



Geometry & Topology

Volume 27 (2023)

Cabling in terms of immersed curves

JONATHAN HANSELMAN

LIAM WATSON



Cabling in terms of immersed curves

JONATHAN HANSELMAN

LIAM WATSON

In joint work with J Rasmussen (Proc. Lond. Math. Soc. (3) 125 (2022) 879–967), we gave an interpretation of Heegaard Floer homology for manifolds with torus boundary in terms of immersed curves in a punctured torus. In particular, knot Floer homology is captured by this invariant (arXiv 1810.10355). Appealing to earlier work of the authors on bordered Floer homology (Geom. Topol. 27 (2023) 823–924), we give a formula for the behaviour of these immersed curves under cabling.

57M25, 57M27

Knot Floer homology, as introduced by Ozsváth and Szabó [18] and Rasmussen [20], provides a categorification: Given a knot K in the three-sphere, this invariant is a bigraded vector space $\bigoplus_{a,m \in \mathbb{Z}} \widehat{HFK}_m(K, a)$ with the property that

$$\sum_{a,m} (-1)^m \dim(\widehat{HFK}_m(K, a)) t^a$$

recovers the (symmetrized) Alexander polynomial. This polynomial knot invariant satisfies natural properties associated with operations on knots; for instance, it is well behaved under cabling. Understanding how this particular property manifests at the categorified level drove some of the early calculations of knot Floer homology; see in particular work of Hedden [7; 8; 9].

Bordered Floer homology provides an essential tool for studying decompositions of three-manifolds along essential tori; see Lipshitz, Ozsváth and Thurston [15]. They laid out a framework of bimodules, of relevance to satellite operations, in [14]. The work of Levine [13], Hom [11], and Petkova [19], for example, puts this to use in an essential way. In the setting of manifolds with a single toroidal boundary component, the relevant bordered invariants have been recast in terms of immersed curves in the once-punctured torus; see Hanselman, Rasmussen and Watson [4; 5]. For the purpose of this note, the examples of interest will be provided by the complement of a knot in the three-sphere; our aim is to establish formulas for how these invariants behave under cabling. Namely,




Figure 1: The Heegaard Floer homology for the $(2, -1)$ - and $(2, 1)$ -cables of the right-hand trefoil; the invariant for the trefoil complement is shown in grey.

for a knot K in S^3 , let $K_{p,q}$ denote the (p, q) -cable of K , and denote the respective knot complements by $M = S^3 \setminus \nu(K)$ and $M_{p,q} = S^3 \setminus \nu(K_{p,q})$; given the immersed multicurve $\widehat{HF}(M)$, we wish to describe $\widehat{HF}(M_{p,q})$ explicitly. For example, immersed curves for two cables of the right-hand trefoil are illustrated in Figure 1; the expert reader already familiar with the passage from $\widehat{HFK}(K)$ to $\widehat{HFK}(K_{p,q})$ should compare these pictures with the detailed calculations of Hedden [7] or Ozsváth, Stipsicz and Szabó [17]. Our calculation makes an explicit appeal to a bordered trimodule calculated by the first author [3], which was reinterpreted combinatorially in work of the authors predating the immersed curves invariant [6]. Indeed, central to this article is the work of translating our merge operation (described in terms of loop calculus) into the language of immersed curves (Section 1); cabling is then seen as a special case of the merge operation (Section 2).

Recall from [4; 5] that, for a (connected, orientable) three-manifold M with torus boundary, the invariant $\widehat{HF}(M)$ takes the form of a collection of immersed curves, possibly decorated with local systems, in the punctured torus $T_\bullet = \partial M \setminus z$, where z is some fixed basepoint in ∂M . If we choose a pair of parametrizing curves (α, β) on ∂M , then T_\bullet can be identified with the square $[0, 1] \times [0, 1]$ with opposite sides identified such that α runs in the positive vertical direction, β runs in the positive horizontal direction, and the puncture z is identified with $(0, 0)$. For a knot complement, there is a preferred choice of parametrizing curves, (μ, λ) , where μ is the meridian and λ is the Seifert longitude. The invariant $\widehat{HF}(M)$ comes equipped with grading data which, among other things, specifies a lift of these curves to the punctured cylinder

$\bar{T}_\bullet = (\mathbb{R}^2 \setminus \mathbb{Z}^2) / \langle \lambda \rangle$; in the standard framing, this can be identified with $(\mathbb{R}/\mathbb{Z}) \times \mathbb{R}$ with punctures at each lattice point $(0, n)$. (Note that this is the point of view taken in the presentation of the invariants in Figure 1: in each rectangle, the sides are identified to form a cylinder.) Thus as a graded object it makes sense to view $\widehat{HF}(M)$ as a collection of closed immersed curves $\gamma = (\gamma_0, \dots, \gamma_n)$ in \bar{T}_\bullet , defined up to homotopy in \bar{T}_\bullet , possibly decorated with local systems. For knots in S^3 , these curves have the property that, possibly after a homotopy, the curve set intersects the vertical line $\{\frac{1}{2}\} \times \mathbb{R}$ exactly once; we will always assume that γ_0 is the curve component containing this intersection. In other words, γ_0 wraps around the cylinder exactly once, while the remaining $\gamma_{i>0}$ can be confined to a neighbourhood of the vertical line through the punctures. We remark also that, while γ may carry nontrivial local systems, γ_0 always carries the trivial 1-dimensional local system (otherwise the rank of \widehat{HF} of the meridional filling of K would be greater than one). Finally, it is sometimes convenient to work in the plane $\tilde{T}_\bullet = \mathbb{R}^2 \setminus \mathbb{Z}^2$ rather than the cylinder, with the multicurve γ lifting to one that is invariant under translation by λ ; note that in this cover γ_0 lifts to a single periodic curve while each $\gamma_{i>0}$ lifts to infinitely many copies of the same curve.

We will show that the (p, q) -cable operation acts on $\widehat{HF}(M)$ by applying a particular diffeomorphism to the plane. Let $g_{p,q}$ be a diffeomorphism of \mathbb{R}^2 defined on the lattice \mathbb{Z}^2 by sliding each lattice point leftward along lines of slope $\frac{q}{p}$ until they first meet a vertical line $x = np$ for some integer n . Note that $g_{p,q}$ does not fix the lattice \mathbb{Z}^2 but rather takes it to $p\mathbb{Z} \times \frac{1}{p}\mathbb{Z}$; let $f_{p,q}$ be the composition of this map with vertical stretching by a factor of p and horizontal compression by a factor of p , so that $f_{p,q}$ takes \mathbb{Z}^2 to \mathbb{Z}^2 , followed by a vertical shift of $\frac{1}{2}(p-1)(q-1)$. We remark that the vertical shift is forced by the symmetry of the curves $\widehat{HF}(M)$ for any M and our convention that these curves are centred at height $\frac{1}{2}$; with this convention understood we will generally ignore the vertical positioning of the curves, but it is sometimes helpful to keep track of this vertical translation explicitly. The map $f_{p,q}$ is not linear, though in some sense it is as close to being linear as possible: it is the composition of linear transformations, which can each be realized as a sequence of plane shears, with a single *fractional plane shear* (defined in Section 2).

Theorem 1 *If γ is the immersed multicurve associated with K , $\gamma_{p,q}$ is the immersed multicurve associated with $K_{p,q}$, and $\tilde{\gamma}$ and $\tilde{\gamma}_{p,q}$ are the corresponding lifts to $\tilde{T}_\bullet = \mathbb{R}^2 \setminus \mathbb{Z}^2$, then $\tilde{\gamma}_{p,q}$ is homotopic to $f_{p,q}(\tilde{\gamma})$.*

Note that $f_{p,q}$ is periodic with period p in the horizontal direction, so it makes sense to view $f_{p,q}$ as a map from the cylinder $p\bar{T} := (\mathbb{R}/p\mathbb{Z}) \times \mathbb{R}$ to $\bar{T} := (\mathbb{R}/\mathbb{Z}) \times \mathbb{R}$ taking




Figure 2: Computation of the immersed curve associated with the $(3, 2)$ -cable of the right-hand trefoil, starting from the trefoil curve pictured on the left. The two middle diagrams are two ways of thinking about the construction starting from three copies of the trefoil curve: we either slide lattice points along lines of slope $\frac{2}{3}$ or we stagger the heights of the three copies of the trefoil curve and then slide lattice points horizontally. Either way the result is the curve on the right.

lattice points to lattice points. With this view, the process of computing $\widehat{HF}(M_{p,q})$ from $\widehat{HF}(M)$ is to lift from \bar{T} to $p\bar{T}$ and then apply $f_{p,q}$. In practice, this amounts to drawing p copies of $\widehat{HF}(M)$ in sequence, perturbing the curve by pushing lattice points along lines of slope $\frac{q}{p}$ until they all lie on the same vertical line, and then scaling vertically by a factor of p . This procedure is depicted in Figure 2 for the case of the $(3, 2)$ -cable of the right-hand trefoil. It is helpful to note that the procedure of pushing lattice points along lines of slope $\frac{q}{p}$ can equivalently be viewed as drawing p copies of the input curve with staggered heights and then translating punctures horizontally. In practice, the process for computing $\widehat{HF}(M_{p,q})$ from $\widehat{HF}(M)$ amounts to a three-step process:

- (1) draw p copies of $\widehat{HF}(M)$ next to each other, each scaled vertically by a factor of p , staggered in height such that each copy of the curve is a height of q units lower than the previous copy;
- (2) connect the loose ends of the successive copies of the curve; and
- (3) translate the pegs horizontally so that they lie in the same vertical line, carrying the curve along with them.

Numerical concordance invariants extracted from curves

As an illustration of Theorem 1 at work, we can revisit the work of Hedden [8; 9] and Van Cott [21], culminating in a result of Hom [11], which establishes the behaviour of the τ -invariant under cabling. Since $\tau(K)$ can be easily extracted from the immersed

multicurve $\widehat{HF}(M)$, we can recover this cabling behaviour from Theorem 1. The same is true for some other numerical invariants. We begin by making an observation that is implicit in earlier work. Let $\boldsymbol{\gamma} = (\gamma_0, \dots, \gamma_n)$ denote the underlying set of immersed curves for $\widehat{HF}(M)$, with γ_0 the unique component which wraps around the cylinder. This component is itself an invariant of K , so it will sometimes be convenient to express it as $\gamma_0(K)$.

