Braid-positive Legendrian links

Tamás Kálmán

March 30, 2007

Abstract

Any link that is the closure of a positive braid has a natural Legendrian representative. These were introduced in an earlier paper, where their Chekanov–Eliashberg contact homology was also evaluated. In this paper we re-phrase and improve that computation using a matrix representation. In particular, we present a way of finding all augmentations of such Legendrians, construct an augmentation which is also a ruling, and find surprising links to LU–decompositions and Gröbner bases.

1 Introduction

I came across a certain set of Legendrian links while searching for examples to illustrate the main theorem of my thesis [16], and they served that purpose very well. Since then I kept returning to them because I could always discover something pretty. This paper is a collection of those findings.

The links in question (see Figure 1), that I call *Legendrian closures of positive braids*, denote by L_β, and represent by front diagrams f_β, are Legendrian representatives of braid-positive links, i.e. link types that can be obtained as the closure of a positive braid β. (These are not to be confused with the more general notion of positive link, i.e. link types that can be represented with diagrams whose geometric and algebraic crossing numbers agree.)

We will concentrate on Legendrian isotopy invariants of L_β. Some of these have been evaluated in [16], of which the present paper is a continuation. It is thus assumed that the reader is familiar with sections 2 (basic notions) and 6 (Legendrian closures of positive braids and their relative contact homology) of [16]. In this paper, we will re-formulate some of those computations, and get new results also, by using what we call the path matrix of a positive braid. This construction is very similar to that of Lindström [18] and Gessel–Viennot [15], which is also included in the volume [1].

The paper is organized as follows. We review some results of [16] in section 2, and discuss elementary properties of the path matrix in section 3. Then, the main results are Theorem 4.4, where we compute a new generating set for the (abelianized) image I of the contact homology differential and its consequence, Theorem 5.3, which gives a quick test to decide whether a given set of crossings is an augmentation. The latter is in terms of an LU–decomposition, i.e. Gaussian elimination of the path matrix. In Theorem 4.5, we point out that the old generators, i.e. the ones read off of the knot diagram, automatically form a Gröbner basis for I. In Theorem 6.3, we construct a subset of the crossings of β which is simultaneously an augmentation and a ruling of L_β. This strengthens the well known relationship between augmentations and rulings. We close the paper with a few examples and speculations.

Acknowledgements: Part of the research was carried out in the summer of 2005 when I visited the Alfréd Rényi Mathematical Institute in Budapest, supported by the European Union Marie Curie TOK program BudAlgGeo. It is a great pleasure to thank the Institute for their hospitality and András Némethi, Endre Szabó, and András Szűcs for stimulating discussions. I am grateful to Alexander Stoimenow for providing me with a list of braid-positive knots up to 16 crossings. Francis Bonahon helped me sort out the mutants that appear in section 7. My conversations with Mikhail Kotchetov were invaluable to the discovery of Theorem 4.5. Last but not least, many thanks to Supap Kirtsaeng for writing the computer implementation of Theorem 5.3.

2 Preliminaries

The goal of this section is to recall some results from [16] relevant to this paper. We will work in the standard contact 3–space $\mathbf{R}^3_{xyz}$ with the kernel field of the 1–form $\mathrm{d}z - y\mathrm{d}x$. We will use the basic notions of Legendrian knot, Legendrian isotopy, front (xz) diagram, Maslov potential, Lagrangian (xy) diagram, resolution [19], Thurston–Bennequin (tb) and rotation (r) numbers, admissible disc and contact homology[1] etc. without reviewing their definitions. We will also assume that the reader is familiar with section 6 of [16], of which this paper is in a sense an extension. For a complete introduction to Legendrian knots and their contact homology, see [8].

We would like to stress a few points only whose treatment may be somewhat non-standard. Crossings a of both front and Lagrangian diagrams are assigned an *index*, denoted by $|a|$, which is an element of $\mathbf{Z}_{2r}$, with the entire assignment known as a *grading*. This is easiest to define for fronts of single-component knots as the difference of the Maslov potentials (upper − lower) of the two intersecting strands. If a Lagrangian diagram is the result of resolution, the old crossings keep their indices and the crossings replacing the right cusps are assigned the index 1. In the multi-component case, the Maslov potential difference becomes ambiguous for crossings between different components. This gives rise to an infinite set of so-called admissible gradings. We consider these as introduced in [19, section 2.5] and not the larger class of gradings described in [3, section 9.1].

Let β denote an arbitrary positive braid word. The Legendrian isotopy class L_β is a natural Legendrian representative of the link which is the closure of β. (All braids and braid words in this paper are positive. The same symbol β may

[1] Because absolute contact homology doesn't appear in the paper, we'll use this shorter term for what may be better known as relative, Legendrian, or Chekanov–Eliashberg contact homology.

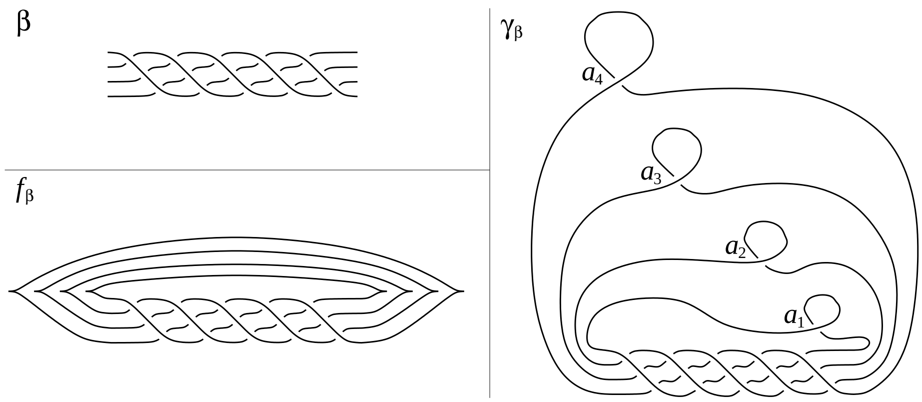

Figure 1: Front (f_β) and Lagrangian (γ_β) diagrams of the closure (L_β) of the positive braid β

sometimes refer to the braid represented by the braid word β.) L_β, in turn, is represented by the front diagram f_β and its resolution, the Lagrangian diagram γ_β (see Figure 1). Considering β drawn horizontally, label the left and right endpoints of the strands from top to bottom with the first q whole numbers (q is the number of strands in β). The crossings of β, labeled from left to right by the symbols $b_1, \ldots, b_w$, are the only crossings of f_β. Due to resolution, γ_β also has the crossings $a_1, \ldots, a_q$.

The differential graded algebra (DGA) $\mathcal{A}$ which is the chain complex for the contact homology of L_β is generated (as a non-commutative algebra with unit) freely over $\mathbf{Z}_2$ by these $q+w$ symbols (w is the word length or exponent sum of β). It's assigned a $\mathbf{Z}$–grading[2] which takes the value 0 on the b_k and the value 1 on the a_n (extended by the rule $|uv| = |u| + |v|$). By Theorem 6.7 of [16], the differential ∂ is given on the generators by the formulas

$$\partial(b_k) = 0 \quad \text{and} \quad \partial(a_n) = 1 + C_{n,n}. \tag{1}$$

(It is extended to $\mathcal{A}$ by linearity and the Leibniz rule.) Here, for any n,

$$C_{n,n} = \sum_{\{\, i_1, \ldots, i_c \,\} \in D_n} B_{n,i_1} B_{i_1,i_2} B_{i_2,i_3} \ldots B_{i_{c-1},i_c} B_{i_c,n}, \tag{2}$$

where two more terms require explanation.

Definition 2.1. A finite sequence of positive integers is called *admissible* if for all $s \geq 1$, between any two appearances of s in the sequence there is a number greater than s which appears between them. For $n \geq 1$, we denote by D_n the set of all admissible sequences whose elements are taken from the set $\{\, 1, 2, \ldots, n-1 \,\}$.

[2]All components of L_β have $r = 0$. If there are multiple components, what we describe here is only one of the admissible gradings.

Note that non-empty admissible sequences have a unique highest element.

Definition 2.2. Let $1 \leq i, j \leq q$. The element $B_{i,j}$ of the DGA of γ_β is the sum of the following products. For each path composed of parts of the strands of the braid (word) β that connects the left endpoint labeled i to the right endpoint labeled j so that it only turns around quadrants facing up, take the product of the labels of the crossings from left to right that it turns at. (We will refer to the paths contributing to $B_{i,j}$ as *paths in the braid.*)

We will also use the following notation: for any $i < j$, let

$$C_{i,j} = \sum_{\{\, i, i_1, \ldots, i_c \,\} \in D_j} B_{i,i_1} B_{i_1,i_2} B_{i_2,i_3} \ldots B_{i_{c-1},i_c} B_{i_c,j}. \tag{3}$$

The expressions $B_{i,j}$ and $C_{i,j}$ are elements of the DGA $\mathcal{A}$. Even though $\mathcal{A}$ is non-commutative, we will refer to them, as well as to similar expressions and even matrices with such entries, as polynomials.

3 The path matrix

The polynomials $B_{i,j}$ are naturally arranged in a $q \times q$ matrix B_β (with entries in $\mathcal{A}$), which we will call the *path matrix of* β.

