension α_1^0, the new 1-forms are necessarily defined by $\alpha_1^\pm = \alpha_1^0 - \alpha^\mp$ on $M^\mp \cap M_1^0$ and $\alpha_1^\pm = \alpha^\pm$ on $M^\pm$. The fact that α_1^0 is closed and that $\alpha^\pm$ is a contact pair implies that $-d\alpha_1^- = d\alpha_1^+ = \pm d\alpha^\pm = \gamma$. Thus the fact that $\alpha_1^0(\alpha^\pm) > 1$ implies that $\alpha_1^\pm \wedge \gamma > 0$. Therefore (α_1^+, α_1^-) is a contact pair. □

In this proof we have also seen that, in order to check that a certain 1-form δ has the property $\delta(R_{\alpha^\pm}) > 0$ we can equivalently check that $\delta \wedge \gamma > 0$. This will be useful in chapter 7.

5.2 Contact pairs and symplectic germs

In the introduction we also defined the notion of a dilation-contraction pair (definition 1.9) and pointed out that a dilation-contraction pair on a 4-manifold X transversely covering a 3-dimensional submanifold M induces a contact pair on M. To prove the glueing result mentioned in the introduction for boundaries which are "partially convex and partially concave" and to prove some other results we start with a lengthy lemma. The uniqueness in this lemma gives the sense in which the condition that $\omega(V^+, V^-) = 0$ in the definition of "dilation-contraction pair" gives the necessary rigidity.

Lemma 5.2. *Given $(M,(\alpha^+,\alpha^-))$, consider the two symplectic manifolds (S^+,ω^+) and (S^-,ω^-), where $S^\pm=\mathbb{R}\times M^\pm$ and $\omega^\pm=\pm d(e^{\pm t}\alpha^\pm)$, and identify $M^\pm$ with $\{0\}\times M^\pm\subset S^\pm$. First of all, if (X,ω) is another symplectic manifold containing M^+ or M^- with a single symplectic dilation V^+ or contraction V^- positively transverse to $M^\pm$ inducing the contact form α^+ or α^-, respectively, then flow along $V^\pm$ starting from $M^\pm$ gives an imbedding Φ of an open subset of S^+ or S^-, respectively, into X such that $\Phi^*\omega=\omega^\pm$ and $D\Phi(\partial_t)=V^\pm$.*

Secondly:

1. *There exists a unique vector field V^- defined on $\mathbb{R}\times M^0$ such that the pair $(V^+=\partial_t,V^-)$ is a dilation-contraction pair on (S^+,ω^+) inducing the contact pair $(\alpha^+,\alpha^-|_{M^0})$ on M^+.*

2. *There exists a unique vector field V^+ defined on $\mathbb{R}\times M^0$ such that the pair $(V^+,V^-=\partial_t)$ is a dilation-contraction pair on (S^-,ω^-) inducing the contact pair $(\alpha^+|_{M^0},\alpha^-)$ on M^-.*

In fact, let $\gamma=\pm d\alpha^\pm$, let $g^\pm=\frac{\alpha^0\wedge\gamma}{\alpha^\pm\wedge\gamma}=\alpha^0(R_{\alpha^\pm})$ and let $\beta^\pm=\alpha^0-g^\pm\alpha^\pm$. Then there exist unique vector fields $Z^\pm\in\ker\alpha^+\cap\ker\alpha^-$ on M^0 such that $\imath_{Z^\pm}(\gamma)=\beta^\pm$ and, in the two cases above we have, respectively:

1. $V^-=(g^+e^{-t}-1)\partial_t+e^{-t}Z^+$

2. $V^+=(g^-e^{t}-1)\partial_t+e^{t}Z^-$

We will call the symplectic manifold (S^+,ω^+) with its dilation-contaction pair the *positive symplectification* of $(M^+,(\alpha^+,\alpha^-|_{M^0}))$ and we will call (S^-,ω^-) with its dilation-contraction pair the *negative symplectification* of $(M^-,(\alpha^+|_{M^0},\alpha^-))$.

Proof. The first result follows from the fact that $\Phi^*\omega$ and $\omega^\pm$ are solutions to the same ordinary differential equations with the same initial conditions.

For the second result, we will show the existence and uniqueness of the vector fields $Z^\pm$, then show that $V^\pm$ satisfy the conditions and then show uniqueness.

There exists a unique $Z^\pm\in\ker\alpha^\pm$ such that $\imath_{Z^\pm}(\gamma)=\beta^\pm$ because contraction with γ gives a linear isomorphism from $\ker\alpha^\pm$ to $\{\beta\mid\beta\wedge\gamma=0\}$ (this depends on working in dimension 3), and $\beta^\pm$ is constructed to be in this latter subspace. But $Z^\pm$ is also in $\ker\alpha^0$ because $0=\gamma(Z^\pm,Z^\pm)=\beta^\pm(Z^\pm)=\alpha^0(Z^\pm)$, and thus $Z^\pm\in\ker\alpha^\mp$.

On S^+, letting $V^+ = \partial_t$ and $V^- = (g^+e^{-t} - 1)\partial_t + e^{-t}Z^+$, we need to show that

$$\mathcal{L}_{V^\pm}(\omega^+) = \pm\omega^+ \tag{5.1}$$

$$\imath_{V^\pm}(\omega^+)|_{t=0} = \alpha^\pm \tag{5.2}$$

$$\omega^+(V^+, V^-) = 0 \tag{5.3}$$

First note that $\omega^+ = e^t(dt\wedge\alpha^+ + \gamma)$. Equation 5.3 is quick: $\omega^+(V^+, V^-) = e^{-t}\omega^+(\partial_t, Z^+) = \alpha^+(Z^+) = 0$. To show equation 5.1 and equation 5.2, note that $\imath_{\partial_t}(\omega^+) = e^t\alpha^+$ and that

$$\begin{aligned}\imath_{V^-}(\omega^+) &= (g^+e^{-t} - 1)\imath_{\partial_t}(\omega^+) + e^{-t}\imath_{Z^+}(\omega^+)\\ &= e^t[(g^+e^{-t} - 1)\alpha^+ + e^{-t}\beta^+]\\ &= -e^t\alpha^+ + \alpha^0\end{aligned}$$

Next we will prove that V^- is the unique vector field on $\mathbb{R} \times M^0$ satisfying these equations. Suppose V_0^- and V_1^- are two solutions. Let $\delta = \imath_{V_1^- - V_0^-}(\omega^+)$; we will prove that $\delta = 0$ and thus conclude that $V_0^- = V_1^-$. Equation 5.1 implies that $d\delta = 0$, equation 5.2 implies that $\delta|_{\{t=0\}} = 0$ and equation 5.3 implies that $\delta(\partial_t) = 0$ everywhere. Thus δ is invariant in the t direction and vanishes when $t = 0$, so $\delta = 0$ everywhere.

On S^- the argument is a mirror image of the argument for S^+. □

Note that the uniqueness argument first proved uniqueness *along* M (since $\delta|_{\{t=0\}} = 0$) and then proved uniqueness for all t. Thus in fact we have also proved:

Lemma 5.3. *In $(\mathbb{R}\times M^0, \omega^+)$ there exists a unique vector field V^- along M^0 (i.e. a section of $T_{M^0}(\mathbb{R} \times M^0)$) positively transverse to M^0 such that $\imath_{V^-}\omega^+ = \alpha^-$ and $\omega^+(\partial_t, V^-) = 0$. Likewise there exists a unique vector field V^+ along M^0 in $(\mathbb{R}\times M^0, \omega^-)$ positively transverse to M^0 such that $\imath_{V^+}\omega^- = \alpha^+$ and $\omega^-(V^+, \partial_t) = 0$.*

We will make use of the following observation for later constructions:

Lemma 5.4. *Given $(M, (\alpha^+, \alpha^-))$, let $(S^\pm, \omega^\pm)$ be the positive and negative symplectifications with the dilation-contraction pairs from lemma 5.2. Let $g^\pm$, $\beta^\pm$ and $Z^\pm$ be as in lemma 5.2. Given a function $h : M^\pm \to \mathbb{R}$, consider its graph $M_h^\pm = \{(h(p), p)\} \subset S^\pm$. Then $V^\pm = \partial_t$ is automatically positively transverse to $M_h^\pm$ and $V^\mp$ is positively transverse to $M_h^\pm$ if and only if $e^{\pm h} < g^\pm - dh(Z^\pm)$ on M^0. The induced contact pair on $M_h^\pm$ is then given by*

$$\alpha_h^\pm = e^h\alpha^\pm,\ \alpha_h^0 = \alpha^0,\ \alpha_h^\mp = \alpha^0 - e^h\alpha^\pm\ .$$