Proposition 2 *The curve $\gamma_0(K)$ is an invariant of the concordance class of K .*

Proof This follows from Hom [12] and the recipe for deriving $\widehat{HF}(M)$ from $CFK^-(K)$ described in [5, Section 4]. The concordance invariant described in [12] is the smallest direct summand of $CFK^-(K)$, up to homotopy equivalence, which supports the homology of S^3 . The set of immersed curves derived from this summand is a subset of the immersed curves $\widehat{HF}(M)$ which necessarily contains $\gamma_0(K)$. This subset of curves, and in particular $\gamma_0(K)$, is thus a concordance invariant. \square

Note that the concordance invariant described in [12] is slightly stronger than γ_0 since some information may be lost when passing from complexes to immersed curves (namely, diagonal arrows are ignored). In fact, $\gamma_0(K)$ carries exactly the same information as the ϵ -equivalence class of K defined in [12]. Any number that can be extracted from γ_0 is automatically a concordance invariant, and several familiar concordance invariants can be defined in this way. The two most common are τ and Hom's ϵ -invariant, which are extracted from γ_0 as follows: Starting on the section of γ_0 which wraps around the back of the cylinder — say, at the unique intersection of γ_0 with the line $x = \frac{1}{2}$ — and moving rightward along γ_0 , let a denote the first intersection of γ_0 with the vertical axis $x = 0$. Then the integer τ records the height of the intersection point a (here we use a discrete notion of height given by the greatest integer less than the y -coordinate of a). Continuing along γ_0 from a , one of three things can happen: γ_0 can turn downwards, it can turn upwards, or it can continue straight to wrap around the cylinder. This is recorded by ϵ , which takes the values $+1$, -1 or 0 in these three cases, respectively. (Both of these observations are made in [5].) Note that if $\epsilon = 0$ then there is only one intersection of γ_0 with the vertical axis, so γ_0 is simply a horizontal curve, which is the immersed curve associated with the complement of the unknot. Now consider the effect of cabling on each of these invariants. Throughout, let $\gamma_0 = \gamma_0(K)$ and let $\gamma'_0 = \gamma_0(K_{p,q})$.

Theorem 3 (Hom [11, Theorem 2]) *If $\epsilon(K) = \pm 1$ then $\epsilon(K_{p,q}) = \epsilon(K)$; and if $\epsilon(K) = 0$ then $\epsilon(K_{p,q}) = \epsilon(T_{p,q})$.*

A quick reproof of Theorem 3 By Theorem 1, γ'_0 is obtained from γ_0 by placing p copies of γ_0 next to each other, with appropriate vertical shifts, and compressing them into one vertical line. The first intersection of γ'_0 with the vertical axis thus comes from the first intersection of the first copy of γ_0 with the vertical axis, and clearly if γ_0 turns upward or downward at this point then γ'_0 does also. On the other hand, if $\epsilon(K) = 0$ then γ_0 is simply a horizontal line, the same as the curve associated with the unknot. It follows that γ'_0 agrees with $\gamma_0(T_{p,q})$, since $T_{p,q}$ is the (p, q) -cable of the unknot, and thus $\epsilon(K_{p,q}) = \epsilon(T_{p,q})$. \square

The value $\epsilon(T_{p,q})$ was also computed in [11, Theorem 2]; we can recover this computation by viewing $T_{p,q}$ as the (p, q) -cable of the unknot. In this case γ_0 is horizontal in \bar{T} and lifts to a horizontal line in $p\bar{T}$. To compute γ'_0 from this we shift the i^{th} column downwards by $\frac{iq}{p}$ and then compress horizontally. If $|q| = 1$ then $T_{p,q}$ is unknotted and we must have $\epsilon(T_{p,q}) = 0$; indeed, in this case every column of lattice points shifts by less than one unit, so it is possible for γ'_0 to remain horizontal despite the shift. On the other hand, if $q > 1$ then the shift causes γ'_0 to turn downwards, so $\epsilon(T_{p,q}) = +1$; similarly, if $q < -1$ then γ'_0 turns upward and $\epsilon(T_{p,q}) = -1$.

Theorem 4 [11, Theorem 1] *If $\epsilon(K) = \pm 1$ then $\tau(K_{p,q}) = p\tau(K) + \frac{1}{2}(p-1)(q \mp 1)$; and if $\epsilon(K) = 0$ then $\tau(K_{p,q}) = \tau(T_{p,q}) = (-1)^{\text{sign}(q)}\frac{1}{2}(p-1)(|q|-1)$.*

A quick reproof of Theorem 4 The first intersection of γ'_0 with the vertical axis clearly comes from the first intersection of the first copy of γ_0 with the vertical axis. This intersection occurs between the lattice points at height $\tau(K)$ and $\tau(K) + 1$; after applying $f_{p,q}$ and the appropriate vertical shift, these lattice points map to heights $h_1 = p\tau(K) + \frac{1}{2}(p-1)(q-1)$ and $h_2 = p\tau(K) + p + \frac{1}{2}(p-1)(q-1)$. Note that there are $p-1$ lattice points between these two heights; whether or not γ'_0 first intersects the vertical axis above or below these points depends on the behaviour of γ_0 just after it crosses the vertical axis, as pictured in Figure 3. If γ_0 turns downward (ie if




Figure 3: Calculating $\tau(K_{p,q})$.

$\epsilon(K) = +1$), then γ'_0 will also turn downward and meet the vertical axis just above height h_1 ; thus $\tau(K_{p,q}) = h_1$. If γ_0 turns upward (ie if $\epsilon(K) = -1$), then γ'_0 will also turn upwards and meet the vertical axis just below height h_2 ; thus $\tau(K_{p,q}) = h_2 - 1 = p\tau(K) + \frac{1}{2}(p-1)(q+1)$. Finally, if $\epsilon(K) = 0$ then γ_0 agrees with the curve invariant of the unknot, so γ'_0 is the curve associated with $T_{p,q}$, and thus $\tau(K_{p,q}) = \tau(T_{p,q})$. In particular, if $q > 1$ then γ'_0 bends down after its first intersection with the vertical axis, and as above $\tau(K_{p,q}) = h_1 = \frac{1}{2}(p-1)(q-1)$. If $q < -1$ then γ'_0 bends upward and $\tau(K_{p,q}) = h_2 - 1 = -\frac{1}{2}(p-1)(-q-1)$, while if $|q| = 1$ then γ'_0 is horizontal and $\tau(K_{p,q}) = 0$. □

Other concordance invariants can be extracted from $\gamma_0(K)$. For instance, for any positive integer i , the invariant $\phi_i(K)$ introduced recently by Dai, Hom, Stoffregen and Truong [1] counts the number of left arcs of γ_0 of length i , where a left arc of length i refers to a segment of γ_0 connecting successive intersections with the vertical axis whose height differ by i which does not wrap around the cylinder and which lies to the left of the vertical axis. These arcs are counted with sign coming from the orientation of γ_0 , with downward oriented arcs counting positively.¹ Like τ , the integers ϕ_i are of particular interest in the study of knot concordance because they are additive under connected sum; that is, they define concordance homomorphisms.

Returning to cabling, the behaviour of the invariants ϕ_i is more complicated. In particular, $\phi_i(K_{p,q})$ does not depend only on $\phi_i(K)$, or even on the collection of invariants $\tau(K)$, $\epsilon(K)$ and $\phi_j(K)$ for all j . In order to express the effect of cabling we need to keep track of how each left arc in $\gamma_0(K)$ behaves at each end. For example, we can define refined invariants ϕ_i^{++} , ϕ_i^{+-} , ϕ_i^{-+} and ϕ_i^{--} encoding the signed count of four different types of length i left arcs in γ_0 . The type is determined by the direction γ_0 turns at each end of the segment; $+$ indicates that γ_0 turns upward and $-$ indicates that γ_0 turns downward, with the first sign indicating the behaviour at the top of the arc and the second sign indicating the behaviour at the bottom of the arc, as in Figure 4. Note that $\phi_i = \phi_i^{++} + \phi_i^{+-} + \phi_i^{-+} + \phi_i^{--}$. With these extra quantities defined, it is possible to derive explicit formulas for $\phi_i(K_{2,1})$. More generally, we could derive explicit formulas for $\phi_i(K_{p,q})$ in a similar way; the key difference is that the notions of turning up or turning down used in defining the invariants $\phi_i^{\pm\pm}$ are dependent on p

¹This is a straightforward translation of the definition of ϕ_i given in [1] to the language of immersed curves. The *standard complex* described in [1] corresponds precisely to the component γ_0 of $\widehat{HF}(M)$. The integers ϕ_i count horizontal arrows of length i in the standard complex, which correspond to length i right arcs in γ_0 . By symmetry, we can equivalently count length i left arcs in γ_0 .




Figure 4: Four cases complete the proof.

and q (here turning up means moving upward vertically or rightward with slope greater than $\frac{q}{p}$). The formulas are cumbersome so, rather than derive the general case, we focus instead on the special case of $(2, 1)$ -cabling.

Proposition 5 For $i > 1$ all the variants of $\phi_i(K_{2,1})$ are either determined by

$$\phi_{2n}^{\pm\pm}(K_{2,1}) = \phi_n^{\pm\pm}(K), \quad \phi_{2n\pm 1}^{\pm\mp}(K_{2,1}) = \phi_n^{\pm\mp}(K)$$

or they are trivial. In particular, for $n \geq 1$,

$$\phi_{2n}(K_{2,1}) = \phi_n^{++}(K) + \phi_n^{--}(K), \quad \phi_{2n+1}(K_{2,1}) = \phi_n^{+-}(K) + \phi_{n+1}^{-+}(K).$$

Proof The curve $\gamma'_0 = \gamma_0(K_{2,1})$ is constructed in three steps: take two consecutive copies of $\gamma_0 = \gamma_0(K)$; scale vertically by a factor of two and shift the second copy of γ_0 down one unit; and compress horizontally (compare Figure 2). Before compressing horizontally, we can divide this curve into two (nonconnected) subcurves which lie to the left and right of vertical line through the first column of lattice points; let γ_L and γ_R denote the images of these subcurves after horizontal compression, so that $\gamma'_0 = \gamma_L \cup \gamma_R$.

