Legendrian Surgeries

Nancy Mae Eagles

January 14, 2025

Contents

1 Topological Surgery

1.1 Dehn surgery

Recall the general procedure of performing Dehn surgery to a topological 3 manifold Y: Choose a knot $K \subseteq Y$, and let $\nu(K)$ denote its tubular neighborhood. Then, choose a meridian and longitude of the knot, considered as curves on $\partial\nu(K)$. The meridian μ should bound a disk in $\nu(K)$, and the longitude λ should be homotopically nontrivial in $\nu(K)$, or equivalently nullhomologous in the $Y \setminus \nu(K)$. Such a choice of μ and λ allows us to identify $\partial\nu(K)$ with $\mathbb{R}^2/\mathbb{Z}^2$, with where μ has slope 0 and λ has slope ∞. To perform Dehn surgery, cut out the tubular neighborhood, and glue it in back by an orientation preserving diffeomorphism of the boundary $S^1 \times S^1$.

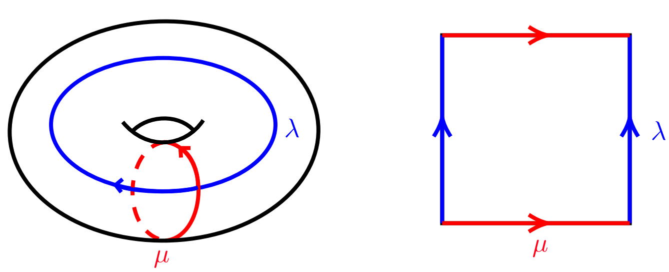

By the above, the space of orientation-preserving diffeomorphisms can be identified with $SL_+(2, \mathbb{Z})$. Let $p/q \in \mathbb{Q}$. We refer to the above procedure sending $\mu \mapsto p\mu + q\lambda$ as rational p/q surgery.

1.2 Dehn surgery as handle attachments

We can also describe integral surgeries as handle attachments to a 4-manifold. Let $X = Y \times [0,1]$. Then X is a 4-manifold with boundary $Y \sqcup -Y$. We may attach a 4-dimensional 2-handle $D^2 \times D^2$ to X along the boundary $Y \times \{1\}$ by specifying an attaching map $f : S^1 \times D^2 \hookrightarrow Y$. Such a map can be described by a knot in Y, along with a framing of the knot identifying its normal bundle with a tubular neighborhood of the knot $\nu(K) \simeq S^1 \times D^2$, and a gluing map of the attaching region of the handle to $\nu(K)$. These gluing maps are indexed by the integers, $S^1 \times \{0\} \mapsto \lambda + k\mu$ and $\{0\} \times S^1 \mapsto \mu$. Here, μ denotes the meridian and λ the longitude chosen to identify $\nu(K)$ with $S^1 \times D^2$.
After attaching the 2-handle, the boundary $Y \times \{1\}$ loses a copy of $S^1 \times D^2$, and gains a copy of $D^2 \times S^1$. The cutting out of $S^1 \times D^2$ swaps the roles of μ and λ, and so under the identification, $\partial D^2 \times \{pt\} \mapsto k\mu + \lambda$. Thus, this operation performs k-surgery on $Y \times \{1\}$.

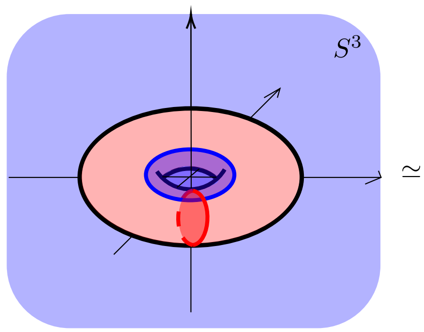

Figure 1: Meridian and longitude swap between knot neighborhood and knot complement.

1.3 Handle-attachings verus Dehn Fillings

When attaching a handle, we identify two **solid** tori with each other via an automorphism of $S^1 \times D^2$. When Dehn-filling, we **close** a boundary $S^1 \times S^1$ component by gluing in a copy of $S^1 \times D^2$. Such surgeries are given up to an automorphism of $S^1 \times S^1$. These processes are related by the following table.

Operation	Picture	Gluing Map	Matrix
Handle Attaching	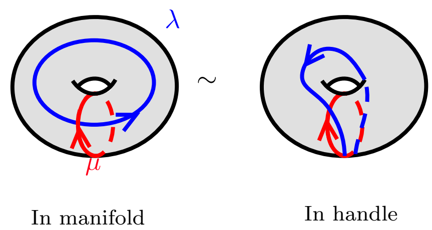	$\mu \mapsto \mu$, $\lambda \mapsto -\mu + \lambda$ Identifying two solid tori	$\begin{bmatrix}1 & -1\\0 & 1\end{bmatrix}$ related by transposition
Dehn Filling	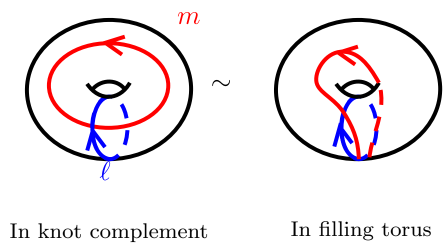	$m \mapsto \mu - \lambda$, $\ell \mapsto \lambda$ Gluing along knot complement	$\begin{bmatrix}1 & 0\\-1 & 1\end{bmatrix}$

2 Contact surgery

We'd like to extend topological knot surgery to a contact/symplectic construction, equipping the surgered region with a contact structure that that in $\nu(K)$ under the chosen gluing. There are several models to do describe this extension. We begin by describing the standard model of a symplectic 2-handle attachment.

2.1 Symplectic 2-handle attachment

This description is used in [HT11]. An alternative description is given in [Gei08, Ch6], and [Etn06]. This construction is originally due to Eliashberg [Eli90], and independently Weinstein [Wei91].