Proof. Everything follows from the explicit expressions $V^{\mp} = (g^{\pm}e^{\mp t} - 1)\partial_t + e^{\mp t}Z^{\pm}$ and $\omega^{\pm} = e^{\pm t}(dt \wedge \alpha^{\pm} + \gamma)$. □

Theorem 5.5. *A contact pair (α^+, α^-) on a 3-manifold M defines a unique 4-dimensional symplectic germ $\mathcal{G}(\alpha^+, \alpha^-)$ along M in the following sense: There exists a symplectic 4-manifold (X, ω) containing M with a dilation-contraction pair transversely covering M inducing (α^+, α^-), and any other symplectic 4-manifold (X_1, ω_1) containing M with the property that $\omega_1|_M = \pm d\alpha^{\pm}$ represents the same germ along M.*

Proof. We will prove existence of (X, ω) with its dilation-contraction pair (V^+, V^-) first. Construct $(S^{\pm}, \omega^{\pm})$ as in lemma 5.2. Let U^+ be an open neighborhood of M^0 in S^+ such that forward and backward flow along V^+ in S^- starting from $p_0 \in M^0$ is defined for all times t with $(t, p_0) \in U^+$. Thus there exists an embedding $\phi^+ : U^+ \hookrightarrow S^-$ such that $\phi^+|_{M^0} = \mathrm{id}$ and $D\phi^+(\partial_t) = V^+$. Since both ∂_t and V^+ induce the same contact forms on M^0 and are both symplectic dilations, we can conclude that $(\phi^+)^*(\omega^-) = \omega^+$. By the uniqueness in the lemma we also know that $D\phi^+(V^-) = \partial_t$, so that ϕ^+ preserves all the relevant structure. Let $U^- = \phi^+(U^+)$ and $\phi^- = (\phi^+)^{-1}$.

Now choose two functions $f^{\pm} : M \to [0, \infty)$ such that $f^{\pm}|_{M^{\mp} \setminus M^0} = 0$ but $f^+ + f^- > 0$ everywhere, let $F^{\pm} = \{(t,p) \mid -f^{\pm}(p) < t < f^{\pm}(p)\} \subset S^{\pm}$ and let $E^{\pm} = F^{\pm} \cap \phi^{\mp}(F^{\mp} \cap U^{\mp})$. Finally let $\psi^{\pm} = \phi^{\pm}|_{E^{\pm}} : E^{\pm} \to E^{\mp}$. If we choose $f^{\pm}$ small enough we can guarantee that

$$X = F^+ \cup_{\psi^+} F^-$$

is Hausdorff. Since $(\psi^+)^*\omega^- = \omega^+$ and $D\psi^+(V^{\pm}) = V^{\pm}$, we know that the symplectic forms and the dilation-contraction pairs patch together define a symplectic form ω on X with a dilation-contraction pair (V^+, V^-) transversely covering M inducing (α^+, α^-). (See figure 5.1.)

For the uniqueness result, we need to construct a bundle isomorphism $\psi : T_M X_1 \to T_M X$ covering the identity and preserving the symplectic forms, and then we apply corollary 3.2. To do this we construct a pair of vector fields (V_1^+, V_1^-) along M in $T_M X_1$ with open domains $M_1^{\pm} \subset M^{\pm}$ covering M, both positively transverse to M, such that $\imath_{V_1^{\pm}}\omega_1|_{M_1^{\pm}} = \alpha^{\pm}|_{M_1^{\pm}}$ and such that $\omega_1(V_1^+, V_1^-) = 0$. Then there exists a unique ψ sending $V_1^{\pm}$ to $V^{\pm}$ by lemma 5.3. To see the existence of the pair (V_1^+, V_1^-), first extend α^+ to a maximal rank 1-form on $T_{M^+}X_1$ to get V_0^+ such that $\imath_{V_0^+}\omega_1|_{M^+} = \alpha^+$ (as in

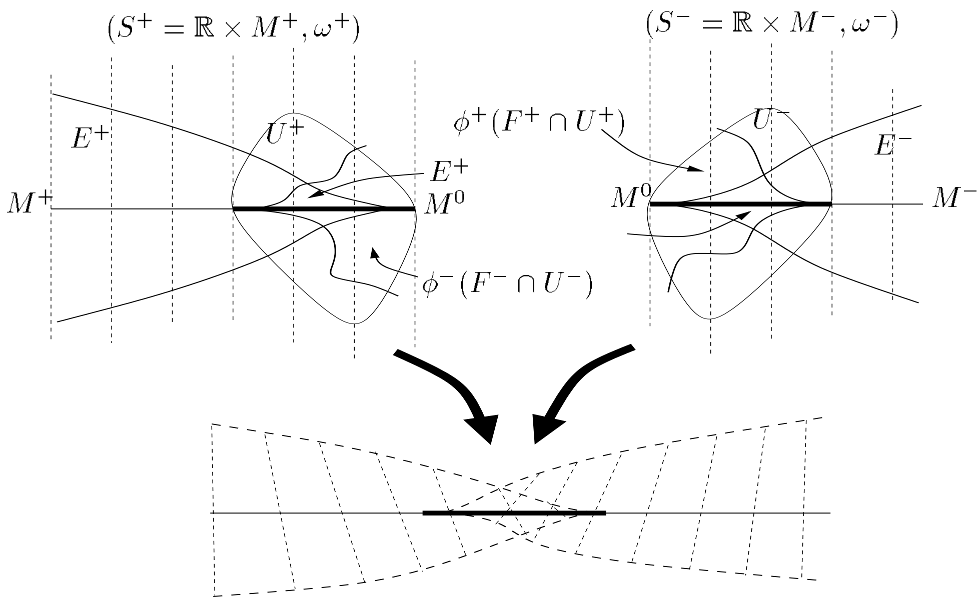

Figure 5.1: Glueing S^+ to S^- after some trimming.

the proof of lemma 3.7). Then by lemma 5.3 there exists a unique V_0^- along M^0 such that $\imath_{V_0^-}\omega_1|_{M^0} = \alpha^-|_{M^0}$ and $\omega_1(V_0^+, V_0^-) = 0$. Let $\tilde{\alpha}_0^- = \imath_{V_0^-}\omega_1$, a maximal rank 1-form on $T_{M^0}X_1$ which extends $\alpha^-|_{M^0}$. Then there exists a maximal rank extension $\tilde{\alpha}_1^-$ of α^- to $T_{M^-}X_1$ which agrees with $\tilde{\alpha}_0^-$ outside a closed set C inside M^- containing $M^- \setminus M^0$ in its interior. Let V_1^- be the corresponding vector field such that $\imath_{V_1^-}\omega_1 = \tilde{\alpha}_1^-$ and let $V_1^+ = V_0^+|_{M^+\setminus C}$.

□

Corollary 5.6. *If* (α_1^+, α_1^-) *and* (α_2^+, α_2^-) *are two different contact pairs on* M *with* $\pm d\alpha_1^\pm = \pm d\alpha_2^\pm$ *then* $\mathcal{G}(\alpha_1^+, \alpha_1^-) = \mathcal{G}(\alpha_2^+, \alpha_2^-)$. *In particular, we have this result if we can get from* (α_1^+, α_1^-) *to* (α_2^+, α_2^-) *by a sequence of extensions and restrictions.*

Note that, for a purely positive pair $(\alpha, 0)$, $\mathcal{G}(\alpha, 0) = \mathcal{G}(\alpha)$. For a purely negative pair $(0, \alpha)$, $\mathcal{G}(0, \alpha) = \mathcal{G}(\alpha)$ as a germ along $-M$.

5.3 Contact pairs and symplectic bordisms

Definition 5.7. *A 4-dimensional symplectic bordism from* $(M_1, \mathcal{G}_1)$ *to* $(M_2, \mathcal{G}_2)$ *is* convex-concave *if* $\mathcal{G}_i = \mathcal{G}(\alpha_i^+, \alpha_i^-)$ *for contact pairs* (α_i^+, α_i^-) *on* M_i, *with* $(\alpha_1^+, \alpha_1^-) = (\alpha_2^+, \alpha_2^-)$ *on*

$M_1 \cap M_2$.