The key observation is that every component of γ_R lies to the right of every even height lattice point, and therefore any left arc on γ'_0 which lies in γ_R must have length 1. It is also clear that any left arc of γ'_0 intersects at most one component of γ_L , since otherwise it contains a full component of γ_R which must lie to the right of some lattice point. Thus each left arc of length greater than 1 in γ'_0 comes from a component of γ_L , which in turn comes from a left arc of the first copy of γ_0 . Conversely, every left arc of length i in the first copy of γ_0 gives rise to exactly one left arc in γ'_0 , which has the same end behaviour. The length of this new arc depends on the end behaviour: it is $2i - 1$ for $-+$ arcs, $2i$ for $++$ or $--$ arcs, and $2i + 1$ for $+-$ arcs (see Figure 4). \square

We will say that γ_0 has a unique maximal-length left arc of type $++$ and length N if $\phi_N^{++}(K) = 1$, $\phi_i^{++}(K) = 0$ for all $i > N$, and $\phi_i^{+-}(K) = \phi_i^{-+}(K) = \phi_i^{--}(K) = 0$ for all $i \geq N$. The following is an immediate consequence of the formulas above:




Figure 5: Immersed curves for the first few iterated $(2, 1)$ -cables of the right-hand trefoil. These are also the distinguished curve γ_0 for the knots K_0, K_1, K_2 and K_3 from Corollary 7. The longest left arc (highlighted) is stretched by a factor of two with each cabling iteration; thus, the length of the longest left arc for K_n is 2^n .

Proposition 6 *If $\gamma_0(K)$ has a unique maximal-length left arc of type $++$ and length N , then $\gamma_0(K_{2,1})$ has a unique maximal-length left arc of type $++$ and length $2N$.*

Consider for example iterated $(2, 1)$ -cables of the right-hand trefoil $T_{2,3}$; the immersed curves for the first few of these knots are shown in Figure 5. The immersed curve $\gamma_0(T_{2,3})$ has only one left arc, which has type $++$ and length 1. If we repeatedly $(2, 1)$ -cable this knot, there is always a single left arc of maximal length, which always has type $++$, and the length of this arc doubles in length with each iteration.

In [1], the concordance invariants ϕ_i were used to identify a \mathbb{Z}^∞ direct summand in the topologically slice smooth concordance group \mathcal{C}_{TS} ; see also [17]. The relevant

infinite family of knots is built from cables of a certain knot D , the untwisted positively clasped Whitehead double of $T_{2,3}$. More precisely, the family of knots is given by $D_{n,n+1} \# -T_{n,n+1}$. Using Proposition 6, we can construct another \mathbb{Z}^∞ summand from D by instead taking iterated $(2, 1)$ -cables. The key properties of D are that

- (i) the Alexander polynomial of D is trivial, and
- (ii) the distinguished component γ_0 associated to D agrees with $\gamma_0(T_{2,3})$.

The knot D can be replaced with any other knot which shares these two properties — an example of a hyperbolic knot with this property is 15n113775.

Corollary 7 *Let $K = K_0$ be a knot for which $\Delta_K(t) = 1$ and $\gamma_0(K) = \gamma_0(T_{2,3})$. For $n \geq 1$ let K_n be the $(2, 1)$ -cable of K_{n-1} . The knots $\{K_n\}_{n=0}^\infty$ span a \mathbb{Z}^∞ summand of \mathcal{C}_{TS} .*

Proof According to a result of Freedman, $\Delta_K(t) = 1$ implies that K is topologically slice [2]. The $(2, 1)$ -cable of a topologically slice knot is topologically concordant to the $(2, 1)$ -cable of the unknot, which is the unknot; thus, by induction, K_n is topologically slice for all n . On the other hand, K_0 has a unique maximal-length left arc of type $++$ and length 1, so Proposition 6 and induction implies that K_n has a unique maximal-length left arc of type $++$ and length 2^n . In particular, for each n we have $\phi_{2^n}(K_n) = 1$ and $\phi_i(K_n) = 0$ for all $i > 2^n$. Since each ϕ_i is a concordance homomorphism, it follows that the knots are linearly independent in the smooth concordance group. Moreover, it is straightforward to see that the homomorphism

$$\bigoplus_{i=0}^\infty \phi_{2^i} : \mathcal{C}_{\text{TS}} \rightarrow \bigoplus_{i=0}^\infty \mathbb{Z}$$

is an isomorphism when restricted to the span of the K_n . Indeed, it follows from the information above that the images of the K_n form a basis for \mathbb{Z}^∞ (see for instance [17, Proposition 6.4]). In fact, by the remark below we also have that $\phi_i(K_n) = 0$ for all $i < 2^n$, so the image of K_n is the standard i^{th} basis vector of \mathbb{Z}^∞ . \square

Remark 8 We leave the behaviour of ϕ_1 under $(2, 1)$ -cabling, which was not needed in the above application, as an exercise to the motivated reader, who will find that

$$\phi_1(K_{2,1}) = -\sum_{j \geq 1} \phi_j(K) + \begin{cases} 1 & \text{if } \tau(K) > 0, \\ 0 & \text{if } \tau(K) \leq 0. \end{cases}$$

Using corresponding formulas for variants of ϕ_1 along with those from Proposition 5 and induction on n , it can be shown that the knots K_n in Corollary 7 in fact satisfy

$\phi_i(K_n) = 1$ if $i = 2^n$ and $\phi_i(K_n) = 0$ otherwise. Thus the map

$$\bigoplus_{i=0}^{\infty} \phi_{2^i} : \text{Span}\{K_n\} \rightarrow \bigoplus_{i=0}^{\infty} \mathbb{Z}$$

defines an explicit isomorphism from the span of the K_n in the topologically slice concordance group to \mathbb{Z}^{∞} .

Curves that do not come from cables

In many cases Theorem 1 provides a simple obstruction to a knot being a nontrivial cable. For example, if the curve set corresponding to a knot contains a figure-eight component enclosing two adjacent lattice points (as is the case, for example, for any nontorus alternating knot; see [19]), then the knot is not a cable of a knot in S^3 . This is because any closed component of $\widehat{HF}(M_{p,q})$ comes from one copy of $\widehat{HF}(M)$ before the transformation of the plane and thus encloses only lattice points with the same height modulo p . If we restrict our attention to γ_0 , we can find knots which are not concordant to a nontrivial cable. An example of this is the knot $12n242$, whose γ_0 is pictured in Figure 6. Because the first left arc has length 2, if this curve comes from a (p, q) -cable then p must be at most 3. If $p = 1$, the cable operation is trivial. If $p = 2$, then the curve should not pass to the left of an odd-height lattice point (grey in Figure 6) after passing to the right of an even-height lattice point (black in the figure), but this clearly happens. Similarly, if $p = 3$ then the curve should not pass to the left of a point whose height is congruent to 2 mod 3, after passing to the right of a lattice point with a different height modulo 3, but this also happens. We thank Tye Lidman for asking us about the existence of such an example.



Figure 6

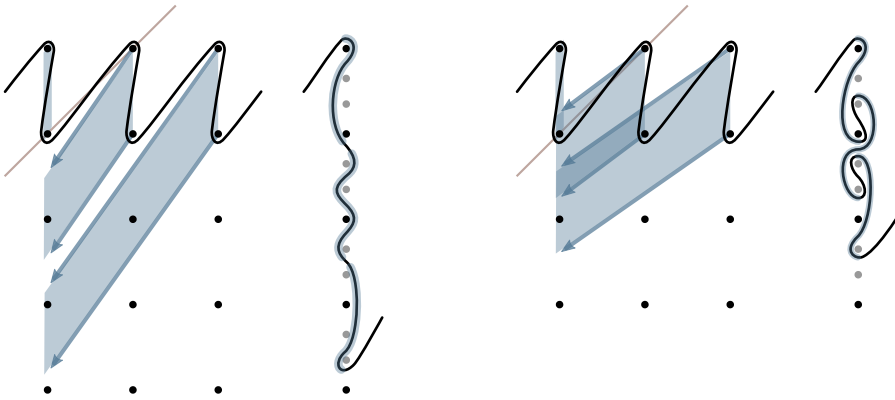


Figure 7: The (3, 4)-cable (left) and the (3, 2)-cable (right) of the right-hand trefoil.

L-space surgeries on cable knots

Other properties of the knot Floer homology of cable knots are made relatively transparent by Theorem 1. The following is a well-known property that was established by Hom [10] (building on work of Hedden [9]):

Theorem 9 [10, Theorem] *For any knot K in S^3 , $K_{p,q}$ admits a positive L -space surgery if and only if K admits a positive L -space surgery and $\frac{q}{p}$ is at least $2g - 1$, where g denotes the Seifert genus of K .*

Note that K admits a positive L -space surgery if and only if $\widehat{HF}(M)$ is a single curve which, apart from the segment that wraps around the cylinder, moves monotonically downward in the neighbourhood of the vertical axis (see [4, Section 7.5]). When this curve is pulled tight in the cylinder \tilde{T}_\bullet (or in the plane \tilde{T}_\bullet), the slope of the nonvertical segment is $2g - 1$. Following Theorem 1, we construct $\widehat{HF}(M_{p,q})$ from p columns of the lift of $\widehat{HF}(M)$ to \tilde{T}_\bullet by translating lattice points along lines of slope $\frac{q}{p}$.