Consider $\mathbb{R}^4$ with coordinates $\{q_1, q_2, p_1, p_2\}$, and symplectic form $\omega = \sum dp_i \wedge dq_i$. The vector field

$$V = \sum_i -p_i \partial_{p_i} + 2q_i \partial_{q_i}$$

is Liouville with respect to this symplectic form, and transverse to the hypersurface $Y = \{p_1^2 + p_2^2 = 1 \subseteq \mathbb{R}^4\}$.
Proof. Since ω is closed, $d\omega = 0$, so $\mathcal{L}_V \omega = d\iota_V \omega$. But

$$\begin{aligned} \iota_V \omega &= -p_1 dq_1 - p_2 dq_2 - 2q_1 dp_1 - 2q_2 dp_2, \\ \implies d(\iota_V \omega) &= -p_1 \wedge dq_1 - dp_2 \wedge dq_2 - 2dq_1 \wedge dp_1 - 2dq_2 \wedge dp_2 \\ &= dp_1 \wedge dq_1 + dp_2 \wedge dq_2 = \omega. \end{aligned}$$

It is clearly also transverse to Y. □
As such, $\lambda = \iota_V \omega$ is a contact for the hypersurface Y, given by

$$\lambda = -p_1 dq_1 - p_2 dq_2 - 2q_1 dp_1 - 2q_2 dp_2.$$

Within Y, we have a model Legendrian knot given by

$$K = \{q_1 = q_2 = 0, p_1^2 + p_2^2 = 1\}$$

Proof. We will switch to coordinates $p_1 = \cos(\theta)$ and $p_2 = \sin(\theta)$ for simplicity. Then our knot K admits the parametrization $\psi : (q_1, q_2, \theta) \mapsto (0, 0, \cos(\theta), \sin(\theta)) \in \mathbb{R}^4$, and our contact form with respect to these coordinates is

$$\lambda = -\cos(\theta)dq_1 - \sin(\theta)dq_2 + 2q_1 \sin(\theta)d\theta - 2q_2 \cos(\theta)d\theta.$$

In particular, when $q_1 = q_2 = 0$, the contact form becomes

$$\lambda = -\cos(\theta)dq_1 - \sin(\theta)dq_2,$$

for which certainly $\partial_\theta \in \ker(\lambda)$. So K is Legendrian. □

We now describe a model of attaching a 2-handle to the region $\{p_1^2 + p_2^1 \geq 1, (q_1, q_2) \in \mathbb{R}^2\}$ as follows. The boundary of this region is a copy of $S^1 \times \mathbb{R}^2$, within which we specify an attaching region to be

$$\{p_1^2 + p_2^2 = 1, q_1^2 + q_2^2 \leq \epsilon\}$$

and consider our handle as

$$D^2 \times D^2 = \{p_1^2 + p_2^2 \leq 1, q_1^2 + q_2^2 \leq \epsilon\}.$$

Remark 1. Let $K : S^1 \hookrightarrow Y$ be a Legendrian knot in a cooriented contact manifold, with contact form λ. Such a contact form induces an orientation on the symplectic vector bundle $(\xi, d\lambda)$, and on the manifold Y, and these orientations are compatible. In a tubular neighborhood of K, there exist local coordinates such that one direction of the contact structure is given by flowing along the knot, and the other is transverse to the knot. Equipped with this induced orientation, there is a well-defined notion of a positive and negative Legendrian push-off of K, $\ell_\pm$. The positive Legendrian push-off gives a canonical framing from the contact structure. That is, we consider K as having $tb(K)$ framing, where $tb(K)$ is the Thurston-Bennequin number of K. Note that the standard Seifert framing differs from the contact framing by $tb(K)$.

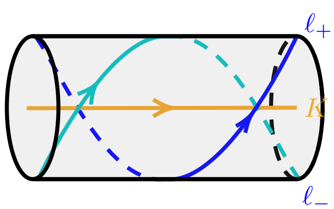

Figure 2: Positive and negative Legendrian push-offs

We now glue on the 2-handle by framing $tb(K) - 1$. That is, we identify the longitude ℓ_+ with a longitude of linking number $tb(K) - 1$ in the handle. The handle attachment reduces the number of right-handed meridional twists of the contact structure in the tubular neighborhood of the knot by 1.

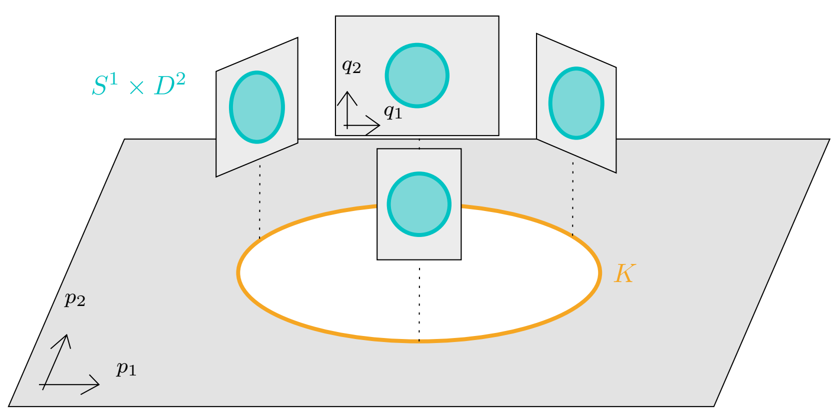

Figure 3: Model handle attaching region. Here we consider $\mathbb{R}^4$ as $\mathbb{R}^2_{q_1,q_2}$ fibered over $\mathbb{R}^2_{p_1,p_2}$.

