Eötvös Loránd University
Institute of Mathematics

The multivariable Alexander polynomial and Thurston norm

Master's Thesis

Author: *Daniel Copeland*

Advisor: *András Stipsicz*

Institutional Consultant: *András Szűcs*

Department of Analysis

Budapest, 2014

Contents

List of Figures

0.1 Introduction

This thesis is concerned with generalizing two classical knot invariants to links, namely the Alexnader polynomial and the knot genus. Their extensions are the multivariable Alexander polynomial and the Thurston norm. We assume the reader is acquainted with the basic theory of knots and links, e.g. the existence of link diagrams, the notion of oriented link equivalence, and Reidemeister's theorem about Reidemeister moves. We only consider tame knots and links, that is smoothly or PL embedded circles in S^3.

The single variable Alexander polynomial is an integral Laurent polynomial that is easily calculable from a knot or link diagram. It is one of the earliest algebraic knot invariants, introduced by Alexander in 1928 [Ale28], in which he used the combinatorial data of a knot diagram to extract a matrix whose determinant is a knot invariant, the Alexander polynomial. This polynomial was quickly seen to satsify a number of properties: for example, it is symmetric, its degree bounds the knot genus, and it is monic when the knot is fibered.

A number of methods of calculating the Alexander polynomial were developed shortly after Alexander's introduction. For example, it is calculable by Seifert forms on the first homology of Seifert surfaces or from a presentation of the knot group via the Fox calculus. When extracted from a knot, the Alexander polynomial can interpreted as an invariant of the first homology of the maximal free abelian cover of the link complement, which in this case is infinite cyclic. Finally, one of the most practical tools for computation is that the Alexander polynomial satisfies a skein relation, and hence can be computed inductively in a simple manner from a knot or link diagram.

Around the 1950s, Fox and others studied a refinement of the Alexander polynomial for links, called the multivariable Alexander polynomial, which is a Laurent polynomial in l variables where l is the number of components of the link. It is perhaps most naturally understood as the generalization of the Alexander polynomial using the computation via maximal free abelian covers mentioned above, and this is the approach we describe in Chapter 2. In any case, each variable corresponds to a generator of the first homology of the link complement, which is freely generated by oriented meridians wrapped around the boundary tori corresponding to each component of the link. One of the primary difficulties of working with this new polynomial is that there is no simple skein rule: indeed, application of the skein rule may increase or reduce the number of components, so there is no canonical identification of variables in each portion of the skein step. However, Murakami has provided a system of local axioms for the multivariable Alexander polynomial analagous to the skein rule [Mur93], but they are considerably more complicated. This view is related to the Conway function invariant, but we shall not discuss it here. In the absence of the skein rule, we do have a different practical tool, which is the Fox calculus. It allows one to compute the multivariable Alexander polynomial from a presentation of the knot group along with the information of the orientation of the link. We discuss this in Chapter 2.

The generalization of the knot genus to links with many components is quite interesting. The main issue is that one may study the complexity of surfaces which bound any number of the components of the link. The key observation is that these surfaces are parametrized by the second homology of the link complement, and the "link genus" should be a function on this entire space. In the 1980s William Thurston pursued this line of thought (actually for arbitrary compact 3-manifolds) and using the general idea that Euler characteristic-maximizing surfaces minimize the topological complexity of a homology class, introduced the Thurston norm, a function on the real second homology of the link complement that in general is a semi-norm defined by a (possibly non-compact) polyhedral unit ball.

McMullen then extended the result that the degree of the Alexander polynomial bounds the knot genus by showing that the Alexander norm on the first cohomology of the link complement, derived from the multivariable Alexander polynomial, bounds the Thurston norm (where the Thurston norm is transferred to cohomology via Poincaré duality). A special case of equality occurs for fibered classes, which are cohomology classes representable by a fibration from the link complement to S^1.

The structure of this thesis is as follows: in Chapter 1 we study a particular presentation of the link group obtained from the data of a grid diagram, and the homology of the link complement in preparation for the later chapters. Chapter 2 defines the multivariable Alexander polynomial and describes how the Fox calculus is used for computation. In Chapter 3 we offer another definition of the multivariable Alexander polynomial obtained from the combinatorial data of a grid diagram, and use this version to prove the symmetry of the polynomial. Chapter 4 addresses the Thurston norm, McMullen's result, and the development of the theory of the Thurston norm and fibered classes. Finally, the Epilogue presents a small view of modern link invariants that are currently active research subjects and whose origins lie in the topics discussed in this thesis.

The content of the thesis is by no means original, but we hope that it may be a guide or supplement for interested students. We have added considerably more explanation to the theorems of Thurston and McMullen (Chapter 4) than is included in the original expositions. Some theorems were hard to find in the literature so we take special care to include complete proofs (e.g. Theorem 1.7 and Theorem 2.20).

Chapter 1

Our basic tool set

1.1 Grid diagrams and the Neuwirth presentation

Definition 1.1. An $n \times n$ **grid diagram** is an $n \times n$ diagram of squares filled with X's and O's such that no square contains both an X and an O, and each column or row contains a single X and a single O.

The squares of a grid diagram are given matrix coordinates, so the (i, j)th grid square sits in the ith row from the top and the jth column from the left.

With every grid diagram G there is an associated oriented link, denoted $L(G)$. This is obtained by connecting the X's and O's of each row and column with vertical and horizontal line segments. This produces a link projection. To designate crossings and orientation, we require all vertical arcs to cross over horizontal arcs. The orientation is determined by requiring that all the vertical arcs are oriented towards the (unique) O marking in the column (see Figure 1.1 for a grid diagram of the $(4, 2)$ torus link).

Conversely, for every link there is a grid diagram representing it. This is easy to see by starting with a link projection, approximating it with a piecewise linear projection consisting only of horizontal and vertical segments, and then adjusting crossings locally so that all over-strands are vertical segments.

There is an analogue of Reidemeister's theorem, which states that two grids represent the same link if and only if they are related by a sequence of elementary modifications

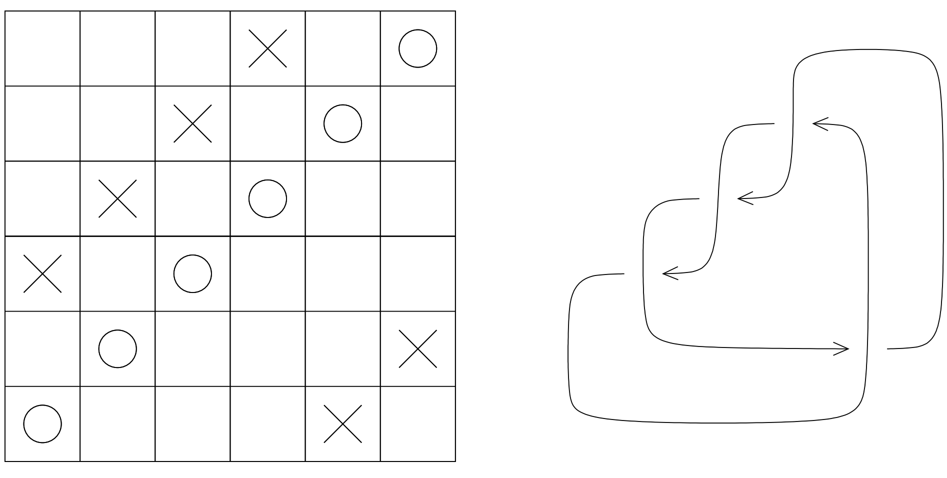

Figure 1.1: Grid diagram for the $(4, 2)$ torus link.

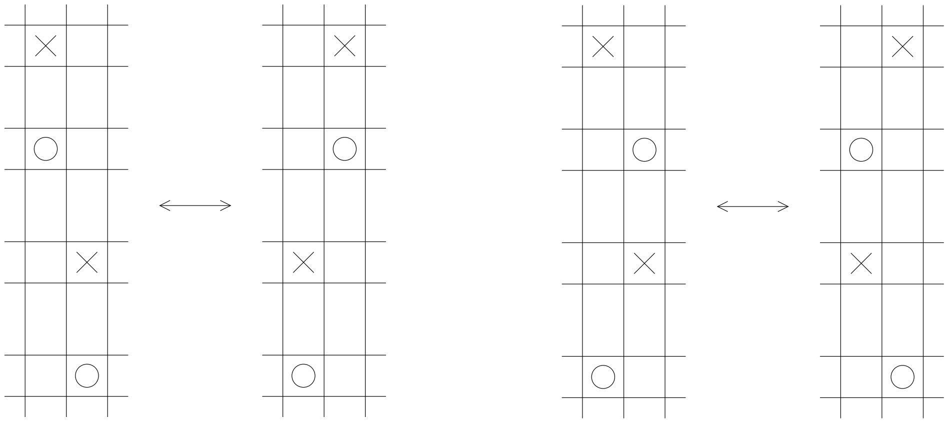

Figure 1.2: The two types of commutation moves.

to the grid diagram. These elementary moves are commutations and stabilizations:

Definition 1.2. Suppose G is a grid diagram with two consecutive columns such the arcs of the link in each column, when projected to the vertical grid line between the columns, are either disjoint or one strand is properly contained in the other. Then a grid G' is obtained by a **commutation move** by swapping the X and O markings in these two columns.

See Figure 1.2 for examples where a commutation movie is applicable. Note that the conditions for a commutation to be applied are not satisfied if the columns have any X or O markings in the same row. It is also clear that the type of the link is not affected by a commutation, since case analysis shows that this move is equivalent to an isotopy of the link or an application of the Reidemeister II move.

Definition 1.3. Suppose G is an $n \times n$ grid diagram. The $(n+1) \times (n+1)$ grid diagram G' is a **stabilization** of G if G' is obtained from G in the following way. Pick a distinguished X or O marking of G. Remove the other X and O markings in the same row and column as the distinguished marking, and replace this row and column by two rows and two columns, so there are now two empty rows and two empty columns whose intersection is 4 grid squares. There are now four ways of inserting markings into this grid to make a new grid diagram, indexed by which of the 4 grid squares does not have a marking. Each of these grid diagrams results in a stabilization of G.

The inverse of a stabilization move is a **destabilization**.

A stabilization is usually described by which marking it modifies and what direction the empty square is. For example, Figure 1.3 shows the local modification of an $X : NW$ stabilization.

Definition 1.4. A **cyclic permutation** of a grid diagram consists of moving one of the rows or columns on the edge of the diagram to the opposite row or column.

This corresponds to simply viewing the grid diagram on the torus in a different manner, and clearly does not change the type of the link. It is an easy exercise to show:

Lemma 1.5. *A cyclic permutation is the result of finitely many commutation, stabilization and destabilization moves.* □

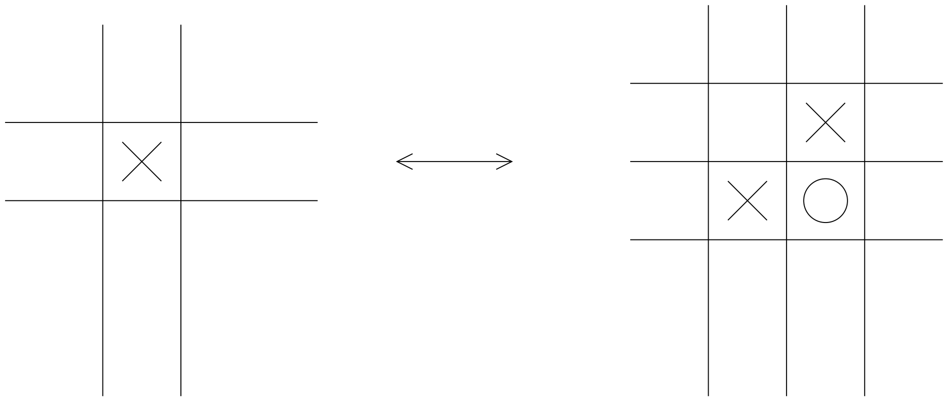

Figure 1.3: An $X : NW$ stabilization move.

Then we have:

Theorem 1.6. *(Cromwell)[Cro95], [OSS] Two grid diagrams represent the same link if and only if there exists a finite sequence of commutation and (de)stabilization moves relating the diagrams.*

Therefore, one way to present a link invariant is to derive an object from the data of a grid diagram and prove it is invariant under the elementary moves. We will do so with the multivariable Alexander polynomial in Chapter 3.

We now turn to the Neuwirth presentation of the link group, given by the combinatorial data of the grid diagram. There are n generators, one for each column of the grid diagram corresponding to a loop around the vertical strand in each given column. There is a relation for each horizontal grid line which is the product of those generators whose corresponding vertical strand crosses the grid line. It is clear that these are valid relations in the link group, since the loop corresponding to each relation may be homotoped to a point by pulling it "behind" the grid diagram. Figure 1.4 shows the setup for the grid diagram of the $(4, 2)$ torus link given above. The following theorem states that no other relations are needed to describe the link group.

Theorem 1.7. *Let G be an $n \times n$ grid diagram of a link L. Then $\pi_1(S^3 - L) \cong \langle x_1, \ldots, x_n | r_1, \ldots, r_{n-1} \rangle$ where r_i is the product of those x_j's for which the vertical strand in the j^{th} column intersects the i^{th} horizontal grid line.*

Proof. We will apply the Seifert-van Kampen theorem to a suitable open cover of the link complement. First, we construct a particular embedding of the link in $\mathbb{R}^3$ given by a grid diagram. The vertical strands are placed on the plane $z = 1$ parallel to the y-axis, and the horizontal strands are placed on the plane $z = 0$ parallel to the x-axis such that the projection of these strands onto the xy-plane give us the link projection associated with the diagram G. To create an embedding of the link, we connect the endpoints of horizontal and vertical strands using segments parallel to the z-axis. Clearly there is one such segment for every appearance of an X or O in the grid diagram. For example, Figure 1.5 shows the embedding of the $(4, 2)$ torus link using the grid diagram from before.

Now consider the regions $A = \{z > 0\} - L$, $B = \{z < 1\} - L$ and $C = A \cap B$. Fix a basepoint $p \in \{z = 1/2\} \subset C$ that is far away from the embedded link, in the positive quadrant of the plane. Then $\mathbb{R}^3 - L = A \cup B$, and since A, B, C are all path-connected we may use the Seifert-van Kampen theorem to calculate $\pi_1(\mathbb{R}^3 - L, p)$. First,

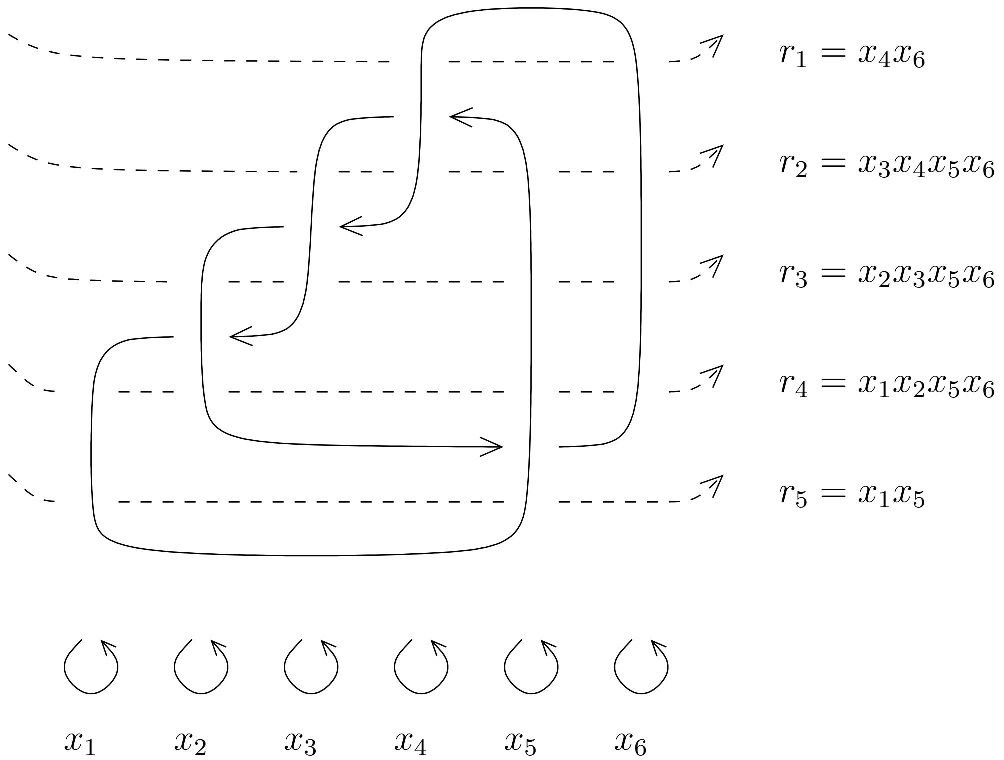

Figure 1.4: A Neuwirth presentation for the $(4, 2)$ torus link.

consider A. It consists of a half space missing some unknotted "arches". It is not hard to see this is homotopy equivalent to a bouquet of n circles. Hence its fundamental group is freely generated by n elements, and these generators, denoted $x_1, \ldots, x_n$ are represented by loops which wind neatly around each vertical strand of the link, with the orientations suggested by Figure 1.4, i.e. counterclockwise in the plane. Similarly, $\pi_1(B, p)$ is freely generated by elements $y_1, \ldots, y_n$ which have representatives that wind neatly c.c.w. around each horizontal strand of the link.

Consider now the region C. It deformation retracts onto a plane minus $2n$ points (in correspondence with the various X's and O's in the grid diagram). Therefore it is freely generated by $2n$ elements, corresponding to loops that wind neatly around each missing segment. However, we shall consider different generators, indexed by the grid lines of G. Let r_i be the class of a loop that wraps around all of the segments connecting X's and O's in the first i horizontal grid lines. Now let s_j be represented by the loop that winds neatly around the segment corresponding to the X in the j^{th} row, approaching between the $j-1$ and j^{th} rows. Some of these generators for the example of the torus link are shown in Figure 1.6 (where C has been projected onto the xy-plane). These loops indeed generate all of $\pi_1(C, p)$, since all of the segments corresponding to X's already have loops aroud them, and a loop around the a O-segment in the j^{th} column may be expressed as the product $r_j r_{j-1}^{-1} s_j^{-1}$, by our choice of how the loops s_j approach the segment that they are wrapped tightly around.

We are ready to apply Seifert-van Kampen. Since there are no relations in the fundamental groups of A and B, we have that

$$\pi_1(\mathbb{R}^3 - L) = \left\langle x_1, \ldots x_n, y_1, \ldots y_n \;\middle|\; \begin{array}{ll} (i_A)_*(r_i) = (i_B)_*(r_i) & 1 \leq i \leq n \\ (i_A)_*(c_j) = (i_B)_*(s_j) & 1 \leq j \leq n \end{array} \right\rangle$$

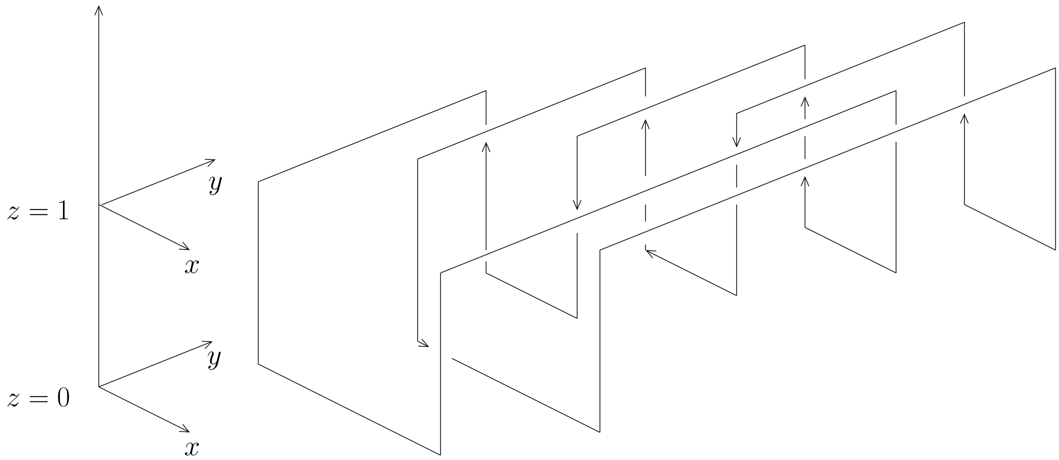

Figure 1.5: Embedding of the $(4, 2)$ torus link from its grid diagram.

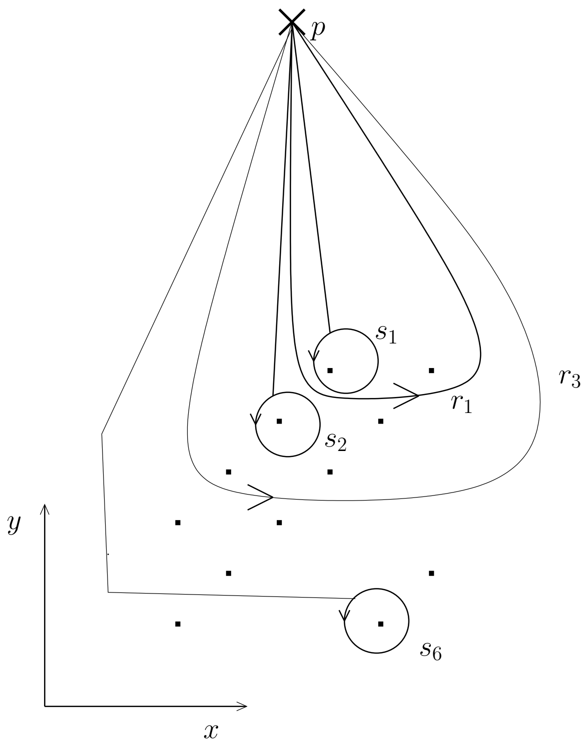

Figure 1.6: Choice of generators for $\pi_1(C, p)$.

where i_A and i_B are the inclusion maps of C into A and B. Clearly $(i_A)_*(s_j) = x_j$ and $(i_B)_*(s_j) = y_j$, so we need only consider the generators x_i and the relations r_j. But $(i_B)_*(r_j) = 1$ since we may pull each loop underneath all the missing strands in B and contract it to a point. On the other hand, $(i_A)_*(r_j)$ is represented by a loop that goes through all the arches corresponding to the missing strands that cross the j^{th} horizontal grid line. Therefore, it may be written as a product of the corresponding generators of A. Finally, $(i_A)_*(r_n) = 1$ since it may be pulled above the arches and contracted to a point. The relations we are left with are exactly those described in the statement of the theorem, so we are done.

□

1.2 Homology of the link complement

In this section we compute the homology and cohomology of the link complement. Let $L \subset S^3$ be an l component link, $X = S^3 - L$ the unbounded link complement and $M = S^3 - \mathring{\nu}(L)$ the (bounded) link complement where $\mathring{\nu}(L)$ denotes a small open tubular neighborhood of L. Note that X is an open 3-manifold and M is a compact 3-manifold whose boundary is a disjoint union of tori. Since M is a deformation retract of X their homologies are the same. However the relative groups $H_*(M, \partial M)$ are different. Whenever coefficients are omitted in the notation, we are using $\mathbb{Z}$.

Proposition 1.8. *The homology of X is as follows: $H_0(X) = \mathbb{Z}$, $H_1(X) = \mathbb{Z}^l$ $H_2(X) = \mathbb{Z}^{l-1}$ and the remining homologies vanish.*

Proof. We prove everything except that $b_2(X) = l - 1$ where b_2 is the second Betti number of X (this is proved below). The neighborhood $\mathring{\nu}(L)$ is a disjoint union of open solid tori, each with the homology of a circle, so

$$H_0(\mathring{\nu}(L)) = \mathbb{Z}^l$$
$$H_1(\mathring{\nu}(L)) = \mathbb{Z}^l$$

and the remaining groups are trivial. Also, the space $X \cap \mathring{\nu}(L)$ deformation retracts onto a disjoint union of l tori, so its nontrivial homology groups are:

$$H_0(X \cap \mathring{\nu}(L)) = \mathbb{Z}^l$$
$$H_1(X \cap \mathring{\nu}(L)) = (\mathbb{Z} \oplus \mathbb{Z})^l$$
$$H_2(X \cap \mathring{\nu}(L)) = \mathbb{Z}^l$$

We apply the Mayer-Vietoris sequence for $S^3 = \mathring{\nu}(L) \cup X$:

$$\cdots \to H_{d+1}(S^3) \to H_d(X \cap \mathring{\nu}(L)) \to H_d(X) \oplus H_d(\mathring{\nu}(L)) \to H_d(S^3) \to H_{d-1}(X \cap \mathring{\nu}(L)) \to \ldots$$

Note that $H_3(X) = 0$ since X is not compact, and the sequence provides the rest. □

Corollary 1.9. $H^i(X) \cong \mathrm{Hom}(H_i(X); \mathbb{Z}) \cong H_i(X)$ *for each* $i \in \mathbb{N}$.