A quick reproof of Theorem 9 If $\widehat{HF}(M)$ is oriented upward at any point apart from the nonvertical segment, it is clear the same will be true at the image of this point on $\widehat{HF}(M_{p,q})$; thus K having a positive L -space surgery is a necessary condition for $K_{p,q}$ to have one. Supposing K has a positive L -space surgery, it is clear that if $\frac{q}{p} > 2g - 1$ then the p copies of the downward-oriented portion of $\widehat{HF}(M)$ miss each other, so the resulting curve moves monotonically downward and $K_{p,q}$ has a positive L -space surgery. On the other hand, if $\frac{q}{p} < 2g - 1$ then these sections of curves overlap, forcing some backtracking in the resulting curve, implying that $K_{p,q}$ has no L -space surgeries. An example is given in Figure 7. □




Figure 8: The $(3, 2)$ -cabling operation interpreted as a plane tiling: three copies of the standard square tile (above) are carried to a new regular tile in $\mathcal{T}_{3,2}$ (below) under the operation $f_{p,q}$ appearing in Theorem 1. To illustrate, the image of the longitude has been included (gradually homotoped to a simpler form moving rightward), which recovers the invariant associated with the right-hand trefoil as expected.

Cabling via tiling

From Theorem 1 it is possible to interpret cabling in terms of plane tilings. That is, in a visual summary of the above discussion, we record the following:

Corollary 10 *For every relatively prime pair (p, q) there is a periodic tiling $\mathcal{T}_{p,q}$ of the plane, unique up to lattice-fixing planar isotopy, such that $\gamma(K_{p,q})$ is the image of $\gamma(K)$ under the transformation taking the lattice \tilde{T}_\bullet to $\mathcal{T}_{p,q}$.*

Proof This is a simple reformulation of Theorem 1: Consider the standard square tiling of the plane \tilde{T}_\bullet defined by the preferred (μ, λ) -framing. The image of p square tiles aligned horizontally, under the application of $f_{p,q}$, gives a tile in a periodic tiling of the plane. □

This is best illustrated in an example, and we have shown the tiling associated with $(3, 2)$ -cabling in Figure 8. Note that this point of view comes with a built-in sanity check: one can check that the image of a longitudinal curve under the transformation to $\mathcal{T}_{p,q}$ is the immersed curve $\gamma(T_{p,q})$. Recalling that, as a polynomial in t , the Alexander polynomial satisfies

$$\Delta(K_{p,q}) = \Delta(K)|_{t^p} \cdot \Delta(T_{p,q}),$$

our formula has $\mathcal{T}_{p,q}$ playing the role of $\Delta(T_{p,q})$ in this formula while replacing t with t^p corresponds to the p repeated copies of $\gamma(K)$.

In general, one expects bimodules in bordered Floer homology (for manifolds with two boundary tori) to be associated with Lagrangians in $T_\bullet \times T_\bullet$. A simple first example of this is the bimodule associated with a diffeomorphism of the torus, where the (embedded) Lagrangian surface is the graph of the diffeomorphism. In that case we can interpret the action of the bimodule as follows: to compute the image of an immersed curve γ , we consider $\gamma \times T_\bullet$, intersect with the Lagrangian surface, and project to the second coordinate. Cabling bimodules provide a first glimpse at how this construction might be generalized to arbitrary bimodules. The diffeomorphism of the plane $f_{p,q}$ does not descend to a diffeomorphism of the torus, but, since $f_{p,q}$ is periodic and is determined by its effect on p consecutive tiles of the plane, it can be viewed as a p -valued function on T_\bullet ; that is, to each point in T_\bullet it associates an unordered tuple of p points in T_\bullet . The graph of this multivalued function is an (immersed) Lagrangian surface in $T_\bullet \times T_\bullet$, and the action of the bimodule on curves can be interpreted geometrically as before.

1 Immersed curves and the merge operation

For any orientable manifold M with torus boundary, the Heegaard Floer homology $\widehat{HF}(M)$ is an immersed multicurve in the marked torus ∂M [4], as introduced above. This view of the Heegaard Floer invariants of M arises from an interpretation of bordered Floer homology [15], and is closely related to the *loop calculus* introduced in [6]. This section builds a glossary between loop calculus [6] and immersed curves [4];




Figure 9: The torus algebra \mathcal{A} as the path algebra of a quiver with relations.

in the former we developed the machinery for understanding gluing pairs of manifolds along essential annuli in their boundaries, which we aim to interpret in terms of the immersed curves in the case of cabling knots in the three-sphere.

1.1 From puzzle pieces to curve segments

Assuming familiarity with some subset of [4; 6; 15], we give a very terse summary of the bordered invariants in order to set up the desired glossary.

The torus algebra \mathcal{A} is obtained as the path algebra of the quiver described in Figure 9. Let $\mathcal{I} \subset \mathcal{A}$ denote the subring of idempotents generated by ι_0 and ι_1 . Working over the two-element field \mathbb{F} , a type D structure over \mathcal{A} is a finite-dimensional left \mathcal{I} -module V together with a map $\delta: V \rightarrow \mathcal{A} \otimes_{\mathcal{I}} V$. This map must satisfy a compatibility condition equivalent to ensuring that $\partial(a \otimes x) = a \cdot \delta(x)$ is a differential on the \mathcal{A} -module $\mathcal{A} \otimes_{\mathcal{I}} V$.

There is a simple interpretation of the above data in terms of decorated graphs: the vertices encode the generating set (these come in two types \bullet and \circ , depending on the idempotents ι_0 and ι_1 , respectively) and, by passing to type D structures that are reduced, the directed edges are labelled by the set $\{1, 2, 3, 12, 23, 123\}$ in order to encode the coefficient maps; see Figure 11. These graphs can be naturally immersed in the marked torus or, more precisely, in the once-punctured torus with a fixed choice of 1-handle cocores cutting the surface into a disk. In our case, these cocores will always coincide with the preferred (μ, λ) -pair, since we are focussed on knots in S^3 . With this data in hand, we can decompose the torus into the familiar square patch with opposite edges identified. The type D structures of interest then are immersed train tracks (in the sense of Thurston [16]) where all of the vertices/switches lie on the horizontal or vertical edges; when such a train track comes from a three-manifold, the classification theorem proved in [4] tells us it is equivalent to an immersed multicurve, possibly decorated with local systems, which we denote by $\widehat{HF}(M)$ [4].

In the case where the local systems are trivial, we recover the class of loop-type manifolds considered in our earlier work [6] (see also [5, Section 1]). Central to




Figure 10: Segments of immersed curve in the cover of the marked torus, labelled to be consistent with the puzzle pieces given in [6]. The integer subscript $k > 0$ indicates the number of \circ generators in the segment. These letters can appear forwards or backwards in a cyclic word, so that \bar{a}_k runs against the direction indicated by the blue arrow. We can also extend our notation by setting $a_{-k} = \bar{a}_k, b_{-k} = \bar{b}_k, c_{-k} = \bar{c}_k, d_{-k} = \bar{d}_k$ and $d_0 = \bar{c}_0 = e$; note that then a segment with subscript k moves upward k units in the plane.

this is the observation that, when the type D structure in question can be represented by a valence 2 graph, it is possible to decompose along \bullet vertices into segments, each of which takes one of five possible forms as described in Figure 10 (compare [6, Figure 1]).² As a result, studying these type D structures amounts to a calculus for manipulating cyclic words in the infinite alphabet $\mathfrak{A} = \{a_k, b_k, c_k, d_k, e\}$ for all positive integers k . The segments corresponding to these letters may appear backwards as we traverse a loop; this is indicated by a *bar*. There are rules governed by the algebra restricting the letters that can be concatenated, which are most easily described by noting that each segment also corresponds to a segment of immersed curve as in Figure 10: if two curve segments share an endpoint, they must lie on opposite sides of the vertical near that point. (In [6], a puzzle piece convention is used to describe these rules.) Note that the a_k and the b_k correspond to the two types of stable chains introduced in [15], while c_k, d_k and e correspond to the three types of unstable chains. In fact, it makes sense to view the three types of unstable chain as part of a single family, and with this in mind we set $c_0 = \bar{e}$ and $c_{-k} = \bar{d}_k$. The example in Figure 11 explains this for the right-hand trefoil exterior.

Now consider a component of $\widehat{HF}(M)$, that is, an immersed curve γ decorated with a local system (V, Φ) of dimension n . Following [4], we can interpret this as a *curve-like*

²There are certain exceptional type D structures that cannot be decomposed in this way; however, these examples are not particularly important in this setting. The interested reader can consult [6] for a *dual* notation that decomposes along \circ vertices.

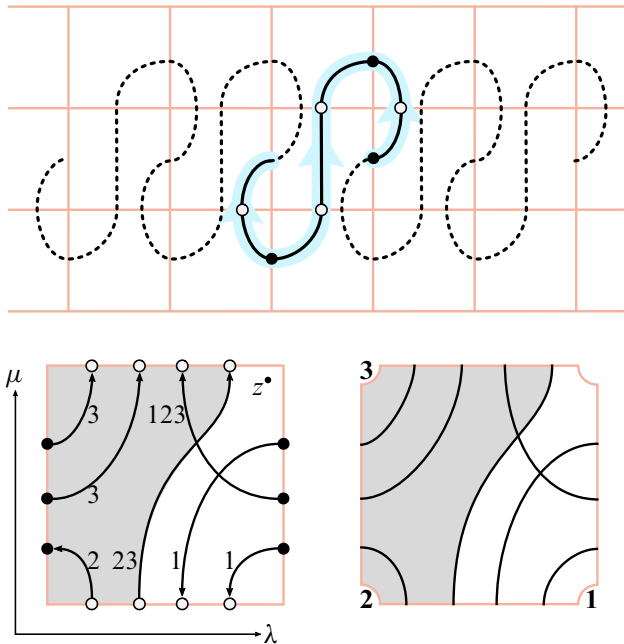


Figure 11: Three different views of the invariant associated with the exterior of the right-hand trefoil. In all three cases, we have fixed the preferred (μ, λ) -framing in order to present the torus boundary. On the lower left, the decorated graph describing the type D structure has been immersed in the marked torus as a train track. This description exhibits the redundancy in the edge labels: as shown in the lower right figure, the idempotents can be recovered from the horizontal and vertical edges while the coefficient maps are determined by which of the labelled corners are traversed by the curve segments (a region indicating the 23 edge is shaded). Finally, lifting the curve to the cover \tilde{T}_\bullet (or, as pictured, \tilde{T}_\bullet) makes obvious the cyclic word $a_1\bar{c}_2b_1$, which in [6] is referred to as a *loop*.

train track, which consists of n parallel copies of γ along with some additional edges that we may assume all lie on a portion of γ corresponding to a single segment (that is, along one letter of \mathfrak{A} as described above). When M is the complement of a knot K in S^3 , we may in fact assume that these edges lie on a segment of type a_k ; this is because the curve $\gamma_0(K)$ does not carry a nontrivial local system, and all other curves are closed in the lift to the plane \tilde{T}_\bullet and thus must contain a type a_k segment. The portion of the train track containing the extra edges is precisely a type a_k segment with local system (V, Φ) . These extra edges determine an $n \times n$ matrix over \mathbb{F} , where the (i, j) entry is nonzero if the curve-like train track contains a copy of the a_k segment




Figure 12: A 3-dimensional local system, expanded at an a_2 to give a train track.

from the i^{th} copy of the initial generator in the segment to the j^{th} copy of the final generator of the segment. By construction, this matrix represents the local system Φ .