We can round out the corners of the construction by making a suitable identification of the corner region with a smoothing hypersurface. On the thickened torus $\{p_1^2 + p_2^2 = 1, \epsilon/2 \leq q_1^2 + q_2^2 \leq \epsilon\}$, we can replace the hypersurface on the handle with one cut out by an equation given

$$f(q_1^2 + q_2^2, p_1^2 + p_2^2) = 0$$

where $\partial_x f > 0$ and $\partial_y f = 0$, so that the attaching looks like

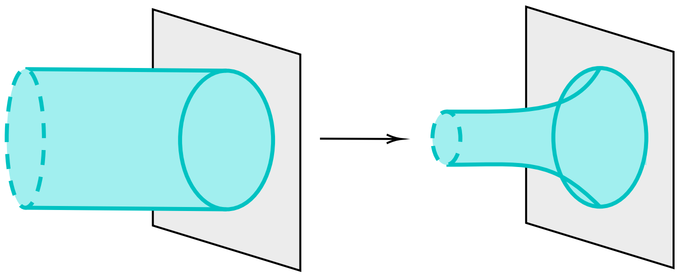

Figure 4: Smooth attaching of handle

Remark 2. The handle attaching amounts to doing -1-surgery on the knot K in Y. This description is useful, as it provides a canonical cobordism between Y before and after surgery. These cobordisms give rise to chain maps in SWH, ECH, and HFH.

2.2 Dehn surgery via 1-jet bundles

This description is used in [Avd23], and also appears in less detail in [BEE12]. The following exposition follows the notation of the former.

We may also describe Dehn surgery at the level of 3-manifolds. Whereas in the previous description we only obtained -1-surgery (and more generally integral surgeries), we can now extend these ideas to perform rational $1/k$-surgeries along Legendrian knots.
As before, let $K : S^1 \hookrightarrow Y$ be a Legendrian knot. We may consider an alternative description of a standard tubular neighborhood of K, identifying it with the 1-jet bundle of S^1, $\mathbb{R}_z \times T^*S^1(q, p)$, equipped with the contact form $\lambda = dz + pdq$. We restrict our attention to the neighborhood $\nu_\epsilon(K) := I_{\epsilon,z} \times I_{\epsilon,p} \times S^1_q$, where $I_\epsilon = [-\epsilon, \epsilon]$. We can identify K with the Legendrian $\{z = p = 0\}$. Moreover, the vector field $\partial_p \in \ker(\lambda)$ is transverse to this curve, and gives a canonical positive and negative push-off as before.

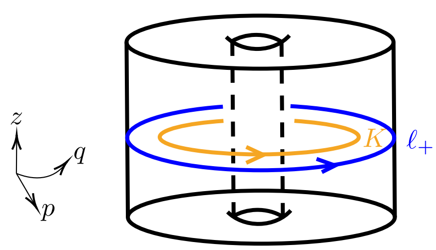

We'd like to construct agluing map that identifies the boundary of $\nu_\epsilon(K)$ with itself, sending the longitude $\ell_+ \mapsto \ell_+ - \mu$, where μ is a standard meridian. We will concentrate the nontriviality of the gluing at the top of the neighborhood. This will make the effect of doing surgery on the Reeb dynamics simpler to see.
Let $f : \mathbb{R} \to S^1$ be a smoothing of the piecewise linear function

$$f(x) = \begin{cases} 0 & \text{if } x < -1/2, \\ x + 1/2 & \text{if } -1/2 \leq x < 1/2, \\ 1 & \text{if } x \geq 1/2, \end{cases}$$

and let $\tau_f \in \mathrm{Diff}^+(\mathbb{R}_p \times S^1_q)$ be the diffeomorphism

$$\tau_f(p,q) = (p, q + f(p))$$

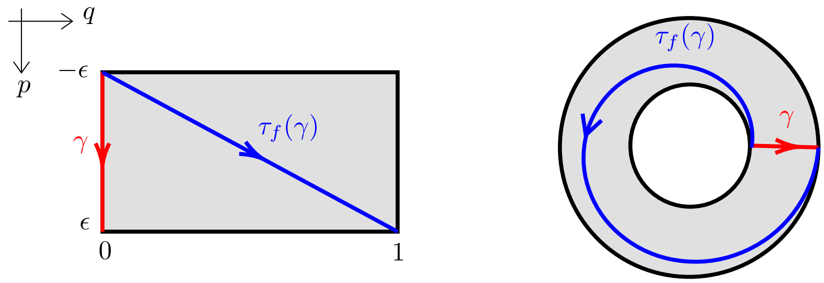

Figure 5: Image of τ_f.

Straightforward computations show that τ_f preserves the symplectic form $dp \wedge dq$, but does not preserve pdq.
Proof.

$$\tau_f^* pdq = pd(q + f(p)) = pdq + p\partial_f(p)dp,$$

Which implies

$$\tau_f * (dp \wedge dq) = dp \wedge dq.$$

Notice that $\tau_f^* pdq$ differs from pdq by the factor $p\partial_f(p)dp$. □
Let $f_\epsilon(p) = f(p/\epsilon)$. We'd like to construct a gluing that preserves the contact form, so that we may identify the contact forms inside and outside of $\nu_\epsilon(K)$ along its boundary after surgery. Hence, let

$$H_\epsilon(p) = \int_{-\infty}^{p} P\partial_p f_\epsilon(P)\ dP.$$

Roughly, H_ϵ measures the failure of τ_f in preserving pdq.
Let $0 < \delta < \epsilon/2$, and define regions T, B, S of an arc of $\nu_\epsilon(K)$ as in the figure below. The thickness of each component is given by δ.

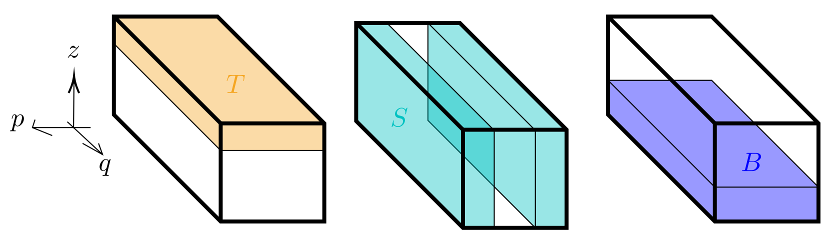

Figure 6: Gluing regions of $\nu_\epsilon(K)$.