Proof. Apply the universal coefficients theorem and note none of the homology groups have torsion. □

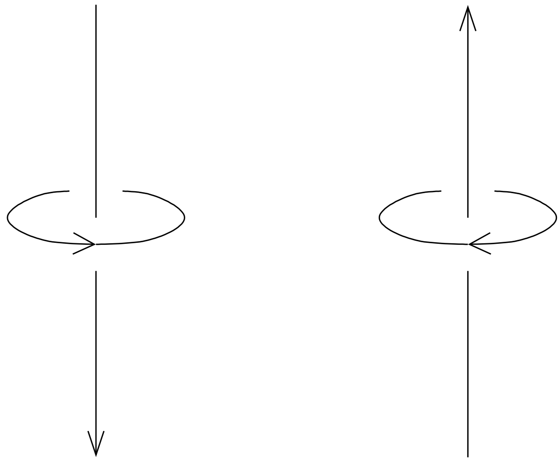

Figure 1.7: Orientation convention for the basis elements of $H_1(X - L)$.

We can find a simple basis for $H_1(X) \cong \mathbb{Z}^l$: for each link component L_i, consider a small loop around it and pick an orientation by the convention in Figure 1.7. The homology class of this loop is called the ***i*th meridian** and will be denoted μ_i. We claim these form a basis for $H_1(X)$. Indeed, consider the portion of the Mayer-Vietoris sequence above:

$$0 \to H_1(X \cap \mathring{\nu}(L)) \xrightarrow{\phi} H_1(X) \oplus H_1(\mathring{\nu}(L)) \to 0$$

where the map ϕ takes a class α to $(\alpha, -\alpha)$. It is clear that $H_1(X \cap \mathring{\nu}(L)$ is generated by $\{\mu_1, \ldots, \mu_l, l_1, \ldots, l_l\}$ where l_i is the ***i*th longitude**, i.e. a generator of $H_1(\mathring{\nu}(L_i))$. Clearly each μ_i is contractible in $\mathring{\nu}(L_i)$. Therefore, ϕ maps the subgroup spanned by the meridians isomorphically onto $H_1(X)$.

The orientation convention for this basis is a "left-hand rule": point your left-hand thumb in the direction of the link, and the curl of your other fingers determines the positive orientation of the meridian loop around that link (see Figure 1.7). We denote the generator corresponding to the i^{th} component by t_i, so $H_1(X - L) \cong \mathbb{Z}^l \cong \mathbb{Z}\langle t_1, \ldots, t_l \rangle$.

Next let's consider the relative homology $(M, \partial M)$. Writing out the LES of the pair $(M, \partial M)$, we have

$$\begin{array}{ccccccc}
 & & & 0 & \longrightarrow & H_3(M,\partial M) & \longrightarrow \\
\longrightarrow & H_2(\partial M) & \longrightarrow & H_2(M) & \longrightarrow & H_2(M,\partial M) & \longrightarrow \\
\longrightarrow & H_1(\partial M) & \longrightarrow & H_1(M) & \longrightarrow & H_1(M,\partial M) & \longrightarrow \\
\longrightarrow & \tilde{H}_0(\partial M) & \longrightarrow & 0 & & &
\end{array}$$

By Poincaré duality, $H_i(M, \partial M) \cong H^i(M)$ and the latter group is isomorphic to $H_i(M)$. Hence the only unknown groups in the sequence above are $H_2(M) \cong H_1(M, \partial M)$. Note that these must be free since every other group in the sequence is free. The only unknown is the value $b_2 = \text{rank}(H_2(M))$. But since the sequence is exact, the alternating sum of the ranks must be 0, so adding from top to bottom (and using the fact that ∂M is a disjoint union of l tori) we get:

$$0 = 1 - l + b_2 - l + 2l - l + b_2 - (l - 1)$$

so $b_2 = l - 1$, as promised in the previous proposition. We've proved:

Proposition 1.10. *The relative homology of the bounded link complement* $(M, \partial M)$ *is* $H_1(M, \partial M) = \mathbb{Z}^{l-1}$, $H_2(M, \partial M) = \mathbb{Z}^l$, $H_3(M, \partial M) = \mathbb{Z}$ *and the remaining homology groups vanish.* □

1.3 De Rham cohomology and Poincaré duality in the link complement

We review the connection between singular and de Rham cohomology in our low dimensional case without details; the reader should consult [BT82] for a full treatment. Consider now the bounded link complement $(M, \partial M)$. Poincaré duality gives an isomorphism $H_i(M, \partial M; \mathbb{Z}) \cong H^{3-i}(M)$ where the Poincaré dual $D(\alpha)$ of $\alpha \in H_i(M, \partial M; \mathbb{Z})$ is characterized by

$$\phi(\alpha) = (\phi \smile D(\alpha))([M])$$

for all $\phi \in H^i(M, \partial M; \mathbb{Z})$ considered as a function on $H_i(M, \partial M : \mathbb{Z})$. $[M]$ denotes the fundamental class of M and $\smile$ denotes the cup product on cohomology.

Now we consider the de Rham cohomology groups, denoted $H^i_{DR}(M; \mathbb{R})$. Given any (singular) cohomology class $\alpha \in H^i(M; \mathbb{Z})$ we can assign a closed i-cycle ω characterized by

$$\int_\sigma \omega = \alpha(\sigma)$$

for every $(n - i)$-chain σ.

The de Rham theorem states that this correspondence is a ring isomorphism between $H^i(M; \mathbb{R})$ equipped with the cup product and $H^i_{DR}(M; \mathbb{R})$ equipped with the wedge product. Now since none of the homology groups with integral coefficients have torsion, we have an inclusion

$$H^i(M; \mathbb{Z}) \hookrightarrow H^i_{DR}(M; \mathbb{R})$$

which assigns to every integral cohomology class a unique de Rham cohomology class.

We focus our attention now on the first cohomology of M. Suppose $\phi \in H^1(M; \mathbb{Z})$. We've seen that the oriented meridians around the boundary tori form a basis of $H_1(M)$ (which we denote $t_1, \ldots, t_l$) and these give us a dual basis $\bar{t_1}, \ldots, \bar{t_l}$ of $H^1(M; \mathbb{Z})$. Therefore we may write $\phi = \sum_{i=1}^{l} \lambda_i \bar{t_i}$ for some integer coefficients λ_i.

Now suppose S is a properly embedded compact oriented surface so its fundamental class is an element of $H_2(M, \partial M; \mathbb{Z})$. Using de Rham cohomology, the Poincaré dual $[\eta_S]$ to the fundamental class of S is the unique class of a closed 1-form η_S which satisfies

$$\int_S i^* \omega = \int_M \omega \wedge \eta_S$$

for all closed 2-forms ω. Since we are working in low-dimensions, we can easily describe a 1-form representing $[\eta_S]$. Let $\nu(S)$ be a smoothly embedded tubular neighborhood of S, so $\nu(S) = S \times (-1, 1) = \{(s, r) : s \in S, r \in (-1, 1)\}$. Let $f(r)$ be a smooth bump function with integral 1 on the interval $(-1, 1)$. We define the 1-form η_S by

$$\eta_S(p) = \begin{cases} f(r)dr & \text{if } p \in \nu(S) \\ 0 & \text{if } p \notin \nu(S) \end{cases}$$

Note that the fundamental class of $[S]$ is an integral homology class, so its Poincaré dual is evidently also an integral class. Thus we would like to identify $[\eta_S]$ as an element of $H^1(M;\mathbb{Z})$ by trying to evaluate $[\eta_S](\sigma)$ where $\sigma \in H_1(M;\mathbb{Z})$. It suffices to consider the case when σ is a meridian t_i, as these form a basis of $H_1(M;\mathbb{Z})$. By following the definitions, we see

$$[\eta_S](t_i) = \int_{t_i} \eta_S = \text{algebraic intersection number of } t_i \text{ and } \partial S$$

where we calculated the integral by assuming t_i intersects $\nu(S)$ as a 'vertical fiber', i.e. a set of the form $\{(s_0, r) : r \in (-1,1)\}$, so the integral of η_S adds ± 1 to the total for every component in the intersection. This assumption is not problematic since we may always homotope t_i within $\nu(S)$ to have such a form. Putting everything together, we have the following propositions:

Proposition 1.11. *Suppose S is an embedded surface Poincaré dual to $\phi \in H^1(M)$. Then for any 1-cycle γ in $H_i(M)$ we have*

$$\phi(\gamma) = \textit{algebraic intersection number of } \gamma \textit{ and } \partial S.$$

Proposition 1.12. *The fundamental class of $[S]$ is Poincaré dual to $\phi = \sum \lambda_i \bar{t_i}$ iff the boundary of S wraps longitudinally around the i^{th} boundary torus λ_i times (where λ_i is positive or negative considering orientations).* □

Chapter 2

The classical multivariable Alexander polynomial

2.1 Algebraic preliminaries

Let R be a UFD and M a finitely generated R-module. A a presentation of M with n generators and r relations (where possibly $|r| = \infty$) is given by an exact sequence:

$$R^r \xrightarrow{A} R^n \to M \to 0$$

Note that A may be written as an $n \times r$ matrix with coefficients in R. The above sequence means M is the cokernel of the homomorphism A.

Definition 2.1. The k**th elementary ideal of** M, denoted $E_k(M)$, is the ideal of R generated by the $n - k \times n - k$ subdeterminants of A if $0 < n - k \leq r$.

If $n - k \leq 0$ then we define $E_k(M) = R$ and if $n - k > r$, we define $E_k(M) = 0$.

It appears as though the definition depends on the presentation of M. However, it does not, so the elementary ideals are indeed module invariants.

Proposition 2.2. *The kth elementary ideals are all invariant under a change of presentations for M.*

Proof. [CF63], [Zas49]. The proof uses elementary linear algebra and Tietze moves. □

Since any determinant can be written as a linear combination of subdeterminants, we have

$$0 = E_0(M) \subset E_1(M) \subset \cdots \subset E_n(M) = R$$

Definition 2.3. The k**th order of the module** M, denoted $\Delta_k(M)$ is the gcd of $E_k(M)$ (in $R^\times$). The **order of** M, denoted $\mathrm{ord}(M)$, is the 0th order of M.

Note that the $\Delta_k(M)$ are defined only up to multiplication by units of R, so we introduce the notation

$$\Delta \doteq \Delta'$$

to mean equality up to multiplication by a unit (or in other terms, Δ and Δ' are associates).

Example 2.4. The trivial module has the identity matrix as a presentation matrix, so $\Delta_k(0) \doteq 1$ for all k. A free module R^n is presented by $0 \to R^n \to R^n \to 0$ so a presentation matrix is the zero matrix. Hence $\Delta_k(R^n) = 0$ for all k.

The sequence of inclusions above give the divisibility relations:

$$1 = \Delta_n(M) \mid \Delta_{n-1}(M) \mid \ \dots \ \mid \Delta_0(M) = \mathrm{ord}(M)$$

For PIDs the order of a finitely generated module is particularly simple:

Example 2.5. Suppose M is a finitely generated module over a PID R. Then M has a decomposition as a direct sum of cyclic modules:

$$M = R/(p_1) \oplus \cdots \oplus R/(p_n)$$

Therefore it has a presentation $R^n \xrightarrow{A} R^n \to M$ where A is an $n \times n$ diagonal matrix with entries $p_1, \dots, p_n$. Taking the determinant, we see $\mathrm{ord}(M) = p_1 \cdots p_n$. Note that $\mathrm{ord}(M) = 0 \iff \mathrm{rank}(M) \neq 0$, which also holds for more general rings, as we shall see.

Later we use modules over the ring of Laurent polynomials in l variables with integer coefficients, denoted Λ_l. This is isomorphic to the group ring of $\mathbb{Z}^l$. We record that Λ_l is a Noetherian UFD, and in fact many of the relevant properties of the orders are deduced at this level of generality. Recall that for a module M over a domain R, the **torsion part** of M is the submodule $\{m \in M \mid \exists r \in R : rm = 0\}$. Equivalently, it is the kernel of the map $M \xrightarrow{Q(R)\otimes -} Q(R) \otimes M$ where $Q(R)$ is the field of fractions of R. A module is **torsion free** if its torsion part is trivial.

Proposition 2.6. *Let Λ be a Noetherian UFD and suppose*

$$0 \to H_1 \to H \to H_2 \to 0$$

is a short exact sequence of Λ-modules. Then

1. $\Delta_0(H) = \Delta_0(H_1)\Delta_0(H_2)$
2. *If H_2 has no torsion part and rank $H_2 = r$, then*

$$\Delta_d(H) = \begin{cases} \Delta_{d-r}(H_1) & \text{if } r \leq d \\ 0 & \text{if } r > d \end{cases}$$

Lemma 2.7. *Suppose H_1 and H_2 have square presentation matrices P_1 and P_2 respectively. Then H has a presentation matrix of the form*

$$\begin{pmatrix} P_1 & * \\ 0 & P_2 \end{pmatrix}$$

where $*$ represents some $n \times m$ matrix, and $* = 0$ iff $H = H_1 \oplus H_2$.

Proof. Suppose P_1 and P_2 correspond to presentations $\langle x_1, \dots, x_n | r_1, \dots, r_n\rangle$ and $\langle y_1, \dots, y_m | s_1, \dots, s_m\rangle$ respectively. Then H is generated by $\{x_1, \dots, x_n, y_1, \dots, y_m\}$. The relations $\{r_1, \dots, r_n\}$ determine the submodule H_1. The generators $\{y_1, \dots, y_m\}$ satisfy the relations $\{s_1, \dots, s_m\}$ only modulo H_1, which accounts for the $*$ in the presentation matrix for H. There are no other relations, since if we are given a relation between the generators $x_1, \dots x_n, y_1, \dots y_m$, then we can use the matrix $\begin{pmatrix} * \\ P_2 \end{pmatrix}$ to reduce it to a relation between the generators in H_1. The final assertion is clear. □

We prove Proposition 2.6 by using localization at prime ideals to compare factorizations of various orders. Briefly we recall the definition of localization. If $\mathfrak{p} \lhd R$ is a prime ideal of the commutative ring R, then the **localization of R at $\mathfrak{p}$**, denoted $R_{\mathfrak{p}}$, is the ring of fractions $\{\frac{a}{t} \mid a \in R, t \notin \mathfrak{p}\}$ with the obvious multiplication. More generally, if M is an R-module, then $M_{\mathfrak{p}}$ is the $R_{\mathfrak{p}}$-module $M \otimes R_{\mathfrak{p}}$. The elements of $M_{\mathfrak{p}}$ may be written as fractions $\{\frac{m}{t} \mid t \notin \mathfrak{p}\}$. It is easily shown that $- \otimes R_{\mathfrak{p}}$ is an exact functor. See [AM69] for details.

Returning to the case where Λ is a Noetherian UFD, let $p \in \Lambda$ be an irreducible element, so (p) is a prime ideal, and $\Lambda_{(p)}$ the corresponding localization. Note that $\Lambda_{(p)}$ is a PID, since (up to multiplication by units) every element is of the form p^k for some non-negative integer k. We can use localization to compare orders due to the following observation:

Lemma 2.8. *Let M be a Λ-module. Then we may write $\Delta_k(M) \doteq p^a q$ where $p \nmid q$ iff $\Delta_k(M \otimes \Lambda_{(p)}) \doteq p^a$ where $M \otimes \Lambda_{(p)}$ is considered a $\Lambda_{(p)}$-module.*

Proof. If P is a presentation matrix for M, then $P \otimes \Lambda_{(p)}$ is a presentation matrix for $M \otimes \Lambda_{(p)}$ as a $\Lambda_{(p)}$-module. This is a matrix whose entries are the images of the entries of P under the canonical homomorphism $\Lambda \to \Lambda_{(p)}$. The gcds of the various minors are preserved under this homomorphism, which proves the claim. □

Proof of 2.6. For (1), see [Kaw90]. For (2), note that since H_2 is torsion free with rank r, we have $H_2 \otimes \Lambda_{(p)} = \Lambda^r_{(p)}$ (since it is a torsion free module over the PID $\Lambda_{(p)}$, hence free, and tensoring by $Q(\Lambda)$ we must get $H_2 \otimes Q(\Lambda) = Q(\Lambda)^r$).

Applying the exact functor $- \otimes \Lambda_{(p)}$ to our original sequence gives us the short exact sequence of $\Lambda_{(p)}$-modules

$$0 \to H_1 \otimes \Lambda_{(p)} \to H \otimes \Lambda_{(p)} \to \Lambda^r_{(p)} \to 0$$

Since $\Lambda^r_{(p)}$ is free, this sequence splits and we arrive at

$$H \otimes \Lambda_{(p)} = (H_1 \otimes \Lambda_{(p)}) \oplus \Lambda^r_{(p)}$$

Now $H_1 \otimes \Lambda_{(p)}$ is a finitely generated module over a PID so has an $n \times n$ presentation matrix P. A presentation matrix for $\Lambda^r_{(p)}$ is given by the $r \times r$ zero matrix. By the lemma, $H \otimes \Lambda_{(p)}$ has an $n + r \times n + r$ presentation matrix A which looks like

$$A = \begin{pmatrix} P & 0 \\ 0 & 0 \end{pmatrix}$$

Now if $d < r$ then any $n + r - d \times n + r - d$ minor contains a row of zeroes, so $\Delta_d(H \otimes \Lambda_{(p)}) = 0$. If $d \geq r$, then the $n+r-d \times n+r-d$ subdeterminants of A are the $n - (d - r)$ subdeterminants of P. Applying this reasoning to each irreducible element of Λ and using Lemma 2.8 completes the proof. □

For modules over a Noetherian UFD Λ, the 0^{th} order of M has a special meaning as a second obstruction to the "vanishing" of M. In other words, if we want to detect whether a module M is non-trivial, we may first calculate its rank. If the rank vanishes, then the order is non-zero, and if $\text{ord}(M) \neq 1$ then we conclude the module is non-trivial.

Proposition 2.9. *Let M be a non-trivial Λ-module. Then* $\text{ord}(M) \neq 0 \iff \text{rank}(M) = 0$.

Proof. We've already noticed this when Λ is a PID. In general, suppose we have a presentation $\Lambda^m \xrightarrow{A} \Lambda^n \to M \to 0$. Applying the exact functor $- \otimes Q(\Lambda)$ we see that $rank(M) = 0$ iff $A \otimes Q(\Lambda)$ is surjective. This occurs if and only if the columns of $A \otimes Q(\Lambda)$ span $Q(\Lambda)^n$, which is equivalent to the existence of an $n \times n$ minor with non-zero determinant. □

The ring Λ_l of Laurent polynomials in l variables belongs to another general class of rings: group rings. We briefly list some important structures associated to every (integral) group ring.

Definition 2.10. Let G be a group. The map $\epsilon : \mathbb{Z}[G] \to \mathbb{Z}$ given by $\epsilon(g) = 1$ for all $g \in G$ is the **augmentation map**. Its kernel is the **augmentation ideal**, denoted ϵ_G.

The augmentation ideal is simply $(\{g - 1 \mid g \in G\}) \lhd \mathbb{Z}[G]$. We shall require the following facts regarding ϵ_F when $F \cong \mathbb{Z}^l$ is a free abelian group:

Proposition 2.11. *Let F be a free abelian group. Then ϵ_F, considered as a $\mathbb{Z}[F]$-module, is torsion free with rank* 1.

Proof. That ϵ_F is torsion free is clear from the fact that it is an ideal of $\mathbb{Z}[F]$, which is a domain (it is a ring of integral Laurent polynomials). This implies its rank is non-zero. On the other hand, its rank cannot exceed 1 since it is a submodule of $\mathbb{Z}[F]$, which has rank 1. □

2.2 The Alexander invariants

The Alexander polynomial is a multivariable Laurent polynomial obtained from the homology of the maximal free abelian cover of the link complement. Our approach, following [McM02] associates Alexander invariants to every free abelian cover of the link complement. This allows us to use single-variable Laurent polynomial techniques and will be useful in comparing the Thurston and Alexander norms later.

Let L denote an l component link and (X, p) its (pointed) complement in S^3 and $G = \pi_1(X)$. Suppose we are given a surjective homomorphism $\phi : G \to F$ to a free abelian group $F \cong \mathbb{Z}^b$. This map corresponds to a covering space $\pi : X_\phi \to X$ whose group of deck transformations is F. The covering space is characterized by $\pi_*(\pi_1(X_\phi)) = \ker\phi \leq G$. Let $\tilde{p} = \pi^{-1}(p)$.

Definition 2.12. [McM02] The **Alexander module** is the $\mathbb{Z}[F]$-module

$$A_\phi(L) = H_1(X_\phi, \tilde{p})$$

where the F-action is given by its action on $\tilde{X}$ by deck transformations. The **Alexander ideal** is the first elementary ideal of the Alexander module:

$$I_\phi(L) = E_1(A_\phi(L)) \lhd \mathbb{Z}[F]$$

The **Alexander polynomial** corresponding to ϕ is the first order of the Alexander module, i.e.

$$\Delta_\phi(L) = \gcd(I_\phi) \in \mathbb{Z}[F]$$

Occasionally we omit reference to the link L in the notation. When ϕ is the Hurewicz/abelianization map $\pi_1(X) \to H_1(X) \cong \pi_1(X)/[\pi_1(X), \pi_1(X)]$, the corresponding covering space is the maximal free abelian cover of X. We abbreviate the above definitions in this case and write $A(L)$, $I(L)$ and $\Delta(L)$, and the maximal free abelian cover is denoted X_∞. Unless the context states otherwise, the **(multivariable) Alexander polynomial** refers to $\Delta(L)$.

Note that the Alexander polynomial is defined only up to multiplication by units in $\mathbb{Z}[F]$. This ring may be viewed as the ring of Laurent polynomials in b variables upon picking a basis for $F \cong \mathbb{Z}^b$. It is often convenient to pick $\phi : G \to \mathbb{Z}$ so that A_ϕ is a Laurent polynomial in just one variable.

It is sometimes useful (and often done in the literature) to instead consider the absolute homology of X_ϕ:

Definition 2.13. The **Alexander invariant** corresponding to ϕ is the module $H_1(X_\phi)$ with F acting as the group of deck transformations.

As usual, if the phrase Alexander invariant is mentioned without a homomorphism ϕ, it is implicit that ϕ is the abelianization map. The Alexander polynomial is then defined to be $\operatorname{ord} H_1(X_\phi)$. This is justified by the following:

Proposition 2.14. *Let ϕ and $(X_\phi, \tilde{p})$ be as above. Then*

$$\Delta(L) = \operatorname{ord}(H_1(X_\phi))$$

Proof. Consider the long exact sequence of homology for the pair $(X_\phi, \tilde{p})$. By the naturality of the sequence, it preserves the action by deck transformations, so it may be considered as a sequence of $\mathbb{Z}F$-modules. The non-trivial part is:

$$0 \to H_1(X_\phi) \to H_1(X_\phi, \tilde{p}) \to \tilde{H}_0(\tilde{p}) \to 0$$

The reduced homology group $\tilde{H}_0(\tilde{p})$ is the kernel of the augmentation map $\epsilon : \mathbb{Z}[F] \to \mathbb{Z}$, which is the augmentation ideal ϵ_F. By Proposition 2.11, ϵ_F is torsion free with rank 1. Hence by Proposition 2.6,

$$\Delta_0(H_1(X_\phi)) = \Delta_1(H_1(X_\phi, \tilde{p}))$$

This proves the claim, as the first quantity is $\operatorname{ord}(H_1(X_\phi))$ and the second is $\Delta(L)$. □

2.3 Algebraic description of the Alexander invariants

As usual, we have a covering map $(X, \hat{p}) \xrightarrow{\pi} (X, p)$ such that $\pi_*(\pi_1(X, \hat{p}) = \ker \phi$. The Alexander invariant corresponding to $\phi : G \to \mathbb{Z}^b$ may be derived algebraically just from the inclusion $\ker \phi \hookrightarrow G$. For this reason many authors refer to the Alexander invariant of a group, and this can be extended to the Alexander invariant of any topological space X by setting the Alexander invariant of X to be that of its fundamental group. In particular, one may define this invariant for any 3-manifold, not just link complements. We will not take advantage of this level of generality but it is good to be aware of it.