By theorem 5.5, we can glue convex-concave bordisms together to make new convex-concave bordisms.

Generally when we work with convex-concave bordisms there will be an understood contact pair on each boundary, and we will label the domains of the respective forms $\partial_i^{\pm}\mathcal{B}$.

A symplectic manifold (X, ω) with boundary transversely covered by a dilation-contraction pair, with induced contact pair (α^+, α^-) on ∂X, is the support of a convex-concave bordism from the empty manifold to $(\partial X, \mathcal{G}(\alpha^+, \alpha^-))$. The glueing mentioned in the introduction then follows from theorem 5.5.

Definition 5.8. *A* product convex-concave bordism *is a product symplectic bordism $\mathcal{B}$ represented by a symplectic patch $(X = \mathbb{R} \times M, X_1, X_2, \omega)$ with a dilation-contraction pair (V^+, V^-) transversely covering both boundaries, with both vector fields defined on all of* $\mathrm{supp}(\mathcal{B})$*, and with either $V^+ = \partial_t$ or $V^- = \partial_t$.*

The following result then follows quickly from lemma 5.4 and the construction of (X, ω) in theorem 5.5. (Note that in the construction of (X, ω), if we are given one of the functions $f^{\pm}$ with the property that $f^{\pm}|_{M^{\mp}\setminus M^0} = 0$ and $f^{\pm}|_{M^{\pm}\setminus M^0} > 0$ then we can always choose the other function $f^{\mp}$ to make the construction work.)

Proposition 5.9. *Given $(M, (\alpha^+, \alpha^-))$, let $g^{\pm}$, $\beta^{\pm}$ and $Z^{\pm}$ be defined as in lemma 5.2. If $h : M \to [0, \infty)$ is a function with $h|_{M^-\setminus M^0} = 0$ and $e^h < g^+ - dh(Z^+)$ on M^0, then there exists a unique product convex-concave bordism $\mathcal{B}_h$ of height h constructed on M for which $\partial_t = V^+$ on* $\mathrm{supp}(\mathcal{B})$ *and such that the dilation-contraction pair induces (α^+, α^-) on M. If instead $h|_{M^+\setminus M^0} = 0$ and $e^{-h} < g^- - dh(Z^-)$ then there exists a unique $\mathcal{B}_h$ for which $\partial_t = V^-$. The contact pair on $\partial_2\mathcal{B}_h$ is given as in lemma 5.4.*

Chapter 6

Skinny Transverse 2-Handles

In this chapter we will prove theorem 1.8 along with more general results underlying the theorem. Theorem 1.8 is really a statement about the existence of special convex-concave bordisms and follows from the existence of certain convex-concave 2-handles and carefully constructed product convex-concave bordisms. We will use a "sloppy glueing" similar to that in section 3.2 except that product convex-concave bordisms are a little more delicate to construct than product convex bordisms. We begin by investigating control on the behavior of contact pairs near knots.

6.1 Well-behaved transverse knots

Throughout this section let ν be a neighborhood of a knot K with a contact pair (α^+, α^-) and polar coordinates (r, μ, λ). Let R_ν be the radius of ν, let $\nu^\pm$ be the domain of $\alpha^\pm$ and let $\nu^0 = \nu^+ \cap \nu^-$. We will be interested in two cases where $\nu^0 = \nu \setminus K$:

- The *positive case* is the case where $\nu^+ = \nu$ while $\nu^- = \nu \setminus K$; in this case the contact pair is essentially positive, and we let $\alpha = \alpha^+$.
- The *negative case* is the case where $\nu^- = \nu$ while $\nu^+ = \nu \setminus K$; in this case the contact pair is essentially negative, and we let $\alpha = \alpha^-$.

As usual we let $\alpha^0 = \alpha^+ + \alpha^-$ on ν^0. Recall the definition of "well-behaved" for contact forms and closed 1-forms in the introduction (definition 1.7).

Definition 6.1. *Such a contact pair* (α^+, α^-) *is* well-behaved *with respect to* (r, μ, λ) *if* α *and* α^0 *are both well-behaved with respect to* (r, μ, λ)*. (See definition 1.7.)*

In other words, associated to a well-behaved contact pair (α^+, α^-) and the coordinates (r, μ, λ) is a quadruple of real numbers (A, B, C, D) such that:

$$R_\alpha = A\partial_\mu + B\partial_\lambda \quad \text{and} \quad \alpha^0 = C d\mu + D d\lambda$$

and subject to the following constraints:

$$B > 0, \quad C > 0, \quad AC + BD > 1 \ .$$

(This last condition is simply the condition that $\alpha^0(R_\alpha) > 1$.) We will call the data (A, B, C, D) the *structural data* associated to (α^+, α^-) and (r, μ, λ). Any quadruple satisfying these constraints will be called a well-behaved quadruple.

Lemma 6.2. *Given any well-behaved quadruple (A, B, C, D) there exists an essentially positive (resp. negative) well-behaved contact pair (α^+, α^-) on $\nu = \mathbb{R}^2 \times S^1$ with coordinates (r, μ, λ) with these structural data. Given any other essentially positive (resp. negative) contact pair (α_1^+, α_1^-) on a neighborhood ν_1 of a knot K well-behaved with respect to coordinates (r_1, μ_1, λ_1) with these same structural data, there exists an embedding $\phi : \nu_1 \hookrightarrow \nu$ of the form $(r_1, \mu_1, \lambda_1) \mapsto (g(r_1), \mu, \lambda)$ for some g such that $\phi^*(\alpha^+, \alpha^-) = (\alpha_1^+, \alpha_1^-)$.*

Proof. In the positive case let $\pm = +$ and let $\mp = -$. In the negative case let $\pm = -$ and let $\mp = +$. With this notation, let $\alpha^\pm = \pm r^2(d\mu - \frac{A}{B}d\lambda) + \frac{1}{B}d\lambda$, let $\alpha^0 = C d\mu + D d\lambda$ and let $\alpha^\mp = \alpha^0 - \alpha^\pm$. This contact pair is well-behaved with the required structural data. Given (α_1^+, α_1^-) with the same structural data, note that, since (α_1^+, α_1^-) is well-behaved, $\alpha_1^\pm$ is completely determined by A, B and the framing function s (defined so that $\ker \alpha_1^\pm$ is spanned by ∂_r and $s(r^2)\partial_\mu + \partial_\lambda$). Also note that $t(r^2) = \frac{1}{s(r^2)}$ extends smoothly across $\{r^2 = 0\}$ with $t(0) = 0$, that $\pm s' > 0$ and that $s(r^2) \neq \frac{A}{B}$. Thus there exists some g such that $s(r^2) = \frac{A}{B} \pm \frac{1}{Bg(r^2)}$, and we define ϕ as in the statement of the lemma. □

If a contact pair is well-behaved with respect to (r, μ, λ) then it is also well-behaved with respect to $(r, \mu - k\lambda, \lambda)$ for any integer k. Since $F_{\mu - k\lambda} = F_\mu + k$, this means that if a contact pair is well-behaved with respect to coordinates realizing one framing then it is well-behaved with respect to coordinates realizing any other framing.

6.2 Description of the handles and associated surgeries

Definition 6.3. *A well-behaved contact pair (α^+, α^-) on ν is* prepared for surgery *with respect to (r, μ, λ) if the associated structural data (A, B, C, D) satisfy the following additional* *constraints:*

$$A = B, \quad D > 0$$

and:

$$AD > 1 \quad \text{in the positive case or} \quad AD < 1 \quad \text{in the negative case.}$$

We will also say that K is prepared for surgery with respect to (α^+, α^-).

Note that when (α^+, α^-) is prepared for surgery, the integral curves of R_α are closed curves linking K and realizing the framing $F_\mu + 1$ of K. This means that, in contrast to the well-behaved condition, there is a specific framing associated to a contact pair which is prepared for surgery.