Define the gluing map

$$\phi(z,p,q) \mapsto \begin{cases} (z + h_\epsilon(p), p, q - f_\epsilon(p)) & (z,p,q) \in T, \\ (z,p,q) & (z,p,q) \in B \cup S. \end{cases}$$

By taking δ to be sufficiently small compared to the smoothing of f_ϵ and H_ϵ, we can assure that the gluing agrees on the intersections $T \cap S$. We can easily visualize the effect of this surgery on the standard meridian: From this, it is clear that the effect of this gluing is performing -1-surgery in the topological sense. Notice then that

$$\begin{aligned} \phi^*(dz + pdq) =& d(z + H_\epsilon(p)) + pd(q - f_\epsilon(p)) \\ =& dz + \partial_p H_\epsilon(p)dp + pdq - p\partial_p f_\epsilon(p)dp \\ =& dz + pdq + (\partial_p H_\epsilon(p) - p\partial_p f_\epsilon(p))dp \\ =& dz + pdq, \end{aligned}$$

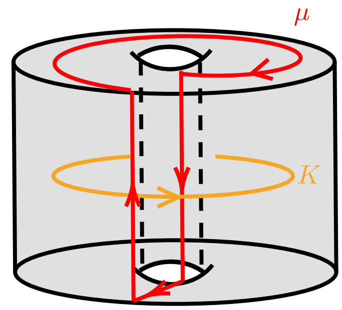

Figure 7: Image of μ after -1-surgery

where we use the fact that f_ϵ, and hence H_ϵ, is compactly supported. Thus, the gluing preserves the contact form. Notice that the Reeb vector field, viewed from outside the neighborhood, twists nontrivially around $\nu(K)$, rather than flowing vertically through the neighborhood as before.

2.3 Aternative Description

This description is used in [Gei08, Ch2, Ch4, Ch6], and in [Hon00]. We follow the conventions of [Gei08].

We now present a third description of Legendrian surgery, following the conventions of [Gei08]. This description uses a tubular neighborhood that is different from the one described in the previous section. The Legendrian push-offs in either case give a topologically different framing of the knot K, however the constructions achieve the same -1-surgered manifold.
Let K be a Legendrian knot. Then there exists a tubular neighborhood of K, $\nu(K)$ contactomorphic to

$$(\mathbb{R}^2_{(x,y)} \times S^1_z, \lambda = \cos(z)dx - \sin(z)dy).$$

As before, $\{(0,0,z) \mid z \in S^1\}$ is identified with our knot in K. From the contact structure, we have two distinguished dividing curves on the boundary of the neighborhood, given by

$$\ell_\pm := \{(\pm\delta sin(z), \pm\delta\cos(z), z) \mid z \in S^1\}$$

We take $\ell_\pm$ to be the positive Legendrian push-off of K.

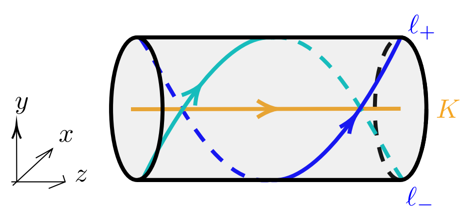

We begin by describing the effect of a general p/q-surgery on the contact form. Let (μ, ℓ_+) be a chosen meridian and longitude as above. Suppose we attempt the gluing given by

$$\mu \mapsto p\mu + q\ell_+ \qquad \ell_+ \mapsto s\mu + t\ell_+$$

Converting to polar coordinates $(x,y) \to (r,\theta)$,

$$\lambda = \cos(z+\theta)dr - r\sin(z+\theta)d\theta$$

Let $(\bar{r}, \alpha, \beta)$ be coordinates on the solid torus we'd like to glue back. Then the gluing gives transition maps

$$\theta = p\alpha + m\beta, \qquad \theta + z = q\alpha + n\beta.$$

And the pullback of the contact form along this gluing is thus

$$\bar{\lambda} = \cos(q\alpha + n\beta)dr - pr\sin(q\alpha + n\beta)d\alpha - mr\sin(q\alpha + n\beta)d\beta.$$

To determine whether this is still contact, it suffices to check that $\bar{\lambda} \wedge d\bar{\lambda}$ is a volume form. After a tedious calculation, we get that

$$\bar{\lambda} \wedge d\bar{\lambda} = r(pn - mq)d\alpha \wedge d\beta \wedge dr.$$

Since our gluing is an orientation-preserving diffeomorphism of $S^1 \times S^1$, we have that $pn - mq = 1 \neq 0$. So this is a well-defined volume form everywhere **except** when $r = 0$, i.e. along K.

Remark 3. We have a viable change of coordinates which, if we glue back along that identification, allows a thickened boundary of the neighborhood to inherit a contact structure.

Question. How can we extend the contact structure over the whole solid torus?

The answer was deduced by Honda [Hon00], using foundational work of Giroux. Giroux deduced that, given any convex surface in a contact manifold, its dividing set encodes all essential behaviour of the contact structure in a neighborhood of the surface. This is known as Giroux's Flexibility Theorem. This allows for the gluing of two contact manifolds along convex surfaces with the same dividing set.

Honda combined these results, along with work by Kanda and Etnyre to produce the following result enumerating the number of tight contact structures admitted by $S^1 \times D^2$:

Theorem 2.1 (Honda, 2000)**.** Consider the tight contact structures on $S^1 \times D^2$ with convex boundary T^2, for which $\#\Gamma = 2$ and $s(T^2) = -p/q$, $p \geq q > 0$, p and q coprime. Fix a characteristic foliation F which is adapted to Γ_{T^2} . There exist exactly $|(r_0 + 1)(r_1 + 1)...(r_{k-1} + 1)(r_k)|$ tight contact structures on $S^1 \times D^2$ with this boundary condition, up to isotopy fixing T^2 . Here, $r0, , rk$ are the coefficients of the continued fraction expansion of $-p$:

$$-\frac{p}{q} = r_0 + \frac{1}{r_1 - \frac{1}{r_2 - \ldots r_k}}$$

Corollary 2.2 (Well-definedness when $p/q = 1/k$)**.** There is a unique tight contact structure on $S^1 \times D^2$ having convex boundary with these dividing curves.