If we substitute 0 for each crossing label of β, then B_β reduces to the matrix of the underlying permutation π of β:

$$B_\beta(0, 0, \ldots, 0) = \left[\delta_{\pi(i),j}\right] =: P_\pi,$$

where δ is the Kronecker delta. Note that B_β depends on the braid *word*, whereas P_π only on the braid itself.

Remark 3.1. When the braid group relation $\sigma_i \sigma_j = \sigma_j \sigma_i$, $|i - j| > 1$ is applied to change β, the diagram γ_β only changes by an isotopy of the plane and the path matrix B_β hardly changes at all. In fact if we don't insist on increasing label indices and re-label the braid as on the right side of Figure 2, then B_β remains the same. Therefore such changes in braid words will be largely ignored in the paper.

3.1 Multiplicativity

The path matrix behaves multiplicatively in the following sense: If two positive braid words β_1 and β_2 on q strands are multiplied as in the braid group (placing β_2 to the right of β_1) to form the braid word $\beta_1 * \beta_2$, then

$$B_{\beta_1 * \beta_2} = B_{\beta_1} \cdot B_{\beta_2}. \tag{4}$$

Note that for this to hold true, β_1 and β_2 have to carry their own individual crossing labels that $\beta_1 * \beta_2$ inherits, too. Otherwise, the observation is trivial: we may group together paths from left endpoint i to right endpoint j in $\beta_1 * \beta_2$ by the position of their crossing over from β_1 to β_2.

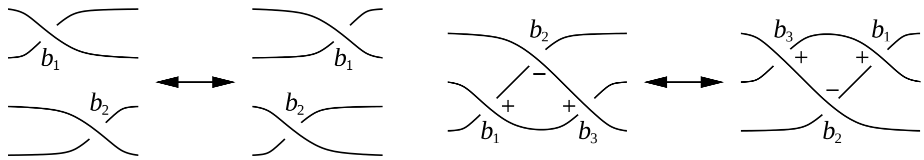

Figure 2: Labels before and after an isotopy and a Reidemeister III move. The signs on the right are the so called Reeb signs.

Remark 3.2. Apart from the technicality of having to view B_{β_1} and B_{β_2} as polynomials of separate sets of indeterminates, there are other problems that so far prevented the author from defining a representation of the positive braid semigroup based on (4). Namely, when we represent the same positive braid by a different braid word, the path matrix changes. This can be somewhat controlled by requiring, as another departure from our convention of increasing label subscripts, that whenever the braid group relation $\sigma_i\sigma_{i+1}\sigma_i = \sigma_{i+1}\sigma_i\sigma_{i+1}$ is applied to change β, the two sets of labels are related as on the right side of Figure 2. Then the path matrix changes

$$\text{from } \begin{bmatrix} b_2 & b_3 & 1 \\ b_1 & 1 & 0 \\ 1 & 0 & 0 \end{bmatrix} \text{ to } \begin{bmatrix} b_2 + b_3 b_1 & b_3 & 1 \\ b_1 & 1 & 0 \\ 1 & 0 & 0 \end{bmatrix}. \tag{5}$$

Notice that this is just an application of Chekanov's chain map [3] relating the DGA's of the diagrams before and after a Reidemeister III move (and the same happens if the triangle is part of a larger braid). Therefore we may hope that the path matrix of a positive braid β, with its *entries viewed as elements of the relative contact homology* $H(L_\beta)$, is independent of the braid word representing β. This is indeed the case because the set of equivalent positive geometric braids (with the endpoints of strands fixed[3]) is contractible, thus it is possible to canonically identify the contact homologies coming from different diagrams. But because there isn't any known relation between the contact homologies of L_{β_1}, L_{β_2}, and $L_{\beta_1 * \beta_2}$, this doesn't help us.

The path matrix of the braid group generator σ_i, with its single crossing labeled b is block-diagonal with only two off-diagonal entries:

$$B_{\sigma_i} = \begin{bmatrix} I_{i-1} & & & \\ & b & 1 & \\ & 1 & 0 & \\ & & & I_{q-i-1} \end{bmatrix}. \tag{6}$$

By (4), all path matrices are products of such elementary matrices.

Example 3.3. Consider the braid β shown in Figure 3. Its path matrix is

$$B_\beta = \begin{bmatrix} B_{1,1} & B_{1,2} \\ B_{2,1} & B_{2,2} \end{bmatrix} = \begin{bmatrix} b_1 & 1 \\ 1 & 0 \end{bmatrix} \begin{bmatrix} b_2 & 1 \\ 1 & 0 \end{bmatrix} \begin{bmatrix} b_3 & 1 \\ 1 & 0 \end{bmatrix} = \begin{bmatrix} b_1 + b_3 + b_1 b_2 b_3 & 1 + b_1 b_2 \\ 1 + b_2 b_3 & b_2 \end{bmatrix}.$$

[3] I.e., conjugation is not allowed here; if it was, the space in question would not be contractible any more, as demonstrated in [16].

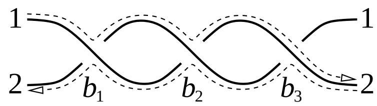

Figure 3: Trefoil braid

(The path contributing b_1b_2 to $B_{1,2}$ is shown.) As $D_1 = \{\, \varnothing \,\}$ and $D_2 = \{\, \varnothing, \{\, 1 \,\}\, \}$, we have $C_{1,1} = B_{1,1} = b_1 + b_3 + b_1b_2b_3$ and $C_{2,2} = B_{2,2} + B_{2,1}B_{1,2} = b_2 + (1 + b_2b_3)(1 + b_1b_2)$. Thus in the DGA of γ_β, the relations $\partial a_1 = 1 + b_1 + b_3 + b_1b_2b_3$ and $\partial a_2 = 1 + b_2 + (1 + b_2b_3)(1 + b_1b_2) = b_2 + b_2b_3 + b_1b_2 + b_2b_3b_1b_2$ hold.

3.2 Inverse matrix

The inverse of the elementary matrix B_{σ_i} is

$$B_{\sigma_i}^{-1} = \begin{bmatrix} I_{i-1} & & & \\ & 0 & 1 & \\ & 1 & b & \\ & & & I_{q-i-1} \end{bmatrix}.$$

Therefore, writing $\beta = \sigma_{i_1}\sigma_{i_2}\cdots\sigma_{i_w}$, from $B_\beta^{-1} = \left(B_{\sigma_{i_1}}B_{\sigma_{i_2}}\cdots B_{\sigma_{i_w}}\right)^{-1} = B_{\sigma_{i_w}}^{-1}\cdots B_{\sigma_{i_1}}^{-1}$ we see that B_β^{-1} is also a path matrix of the same braid word β, but in a different sense. This time, the (i, j)–entry is a sum of the following products: For each path composed of parts of the strands of β that connects the *right* endpoint labeled i to the *left* endpoint labeled j so that it only turns at quadrants facing *down*, take the product of the crossings from right to left that it turns at. So it's as if we turned β upside down by a 180° rotation while keeping the original labels of the crossings and of the endpoints of the strands.

That operation on the braid word produces a Legendrian isotopic closure (where by closure we mean adding strands above the braid, as in Figure 1). This is seen by a two-step process. First, apply 'half-way' the conjugation move of [16] (as in Figure 4) successively to each crossing of β from left to right. This turns β upside down, but now the closing strands are underneath.

Then, repeat q times the procedure shown in Figure 5, which we borrow from [13]. The box may contain any front diagram. Before the move represented by the third arrow, we make the undercrossing strand on the left steeper than all slopes that occur inside the box, so that it slides underneath the entire diagram

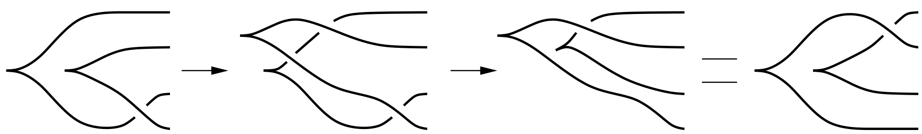

Figure 4: First half of a conjugation move (sequence of two Reidemeister II moves)

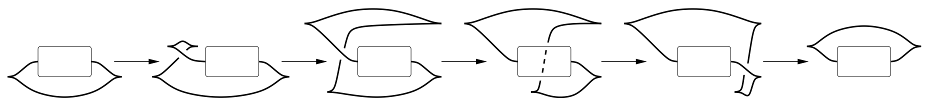

Figure 5: Moving a strand to the other side of a front diagram.

without a self-tangency moment. (In 3–space, increasing the slope results in a huge y–coordinate. Recall that fronts appear on the xz–plane, in particular the y–axis points away from the observer. So the motion of the strand happens far away, way behind any other piece of the knot.)

Example 3.4. The inverse of the matrix from the previous example is

$$B_\beta^{-1} = \begin{bmatrix} 0 & 1 \\ 1 & b_3 \end{bmatrix} \begin{bmatrix} 0 & 1 \\ 1 & b_2 \end{bmatrix} \begin{bmatrix} 0 & 1 \\ 1 & b_1 \end{bmatrix} = \begin{bmatrix} b_2 & 1 + b_2 b_1 \\ 1 + b_3 b_2 & b_3 + b_1 + b_3 b_2 b_1 \end{bmatrix}.$$

In Figure 3, the path contributing $b_3 b_2 b_1$ to the $(2,2)$ entry is shown.