For the algebraic description of the Alexander invariant corresponding to ϕ, we note that as a group $H_1(X_\phi) \cong \ker \phi/(\ker \phi)'$ by the Hurewicz theorem, where we are factoring by the commutator of $\ker \phi$. We will examine the actual map giving the isomorphism in more detail below. Also, the group F of deck transformations of X_ϕ is $G/\ker \phi$.

Definition 2.15. The **conjugation action** of $F = G/\ker\phi$ on $H = \ker\phi/(\ker\phi)'$ is the following (left) group action: let $x \in \ker\phi$ represent $[x] \in H$ and let $f \in G$ represent $[f] \in F$. Then

$$[f] \cdot [x] = [fxf^{-1}]$$

We will show this action is well defined and that $H_1(X_\phi) \cong \ker\phi/(\ker\phi)'$ as F-modules. Since the conjugation action doesn't use any topological information, we will have shown that the Alexander invariants are algebraic invariants derived from the fundamental group of the link complement.

To show the conjugation action is well defined, suppose $x, y \in \ker\phi$ are congruent modulo $(\ker\phi)$ and $f, g \in G$ are congruent modulo $\ker\phi$. We must check that $fxf^{-1} \equiv gyg^{-1}$ modulo $(\ker\phi)'$ (note that these elements are clearly in $\ker\phi$ since it is a normal subgroup). To that end, we may write $y = cx$ and $g = tf$ for $c \in (\ker\phi)'$ and $t \in \ker\phi$. Then we have:

$$\begin{aligned} fxf^{-1} \equiv gyg^{-1} \bmod (\ker\phi)' &\iff fxf^{-1}gy^{-1}g^{-1} \in (\ker\phi)' \\ &\iff fxf^{-1}tfx^{-1}c^{-1}f^{-1}t^{-1} \in (\ker\phi)' \\ &\iff (fc^{-1}f^{-1})(fcxf^{-1}tfx^{-1}c^{-1}f^{-1}t^{-1}) \in (\ker\phi)' \end{aligned}$$

The first term in parentheses is in $(\ker\phi)'$ since $c \in (\ker\phi)'$ and the commutator subgroup is a characteristic subgroup (preserved under automorphism). The second term is just the commutator $[fcxf^{-1}, t]$ which is also in $(\ker\phi)'$.

Next we want to compare the action of F on $H_1(X_\phi)$ and $\ker\phi/(\ker\phi)'$. In the definition of the conjugation action, all elements are chosen from the ambient group G. In the topological setup, G is $\pi_1(X, p)$ and so to establish an equivalence of the two F-actions we translate the action of F by deck transformations into multiplication in $\pi_1(X, p)$. Concretely, a homology class $\alpha \in H_1(X_\phi)$ is represented by a loop γ with basepoint $\hat{p}$ (this is much of the content of the Hurewicz theorem). Thus α is the image of γ under the Hurewicz homomorphism $h : \pi_1(X_\phi, \hat{p}) \to H_1(X_\phi)$. Now $[\gamma]$ sent into the ambient group $G = \pi_1(X, p)$ via the injection π_*. On the other hand, an element $f \in F$ is also represented by a loop $[f] \in \pi_1(X, p)$. Our goal is to show $f_*(\alpha) = h([f][\gamma][f]^{-1})$ where f_* is the induced map on homology, and h is shorthand for the restriction of π_*^{-1} to $\ker\phi$ followed by h.

Let us calculate $\pi_*(f_*\alpha)$. First, the class $f_*(\alpha)$ is represented by $f \circ \gamma$, which is a loop in X_ϕ with basepoint $f(\hat{p})$. We want to apply the Hurewicz map to calculate the homology class of $f \circ \gamma$, but the Hurewicz map is defined only on loops with basepoint $\hat{p}$. Therefore, we modify $f \circ \gamma$: if δ is a path from $\hat{p}$ to $f(\hat{p})$, then $\delta \star (f \circ \gamma) \star \bar{\delta}$ is a loop with basepoint $\hat{p}$ that is homologous to $f \circ \gamma$ and now $h([\delta \star (f \circ \gamma) \star \bar{\delta}]) = f_*(\alpha)$ (here $\star$ is the normal concatenation operation for paths and $\bar{\delta}$ is the opposite path of δ). Moving to $\pi_1(X, p)$ via the projection, we have $\pi_*([\delta \star (f \circ \gamma) \star \bar{\delta}]) = [f][\pi(\gamma)][f^{-1}] \in \ker\phi \subset \pi_1(X, p)$. Applying the Hurewicz map to both sides gives $f_*(\alpha) = h([f][\pi_*\gamma][f^{-1}])$ as desired.

We have proved (Cf. [Rol76], pp. 174-175):

Theorem 2.16. *Suppose $(X_\phi, \hat{p}) \xrightarrow{\pi} (X, p)$ is a covering map with $\pi_*(\pi_1(X_\phi, \hat{p})) = \ker\phi$ and denote $F = G/\ker\phi$. Let*

$$\mathfrak{h} : H_1(X_\phi) \to \ker\phi/(\ker\phi)'$$

be the inverse of the Hurewicz isomorphism $h : \ker\phi/(\ker\phi)' \to H_1(X_\phi)$. Then $\mathfrak{h}$ is an F-module isomorphism where $\ker\phi/(\ker\phi)'$ is equipped with the conjugation action. □

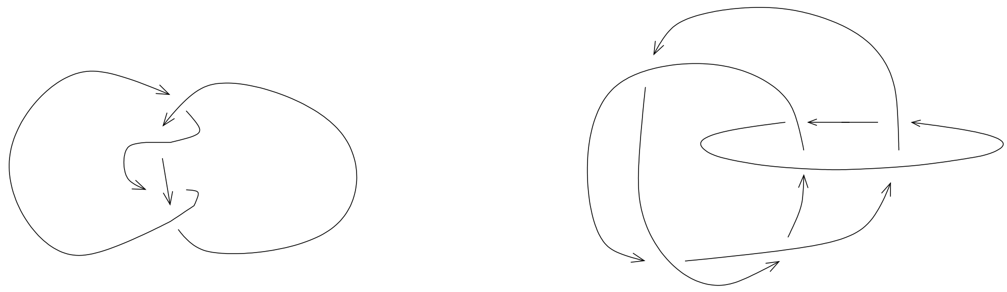

Figure 2.1: Two links with homeomorphic complements and different Alexander polynomials

Corollary 2.17. *The* $\mathbb{Z}[\mathbb{Z}^b]$*-module structure of the Alexander invariant corresponding to* ϕ *depends only on the injection* $\ker\phi \subset \pi_1(X, p)$. □

There is a subtlety in the above statement that we will explore in the next section. This is the fact that the orientation of a link provides us with a canonical choice of basis of $\mathbb{Z}^l$, corresponding to the oriented meridians of the boundary of the link complement. This is extra information (it is called "peripheral" data in the literature) and is not included in the data of the fundamental group.

2.4 The choice of meridians

Here we outline an example pointed out by Rolfsen ([Rol76], pp. 195 - 196, to which we refer the reader for details. Consider the two component links given in Figure 2.1.

Link A in Figure 2.1 is the (4,2) torus link, whose Alexander polynomial is computed later to be $1 + t_1t_2$ (see Example 3.5). From [Rol76], we see the Alexander polynomial of Link B is $1 + t_1^3t_2$.

On the other hand, Rolfsen shows that these two links have homeomorphic complements. Therefore, by the last corollary, their Alexander invariants have the same $\mathbb{Z}[\mathbb{Z}^2]$-module structure. However, this discrepancy is resolved by realizing that the oriented meridians, providing bases for $\mathbb{Z}^2$, do not coincide with the induced isomorphism of fundamental groups of the complements. This shows that the peripheral information of the choice of meridians gives us more data than is included in the fundamental group.

For knots, this problem never arises, due to the celebrated theorem of Gordon and Luecke which states that knots are determined by the homeomorphism type of their complement [GL89].

2.5 Construction of infinite cyclic covers

In this section we provide a geometric construction for for infinite cyclic covers that often allows us to compute the corresponding Alexander polynomial. This is a technique used in the early days of the Alexander polynomial, based on Seifert surfaces. We consider now the bounded link complement $M = S^3 - \mathring{\nu}(L)$ with a choice of oriented meridians t_i circling the i^{th} boundary torus. Suppose $\phi : G = \pi_1(M) \to \mathbb{Z}$ is surjective. Then ϕ

represents a cohomology class since it factors through $G/G' = H_1(M;\mathbb{Z})$. Suppose also that S is an embedded connected compact oriented surface that is Poincaré dual to ϕ. By Proposition 1.12, the boundary of S wraps $\phi(t_i)$ times longitudinally around the i^{th} boundary component.

The assumption that S is connected allows us to geometrically construct the cover $M_\phi \to M$ corresponding to ϕ. Let $\nu(S)$ be a tubular neighborhood of S and let S^+ and S^- denote the positive and negative sides of $\nu(S)$ according to the orientation of S and M, so that we have a disjoint union

$$\nu(S) = S^- \cup S \cup S^+$$

Now take countably many disjoint copies of $M - S$ labelled N_i. We glue these together using by attaching $\nu(S)$ to N_i and N_{i+1} by identifying S^- and S^+ in each piece. The result is a 3-manifold $\tilde{M}$ equipped with an obvious covering map $\pi : (\tilde{M}, \hat{p}) \to (M, p)$ with $\mathbb{Z}$ as the group of deck transformations. Here we pick $p \in M - S$ and $\hat{p} \in N_0$, and identify $\mathbb{Z} = \langle t \rangle$ where t acts by moving N_i to N_{i+1}. We claim this covering space is just M_ϕ, the covering space corresponding to ϕ.

Proposition 2.18. *The space $\tilde{M}$ constructed above is equal to M_ϕ as a covering space.*

Proof. It is enough to show that $\pi_*(\pi_1(\tilde{M})) = \ker\phi$ as this is a characterizing property of M_ϕ. Equivalently, on the level of deck transformations, we show that $\gamma \in \pi_1(M, p)$ lifts to a path from $\hat{p}$ to $t \cdot \hat{p}$ iff $\phi(\gamma) = 1$. By Proposition 1.11, $\phi(\gamma) = 1$ iff the algebraic intersection of S and γ is exactly 1. This is equivalent to the lift of γ ending at $\hat{t}$, since this means that the lift of γ passes from N_0 to N_1, where $t \cdot \hat{p}$ lies. □

Examples in which the Alexander polynomial are computed for some simple links by examining these covers are given in [Rol76].

2.6 The Fox free calculus

The Fox calculus is a handy tool that lets us calculate the Alexander polynomial from a presentation of the fundamental group of the link complement.

First we recall and fix some notation: Λ_k denotes the group ring $\mathbb{Z}[\mathbb{Z}^k]$. For a commutative ring R, $R\langle a_1, \dots, a_k\rangle$ denotes the free R-module generated by $a_1, \dots, a_k$. Let F^n denote the free group generated by n elements, typically labelled $\{x_1, \dots, x_n\}$.

Fox proved that for each $1 \le i \le n$ there exists a unique map

$$\frac{\partial}{\partial x_i} : F^n \to \mathbb{Z}F^n$$

satisfying the identities:

$$\frac{\partial x_j}{\partial x_i} = \delta_{i,j}$$
$$\frac{\partial uv}{\partial x_i} = \frac{\partial u}{\partial x_i} + u\frac{\partial v}{\partial x_i}$$

for any two words $u, v \in F^n$.

We shall show how to compute the Alexander polynomial from a presentation of the link group. Suppose $\langle x_1, \dots x_n | r_1, \dots r_m\rangle$ is a presentation of $\pi_1(X)$. Given a homomorphism $\phi : \pi_1(X) \to \mathbb{Z}^r$ we may extend it to a map $\bar{\phi} : \mathbb{Z}[F^n] \to \mathbb{Z}[\mathbb{Z}^r] = \Lambda_r$ by composing ϕ with the map $F^n \to \pi_1(M)$ given by the presentation and extending linearly.

Definition 2.19. The **Alexander matrix** corresponding to ϕ is the $n \times m$ matrix $\left(\bar{\phi} \circ \frac{\partial}{\partial x_j}(r_i)\right)$, where the relations r_i are considerd as words in $F\langle x_1, \ldots, x_n\rangle$.

As usual, the Alexander matrix referenced without mention of a particular homomorphism means the matrix corresponding to the abelianization/Hurewicz map $\pi_1(X) \to H_1(X) \cong \mathbb{Z}^l$.

Theorem 2.20. *The Alexander polynomial corresponding to* ϕ *is given by the* gcd *of the* $n-1 \times n-1$ *determinants of the Alexander matrix corresponding to* ϕ*.*

In particular, the gcd of the ideal generated by the $n-1 \times n-1$ determinants is invariant of the choice of presentation used to construct the Alexander matrix. Crowell and Fox [CF63] confine their discussion to the Alexander matrix and prove the invariance of this gcd using Tietze moves and linear algebraic methods. Instead, we shall prove Theorem 2.20, which will establish this invariance by the invariance of the Alexander polynomial.

The heart of Theorem 2.20 lies in the topological interpretation of the Fox calculus, which associates the Alexander matrix with the second CW-boundary map of a CW complex that is built from the data of a presentation of $\pi_1(X)$. Explicitly, we consider a simplified topological space that has the same fundamental group as X (and hence the Alexander invariants are the same). To that end, we let (W, b) be the pointed two dimensional CW realization of $\pi_1(X)$, using the presentation $\langle x_1, \ldots x_n | r_1, \ldots r_m\rangle$. The 0-skeleton of (W, b), denoted W^0, consists of a single point $\{b\}$, the 1-skeleton, denoted W^1, is a bouquet of circles corresponding to the generators $\{x_i\}$, and the 2-skeleton, denoted W^2, consists of m $2-$cells corresponding to the relations r_j, each attached to W^1 according to the word of the relation. This construction ensures that $\pi_1(W) = \pi_1(X)$.

Now let $W_\phi \xrightarrow{p} W$ be the free abelian cover corresponding to ϕ with group of deck transformations $\pi_1(X)/\ker\phi \cong \mathbb{Z}^r$. As a covering space, W_ϕ has a natural CW-decomposition composed of lifts of cells in W. By fixing a basepoint $\hat{b} \in p^{-1}(b)$, a cell in W_ϕ consists of a choice of a cell in W that it covers, and a deck transformation that determines the basepoint (which is the image of $\hat{b}$ under this deck transformation). This gives the CW-chain groups a Λ_r-module structure via the identifications:

$$\begin{array}{ccccccccc} 0 & \longrightarrow & C_2(W_\phi) & \xrightarrow{\partial_2} & C_1(W_\phi) & \xrightarrow{\partial_1} & C_0(W_\phi) & \longrightarrow & 0 \\ & & \| & & \| & & \| & & \\ & & \Lambda_r\langle r_1, \ldots r_m\rangle & \xrightarrow{\exists\ d_2} & \Lambda_r\langle x_1, \ldots x_n\rangle & \xrightarrow{\exists\ d_1} & \Lambda_r\langle b\rangle & & \end{array}$$

The maps d_i are Λ_r-module homomorphisms, since the Λ_r action comes from the action of deck transformations on $C_i(W_\phi)$, which commutes with boundary maps.

Then we have:

Theorem 2.21 (Topological Interpretation of the Fox Calculus)**.** *(Cf. [Kaw90], [Tur00], [Hir97]) The Alexander matrix* $\left(\bar{\phi} \circ \frac{\partial}{\partial x_i}(r_j)\right)$ *is the matrix of* d_2 *with respect to the bases in the diagram above.*

Proof. The idea of the proof is relatively simple, although the technicalities may be confusing. Intuitively, we want to define a map $F^n \to C_1(W_\phi)$ that coincides with d_2 on the generators $\{r_1, \ldots, r_m\}$. On these generators, the boundary map d_2 is actually a lifting map between W and W_ϕ, and we will extend this lifting map to arbitrary words of F^n. From lifting properties, it will be clear that this map shares the characterizing properties of the Fox derivatives.

First some notational preparation. The composition $F^n \to \pi_1(M) \to \mathbb{Z}^r$ factors through the abelianization of F^n via a map $\psi \circ \alpha$:

$$\begin{array}{ccc} F^n & \xrightarrow{\alpha} & \mathbb{Z}^n \\ \downarrow & & \psi\downarrow \\ \pi_1(M) & \xrightarrow{\phi} & \mathbb{Z}^r \end{array} \tag{2.1}$$

We label a basis $\{x_1, \ldots, x_n\}$ of $\mathbb{Z}^n$ so that $\alpha(x_i) = x_i$. Upon applying the integral group ring functor, we obtain a diagram containing $\bar{\phi}$:

$$\begin{array}{ccc} \mathbb{Z}[F^n] & \xrightarrow{\alpha_*} & \mathbb{Z}[\mathbb{Z}^n] = \Lambda^n \\ & \searrow^{\bar{\phi}} & \downarrow \psi_* \\ & & \mathbb{Z}[\mathbb{Z}^r] = \Lambda_r \end{array}$$

We must compute the CW-boundary map d_2 on the generating elements $r_1, \ldots, r_m$ of $C_2(W_\phi)$. By the identifications given above, it is clear that the boundary of the 2-cell corresponding to a generator r_i is the lift of the boundary of the 2-cell corresponding to r_i in the original CW-complex W. In general, we can compute the lift of a 1-cell of W as follows.

Note that W^1 is a bouquet of n-circles, and the covering map $W_\phi \to W$ restricts to a covering map on the 1-skeletons: $W^1_\phi \to W^1$. Let $(L_n, b_\infty) \to (W^1, b)$ denote the maximal abelian cover of W^1. This is a lattice on n generators. These two coverings induce a covering map $L_n \xrightarrow{\eta} W^1$ making the following diagram commute:

$$\begin{array}{ccc} L_n & \xrightarrow{\eta} & W^1_\phi \\ \downarrow & \swarrow & \\ W^1 & & \end{array}$$

A lift of a word $r \in \pi_1(W^1) = F_n$ is obtained by first lifting to the lattice L_n and then applying η. Actually, we are interested in lifting at the level of 1-cells so we consider the lifting map $\pi_1(W^1) \xrightarrow{l} C_1(L_n)$. As we did for $C_1(W_\phi)$, we can naturally identify the 1-chain group $C_1(L_n)$ with $\Lambda_n\langle x_1, \ldots, x_n\rangle$ and η induces a map η_* on chain groups. Then we have:

$$d_2(r_i) = \eta_* \circ l(r_i) \tag{2.2}$$

We will look at these two maps more closely. First examine the map $\eta_* : \Lambda_n\langle x_1, \ldots, x_n\rangle \to \Lambda_r\langle x_1, \ldots, x_n\rangle$. It is a $\mathbb{Z}$-module homomorphism, so to understand it we need only compute $\eta_*(tx_i)$ for $t \in \mathbb{Z}^n$. First, by our identifications, the element tx_i is the 1-cell which is a lift of $x_i \in \pi_1(W^1)$ starting at the basepoint t. This basepoint is determined as the endpoint of a lift of a loop γ in $\pi_1(W^1)$ to L_n which lies in the coset $\alpha^{-1}(t)$ (recall $\alpha : F^n \to \mathbb{Z}^n$ is the abelianization map). Now the 1-cell $\eta_*(tx_i)$ is obtained as the lift of x_i whose basepoint is the endpoint of a lift of γ in the covering $W_\phi \xrightarrow{p} W$. This is determined by ϕ. Therefore, by the commutativity of Diagram 2.1, $\eta_*(tx_i) = \psi(t)x_i$.

Concisely, this means we can express η_* as an $n \times n$ diagonal matrix

$$\eta_* = \begin{pmatrix} \psi_* & & \\ & \ddots & \\ & & \psi_* \end{pmatrix} \tag{2.3}$$

Next, we consider the lifting map $l : F^n \to \Lambda_n\langle x_1, \ldots, x_n\rangle$ where we have identified $\pi_1(W^1) = F^n$ and $C_1(L_n) = \Lambda_n\langle x_1, \ldots, x_n\rangle$. Let l_i be the ith component function. Then clearly $l_i(x_i) = 1$ since $l(x_i) = x_i$. Suppose $u, v \in F^n$. Then $l(uv)$ as a path in the lattice L_n consists of a lift of u with basepoint p_∞ followed by a lift of v with basepoint $\alpha(u)$, since α is interpreted as the map taking loops to deck transformations. Therefore $l(uv) = l(u) + \alpha(u)l(v)$. This relation descends to the components functions of l_i. Hence the components l_i satsify $l_i(x_j) = \delta_{i,j}$ and $l_i(uv) = l_i(u) + \alpha(u)l_i(v)$. The map $\alpha_* \circ \frac{\partial}{\partial x_i}$ also satisfies these properties; hence they agree on every element of F^n by induction. Therefore we may write

$$l = \begin{pmatrix} \alpha_* \circ \frac{\partial}{\partial x_i} \\ \vdots \\ \alpha_* \circ \frac{\partial}{\partial x_n} \end{pmatrix} \tag{2.4}$$

Finally, we combine Equations 2.2, 2.3, and 2.4 to get

$$\begin{aligned} d_2(r_j) = \eta_* \circ l(r_j) &= \begin{pmatrix} \psi_* & & \\ & \ddots & \\ & & \psi_* \end{pmatrix} \begin{pmatrix} \alpha_* \circ \frac{\partial}{\partial x_1} \\ \vdots \\ \alpha_* \circ \frac{\partial}{\partial x_n} \end{pmatrix} (r_j) \\ &= \begin{pmatrix} \bar{\phi} \circ \frac{\partial}{\partial x_1}(r_j) \\ \vdots \\ \bar{\phi} \circ \frac{\partial}{\partial x_n}(r_j) \end{pmatrix} \end{aligned}$$

Therefore d_2 has the matrix $\left(\bar{\phi} \circ \frac{\partial}{\partial x_i}(r_j)\right)$ as desired. □

This remarkable interpretation lets us easily prove Thm 2.20 and gives us insight into the Fox calculus. Note that the topological interpretation does not require us to start with a topological space; all we used was a presentation of a group. So this theorem gives a concrete bridge between purely algebraic and topological concepts.

Proof of 2.20. To determine the Alexander polynomial $\Delta_\phi(L)$ we must consider the Λ_r-module $H_1(W_\phi, p^{-1}(b))$ and calculate its first order. Fortunately, a presentation of $H_1(W_\phi, p^{-1}(b))$ is given by the CW-structure of W_ϕ, from the long exact sequence of the triple $(W_\phi, W_1, p^{-1}(b))$. Indeed, denoting the CW-chain complex of W_ϕ by (C_*, ∂_*) we have the presentation

$$\mathcal{C}_2 \xrightarrow{\partial_2} \mathcal{C}_1 \xrightarrow{\partial_1} H_1(W_\phi, p^{-1}(b)) \longrightarrow 0$$

This is a presentation of abelian groups, but under the identifications given in Theorem 2.21 may also be considered an Λ_l-module presentation. Therefore, we may use this sequence to calculate the first order of $H_1(W_\phi, p^{-1}(b))$. By the topological interpretation of the Fox calculus, the map δ_2 is represented by the (transposed) Alexander matrix in the usual bases. Therefore the first order of $H_1(W_\phi, p^{-1}(b))$ is given by the gcd of the $n-1 \times n-1$ subdeterminants of the Alexander matrix. □

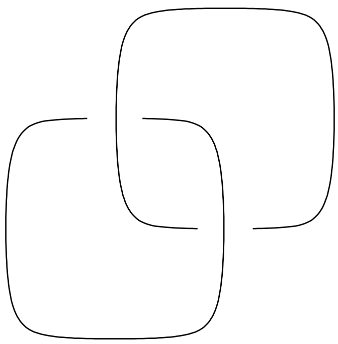

Figure 2.2: The Hopf link.