Corollary 2.3. The construction above when $p/q = -1$ agrees with that of previous surgery descriptions. Indeed, the tubular neighborhoods of K in the handle description are related to the ones here by the contactomorphism sending $(q_1, q_2, \theta) \mapsto (-x, y, z)$.

3 Arnold Chord Conjecture

The Arnold Chord Conjecture in 3-dimensions was proved by Hutchings in [HT11]. The original conjecture is stated in [Arn86]. The proof uses a result of Kronheimer and Mrowka in [KM07] on cobordism maps in monopole Floer homology.

3.1 Statement and Outline of Proof

Let (Y, ξ) be a coorientable contact 3-manifold, with contact form λ. We denote by the Reeb vector field of this contact form by R. A *Reeb chord* of a Legendrian knot K in Y is an integral curve $\gamma : [0, 1] \to Y$ of R such that $\gamma(0), \gamma(1) \in K$. A result of Hutchings, 2011, previously a conjecture of Arnold, 1986, on the existence of Reeb chords is the following.

Theorem 3.1 ([HT11])**.** Let Y be a closed oriented 3-manifold with a contact form λ. Then every Legendrian knot in (Y, λ) has a Reeb chord.

The proof makes use of cobordism maps arising from monopole Floer homology, the isomorphism between monopole Floer homology and embedded contact homology, and changes in embedded contact homology generators under Legendrian surgery. Roughly speaking, we can detect Reeb chords by understanding the effect of Legendrian surgery on Reeb dynamics. As shown in the previous section, the induced contact structure on the surgered region has Reeb flow that can cause certain Reeb chords to close up (or concatenate and close up), generating new Reeb orbits. Reeb orbits (in some sense) form the cycles in ECH, and the boundary map counts J-curves flowing between them. The addition of these new Reeb orbits, assuming sufficiently well-behaved/understood boundary maps, can nontrivially change the ECH groups, so that they are no longer isomorphic to the original ones. The proof of the Arnold chord conjecture is motivated by this reasoning. An outline of the proof is offered below.

1. Construct an explicit model for -1-surgery on Y in terms of a 2-dimensional handle attachment to its symplectization $Y \times [0, 1]$.

2. Analyze the cobordism between $Y_0 := Y$ and $Y_1 := Y(-1)$ arising from this handle attaching. Specifically, by restricting to a filtered view of Reeb orbits via action, describe the chain maps induced by the cobordisms.

3. Define measures of the failure of the induced maps on homology to be isomorphisms. Relate these measures to the existence of Reeb chords in the sense $ECH(X)_* : ECH(Y_1) \to ECH(Y_0)$ not an isomorphism $\implies$ Y_0 has a Reeb chord (of a certain finite action).

4. Invoke the following result from Seiberg-Witten Theory:

   Theorem 3.2. If Y_1 is obtained from a closed oriented 3-manifold Y_0 by surgery along a knot K, and if X denotes the corresponding smooth cobordism from Y_1 to Y_0, then the induced map on Seiberg-Witten Floer cohomology with $\mathbb{Z}/2$ coefficients, then

   $$\widehat{HM}^*(X) : \widehat{HM}^*(Y_1) \to \widehat{HM}^*(Y_0)$$

   is not an isomorphism.

3.2 Embedded Contact Homology

Recall the following information about Embedded Contact Homology. Let (Y, ξ) be a contact 3-manifold, with contact form λ. A *Reeb orbit* is a closed trajectory of the Reeb vector field associated to λ, the unique vector field satisfying $\iota_R d\lambda = 0$ and $\lambda(R) = 1$. For a Reeb orbit γ of period T, the linearized return map of γ is the endomorphism $P_\gamma : (\xi_{\gamma(0)}, d\lambda_{\gamma(0)}) \to (\xi_{\gamma(0)}, d\lambda_{\gamma(0)})$ generated by flowing along R. We assume that R has no degenerate orbits - that is, orbits whose linearized return map has eigenvalue 1.
We define the embedded contact homology of (Y, ξ) as follows. The generators are given by orbit sets $\alpha := \{(\alpha_i, m_i)\}$, where α_i are (nondegenerate) Reeb orbits, and $m_i \in \mathbb{N}_{>0}$. We impose the condition that

$m_i = 1$ if α_i is a hyperbolic orbit. We generate the chains freely over $\mathbb{Z}/2$ by these orbit sets, and the differential counts certain pseudoholomorphic curves in the symplectization of Y with asymptotic ends in orbit sets. More precisely, if we denote by

$$\mathcal{M}(\alpha,\beta) := \{u : (\Sigma, j_\sigma) \to (Y \times R, d(e^s\lambda)) \mid \lim_{\mathbb{R}\to+\infty} u = \alpha, \lim_{\mathbb{R}\to-\infty} u = \beta, I(u) = 1\},$$

then $\langle \partial\alpha, \beta\rangle = \#(\mathcal{M}(\alpha,\beta)/\mathbb{R})$. Here, $I(u)$ denotes the ECH index of u. Note that the index I depends only on the homology class of u, not the map itself.