Proposition 6.4. *Given $(M, (\alpha^+, \alpha^-))$ with a knot $K \subset M$ which is prepared for surgery with respect to coordinates (r, μ, λ) on a neighborhood ν, there exists a convex-concave 2-handle $\mathcal{H}$ with $\mathcal{G}(\partial_i \mathcal{H}) = \mathcal{G}(\alpha_i^+, \alpha_i^-)$ and an embedding $\psi : \partial_1 \mathcal{H} \hookrightarrow \nu$ such that $\psi(K) = K_1$, ψ takes the handle framing to F_μ and $\psi^*(\alpha^+, \alpha^-) = (\alpha_1^+, \alpha_1^-)$. In the positive case we have*

$$\partial_1^+ \mathcal{H} = \partial_1 \mathcal{H}, \quad \partial_1^- \mathcal{H} = \partial_1 \mathcal{H} \setminus K_1, \quad \partial_2^+ \mathcal{H} = \partial_2 \mathcal{H} \setminus K_2, \quad \partial_2^- \mathcal{H} = \partial_2 \mathcal{H},$$

while in the negative case we have

$$\partial_1^- \mathcal{H} = \partial_1 \mathcal{H}, \quad \partial_1^+ \mathcal{H} = \partial_1 \mathcal{H} \setminus K_1, \quad \partial_2^- \mathcal{H} = \partial_2 \mathcal{H} \setminus K_2, \quad \partial_2^+ \mathcal{H} = \partial_2 \mathcal{H} \ .$$

Furthermore there exists a positive $R < R_\nu$, a nonincreasing function $h : [R^2, R_\nu^2] \to [0, \infty)$ and a diffeomorphism $\phi : \partial_1 \mathcal{H} \setminus \{r \leq \delta\} \to \partial_2 \mathcal{H} \setminus K_2$ such that, letting $\Phi = \phi \circ \psi^{-1}$, we have the following properties:

- *In the positive case:*

$$\Phi^* \alpha_2^+ = e^{h(r^2)} \alpha_1^+, \quad \Phi^* \alpha_2^0 = \alpha_1^0, \quad \Phi^* \alpha_2^- = \alpha_1^0 - e^{h(r^2)} \alpha_1^+ \ .$$

- *In the negative case:*

$$\Phi^* \alpha_2^- = e^{-h(r^2)} \alpha_1^-, \quad \Phi^* \alpha_2^0 = \alpha_1^0, \quad \Phi^* \alpha_2^+ = \alpha_1^0 - e^{-h(r^2)} \alpha_1^- \ .$$

We will use this result in the positive case to prove theorem 1.8. Such a handle can be attached along a knot with a given framing as long as the contact pair near the knot is prepared for surgery with respect to coordinates realizing that given framing. We call these "skinny" 2-handles because the neighborhood ν can be chosen arbitrarily small.

Then we can describe the "contact-pair surgery" that results as follows: We start with a knot K in a neighborhood ν with coordinates (r, μ, λ) and a contact pair (α^+, α^-). A small neighborhood of radius R of K is removed, the contact form α that did extend across K is replaced by $e^{\pm h}\alpha$ for some bump function h and α^0 is kept constant. The other contact form is replaced by $\alpha^0 - e^{\pm h}\alpha$. Now a solid torus T with coordinates (r, μ, λ) is glued in via a map $\phi : T \setminus \{r = 0\} \to \nu \setminus \{r \leq R\}$ with $\mu \circ \phi = \lambda$ and $\lambda \circ \phi = -\mu$. Then we notice that h was chosen carefully enough so that $\alpha^0 - e^{\pm h}\alpha$ extends across $\{r = 0\} \subset T$. In fact, we will see from the proof that h is constant near $\{r = R\}$.

6.3 Construction of the handles

Proof of proposition 6.4. We will parallel the construction of 2-handles in the proof of proposition 4.3 as much as possible except that now we will have two vector fields to keep track of.

Consider $\mathbb{R}^4$ with polar coordinates $(r_1, \theta_1, r_2, \theta_2)$, with the symplectic form $\omega = r_1 dr_1 \wedge d\theta_1 + r_2 dr_2 \wedge d\theta_2$ and with the Morse function $f = -r_1^2 + r_2^2$. We will construct a dilation-contraction pair (V^+, V^-) on this symplectic manifold which is gradient-like for f as long as $|f|$ is small enough. In the positive case, V^+ will blow up along Z^{asc} but be defined everywhere else and V^- will blow up along Z^{desc} but be defined everywhere else. In the negative case the roles will be reversed. Then we will construct $\mathcal{H}$ as in the proof of proposition 4.3, taking care to arrange that $\partial_2\mathcal{H}$ is transverse to both vector fields. We will then calculate the contact pair induced on $\partial_1\mathcal{H}$ and show that it is prepared for surgery, that supp(H) can be made arbitrarily small, and that our construction can realize every set of constants $(A = B, C, D)$ meeting the "prepared for surgery" criteria. Lemma 6.2 will provide the embedding ψ and the diffeomorphism ϕ will be given by flow along either V^+, in the positive case, or V^-, in the negative case.

More specifically, we will compute $\alpha_1^{\pm} = \imath_{V^\pm}\omega|_{f^{-1}(\epsilon_1)}$ in terms of the coordinates (r, μ, λ) on $f^{-1}(\epsilon_1)$ described in section 2.2. (In fact we will compute α_1^0 and either α_1^+ or α_1^-). Then we will describe the flow along either V^+ or V^- from $f^{-1}(\epsilon_1)$ as an imbedding

Φ of (some subset of) $\mathbb{R} \times f^{-1}(\epsilon_1)$ into $\mathbb{R}^4$ sending $\{0\} \times f^{-1}(\epsilon_1)$ to $f^{-1}(\epsilon_1)$ and the vector field ∂_t to either V^+ or V^-. We will show that, given a small enough $\epsilon_2 > 0$, there exists a radius R and a positive constant T such that Φ imbeds all of

$$B = \{(t,p) \in \mathbb{R} \times f^{-1}(\epsilon_1) \mid r(p) > R, 0 \leq t \leq T\}$$

into $\mathbb{R}^4$ sending the subset $\{(T,p) \in B\}$ onto $f^{-1}(\epsilon_2) \setminus K_2$. We will also show that we can make R arbitrarily small by choosing ϵ_2 small enough. We then choose the three radii $R_3 > R_2 > R_1 > R$ and construct $\mathcal{H}$ with $U = \{r < R_3\}$, $C' = \{r \leq R_2\}$ and $C = \{r \leq R_1\}$.

The construction involves a choice of the height function $h : f^{-1}(\epsilon_1) \setminus \{r \leq R\} \to [0,T]$. We will choose h to be a nonincreasing function of r^2 with $\operatorname{supp}(h) \subset \{r^2 \leq R_2^2\}$ and $h \equiv T$ on $\{r^2 \leq R_1^2\}$. In the positive case we will show that, if R_2 is chosen small enough, then for any such h, both vector fields will be transverse to both boundary components. In the negative case we will show that for any R_2 and any such h the required transversality will hold.

We now assume we are given $(A = B, C, D)$ and look at the two cases separately.

In the positive case, define (V^+, V^-) as follows:

$$V^+ = (\frac{1}{2}r_1 - \frac{C}{r_1})\partial_{r_1} + \frac{1}{2}r_2\partial_{r_2}$$
$$V^- = -\frac{1}{2}r_1\partial_{r_1} - (\frac{1}{2}r_2 - \frac{D}{r_2})\partial_{r_2}$$

and calculate that (V^+, V^-) is a dilation-contraction pair with:

$$df(V^+) = f + 2D, \quad df(V^-) = -f + 2C$$

Thus V^+ and V^- are both gradient-like when $-2D < f < 2C$. Let $\epsilon_1 = \frac{2}{A} - 2D$ and note that $-2D < \epsilon_1 < 0$ (because $AD > 1$ and $A > 0$). Choose any ϵ_2 with $0 < \epsilon_2 < 2C$. We compute the contact pair induced by (V^+, V^-) on $f^{-1}(\epsilon_1)$ using the coordinates (r, μ, λ) to be given by:

$$\alpha_1^+ = \frac{1}{2}[r^2(d\mu - d\lambda)] + \frac{1}{A}d\lambda, \quad \alpha_1^0 = C d\mu + D d\lambda$$

and note that $R_{\alpha_1^+} = A(\partial_\mu + \partial_\lambda)$.