3.3 Permutation braids

As an illustration, we examine the path matrices of permutation braids, which are positive braids in which every pair of strands crosses at most once. They are in a one-to-one correspondence with elements of the symmetric group S_q and they play a crucial role in Garside's solution [14] of the word and conjugacy problems in the braid group B_q.

It is always possible to represent a braid with a braid word in which the product $\sigma_i \sigma_{i+1} \sigma_i$ doesn't appear for any i. (That is, all possible triangle moves in which the "middle strand is pushed down," as in Figure 2 viewed from the right to the left, have been performed.) Such *reduced braid words* for permutation braids (up to the relation $\sigma_i \sigma_j = \sigma_j \sigma_i$, $|i - j| > 1$; see Remark 3.1) are unique.

Proposition 3.5. *Let $\pi \in S_q$. The path matrix B_π associated to its reduced permutation braid word is obtained from the permutation matrix P_π as follows. Changes are only made to entries that are above the* 1 *in their column and to the left of the* 1 *in their row. At each such position, a single crossing label appears in B_π.*

In particular, the positions that carry different entries in P_π and B_π are in a one-to-one correspondence with the inversions of π.

Proof. Starting at the left endpoint labeled i, our first "intended destination" (on the right side of the braid) is $\pi(i)$. Whenever we turn along a path in the braid, the intended destination becomes a smaller number because the two strands don't meet again. This shows that entries in B_π that are to the right of the 1 in their row are 0. Traversing the braid from right to left, we see that entries under the 1 in their column are 0, too. Either one of the two arguments shows that the 1's of P_π are left unchanged in B_π. (This part of the proof is

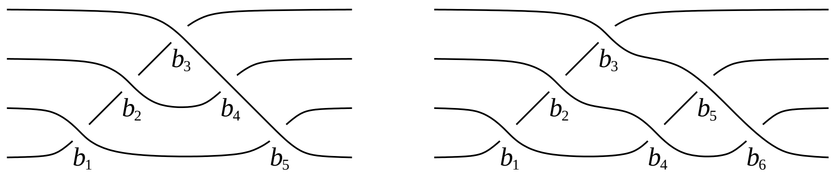

Figure 6: Permutation braids of (14) and of (14)(23) (the latter also known as the Garside braid Δ_4).

valid for any positive braid word representing a permutation braid; cf. Figure 2 and equation (5).)

We claim that any path in the braid contributing to any $B_{i,j}$ can contain at most one turn. Assume the opposite: then a strand s crosses under the strand t_1 and then over the strand t_2, which are different and which have to cross each other as well. This contradicts our assumption that the braid word is reduced, for it is easy to argue that (in a permutation braid) the triangle s, t_1, t_2 that we have just found must contain an elementary triangle as on the right side of Figure 2.

So the paths we have not yet enumerated are those with exactly one turn. Because strands cross at most once, these contribute to different matrix entries. Finally, if (i, j) is a position as described in the Proposition, then $\pi(i) > j$ and $\pi^{-1}(j) > i$. This means that the strand starting at i has to meet the strand ending at j, so that the label of that crossing becomes $B_{i,j}$. □

Example 3.6. The transposition (14) of S_4 is represented by the reduced braid word shown in Figure 6. It contains 5 inversions, corresponding to the 5 crossings of the braid. Its path matrix is $B_{(14)} = \begin{bmatrix} b_3 & b_4 & b_5 & 1 \\ b_2 & 1 & 0 & 0 \\ b_1 & 0 & 1 & 0 \\ 1 & 0 & 0 & 0 \end{bmatrix}$. The path matrix of the Garside braid (half-twist) Δ_4, also shown in Figure 6, is $\begin{bmatrix} b_3 & b_5 & b_6 & 1 \\ b_2 & b_4 & 1 & 0 \\ b_1 & 1 & 0 & 0 \\ 1 & 0 & 0 & 0 \end{bmatrix}$.

The latter pattern obviously generalizes to Δ_n for any n.

3.4 Row reduction

There is yet another way to factorize the path matrix. Let $\tau_i \in S_q$ denote the underlying permutation (transposition) of the elementary braid $\sigma_i \in B_q$.

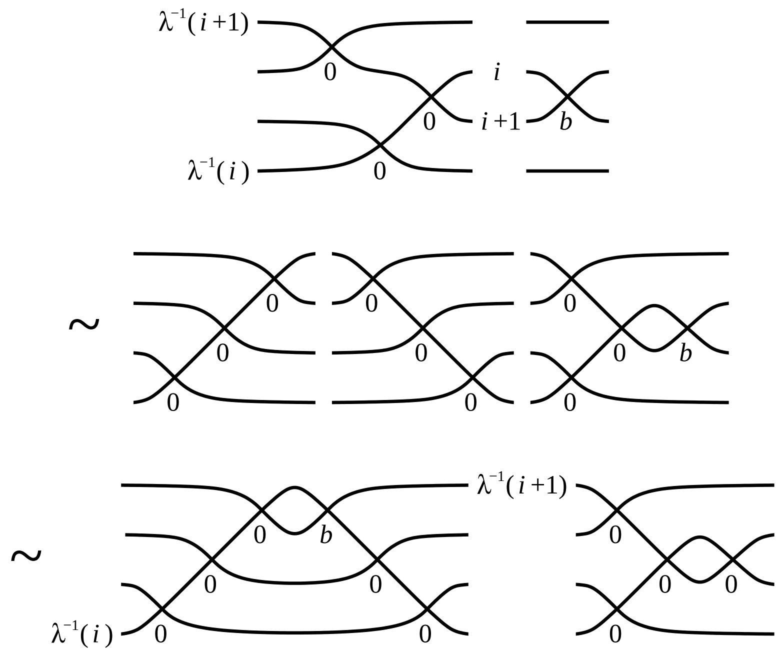

Figure 7: Braids so decorated that they have the same path matrix

Lemma 3.7. *Let $\lambda \in S_q$ be an arbitrary permutation. Then for all i,*

$$\begin{bmatrix} \textit{matrix} \\ \textit{of } \lambda \end{bmatrix} \cdot \begin{bmatrix} I_{i-1} & & & \\ & b & 1 & \\ & 1 & 0 & \\ & & & I_{q-i-1} \end{bmatrix} = \begin{bmatrix} 1 & & & \\ & 1 & & \\ & b & \ddots & \\ & & & 1 \end{bmatrix} \cdot \begin{bmatrix} \textit{matrix} \\ \textit{of } \tau_i \circ \lambda \end{bmatrix}, \quad (7)$$

where in the first term of the right hand side, the single non-zero off-diagonal entry b appears in the position $\lambda^{-1}(i), \lambda^{-1}(i+1)$.

Proof. The essence of the proof is in Figure 7. It will be crucial that the path matrix depends on how the braid is decorated with labels. On the other hand, for the purposes of the argument, over- and undercrossing information in the braids is irrelevant. In fact, although we will not change our terminology, we will actually think of them (in particular, when we take an inverse) as words written in the generators $\tau_1, \ldots, \tau_{q-1}$ of S_q.

Take the permutation braid for λ (or choose any other positive braid word with this underlying permutation) and label its crossings with zeros. (In Figure 7 we used $\lambda = (1342) \in S_4$ as an example.) Add a single generator σ_i, with its crossing labeled b to it (Figure 7 shows $i = 2$). The left hand side of (7) is the path matrix of this braid β.

Next, choose any positive braid word μ in which the strands with right endpoints $\lambda^{-1}(i)$, $\lambda^{-1}(i+1)$ cross (say exactly once) and form the product $\mu^{-1} * \mu * \beta$. Label the crossings of μ^{-1} and μ with zeros, as in the middle of Figure 7. This way, the path matrix does not change.

Now, it does not matter for the path matrix where exactly the single non-zero label b appears in the braid as long as that crossing establishes a path between the same two endpoints. In other words, we may move the label from the first (from the right) to the third, fifth etc. crossing of the same two strands. By construction, one of those crossings is either in μ (if $\lambda^{-1}(i) > \lambda^{-1}(i+1)$, as is the case in Figure 7) or in μ^{-1}, and we move the label there (bottom of Figure 7). When we read off the path matrix from this form, we obtain the right hand side of (7): The path matrix of $\mu^{-1} * \mu$ is I_q except for the single b that establishes a path from $\lambda^{-1}(i)$ to $\lambda^{-1}(i+1)$, and the path matrix of β, now labeled with only zeros, is $P_{\tau_i \circ \lambda}$. ☐

Next, for the positive braid word $\beta = \sigma_{i_1}\sigma_{i_2}\cdots\sigma_{i_w}$ with crossings labeled $b_1, b_2, \ldots, b_w$, we'll introduce a sequence of elementary matrices. The underlying permutation is $\pi = \tau_{i_w}\ldots\tau_{i_2}\tau_{i_1}$. Let us denote the "permutation up to the k'th crossing" by $\pi_k = \tau_{i_k}\ldots\tau_{i_1}$, so that $\pi_0 = \text{id}$ and $\pi_w = \pi$. Let A_k be the $q \times q$ identity matrix with a single non-zero off-diagonal entry of b_k added in the position $\pi_{k-1}^{-1}(i_k), \pi_{k-1}^{-1}(i_k+1)$. Note that because we work over $\mathbf{Z}_2$, $A_k^2 = I_q$ for all k.