Let $h : \pi_1(X) \to \mathbb{Z}^l$ denote the abelianization map. Then any map $\phi : \pi_1(X) \to \mathbb{Z}^r$ factors through h, and the factor extends linearly to a map denoted $\phi_* : \mathbb{Z}[\mathbb{Z}^l] \to \mathbb{Z}[\mathbb{Z}^r]$. Clearly the Alexander matrix corresponding to ϕ is obtained by applying ϕ_* to the entries of the Alexander matrix. Therefore we have:

Corollary 2.22. *The Alexander ideal corresponding to ϕ may be computed from the Alexander ideal by*

$$I_\phi = \phi_*(I)$$

This reflects the naturality of the covering constructions. We remark that this relation does not extend to the level of gcds, since the gcd operation is not functorial. For example, suppose $\phi : \mathbb{Z}langlet_1, t_2\rangle \to \mathbb{Z}\langle t\rangle$ takes t_1 and t_2 to t. Then

$$\phi(\gcd(t_1 - 1, t_2 - 1)) = \phi(1) = 1$$

whereas

$$\gcd(\phi(t_1 - 1), \phi(t_2 - 1)) = \gcd(t - 1, t - 1) = t - 1$$

Example 2.23. Let L be the Hopf link given in Figure 2.2. The complement of one component is a solid torus, so the complement of the entire link is a torus. Its fundamental group is $\mathbb{Z} \oplus \mathbb{Z}$, with a presentation given by

$$\pi_1(S^3 - L) = \langle x, y \mid xyx^{-1}y^{-1}\rangle$$

Note that x and y are represented by oriented meridian loops around the components of the link. We compute the Fox derivatives of the single relation as follows:

$$\frac{\partial}{\partial x}(xyx^{-1}y^{-1}) = \frac{\partial}{\partial x}(x) + x\frac{\partial}{\partial x}(yx^{-1}y^{-1}) = 1 + xy\frac{\partial}{\partial x}(x^{-1}y^{-1}) = 1 - xyx^{-1}$$

where we used the identity $\frac{\partial}{\partial x}(x^{-1}) = -x^{-1}$, which is easily derivable by noting $\frac{\partial}{\partial x}(1) = 0$ and applying the product rule. Similarly,

$$\frac{\partial}{\partial y}(xyx^{-1}y^{-1}) = x - xyx^{-1}y^{-1}$$

Now we compute the Alexander polynomial using the map $\alpha_* : \mathbb{Z}[F^n] \to \Lambda_2$ induced by the Hurewicz homomorphism $\alpha : \pi_1(X) \to H_1(X; \mathbb{Z}) \cong \mathbb{Z}^2$. Depending on the orientations of the link components and the choice of generators x and y, the map α takes x to $t_1^{\pm 1}$ and y to $t_2^{\pm 1}$. Therefore the Alexander matrix has the form:

$$\begin{pmatrix} 1 - t_2^{\pm 1} & t_1^{\pm 1} - 1 \end{pmatrix}$$

In any case, the gcd of the 1×1 subdeterminants is 1. Therefore the Hopf link has trivial Alexander polynomial.

The topological interpretation also gives us a very nice linear relation between the columns of the Alexander matrix, due to the equation $d_1 \circ d_2 = 0$:

Proposition 2.24. *Let C_i denote the i^{th} column of the Alexander matrix (i.e. the column corresponding to x_i). Then*

$$\sum_{i=1}^{n} (\alpha(x_i) - 1)C_i = 0$$

Proof. Examine the action of d_1 on a basis element $x_i \in C_1(W_\infty)$. By our identification, the cell represented by x_i corresponds to a lift of the i^{th} circle in the 1-skeleton of W, with basepoint $\hat{b}$. The boundary map is the formal difference of the endpoint minus the basepoint, which are $\alpha(x_i)\hat{b}$ and $\hat{b}$ respectively. Therefore $d_1(x_i) = (\alpha(x_i) - 1)b \in C_0(W_\infty)$. Now translating the relation $d_1 \circ d_2 = 0$ into matrix multiplication completes the proof. □

Remarkably, we have derived a purely algebraic property of the Fox calculus using topological methods. This relation is useful on many occasions when we have a presentation consisting of one less relation than generator, so that the Alexander matrix is $n \times n - 1$. The relation between the columns allow us to calculate the gcd of the $n-1 \times n-1$ minors using only the determinant of a single minor. This fact is expressed without referring to the particular minors below.

Proposition 2.25. *For a link with more than one component, the Alexander ideal I is $\epsilon_{\Lambda_l} \cdot (\Delta)$.*

Recall that ϵ_{Λ_l} is the augmentation ideal generated by $(t_1 - 1, \ldots, t_l - l)$. Compare [McM02], Theorem 5.1.

Proof. The Neuwirth presentation given in Theorem 1.7 gives us a presentation of $\pi_1(S^3 - L)$ with n generators and $n - 1$ relations, so the Alexander matrix is of size $n-1 \times n$. Let M_i denote the i^{th} column of the Alexander matrix and M_i the determinant of the minor obtained by removing C_i. By definition $\Delta = \gcd(M_1, \ldots, M_n)$. For $i \neq j$, using the previous proposition we calculate:

$$\begin{aligned}
(\alpha(x_j) - 1)M_i &= \det(C_1, \ldots, (\alpha(x_j) - 1)C_j, \ldots, \hat{C}_i, \ldots, C_n) \\
&= \det(C_1, \ldots, -\sum_{k \neq j} (\alpha(x_k) - 1)C_k, \ldots, \hat{C}_i, \ldots, C_n) \\
&= \det(C_1, \ldots, -(\alpha(x_i) - 1)C_i, \ldots, \hat{C}_i, \ldots, C_n) \\
&= \pm \det(C_1, \ldots, \hat{C}_j, \ldots, (\alpha(x_i) - 1)C_i, \ldots, C_n) = \pm(\alpha(x_i) - 1)M_j
\end{aligned}$$

So we have the following relation among the minors:

$$\frac{M_i}{\alpha(x_i) - 1} = \pm \frac{M_j}{\alpha(x_j) - 1}$$

Now for the Neuwirth presentation in particular, $\alpha(x_i)$ is exactly $t_{a(i)}$ where $a(i)$ computes the index of the link that the loop x_i wraps around. Note that every generator t_k is the image under α of some x_i. Furthermore, for $k \neq j$, $t_k - 1$ and $t_j - 1$ share no common factors in Λ_l. Therefore:

$$\Delta \doteq \gcd(M_1, \ldots, M_n) = \frac{M_i}{\alpha(x_i) - 1} \qquad \forall i = 1, \ldots, n$$

From this equation it is clear that $\epsilon_{\Lambda_l} \cdot (\Delta) \supseteq (M_1, \ldots, M_n)$. On the other hand, by the remark above, for any t_k we have $\alpha(x_i) = t_k$ for some x_i and so $(t_k - 1)\Delta = M_k$ which shows the reverse inclusion. □

Remark. For knots, $\Lambda_l = \mathbb{Z}[t, t^{-1}]$ is a PID so the Alexander ideal is simply (Δ). The above argument breaks down for knots because in this case $\alpha(x_i)$ is always a power of t, so $\alpha(x_i) - 1$ is always divisible by $(t - 1)$. In particular these elements are not relatively prime.

Chapter 3

The multivariable Alexander polynomial via grid diagrams

Let G denote a fixed grid diagram. We discuss the calculation of the multivariable Alexander polynomial (of a link) from a grid diagram. The multivariable Alexander polynomial may be retrieved as the Euler characteristic of a suitable grid homology theory (denoted $\widehat{CL}(G)$ in [MOST07]). Defined in this way, the status of the polynomial as a knot invariant follows from the invariance of the entire homology theory. Here we instead take a more elementary approach and define the multivariable Alexander polynomial solely via the grid diagram, and directly prove its invariance under commutation and stabilization moves. Cromwell's Theorem then asserts that it is a knot invariant. Let us define some quantities obtained from a grid diagram.

The **winding matrix W**(G) is the matrix whose (i, j)-th entry is the l-component vector

$$(w_1(i, j), \ldots, w_l(i, j))$$

where $w_k(i, j)$ is the winding number of the k-th component of the link around the vertex corresponding to the (i, j)-th entry of the matrix. We throw out the bottom row and right hand column in accordance with our interpretation of the grid diagram as living on a torus. We adopt the convention that we label vertices of a grid as we label the entries of the winding matrix, so we count rows from top to bottom starting at 1 and count columns from left to right starting at 1. The **grid matrix M**(G) has (i, j)-th entry $t^{-w_{i,j}}$, using multi-index notation. In other words, this entry is the product $t_1^{-w_1(i,j)} \cdots t_l^{-w_l(i,j)}$.

We denote the set of O's as $\mathbb{O}$ and the X's as $\mathbb{X}$. We sometimes consider these sets as bijections between columns and rows. In this sense, we define $\epsilon(G)$ to be the sign of the permutation connecting $\mathbb{O}$ with the configuration of O's running from the top left corner of the grid to the bottom right corner (we call this the **downwards diagonal configuration**. The quantity n_i is equal to the number of vertical segments in the i-th component of the link. Let P be the set of those vertices which are at the corner of an X or O marking. The quantity a_i is defined to be

$$a_i = \frac{1}{8} \sum_{p \in P} \{\text{winding numbers of the } i\text{-th component at the vertex } p\}$$

. Let n_i denote the number of vertical strands corresponding to the i^{th} component of the link L_i. Then we define:

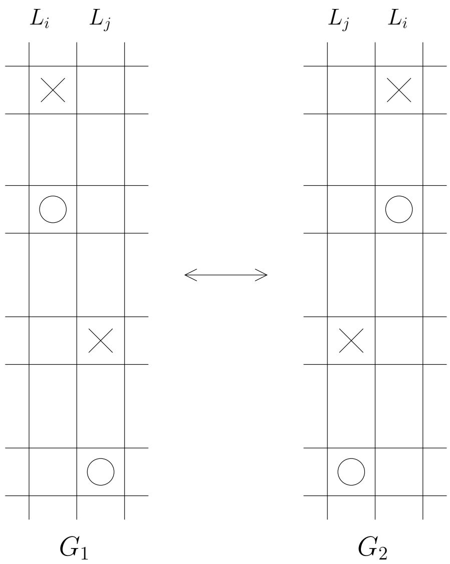

Figure 3.1: Commutation of the first type.

Definition 3.1. The **multivariable Alexander polynomial** associated to a grid diagram G of an oriented link $\vec{L}$ with l components is the polynomial in l variables

$$\Delta_G(t_1, \ldots, t_l) = \epsilon(G) \det \mathbf{M}(G) \prod_{i=1}^{l} (1 - t_i)^{-n_i} t_i^{a_i + \frac{n_i}{2}}$$

3.1 Link invariance

We would like to show Δ_G is a link invariant, i.e. is independent of the choice of grid diagram representing L. Cromwell's Theorem asserts that it is sufficient to show Δ_G is invariant under commutation and stabilization moves.

Lemma 3.2. *$\Delta_G(t_1, \ldots, t_l)$ is invariant under commutations.*

Proof. We consider two cases. First, suppose a commutation is such that the vertical portions of each link are non-overlapping (Figure 3.1). The associated matrices $\mathbf{M}(G_1)$ and $\mathbf{M}(G_2)$ differ only in the middle column of the commutation. To compare determinants, we subtract the left column from the middle column in $\mathbf{M}(G_1)$ and the right column from the middle column in $\mathbf{M}(G_2)$. Let w_i and w_j denote the winding numbers of the respective links around the vertex lying at the SE corner of the O marking

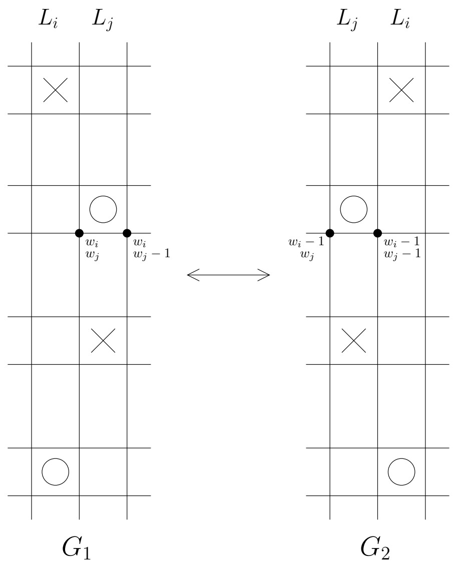

Figure 3.2: Commutation of the second type, with winding numbers indicated.

associated to L_i. Then we have (in the configuration of Figure 3.1):

$$\mathbf{M}(G_1) = \begin{pmatrix} & 0 & \\ & \vdots & \\ & 0 & \\ & t_1^{-w_i}t_2^{-w_j} - t_1^{-w_i-1}t_2^{-w_j} & \\ * & & * \\ & \vdots & \\ & t_1^{-w_i}t_2^{-w_j} - t_1^{-w_i-1}t_2^{-w_j} & \\ & 0 & \\ & \vdots & \\ & 0 & \end{pmatrix} \qquad \mathbf{M}(G_2) = \begin{pmatrix} & 0 & \\ & \vdots & \\ & 0 & \\ & t_1^{-w_i-1}t_2^{-w_j} - t_1^{-w_i}t_2^{-w_j} & \\ * & & * \\ & \vdots & \\ & t_1^{-w_i-1}t_2^{-w_j} - t_1^{-w_i}t_2^{-w_j} & \\ & 0 & \\ & \vdots & \\ & 0 & \end{pmatrix}$$

where $*$ denotes that the remaining entries are equal in both matrices. Hence we have $\det \mathbf{M}(G_1) = -\det \mathbf{M}(G_2)$ Clearly the quantities a_i, a_j, n_i and n_j are unchanged, and ϵ changes sign, so we have $\Delta_{G_1} = \Delta_{G_2}$. The other configurations of X's and O's are very similar.

Now suppose the commutation is such that one vertical segment is "contained" in the other (see Figure 3.2). Again the matrices $\mathbf{M}(G_1)$ and $\mathbf{M}(G_2)$ differ only in the middle column of the commutation. Now let w_i and w_j denote the winding numbers of the respective links at the SW corner of the O marking associated to L_j, as in Figure 3.2. We subtract the right column from the middle column in $\mathbf{M}(G_1)$ and the left column from the middle columin in $\mathbf{M}(G_2)$. In general, we always subtract the column with the shorter vertical segment from the middle segment. Now the middle columns of each matrix have nonzero entries in the same positions, and the non-zero entries are constant in each matrix, as in the previous case. In the case of Figure 3.2, the non-zero entries of $\mathbf{M}(G_1)$ are $t_i^{-w_i}t_j^{-w_j} - t_i^{-w_i}t_j^{-w_j+1}$ and the non-zero entries of $\mathbf{M}(G_2)$ are $t_i^{-w_i+1}t_j^{-w_j+1} - t_i^{-w_i+1}t_j^{-w_j}$. Therefore:

$$\det \mathbf{M}(G_1) = -t_i \det \mathbf{M}(G_2)$$

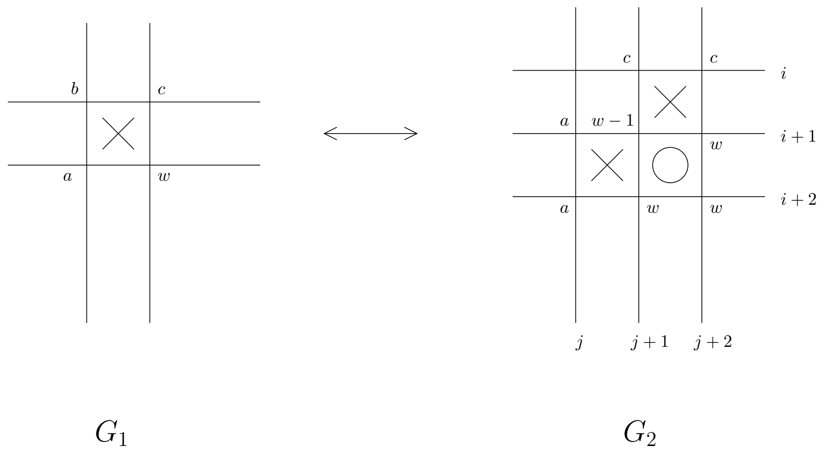

Figure 3.3: $X : NW$ stabilization, with winding numbers indicated.

Now in the expression for Δ_G, the n_k quantities remain the same and ϵ switches sign from G_1 to G_2. The only changes in the a_k quantities occur at the corners of the X and O marking of L_j in the commutation. Thus $a_j(G_2) = a_j(G_1)$ and $a_i(G_2) = -1 + a_i(G_1)$. Hence the $-t_i$ difference in determinants is absorbed by the chainge in ϵ and a_i. For other orientations, one may explicitly repeat the argument, or notice that the point of the calculation is that the determinant of $\mathbf{M}(G_2)$ is always a multiple of $\mathbf{M}(G_1)$ by $-t_i^{\pm 1}$. The overall sign is absorbed by ϵ, and the sign of the exponent is determined by the orientation of the vertical strand in L_i; this is exactly compensated by the change in a_i. □

Lemma 3.3. $\Delta_L(t_1, \dots, t_l)$ *is invariant under stabilizations.*

Proof. For concreteness, first consider an $X : NW$ stabilization (Figure 3.3) where we are stabilizing the ith component of the link, L_i. In Figure 3.3 we've labelled the winding numbers (with respect to L_i) of some of the corners by a, b, c, w, and the position of the stabilization in G_2 by i and j. Note that $w + b - a - c = 1$, since all but one of these quantities are equal, and the other differs by 1. If instead there was an O marking in the box, the sum would be -1.

Now subtract the $(j+2)^{nd}$ column from the $(j+1)^{st}$ column in $\mathbf{M}(G_2)$. The result has only one nonzero entry in the $j+1$ column, which is $t_i^{-w+1} - t^{-w}$, in the $i+1$ row. Clearly the corresponding minor is exactly $\mathbf{M}(G_1)$. Therefore

$$\det \mathbf{M}(G_2) = (-1)^{(i+1)+(j+1)}(t_i^{-w+1} - t_i^{-w}) \det \mathbf{M}(G_1)$$

Now calculate the change in a_i:

$$a_i(G_2) - a_i(G_1) = \frac{1}{8}(7w - 3 + a + c - b) = \frac{1}{8}(8w - 4) = w - 1/2$$

where the second equality is due to $w + b - a - c = 1$. Finally, the additional vertical segment in G_2 is incorporated in $n_i(G_2) = n_i(G_1) + 1$. Substituting these calculations into the formula for Δ_{G_2} we obtain:

$$\begin{aligned}\Delta_{G_2} &= \epsilon(G_2)\epsilon(G_1)(-1)^{i+j}(t_i^{-w+1} - t_i^{-w})(1 - t_i)^{-1} t_i^{w} \Delta_{G_1} \\ &= \epsilon(G_2)\epsilon(G_1)(-1)^{i+j}\big(t_i^{-w}(t_i - 1)\big)(1 - t_i)^{-1} t_i^{w} \Delta_{G_1} \\ &= \epsilon(G_2)\epsilon(G_1)(-1)^{i+j+1}\Delta_{G_1}\end{aligned}$$

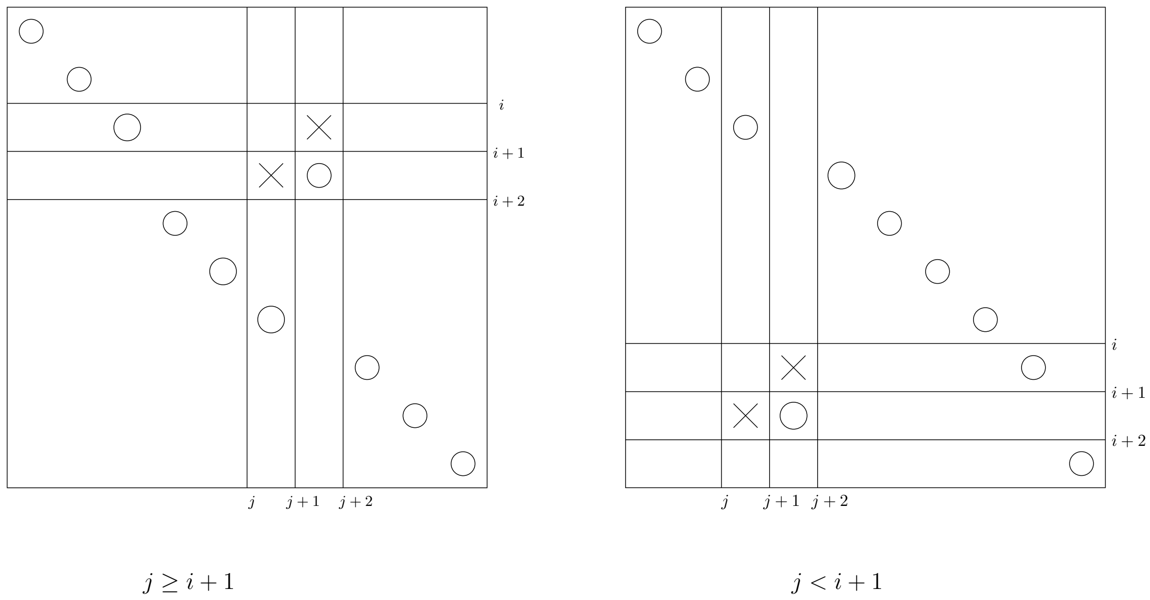

Figure 3.4: Result of applying $\tilde{\pi}$ to G_2.

Thus it remains to show $\epsilon(G_1)\epsilon(G_2) = (-1)^{i+j+1}$. By the definition of ϵ, there is some permutations π of columns in G_1 taking the O markings to the downwards diagonal configuration of O's and $\operatorname{sgn}\pi = \epsilon(G_1)$. We extend π to a permutation $\tilde{\pi}$ on the columns of G_2 by requiring $\tilde{\pi}$ to fix the right column in the stabilization (between the $j+1$ and $j+2$ arcs), and act as π on the remaining columns. Then $\operatorname{sgn}\tilde{\pi} = \operatorname{sgn}\pi = \epsilon(G_1)$. We have two cases: $j \geq i+1$ and $j < i+1$ according to whether the stabilization occurs above or below the downwards diagonal line. These cases are laid out schematically in Figure 3.4, where we show the effect of $\tilde{\pi}$ on G_2. In the first case, $j \geq i+1$, we may achieve a downwards diagonal configuration by successively moving the column of the stabilization with the O marking to the left; this requires $(j+1)-(i+2)$ transpositions. When $j < i+1$, we move this column towards the right, requiring $(i+1)-(j+2)$ transpositions to achieve a downwards diagonal configuration. In either case, the sign of the product of these permutations is $(-1)^{i+j+1}$. This shows Δ_G is invariant under an $X : NW$ stabilization. The other cases are omitted as they are similar, but we mention that since we proved commutation invariance, it is only necessary to examine three other cases; for example the other X stabilizations. □

Therefore, by Cromwell's Theorem, we have:

Corollary 3.4. *$\Delta_{\vec{L}}(t_1, \ldots, t_l)$ is a well defined link invariant.*

3.2 Equivalence of definitions

We will now justify the nomer "Alexander polynomial" that we've given to this invariant by proving that it is equal to the Alexander polynomial we've discussed in previous chapters. Of course, there the Alexander polynomial was defined only up to units in Λ_l, so this is the most we can prove. In other words, the polynomial obtained from a grid diagram is a particular normalization of the Alexander polynomial.

The details of the proof are a bit confusing, so it may help to look at the example following the proof to understand the steps.