We calculate the forward flow along V^+ as a map Φ from some subset of $\mathbb{R} \times f^{-1}(\epsilon_1)$ into $\mathbb{R}^4$. The angle coordinates (θ_1, θ_2) are invariant under this flow and the radial coordinates (r_1, r_2) obey the following equations:

$$r_1^2(t) = (r^2 - \frac{2}{A})e^t + 2D, \quad r_2^2(t) = r^2 e^t$$

We solve for t when $f = \epsilon_2$ to get that the flow lines encounter $f^{-1}(\epsilon_2)$ when

$$t = T = \log(A(D + \frac{\epsilon_2}{2}))$$

($T > 0$ because $AD > 1$) and thus that the diffeomorphism ϕ is given by:

$$r^2 \circ \phi = A(D + \frac{\epsilon_2}{2})r^2 - \epsilon_2$$

$$\mu \circ \phi = -\lambda, \quad \lambda \circ \phi = \mu$$

(where the coordinates on the left are coordinates on $f^{-1}(\epsilon_2)$ and on the right are on $f^{-1}(\epsilon_1)$). Note that ϕ is defined on $\{r^2 > \frac{\epsilon_2}{A(D+\epsilon_2/2)}\}$ so that we take $R = \sqrt{\frac{\epsilon_2}{A(D+\epsilon_2/2)}}$. In particular we can make R arbitrarily small by choosing ϵ_2 small enough.

When we build the handle using a height function h, the only part of $\partial_2\mathcal{H}$ which might fail to be transverse to both vector fields is $\Gamma_h = \Phi(\{(h(r(p)^2), p) \mid R_1 \leq r(p) \leq R_2\})$. Using lemma 5.4 we can write down conditions for Γ_h to be transverse to V^- Using the notation from lemma 5.2 we get:

$$\gamma = r dr \wedge (d\mu - d\lambda)$$

$$g^+ = A(C + D)$$

$$\beta^+ = (C - A(C + D)r^2)(d\mu - d\lambda)$$

$$Z^+ = \frac{1}{r}(C - A(C + D)r^2)\partial_r$$

Thus h needs to satisfy

$$e^{h(r^2)} < A(C + D) - 2h'(r^2)(C - A(C + D)r^2) \ .$$

Since we will have $h' \leq 0$, $1 \leq e^h \leq A(D + \frac{\epsilon_2}{2})$ and $\epsilon_2 < 2C$, this is fine as long as

$$r^2 < \frac{C}{A(C + D)} \ .$$

Therefore we need to build $\mathcal{H}$ with $R_2 < \sqrt{\frac{C}{A(C+D)}}$, and we have completed the proof in the positive case.

In the negative case define (V^+, V^-) as follows:

$$V^+ = \frac{1}{2}r_1\partial_{r_1} + (\frac{1}{2}r_2 + \frac{C}{r_2})\partial_{r_2}$$

$$V^- = -(\frac{1}{2}r_1 + \frac{D}{r_1})\partial_{r_1} - \frac{1}{2}r_2\partial_{r_2}$$

and calculate that (V^+, V^-) is a dilation-contraction pair with:

$$df(V^+) = f + 2D, \quad df(V^-) = -f + 2C$$

Thus V^+ and V^- are both gradient-like when $-2C < f < 2D$. Let $\epsilon_1 = 2D - \frac{2}{A}$ and note that $-2C < \epsilon_1 < 0$ (because $AD < 1$, $A > 0$ and $A(C+D) > 1$). Choose any ϵ_2 with $0 < \epsilon_2 < 2D$. We compute the contact pair induced by (V^+, V^-) on $f^{-1}(\epsilon_1)$ using the coordinates (r, μ, λ) to be given by:

$$\alpha_1^- = \frac{1}{2}[r^2(d\lambda - d\mu)] + \frac{2}{A}d\lambda, \quad \alpha_1^0 = Cd\mu + Dd\lambda$$

and note that $R_{\alpha_1^-} = A(\partial_\mu + \partial_\lambda)$.

Now the forward flow along V^- as a map Φ from some subset of $\mathbb{R} \times f^{-1}(\epsilon_1)$ into $\mathbb{R}^4$ is given by:

$$r_1^2(t) = (r^2 + \frac{2}{A})e^{-t} - 2D$$
$$r_2^2(t) = r^2 e^{-t}$$

Solve for t when $f = \epsilon_2$ to get that the flow lines encounter $f^{-1}(\epsilon_2)$ when

$$t = T = -\log(A(D - \frac{\epsilon_2}{2}))$$

and thus that the diffeomorphism ϕ is given by:

$$r^2 \circ \phi = A(D - \frac{\epsilon_2}{2})r^2 - \epsilon_2$$
$$\mu \circ \phi = -\lambda, \quad \lambda \circ \phi = \mu \ .$$

Thus we take $R = \sqrt{\frac{\epsilon_2}{A(D - \epsilon_2/2)}}$. Again we can make R arbitrarily small by choosing ϵ_2 small enough.

Now we write down conditions for Γ_h to be transverse to V^+:

$$\gamma = r dr \wedge (d\mu - d\lambda)$$
$$g^- = A(C + D)$$
$$\beta^- = (D + A(C+D)r^2)(d\mu - d\lambda)$$
$$Z^- = \frac{1}{r}(D + A(C+D)r^2)\partial_r$$

Thus h needs to satisfy

$$e^{-h(r^2)} < A(C+D) - 2h'(r^2)(D + A(C+D)r^2) .$$

Since we will have $h' \leq 0$, $e^{-h} \leq 1$, $D > 0$ and $A(C+D) > 1$ this can in fact be achieved for any such h and any r^2, and we have completed the proof in the negative case. □

6.4 How to prepare transverse knots for skinny surgery

The "prepared for surgery" condition needed to attach a handle from propositon 6.4 is too restrictive to prove theorem 1.8, so now we discuss product convex-concave bordisms that can "fill the gap" between a general well-behaved contact pair and a contact pair which is prepared for surgery.

Proposition 6.5. *If (α^+, α^-) is well-behaved with respect to coordinates (r, μ, λ) on ν with structural data (A, B, C, D), suppose that the coordinate framing F_μ is* positive *with respect to α^0 and that*

- *in the positive case, $BD > 1$ whereas*
- *in the negative case, $BD < 1$.*

Then there exists a height function h on ν supported inside an arbitrarily small neighborhood of K satisfying the hypotheses of proposition 5.9 yielding a product convex bordism $\mathcal{B}_h$ such that, letting (r, μ, λ) be the induced coordinates on $\partial_2 \mathcal{B}_h$, (α_h^+, α_h^-) is prepared for surgery with respect to (r, μ, λ) on a neighborhood of $K \subset \partial_2 \mathcal{B}_h$.

Proof. In this proof we will use $\pm$ and $\mp$, where $\pm = +$ and $\mp = -$ in the positive case while $\pm = -$ and $\mp = +$ in the negative case. By lemma 6.2 we may assume that $\alpha^\pm = \frac{1}{B \pm Ar^2}(\pm r^2 d\mu + d\lambda)$. Note that the framing condition tells us that C and D are both positive.

Given an $\epsilon > 0$ we will construct h as a function of r^2 with $h(r^2) = 0$ when $r^2 \geq \epsilon$. Our goal is to find such an $h \geq 0$, a positive $\delta < \epsilon$ and a constant A_0 with $A_0(C+D) > 1$ and, in the positive case, $A_0 D > 1$ or, in the negative case, $A_0 D < 1$, such that the following properties are satisfied:

1. Letting $\alpha_h^\pm = e^{\pm h(r^2)}\alpha^\pm$, we need $\alpha_h^\pm = \frac{1}{A_0(1\pm r^2)}(\pm r^2 d\mu + d\lambda)$ for $0 < r < \delta$. Equivalently, we need
$$e^{\pm h(x)} = \frac{B \pm Ax}{A_0(1 \pm x)}$$
for all $x \in (0, \delta^2)$.

2. Letting $g^\pm = \alpha^0(R_{\alpha^\pm}) = AC + BD$, letting $\beta^\pm = \alpha^0 - g^\pm \alpha^\pm$ and letting $Z^\pm$ be the unique vector field in $\ker \alpha^0 \cap \ker \alpha^\pm$ with the property that $\imath_{Z^\pm}\gamma = \beta^\pm$, we need that $e^{\pm h(r^2)} < g - dh(Z^\pm)$. Since $\ker \alpha^0 \cap \ker \alpha^\pm$ is spanned by ∂_r, $Z^\pm$ can be expressed as $Z^\pm = \frac{f(r^2)}{2r}\partial_r$ for some function f, so that $dh(Z^\pm) = h'(r^2)f(r^2)$. We compute $Z^\pm$ and get that $f(x) = (C \mp Dx)(B \pm Ax)$. Thus this condition is equivalent to the condition that:
$$e^{\pm h(x)} < AC + BD - h'(x)(C \mp Dx)(B \pm Ax)$$
for all $x \in (0, \epsilon^2)$. As long as ϵ is small enough we will satisfy this condition if we satisfy the simpler condition that:
$$e^{\pm h(x)} < AC + BD - h'(x)CB \ .$$

The first condition determines h on $[0, \delta^2]$ so first we check that h so defined is not negative and also satisfies the second condition on $[0, \delta^2]$. Since $e^{\pm h(x)} = \frac{B \pm Ax}{A_0(1\pm x)}$, we will have $h \geq 0$ as long as, in the positive case, $B \pm Ax > A_0(1 \pm x)$ or, in the negative case, $B \pm Ax < A_0(1 \pm x)$. For x small enough we need only require that $B > A_0$ or $B < A_0$, respectively.