It is relatively straightforward to extend this language to cases admitting a nontrivial local system. Recall that each letter in \mathfrak{A} corresponds to a (portion of a) type D structure that is a linear chain of arrows. We allow a letter in \mathfrak{A} to be decorated by a local system, as follows. Let V be a vector space over \mathbb{F} of dimension n , and let $\Phi: V \rightarrow V$ be an endomorphism. Decorating a segment with (V, Φ) amounts to taking n parallel copies of the appropriate chain, with the n parallel copies of any one arrow in the chain replaced with a collection of arrows determined by Φ ; see Figure 12 for an example. Suppose the relevant arrow in the chain connects generators (ie vertices) x and y , with x occurring first in the chain (ie there is an edge connecting x to y). Fix bases $\langle x_1, \dots, x_n \rangle$ and $\langle y_1, \dots, y_n \rangle$ for V and consider an $n \times n$ matrix over \mathbb{F} representing Φ ; these arrows connect the i^{th} copy of x (ie x_i) to the j^{th} copy of y (ie y_j) if and only if the (i, j) entry of the matrix is nonzero (again, see Figure 12). If all letters in a cyclic word carry a local system (each having the same dimension), then the local system on the cyclic word is determined by composing the endomorphisms. Note that a letter decorated by the trivial local system of dimension n corresponds to n parallel copies of the relevant curve segment.

To summarize, given a knot K in S^3 , the invariant $\widehat{HF}(M)$ is an immersed multicurve $\boldsymbol{\gamma}(K) = (\gamma_0, \gamma_1, \dots, \gamma_n)$ where each $\gamma_{i>0}$ carries a (possibly trivial) local system. By the above discussion, we can assume that each component of $\boldsymbol{\gamma}(K)$ is represented by a cyclic word in \mathfrak{A} , possibly with a nontrivial local system on a single a_k segment.

1.2 The merge operation

Given type D structures $\boldsymbol{\vartheta}$ and $\boldsymbol{\gamma}$, we describe a new type D structure $M(\boldsymbol{\gamma}, \boldsymbol{\vartheta})$. This follows the notation set out in [6], where we showed that this type D structure agrees

with $\widehat{HF}(\mathcal{M}(M_1, M_2))$ in the case where $\boldsymbol{\gamma}$ and $\boldsymbol{\vartheta}$ correspond to $\widehat{HF}(M_1)$ and $\widehat{HF}(M_2)$, respectively.³ The operations M and \mathcal{M} are referred to as merges; the latter glues two manifolds along essential annuli in their torus boundaries. We will first describe the operator M algebraically, and then explain the gluing conventions for \mathcal{M} in the next section in the context of cabling.

Some simplifications are possible in the present setting. First, we assume that $\boldsymbol{\gamma}$ is a loop consisting only of some c_k for integers k . This assumption holds in particular when M_1 is a solid torus, in which case $\mathcal{M}(M_1, -)$ will give rise to a cabled knot.⁴ Further, as described above, we assume that $\boldsymbol{\vartheta}$ is represented as a curve-like train track. This may consist of several disjoint components, but we can restrict to connected train tracks without loss of generality: if $\boldsymbol{\vartheta} = (\vartheta_0, \dots, \vartheta_n)$ then $M(\boldsymbol{\gamma}, \boldsymbol{\vartheta}) = (M(\boldsymbol{\gamma}, \vartheta_0), \dots, M(\boldsymbol{\gamma}, \vartheta_n))$.

The main tool used in this paper is a distilled version of [6, Proposition 6.4]:

Proposition 11 *Let $\boldsymbol{\vartheta}$ be a type D structure represented by a single cyclic word in \mathfrak{A} and let $\boldsymbol{\gamma}$ be a word containing only the c_k . If the local system on $\boldsymbol{\vartheta}$ is trivial then the type D structure $M(\boldsymbol{\gamma}, \boldsymbol{\vartheta})$ is obtained by applying the rules*

$$M(c_k, a_j) = a_j, \quad M(c_k, b_j) = b_j, \quad M(c_k, c_j) = c_{j+k}$$

to every letter in $\boldsymbol{\vartheta}$, ranging over all letters for $\boldsymbol{\gamma}$, and assembling the result together using a toroidal grid to match up the endpoints, as described in Figure 13.

The proof of this result is contained in [6]; however, because nontrivial local systems are not handled there, we want to be precise about how to extend the result based on the material in our earlier work.

Proposition 12 *Proposition 11 holds when $\boldsymbol{\vartheta}$ carries a local system, where, for each letter u in the word representing $\boldsymbol{\vartheta}$ and each c_k in the word representing $\boldsymbol{\gamma}$, $M(c_k, u)$ carries the same local system as u .*

³More specifically, treating $\boldsymbol{\gamma}$ and $\boldsymbol{\vartheta}$ as type D structures, in [6] we use $M(\boldsymbol{\gamma}, \boldsymbol{\vartheta})$ as a shorthand for the type D structure $\widehat{CFDAA}(\mathcal{P} \times S^1) \boxtimes (\boldsymbol{\gamma}, \boldsymbol{\vartheta})$, where $\widehat{CFDAA}(\mathcal{P} \times S^1)$ is the bordered trimodule calculated in [3] and the three-manifold $\mathcal{P} \times S^1$ is a circle bundle over a pair of pants (this plays a key role in the next section).

⁴In fact, everything we do works in a much more general setting: Any manifold admitting L -space surgeries has a type D structure that, relative to a slope corresponding to an L -space filling, can be expressed in terms of only letters c_k . We have opted to simplify matters and focus on a well-known construction with well-established conventions in order to illustrate the key principle. More general cases follow the same lines, and can be extracted from [4; 6].

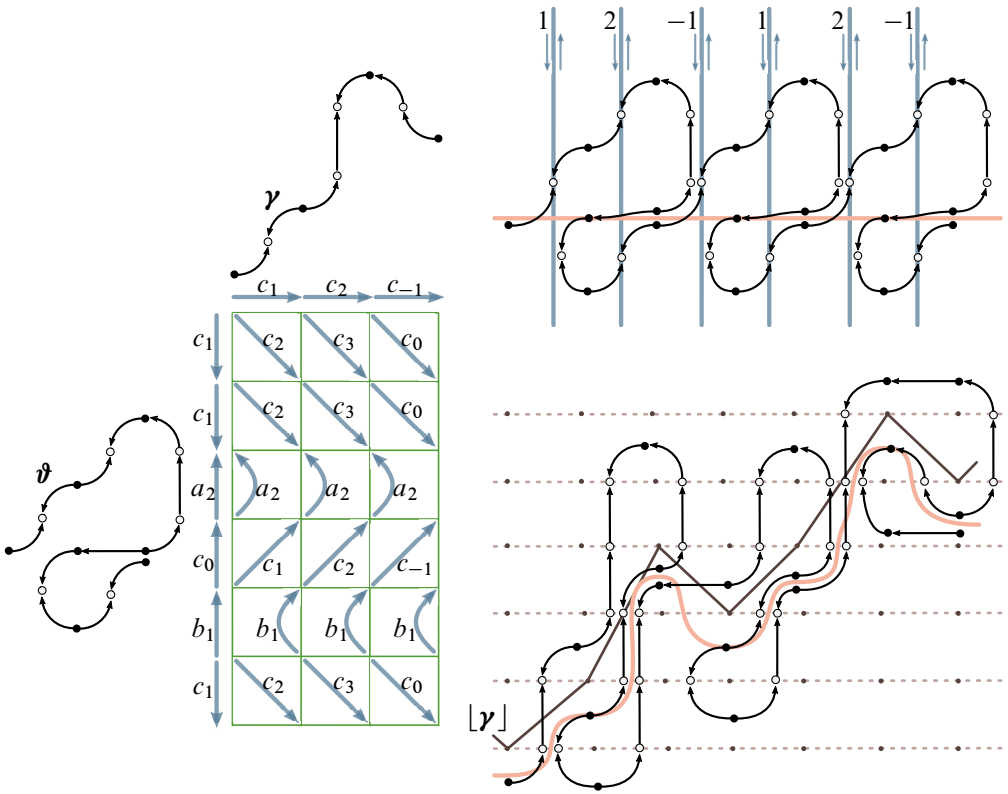


Figure 13: Merging a pair of curves, as in Proposition 11, described graphically: On the left-hand side of the diagram, the output curve is interpreted on a toroidal grid, where the c_k from γ (written on the horizontal) act on the letters in ϑ (written on the vertical). On the right, this process is interpreted in terms of curves, where the top right figure gives a section of the (periodic) curve in \tilde{T}_\bullet while the bottom right figure is the result of the merge. Note that the horizontal is moved to the key curve $[\gamma]$.