Theorem 3.5. *[MOST07] Let G be a grid diagram for $\vec{L}$. Let Δ_G denote the polynomial obtained from the diagram G while Δ_L is the classical Alexander polynomial. Then $\Delta_G \doteq \Delta_L$.*

Proof. The proof is a standard application of the Fox calculus via Theorem 2.20. In the following $\sigma(k)$ will denote the index of the link component which is in the k^{th} column of the grid diagram. We use the Neuwirth presentation coming from the grid diagram to calculate Δ_L, and compare the determinant of the Alexander matrix with the determinant of the winding matrix. The Neuwirth presentation has n generators $x_1, \dots x_n$ corresponding to the columns of G. There are $n-1$ relations $r_1, \dots, r_{n-1}$, where r_j is the product of those x_i's for which the vertical strand in the i^{th} column crosses the j^{th} horizontal grid line. These relations are particularly nice because each generator appears at most once, with exponent 1. Explicitly, r_j has the form $x_{j_1} \dots x_{j_k}$ where each generator appears at most once in the sequence. Then we have:

$$\frac{\partial}{\partial x_i} r_j = \begin{cases} x_{j_1} \dots x_{j_{l-1}} & \text{if } x_{j_l} = x_i \\ 0 & \text{if } x_i \notin r_j \end{cases}$$

where the product is defined to be 1 if it is over an empty set (which happens when $x_{j_1} = x_i$). When we apply the abelianization map α, each x_i is mapped to $t_{\sigma(i)}^{\pm 1}$ where the exponent is determined by the orientation convention given in Figure 1.7. We see that x_i is sent to $t_{\sigma(i)}$ if the strand in the i^{th} column points upwards, and $t_{\sigma(i)}^{-1}$ if the strand points downward The result of the abelianzation is the Alexander matrix, denoted A with entries $A(i,j)$. Consider the j^{th} column of A. If the strand in the j^{th} column of the grid diagram does not cross the i^{th} horizontal grid line, then $A(i,j) = 0$. Now suppose this is not the case and $k \neq \sigma(j)$. Then the exponent of t_k in $A(i,j)$ is the number of upwards strands of L_k appearing in columns with index less than j minus the number of downwards such strands, which corresponds with $-w(i+1, j+1)$. If $k = \sigma(j)$ then the exponent of t_k is $-w(i+1, j+1) \pm 1$ depending whether the strand in the i^{th} column is oriented upwards or downwards. In summary, we have

$$A(i,j) = \frac{\prod_{k=1}^{l} t_k^{-w_k(i+1,j+1)}}{t_{\sigma(j)}^{\pm 1}}$$

Since we need only calculate the determinant up to multiplication by units, we can multiply each column by $t_{\sigma(j)}^{\pm 1}$ so that the non-zero entries correspond to the entries of the $n-1 \times n-1$ lower right minor of $\mathbf{M}(G)$.

Now examine the grid matrix $\mathbf{M}(G)$. We subtract the first column from the second, the second from the third, and so on. The top row is all 0's except for the first entry, which is a 1, so we only examine the lower right $n-1 \times n-1$ minor. The only non-zero entries in each column occur where there was a vertical strand in the previous column, and up to multiplication by a unit in Λ_l, each of these entires share a factor of $t_i - 1$ or $(t_i^{-1} - 1)$, where i is the component on the link appearing in that vertical strand. These are factored out of the determinant for a total factor of $((t_1 - 1)^{n_1} \cdots (t_n - 1)^{n_l}) (t_{\sigma(n)} - 1)^{-1}$ times some unit of Λ_l (this multiple comes from $(t_i^{-1} - 1) = -t^{-1}(t_i - 1)$. Note that we only have $n_{\sigma(n)} - 1$ such factors of $(t_{\sigma(n)} - 1)$ since we have not accounted for the strand in the last vertical column of the grid diagram.

But now we may multiply all the entries of the j^{th} column by $t_{\sigma(j)}^{\pm 1}$ so that the non-zero entries again encode the winding numbers in the exponents. In other words, the non-zero entries are equal to the entries of the lower right $n-1 \times n-1$ minor of $\mathbf{M}(G)$. The matrix derived from the Fox derivatives has zero entries in the same positions, and now we've seen that the non-zero entries can be made to agree through multiplication by units. Therefore, the determinants obtained from the two procedures are equivalent.

Finally, our computations for the Alexander polynomial come out to be the same because due to the linear relation between columns of the Alexander matrix from Proposition 2.24 (cf. the proof of Proposition 2.25), the Alexander polynomial is calculated from the minor we are using by dividing the determinant by $(t_{\sigma(n)} - 1)$. This accounts for the missing factor when we divided such elements from the matrix obtained from $\mathbf{M}(G)$. □

Example 3.6. We illustrate the proof with a calculation of the Alexander polynomial of the $(4, 2)$ torus link.

Figure 3.5 shows the calculations described in Theorem 3.5. The left column follows the definition of the polynomial taken from the grid diagram, while the right column is derived from the Neuwirth presentation and the Fox calculus. The resulting matrices differ only by unit multiples in the columns (in this case, just by a factor of t_1^{-1} in the last column). Of course this doesn't affect the unnormalized Alexander polynomial, which we calculate by taking the determinant of either matrix in the end and dividing by $(t_2 - 1)$. We conclude

$$\Delta(\text{(4,2) torus link}) \doteq \frac{1 - t_2^{-1} + t_1^{-1}t_2^{-1} - t_1^{-1}t_2^{-2}}{t_2 - 1} \doteq 1 + t_1 t_2$$

3.3 Symmetry of the multivariable Alexander polynomial via grid diagrams

Let $\vec{L} = L_1 \cup \cdots \cup L_l$ be an oriented l component link. The classical multivariable Alexander polynomial is known to satisfy certain symmetries upon inverting some of the variables, corresponding geometrically to reversing the orientation of some of the link components. These were first explained in the multivariable case by Torres [Tor53]. We examine an important symmetry through the grid diagram approach, inspired by material in the forthcoming book [OSS].

Theorem 3.7. $\Delta_L(t_1^{-1}, \ldots, t_l^{-1}) = \Delta_L(t_1, \ldots, t_l)$

To prove this, we modify a grid diagram G of $\vec{L}$ to obtain a grid diagram for $-\vec{L}$ two ways:

Lemma 3.8. *Let G be a grid diagram for $\vec{L}$. Then a grid diagram for $-\vec{L}$ may be obtained by switching all X and O markings in G. Alternately, a grid diagram for $-\vec{L}$ may be obtained by reflecting G about the downwards diagonal.*

Proof. The first statement is clear. For the second, let G' be the grid diagram obtained by reflecting G about the downward diagonal, and let P' be the link diagram obtained from G'. If we now rotate P' about the downward diagonal (thinking of P' embedded in $\mathbb{R}^3$), we get a new diagram of the same link, since this corresponds to an ambient isotopy. Now this diagram is the original link diagram with the orientations reversed. □

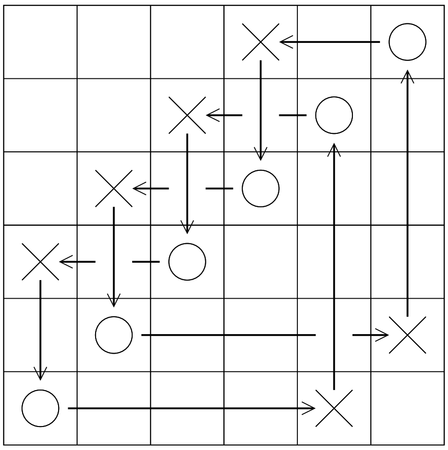

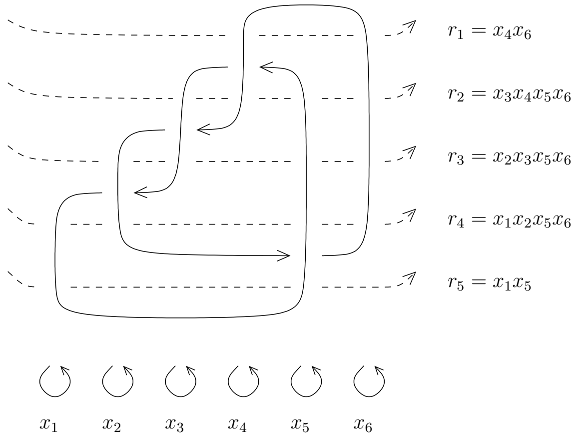

Calculate $\mathbf{M}(G)$:

$$\begin{pmatrix} 1 & 1 & 1 & 1 & 1 & 1 \\ 1 & 1 & 1 & 1 & t_2^{-1} & t_2^{-1} \\ 1 & 1 & 1 & t_1^{-1} & t_1^{-1}t_2^{-1} & t_2^{-1} \\ 1 & 1 & t_2^{-1} & t_1^{-1}t_2^{-1} & t_1^{-1}t_2^{-1} & t_2^{-1} \\ 1 & t_1^{-1} & t_1^{-1}t_2^{-1} & t_1^{-1}t_2^{-1} & t_1^{-1}t_2^{-1} & t_2^{-1} \\ 1 & t_1^{-1} & t_1^{-1} & t_1^{-1} & t_1^{-1} & 1 \end{pmatrix}$$

Subtract columns from each other and examine the lower right minor:

$$\begin{pmatrix} 0 & 0 & 0 & t_2-1 & 0 \\ 0 & 0 & t_1^{-1}-1 & t_1^{-1}(t_2^{-1}-1) & t_2^{-1}(1-t_1^{-1}) \\ 0 & t_2^{-1}-1 & t_2^{-1}(t_1^{-1}-1) & 0 & t_2^{-1}(1-t_1^{-1}) \\ t_1^{-1}-1 & t_1^{-1}(t_2^{-1}-1) & 0 & 0 & t_2^{-1}(1-t_1^{-1}) \\ t_1^{-1}-1 & 0 & 0 & 0 & 1-t_1^{-1} \end{pmatrix}$$

Factor:

$$\pm t_1^{-6}t_2^{-4}(t_1-1)^3(t_2-1)^2 \begin{pmatrix} 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 1 & t_1^{-1} & t_2^{-1} \\ 0 & 1 & t_2^{-1} & 0 & t_2^{-1} \\ 1 & t_1^{-1} & 0 & 0 & t_2^{-1} \\ 1 & 0 & 0 & 0 & 1 \end{pmatrix}$$

Take Fox derivatives:

$$\begin{pmatrix} 0 & 0 & 0 & 1 & 0 & x_4 \\ 0 & 0 & 1 & x_3 & x_3x_4 & x_3x_4x_5 \\ 0 & 1 & x_2 & 0 & x_2x_3 & x_2x_3x_5 \\ 1 & x_1 & 0 & 0 & x_1x_2 & x_1x_2x_5 \\ 1 & 0 & 0 & 0 & x_1 & 0 \end{pmatrix}$$

Abelianize and remove the last column:

$$\begin{pmatrix} 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 1 & t_1^{-1} & t_1^{-1}t_2^{-1} \\ 0 & 1 & t_2^{-1} & 0 & t_1^{-1}t_2^{-1} \\ 1 & t_1^{-1} & 0 & 0 & t_1^{-1}t_2^{-1} \\ 1 & 0 & 0 & 0 & t_1^{-1} \end{pmatrix}$$

Figure 3.5: Two calculations of the Alexander polynomial for the $(4, 2)$ torus link.

Proof of Theorem 3.7. Let G be an $n \times n$ grid diagram representing $\vec{L}$. First, let G_1 be the grid diagram representing $-\vec{L}$ by switching the O and X markings. We wish to compare the Alexander polynomial obtained from this diagram with the formal substitution $\Delta_{\vec{L}}(t_1^{-1}, \ldots, t_l^{-1})$. In the calculation below we write $\det \mathbf{M}(G_1)(t_1, \ldots, t_l)$ to emphasize the determinant is a polynomial in these variables.

$$\begin{aligned}\Delta_{-\vec{L}}(t_1, \ldots, t_l) &= \epsilon(G_1) \det \mathbf{M}(G_1)(t_1, \ldots, t_l) \prod_{i=1}^{l} (1 - t_i)^{-n_i(G_1)} t_i^{a_i(G_1) + \frac{n_i(G_1)}{2}} \\ &= \epsilon(G_1) \det \mathbf{M}(G)(t_1^{-1}, \ldots, t_l^{-1}) \prod_{i=1}^{l} (1 - t_i^{-1})^{-n_i} (-t_i^{-n_i}) t_i^{-a_i(G) + \frac{n_i(G)}{2}} \\ &= \epsilon(G_1)\epsilon(G)(-1)^n \Delta_{\vec{L}}(t_1^{-1}, \ldots, t_l^{-1})\end{aligned}$$

Note that the O markings of G_1 may be permuted to the O markings of G by applying n row transpositions (interchanging the O's and X's in each column). Hence

$$\Delta_{-\vec{L}}(t_1, \ldots, t_l) = \Delta_{\vec{L}}(t_1^{-1}, \ldots, t_l^{-1}) \tag{3.1}$$

Now let G_2 denote the grid diagram representing $-\vec{L}$ by reflecting G about the downwards diagonal. We calculate $\Delta_{-\vec{L}}$ using this diagram. The reflection induces a bijection between vertices of G and vertices of G_1, which we denote by $\phi : (i, j) \mapsto (j, i)$. We claim that ϕ preserves winding numbers: that is, $w_i(\phi(v)) = w_i(v)$ for each vertex v of G. To see this, we compute the winding number of v (of some L_i) using a horizontal ray pointing towards the left, and compare this to the winding number of $\phi(v)$ computed using a vertical ray pointing upwards in G_2. These rays intersect the same segments in the same orientation, hence ϕ preserves the winding numbers. Moreover, when calculating determinants, the sign of the minors of v and $\phi(v)$ are equal. Hence $\det \mathbf{M}(G) = \det \mathbf{M}(G_2)$, and furthermore:

$$\Delta_{\vec{L}} = \epsilon(G)\epsilon(G_2)\Delta_{-\vec{L}}$$

However, if we may permute the O markings of G to the downwards diagonal configuration using N column transpositions, then we may also permute the O markings of G_2 to the downwards diagonal configuration using N row transpositions. Hence $\Delta_{\vec{L}} = \Delta_{-\vec{L}}$. Comparing with Equation 3.1 gives us the desired equality. □

From the proof we also see that

Corollary 3.9. *$\Delta_{\vec{L}}$ does not detect the overall orientation of $\vec{L}$.*

Chapter 4

The Thurston norm

The Thurston norm is a function on the second (real) homology of a 3-manifold that enjoys several nice algebraic properties (which, although they are not enough to guarantee it is a norm, makes it a semi-norm). It measures the complexity of an embedded oriented surface representing a given homology class. Oriented surfaces are determined by their genus, so this notion of complexity is essentially a "minimal genus" function, i.e. it takes a homology class and returns the minimal genus of a surface representing it. We will see that the semi-norm vanishes exactly on the subspace spanned by integral classes represented by surfaces of non-negative Euler characteristic. Furthermore, the unit ball is seen to be a (possibly non-compact) polyhedron defined by finitely many linear equalities. The Thurston norm is completely determined by the degenerate space and the compact polyhedron which is the unit ball of the norm obtained by descending to the quotient by the degenerate subspace.

Although the Thurston norm is defined for all compact oriented 3-manifolds, we are most interested in link complements, and in this scenario the Thurston norm generalizes the knot genus. The knot genus is the minimal genus of a Seifert surface, an embedded oriented surface whose boundary is the given knot. The knot genus, among many other uses, detects the unknot, a surprisingly difficult task. Alexander noticed that the degree of the single-variable Alexander polynomial (which is the difference of the highest and lowest degrees in a Laurent polynomial) provides a lower bound for the knot genus. Analogously, McMullen proved that the Thurston norm is bounded below by the Alexander norm, a number easily calculable from the multivariable Alexander polynomial that is related to the degree. We shall discuss this below in detail.

Aside from its definition, the Thurston norm encodes some geometric information about the link complement. For example, in section 4.4 we will see that all homology classes which are represented as the fiber of a fibration of the link complement over S^1 lie in the cone of one of a selected number of top-dimensional faces of the unit ball, called fibered faces. Conversely, any class lying in the cone of a fibered face corresponds to a fiber of some fibration. These "fibered classes" are interesting because it allows us to describe the 3-manifold structure of the link complement through lower-dimensional objects, namely the fiber surface and the circle.

We will not mention the myriad applications of the Thurston norm provided in the last 20 years. Suffice to say, the Thurston norm is a prevalent tool in low-dimensional topology and is routinely compared to other 3-manifold invariants which aim to measure topological complexity.

4.1 Existence and shape of the unit ball

Proposition 4.1. *[Thu86] Let M be a compact oriented 3-manifold. Then every element in $H_2(M, \partial M; \mathbb{Z})$ is represented by the fundamental class of an embedded oriented surface S. If a is divisible by $k \in \mathbb{N}$, then S is a union of k subsurfaces, each representing a/k.*

Proof. (Cf. [Bre93], Theorem VI.11.16.) Suppose a is an element of $H_2(M, \partial M; \mathbb{Z})$ and let $D_M(a) \in H^1(M)$ denote its Poincaré dual. Using either Hopf's Theorem or the fact that S^1 is a $K(\mathbb{Z}, 1)$ space, there exists a continuous $f_a : M \to S^1$, unique up to homotopy, such that $f_a^*(u) = D_M(a)$, where u generates $H^1(S^1)$. By homotopy invariance we may assume f_a is smooth. Let $x \in S^1$ be any regular value for f_a. We claim that $N = f_a^{-1}(x)$ works. Thus is a codimension 1 framed submanifold of an oriented 3-manifold (the framing being given by u and f_a), hence it is an embedded oriented surface.

We must show $[N] = a$, where $[N]$ is the (inclusion of the) fundamental class of N. It suffices to show $D_M([N]) = D_M(a)$. Now u is the Thom class of the normal bundle of $\{x\}$ in S^1. The normal bundle of N is induced by f_a, so its Thom class is exactly $f_a^*(u) = D_M(a)$. But the Thom class of the normal bundle of N is the Poincaré dual of $[N]$ as required.

For the second part, suppose $a = kb \in H_2(M, \partial M)$ with corresponding functions f_a and f_b. Let $p : S^1 \to S^1$ be a k-fold covering map. Then

$$f_a^*(u) = D_M(a) = kD_M(b) = f_b^*(ku) = (p \circ f_b)^*(u)$$

By homotopy uniqueness, f_a and $p \circ f_b$ are homotopic. By the covering homotopy property, we may homotope f_b so that $f_a = p \circ f_b$. Then the k preimages $x_1, \ldots, x_k$ of x are regular values of f_b. Then $f_a^{-1}(x) = f_b^{-1}(x_1) \cup \cdots \cup f_b^{-1}(x_k)$ which is a disjoint union of surfaces, each representing b. □

Remark. We note that we never needed $\dim M = 3$ except when we used the word "surface", and in fact any codimension 1 homology class in a smooth manifold is represented by an embedded submanifold. The same is true for codimension 2, which can be seen by the same argument using that $\mathbb{C}P^\infty$ is a $K(\mathbb{Z}, 2)$ space and also a smooth manifold. In general it is not true. Inquiry of this sort originated with the work of Poincaré and later Thom, and remains interesting to this day.

For a connected surface S, let $\chi_-(S)$ denote the negative part of the Euler characteristic:

$$\chi_-(S) = \max(0, -\chi(S))$$

We sometimes call $\chi_-(S)$ the complexity of S. For a non-connected surface, $\chi_-(S)$ denotes the sum of the negative parts of the Euler characteristic of each connected component. Part of Thurston's ingenuity lies in "forgetting" low genus surfaces: note that the surfaces for which $\chi_-(S) \neq -\chi(S)$ are exactly spheres and discs. Clearly $\chi_-(S_1 \# S_2) \geq \chi_-(S_1) + \chi_-(S_2)$. First we define the Thurston norm on the integer homology.

Definition 4.2. Let $(M, \partial M)$ be a compact, oriented 3-manifold. The (integral) **Thurston norm** is a function $x_T : H_2(M, \partial M; \mathbb{Z}) \to \mathbb{Z}_{\geq 0}$ defined by

$$x_T(a) = \inf\{\chi_-(S) \mid S \text{ is an embedded surface representing } a\}$$

Theorem 4.3. *[Thu86] The integral Thurston norm satisfies $x_T(ka) = |k|x_T(a)$ and $x_T(a+b) \leq x_T(a) + x_T(b)$ for all $a, b \in H_2(M, \partial M; \mathbb{Z})$ and $k \in \mathbb{Z}$.*

Proof. First, note that $x_T(-a) = x_T(a)$, since if we are given a surface representing a, we may represent $-a$ with the same surface but with reversed orientation. Now the second part of the previous proposition shows $x_T(ka) \geq |k|x_T(a)$. On the other hand, if S represents a, then we may take $|k|$ parallel, pairwise disjoint copies of S (with the orientations reversed if k is negative) to represent ka, so $x_T(ka) \leq |k|x_T(a)$, which establishes the first claim.

Now take $a, b \in H_2(M, \partial M)$ and let S_a and S_b be surfaces representing these classes such that $\chi_-(S_a) = x_T(a)$ and $\chi_-(S_b) = x_T(b)$. We may assume they intersect transversely. Now $S_a \cup S_b$ represents the class $a + b$, which can be seen by considering triangulations of either surface. We wish to perform surgery on $S_a \cup S_b$ to end up with an embedded surface representing the same homology class whose complexity is equal to $\chi_-(S_a) + \chi_-(S_b)$. This is sufficient to show $x_T(a + b) \leq x_T(a) + x_T(b)$.

The intersection $S_a \cap S_b$ consists of finitely many circles and arcs. First we eliminate those circles which bound a disc on either surface and those arcs which are homotopic (rel. endpoints) to a portion of a boundary component. Suppose $C \subseteq S_a \cap S_b$ is a circle bounding a disc on, say, S_a. If this disc contains another circle in the intersection, then we consider that circle instead. Therefore we may suppose C is an "innermost" circle that bounds a disc on S_a. Then we perform surgery on S_b by removing a small tubular neighborhood of C. This operation preserves the homology class of the total surface (since we have modified it by the boundary of a cylinder). Furthermore this operation doesn't change $\chi_-(S_b)$, since it certainly cannot increase the complexity, and it cannot decrease it because S_b was chosen with minimal complexity. Repeating this step we may assume that no component of $S_a \cap S_b$ bounds a disc on either surface.

Now suppose $I \subseteq S_a \cap S_b$ is homotopic rel. endpoints to a portion of the boundary of, say, S_a. This homotopy covers a disc in S_a. As before, we may assume that no other component of $S_a \cap S_b$ lies in this disc (it cannot contain a circle because that circle would bound a disc, and we have eliminated these already). We modify S_b to S_b' by removing a tubular neighborhood of I and attaching two disks parallel to the disc on S_a. This does not alter the homology class of S_b since we have modified it by the boundary of a cylinder (modulo ∂M). To check that the complexity does not increase, we note that the modification increases the Euler characteristic by 1, and here the complexity agrees with the negative of the Euler characteristic unless one of the components of S_b' is a disc. But then the other component is clearly homeomorphic to S_b and we have $\chi_-(S_b') = \chi_-(S_b)$ as needed.

Thus we assume S_a and S_b have only essential intersections (in the sense described above), so none of the components with positive Euler characteristic intersect (the spheres or discs). At every component of the intersection we perform an oriented sum operation to form a new oriented embedded surface S. The local picture for this operation is shown in Figure 4.1. The new surface represents the same homology class as $S_a \cup S_b$ since these surfaces differ by the boundary of a collection of wedges.