Proposition 3.8. *For the positive braid word $\beta = \sigma_{i_1}\sigma_{i_2}\cdots\sigma_{i_w}$ with underlying permutation π, we have $B_\beta = A_1 A_2 \ldots A_w P_\pi$, where P_π is the permutation matrix.*

Proof. If $\lambda = \pi_{k-1}$, $i = i_k$, and $b = b_k$, then equation (7) reads $P_{\pi_{k-1}} B_{\sigma_{i_k}} = A_k P_{\pi_k}$. Starting from $B_\beta = (P_{\pi_0} B_{\sigma_{i_1}}) B_{\sigma_{i_2}} \ldots B_{\sigma_{i_w}}$, we apply Lemma 3.7 w times. ☐

Read in another way, this result shows that B_β reduces to P_π by applying a particular sequence of elementary row operations: $A_w \ldots A_1 B_\beta = P_\pi$. This works in the non-commutative sense.

Example 3.9. For the braid β of Figure 3, we have

$$B_\beta = \begin{bmatrix} 1 & b_1 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} 1 & 0 \\ b_2 & 1 \end{bmatrix} \begin{bmatrix} 1 & b_3 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix},$$

that is

$$\begin{bmatrix} 1 & b_3 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} 1 & 0 \\ b_2 & 1 \end{bmatrix} \begin{bmatrix} 1 & b_1 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} b_1 + b_3 + b_1 b_2 b_3 & 1 + b_1 b_2 \\ 1 + b_2 b_3 & b_2 \end{bmatrix} = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}.$$

4 Algebraic results

In this section, we treat (re-define, if you like) the symbols $B_{i,j}$ as independent variables. Instead of $\mathbf{Z}_2$–coefficients, we will work in the free non-commutative unital ring generated by these symbols (where $1 \le i, j \le q$) over $\mathbf{Z}$. (The change of coefficients doesn't just make the results of this section slightly more general, but it is also hoped that it will facilitate the generalization of theorems in other sections to the context of contact homology with group ring coefficients [10].) After the first set of statements, we will abelianize so that we can consider determinants.

Note that the $C_{i,j}$ (equation (3)) are polynomials in the $B_{i,j}$. To state our results, we will need a similar family of polynomials whose definition is based on the notion of admissible sequence (Definition 2.1).

Definition 4.1. For any $1 \le i, j \le q$, let

$$M_{i,j} = \sum_{\{ i_1,\dots,i_c \} \in D_{\min\{i,j\}}} B_{i,i_1} B_{i_1,i_2} B_{i_2,i_3} \dots B_{i_{c-1},i_c} B_{i_c,j}.$$

Note that $M_{1,j} = B_{1,j}$, $M_{i,1} = B_{i,1}$, $M_{n,n} = C_{n,n}$, and $M_{i-1,i} = C_{i-1,i}$, whenever these expressions are defined.

Lemma 4.2.

$$\begin{bmatrix} 1 & C_{1,2} & C_{1,3} & \cdots & C_{1,q} \\ & 1 & C_{2,3} & \cdots & C_{2,q} \\ & & 1 & \cdots & C_{3,q} \\ & & & \ddots & \vdots \\ & & & & 1 \end{bmatrix} \cdot \begin{bmatrix} 1 & -M_{1,2} & -M_{1,3} & \cdots & -M_{1,q} \\ & 1 & -M_{2,3} & \cdots & -M_{2,q} \\ & & 1 & \cdots & -M_{3,q} \\ & & & \ddots & \vdots \\ & & & & 1 \end{bmatrix}$$

$$= \begin{bmatrix} 1 & -M_{1,2} & -M_{1,3} & \cdots & -M_{1,q} \\ & 1 & -M_{2,3} & \cdots & -M_{2,q} \\ & & 1 & \cdots & -M_{3,q} \\ & & & \ddots & \vdots \\ & & & & 1 \end{bmatrix} \cdot \begin{bmatrix} 1 & C_{1,2} & C_{1,3} & \cdots & C_{1,q} \\ & 1 & C_{2,3} & \cdots & C_{2,q} \\ & & 1 & \cdots & C_{3,q} \\ & & & \ddots & \vdots \\ & & & & 1 \end{bmatrix} = I_q,$$

and a similar statement can be formulated for lower triangular matrices.

Note that the two claims don't imply each other because we work over a non-commutative ring.

Proof. We need that for all $1 \le i < j \le q$,

$$-M_{i,j} - C_{i,i+1} M_{i+1,j} - C_{i,i+2} M_{i+2,j} - \dots - C_{i,j-1} M_{j-1,j} + C_{i,j} = 0$$

and that

$$C_{i,j} - M_{i,i+1} C_{i+1,j} - M_{i,i+2} C_{i+2,j} - \dots - M_{i,j-1} C_{j-1,j} - M_{i,j} = 0.$$

We may view both of these equalities as identities for $C_{i,j}$. The first one groups the terms of $C_{i,j}$ according to the highest element of the admissible sequence. The second groups them according to the first element which is greater than i. The lower triangular version is analogous. $\square$

Lemma 4.3. *For all $1 \leq n \leq q$,*

$$\begin{bmatrix} -1 & & & & \\ M_{2,1} & -1 & & & \\ M_{3,1} & M_{3,2} & -1 & & \\ \vdots & \vdots & \vdots & \ddots & \\ M_{n,1} & M_{n,2} & M_{n,3} & \cdots & -1 \end{bmatrix} \cdot \begin{bmatrix} 1 & -M_{1,2} & -M_{1,3} & \cdots & -M_{1,n} \\ & 1 & -M_{2,3} & \cdots & -M_{2,n} \\ & & 1 & \cdots & -M_{3,n} \\ & & & \ddots & \vdots \\ & & & & 1 \end{bmatrix}$$

$$= \begin{bmatrix} -1 & B_{1,2} & \cdots & B_{1,i} & \cdots & B_{1,n} \\ B_{2,1} & -1 - B_{2,1}B_{1,2} & \cdots & B_{2,i} & \cdots & B_{2,n} \\ \vdots & \vdots & \ddots & \vdots & & \vdots \\ B_{i,1} & B_{i,2} & \cdots & B_{i,i} - C_{i,i} - 1 & \cdots & B_{i,n} \\ \vdots & \vdots & & \vdots & \ddots & \vdots \\ B_{n,1} & B_{n,2} & \cdots & B_{n,i} & \cdots & B_{n,n} - C_{n,n} - 1 \end{bmatrix}. \tag{8}$$

Proof. For entries above the diagonal ($i < j$), the claim is that

$$B_{i,j} = -M_{i,1}M_{1,j} - M_{i,2}M_{2,j} - \ldots - M_{i,i-1}M_{i-1,j} + M_{i,j}.$$

Viewing this as an identity for $M_{i,j}$, we see that it holds because terms are grouped with respect to the highest element in the admissible sequence. The reasoning is the same for positions below the diagonal. For the diagonal entries, we need to show that

$$B_{i,i} - C_{i,i} - 1 = -M_{i,1}M_{1,i} - M_{i,2}M_{2,i} - \ldots - M_{i,i-1}M_{i-1,i} - 1.$$

Isolating $C_{i,i}$ this time, we again see a separation of its terms according to the highest element of the admissible sequence. □

For the rest of the section, we will work in the *commutative* polynomial ring generated over $\mathbf{Z}$ by the $B_{i,j}$, so that we can talk about determinants.

Theorem 4.4. *The ideal I' generated by the polynomials*

$$1 + C_{1,1}, \quad 1 + C_{2,2}, \quad \ldots, \quad 1 + C_{q,q}$$

agrees with the ideal I generated by the polynomials

$$L_1 = B_{1,1} + 1, \;\; L_2 = \begin{vmatrix} B_{1,1} & B_{1,2} \\ B_{2,1} & B_{2,2} \end{vmatrix} - 1, \;\; \ldots, \;\; L_q = \begin{vmatrix} B_{1,1} & \cdots & B_{1,q} \\ \vdots & \ddots & \vdots \\ B_{q,1} & \cdots & B_{q,q} \end{vmatrix} - (-1)^q.$$

Proof. Let $n \leq q$ and take determinants of both sides of equation (8): $(-1)^n$, on the left hand side, agrees with $\begin{vmatrix} B_{1,1} & \cdots & B_{1,n} \\ \vdots & \ddots & \vdots \\ B_{n,1} & \cdots & B_{n,n} \end{vmatrix}$ plus an element of I' on the right hand side. Thus, $L_n \in I'$ for all n.