Figure 14: The PL key curves approximating γ .

Proof There is nothing to check for trivial local systems, as these are just disjoint copies of some curve and Proposition 11 applies. For a nontrivial local system, we need to carry out the computation in [6, Figures 10, 11 and 12], replacing the simple segments for the ℓ_2 input with segments carrying an arbitrary local system as in, for example, Figure 12. This is a straightforward computation. Note that we do not need to check this computation for type b_k pieces, since we may assume a local system on a loop is concentrated on any one letter and a loop containing a b_k must also contain an a_j . (Furthermore if ϑ is the multicurve corresponding to a knot in S^3 , as in this paper, it is enough to check the computation for a_k pieces since any component with a nontrivial local system must contain an a_k piece.) \square

The takeaway from Proposition 11 (and its extension to nontrivial local systems in Proposition 12) is a graphical calculus used to determine the merge of two curves when one contains only c_k segments; this is the content of Figure 13. Consider a word (in the c_k) representing γ , and write this along the top of a rectangular grid; consider a word in \mathfrak{A} representing ϑ , and write this along the side of the rectangular grid. Then, following the letter-by-letter instructions in Proposition 11, the new word $M(\gamma, \vartheta)$ can be obtained by running through the grid, starting at the top left. As shown in Figure 13, the $M(c_j, c_k) = c_{j+k}$ run diagonally, while the $M(c_j, a_k) = a_k$ and $M(c_j, b_k) = b_k$ change direction. The sides of this grid are identified to form a toroidal grid, and this connects up the endpoints of the segments to form the new loop $M(\gamma, \vartheta)$.

Note that, as the resulting loop is traversed, horizontal motion in the grid corresponds exactly to horizontal motion of the corresponding curve in the plane. In particular, each vertical line of lattice points in the plane corresponds to a column of the grid containing some c_k in γ , and the effect of merging on the curve for ϑ is to shear the plane along that vertical line by k .

The graphical shorthand from Figure 13 suggests an interpretation of the merge operation in terms of immersed curves, which we can think of as γ acting on ϑ . To describe this, it is useful to have a piecewise-linear representative of the curve γ . Let γ be expressible as a word in only the c_k , so that, viewed in the plane \tilde{T}_\bullet , γ is a graph. Let $[\gamma]$ be the curve consisting of linear segments that, at each integer in the horizontal direction, intersects the lattice point immediately below γ . The curve $[\gamma]$ is defined similarly, by instead pushing up to the lattice points immediately above γ ; see Figure 14. Note that this is closely related to the *pegboard diagrams* introduced in [4].

Recall that the immersed multicurve $\widehat{HF}(M)$ coming from a bordered 3-manifold M lives in the cylinder $(\mathbb{R}^2 \setminus \mathbb{Z}^2)/\langle \lambda \rangle$, where λ corresponds to the homological longitude of M . Equivalently, we think of this as a multicurve in $\mathbb{R}^2 \setminus \mathbb{Z}^2$ which is invariant under the action of λ . We will say that such a curve has *horizontal period* p if translation by λ moves p units in the horizontal direction.

Corollary 13 *Let ϑ be an immersed multicurve with local systems in \widetilde{T}_\bullet with horizontal period q , and let γ be a curve in \widetilde{T}_\bullet with no vertical tangent lines (ie the graph of a function) with horizontal period p , with p and q relatively prime. Then the immersed multicurve for $M(\gamma, \vartheta)$ is obtained by adding $\lfloor \gamma \rfloor$ to ϑ vertically. That is, we find the image ϑ under the transformation of \widetilde{T}_\bullet which translates along each vertical line to take the horizontal axis to $\lfloor \gamma \rfloor$.*

Proof This is the main thrust of Figure 13: For a component homologous to λ , the new cyclic word moves q columns to the right in the grid each time it traverses the grid vertically. Since p and q are relatively prime, the new word makes p vertical passes, tracing out the entire grid, before returning to the starting point. The new word is p copies of the word representing ϑ , with the indices on type c letters shifted according to the column in the grid; this corresponds to p copies of the fundamental region in $\widetilde{\vartheta}$, each of which moves q units to the right, with a plane shear applied along each column of lattice points. The magnitude of each shear is determined by the index of the corresponding letter in γ , which amounts to shifting each column upwards by the height of $\lfloor \widetilde{\gamma} \rfloor$ in that column. The resulting curve has horizontal period pq . For a nullhomologous component, the grid gives rise to p separate cyclic words, each traversing the grid vertically once starting in a different column. Each word is a copy of ϑ with shifted indices on type c letters. The nullhomologous component of ϑ lifts to infinitely many copies of the same closed curve in $\widetilde{\vartheta}$, which are translations of each other by multiples of λ . Taking p consecutive copies corresponds to the p cyclic words in the grid, and adding $\lfloor \widetilde{\gamma} \rfloor$ corresponds to the required shifts in indices. □

As mentioned previously, the setup of Corollary 13 is more general than we need for cabling; we will only need the case that γ is in fact a straight line of some rational slope. Note that, when γ is a line of slope 1, the transformation taking the horizontal axis to $\lfloor \gamma \rfloor$ is a lift to \widetilde{T}_\bullet of a Dehn twist in T_\bullet . This is a linear transformation of the plane, which we refer to as a *plane shear* in the vertical direction. The case that γ is a line of rational slope is a mild generalization of this, which we call a *fractional plane shear*.

2 The proof of Theorem 1

In order to complete the proof of our theorem, we need to connect the operation described in Corollary 13 to the specific context of cabling. To do this we first set our conventions.

2.1 Cabling conventions

Recall that, fixing a knot K , we let M denote the complement $S^3 \setminus \nu(K)$ and $M_{p,q}$ denote the complement of the cable $C_{p,q}(K)$. Let \mathcal{P} denote a two-sphere with three disks removed (so that \mathcal{P} is homeomorphic to a pair of pants). The manifold $M_{p,q}$ can be obtained by gluing M into one boundary component of $\mathcal{P} \times S^1$ and an appropriately framed solid torus $D^2 \times S^1$ into another boundary component. We will briefly review this construction, paying particular attention to framing conventions.

Each torus boundary in this construction has a natural choice of parametrizing curves. For ∂M , we use a meridian μ and the Seifert longitude λ , fixing orientations on these curves with the convention that $\lambda \cdot \mu = +1$. For $\partial(D^2 \times S^1)$, we let m be a meridian $\partial D^2 \times \{\text{pt}\}$ and let ℓ be the longitude $\{\text{pt}\} \times S^1$, with the orientation convention that $m \cdot \ell = +1$. For $i \in \{1, 2, 3\}$, the i^{th} boundary component of the S^1 -bundle $\mathcal{P} \times S^1$ is parametrized by a fibre $f_i = \{\text{pt}\} \times S^1$ and $b_i = \partial_i(\mathcal{P}) \times \{\text{pt}\}$, where $\partial_i(\mathcal{P})$ denotes the i^{th} boundary component of the base surface \mathcal{P} . We set orientations on these curves so that $b_i \cdot f_i = +1$.

The third boundary component of $\mathcal{P} \times S^1$ will ultimately become the boundary of $M_{p,q}$; however, it is helpful for the moment to fill this third boundary component in with a solid torus in a trivial way so that $\mathcal{P} \times S^1$ becomes $A \times S^1$, where A is an annulus. This solid torus can be removed later by deleting a neighbourhood of a fibre of $A \times S^1$. We glue $D^2 \times S^1$ to the first boundary component of $\mathcal{P} \times S^1$ (now $A \times S^1$) such that f_1 is identified with $p\ell + qm$; this means that b_1 is identified with $r\ell + sm$ for some integers r and s with $ps - qr = -1$ (we can choose r and s arbitrarily subject to this condition, but the choice affects the framings on the resulting boundary components). The result of this gluing is a solid torus, equipped with a Seifert fibration in which the core of the solid torus is a singular fibre and the regular fibres wind p times longitudinally and q times meridionally. This solid torus is glued to the knot complement M such that the result is S^3 and the core of the solid torus is identified with K . As a result, a regular fibre of $D^2 \times S^1 \cup A \times S^1$ is the cable $C_{p,q}(K)$, and removing a neighbourhood of

one of these (or, equivalently, not filling in the third boundary of $\mathcal{P} \times S^1$) yields the complement $M_{p,q}$.

Note that inserting $A \times S^1 \cong T^2 \times [0, 1]$ between M and $D^2 \times S^1$ amounts to a change of framing and, in particular, f_1 can be identified with f_2 and b_1 can be identified with $-b_2$. To recover S^3 , we want μ to be identified with m and λ to be identified with ℓ . It follows that f_2 glues to $p\lambda + q\mu$ and b_2 glues to $-r\lambda - s\mu$. To summarize, we have

$$\begin{bmatrix} f_1 \\ b_1 \end{bmatrix} = \begin{bmatrix} p & q \\ r & s \end{bmatrix} \begin{bmatrix} \ell \\ m \end{bmatrix}, \quad \begin{bmatrix} f_2 \\ b_2 \end{bmatrix} = \begin{bmatrix} p & q \\ -r & -s \end{bmatrix} \begin{bmatrix} \lambda \\ \mu \end{bmatrix}$$

with $ps - qr = -1$. Inverting these matrices,

$$\begin{bmatrix} \ell \\ m \end{bmatrix} = \begin{bmatrix} -s & q \\ r & -p \end{bmatrix} \begin{bmatrix} f_1 \\ b_1 \end{bmatrix}, \quad \begin{bmatrix} \lambda \\ \mu \end{bmatrix} = \begin{bmatrix} -s & -q \\ r & p \end{bmatrix} \begin{bmatrix} f_2 \\ b_2 \end{bmatrix}.$$

If we do not fill in the third boundary of $\mathcal{P} \times S^1$ in the construction above, the resulting manifold with torus boundary is $M_{p,q}$; it is clear that the meridian μ_C of the cable knot is given by b_3 . While not required, as will follow from the computation below, one can check that the Seifert longitude λ_C of the cable knot is given by $-f_3 + q^2b_3$.