To see that the second condition is satisfied on $[0, \delta^2]$, we solve for h to get:
$$h(x) = \pm \log(B \pm Ax) \mp \log(A_0) \mp \log(1 \pm x) \ .$$
Solving for $h'(x)$ and plugging into the second condition, we find that we need to be able to find $\delta > 0$ such that
$$B \pm Ax < A_0(D + C)(B \pm Ax)$$
for all $x \in [0, \delta^2]$. But for x small enough, this is satisfied as long as $B < A_0(D + C)(B)$, and since we are requiring that $A_0(D + C) > 1$ anyway, we can always find such a δ.

Now note that if h is defined and satisfies both conditions (and $h \geq 0$) on $[0, \delta^2]$ then, after perhaps making δ smaller, using the facts that $AC + BD > 1$ and C and B are positive, we can extend h to meet the second condition (and $h \geq 0$) on $[0, \epsilon^2]$ with $h = 0$ near ϵ^2. (We make sure that $h'(x) < h'(\delta^2)$ for all $x > \delta^2$ and that $h'(x) < 0$ for all $x > \delta^2 + \delta_1$, for some small $\delta_1 > 0$.)

Thus, we can always find h and δ as long as we can find A_0 satisfying the following conditions:

- In the positive case we require that $\frac{1}{C+D} < A_0 < B$ and that $\frac{1}{D} < A_0$. This can be done as long as $BD > 1$.
- In the negative case we require that $\max(B, \frac{1}{C+D}) < A_0$ and that $A_0 < \frac{1}{D}$. This can be done as long as $BD < 1$.

□

6.5 Putting the pieces together

Now we are ready for the

Proof of Theorem 1.8. First realize X as the support of a convex bordism $\mathcal{X}$ with $\partial_1 \mathcal{X} = \emptyset$, $\partial_2 \mathcal{X} = \partial X$ and with germ $\mathcal{G}(\partial_2 \mathcal{X}) = \mathcal{G}(\alpha)$ where α is the induced positive contact form on ∂X. Pick well-behaved coordinates (r, μ, λ) near each component K of L with $F_\mu = F_K$. Since ∂X is compact we may rescale α^0 to get $\alpha^0(R_\alpha) > 1$ and to make sure that the structural data near each K satisfy $BD > 1$. Then let $\alpha^+ = \alpha$ and $\alpha^- = \alpha^0 - \alpha$. By lemma 5.1, (α^+, α^-) is an extension of the purely positive pair $(\alpha, 0)$, so $\mathcal{G}(\alpha^+, \alpha^-) = \mathcal{G}(\alpha)$. Now, near each component K we find that (α^+, α^-) satisfies the conditions of proposition 6.5 so that we can attach product convex-concave bordisms along each K to get a new convex-concave bordism $\mathcal{Z}$ with $\partial_1 \mathcal{Z} = \emptyset$ and a diffeomorphism $\phi : \partial_2 \mathcal{Z} \to \partial_2 \mathcal{X}$. We have that $\mathcal{G}(\partial_2 \mathcal{Z}) = \mathcal{G}(\alpha_Z^+, \alpha_Z^-)$ where $\phi^* \alpha_Z^+ = e^h \alpha^+$ for some h compactly supported in a neighborhood of L, $\phi^* \alpha_Z^0 = \alpha^0$ and (α_Z^+, α_Z^-) is prepared for surgery with respect to the framing $\phi(F_K)$ near each knot $\phi(K)$. Now attach 2-handles as constructed by proposition 6.4 (positive case) along $\phi(L)$ with framing $\phi(F)$ to construct a convex-concave bordism $\mathcal{Y}$ with $\partial_1 \mathcal{Y} = \emptyset$ and $\partial_2 \mathcal{Y}$ the result of surgery on $\partial_2 \mathcal{X}$ along L with framing F. Note that $\mathcal{G}(\partial_2 \mathcal{Y}) = \mathcal{G}(\alpha_Y^+, \alpha_Y^-)$, where α_Y^- is now defined on all of

$\partial_2\mathcal{Y}$. So in fact (α_Y^+, α_Y^-) is essentially negative and $\mathcal{G}(\partial_2\mathcal{Y}) = \mathcal{G}(0, \alpha_Y^-)$ or, in other words, $Y = \mathrm{supp}(\mathcal{Y})$ has concave boundary with induced *negative* contact form α_Y^- on ∂Y. □

The new 3-manifold ∂_Y has a link L_Y made up of the ascending circles of the 2-handles. Note that the negative contact form α_Y^- is characterized as follows: There exists a closed tubular neighborhood τ of L and a diffeomorphism ϕ from $\partial X \setminus \tau$ to $\partial_Y \setminus L_Y$ such that $\alpha^0 - \phi^*(\alpha_Y^-) = e^h\alpha$ for some function $h : M_1 \setminus \tau \to [0, \infty)$.

One might imagine that there exists a parallel theorem describing how to attach handles constructed by proposition 6.4 in the *negative* case to a concave boundary to yield a convex boundary. Two complications arise: The first is that, to rescale α^0, we in general want to multiply by a constant larger than 1 to get $\alpha^0(R_\alpha) > 1$ but to multiply by a constant less than 1 to get $BD < 1$, to satisfy proposition 6.5 in the negative case. We can state a theorem that assumes that α^0 already satisfies these conditions but then it turns out to be useless because one can show that there never exists a negative contact form and a closed 1-form on a compact 3-manifold satisfying the required conditions.

We have alot more freedom in attaching these 2-handles if we do not try to make the final contact pair essentially pure and if we attach the handles along knots in pure regions. Thus we now consider the case where (α^+, α^-) is either a purely positive contact pair on ν, in which case we have $\alpha = \alpha^+$ and we have no α^- or α^0, or a purely negative contact pair on ν, in which case we have $\alpha = \alpha^-$ and we have no α^+ or α^0.

Definition 6.6. *A contact form α on ν is almost prepared for surgery with respect to (r, μ, λ) if it is well-behaved with respect to (r, μ, λ) with constants $A = B$, i.e. if*

$$R_\alpha = A(\partial_\mu + \partial_\lambda)$$

for some constant $A > 0$.

Proposition 6.7. *Given a single contact form α on ν, suppose that K is transverse to $\xi = \ker\alpha$. Apply lemma 1.3 and shrink ν if necessary to get a coordinate system (r, μ, λ) on ν, realizing some chosen framing of K, with radius R_ν and framing function s. Let R be any radius with $R < R_\nu$ and with $s(R^2) > 1$ if α is negative or with $s(R^2) < 1$ if α is positive. Then there exists a product convex bordism $\mathcal{B}_h$ from $(\nu, \mathcal{G}(\alpha))$ to $(\nu, \mathcal{G}(\alpha_1))$ for a contact form α_1 which is almost prepared for surgery with respect to (r, μ, λ) on $\{r < R\}$.*

Proof of proposition 6.7. Lemma 1.3 gives us the coordinate system and the framing function, so that $\ker\alpha = \ker\alpha_1$ for some α_1 which is almost prepared for surgery (use the existence in lemma 6.2). Thus $\alpha_1 = e^h\alpha$ for some h. If necessary replace α_1 by $k\alpha_1$ to make $h > 0$ on $\{r \leq R + \delta < R_\nu\}$ for some positive δ, then multiply h by a bump function to give it compact support. Let $\mathcal{B}_h$ be the product convex bordism of height h. $\square$

This allows us to prove:

Theorem 6.8. *Given a convex-concave bordism $\mathcal{X}$ with $\mathcal{G}(\partial_2\mathcal{X}) = \mathcal{G}(\alpha^+, \alpha^-)$ and with the domains of the 1-forms labelled $\partial_2^\pm\mathcal{X}$, if K is a* transverse *knot in the interior of $\partial_2^\pm\mathcal{X} \setminus \partial_2^0\mathcal{X}$, one can attach a convex-concave 2-handle to $\partial_2\mathcal{X}$ along K with* any *chosen framing.*

Proof. Let α be whichever form in the contact pair is defined near K. Use proposition 6.7 to get α almost prepared for surgery with respect to coordinates (r, μ, λ) near K realizing the desired framing. Then extend α^0 to be equal to $Cd\mu + Dd\lambda$ on a small neighborhood of K well inside $\partial_2^\pm\mathcal{X} \setminus \partial_2^0\mathcal{X}$, for positive constants C, D chosen so that $\alpha^\pm$ is prepared for surgery near K. Now attach a 2-handle. $\square$

Note that it becomes quite messy to state properties that characterize the new contact pair in this theorem.