The proof of the other containment relation is also based on equation (8) and goes by induction on n. Note that $1+C_{1,1}=L_1$ and assume that $1+C_{1,1}, 1+C_{2,2}, \ldots, 1+C_{n-1,n-1}$ are all in I (actually, they are in the ideal generated by $L_1, L_2, \ldots, L_{n-1}$). Re-writing the determinant of the matrix on the right hand side of (8), we find that

$$(-1)^n = \begin{vmatrix} -1 & B_{1,2} & \cdots & B_{1,n-1} & B_{1,n} \\ B_{2,1} & -1-B_{2,1}B_{1,2} & \cdots & B_{2,n-1} & B_{2,n} \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ B_{n-1,1} & B_{n-1,2} & \cdots & B_{n-1,n-1}-C_{n-1,n-1}-1 & B_{n-1,n} \\ B_{n,1} & B_{n,2} & \cdots & B_{n,n-1} & B_{n,n} \end{vmatrix}$$
$$- \begin{vmatrix} -1 & B_{1,2} & \cdots & B_{1,n-1} & B_{1,n} \\ B_{2,1} & -1-B_{2,1}B_{1,2} & \cdots & B_{2,n-1} & B_{2,n} \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ B_{n-1,1} & B_{n-1,2} & \cdots & B_{n-1,n-1}-C_{n-1,n-1}-1 & B_{n-1,n} \\ 0 & 0 & \cdots & 0 & 1+C_{n,n} \end{vmatrix}.$$

Notice that the second determinant is $(-1)^{n-1}(1+C_{n,n})$ (by Lemma 4.3), while the first is $\begin{vmatrix} B_{1,1} & \cdots & B_{1,n} \\ \vdots & \ddots & \vdots \\ B_{n,1} & \cdots & B_{n,n} \end{vmatrix}$ plus an element of I', but the latter, by the inductive hypothesis, is also in I. Isolating $1+C_{n,n}$, we are done. □

So we see that the ideal I defined in terms of the upper left corner subdeterminants of the general determinant is also generated by the polynomials $1+C_{n,n}$, which arise from contact homology (counting holomorphic discs). In fact much more is true: the $1+C_{n,n}$ form the reduced Gröbner basis for I. Of course this can only be true for certain term orders that we'll describe now.

In the (commutative) polynomial ring $\mathbf{Z}[B_{i,j}]$, take any order $\prec$ of the indeterminates where any diagonal entry $B_{i,i}$ is larger than any off-diagonal one. Extend this order to the monomials lexicographically. (But not degree lexicographically! For example, $B_{2,2} \succ B_{2,1}B_{1,2}$.) This is a multiplicative term order.

Theorem 4.5. *The polynomials $1+C_{n,n}$, $n=1,\ldots,q$ (defined in equation* (2)*), form the reduced Gröbner basis for the ideal*

$$I = \left\langle B_{1,1}+1, \quad \begin{vmatrix} B_{1,1} & B_{1,2} \\ B_{2,1} & B_{2,2} \end{vmatrix} - 1, \quad \ldots, \quad \begin{vmatrix} B_{1,1} & \cdots & B_{1,q} \\ \vdots & \ddots & \vdots \\ B_{q,1} & \cdots & B_{q,q} \end{vmatrix} - (-1)^q \right\rangle$$

under any of the term orders $\prec$ described above.

Proof. This is obvious from the definitions (see for example [2]), after noting that the initial term of $1+C_{n,n}$ is $B_{n,n}$ and that by the definition of an admissible sequence, no other term in $1+C_{n,n}$ contains any $B_{i,i}$. (The initial ideal of I is that generated by the $B_{n,n}$.) □

5 Augmentations

Definition 5.1. Let γ be a Lagrangian diagram of a Legendrian link L. If L has more than one components, we assume that an admissible grading of the DGA of γ has been chosen, too. An *augmentation* is a subset X of the crossings (the *augmented crossings*) of γ with the following properties.

- The index of each element of X is 0.
- For each generator a of index 1, the number of admissible discs with positive corner a and all negative corners in X is even.

Here, an admissible disc is the central object of Chekanov–Eliashberg theory: These discs determine the differential ∂ of the DGA $\mathcal{A}$, and thus contact homology $H(L)$. Unlike most of the literature, we expand the notion of augmentation here (in the multi-component case) by allowing 'mixed' crossings between different components to be augmented, as long as they have index 0 in the one grading we have chosen. Such sets of crossings would typically not be augmentations for other admissible gradings because it's exactly the index of a mixed crossing that is ambiguous. Our motivation is that γ_β, even if it is of multiple components, has the natural admissible grading introduced in section 2.

The evaluation homomorphism (which is defined on the link DGA, and which is also called an augmentation) $\varepsilon_X \colon \mathcal{A} \to \mathbf{Z}_2$ that sends elements of X to 1 and other generators to 0, gives rise to an algebra homomorphism $(\varepsilon_X)_* \colon H(L) \to \mathbf{Z}_2$. In fact, the second requirement of Definition 5.1 is just an elementary way of saying that ε_X vanishes on $\partial(a)$ for each generator a of index 1, while for other indices this is already automatic by the first point and the fact that ∂ lowers the index by 1.

Remark 5.2. As a preview of a forthcoming paper, let us mention that augmentations do define a Legendrian isotopy invariant in the following sense: the set of all induced maps $(\varepsilon_X)_* \colon H(L) \to \mathbf{Z}_2$ depends only on L. (The correspondence between augmentations of different diagrams of L is established using pull-backs by the isomorphisms constructed in Chekanov's proof of the invariance of $H(L)$.) The number of augmentations in the sense of Definition 5.1 may however change by a factor of 2 when a Reidemeister II move or its inverse, involving crossings of index 0 and -1, is performed.

In practice, finding an augmentation means solving a system of polynomial equations (one equation provided by each index 1 crossing) over $\mathbf{Z}_2$. In this sense, augmentations form a variety. In this section we prove a few statements about the variety associated to γ_β.

The main result is the following theorem, which allows for an enumeration of all augmentations of γ_β. The author is greatly indebted to Supap Kirtsaeng, who wrote a computer program based on this criterion. It may first seem ineffective to check all subsets of the crossings of β, but it turns out that a significant portion of them are augmentations (see section 7).

Let Y be a subset of the crossings of β. Let $\varepsilon_Y \colon \mathcal{A} \to \mathbf{Z}_2$ be the evaluation homomorphism that sends elements of Y to 1 and other generators to 0. In particular, we may talk of the 0-1–matrix $\varepsilon_Y(B_\beta)$. (This could also have been denoted by $B_\beta(\chi_Y)$, where the 0-1–sequence χ_Y is the characteristic function of Y.)

Theorem 5.3. *Let Y be a subset of the crossings of the positive braid word β. Y is an augmentation of γ_β if and only if the* 0-1*–matrix $\varepsilon_Y(B_\beta)$ is such that every upper left corner square submatrix of it has determinant* 1*.*

It is then a classical theorem of linear algebra that the condition on $\varepsilon_Y(B_\beta)$ is equivalent to the requirement that it possess an *LU*–decomposition and also to the requirement that Gaussian elimination can be completed on it without permuting rows.

Proof. In our admissible grading, each crossing of β has index 0. Therefore Y is an augmentation if and only if ε_Y vanishes on $\partial(a)$ for each index 1 DGA generator a. This in turn is clearly equivalent to saying that ε_Y vanishes on the two-sided ideal generated by these polynomials. In fact because ε_Y maps to a commutative ring ($\mathbf{Z}_2$), we may abelianize $\mathcal{A}$ and say that the condition for Y to be an augmentation is that ε_Y vanishes on the ideal generated by the expressions $\partial(a_1), \ldots, \partial(a_q)$, which are now viewed as honest polynomials in the commuting indeterminates $b_1, \ldots, b_w$.

In [16], section 6, we computed these polynomials and found that they really were polynomials of the polynomials $B_{i,j}$, as stated in equation (1). Now by (the modulo 2 reduction of) Theorem 4.4, the ideal generated by the $\partial(a_n)$ is also generated by the polynomials

$$B_{1,1}+1, \quad \begin{vmatrix} B_{1,1} & B_{1,2} \\ B_{2,1} & B_{2,2} \end{vmatrix} + 1, \quad \ldots, \quad \begin{vmatrix} B_{1,1} & \cdots & B_{1,q} \\ \vdots & \ddots & \vdots \\ B_{q,1} & \cdots & B_{q,q} \end{vmatrix} + 1,$$

which implies the Theorem directly. □

Remark 5.4. Notice that for a path matrix B_β, a quick look at (6) with formula (4) implies that we always have $\det(B_\beta) = 1$. Therefore the condition on the $q \times q$ subdeterminant is vacuous: if a subset of the crossings of β "works as an augmentation" for $a_1, \ldots, a_{q-1}$, then it automatically works for a_q as well.