3.3 Surgery Constructions

We follow the surgery construction laid out in Section 2.1. Consider the symplectization $Y \times [0,1]$ of Y, equipped with the standard symplectic structure $d(e^s\lambda)$. We can identify a neighborhood of $K \times \{1\} \subset Y \times [0,1]$ with the neighborhood $\{p_1^2 + p_2^2 \geq 1, (q_1, q_2) \in \mathbb{R}^2\} \subseteq \mathbb{R}^4$. Under this identification, $K \times \{1\} \sim \{p_1^2 + p_2^2 = 1, q_1 = q_2 = 0\}$, and locally $Y \sim \{p_1^2 + p_2^2 = 1\}$. By choosing $\epsilon > 0$ to be sufficiently small, the attaching region is identified with the standard tubular neighborhood of K in Y.
We now describe the important parts of the arising cobordisms, and investigate the image of the Reeb flow along these components.
The cobordism $X : Y_0 \to Y_1$ comprises $Y \times [0,1]$, along with an attached symplectic 2-handle $D^2 \times D^2$. Inside of X, there are nested cobordisms $X_{1/n}$, created by attaching a thinner and longer handle to $Y \times [0, 1/n]$. We think of the handles as

$$p_1^2 + p_2^2 \leq e^{2(1-1/n)}, \qquad q_1^2 + q_2^2 \leq 2^{1-n}\epsilon.$$

Let $\tilde{U}_{1/n}$ be the part of $Y_{1/n}$ in the handle. The Liouville vector field $V = \sum_i -p_i\partial_{p_i} + q_i\partial_{q_i}$ expands p and contracts q in positive time.

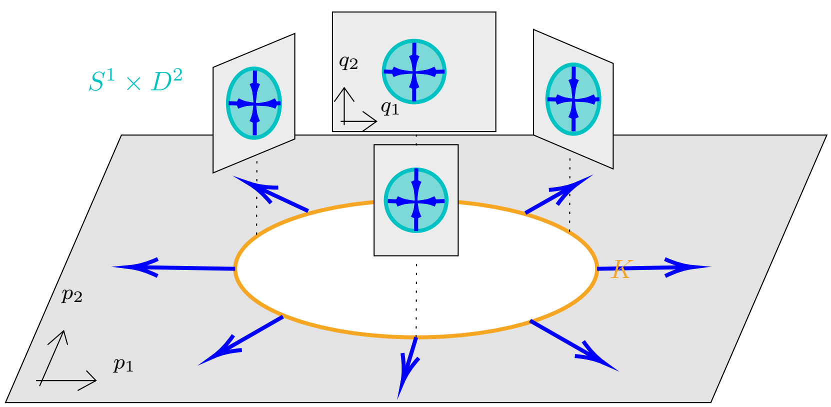

Figure 8: A (more accurate) picture of the flow of V.

More specifically, the negative-time flow ϕ of V gives a diffeomorphism $Y_{1/n} \to Y_{1/(n+1)}$. If we identify Y_0 with the hypersurface $\{p_1^2 + p_2^2 = e\}$, then we have the following lemma.

Lemma 3.3. The flow $\phi_{-1/n}$ of the Liouville vector field maps $Y_{1/n} \setminus \tilde{U}_{1/n} \hookrightarrow Y_0$.

Proof. Let $p_1 = r\cos(\theta)$ and $p_2 = r\sin(\theta)$. Then with respect to these cylindrical coordinates, along p_1, p_2 we have that

$$\begin{aligned} V_{r,\theta} &= -r\cos(\theta)(\partial_r\cos(\theta) - r\sin(\theta)\partial_\theta) - r\sin(\theta)(\partial_r\sin(\theta) + r\cos(\theta)\partial_\theta) \\ &= -r\cos^2(\theta)\partial_r - r\sin^2(\theta)\partial_r \\ &= -r\partial_r. \end{aligned}$$

To determine the flow of V along the coordinates $(p_1, p_2) \sim (r, \theta)$, let $\gamma(t) = (r(t), \theta(t))$ be an integral curve. Then

$$\dot{\theta}(t) = 0, \qquad \text{and} \qquad \dot{r}(t) = -r(t),$$

so $r(t) = Ce^{-t}$, where $C = r(0) = e^{1-1/n}$. Then $r(-1/n) = e^1 = e$. □

Remark 4. In this model, we local identify the cobordism with $\{1 \leq p_1^2 + p_2^2 \leq \epsilon\}$, so that $Y \times \{0\} \sim \{p_1^2 + p_2^2 = e\}$, and $Y \times \{1\} \sim \{p_1^2 + p_2^2 = 1\}$.

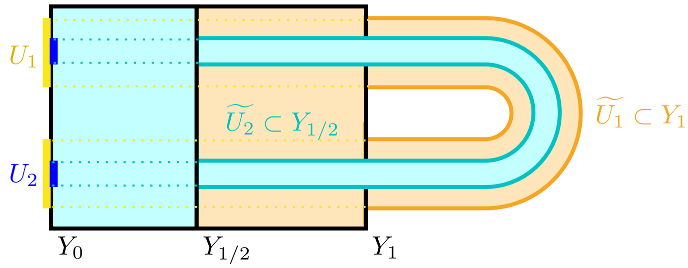

Figure 9: Nested handle attachings schematic.

Denote by $X_{1/n}$ the cobordism obtained by attaching the thinner handle to $X \times [0, 1/n]$. Then we obtain a sequence of nested cobordisms, with $X_{1/n} \simeq X_{1/(n+1)}$ for all n.

3.4 Existence of Reeb Chords

General Idea. The basic premise is that we'd like to understand how Reeb orbits are created via surgery, but if we work in too much generality, the cobordism maps are unhelpful and practically impossible to write down. Instead, we can help our problems by using filtered ECH. Filtering ECH by the action of the Reeb orbits allows us to get a better handle on the cobordism maps, and create viable ways of guaranteeing existence of Reeb orbits. From Seiberg-Witten theory, we know that the cobordism map induced on ECH is not an isomorphism. So it must either fail to be subjective somewhere, or fail to be injective somewhere. We can measure the point of where this failure occurs by quantities A and B. Briefly, A is the least action where ϕ fails to "capture" all of $H_*^L(Y_0, \lambda_0)$ (is not surjective). B is the least action where ϕ fails to "preserve" all of $H_*^L(Y_0, \lambda_0)$ (is not injective). If ϕ is not an isomorphism, then either $A < \infty$ or $B < \infty$. On the other hand, if there is no Reeb chord of action $< L$, then $A > L$ and $B > L$. In the case that there are no Reeb chords *at all*, then we may let L get arbitrarily large, so that $A = \infty$ or $B = \infty$. So we get a contradiction.