We have $\chi(S) = \chi(S_a) + \chi(S_b)$, since the oriented sum operation involves the same type of cutting and pasting that is required to dissemble $S_a \cup S_b$ into a disjoint union. Also, none of the components of S with positive Euler characteristic were involved with in oriented sum operation, since otherwise it would have involved a component of S_a or S_b which also has positive Euler characteristic, by the additivity of the Euler characteristic with respect to the oriented sum operation. By our previous surgeries,

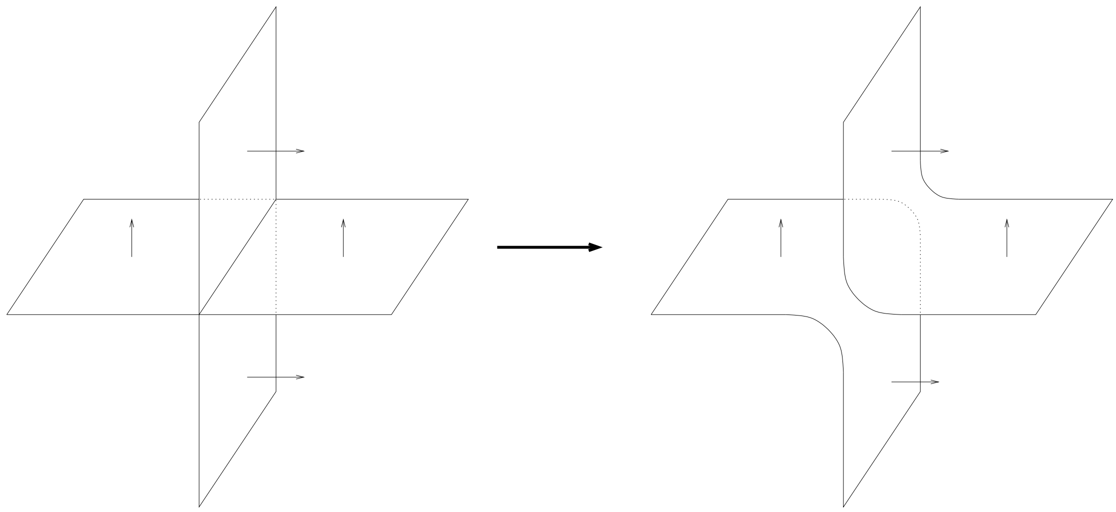

Figure 4.1: The oriented sum operation.

this is impossible. Therefore in the equation $\chi(S) = \chi(S_a) + \chi(S_b)$ we may cancel from both sides the contributions from spheres and discs to achieve the equality $\chi_-(S) = \chi_-(S_a) + \chi_-(S_b)$. □

Next we extend this norm to $H_2(M, \partial M; \mathbb{R})$. First, we define x_T on the rational subspace $H_2(M, \partial M; \mathbb{Q})$ by linearity on the rays. For $s, t \in \mathbb{Q}$ and $a, b \in H_2(M, \partial M; \mathbb{Q})$ we have

$$x_T(sa + tb) \leq |s|x_T(a) + |t|x_T(b)$$

by the previous theorem. Since the rational points are dense in $H_2(M, \partial M; \mathbb{R})$, there is at most one continuous extension of x_T. Using the equation above it is clear that the continuous extension exists (constructed, for example, through Cauchy sequences) and is a non-negative convex function linear on rays through the origin. In other words, x_T is a semi-norm, the only condition preventing it from being a norm is that it may vanish on some subspace. This subspace is addressed as follows:

Proposition 4.4. *The semi-norm x_T on $H_2(M, \partial M; \mathbb{R})$ vanishes exactly on the subspace spanned by those points in the integer lattice which are represented by surfaces of non-negative Euler characteristic.*

Proof. Let K_x denote the set on which x_T vanishes. By linearity K_x is a linear subspace. It is clear that the subspace spanned by such points is contained in K_x. For the converse, suppose $x_T(v) = 0$ for some $v \in H_2(M, \partial M; \mathbb{R})$. Then x_T vanishes on the entire ray through v, and this ray passes arbitrarily close to some integer lattice points. But if for some $s \in \mathbb{R}$, the point sv is arbitrarily close to an integer lattice point z, we must have that the integer $x_T(z) = 0$ by continuity, which means that z is represented by a surface of non-negative Euler characteristic. Therefore every element of K_x is approximated by such integer lattice points, so K_x is contained in the closure of the linear span of those points, and the conclusion follows. □

Now note that by linearity and convexity, x_T is constant on cosets of K_x. Therefore x_T descends to a function on $H_2(M, \partial M; \mathbb{R})/K_x$ that is convex, linear on rays, and vanishes only at the origin. Therefore it is a norm on this space. In particular, if every non-trivial homology class is represented by a surface of negative Euler characteristic, then x_T is a norm on $H_2(M, \partial M; \mathbb{R})$.

Associated to any norm or semi-norm is a unit ball, which may be any convex body symmetric about the origin. By linearity, the unit ball describes the (semi-)norm exactly. It turns out that the unit ball of the Thurston norm is a polyhedron with vertices in the integer lattice. This follows from the purely formal result on norms on vector spaces which take on integral values on a canonical integer lattice:

Theorem 4.5. *[Cal07], [Thu86]. Suppose x is a norm on a finite-dimensional vector space H, where H contains a canonical $\mathbb{Z}$ lattice and on this lattice x takes integral values. Then the unit ball of x is a compact polyhedron.*

When x_T is not a norm, then the unit ball is not compact, but it does have the form $P \times K_x$ where P is a compact polyhedron of codimension $\dim K_x$, by our observation above that x_T descends to a norm on the cosets of K_x.

4.2 The Alexander norm on cohomology

The Alexander norm on $H^1(M)$ is a multivariable generalization of the degree of the Alexander polynomial. The Alexander polynomial may be written as

$$\Delta(L) \doteq \sum \lambda_\alpha t^\alpha$$

where $\alpha = (\alpha_1, \ldots, \alpha_l)$ is a multi-index and $\lambda_\alpha \in \mathbb{Z}$. Since each t_i corresponds to a homology class, each multi-index α may be considered as an element of $H_1(M; \mathbb{R})$, namely $\alpha = \alpha_1 t_1 + \cdots + \alpha_l t_l$.

Definition 4.6. The **Newton polytope** $N(\Delta_L)$ is the convex hull in $H_1(M; \mathbb{R})$ of the points α, ranging over the multi-indices that appear in $\Delta(L)$.

The Newton polytope is only defined up to translation. By Theorem 3.7 the Alexander polynomial is symmetric, so the Newton polytope may be translated to be symmetric about the origin.

We now define the Alexander norm on $H^1(M)$. Since $H_0(M; \mathbb{Z})$ is free abelian, we have by the universal coefficient theorem $H^1(M; \mathbb{Z}) = \mathrm{Hom}(H_1(M; \mathbb{Z}), \mathbb{Z}) = \mathrm{Hom}(\pi_1(M), \mathbb{Z})$ where the second equality follows from $H_1(M; \mathbb{Z}) = \mathrm{ab}(\pi_1(M))$.

Definition 4.7. The **Alexander norm** on $H^1(M; \mathbb{Z})$ or $H^1(M; \mathbb{R})$ is given by

$$\|\phi\|_A = \sup \phi(\alpha - \beta)$$

where α and β are multi-indices appearing in Δ.

The Alexander norm is clearly linear on rays and convex, making it a semi-norm. It also takes integral values on the canonical integral lattice, so by Theorem 4.5, the unit ball is a (possibly non-compact) polyhedron. An equivalent formulation of the norm is

$$\|\phi\|_A = \mathrm{length}(\phi(N(\Delta)))$$

In other words, the Alexander norm of a cohomology class ϕ is the length of the projection of the Newton polytope under ϕ. Clearly then the unit ball of the Alexander norm is dual to the Newton polytope with a scaling factor of 1/2.

4.3 The Alexander norm bounds the Thurston norm

We now present McMullen's argument that $\|\phi\|_A \leq x_T(\phi)$ [McM02]. We've implcitly transferred the Thurston norm to $H^1(M)$ by Poincaré duality. The ingenuity in McMullen's approach is to relate the Alexander norm to the first Betti numbers of free abelian covers corresponding to surjective homomorphisms $\phi : \pi_1(X) \to \mathbb{Z}$. The point is that we may use single-variable techniques to analyse this cover. As usual, we consider a bounded link complement M with $G = \pi_1(M)$.

The **degree** of a non-zero single-variable Laurent polynomial is the difference between the highest and lowest exponents, and the degree of 0 is ∞.

Definition 4.8. A cohomology class $\phi \in H^1(M;\mathbb{Z})$ is **primitive** if ϕ is surjective as a homomorphism from $H_1(M;\mathbb{Z}) \to \mathbb{Z}$.

For primitive classes, there is an associated free abelian cover $X_\phi \to X$. We can extract some information about the space X_ϕ by considering the Alexander polynomial corresponding to ϕ:

Proposition 4.9. *Suppose $\phi \in H^1(M;\mathbb{Z})$ is a primitive class and let $X_\phi \to X$ be the corresponding free abelian cover. Then*

$$b_1(X_\phi) = \deg \Delta_\phi$$

Proof. We are considering the $\mathbb{Z}[\mathbb{Z}]$-module $H_1(X_\phi;\mathbb{Z})$. The ring $\Lambda = \mathbb{Z}[\mathbb{Z}]$ is a PID and $H_1(X_\phi)$ is finitely generated over Λ, so

$$H_1(X_\phi) = \Lambda/(f_1) \oplus \cdots \oplus \Lambda/(f_n)$$

for some polynomials satisfying $f_1 \cdots f_n = \Delta_\phi$ (see Example 2.5). Note that $\deg \Delta_\phi = \deg f_1 + \cdots + \deg f_n$. Now we calculate:

$$\begin{aligned} b_1(X_\phi) &= \dim(H_1(X_\phi) \otimes \mathbb{R}) = (\Lambda/(f_1) \oplus \cdots \oplus \Lambda/(f_n)) \otimes \mathbb{R} \\ &= (\Lambda/(f_1) \otimes \mathbb{R}) \oplus \cdots \oplus (\Lambda/(f_n) \otimes \mathbb{R}) \\ &= \mathbb{R}[\mathbb{Z}]/(f_1) \oplus \cdots \oplus \mathbb{R}[\mathbb{Z}]/(f_n) \end{aligned}$$

Since $\mathbb{R}[\mathbb{Z}]$ is a field (the field of single variable Laurent polynomials over a field), the $\mathbb{R}$-vector space $\mathbb{R}[\mathbb{Z}]/(f)$ has dimension $\deg f$. Therefore

$$b_1(X_\phi) = \deg f_1 + \cdots + \deg f_n = \deg \Delta_\phi$$

as desired. □

Theorem 4.10. *([McM02], Theorem 4.1.) Suppose ϕ is a primitive class contained in the cone of an open face in the unit ball of the Alexander norm. Then*

$$b_1(X_\phi) = \|\phi\|_A + 1$$

Proof. Consider the map $\phi_* : \mathbb{Z}[\mathbb{Z}^l] \to \Lambda = \mathbb{Z}[t, t^{-1}]$ which is the linear extension of $\phi : G \to \mathbb{Z}$. When we apply it to the Alexander polynomial $\Delta \in \mathbb{Z}[\mathbb{Z}^l]$ we get the sum

$$\phi_*(\Delta) = \sum \lambda_\alpha t^{\phi(\alpha)}$$

That is, we apply ϕ to each multi-index $\alpha \in \mathbb{Z}[\mathbb{Z}^l]$ to find the exponents of t. Therefore, if the largest and smallest values of $\phi(\alpha)$ appear in this sum (with non-zero coefficient), we have

$$\|\phi\|_A = \deg \phi_*(\Delta)$$

This occurs, for example, under our assumption that ϕ is contained in the cone of an open face of the norm ball, since in this case the largest and smallest values of $\phi(\alpha)$ appear exactly once (so they cannot be cancelled away in the sum). (Recall that $\|\phi|_A$ is the length of the image of the Newton polytope under ϕ. If ϕ is contained in the cone of an open face of the norm ball, then the vertices of $N(L)$ all have distinct images under ϕ.)

Next we relate $\phi_*(\Delta)$ and Δ_ϕ. By Theorem 2.25 the Alexander ideal $I = \epsilon_{\Lambda_l} \cdot (\Delta)$ and by Corollary 2.22 $I_\phi = \phi_*(I)$. Therefore:

$$(\Delta_\phi) = (I_\phi) = \phi_*(I) = \phi_*(\epsilon_{\Lambda_l} \cdot (\Delta)) = ((t-1)\phi_*(\Delta))$$

where we used the relation $\phi_*(\epsilon_{\Lambda_l}) = (t-1)$, which holds since if ϕ is primitive then there are t_i and t_j such that $\phi(t_i) = t^{k_i}$ and $\phi(t_j) = t^{k_j}$ where k_i and k_j are relatively prime, so then $\gcd(t^{k_i}-1, t^{k_j}-1) = t-1$. Therefore the polynomials Δ_ϕ and $(t-1)\phi_*(\Delta)$ are the same up to multiplication by a unit, so their degrees agree. Finally, applying the previous proposition yields

$$b_1(X_\phi) = \deg \Delta_\phi = \deg\ (t-1)\phi_*(\Delta) = 1 + \deg \phi_*(\Delta)$$

so $b_1(X_\phi) = 1 + \|\phi\|_A$ by our previous comments. □

Next we would like to relate $b_1(X_\phi)$ to the Thurston norm of ϕ. To that end, we refine the construction described in the proof of the existence of the Thurston norm to show that for certain classes, a Thurston norm-minimizing surface may be found that has particular constraints on its Betti numbers.

Theorem 4.11. *([McM02], Proposition 6.1.) Let ϕ be a primitive class such that $b_1(X_\phi)$ is finite. Then there exists a Thurston norm-minimizing surface $S \subset M$ Poincaré dual to ϕ that is connected.*

Proof. Let S be a surface with $\chi_-(S) = x_T(\phi)$ that minimizes b_0 over all possible surfaces. We want to show S is connected. If S' is a component of S, then a tubular neighborhood $\nu(S')$ intersects at most 2 components of $M - S$ (since $\nu(S') \cap (M-S) = (S')^+ \cup (S')^-$ is a decomposition into connected components). We define an oriented graph Γ whose vertices are the components of $M - S$ and there is an edge between components M_i and M_j if they are joined by a component of S. The edges are oriented by the orientations of S and M (for example, an edge corresponding to a component S' points towards the component of M containing $(S')^+$).

There is a continuous map $M \xrightarrow{f} \Gamma$ which sends the tubular neighborhoods of each component of S to the edges, and the remainder of M is sent to the vertices. On the other hand, by selecting points in each component of $M - S$ and paths between them intersecting every dividing component of S once, we can embed $\Gamma \hookrightarrow M$. Then the composition

$$\Gamma \hookrightarrow M \to \Gamma$$

is homotopy equivalent to the identity. Furthermore, there is a natural map

$$\Gamma \xrightarrow{g} S^1$$

which sends each directed edge positively once around S^1. This map is such that ϕ is the pullback of the generator of $H^1(S^1)$ under the composition

$$M \to \Gamma \to S^1 \tag{4.1}$$

To see this, note that a generator of $H^1(S^1)$ (call it ω) is represented by a bump 1-form with total integral 1 in de Rham cohomology. The pullback of ω to M consists of bump 1-forms on fibers of the tubular neighborhood of S. This is exactly the representative of the Poincaré dual of S described in Section 1.3.

Next we claim that $b_1(\Gamma) = 1$, or in other words Γ contains a single cycle. Suppose first that $b_1(\Gamma) > 1$. To find a contradiction, pullback Equation (4.1) to the universal cover of S^1 to obtain a sequence of covering spaces each with fiber $\mathbb{Z}$:

$$\begin{array}{ccccc} \tilde{M} & \longrightarrow & \tilde{\Gamma} & \longrightarrow & \mathbb{R} \\ {\scriptstyle \pi_M}\big\downarrow & & {\scriptstyle \pi_\Gamma}\big\downarrow & & {\scriptstyle \pi}\big\downarrow \\ M & \xrightarrow{f} & \Gamma & \xrightarrow{g} & S^1 \end{array}$$

In fact the covering space $\tilde{M} \xrightarrow{\pi_M} M$ is equivalent to $M_\phi \to M$. To check this, it is enough to show that $\gamma \in \pi_1(M)$ acts as $1 \in \mathbb{Z}$ (considered as a deck transformation) if and only if $\phi(\gamma) = 1$, as in Proposition 2.18. By the construction of the pullback bundle, γ acts as 1 iff $f \circ g \circ \gamma$ acts as 1 on the covering space $\mathbb{R} \xrightarrow{\pi} S^1$, which occurs iff $\omega(f \circ g \circ \gamma) = 1$. But ϕ is the pullback of ω under $f \circ g$, so this is equivalent to $\phi(\gamma = 1)$.

Now the composition $\Gamma \to M \to \Gamma$ lifts to the covering spaces to a map $\tilde{\Gamma} \to M_\phi \to \tilde{\Gamma}$ also homotopic to the identity. Therefore, applying the functor $H^1(-, \mathbb{R})$ we have $b_1(M_\phi) \geq b_1(\tilde{\Gamma})$. If $b_1(\Gamma) > 1$, then $\tilde{\Gamma}$ has infinitely many loops (the covering action is $\mathbb{Z}$, so it can "kill" at most one cycle of Γ). Then $b_1(\tilde{\Gamma}) = \infty$, which contradicts our assumption that $b_1(M_\phi)$ is finite.

We conclude that there is at most one cycle in Γ, so either Γ is a tree, or a cycle with trees coming off the vertices of that cycle. However, there cannot be any vertex of degree 1, since this mean that there is a component of S lying in the boundary of M; hence it is homologically trivial (since we are working relative to ∂M) so we may exclude it from S, but this contradicts the minimality of $b_0(S)$. Therefore Γ is a cycle and in particular every vertex has degree 2, ignoring orientations. Continuing, note that there cannot be a vertex with two edges pointing towards it, since this means there is a component of M whose boundary consists of components of S, say S_i and S_j. This means that $[S_i] + [S_j]$ is trivial in homology, so we may exclude these from S, again contradicting the minimality of $b_0(S)$.

Therefore Γ is an oriented cycle with $b_0(S)$ many vertices, so the map $\Gamma \to S^1$ is a covering map of degree $b_0(S)$. But since ϕ is primitive, we must have $b_0(S) = 1$ and S is connected. □

Now that we've established that there is a connected representative surface realizing the Thurston norm, we can examine the cover M_ϕ via the geometric construction described in Section 2.5. It is clear from our earlier propositions we wish to relate $b_1(M_\phi)$ to the Thurston norm of ϕ since Theorem 4.10 connects $b_1(M_\phi)$ with the Alexander norm of ϕ. First a general lemma.

Lemma 4.12. *Suppose $X = A \cup B$ is an open cover of a topological space such that $H_i(X)$ is generated both by $H_i(A)$ and $H_i(B)$ for some $i \in \mathbb{N}$. Then $H_i(X)$ is generated by $H_i(A \cap B)$.*

Proof. Consider the Mayer-Vietoris sequence for $X = A \cup B$. A portion of it looks like

$$\to H_i(A \cap B) \xrightarrow{(c,-c)} H_i(A) \oplus H_1(B) \xrightarrow{(a+b)} H_i(X) \to$$

By assumption $H_i(A) \cong H_i(B) \cong H_i(X)$. Therefore, for each $x \in H_i(X)$, the pair $(x, -x) \in H_i(A) \oplus H_i(B)$. This element is in the kernel of the sum map on the right, so by exactness is in the image of the difference map on the left. Therefore each $x \in X$ is the image of an element of $H_i(A \cap B)$ under inclusion. □

Proposition 4.13. *Let ϕ be a primitive class such that $b_1(M_\phi)$ is finite. Then there exists an embedded oriented surface S Poincaré dual to ϕ minimizing the Thurston norm such that*

$$\begin{aligned} &b_0(S) = 1 && \textit{(S is connected)} \\ &b_1(S) \geq b_1(M_\phi) && \\ &b_2(S) = 0 && \textit{(S has non-empty boundary)} \end{aligned}$$

Proof. The previous theorem allows us to pick a connected surface S. Let M_ϕ be written as $\bigcup_{i\in\mathbb{Z}} N_i$ where $N_i \cap N_{i+1} = S$ (see Section 2.5). Now the first homology of M_ϕ is finitely generated by assumption. Therefore, it is generated by some compact subspace which we label $N_0 \cup \cdots \cup N_k$. Now the deck transformation generating $\mathbb{Z}$ induces an automorphism of $H_1(M_\phi)$ (since it is invertible) and acts via translation on the N_i's. Therefore the subspace $N_{-k-1} \cup \cdots \cup N_{-1}$ also generates $H_1(M_\phi)$. By the lemma, $H_1(S)$ generates $H_1(M_\phi)$ and consequently $b_1(S) \geq b_1(M_\phi)$.

Clearly S must have boundary, since if it did not it would not intersect any of the meridians of the boundary tori, which implies $\phi = 0$ by Proposition 1.12. □

We are ready to prove the main result, that the Alexander norm provides a lower bound for the Thurston norm.

Theorem 4.14. *([McM02], Theorem 1.1) Let $\phi \in H^1(M;\mathbb{R})$. Then $\|\phi\|_A \leq x_T(\phi)$.*

Proof. First we may assume $\Delta \neq 0$ since then the Alexander norm is 0. It also suffices to prove the inequality for primitive classes contained in the cone of an open face of the Alexander unit ball, since the Alexander and Thurston norms are both linear and continuous (by linearity this proves the inequality for all classes outside a finite set of hyperplanes, so by continuity the inequality is satisfied everywhere).

For such a class we have $b_1(M_\phi) = \|\phi\|_A + 1$ by Theorem 4.10. Let S be a Thurston norm minimizing surface provided in the previous proposition. If S has positive Euler characteristic (i.e. it is a disk), then these equations imply $\|\phi\|_A = 0 = \chi_-(S) = x_T(\phi)$, so suppose $\chi(S) \leq 0$. Then we have:

$$\begin{aligned} x_T(\phi) = \chi_-(S) &= -\chi(S) = -b_2(S) + b_1(S) - b_0(S) = b_1(S) - 1 \\ &\geq b_1(M_\phi) - 1 = \|\phi\|_A \end{aligned}$$

□

4.4 Fibered classes

We consider now fibrations of the link complement over S^1. Any such fibration may be associated with a cohomology class $\alpha \in H^1(M)$ which is the pullback of a generator of $H^1(S^1)$. It turns out that all such cohomology classes with unit Thurston norm belong to certain top-dimensional faces of the Thurston unit ball, called fibered faces. On these faces, the Alexander and Thurston norm agree.

It is useful to study knots and links whose complements admit fibrations since fibered manifolds are described by lower dimensional submanifolds (the fibers) and have more easily understandable topology than a generic 3-manifold.

Definition 4.15. A **fibered 3-manifold** M is an oriented 3-manifold with the structure of a smooth fiber bundle over S^1. Equivalently (if M is compact), M is equipped with a surjective submersion $\pi : M \to S^1$. The inverse image of any point is a **fiber** of M.

Any two fibers are diffeomorphic compact oriented surfaces since locally we may project fibers onto each other and these projections are diffeomorphisms. Furthermore, any two fibers are homologous since they are the boundary of the inverse image of an interval in S^1. Thus it makes sense to speak of the homology class of the fiber. The Poincaré dual of this class is an element of $H^1(M)$ and is the pullback of a generator of $H^1(S^1)$.

Definition 4.16. A class $\phi \in H^1(M)$ is a **fibered class** if there exists a fiber bundle $\pi : M \to S^1$ such that ϕ is the pullback under π of the generator of $H^1(S^1)$.

Equivalently, as maps $H_1(M) \to H_1(S^1)$ we have $\pi_* = \phi$.

Example 4.17. Consider the complement of the unknot. Then M is a solid torus equipped with a trivial fibration $D^2 \times S^1 \to S^1$ whose fibers are discs. The pullback of a generator of $H^1(S^1)$ is dual to a longitude of the M, which coincides with a meridian of the boundary torus, which is a canonical generator of $H^1(M)$.

When M is the complement of a knot K that fibers over the circle such that the corresponding cohomology class is a generator of $H^1(M) \cong \mathbb{Z}$, we say that K is a **fibered knot**.

If we pullback the map $\pi : M \to S^1$ of a fibered 3-manifold with fiber F by the universal cover $\mathbb{R} \to S^1$ we obtain a covering map $\tilde{M} \to M$ with group of deck transformations $\mathbb{Z}$. Here $\tilde{M}$ is also a fiber bundle over $\mathbb{R}$, which is contractible, hence $\tilde{M}$ is trivial with fiber F. The space $\tilde{M}$ is constructed by gluing together countably many copies of $F \times [0, 2\pi]$ according to some homeomorphism h of F (cf. Section 2.5). Note that since $\tilde{M}$ is a trivial fibration, it has the homotopy type of the fiber, and in particular, its homology is the same as that of the fiber. The map h is understood by noting that if we remove a fiber of M, we are left with a trivial fibration $M - F \cong F \times [0, 2\pi)$. Then M can be recovered by taking the quotient of $F \times [0, 2\pi]$ by identifying $F \times 0$ and $F \times 1$ via the map h.

Remark. This construction is generally known as a mapping torus. It is well known that the homeomorphism type of a mapping torus constructed from an orientation preserving diffeomorphism h depends only on the isotopy class of h through orientation preserving diffeomorphisms. A dependency of this type means that the topology of a mapping torus depends only on the **mapping class** of h. The isotopy classes of orientation preserving diffeomorphisms of a surface form a group called the **mapping class group** that is a popular object of study (see for example [FM11]).

Definition 4.18. The map $h : F \to F$ is the **monodromy** of the fiber bundle. The induced map $h_* : H_1(F) \to H_1(F)$ is the **homological monodromy**.