Chapter 7

Examples and Questions

The purpose of this chapter is to develop some specific examples where theorem 1.8 applies, to sketch some ideas ripe for future development, and to present some interesting questions.

7.1 Examples

First we will show that S^3 with surgery on either the unknot with any framing $F > 0$ or the Hopf link with any framing $F \geq 0$ can be realized as the concave boundary of a symplectic 4-manifold.

Let $(r_1, \theta_1, r_2, \theta_2)$ be polar coordinates on $\mathbb{R}^4$. Consider $S^3 = \partial B^4 \subset (\mathbb{R}^4, \omega)$ where $\omega = r_1 dr_1 \wedge d\theta_1 + r_2 dr_2 \wedge d\theta_2$ is the standard symplectic form. Then $V = \frac{1}{2}(r_1\partial_{r_1} + r_2\partial_{r_2})$ is a symplectic dilation transverse to S^3 inducing the standard positive contact form $\alpha = \frac{1}{2}(r_1^2 d\theta_1 + r_2^2 d\theta_2)$ on S^3. We compute that $R_\alpha = \partial_{\theta_1} + \partial_{\theta_2}$. Now consider two cases:

1. $L = K = \{r_1 = 0\} \subset S^3$, the standard unknot. Consider the fibration $P = \theta_1 : S^3 \setminus L \to S^1$. Then $dP(R_\alpha) > 0$. Polar coordinates near K are given by $(r = r_1, \mu = \theta_1, \lambda = \theta_2)$, from which we can verify that α and $\alpha^0 = dP$ are well-behaved near K. With respect to these coordinates, F_μ is the standard 0-framing of K, and since $\alpha^0 = d\mu$, the condition that a framing F is positive with respective to α^0 is simply the condition that $F > 0$. Thus we can attach a single "convex-to-concave" handle along the unknot with any framing $F \geq 1$ to make a symplectic manifold with concave boundary.

2. $L = K_1 \cup K_2$ where $K_i = \{r_i = 0\} \subset S^3$, the standard Hopf Link. Now consider the fibration $P = \theta_1 + \theta_2 : S^3 \setminus L \to S^1$. Again, $dP(R_\alpha) > 0$. Oriented polar coordinates near K_1 are as above and oriented polar coordinates near K_2 are given by $(r = r_2, \mu = \theta_2, \lambda = \theta_1)$ and again we verify that α and α^0 are well-behaved near each K_i. Also F_μ is the standard 0-framing of each K_i. Now, however, $\alpha^0 = d\mu + d\lambda$ near each K_i, so the condition that a framing F is positive with respect to α^0 is the condition that $F_{K_i} > -1$ for each K_i. Thus we can attach a pair of "convex-to-concave" handles along the Hopf link as long as each handle is framed with framing 0 or larger, and the result is a symplectic manifold with concave boundary.

The first example generalizes using lemma 3.12. Given an n-punctured surface Σ with a proper Morse function f_Σ as in the lemma, apply the lemma to get a symplectic form ω_Σ and a gradient-like symplectic dilation V_Σ such that the structure is "standard" on, say, $f^{-1}(\frac{1}{2}, \infty)$. In other words, $f^{-1}(\frac{1}{2}, \infty)$ looks like n copies of $\mathbb{R}^2 \setminus \{r^2 \leq \frac{1}{2}\}$ with its standard symplectic form, dilation and Morse function. Consider the symplectic 4-manifold $(\Sigma \times \mathbb{R}^2, \omega = \omega_\Sigma + \omega_{\mathbb{R}^2})$ with the Morse function $f = f_\Sigma + r^2$. Note that $V = V_\Sigma + \frac{1}{2} r\partial_r$ is a gradient-like symplectic dilation for f. Let $X = f^{-1}[0,1]$ and let $M = \partial X = f^{-1}(1)$; M is the convex boundary of (X, ω) with induced contact form $\alpha = \alpha_\Sigma + \frac{1}{2} r^2 d\theta$, where $\alpha_\Sigma = \imath_{V_\Sigma} \omega_\Sigma$. Note that M is diffeomorphic to the "boundary with smoothed corners" of the product of a disk and a compact surface of genus g with n boundary components. Alternately, if the handle decomposition of Σ is the usual one with one 0-handle and $(2g + n-1)$ 1-handles, then X is diffeomorphic to B^4 with $(2g+n-1)$ 1-handles attached, and M is diffeomorphic to S^3 with 0-surgery on $(2g+n-1)$ unlinked unknots. We can decompose M into two open sets: $U = \{f_\Sigma < 1, r^2 = 1 - f_\Sigma\}$ and $V = \{0 \leq r^2 < \frac{1}{2}, \frac{1}{2} < f_\Sigma = 1 - r^2 \leq 1\}$. The set U is actually the complement of a link $L = \{r^2 = 0, f_\Sigma = 1\}$ of n components, and the coordinate function $\theta : U \to S^1$ is a fibration with fiber diffeomorphic to Σ. To see that $d\theta(R_\alpha) > 0$, recall that this is equivalent to the requirement that $d\theta \wedge d\alpha > 0$ (see lemma 5.1). But $d\theta \wedge d\alpha = d\theta \wedge d\alpha_\Sigma = d\theta \wedge \omega_\Sigma > 0$. On V, the entire structure is identical to n copies of the structure described in the earlier example on S^3 in a neighborhood of the standard unknot. Thus we can attach handles along L with any framing larger than the "0-framing" determined by the fibration to create a concave symplectic 4-manifold. (See figure 7.1 for an example of a framed link diagram associated with this construction. For background on interpreting these diagrams see [8].)

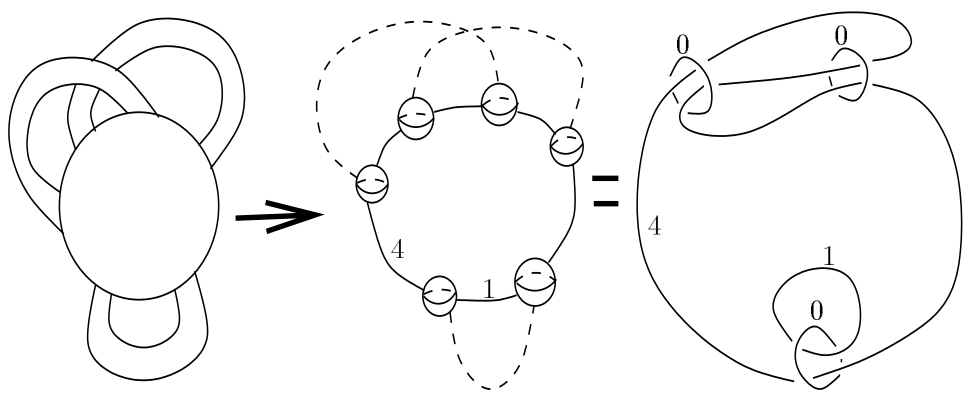

Figure 7.1: A handlebody decomposition of a surface giving a framed link description of a concave 3-manifold boundary.

In these constructions, the characteristic foliation on each fiber is given by the flow lines of the vector field V_Σ. In general there will be some closed leaves of this singular foliation (containing singular points). Each such closed leaf is a Legendrian knot the Thurston-Bennequin framing of which is the framing given by the fiber. Note that, after attaching the 2-handles and getting a concave boundary, the new negative contact form on the boundary induces the same characteristic foliation on the fibers as the original positive contact form did, and still has the property that its Reeb vector field is transverse to the fibers (because the two contact forms form a contact pair (α^+, α^-) with $\alpha^0 = \alpha^+ + \alpha^- = kd\theta$ for some constant $k > 0$). Thus these closed leaves are also Legendrian knots with respect to the negative contact form and their Thurston-Bennequin framings are again given by the fibers. Using theorem 1.2 we can attach symplectic 2-handles along these knots to build concave manifolds. This increases our class of examples of 3-manifolds which bound concave symplectic 4-manifolds; first perform any positive surgeries along the original n-component link L, then perform -1 surgeries on arbitrarily many copies of each closed leaf of the singular foliation (use disjoint fibers to get the different copies). (See figure 7.2 for a random example of this more general construction.)