The well-behavedness of Reeb orbits up to a certain action is captured in the following lemma.

Lemma 3.4. Suppose that K does **not** have a Reeb chord of action $\leq L$. Then if n is sufficiently large,

1. The Reeb orbits of $\lambda_{1/n}$ with action $< e^{1/n}L$ avoid the region $\tilde{U}_{1/n}$.

2. $\phi_{-1/n}$ defines a bijection from the Reeb orbits of $\lambda_{1/n}$ with action $e^{1/n}L$ to the Reeb orbits of λ_0 with action $< L$.

Proof.

1. We suppose that the opposite holds: there is no n sufficiently large such that the Reeb orbits of $\lambda_{1/n}$ with action $< e^{1/n}L$ avoid the region $\tilde{U}_n$. Hence, choose a sequence of positive integers $\{n_k\}_{k\in\mathbb{N}}$, ordered strictly increasingly, such that for each k, there exists a Reeb orbit γ_k of λ_{1/n_k} with action $< e^{1/n_k}L$ that passes through the region $\tilde{U}_{n_k}$. Then $\phi_{-1/n}(\gamma_k \backslash \tilde{U}_{n_k})$ is a Reeb trajectory in Y_0 with ends on U_{n_k}. In particular, each γ_k has total action $\leq L$. Recall that each U_{n-k} is a tubular neighborhood of K, with decreasing thickness as $k \to \infty$. So, potentially passing to a subsequence of k and taking $k \to \infty$, the sequence $\{\gamma_k\}_{k\in\mathbb{N}}$ converges to a Reeb chord on K. Since each $\mathcal{A}(\gamma_k) \leq L$, $\mathcal{A}(\gamma) \leq L$. This contradicts our assumption, that K does **not** have a Reeb chord of action $\leq L$.

2. With 1. proved, we can choose n sufficiently large so that its result holds. Then for any γ a Reeb orbit of $\lambda_{1/n}$ of action $< e^{1/n}L$, γ avoids $\tilde{U}_{1/n}$. But outside of $U_{1/n}$, the cobordism induced by the surgery is trivial, and so the flow $\phi_{-1/n}$ sends Reeb trajectories to Reeb trajectories. Flowing by time $-1/n$ changes the action of the Reeb chords by a factor $e^{-1/n}$.

□

With this in hand, we begin by exploring cobordism maps on the level of filtered ECH. Recall that each cobordism $X_{1/n} : Y_0 \to Y_{1/n}$ induces a map on filtered ECH, which we denote by

$$\Phi^L(X_{1/n}, \lambda) : ECH_*^L(Y_{1/n}, \lambda_{1/n}) \to ECH_*^L(Y_0, \lambda_0).$$

This map satisfies several useful properties. In our circumstances, we can combine these properties to show that the cobordism maps are also well-behaved at a certain action level.

Lemma 3.5. Suppose that K does **not** have a Reeb chord of action $\leq L$. Then for n sufficiently large, the map

$$\Phi^L(X_{1/n}, \lambda) : ECH_*^L(Y_{1/n}, \lambda_{1/n}) \to ECH_*^L(Y_0, \lambda_0)$$

is a composition of an isomorphism

$$\Psi^{e^{1/n}L} : ECH_*^{e^{1/n}L}(Y_{1/n}, \lambda_{1/n}) \to ECH_*^L(Y_0, \lambda_0)$$

with the inclusion map $\iota^{L, e^{1/n}L}$.

Proof. If we choose n to be sufficiently large, then by the previous lemma, the Reeb orbits in $Y_{1/n}$ of action $\leq e^{1/n}L$ are in bijection with the Reeb orbits in Y_0 of action $\leq L$. Moreover, the Reeb orbits of $\lambda_{1/n}$ of action $\leq e^{1/n}L$ avoid the region of $Y_{1/n}$ that is contained in the handle.
In theory, when defining the induced cobordism maps, we should count pseudoholomorphic curves from from one orbit set at the positive end to another at the negative end in the completion of the cobordism. We will show that, restricting to an action level, only certain kinds of pseudoholomorphic curves are counted in the differential.
Let $X_{1/n}$ be as before, and consider the image of the hypersurface $Y_{1/n}$ under the negative flow of the Liouville vector field. If we flow for time $-1/n$, we obtain an embedding of $X^0_{1/n} := [0, 1/n] \times Y_{1/n} \hookrightarrow X_{1/n}$, where we identify $\{1/n\} \times Y_{1/n}$ with our original hypersurface $Y_{1/n}$, and $\{0\} \times (Y_{1/n} \setminus \tilde{U}_n)$ with $Y_0 \setminus U_n$.

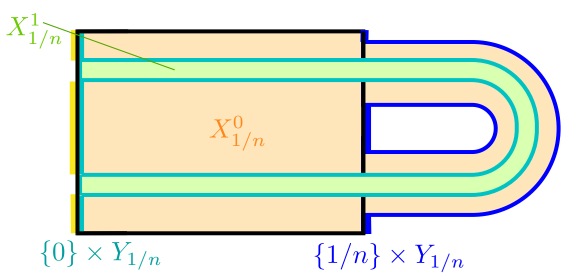

That is, the completion of $X_{1/n}$ splits up into the following parts: $(-\infty, 0] \times Y_0$, $X^1_{1/n} := X_{1/n} \setminus X^0_{1/n}$, and $[0, \infty) \times Y_{1/n}$ as in the figure below,