Now suppose $\phi \in H^1(M)$ is a primitive fibered class obtained from the fiber bundle $\pi : M \to S^1$ with fiber F and homological monodromy h_*. We claim that the space $\tilde{M}$ obtained by gluing countably many copies of $F \times [0, 2\pi]$ together via h is the same as the $\mathbb{Z}$-covering space M_ϕ. Since M is given by taking the product $F \times [0, 2\pi]$ and identifying the fibers $F \times \{0\}$ and $F \times \{1\}$ via h, it is clear that the construction of M_ϕ given in Section 2.5 coincides with that of M_ϕ. As a corollary, we have

Proposition 4.19. *For a fibered class ϕ, the cover M_ϕ has the homotopy type of a fiber.*

The following statements generalize the well-known fact that the Alexander polynomial of a fibered knot is monic.

Proposition 4.20. *Suppose M is a link complement and $\phi \in H^1(M; \mathbb{Z})$ is a primitive fibered class obtained from the fiber bundle $\pi : M \to S^1$ with fiber F and homological monodromy h_*. Pick a basis of $H^1(F)$ and let M be the (integral) matrix of h_* in this basis. Let t be the generator of $\mathbb{Z}$ corresponding to the deck transformation describing h_*. Then the matrix $M - tI$ is a presentation matrix for the Alexander invariant $H^1(M_\phi; \mathbb{Z})$.*

Proof. As we've seen, M_ϕ is obtained by gluing together countably many copies of $F \times [0, 2\pi]$ via h. Furthermore, the homology of M_ϕ is just that of the fiber; in particular, $H_1(M_\phi; \mathbb{Z})$ is finitely generated over $\mathbb{Z}$. As in Proposition 4.13, by Lemma 4.12 we have that $H_1(F; \mathbb{Z})$ generates $H_1(M_\phi; \mathbb{Z})$. As a $\mathbb{Z}[t, t^{-1}]$ module, $H_1(M_\phi; \mathbb{Z})$ is completely described by how t acts on generators. Therefore, if $\alpha_1, \ldots, \alpha_n$ generate $H_1(F)$, then a full set of relations is given by

$$t\alpha_i = M\alpha_i$$

since h_* describes how t maps the homology of one fiber into another. Consequently, the matrix $M - tI$ presents $H_1(M_\phi; \mathbb{Z})$. □

Corollary 4.21. *The Alexander polynomial corresponding to a primitive fibered class is monic (the coefficient of the largest power of t is a unit in $\mathbb{Z}$).*

Proof. Since $M - tI$ is a square matrix, the Alexander polynomial is simply the determinant, which is clearly monic. In fact this is the characteristic polynomial of M. □

We are about to relate fibered class to the Thurston norm ball. The Thurston norm is just the minimum negative Euler characteristic over surfaces representing a given class, unless such a class is represented by a sphere or disc. However, this irregularity does not appear for fibered classes of link complements, due to the following observation:

Proposition 4.22. *Suppose ϕ is a fibered class of a link complement M corresponding to a fibration $M \to S^1$ with disk fibers. Then M is homeomorphic to a solid torus $D^2 \times S^1$.*

Proof. It is well known that any orientation preserving diffeomorphism of a disk is smoothly isotopic to the identity map. Applying this to the monodromy of a fibration, we see that M is homeomorphic to the solid torus (see the remark just before Definition 4.18). □

The complement of a link is a solid torus if and only if the link is the (1-component) unknot. Therefore, we exclude this trivial case in the following discussion and always assume that the fiber of a fibration corresponding to a fibered class has non-positive Euler characteristic.

We now turn to the results of Thurston [Thu86] concerning fibered faces of the Thurston unit ball. Our exposition uses material from [Thu86], [Cal07] and [CC03]. As the treatment relies on some topological ideas outside the scope of this thesis, we offer only sketches and citations of some statements. We must introduce several new concepts, primarily that of a foliation, which generalizes smooth fiber bundles.

Definition 4.23. [Cal07] A (smooth codimension 1) **foliation** of a 3-manifold M consists of an open covering by 3-balls $\{U_i\}$ together with smooth trivializations of each U_i as a product

$$U_i = D^2 \times [0, 1]$$

such that on the overlap of two charts, the sets $D^2 \times \{\text{point}\}$ agree.

Informally, this definition states that locally M may be described as a stack of planes. A chosen plane near a point corresponds uniquely to planes in neighboring charts. A maximal path-connected union of such planes is a **leaf** of the foliation. Clearly M is the union of all the leaves. Each leaf is a (possibly non-compact) surface.

We make additional important assumptions on subsequent foliations. The first is that each foliation is **transversely oriented**, meaning there is a continuous choice of normal vector to a leaf at each point of M. Furthermore, we assume that all foliations meet the boundary of M transversely. It is clear that these conditions are met by a fibration corresponding to a non-zero fibered class.

Example 4.24. A smooth fibration $M \to S^1$ induces a foliation whose leaves are the preimages of points in S^1.

As for fibrations, it turns out that foliations which admit compact leaves of positive Euler characteristic must be trivial. This is a special case of a general principle called Reeb stability.

Theorem 4.25. *(Reeb Stability [CC03], [Cal07]). Suppose $\mathfrak{F}$ is a transversely oriented foliation of a connected oriented 3-manifold M. If one of the leaves L of $\mathfrak{F}$ is a disk or sphere, then M is homeomorphic to $L \times S^1$ and the foliation corresponds to the foliation induced by the trivial fibration $L \times S^1 \to S^1$.*

This means that we will consider only foliations whose compact leaves all have non-positive Euler characteristic. Furthermore, we will be concerned only with a certain class of foliations.

Definition 4.26. A **taut foliation** of M is a foliation such that for every leaf there is a closed curve in M intersecting the leaf transversely. If $\partial M \neq \emptyset$ then the foliation is taut if every leaf meets either a closed transverse curve or a transverse arc connecting components of ∂M.

This innocuous requirement is actually quite useful, since the existence of such transverse arcs allow one to modify M in various ways along the arc. It is actually the case (for M compact) that a taut foliation admits a single closed curve intersecting every leaf transversely ([Cal07], Lemma 4.26). Most importantly, we have

Proposition 4.27. *A foliation corresponding to a fibration $M \to S^1$ is taut.*

Proof. First pick a point p belonging to the fiber F and lift a path going around S^1 once to a path γ starting at p. The result is an arc transverse to each fiber excluding F. The endpoint of this path also belongs to this fiber. Now a tubular neighborhood of F looks like $F \times (-\epsilon, \epsilon)$. Now if we modify γ by connecting the points $\gamma \cap F \times \{-\epsilon\}$ and $\gamma \cap F \times \{\epsilon\}$ by a straight line, we are left with a closed curve transverse to every leaf. □

Given a foliation $\mathfrak{F}$, we may consider the rank-2 vector bundle of tangent planes to the leaves. This is a special case of a (rank-2) **distribution**, which consists of a smooth choice of planes in the tangent space at every point. Equivalently, a distribution is any rank-2 smooth subbundle of the tangent bundle. A number of examples of distributions come from **non-singular 1-forms**, which is a 1-form that vanishes nowhere. Then the pointwise kernel of such a 1-form induces a distribution. The question of whether a non-singular 1-form corresponds to a distribution induced by a foliation is answered by the following classical theorem of Frobenius:

Theorem 4.28. *(Frobenius) The distribution induced by a non-singular 1-form α corresponds to a foliation if and only if*

$$\alpha \wedge d\alpha = 0.$$

In particular, any closed non-singular 1-form defines a foliation on M. Conversely, since we assume foliations are transversely oriented, it is not hard to construct a closed 1-form whose kernel determines the distribution of a given foliation using a partition of unity. We can improve on this by noting that certain closed non-singular 1-forms actually correspond to fibrations over S^1.

Theorem 4.29. *(Tischler, [Tis70]) Suppose α is an integral closed non-singular 1-form. Then the distribution induced by α corresponds to a fibration $M \to S^1$.*

Proof. (Sketch.) Fix a point $p \in M$. Then the map

$$q \mapsto \int_\gamma \alpha \mod \mathbb{Z}$$

where γ is a path from p to q is a well-defined submersion $M \to S^1$ by the assumption that α is integral. Local examination shows that the distribution induced by this fibration coincides with that of α. □

A computational tool that will aid us momentarily is the Euler class of an oriented vector bundle.

Definition 4.30. Let τ denote a smooth oriented vector bundle over an oriented compact manifold $E \to S$. The **Euler class** $e(\tau)$ is the Poincaré dual in $H^*(S)$ of the zero locus of a generic section of E.

The above definition is to be interpreted in the following way: S is naturally embedded as the zero section of E. A generic smooth section $s : S \to E$ intersects the zero section transversely. This intersection is a submanifold of the zero section, which we identify with S. Then the Euler class is the Poincaré dual of the fundamental class of this submanifold.

The Euler class is an example of a characteristic class associated to an oriented vector bundle and as such enjoys several properties; the only one we will use is that it is natural with respect to the pullback operation of vector bundles. Proper treatment is given in [MS74], [Bre93] and [BT82]. The Euler class is the "primary obstruction" to a non-zero section of τ. For example, we have:

Proposition 4.31. *If τ admits a non-zero section, then $e(\tau)$ is zero.*

Proof. In this case, a generic section is chosen with empty zero locus. Hence its Poincaré dual is 0. □

We will perform calculations of the Euler class only for rank-2 bundles over compact oriented surfaces. Every bundle we encounter will admit a section at the boundary of the surface. Then the Euler class is an element of $H^2(S, \partial S; \mathbb{Z})$ and consequently may be expressed as a multiple of the fundamental class of S. This multiple is exactly the Euler characteristic of S.

Proposition 4.32. *Let TS denote the tangent bundle of a compact oriented surface S. Then $e(TS)([S]) = \chi(S)$ where $[S] \in H_2(S, \partial S; \mathbb{Z})$ is the fundamental class of S.*

Proof. (Sketch.) The zero locus of a generic section of TS is a finite collection of points (away from the boundary). The Euler class counts these points with sign ± 1 according to the orientations of the intersecting sections. On the other hand, a generic section may be considered as a smooth vector field on S with only isolated singularities. Examination shows that the indices of this vector field coincide with the signs determined by the orientation of the intersection. Thus the result follows by the Poincaré-Hopf index theorem. □

Our next goal is to show that compact leaves of a taut foliation minimize the Thurston norm amongst representatives in that homology class. First some more general definitions.

Definition 4.33. A 3-manifold is **irreducible** if every embedded S^2 bounds an embedded ball.

Irreducibility implies that each embedded S^2 is homotopically trivial. A link complement M is irreducible if and only if the link is not split. A link $L \subset S^3$ is **split** if there exists a smoothly embedded $S^2 \subset S^3$ disjoint from the link that separates some components of the link. Alexander's Theorem [Ale24], a well-known result in 3-manifold topology, states that any smoothly embedded S^2 in S^3 bounds a smoothly embedded ball. This establishes that a link is split if and only if the link complement is irreducible.

Definition 4.34. A properly embedded compact surface S in a 3-manifold M is **compressible** if there exists an embedded disk $D^2 \subset M$ such that $D^2 \cap S = \partial D^2$ and ∂D^2 does not bound a disk on S. If no such disk exists, then S is **incompressible**.

By the Loop Theorem of Papakyriakopoulos (see [Hem76]), a properly embedded compact surface S in an oriented compact 3-manifold is incompressible if and only if $\pi_1(S) \to \pi_1(M)$ is an injection. It is easy to see then, that fibers of a fibered 3-manifold are incompressible, since by the long exact sequence of homotopy groups from a fibration, the inclusion of the fundamental group of S is an injection. Similarly, it is true that compact leaves of a taut foliation are incompressible, but this is not as easy to see [Nov65], [Thu86].

Proposition 4.35. *Suppose M is an irreducible link complement and S is a Thurston norm minimizing surface representing a non-trivial homology class. Then either S is incompressible, or S is homologous to a union of two disks.*

Proof. Suppose S is compressible. Then there is a disk $D^2 \subset M$ whose boundary lies on S but does not bound a disk on S. We may perform the same surgery as we did when removing non-essential intersections in Theorem 4.3. That is, we remove from S the intersection of S with a tubular neighborhood of D^2 and glue back in two disks on either side of the original disk. This operation results in a new surface S' whose Euler characteristic has increased by 2.

If this surgery does not increase the number of components of S, then S' must be a sphere since it is Thurston norm minimizing. But M is irreducible, so this is impossible (it implies that S is homologically trivial). Therefore the operation results in disjoint surfaces $S' = S_1 \cup S_2$ and we have

$$\chi(S_1) + \chi(S_2) = \chi(S) + 2$$

The only way this equation does not contradict the norm-minimality of S is if $\chi(S) = 0$ and either one of S_1 or S_2 is a sphere, or both are disks. But neither S_1 or S_2 can be a sphere since this would imply that ∂D^2 bounded a disk on S. □

Now consider a foliation $\mathfrak{F}$ of M and a properly embedded surface S. If S is not a leaf of $\mathfrak{F}$, then we consider the intersection of leaves of $\mathfrak{F}$ with S. For a generic S, this intersection produces a foliation of S on the complement of some set of singular points. These are the points where the tangent space of S corresponds with a tangent plane of the foliation. Although we will not prove this, these singularities fall into three categories: center singularities, which locally look like maxima and minima, saddle singularities, which locally appear as saddles (i.e. in Morse theoretic language they are index 1 isolated singularities), and circle singularities, which look like a circular ridge tops. The critical theorem of Thurston and Roussarie states that under certain circumstances, we may remove all center and circle tangencies by isotopy:

Theorem 4.36. *(Thurston, Roussarie). Suppose $\mathfrak{F}$ is a taut foliation and S is a properly embedded incompressible surface. Then S can either be isotoped into a leaf, or isotoped to intersect $\mathfrak{F}$ in only saddle tangencies. Furthermore every boundary component of S is either contained in a leaf or is transverse to $\mathfrak{F}$.*

Remark. Separate arguments for the proof are given in [CC03] and [Cal07].

Suppose now that S has been isotoped in M to have only saddle tangencies with $\mathfrak{F}$. We can give a sign to every singularity, depending on whether the normal orientation of S agrees with the transverse orientation of $\mathfrak{F}$. We let $I_p(S)$ denote the number of positive saddle intersections (where these orientations agree) and $I_n(S)$ denote the number of negative saddle intersections.

Proposition 4.37. *Let τ denote the tangent subbundle of the foliation $\mathfrak{F}$. Then we have*

$$e(\tau)([S]) = I_n - I_p$$
$$\chi(S) = I_p + I_n$$

Proof. (Sketch.) First we calculate $e(\tau)([S])$. By the naturality of the Euler class and its definition, this is the number of points in the intersection of a generic section of τ restricted to S with the zero locus, counted with sign. Conveniently, a generic section of τ is provided by the intersection of the foliation with S. Explicitly, at non-singular points, τ intersects the tangent plane of S in a 1-dimensional subspace, which we orient according to the orientations of S and τ. Examination of this vector field around singularities show that the intersection has positive sign when the saddle singularity is of negative type, and vice versa. This establishes the first equality.

On the other hand, the vector field index at such a singularity is always -1. Hence the second equality follows by the Poincaré-Hopf theorem. □

Corollary 4.38. *Suppose S, $\mathfrak{F}$, and τ are as above. Then*

$$|e(\tau)([S])| \leq -\chi(S)$$

with equality if and only if S is isotopic to a leaf of $\mathfrak{F}$ or S is isotopoic to an embedding that admits only positive type saddle singularities with respect to $\mathfrak{F}$. □

This relation is just what we need to show that compact leaves minimize the Thurston norm.

Theorem 4.39. *Suppose M is an irreducible compact oriented 3-manifold equipped with a taut foliation $\mathfrak{F}$. Then compact leaves of $\mathfrak{F}$ minimize the Thurston norm.*

Proof. By Proposition 4.35, since M is irreducible every non-trivial homology class has an incompressible surface S representing it. Therefore if S is such a surface in the homology class of a fiber, it satisfies the above inequality (note that a compact leaf of a taut foliation is always homologically nontrivial, since it has non-zero intersection number with some closed transversal, by the definition of tautness). But the left side is exactly the Thurston norm of a leaf (remember that in our situation, the leaves must have non-positive Euler characteristic by Reeb stability, so the Thurston norm agrees with the negative Euler characteristic). □

In particular, a fiber of a fibered link complement minimizes the Thurston norm.

Corollary 4.40. *If ϕ is a fibered class of a link complement, then $x_T(\phi) = \|\phi\|_A$.*

Proof. Examine the proof of Theorem 4.14. By the above theorem, a fiber F minimizes the Thurston norm. By Proposition 4.19, M_ϕ has the homotopy type of F so in particular F is connected and $b_1(F) = b_1(M)$. Hence for the surface S used in Theorem 4.14 we may use F. The inequality between the Alexander and Thurston norms comes from $b_1(S) \geq b_1(M_\phi)$, but in this case there is equality. □

Finally, we wish to associate certain top-dimensional faces of the Thurston ball with fibered classes. Suppose we are given a fibration $M \to S^1$. Let τ denote the subbundle of TM corresponding to the fibration. The boundary of M is transverse to the leaves of the foliation, which means that in a neighborhood of ∂M there is a nowhere zero section of τ (pointing away from ∂M). This implies that the Euler class of τ restricted to ∂M is 0. Consequently, we may consider $e(\tau)$ as a relative cohomology class in $H^2(M, \partial M; \mathbb{Z})$.

Proposition 4.41. *(Thurston) Let M be a link complement which fibers over S^1 with fiber F a surface of negative Euler characteristic. Let τ denote the tangent subbundle corresponding to the fibration. Then*

$$x(\alpha) = |e(\tau)(\alpha)|$$

in a neighborhood of $[F]$ in $H_2(M, \partial M; \mathbb{R})$.

Proof. First let us evaluate $e(\tau)([F])$. The bundle τ restricts to the tangent bundle of F, so by the naturality of the Euler class and the previous proposition, this evaluation equals the Euler characteristic of F. Since F has negative Euler characteristic, we have

$$|e(\tau)([F])| = \chi(F) = \chi_-(F) = x_T([F])$$

We must find a neighborhood of $[F]$ in which this equality holds. Let $\alpha_1, \ldots, \alpha_l$ be closed 1-forms representing a basis of $H_1(M; \mathbb{R})$ and let α be the closed non-singular 1-form obtained as the pullback via the fibration of the volume form on S^1. For small constants ϵ_i, the closed 1-form $\omega = \alpha + \epsilon_1\alpha_1 + \cdots + \epsilon_l\alpha_l$ is still non-singular. Therefore it defines a foliation. Furthermore, if ω belongs to the rational subspace $H^1(M; \mathbb{Q})$, then it corresponds to a fibration, since an integer multiple of ω is an integral form and we apply Theorem 4.29, and any multiple of a form clearly induces the same distribution.

Suppose ω belongs to the rational subspace and corresponds to a fibration with fiber F', and let τ' denote the tangent plane bundle corresponding to the foliation induced by ω. By the calculation above, $|e(\tau')([F'])| = x_T([F'])$. But if ω is sufficiently close to α, the bundles τ and τ' are isomorphic via small rotations in TM at every point. Therefore the formula

$$|e(\tau)([F'])| = x_T([F'])$$

holds in a dense neighborhood of $[F]$, hence it holds in the whole neighborhood by the continuity of x_T. □

Since both the Thurston norm and evaluation of $e(\tau)$ are linear, we actually get the stronger statement that the equality $e(\tau)(\alpha) = -x(\alpha)$ holds in a cone neighborhood of the ray containing $[F]$. Furthermore, this implies that the intersection of this cone with the unit ball of the Thurston norm satisfies the linear relation $e(\tau)(-) = -1$. Therefore this intersection lies in a top-dimensional face, and every element of the top-dimensional face satisfies this relation.

Theorem 4.42. *(Thurston)[Cal07] The set of rays in $H_2(M, \partial M; \mathbb{R})$ corresponding to fibrations of M over S^1 is exactly the set of rational rays intersecting union of some of the top-dimensional faces of the Thurston norm unit ball.*

Proof. Suppose $\mathfrak{F}$ is a foliation obtained from some foliation that intersects a top-dimensional face Δ. As we've mentioned, since $\mathfrak{F}$ is transversely oriented, it is represented by a non-singular closed 1-form α which is Poincaré dual to a fiber. Next let S be a norm-minimizing surface lying on a ray that intersects Δ. By Tischler's theorem, it is enough to show that every ray intersecting Δ is represented by a non-singular closed 1-form (then the rational rays, by our previous remarks, correspond to fibrations). To that end, let β be a closed 1-form Poincaré dual to S whose support is contained in a small neighborhood of S. Then it suffices to show that for all $t \geq 0$, $u > 0$ such that $t\beta + u\alpha$ is contained in the rational subspace, this element corresponds to a fibration.

By the previous paragraph, we have $e(\tau)([S]) = -x_T(S)$ where τ is the tangent subbundle corresponding to $\mathfrak{F}$. Therefore, by Corollary 4.38, S is either isotopic to a leaf, or to a properly embedded surface transverse to $\mathfrak{F}$ except at positive saddle singularities. In the first case, $\alpha = [S]$ so any combination $t[S] + u\alpha$ already belongs to the ray of α. In the second case, we may find a vector field X in a neighborhood of S that is transverse to both $\mathfrak{F}$ and S and such that $\alpha(X) > 0$: when $\mathfrak{F}$ and S are transverse such a choice is obvious, and since the only singularities are positive saddle singularities, we can coherently choose X around these points.

Now consider $t\beta + u\alpha$: it is non-singular, since around S we have $(t\beta + u\alpha)(X) > 0$, and away from S it is just $u\alpha$. Therefore, if this element belongs to the rational subspace, it corresponds to a fibration by our previous comments. □

We have seen that on fibered faces, the Alexander and Thurston norm agree. One may ask that if a link complement fibers in some way over S^1, then do the Alexander and Thurston norm agree everywhere? Dunfield answered this in the negative, providing a counterexample in [Dun01]. In fact, the author exhibited a link whose Thurston unit ball has a fibered face properly contained in a top-dimensional face of the Alexander unit ball.

Chapter 5

Epilogue

In conclusion, we offer a glimpse into two modern theories of link invariants that generalize some of the ideas of this thesis.

5.1 Link Floer homology

Link floer homology is a 3-manifold invariant developed by Ozsváth and Szabó and is closely related to Heegaard-Floer homology of 3-manifolds. The original construction of the chain complex from which homologies are derived come from a Heegaard diagram of a 3-manifold with varoius extra information attached to it; this is the holomorphic construction. The steps needed to define the chain complexes, particularly justifying the existence of a suitable boundary map, are quite technical and will not be discussed. There is a large amount of literature expositing this construction, e.g. [OS08a] and the references therein.

More recently, a more elementary combinatorial construction of the chain complexes is given by the data of a grid diagram, see [MOST07] and the forthcoming book by Ozsváth, Stipsicz and Szabó [OSS]. Our discussion in Chapter 3 of the Alexander polynomial derived from a grid diagram is intended to hint towards the sort of computations one encounters in this theory.

In both the combinatorial link Floer homology and the holomorphic version, one ends up with a collection of slightly different chain complexes. We focus our attention on $\widehat{\mathrm{HFL}}(\vec{L})$, the **link Floer homology**. This is a multi-graded $\mathbb{F}$ module where $\mathbb{F} = \mathbb{Z}/2\mathbb{Z}$. The $\mathbb{F}$-vector space of the module is usually relatively simple, but lots of information is contained in the gradings. The multi-grading consists of two parts: one component of the grading is an element $d \in \mathbb{Z}$ called the Maslov grading. The other part is a multi-index that is an element of the half-integer lattice of the link complement. Precisely, one defines a subset $\mathbb{H}(\vec{L}) \subset H_1(S^3 - L; \mathbb{Q})$ called the **Alexander grading set**. Denoting the meridional basis of $H_1(S^3 - L; \mathbb{Z})$ by $\mu_1, \ldots, \mu_l$, then $\mathbb{H}(\vec{L})$ consists of those elements of the form

$$\sum_{i=1}^{l} h_i \mu_i$$

such that $2h_i + \mathrm{lk}(L_i, L - L_i)$ is an even integer. Here L_i is the ith component of the link L and lk denotes the linking number of closed curves. With the grading in hand we may express the link Floer homology as

$$\widehat{\mathrm{HFL}}(\vec{L}) = \bigoplus_{d\in\mathbb{Z},\ h\in\mathbb{H}(\vec{L})} \widehat{HFL}_d(\vec{L}, h)$$

The **Poincaré polynomial** associated to such a multi-graded $\mathbb{F}$-module is

$$P(q, t_1, \dots, t_l) = \sum_{d\in\mathbb{Z},\ h\in\mathbb{H}(\vec{L})} q^d t_1^{h_1} \dots t_l^{h_l} \operatorname{rank} \widehat{HFL}_d(\vec{L}, h)$$

where $h = \sum_{i=1}^{l} h_i\mu_i$. Substituting $q = -1$ returns the Euler characteristic of the homology, which turns out to be the multivariable Alexander polynomial as defined in Chapter 3. However, the Poincaré polynomial includes much more information than the Alexander polynomial, and various extra information can be extracted by substituting other values for q.