2.2 Applying the merge operation

We are interested in obtaining the immersed curve set $\widehat{HF}(M_{p,q})$ from the immersed curve set $\widehat{HF}(M)$. We can do this by applying the merge operation to $\widehat{HF}(M)$ and $\widehat{HF}(D^2 \times S^1)$, keeping in mind the framings discussed above. Following Section 1 and the conventions in [6], the first step is to draw lifts of both curve sets in the plane with respect to the parametrization by f_i and b_i (or, more precisely, by the curves in $\partial(D^2 \times S^1)$ or ∂M which are identified with f_i and b_i) such that b_i is the horizontal direction and f_i is the vertical direction. Recall our convention is that $b_i \cdot f_i = 1$, but we are now considering the plane as a lift of $\partial(D^2 \times S^1)$ or of ∂M , which are identified with boundary tori of $\mathcal{P} \times S^1$ by an orientation-reversing diffeomorphism, so if we take b_i to be the positive horizontal axis then f_i is the negative vertical axis. Note that we could instead choose the opposite orientation for both axes, but this ambiguity can be ignored since immersed curves for bordered invariants are symmetric under the elliptic involution of the torus by [5, Theorem 7]. Since $\widehat{HF}(D^2 \times S^1)$ is the meridian $m = rf_1 - pb_1$, this curve is simple to describe in the relevant basis: it is a line of slope $\frac{r}{p}$ [4]. The second step is to apply Corollary 13 by taking the vertical sum of $\widehat{HF}(M)$ and $[m]$; note that the conditions of the corollary are satisfied because

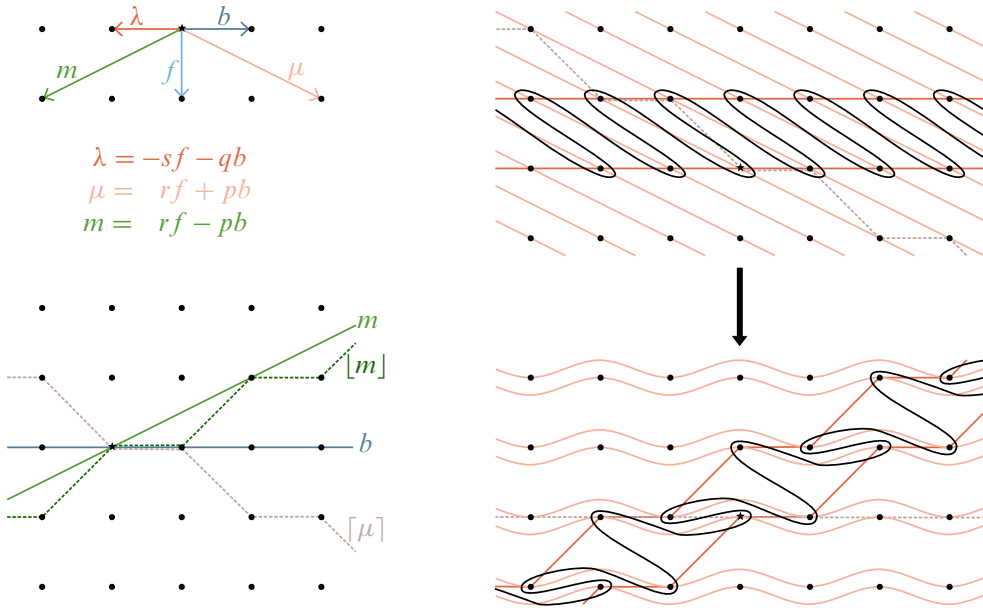


Figure 15: The fractional plane shear in the vertical direction associated with computing a (p, q) -cable, viewed with respect to the $(-f, b)$ -framing. Other relevant curves are shown, with respect to this framing, in the top left. The bottom left shows a copy of m through the origin and the corresponding curve $[m]$ obtained by dropping down to the highest peg below m in each column. This curve serves as a “key” for the plane shear — that is, we shift each column of pegs upward by the height of $[m]$ in that column. Thus the plane shear is determined by the fact that it takes b to $[m]$, or equivalently that it takes $[\mu]$ to b . The right shows the effect of this shear on the curve for the right-hand trefoil. For the concrete example in the figure, $(p, q) = (2, 1)$ and $(r, s) = (1, 0)$.

components $\widehat{HF}(M)$ are homologous to zero or to the rational longitude λ , which moves horizontally by q units, while m moves horizontally by p units. The result is (a lift to the plane of) $\widehat{HF}(M_{p,q})$, though given with respect to the framing (f_3, b_3) rather than the usual (μ_C, λ_C) ; see Figure 15.

While the previous paragraph gives a complete procedure for computing $\widehat{HF}(M_{p,q})$, performing the change of basis to draw the curve set $\widehat{HF}(M)$ with respect to the (f, b) -framing can be cumbersome. Instead, we can follow the same operation but view the plane with respect to (μ, λ) , the preferred framing for ∂M , throughout the process. Now, instead of shifting pegs in each vertical column, we shift along lines parallel to the fibre direction; since $f_2 = p\lambda + q\mu$, this is a line of slope $\frac{q}{p}$. To keep

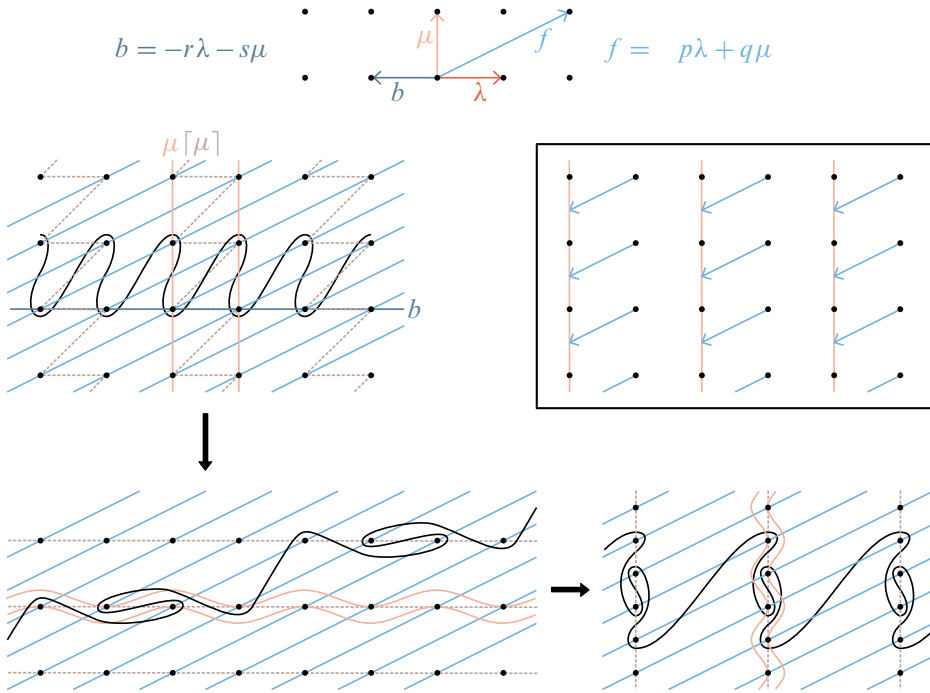


Figure 16: Starting with the curve $\widehat{HF}(M)$ drawn in the plane with respect to the standard (μ, λ) -framing, the fractional plane shear in the f direction which takes $[\mu]$ to b produces the curve $\widehat{HF}(M_{p,q})$, though not in terms of a convenient parametrization. shearing back partially along f gives the curve with the standard parametrization, up to rescaling the lattice. These two steps can be combined into one, as shown in the box: each lattice point is translated leftward along lines of slope $\frac{q}{p}$ until its x -coordinate is a multiple of p . Note that p copies of the curve $\widehat{HF}(M)$ are involved in each copy of $\widehat{HF}(M_{p,q})$. The figure shows the case of the $(2, 1)$ -cable of the right-hand trefoil.

track of how much to shift along each line of slope $\frac{q}{p}$, we can draw a copy of the piecewise linear curve $[\mu]$; note that this is obtained from a vertical line μ through the origin by pushing each point $(0, \frac{n}{p})$ on μ leftward along a line of slope $\frac{q}{p}$ to the first lattice point it encounters (see Figure 16, top left). To perform the cable operation, we shear along lines of slope $\frac{q}{p}$ to bring this curve $[\mu]$ to b (see Figure 16, bottom left).

We can now start with the curve $\widehat{HF}(M)$ represented in terms of its standard framing (μ, λ) and produce the immersed curve $\widehat{HF}(M_{p,q})$ in one simple step. However, as before, the output is not given with respect to the standard framing by (μ_C, λ_C) . Of course, it is straightforward to determine the slopes of μ_C and λ_C in the output picture

and then we simply need to change basis applying a linear map to the plane which takes these to the vertical and horizontal directions, respectively. This can always be accomplished by a sequence of (integral) plane shears in the horizontal and vertical directions. However, this too is cumbersome, so we will describe a shortcut to this reparametrization making use of a linear transformation of the plane which does not preserve the lattice. More precisely, consider the linear transformation which fixes f and takes b to μ ; this can be understood as translating each lattice point on b along a line of slope $\frac{q}{p}$ until it reaches the vertical line μ (see Figure 16, bottom right). Note that the lattice \mathbb{Z}^2 is not mapped to itself under this transformation, but rather its image is $p\mathbb{Z} \times \frac{1}{p}\mathbb{Z}$. Even so, in this new deformed lattice the directions corresponding to μ_C and λ_C are vertical and horizontal, as desired, and we can recover the usual lattice by ending with another linear transformation which scales and compresses by a factor of p in the vertical direction and horizontal direction, respectively.