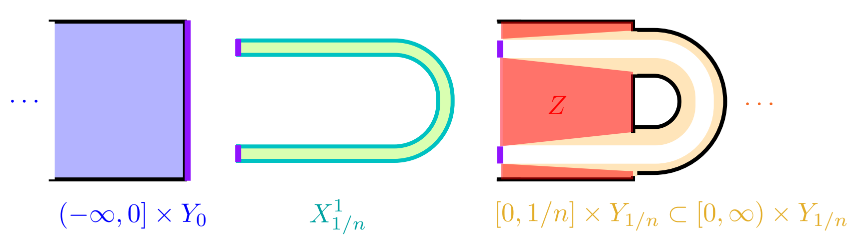

where we include $[0, 1/n] \times Y_{1/n}$ into this third component. Note that the cobordism $\overline{X_{1/n}}$ contains a product region $Z := [0, 1/n] \times (Y_{1/n} \setminus \tilde{U}_n)$. Let α be an ECH generator for $(Y_{1/n}, \lambda_{1/n})$ with $\mathcal{A}(\alpha) \leq e^{1/n}L$. We know that by lemma 3.4 that there exists a J-holomorphic curve from α to $\phi_{-1/n}(\alpha)$, an ECH generator for (Y_0, λ_0) with $\mathcal{A}(\phi_{-1/n}(\alpha)) \leq L$ that is just a union of covers of product cylinders. So in coordinates with

respect to the bijection between generators, the diagonal entries of the induced chain map $\psi^{e^{1/n}L}(X_{1/n})$: $ECC_*^{e^{1/n}L}(Y_{1/n}, \lambda_{1/n}) \to ECC_*^{L}(Y_0, \lambda_0)$ are strictly positive.
Claim (†). Suppose that C is any other J-holomorphic curve in $\overline{X_{1/n}}$ from α in $Y_{1/n}$ to some $\beta \neq \alpha$ in Y_0. Then

$$e^{-1/n}\mathcal{A}(\beta) > \mathcal{A}(\alpha.)$$

Proving this claim will show that there are **no** J-holomorphic curves from α to β **in the cobordism chain map** $\psi^{e^{1/n}L}(X_{1/n})$ with $\mathcal{A}(\alpha) \leq e^{1/n}L$ and $\mathcal{A}(\beta) \leq L$. In other words, the map $\psi^{e^{1/n}}(X_{1/n})$ is an upper triangular chain map, and so induces an isomorphism on homology. So let's prove it.
If such a C were to exist, then we can compare the actions by decomposing C in the cobordism above. Assuming that our choice of J is cobordism-admissible, the Liouville form λ on $\bar{X}_{1/n}$ splits as $e^s\lambda_0$ on $(-\infty, 0] \times Y_0$, λ on $X^1_{1/n}$, and $e^{s-1/n}\lambda_{1/n}$ on $[0, \infty) \times Y_{1/n}$. Since we are integrating in total over an **exact** symplectic cobordism, by Stokes theorem we have that

$$\begin{aligned} e^{-1/n}\mathcal{A}(\beta) - \mathcal{A}(\alpha) = &\int_{(-\infty,0]\times Y_0 \,\cap\, C} d(\lambda_0) \\ &+ \int_{X^1_{1/n} \,\cap\, C} d(\lambda) \\ &+ \int_{[0,\infty)\times Y_{1/n} \,\cap\, C} d(e^{-1/n}\lambda_{1/n}) \end{aligned}$$

The first and third integrals are pointwise nonnegative, and the second is pointwise positive. Thus,

$$e^{-1/n}\mathcal{A}(\beta) - \mathcal{A}(\alpha) \geq 0.$$

The action achieves equality iff C is a union of covers of product cylinders. □

4 The Possibility of Two Reeb Chords

Let Y_0 be a contact 3-manifold with contact form λ_0, containing a legendrian knot K. Suppose there exists exactly one Reeb chord c of K, which under Legendrian surgery closes it up to give a Reeb orbit h of action $L' := \mathcal{A}(h)$.

Question. What does the chain map look like in this case?

References

[Arn86] V. I. Arnold. The first steps of symplectic topology. *Uspekhi Mat. Nauk*, 41(6(252)):3–18, 229, 1986.

[Avd23] Russell Avdek. Combinatorial Reeb dynamics on punctured contact 3-manifolds. *Geom. Topol.*, 27(3):953–1082, 2023. doi:10.2140/gt.2023.27.953.

[BEE12] Frédéric Bourgeois, Tobias Ekholm, and Yasha Eliashberg. Effect of Legendrian surgery. *Geom. Topol.*, 16(1):301–389, 2012. doi:10.2140/gt.2012.16.301. With an appendix by Sheel Ganatra and Maksim Maydanskiy.

[Eli90] Yakov Eliashberg. Topological characterization of Stein manifolds of dimension > 2. *Internat. J. Math.*, 1(1):29–46, 1990. doi:10.1142/S0129167X90000034.

[Etn06] John B Etnyre. Lectures on contact geometry in low-dimensional topology. *arXiv preprint math/0610798*, 2006.

[Gei08] Hansjörg Geiges. *An introduction to contact topology*, volume 109 of *Cambridge Studies in Advanced Mathematics*. Cambridge University Press, Cambridge, 2008. doi:10.1017/CBO9780511611438.

[Hon00] Ko Honda. On the classification of tight contact structures. I. *Geom. Topol.*, 4:309–368, 2000. doi:10.2140/gt.2000.4.309.

[HT11] Michael Hutchings and Clifford Henry Taubes. Proof of the Arnold chord conjecture in three dimensions 1. *Math. Res. Lett.*, 18(2):295–313, 2011. doi:10.4310/MRL.2011.v18.n2.a8.

[KM07] Peter Kronheimer and Tomasz Mrowka. *Monopoles and three-manifolds*, volume 10 of *New Mathematical Monographs*. Cambridge University Press, Cambridge, 2007. doi:10.1017/CBO9780511543111.

[Wei91] Alan Weinstein. Contact surgery and symplectic handlebodies. *Hokkaido Math. J.*, 20(2):241–251, 1991. doi:10.14492/hokmj/1381413841.