This connection between the Alexander polynomial and link homology is roughly understood by the following description of the chain complex obtained from a grid diagram. One first considers the set of **grid states**, which are bijections between the rows and columns of the grid diagram, as usual excluding the lower row and the right hand column of vertices. The grid states freely generate the chain groups, and each grid state is assigned a Maslov grading and an Alexander grading.

It is clear that grid states identify a summand in the Leibniz formula for the determinant of the grid matrix, since such a summand is precisely a bijection between the rows and columns of the matrix. As we've seen in Chapter 3, the Alexander polynomial is obtained from this determinant. Now the Maslov grading essentially encodes the sign of the permutation in each summand (whether this term is added to the determinant with a plus or minus according to the Leibniz formula). The Alexander grading encodes the sum of the winding numbers of those vertices including in a particular grid state. Now the Euler characteristic of the homology may be computed from the chain complex by summing over these grid states, with the sign and appropriate powers of $t_1, \dots, t_l$ determined by the Maslov and Alexander gradings. Careful inspection of the gradings reveals that the resulting polynomial is exactly the Alexander polynomial defined in Chapter 3. Link floer homology is therefore said to categorify the Alexander polynomial, which is a general concept meaning that the invariant is captured as the Euler characteristic of a homology theory.

This shows that the invariant $\widehat{HFK}(\vec{L})$ is at least as strong as the Alexander polynomial; in fact it is much stronger. For example, we may generalize the Alexander norm to the **link homology norm** on $H^1(S^3 - L; \mathbb{R})$ defined by

$$y(h) = \max_{s\in\mathbb{H}(\vec{L})|\widehat{HFL}(\vec{L},s)\neq 0} |h(s)|$$

where $h(s)$ is the Kronecker evaluation. Then it was shown that

Theorem 5.1. *[OS08b] For a link $\vec{L}$ with no trivial components and for each $h \in H^1(S^3 - L; \mathbb{Z})$ we have*

$$x_T(h) + \sum_{i=1}^{l} |h(\mu_i)| = 2y(h)$$

Therefore the link homology determines the Thurston norm exactly.

Another amazing application of link floer homology extending the results of the thesis is the ability of this invariant to detect fibered links. Precisely, we have the following theorem due to Ni:

Theorem 5.2. *[Ni07] Suppose L is a fibered link. If $\Delta \subset H^1(S^3 - L; \mathbb{Z})$ corresponds to a fibered face of the Thurston norm with corresponding extremal point $h \in H_1(S^3 - L; \mathbb{Q})$ then h belongs to the Alexander grading set and the $\mathbb{F}$-module $\widehat{HFL}(\vec{L}, h)$ is one-dimensional.*

Conversely, suppose $h \in H_1(S^3 - L; \mathbb{Q})$ is an extremal point corresponding to a face Δ of the Thurston norm and $\widehat{HFL}(\vec{L}, h)$ is one-dimensional. Then L is fibered and Δ is a fibered face.

The first part of the theorem corresponds to the fact that the single-variable Alexander polynomials corresponding to fibered classes are monic. The more surprising result is the converse statement, of which there is no analogue presented in this thesis.

The results summarized here only hint at the efficacy of the link Floer invariants and were chosen because of their relevance to this paper. On the other hand, while it is an exceptional theoretical tool, the computation of link Floer homology (particularly using the holomorphic construction) can become quite unwieldy for practical purposes. This displays well the trade-off between power and ease of use (or computability) any topologist must deal with when choosing between, or inventing, modern invariants.

5.2 Twisted Alexander polynomials and Reidemeister torsion

While the link Floer homology is related to the presentation of the Alexander polynomial given in Chapter 3, the twisted Alexander polynomials are generalizations of the techniques described in Chapter 2. Research in torsion invariants began with the work of Reidemeister, who used them to classify lens spaces, which are 3-manifolds well known to be homotopy equivalent but not homeomorphic. Later Milnor realized the connection between the Alexander polynomial and this so-called Reidemeister torsion. More recently, research by Turaev and others has developed the torsion invariants into a collection of invariants called the twisted Alexander polynomials. Like link Floer homology, twisted Alexander polynomials offer a very powerful collection of invariants, but with the caveat that one must consider all representations of the fundamental group of the link complement at once, making it a somewhat difficult tool to use in practice. Since the definitions of the twisted torsion invariants is relatively simple, we outline them in full, following [FV11a], [Tur00] and [Mas08].

We briefly sketch the ideas behind the twisted Alexander polynomials. The fundamental tool is the torsion of a based acyclic chain complex, which is a finite-dimensional chain complex over a field $\mathbb{F}$ with trivial homology and a pre-determined choice of basis. Hence we have a complex

$$C = 0 \to C_m \xrightarrow{\partial_{m-1}} C_{m-1} \xrightarrow{\partial_{m-2}} \dots \xrightarrow{\partial_0} C_0 \to 0$$

with a choice of basis c_i for each vector space C_i. By acyclicity, the sequence is made up of short exact sequences

$$0 \to B_i \to C_i \to B_{i-1} \to 0$$

where $B_i = \operatorname{im} \partial_i$. We pick bases b_i of each B_i. Then from the above sequence, we can lift the basis b_{i-1} and append it with b_i to get a basis of C_i, denoted $b_i\widetilde{b_{i-1}}$. Then there is a change of basis matrix A expressing $b_i\widetilde{b_{i-1}}$ in the basis c_i. The quantity $[b_i\widetilde{b_{i-1}}/c_i]$

is defined to be $\det(A)$. We define an equivalence relation on bases of a vector space by $b \equiv c$ if the determinant of the change of basis matrix is 1.

Definition 5.3. The **torsion** of the based acyclic complex C is the quantity

$$\tau(C) = \prod_{i=0}^{m} [b_i \widetilde{b_{i-1}} / c_i]$$

It is a routine check to show that the torsion does not depend on the choice of b_i or the lift $\widetilde{b_{i-1}}$. It does, however, depend on the original bases c_i. Note that the torsion is a sort of secondary invariant to the Euler characteristic, since it is defined only for acyclic complexes, for which the Euler characteristic vanishes. The following proposition is proved with elementary linear algebra, which establishes the torsion as a multiplicative analogue of the Euler characteristic:

Proposition 5.4. *1. (Multiplicativity.) If $0 \to C' \to C \to C''$ is a short exact sequence of based acyclic complexes such that the basis of C is equivalent to one given by the basis of C' and a lift of C'', then*

$$\tau(C) = \pm\tau(C')\tau(C'')$$

2. (Duality.) If $C = 0 \to C_m \to \cdots \to C_0 \to 0$ is a based acyclic chain complex, then its dual C^* obtained by applying $\mathrm{Hom}(-, \mathbb{F})$ is a based acyclic chain complex and

$$\tau(C^*) = \pm\tau(C)^{(-1)^{(m+1)}}$$

3. (Homological computation. [Mas08]) Suppose Λ is a Noetherian UFD and C is a finitely generated free chain complex over Λ which is based and satisfies $\operatorname{rank} H_i(C) = 0$ for all $i = 0, \ldots, m$. Let $Q(\Lambda)$ denote the field of fractions of Λ. Then we have:

$$\tau(C \otimes Q(\Lambda)) = \prod_{i=0}^{m} (\operatorname{ord} H_i(C))^{(-1)^{(i+1)}}$$

Now suppose M is a compact oriented 3-manifold that is either closed or has toroidal boundary, for example a link complement in S^3. Let F be a free abelian group and R a commutative domain, so $R[F]$ is the group algebra over R with field of fractions $Q(R[F])$. Suppose we are given a representation $\alpha : \pi_1(M) \to \mathrm{GL}(k, R[F])$.

Fix a CW-decomposition of M. This induces a CW-decomposition of the universal cover $\tilde{M} \to M$ given by the lifts of cells of M. The CW-chain complex $C_*(\tilde{M})$ is given a right $\mathbb{Z}[\pi_1(M)]$-module structure by $\sigma \cdot g = g^{-1}(\sigma)$ for all chains σ and g is acting by deck transformations. On the other hand, the representation α gives rise to a left group action of $\pi_1(M)$ on $Q(R)[F]^k$. Therefore we can consider the following chain complex over $Q(R)[F]$:

$$C_*(\tilde{M}) \otimes_{\mathbb{Z}[\pi_1(M)]} Q(R)[F]^k$$

Furthermore, we can endow this chain complex with a basis by using cells of $\tilde{M}$ coming from the cell decomposition of M, and picking the canonical basis of $Q(R)[F]^k$.

Definition 5.5. The **twisted Reidemeister torsion** of M corresponding to α is 0 if the above complex is not acyclic, and is equal to

$$\tau(M, \alpha) = \tau(C_*(\tilde{M}) \otimes_{\mathbb{Z}[\pi_1(M)]} Q(R)[F]^k)$$

otherwise.

A theorem of Chapman states that this quantity is determined by M and α up to multiplication by an element of $\{\det \alpha(g) \mid g \in \pi_1(M)\}$. However, the Reidemeister torsion is not a homotopy invariant in general, although certain choices of representation α do result in a homotopy invariant. An important example is the **Milnor torsion** derived from the representation $\mu : \pi_1(M) \to Q(\mathbb{Z}[H_1(M;\mathbb{Z})])$ induced by the Hurewicz homomorphism. The following proposition, proved using the homological computation property of the torsion, relates the Milnor torsion with the classical (untwisted) multivariable Alexander polynomial:

Proposition 5.6. *[Mil62] The Milnor torsion of a (greater than* 1 *component) link with complement $M = S^3 - L$ is calculated by*

$$\tau(M, \mu) \doteq \Delta(L)$$

Milnor actually proved the more general statement that the Milnor torsion is equivalent to the Alexander function of a finite CW-complex:

Definition 5.7. The **Alexander function** $A(X)$ of a finite CW-complex X is

$$A(X) = \prod_{i \geq 0} (\mathrm{ord}(H_i(X_\infty;\mathbb{Z}))^{(-1)^{i+1}}$$

where X_∞ is the maximal free abelian cover of X and the orders are calculated considering $H_i(X_\infty;\mathbb{Z})$ as $\mathbb{Z}[H_1(X;\mathbb{Z})]$-modules.

It is easy to show that for link complements, the higher homologies of X_∞ vanish so the Alexander function coincides with the Alexander polynomial (recall it is the order of $H_1(X_\infty;\mathbb{Z})$, so the correspondence between the Milnor torsion and Alexander function clearly gives us the previous proposition. This correspondence for the representation μ invites a generalization of the Alexander modules to other homomorphisms:

Definition 5.8. Let M be a link complement and $\alpha : \pi_1(M) \to \mathrm{GL}(k, R[F])$ be a representation where R is a Noetherian UFD and $[F]$ is a free abelian group. For $i \geq 0$ the **ith twisted Alexander module of** (M, α) is the R[F]-module

$$H_i(M; R[F]^k) := H_i(C_*(\tilde{M}) \otimes_{\mathbb{Z}[\pi_1(M)]} R[F]^k)$$

The **ith twisted Alexander polynomial** is $\mathrm{ord}(H_i(M; R[F]^k)) \in R[F]$, denoted $\Delta^\alpha_{M,i}$.

We see now that the twisted Alexander modules and polynomials are direct generalizations of the classical definitions given in Chapter 2. It is possible to compute the twisted Alexander polynomials using a variation of the Fox calculus (see [Tur00]). The analogue of Milnor's theorem relating the twisted Reidemeister torsion to the Alexander polynomial for arbitrary representations is

Theorem 5.9. *Let M be a link complement and $\alpha : \pi_1(M) \to \mathrm{GL}(k, R[F])$ be a representation. If $\Delta^\alpha_{M,i} \neq 0$ for $i = 0, 1, 2$ then*

$$\tau(M, \alpha) = \prod_{i=0}^{2} (\Delta^\alpha_{M,i})^{(-1)^{i+1}}$$

We are left with a large collection of link invariants corresponding to representations of the fundamental group. We sample some of the recent results obtained using twisted Alexander polynomials.

Theorem 5.10. *[FV12] Let M be an irreducible link complement and let $\phi \in H^1(M;\mathbb{Z}) = Hom(H_1(M;\mathbb{Z},\mathbb{Z})$. Then there exists a representation $\alpha : \pi_1(M) \to \mathrm{GL}(k,\mathbb{C})$ such that the degree of the twisted Reidemeister torsion $\tau(M, \alpha\otimes\phi)$ determines the Thurston norm of ϕ.*

Here $\alpha \otimes \phi$ is the tensor product representation $\pi_1(M) \to \mathrm{GL}(k, \mathbb{C}[t, t^{-1}])$. The theorem does not seem to give a terminating procedure for determining the Thurston norm, since it alludes only to the existence of a suitable representation. However, the authors of [FV12] do describe an algorithm for calculating the Thurston norm exactly.

As in the case of link Floer homology, twisted Alexander polynomials also detect fibered classes:

Theorem 5.11. *[FV11b] Let M be a link complement. Then there is a necessary and sufficient condition on the set of twisted Alexander polynomials of the form $\Delta_M^{\alpha\otimes\phi}$ where $\phi \in H^1(M;\mathbb{Z})$ and $\alpha : \pi_1(M) \to G$ is a representation of $\pi_1(M)$ in a finite group G, to determine whether ϕ is a fibered class.*

While we did not write down the condition of the theorem, we remark that it involves the comparison of the degree of $\Delta_M^{\alpha\otimes\phi}$ with the Thurston norm of ϕ. Therefore if the Thurston norm of ϕ is unknown it may not be easy to check this condition, but by the preceding result it is theoretically possible to compute the Thurston norm in finite time and subsequently apply this theorem. Finally, we note that in order to show ϕ is a fibered class, it suffices to provide a single representation of $\pi_1(M)$ to a finite group satisfying the condition of the theorem, and conversely, to show ϕ is not a fibered class it is sufficient to exhibit a single representation that doesn't satisfy the condition.

It is unknown whether twisted Alexander polynomials give a complete knot invariant, but it has been shown that these invariants are able to detect the unknot, trivial links, the Hopf link, the trefoil knot, and the figure-8 knot [FV13].

We remark that both the link Floer homology and the theory of torsion invariants are connected to Seiberg-Witten theory, and an explicit connection between some of the information derived from both theories is given in ([OS04], Theorem 1.2).

5.3 A final example

To conclude, we provide an example of knots which are indistinguishable by the known homological knot invariants (in particular, the link Floer homology) but are distinguishable with the methods presented in this thesis. These knots are the Kanenobu knots, parametrized by two integers p, q and denoted $K_{p,q}$. They are given in Figure 5.1. A box labeled with an integer n means that n full twists occur there (where, if we imagine the strands bounding a ribbon, a positive twist corresponds to twisting the ribbon counterclockwise). In fact, the Kanenobu knots are examples of a general class of knots called ribbon knots, which roughly speaking, are formed by connecting trivial links with embedded ribbons and taking the resulting boundary.

Kanenobu [Kan86] shows that $K_{p,q}$ and $K_{p',q'}$ are equivalent knots if and only if $p, q = p', q'$ as unordered pairs. However, it is known that for pairs (p, q) and (p', q') such that $p+q = p'+q'$ and $pq \equiv p'q' \mod 2$, all the known homological knot invariants coincide for $K_{p,q}$ and $K_{p',q'}$ [Lob14]. This includes the link Floer homology discussed above, and the Khovanov homology, which categorifies the Jones polynomial. As a result none of the known polynomial invariants can distinguish these knots.

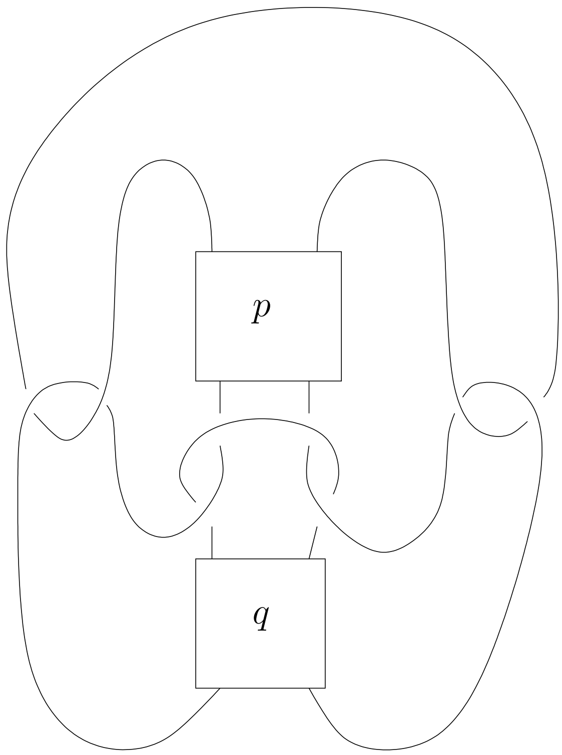

Figure 5.1: The Kanenobu knots $K_{p,q}$.

On the other hand, Kanenobu gives the following presentation for the Alexander module:

$$\begin{pmatrix} t^2 - 3t + 1 & (p-q)t \\ 0 & t^2 - 3t + 1 \end{pmatrix}$$

A short calculation shows that the knots $K_{p,q}$ and $K_{p',q'}$ have isomorphic Alexander modules if and only if $|p-q| = |p'-q'|$. Therefore there are infinitely many examples of knots which are indistinguishable even with the full force of link Floer homology, but are easily distinguished by examining the second elementary ideal of the Alexander module.

We hope this convinces the reader that despite the power of modern invariants, it is still useful to have a thorough understanding of the more elementary tools presented in this thesis.

Bibliography

[Ale24] James W. Alexander. On the subdivision of 3-space by a polyhedron. *Proc. Nat. Acad. Sci. USA*, 10(1):6–8, 1924.

[Ale28] James W. Alexander. Topological invariants of knots and links. *Trans. Amer. Math. Soc.*, 30(2):275–306, 1928.

[AM69] M. F. Atiyah and I. G. Macdonald. *Introduction to Commutative Algebra.* Addison-Wesley Series in Mathematics. Westview Press, 1969.

[Bre93] Glen E. Bredon. *Geometry and Topology.* Springer-Verlag, 1993.

[BT82] Raoul Bott and Loring W. Tu. *Differential forms in algebraic topology.* Springer-Verlag, 1982.

[Cal07] Danny Calegari. *Foliations and the geometry of 3-manifolds.* Oxford Mathematical Monographs. Oxford University Press, 2007.

[CC03] Alberto Candel and Lawrence Conlon. *Foliations II*, volume 60 of *Graduate Studies in Mathematics.* American Mathematical Society, 2003.

[CF63] Richard H. Crowell and Ralph H. Fox. *Introduction to Knot Theory.* Springer-Verlag, 1963.

[Cro95] Peter R. Cromwell. Embedding knots and links in an open book I: Basic properties. *Topology and its applications*, 64(1):37–58, 1995.

[Dun01] Nathan M. Dunfield. Alexander and Thurston norms of fibered 3-manifolds. *Pacific J. Math*, 200:43–58, 2001.

[FM11] Benson Farb and Dan Margalit. *A Primer on Mapping Class Groups.* Princeton University Press, 2011.

[FV11a] Stefan Friedl and Stefano Vidussi. A survey of twisted alexander polynomials. In *The Mathematics of Knots*, volume 1 of *Contributions in Mathematical and Computational Sciences*, pages 45–94. Springer Berlin Heidelberg, 2011.

[FV11b] Stefan Friedl and Stefano Vidussi. Twisted Alexander polynomials detect fibered 3-manifolds. *Ann. of Math.*, 173(3):1587–1643, 2011.

[FV12] Stefan Friedl and Stefano Vidussi. The Thurston norm and twisted Alexander polynomials, 2012. Available at http://arxiv.org/pdf/1204.6456v2.pdf.

[FV13] Stefan Friedl and Stefano Vidussi. Twisted Alexander polynomials detect trivial links, 2013. Available at http://arxiv.org/abs/1306.3487.pdf.

[GL89] C. McA. Gordon and J. Luecke. Knots are determined by their complements. *Bull. Amer. Math. Soc. (N.S.)*, 20(1):83–87, 1989.

[Hem76] John Hempel. *3-Manifolds*. Princeton University Press, 1976.

[Hir97] Eriko Hironaka. Alexander stratifications of character varieties. *Annales de l'institut Fourier*, 47(2):555–583, 1997.

[Kan86] Taizo Kanenobu. Infinitely many knots with the same polynomial invariant. *Proc. Amer. Math. Soc.*, 97:158–162, 1986.

[Kaw90] Akio Kawauchi. *A Survey of Knot Theory*. Springer-Verlag, 1990.

[Lob14] Andrew Lobb. The Kanenobu knots and Khovanov-Rozansky homology. *Proc. Amer. Math. Soc.*, 142(4), 2014.

[Mas08] Gwénaël Massuyeau. A short introduction to the Alexander polynomial. http://www-irma.u-strasbg.fr/∼massuyea/talks/Alex.pdf, 2008.

[McM02] Curtis T. McMullen. The Alexander polynomial of a 3-manifold and the Thurston norm on cohomology. *Ann. Sci. Ecole Norm. Sup. (4)*, 35(2):153, 2002.

[Mil62] John Milnor. A duality for Reidemeister torsion. *Ann. of Math. (2)*, 76:137–147, 1962.

[MOST07] Ciprian Manolescu, Peter Ozsváth, Zoltán Szabó, and Dylan Thurston. On combinatorial link Floer homology. *Geometry and Topology*, 11:2339–2412, 2007.

[MS74] John Milnor and John Stasheff. *Characteristic Classes*. Number 76 in Annals of Math. Studies. Princeton University Press, 1974.

[Mur93] Jun Murakami. A state model for the multivariable Alexander polynomial. *Pacific J. Math.*, 157(1):109–135, 1993.

[Ni07] Yi Ni. Knot Floer homology detects fibred knots. *Inventiones mathematicae*, 170(3):577–608, 2007.

[Nov65] S. P. Novikov. Topology of foliations. *Trans. Moscow Math. Soc.*, 14:268–305, 1965.

[OS04] Peter Ozsváth and Zoltán Szabó. Holomorphic disks and three-manifold invariants: Properties and applications. *Ann. of Math.*, 159(3):1159–1245, 2004.

[OS08a] Peter Ozsváth and Zoltán Szabó. Holomorphic disks, link invariants, and the multi-variable Alexander polynomial. *Algebraic and Geometric Topology*, 8(2):615–692, 2008.

[OS08b] Peter Ozsváth and Zoltán Szabó. Link Floer homology and the Thurston norm. *J. Amer. Math. Soc.*, 21:671–709, 2008.

[OSS] Peter S. Oszváth, András Stipsicz, and Zoltán Szabó. Grid homology for knots and links. unpublished.

[Rol76] Dale Rolfsen. *Knots and Links.* Publish or Perish, Inc., 1976.

[Thu86] W.P. Thurston. A norm for the homology of 3-manifolds. *Mem. Amer. Math. Soc.*, 339:99–130, 1986.

[Tis70] D. Tischler. On fibering certain foliated manifolds over S^1. *Topology*, 9(2):153–154, 1970.

[Tor53] Guillermo Torres. On the Alexander polynomial. *Ann. of Math.*, 57(1):57–89, 1953.

[Tur00] Vladimir Turaev. *Introduction to Combinatorial Torsions.* Birkhäuser Verlag, 2000.

[Zas49] Hans Zassenhaus. *The Theory of Groups.* Chelsea Publishing Co., 1949.

STATEMENT

Name:
ELTE Faculty of Science,
Neptun identity code
Title of the MSc Thesis Work:

I hereby declare, as the author of this thesis, that it is a product of my own and that the structure and the mode of presentation is based on my own ideas. In my thesis I use the standard rules for references and quotations consistently; I do not borrow and write down any part from the work of other people without proper notification and citation.

20 ________________________________, Budapest

the signature of the student