Let us give a geometric explanation of the appearance of *LU*–decompositions. Figure 8 shows another Lagrangian diagram of L_β that is obtained from the front diagram f_β by pushing all the right cusps to the extreme right and then applying resolution. This has the advantage that all admissible discs are embedded. Label the $q(q-1)$ new crossings as in Figure 8. Our preferred grading is extended to the new crossings by assigning 0 to the $c_{i,j}$ and 1 to the $s_{i,j}$. This implies $\partial(c_{i,j}) = 0$, while the index 1 generators are mapped as follows:

$$\partial(a_n) = 1 + c_{n,1}B_{1,n} + \ldots + c_{n,n-1}B_{n-1,n} + B_{n,n}$$

and

$$\partial(s_{i,j}) = c_{i,1}B_{1,j} + \ldots + c_{i,i-1}B_{i-1,j} + B_{i,j}.$$

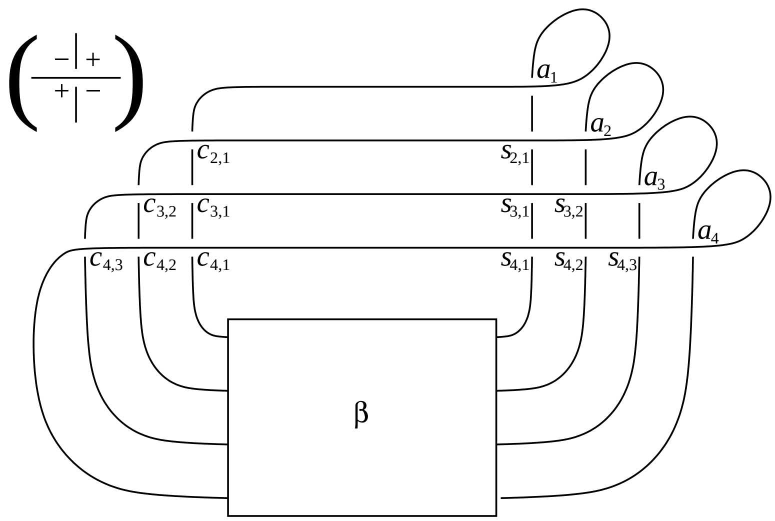

Figure 8: Another Lagrangian diagram of L_β

Setting the latter $q+\binom{q}{2}$ expressions equal to 0 is equivalent to saying that the matrix product

$$\begin{bmatrix} 1 & 0 & 0 & \cdots & 0 \\ c_{2,1} & 1 & 0 & \cdots & 0 \\ c_{3,1} & c_{3,2} & 1 & \cdots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ c_{q,1} & c_{q,2} & c_{q,3} & \cdots & 1 \end{bmatrix} \begin{bmatrix} B_{1,1} & B_{1,2} & B_{1,3} & \cdots & B_{1,q} \\ B_{2,1} & B_{2,2} & B_{2,3} & \cdots & B_{2,q} \\ B_{3,1} & B_{3,2} & B_{3,3} & \cdots & B_{3,q} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ B_{q,1} & B_{q,2} & B_{q,3} & \cdots & B_{q,q} \end{bmatrix}$$

is unit upper triangular. Thus an augmentation evaluates B_β to an *LU*–decomposable 0-1–matrix and the converse is not hard to prove either.

6 Rulings

Definition 6.1. An *ungraded ruling* is a partial splicing of a front diagram where certain crossings, called *switches*, are replaced by a pair of arcs as in Figure 9 so that the diagram becomes a (not necessarily disjoint) union of standard unknot diagrams, called *eyes*. (An eye is a pair of arcs connecting the same two cusps that contain no other cusps and that otherwise do not meet, not even at switches.) It is assumed that in the vertical ($x=$ const.) slice of the diagram through each switch, the two eyes that meet at the switch follow one of the three configurations in the middle of Figure 9.

Let us denote the set of all ungraded rulings of a front diagram f of a Legendrian link by $\Gamma_1(f)$. We get 2–graded rulings, forming the set $\Gamma_2(f)$, if we require that the index of each switch be even. $\mathbf{Z}$–graded rulings (set $\Gamma_0(f)$) are those where each switch has index 0.

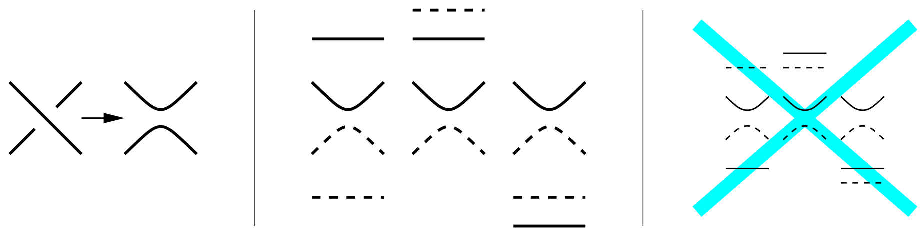

Figure 9: Allowed and disallowed configurations for switches of rulings

Γ_1 is of course grading-independent. For multi-component oriented link diagrams, Γ_2 doesn't depend on the chosen grading, but Γ_0 might.

Rulings can also be classified by the value

$$\theta = \text{number of eyes} - \text{number of switches}.$$

The counts of ungraded, 2–graded, and **Z**–graded rulings with a given θ are all Legendrian isotopy invariants[4] [5, 11]. (In particular, the sizes of the sets $\Gamma_i(f)$, $i = 0, 1, 2$, don't depend on f, only on the Legendrian isotopy class.) We may arrange these numbers as coefficients in the ruling polynomials[5]

$$R_i(z) = \sum_{\rho \in \Gamma_i} z^{1-\theta(\rho)}.$$

Sabloff [22] notes that the existence of a 2–graded ruling implies $r = 0$. Let us add that if we treat the eyes as discs and join them by twisted bands at the switches, then a 2–graded ruling becomes an orientable surface. The number θ is its Euler characteristic and thus $\theta + \mu$, where μ is the number of the components of the Legendrian, is even. In particular, θ is odd for any 2–graded ruling of a Legendrian knot.

There is a marked difference between **Z**–graded rulings and the two less restrictive cases. R_1 and R_2 only depend on the smooth type of the Legendrian and its Thurston–Bennequin number. In fact, Rutherford [21] proved that for any link, $R_1(z)$ is the coefficient of a^{-tb-1} in the Dubrovnik version of the Kauffman polynomial, and that $R_2(z)$ is the coefficient of v^{tb+1} in the Homfly polynomial. On the other hand, R_0 is more sensitive: Chekanov [4] constructed two Legendrian knots of type 5_2, both with $tb = 1$ and $r = 0$, so that one has $R_0(z) = 1 + z^2$ and the other has $R_0(z) = 1$.

Because f_β only contains crossings of index 0, any ungraded ruling is automatically 2–graded and **Z**–graded in this case. Thus we may talk about a single ruling polynomial. By Rutherford's theorems, this implies that the coefficients

[4]In the **Z**–graded case, we may have to assume that the Legendrian has a single component.

[5]These are honest polynomials for knots, but for multi-component links, they may contain negative powers of z. It may seem unnatural first to write them the way we do, but there are two good reasons to do so: One is Rutherford's pair of theorems below, and the other is that rulings can also be thought of as surfaces, in which case θ becomes their Euler characteristic and (in the one-component and 2–graded case) $1 - \theta$ is twice their genus.

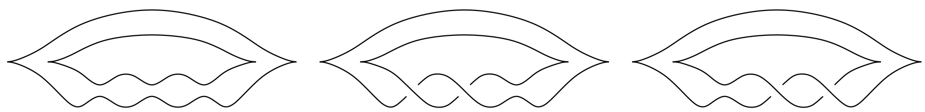

Figure 10: The Seifert ruling and the other two rulings of the positive trefoil

of the terms with minimum v–degree in the Homfly and Kauffman polynomials (for the latter, replace a with v^{-1} in its Dubrovnik version) of a braid-positive link agree. In fact, using Tanaka's results [24], the same can be said about arbitrary positive links. (See [17] for more.) This, without any reference to Legendrians yet with essentially the same proof, was first observed by Yokota [25].

Example 6.2. The positive trefoil knot that is the closure of the braid in Figure 3 has one ruling with $\theta = -1$ and two with $\theta = 1$, shown in Figure 10. The numbers 1 and 2 (i.e., the ruling polynomial $R(z) = 2 + z^2$) appear as the leftmost coefficients in the Homfly polynomial

$$\begin{array}{cc} z^2v^2 & \\ & \\ 2v^2 & -v^4 \end{array}$$

and also in the Kauffman polynomial

$$\begin{array}{cccc} z^2v^2 & & -z^2v^4 & \\ & -zv^3 & & +zv^5 \\ 2v^2 & & -v^4. & \end{array}$$

The diagram f_β admits many rulings. The one that is easiest to see is what we will call the *Seifert ruling*, in which the set of switches agrees with the set of crossings in β. This is the only ruling with the minimal value $\theta = q - w$. Another ruling, that one with the maximum value $\theta = \mu$, will be constructed in Theorem 6.3.

The second lowest possible value of θ for a ruling of f_β is $q - w + 2 = -tb(L_\beta) + 2$. It is easy to see that in such a ruling, the two crossings of β that are not switches have to be 'on the same level' (represented by the same braid group generator) without any other crossing between them on that level, and also that any such arrangement works. Thus, assuming that each generator occurs in the braid word β (i.e., that f_β is connected), the number of such rulings is $w - (q-1) = tb(L_\beta) + 1$. In all of the examples known to the author, the next value of θ, that is $\theta = q - w + 4 = -tb + 4$, is realized by exactly $\binom{w-q}{2} = \binom{tb}{2}$ rulings (but I don't know how to prove this). At $\theta = -tb + 6$ and higher, dependence on the braid occurs (see section 7).

It would be very interesting to have a test, similar to Theorem 5.3, that decides from the path matrix whether a given crossing set of f_β is a ruling.