Finally, we mention that there are now two steps which involve shearing along the lines of slope $\frac{q}{p}$: the fractional plane shear taking $[\mu]$ to b (this transformation is not linear), and the linear transformation taking b to μ . These steps can be combined in one by shearing along lines of slope $\frac{q}{p}$ to push $[\mu]$ onto μ . In other words, every p^{th} vertical column of lattice points is fixed, while all other points are pushed leftward along lines of slope $\frac{q}{p}$ until they reach a vertical line containing one of the fixed columns (see the boxed portion of Figure 16). This proves Theorem 1.

Acknowledgements Hanselman was partially supported by NSF grant DMS-1812527; Watson was partially supported by an NSERC discovery/accelerator grant.

References

- [1] **I Dai, J Hom, M Stoffregen, L Truong**, *More concordance homomorphisms from knot Floer homology*, *Geom. Topol.* 25 (2021) 275–338 MR Zbl
- [2] **M H Freedman**, *The topology of four-dimensional manifolds*, *J. Differential Geometry* 17 (1982) 357–453 MR Zbl
- [3] **J Hanselman**, *Bordered Heegaard Floer homology and graph manifolds*, *Algebr. Geom. Topol.* 16 (2016) 3103–3166 MR Zbl
- [4] **J Hanselman, J Rasmussen, L Watson**, *Bordered Floer homology for manifolds with torus boundary via immersed curves*, preprint (2016) arXiv 1604.03466
- [5] **J Hanselman, J Rasmussen, L Watson**, *Heegaard Floer homology for manifolds with torus boundary: properties and examples*, *Proc. Lond. Math. Soc.* (3) 125 (2022) 879–967 MR

- [6] **J Hanselman, L Watson**, *A calculus for bordered Floer homology*, *Geom. Topol.* 27 (2023) 823–924
- [7] **M Hedden**, *On knot Floer homology and cabling*, PhD thesis, Columbia University (2005) MR Available at <https://www.proquest.com/docview/305015665>
- [8] **M Hedden**, *On knot Floer homology and cabling*, *Algebr. Geom. Topol.* 5 (2005) 1197–1222 MR Zbl
- [9] **M Hedden**, *On knot Floer homology and cabling, II*, *Int. Math. Res. Not.* 2009 (2009) 2248–2274 MR Zbl
- [10] **J Hom**, *A note on cabling and L -space surgeries*, *Algebr. Geom. Topol.* 11 (2011) 219–223 MR Zbl
- [11] **J Hom**, *Bordered Heegaard Floer homology and the tau-invariant of cable knots*, *J. Topol.* 7 (2014) 287–326 MR Zbl
- [12] **J Hom**, *A survey on Heegaard Floer homology and concordance*, *J. Knot Theory Ramifications* 26 (2017) art. id. 1740015 MR Zbl
- [13] **A S Levine**, *Knot doubling operators and bordered Heegaard Floer homology*, *J. Topol.* 5 (2012) 651–712 MR Zbl
- [14] **R Lipshitz, P S Ozsváth, D P Thurston**, *Bimodules in bordered Heegaard Floer homology*, *Geom. Topol.* 19 (2015) 525–724 MR Zbl
- [15] **R Lipshitz, P S Ozsvath, D P Thurston**, *Bordered Heegaard Floer homology*, *Mem. Amer. Math. Soc.* 1216, Amer. Math. Soc., Providence, RI (2018) MR Zbl
- [16] **L Mosher**, *What is a train track?*, *Notices Amer. Math. Soc.* 50 (2003) 354–355
- [17] **P S Ozsváth, A I Stipsicz, Z Szabó**, *Concordance homomorphisms from knot Floer homology*, *Adv. Math.* 315 (2017) 366–426 MR Zbl
- [18] **P Ozsváth, Z Szabó**, *Holomorphic disks and knot invariants*, *Adv. Math.* 186 (2004) 58–116 MR Zbl
- [19] **I Petkova**, *Cables of thin knots and bordered Heegaard Floer homology*, *Quantum Topol.* 4 (2013) 377–409 MR Zbl
- [20] **J A Rasmussen**, *Floer homology and knot complements*, PhD thesis, Harvard University (2003) MR arXiv math/0306378
- [21] **C A Van Cott**, *Ozsváth–Szabó and Rasmussen invariants of cable knots*, *Algebr. Geom. Topol.* 10 (2010) 825–836 MR Zbl

*Department of Mathematics, Princeton University
Princeton, NJ, United States*

*Department of Mathematics, University of British Columbia
Vancouver BC, Canada*

`jh66@math.princeton.edu, liam@math.ubc.ca`

Proposed: Ciprian Manolescu

Received: 12 August 2019

Seconded: Peter Ozsváth, András I Stipsicz

Revised: 22 November 2019

GEOMETRY & TOPOLOGY

msp.org/gt

MANAGING EDITOR

András I. Stipsicz Alfréd Rényi Institute of Mathematics
stipsicz@renyi.hu

BOARD OF EDITORS

Dan Abramovich	Brown University dan_abramovich@brown.edu	Mark Gross	University of Cambridge mgross@dpms.cam.ac.uk
Ian Agol	University of California, Berkeley ianagol@math.berkeley.edu	Rob Kirby	University of California, Berkeley kirby@math.berkeley.edu
Mark Behrens	Massachusetts Institute of Technology mbehrens@math.mit.edu	Frances Kirwan	University of Oxford frances.kirwan@balliol.oxford.ac.uk
Mladen Bestvina	Imperial College, London bestvina@math.utah.edu	Bruce Kleiner	NYU, Courant Institute bkleiner@cims.nyu.edu
Martin R. Bridson	Imperial College, London m.bridson@ic.ac.uk	Urs Lang	ETH Zürich urs.lang@math.ethz.ch
Jim Bryan	University of British Columbia jbryan@math.ubc.ca	Marc Levine	Universität Duisburg-Essen marc.levine@uni-due.de
Dmitri Burago	Pennsylvania State University burago@math.psu.edu	John Lott	University of California, Berkeley lott@math.berkeley.edu
Ralph Cohen	Stanford University ralph@math.stanford.edu	Ciprian Manolescu	University of California, Los Angeles cm@math.ucla.edu
Tobias H. Colding	Massachusetts Institute of Technology colding@math.mit.edu	Haynes Miller	Massachusetts Institute of Technology hrm@math.mit.edu
Simon Donaldson	Imperial College, London s.donaldson@ic.ac.uk	Tom Mrowka	Massachusetts Institute of Technology mrowka@math.mit.edu
Yasha Eliashberg	Stanford University eliash-gt@math.stanford.edu	Walter Neumann	Columbia University neumann@math.columbia.edu
Benson Farb	University of Chicago farb@math.uchicago.edu	Jean-Pierre Otal	Université d'Orleans jean-pierre.otal@univ-orleans.fr
Steve Ferry	Rutgers University sferry@math.rutgers.edu	Peter Ozsváth	Columbia University oszvath@math.columbia.edu
Ron Fintushel	Michigan State University ronfint@math.msu.edu	Leonid Polterovich	Tel Aviv University polterov@post.tau.ac.il
David M. Fisher	Rice University davidfisher@rice.edu	Colin Rourke	University of Warwick gt@maths.warwick.ac.uk
Mike Freedman	Microsoft Research michaelf@microsoft.com	Stefan Schwede	Universität Bonn schwede@math.uni-bonn.de
David Gabai	Princeton University gabai@princeton.edu	Peter Teichner	University of California, Berkeley teichner@math.berkeley.edu
Stavros Garoufalidis	Southern U. of Sci. and Tech., China stavros@mpim-bonn.mpg.de	Richard P. Thomas	Imperial College, London richard.thomas@imperial.ac.uk
Cameron Gordon	University of Texas gordon@math.utexas.edu	Gang Tian	Massachusetts Institute of Technology tian@math.mit.edu
Lothar Götsche	Abdus Salam Int. Centre for Th. Physics gotsche@ictp.trieste.it	Ulrike Tillmann	Oxford University tillmann@maths.ox.ac.uk
Jesper Grodal	University of Copenhagen jg@math.ku.dk	Nathalie Wahl	University of Copenhagen wahl@math.ku.dk
Misha Gromov	IHÉS and NYU, Courant Institute gromov@ihes.fr	Anna Wienhard	Universität Heidelberg wienhard@mathi.uni-heidelberg.de

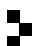
See inside back cover or msp.org/gt for submission instructions.

The subscription price for 2023 is US \$740/year for the electronic version, and \$1030/year (+ \$70, if shipping outside the US) for print and electronic. Subscriptions, requests for back issues and changes of subscriber address should be sent to MSP. Geometry & Topology is indexed by Mathematical Reviews, Zentralblatt MATH, Current Mathematical Publications and the Science Citation Index.

Geometry & Topology (ISSN 1465-3060 printed, 1364-0380 electronic) is published 9 times per year and continuously online, by Mathematical Sciences Publishers, c/o Department of Mathematics, University of California, 798 Evans Hall #3840, Berkeley, CA 94720-3840. Periodical rate postage paid at Oakland, CA 94615-9651, and additional mailing offices. POSTMASTER: send address changes to Mathematical Sciences Publishers, c/o Department of Mathematics, University of California, 798 Evans Hall #3840, Berkeley, CA 94720-3840.

GT peer review and production are managed by EditFLOW[®] from MSP.

PUBLISHED BY

 **mathematical sciences publishers**
nonprofit scientific publishing

<http://msp.org/>

© 2023 Mathematical Sciences Publishers

GEOMETRY & TOPOLOGY

Volume 27 Issue 3 (pages 823–1272) 2023

- A calculus for bordered Floer homology 823
JONATHAN HANSELMAN and LIAM WATSON
- Cabling in terms of immersed curves 925
JONATHAN HANSELMAN and LIAM WATSON
- Combinatorial Reeb dynamics on punctured contact
3-manifolds 953
RUSSELL AVDEK
- Unexpected Stein fillings, rational surface singularities and
plane curve arrangements 1083
OLGA PLAMENEVSKAYA and LAURA STARKSTON
- A smooth compactification of the space of genus two curves in
projective space: via logarithmic geometry and Gorenstein
curves 1203
LUCA BATTISTELLA and FRANCESCA CAROCCI