These examples lead us to think about a special class of contact pairs on 3-manifolds. Consider a contact pair (α^+, α^-) on a 3-manifold M with the property that $M \setminus M^0$ is a link L, each component of which is well-behaved with respect to (α^+, α^-), and such that $\alpha^0 = kdP$ for some constant $k > 0$ and some fibration $P : M^0 \to S^1$. We might call such a pair together with the fibration a *fibered contact pair*. Let $\Sigma = P^{-1}(0)$, a punc-

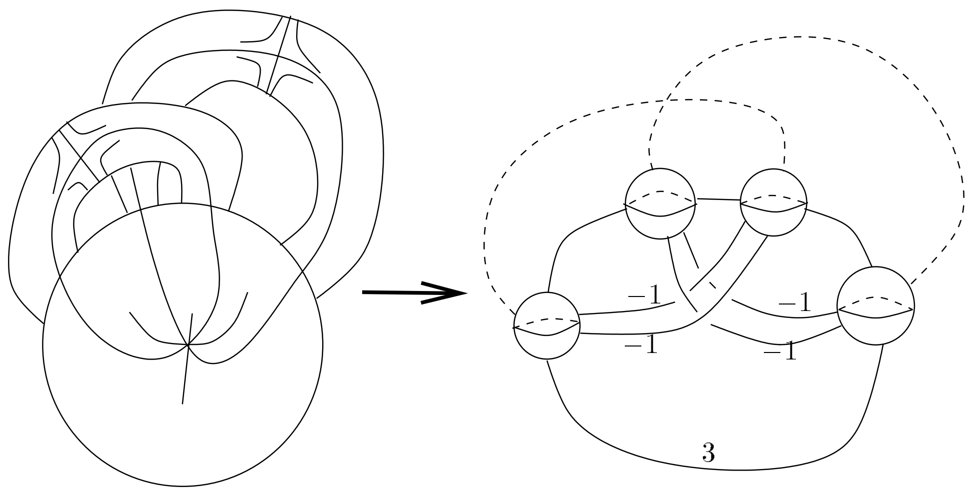

Figure 7.2: A handlebody decomposition of a surface with closed leaves of the characteristic foliation giving a framed link description of a concave 3-manifold boundary using Legendrian surgery.

tured surface and let $\alpha = \alpha^+$ or α^-. If we also assume that Σ intersects each well-behaved neighborhood of each component of L in a single annulus then the number of punctures of Σ is equal to the number of components of L and flow along R_α defines a return map $\psi : \Sigma \to \Sigma$ which describes the topology of M^0. In fact, if we embed Σ in a closed surface $\hat{\Sigma}$ by filling in each puncture with a disk, ψ extends across $\hat{\Sigma}$ because of the "well-behaved" condition.

We also get a "time-till-return map" $T : \Sigma \to (0, \infty)$ (the time it takes to get back to Σ by flowing along R_α. This, together with the 1-form $\beta = \alpha|_\Sigma$ determines α on M^0. We do not quite have enough information to determine both contact forms $\alpha^\pm$ unless we assume that $dP(R_\alpha)$ is invariant under flow along R_α, but at the very least we know that $dP(R_\alpha)$ is constant near L (due to "good behavior"). To determine the entire structure of $(M, (\alpha^+, \alpha^-))$ we need to know the Dehn fillings that create M from M^0 and we need to know how to extend one or the other contact form over each component of L. In fact, given the decomposition $L = L^+ \cup L^-$, where $L^\pm = M^\pm \setminus M^0$, the Dehn fillings are completely determined by the slopes of the contact structures in the neighborhoods of each component of L.

Thus we should be able to build a theory with very explicit control on such contact pairs, describe exactly the results of attaching handles along the link L and use it for the construction of more subtle symplectic 4-manifolds.

7.2 Questions

Comparing the Legendrian surgeries due to Weinstein with the "fat transverse surgeries" described in chapter 4 raises the following questions:

1. It seems likely that the two surgeries yield the same contact structures, when proposition 1.6 is used to generate transverse knots near Legendrian knots which are fat with respect to the Thurston-Bennequin framings minus 1. However it is difficult to make this explicit and it would be useful to find a clear explanation of why this should be true.

2. There do exist transverse knots which are fat with respect to a certain framing F such that no Legendrian knot K in the same isotopy class has $\mathrm{tb}(K) - 1 = F$. The immediate example is the standard unknot in S^3, which is fat with respect to -1. This does not obviously lead to interesting new tight contact structures because the transverse surgery "uses up" the fat neighborhood of the knot. An important question is whether there are any fat transverse surgeries that can be used to create new contact structures by virtue of the fact that they realize framings unachievable by Legendrian surgeries.

The general theory of contact pairs and symplectic germs raises these questions:

1. Does this theory capture all the possible local behavior in a neighborhood of a 3-manifold in a symplectic 4-manifold? In other words, is every 4-dimensional symplectic germ along a 3-manifold equal to the germ determined by a contact pair?

2. In order for this to be true, we would need to answer the more general question of whether, given an arbitrary nowhere-zero 2-form γ on a 3-manifold M, there always exists a contact pair $\alpha^{\pm}$ such that $\gamma = \pm d\alpha^{\pm}$.

3. Is there a good method, given two 4-dimensional symplectic germs along the same 3-manifold, of deciding whether they can be connected by a product symplectic bordism?

Using the skinny transverse 2-handles as building blocks for symplectic 4-manifolds, one naturally considers the following questions:

1. Putting the skinny transverse 2-handles together with the 1-handles and 2-handles described by Weinstein, can we build every symplectic 4-manifold? Is there a theory of symplectic handle-cancellation and handle-slides in this context? Can we put an arbitrary Morse function on a symplectic 4-manifold and, after some perturbations, expect it to give us nice symplectic handle structures as in this paper?

2. Can every symplectic structure on $I \times M$ be constructed using these methods? Can this be used to construct interesting symplectic structures on $S^1 \times M$?

3. Using theorem 1.8 to create symplectic 4-manifolds with concave boundaries, sometimes the induced negative contact structures will be tight because the symplectic 4-manifold constructed in fact embeds in a closed symplectic 4-manifold. In general there is no reason to expect this to be true. Are there examples where this construction yields overtwisted contact structures?

Lastly, one naturally wonders to what extent the material developed here extends to higher dimensions.

Bibliography

[1] Ralph Abraham, Jerrold E. Marsden, and Tudor Ratiu. *Manifolds, tensor analysis, and applications.* Addison-Wesley Publishing Co., Reading, Mass., 1983.

[2] Yakov Eliashberg. Filling by holomorphic discs and its applications. In *Geometry of low-dimensional manifolds, 2 (Durham, 1989)*, pages 45–67. Cambridge Univ. Press, Cambridge, 1990.

[3] Yakov Eliashberg. Legendrian and transversal knots in tight contact 3-manifolds. In *Topological methods in modern mathematics (Stony Brook, NY, 1991)*, pages 171–193. Publish or Perish, Houston, TX, 1993.

[4] Yasha Eliashberg. Unique holomorphically fillable contact structure on the 3-torus. *International Mathematics Research Notices*, 1996(2):77–82.

[5] John Etnyre. *Symplectic Constructions on 4-Manifolds.* PhD thesis, The University of Texas at Austin, 1996.

[6] John Etnyre and Robert Ghrist. Tight contact structures via dynamics. *Proceedings of the American Mathematical Society.* To appear.

[7] Robert E. Gompf. Handlebody construction of Stein surfaces. *Annals of Mathematics. Second Series*, 148(2):619–693, 1998.

[8] Robion C. Kirby. *The topology of* 4*-manifolds.* Springer-Verlag, Berlin, 1989.

[9] Antoni A. Kosinski. *Differential manifolds.* Academic Press Inc., Boston, MA, 1993.

[10] Dusa McDuff and Dietmar Salamon. *Introduction to symplectic topology.* Oxford University Press, Oxford, second edition, 1998.

[11] Alan Weinstein. Contact surgery and symplectic handlebodies. *Hokkaido Mathematical Journal*, 20(2):241–251, 1991.