From work of Fuchs, Ishkhanov [11, 12], and Sabloff [22], we know that $\mathbf{Z}$–graded rulings for a Legendrian exist if and only if augmentations do. Ng and

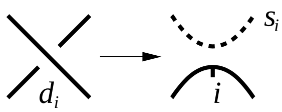

Figure 11: Splicing a crossing to create the path s_i; after removing s_i, a marker i is left on the remaining diagram.

Sabloff also worked out a surjective correspondence [20] that assigns a **Z**–graded ruling to each augmentation. In that correspondence, the size of the preimage of each **Z**–graded ruling ρ of the front diagram f is the number $2^{(\theta(\rho)+\chi^*(f))/2}$, where

$$\chi^*(f) = -\left(\sum_{\text{crossings } a \text{ of } f \text{ with } |a|<0} (-1)^{|a|}\right) + \left(\sum_{\text{crossings } a \text{ of } f \text{ with } |a|\geq 0} (-1)^{|a|}\right) - \text{number of right cusps.}$$

In particular, the number of augmentations belonging to ρ depends on $\theta(\rho)$ and the diagram only. (Note that χ^* has the same parity as tb, and because $r = 0$ is even, it also has the same parity as μ.) Thus the total number of augmentations is

$$R_0(z) \cdot z^{-1-\chi^*}\Big|_{z=2^{-1/2}}. \tag{9}$$

For the diagram f_β, which is without negatively graded crossings, we have $\chi^*(f_\beta) = tb(L_\beta) = w - q$. Thus among the rulings of f_β, the zeroth power of 2 corresponds only to the Seifert ruling. Therefore the number of augmentations of f_β is odd.

The next theorem further illuminates the relationship between augmentations and rulings.

Theorem 6.3. *For any positive braid word β, there exists a subset of its crossings which is (the set of switches in) a ruling of f_β and an augmentation of γ_β at the same time.*

The set we will construct is not, however, a fixed point of Ng and Sabloff's many-to-one correspondence.

Proof. The set X is constructed as follows: In β, the strands starting at the left endpoint 1 and ending at the right endpoint 1 either agree or intersect for an elementary geometric reason. In the latter case, splice/augment their first crossing from the left, d_1. In either case, remove the path s_1 connecting 1 to 1 from the braid. (If splicing was necessary to create s_1, then leave a marker 1 on the lower strand as shown in Figure 11.) Proceed by induction to find the paths $s_2, \ldots, s_q$ and for those s_i that were the result of splicing, leave a marker i and place the spliced crossing d_i in X.

The components of L_β are enumerated by the cycles in the permutation π that underlies β and the construction treats these components independently of one another. The number of elements in X is q minus the number μ of these cycles/components: it is exactly the largest element i of each cycle of π whose corresponding path s_i 'exists automatically,' without splicing. A way to see this is the following. Suppose π contains a single cycle. Unless $q = 1$, d_1 exists. When s_1 is removed from the braid, 1 is 'cut out' of π: in the next, smaller braid, the underlying permutation takes $\pi^{-1}(1)$ to $\pi(1)$. In particular, we still have a single cycle. Unless 2 is its largest (and only) element, d_2 will exist and the removal of s_2 cuts 2 out of the permutation. This goes on until we reach q, at which stage the braid is a single strand and more splicing is neither possible nor necessary.

We define an oriented graph G_β on the vertex set $\{1, 2, \ldots, q\}$ by the rule that an oriented edge connects i to j if s_j contains the marker i. Note that $i < j$ is necessary for this and that each i can be the starting vertex of at most one edge. For that reason, G_β doesn't even have unoriented cycles (consider the smallest number in a supposed cycle). Thus, G_β is a μ–component forest. The largest element of each tree is its only sink.

X is a ruling with the ith eye partially bounded by the path s_i. These are easily seen to satisfy Definition 6.1: if the ith and jth eyes meet at the switch d_i, then an edge connects i to j in G_β, thus $i < j$ and we see that in the vertical slice through d_i, we have the second of the admissible configurations of Figure 9. The value of θ for this ruling is μ.

To prove that X is also an augmentation, we'll check it directly using the analysis of admissible discs in γ_β from p. 2056 of [16]. Note that each of $a_1, \ldots, a_q$ (Figure 1) has a trivial admissible disc contributing 1 to its differential, so it suffices to show that for each j, there is exactly one more admissible disc with positive corner at a_j and all negative corners at crossings in X. In fact we will use induction to prove the following:

- For each j, this second disc Π_j will have either no negative corner or, if d_j exists, then exactly one negative corner at d_j.

- In the admissible sequence corresponding to Π_j, i appears if and only if G_β contains an oriented path from i to j, and each such i shows up exactly once.

The path s_1 completes the boundary of an admissible disc with positive corner at a_1. Because s_1 is removed in the first stage, no crossing along s_1 other than d_1 will be in X.

Now, assume that for each $j < n$, a unique disc Π_j exists with the said properties. Building a non-trivial admissible disc with positive corner at a_n, we start along the path s_n. (We will concentrate on the boundary of the admissible disc. Proposition 6.4 of [16] classifies, in terms of admissible sequences, which of the possible paths correspond to admissible discs.) When we reach a marker j, we are forced to enter $\partial\Pi_j$. Then by the inductive hypothesis, we have no other choice but to follow $\partial\Pi_j$ until we reach a_j. There, we travel around the jth trivial disc and continue along $\partial\Pi_j$, back to d_j and s_n. By the hypothesis, each a_i is visited at most once, so their sequence is admissible. At the next marker along s_n, a similar thing happens but using another, disjoint branch of G_β, so the sequence stays admissible.

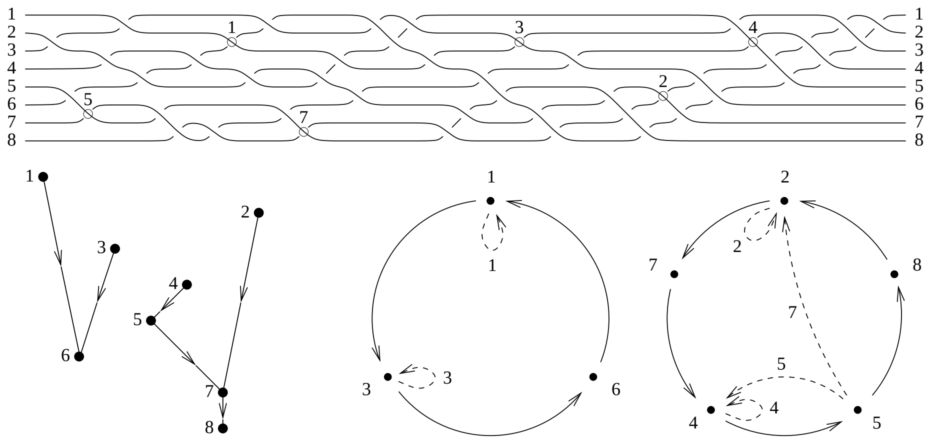

Figure 12: A braid β with an augmentation X (marked crossings) which is also a ruling. The forest graph is G_β and the other two graph components constitute the 'graph realized by X,' as in [16].

If d_n exists, then upon reaching it, we seemingly get a choice of turning or not. If we do turn, i.e. continue along s_n, then after a few more markers, we successfully complete the construction of Π_n. Because all markers along s_n were visited, it has both of the required properties.

We still have to rule out the option of not turning at d_n. Suppose that's what we do. Then we end up on a path s_m, where m is the endpoint of the edge of G_β starting at n; in particular, $m > n$. We may encounter markers along s_m, but the previous analysis applies to them and eventually we always return to s_m and exit the braid at the right endpoint m (or at an even higher number, in case we left s_m at d_m). But this is impossible by Lemma 6.2 of [16]. □

Remark 6.4. In [16], we used a two-component link of the braid-positive knots 8_{21} and $16n_{184868}$ to illustrate a different construction of an augmentation. Comparing Figure 12 to Figure 15 of [16], we see that the set X constructed in the above proof is indeed different from that of Proposition 7.11 of that paper. Also, the graph realized by this 'new' X (in the sense of Definition 7.9 in [16]) is different from what we called the augmented graph of the underlying permutation of β there. In the example, these are both due to the fact that the position of the augmented crossing '3' has changed.

7 Examples

The following proposition is easy to prove either using skein relations of Homfly and/or Kauffman polynomials, or by a straightforward induction proof:

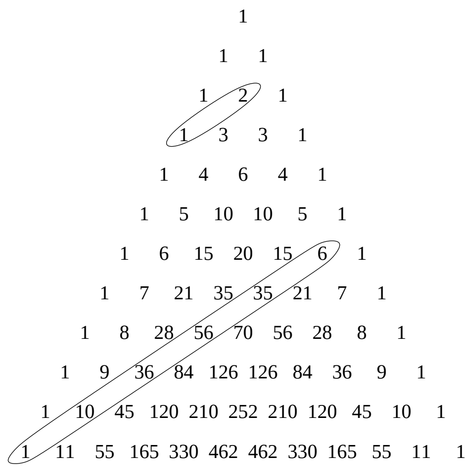

Figure 13: Ruling invariants of the $(3,2)$ and $(11,2)$ torus knots in Pascal's triangle.

Proposition 7.1. *The ruling polynomial of the $(p,2)$ torus link is $R(z) =$*

$$z^{p-1}+(p-1)z^{p-3}+\binom{p-2}{2}z^{p-5}+\binom{p-3}{3}z^{p-7}+\ldots+\binom{p-\lfloor p/2\rfloor}{\lfloor p/2\rfloor}z^{p-2\lfloor p/2\rfloor-1}.$$

The total number of rulings is $R(1) = f_p$, the p'th Fibonacci number. The total number of augmentations is $R(2^{-1/2})2^{(\chi^*+1)/2} = (2^{p+1} - (-1)^{p+1})/3$.

In particular, these ruling polynomials can be easily read off of Pascal's triangle, as shown in Figure 13. For example for $p = 11$, we get the ruling polynomial $R(z) = z^{10} + 10z^8 + 36z^6 + 56z^4 + 35z^2 + 6$. It seems likely that among Legendrian closures of positive braids with a given value of tb, the $(p,2)$ torus link with $p = tb+2$ has the least number of rulings for all values of θ. For $tb = 9$, the braid-positive knots with the largest number of rulings (for each θ) are the mutants $13n_{981}$ and $13n_{1104}$. These have $R(z) = z^{10} + 10z^8 + 36z^6 + 60z^4 + 47z^2 + 14$.

Mutant knots share the same Kauffman and Homfly polynomials, thus mutant braid-positive knots cannot be distinguished by their ruling polynomials. The braid-positive knots $12n_{679}$ and $13n_{1176}$ are not mutants yet they share the same ruling polynomial $R(z) = z^{10} + 10z^8 + 36z^6 + 58z^4 + 42z^2 + 11$ (their Kauffman and Homfly polynomials are actually different, but they agree in the coefficients that mean numbers of rulings).

Proposition 7.1 shows that for the $(p,2)$ torus link, roughly two thirds of the 2^p subsets of its crossings are augmentations. This ratio depends above all on the number of strands in the braid and goes down approximately by a factor of two every time the latter increases by one. When the number of strands is low, the ratio is quite significant[6]. This phenomenon seems to be unique to braid-positive links. (It may be worthwhile to compare to Chekanov's 5_2 diagrams, where out of the 64 subsets, only 3, respectively 2, are augmentations.)

[6] Thus the relatively complicated nature of the proof of Theorem 6.3 and of the construction in section 7 of [16] is somewhat misleading.

Example 7.2. The following were computed using a computer program written by Supap Kirtsaeng, based on Theorem 5.3. (Note that mere numbers of augmentations can also be determined from the Homfly or Kauffman polynomials using formula (9).) The braid word $(\sigma_1\sigma_2)^6$, corresponding to the $(3,6)$ torus link, yields 1597 augmentations (about 39% of all subsets of its crossings). The knot $12n_{679}$ (braid word $\sigma_1^3\sigma_2^2\sigma_1^2\sigma_2^5$) has 1653 augmentations (appr. 40%). $13n_{1176}$ also has 1653 augmentations, but its braid index is 4; for the braid word $\sigma_1\sigma_2^2\sigma_3\sigma_1^2\sigma_2^2\sigma_3^2\sigma_2\sigma_1\sigma_2$, the augmentations account for only 20% of all subsets of crossings. The knots $13n_{981}$ (closure of $\sigma_1\sigma_2^3\sigma_3\sigma_1\sigma_3\sigma_2^3\sigma_3^3$) and $13n_{1104}$ ($\sigma_1\sigma_2^2\sigma_3\sigma_1\sigma_3\sigma_1^2\sigma_2^3\sigma_3\sigma_1$) both have 1845 augmentations (i.e., 23% of all possibilities work). About the following two knots, Stoimenow [23] found that their braid index is 4, but in order to obtain them as closures of positive braids, we need 5 strands. $16n_{92582}$ (braid word $\sigma_1\sigma_2^2\sigma_3\sigma_4\sigma_3\sigma_1^2\sigma_2^2\sigma_3^2\sigma_2\sigma_4\sigma_3^2$) has 7269 augmentations, which is only about 11% of all possibilities. $16n_{29507}$ ($\sigma_1\sigma_2^2\sigma_3\sigma_1\sigma_3\sigma_4\sigma_1\sigma_2\sigma_4\sigma_2\sigma_3^3\sigma_4\sigma_2$) has 8109 (12%).

8 Concluding remarks

After devoting an entire paper to their study, let me attempt to justify why Legendrian closures of positive braids are important. I conjecture that L_β is essentially the only Legendrian representative of its link type, in the following sense.

Conjecture 8.1. *Any braid-positive Legendrian link is a stabilization of the corresponding Legendrian closure shown in Figure 1. In particular, braid-positive links are Legendrian simple.*

To support this conjecture, let us mention that Etnyre and Honda [9] proved it for positive torus knots, that the set of links treated by Ding and Geiges [7] includes many two-component braid-positive links (for example, positive $(2k,2)$ torus links). Chekanov's example [3] of a non-Legendrian simple knot type is 5_2, which is the smallest positive, but not braid-positive knot. Also, by Rutherford's work [21], the Thurston–Bennequin number of L_β is maximal in its smooth isotopy class because the front diagram f_β admits rulings, as we saw in section 6. Because some (actually, all) of those rulings are 2–graded, the maximal Thurston–Bennequin number for braid-positive links is only attained along with rotation number 0.

Finally, the computations in this paper can also be viewed as first steps in a program to calculate contact homology invariants of all cabled knot types and to relate them to the companion knot (which in the case of braid-positive links is the unknot).

References

[1] M. Aigner and G. M. Ziegler: *Proofs from THE BOOK*, Springer-Verlag, 2004.

[2] B. Sturmfels: *What is a Gröbner basis?*, Notices of the AMS 52 (2005), no. 10, 1199–1200.

[3] Y. Chekanov: *Differential algebra of Legendrian links*, Invent. Math. 150 (2002), 441–483.

[4] Y. Chekanov: *New invariants of Legendrian knots*, Third European Congress of Mathematics, vol. II (Barcelona, 2000), 525–534.

[5] Y. Chekanov and P. Pushkar: *Combinatorics of fronts of Legendrian links, and Arnold's 4-conjectures*, Russian Math. Surveys 60 (2005), no. 1, 95–149.

[6] P. Cromwell: *Knots and Links*, Cambridge University Press, 2004.

[7] F. Ding and H. Geiges, *Legendrian knots and links classified by classical invariants*, Commun. Contemp. Math., to appear.

[8] J. Etnyre: *Legendrian and transversal knots*, in Handbook of Knot Theory, Elsevier B. V., Amsterdam, 2005, 105–185.

[9] J. Etnyre and K. Honda: *Knots and contact geometry I: torus knots and the figure eight knot*, J. Symplectic Geom. 1 (2001), no. 1, 63–120.

[10] J. Etnyre, L. Ng, and J. Sabloff: *Invariants of Legendrian Knots and Coherent Orientations*, J. Symplectic Geom. 1 (2002), no. 1, 321–367.

[11] D. Fuchs: *Chekanov–Eliashberg invariants of Legendrian knots: existence of augmentations*, J. Geom. Phys. 47 (2003), no. 1, 43–65.

[12] D. Fuchs and T. Ishkhanov: *Invariants of Legendrian knots and decompositions of front diagrams*, Mosc. Math. J. 4 (2004), no. 3, 707–717.

[13] D. Fuchs and S. Tabachnikov, *Invariants of Legendrian and transverse knots in the standard contact space*, Topology 36 (1997), no. 5, 1025–1053.

[14] F. A. Garside: *The braid group and other groups*, Quart. J. Math. Oxford (2) 78 (1969) 235–254.

[15] I. M. Gessel and G. Viennot: *Binomial determinants, paths, and hook length formulae*, Advances in Math. 58 (1985), 300–321.

[16] T. Kálmán: *Contact homology and one parameter families of Legendrian knots*, Geom. Topol. 9 (2005), 2013–2078.

[17] T. Kálmán, *Maximal Thurston–Bennequin number of +adequate links*, in preparation.

[18] B. Lindström: *On the vector representation of induced matroids*, Bull. London Math. Soc. 5 (1973), 85–90.

[19] L. Ng: *Computable Legendrian invariants*, Topology 42 (2003), no. 1, 55–82.

[20] L. Ng and J. Sabloff: *The correspondence between augmentations and rulings for Legendrian knots*, math.SG/0503168.

[21] D. Rutherford: *The Bennequin number, Kauffman polynomial, and ruling invariants of a Legendrian link: the Fuchs conjecture and beyond*, Int. Math. Res. Not. 2006, Art. ID 78591.

[22] J. Sabloff: *Augmentations and rulings of Legendrian knots*, Int. Math. Res. Not. 2005, no. 19, 1157–1180.

[23] A. Stoimenow: *On the crossing number of positive knots and braids and braid index criteria of Jones and Morton–Williams–Franks*, Trans. Amer. Math. Soc. 354 (2002), 3927–3954.

[24] T. Tanaka: *Maximal Bennequin numbers and Kauffman polynomials of positive links*, Proc. Amer. Math. Soc. 127 (1999), 3427–3432.

[25] Y. Yokota: *Polynomial invariants of positive links*, Topology 31 (1992), no. 4, 805–811.