Floer-Theoretic and Combinatorial Approaches to Knots



Soheil Azarpendar

Linacre College

University of Oxford

A thesis submitted for the degree of

Doctor of Philosophy

2026

To Mum and Dad.

A drop compared to the infinite sea of your love and sacrifice.

Acknowledgements

O you covered up (in your clothes)! Arise and warn (all). Revere your Lord (alone). Purify your garments. (Continue to) shun idols. Do not do a favour expecting more (in return). And persevere for (the sake of) your Lord.

- Al-Muddaththir (74:1–7)

For the unfortunate wanderer who has somehow found their way here, let me first say that you are welcome. I hope you enjoy what is about to come and that this brief detour leaves a pleasant memory behind.

That said, I should add something. This might be the first part of the thesis that you read, but it is the last part that I am writing. In that sense, it serves as an ending. Everyone says that all good stories need an ending, but no one tells you how difficult it is to write one. So, as much as one might strive for originality, I will have to follow the usual routine here. However, although what follows may seem a little formulaic, I assure you that it comes from the heart. So, here we go.

First, I would like to thank my supervisor, András Juhász, who went above and beyond in guiding me from my very first day. He was the person who introduced me to low-dimensional topology and, throughout that process, showed remarkable patience in response to my inexperience. Without his broad knowledge and his calm, clear exposition, finding my way into the subject would have been much more difficult.

When it came to research, I am truly grateful for his willingness to engage seriously with my ideas, including the naive ones. He always took the time to explain why something worked, or why it did not, and this was invaluable for a young researcher who often lacked confidence in his own judgement. I relied heavily on his guidance when exploring new directions and learning how to evaluate my own progress.

Over the years, he did his best to turn me into an independent mathematician. Even so, the prospect of a tomorrow in which I will no longer knock on his office door every week is a rather frightening one. I am deeply grateful for everything he has done for me, and I hope to make him proud as I take my next steps.

As a final remark, I would like to repeat a sentence that concluded all of my term reports during my doctoral studies, and a motto that I will carry with me for the rest of my mathematical life:

"In the end, if Professor Juhász is satisfied with my progress, so am I."

I was also fortunate to benefit from the mentorship of several senior researchers throughout my doctoral studies. I am deeply grateful that they took an interest in my work and were always generous with their time, advice, and expertise.

Among them, I would especially like to thank Tamás Kálmán and Claudius Zibrowius. Their insights were invaluable, not only in helping me learn new bits of mathematics, but also in sustaining my motivation through their genuine interest in my work and ideas. I also burdened them more than once during the academic job application process, yet they invariably responded with remarkable patience, support, and kindness. I am deeply grateful for all their help.

I am also grateful to Livio Ferretti, Marco Golla, Matthew Hedden, Sunkyung Kang, Marco Marengon, Brendan Owens, Jacob Rasmussen, and Vera Vértesi for their generosity with their time, their interest in my work, and the many helpful conversations we shared over the years.

I would also like to thank Marc Lackenby and Steven Sivek for agreeing to examine my thesis. I particularly appreciate their careful reading, thoughtful comments, and kind and encouraging words.

The Oxford Mathematical Institute became a third home to me during my DPhil. I have been fortunate to spend these years in a department that is not only mathematically outstanding, but also welcoming and supportive. I have many fond memories of my time there, and I will miss it dearly.

In particular, I would like to thank my colleagues in the topology group for helping to create a healthy, warm, and mathematically rich environment.

The many conversations, seminars, reading groups, lunches, and shared struggles made these years far more enjoyable than they otherwise would have been. Among the many people who contributed to this atmosphere, I would especially like to thank Adam, Adele, Adrià, Cameron, Colin, Filippo, Filippos, Misha, Paweł, Sam, and Zsombor.

Among these names, I would like to give special thanks to two people.

To Colin, my younger academic brother: it was a pleasure to have someone around with whom I share so many mathematical interests and a similar mathematical taste. Your unwavering enthusiasm and energy were always something I looked forward to. I am especially glad that we eventually had the opportunity to collaborate formally, and I hope that our mathematical connection continues to grow stronger in the years to come.

To Misha, thank you for enduring me as your office mate. I am very glad that we went through the DPhil process together, and I am grateful for your constant willingness to answer my questions about paperwork, academic life, and everything in between. I hope you enjoyed our shared time in the office as much as I did. Our office eventually became the office of unusual working hours, caffeine, sweets, and gossip, as captured by our unofficial (and highly inaccurate) slogan:

"NEC ORA NEC LABORA."

I would like to express my gratitude to the Clarendon Scholarship and the Soudavar Fund Scholarship for supporting me financially throughout my doctoral studies. I am equally grateful to Linacre College for welcoming me into its community. In particular, I would like to thank our Accommodation Manager, Karen Morris, who was extraordinarily kind and flexible, and helped remove a great deal of stress from my life. Her kindness never went unnoticed.

I would like to thank our Senior Tutor, Jane Hoverd, who went above and beyond in supporting me during a particularly difficult period of my DPhil. As I have told many people over the years, Jane is the sweetest person I have met in Oxford. I will certainly miss our conversations, and I remain deeply grateful for her compassion, encouragement, and support.

My thanks also go to the Oxford University Counselling Service, and in particular to Sappho Morgan, for going above and beyond in helping me through that difficult period. Their support and professionalism made a genuine difference at a time when I needed it most.

My gratitude also goes to Arash Rastegar and Amir Jafari, who were my mentors during my undergraduate years at Sharif University of Technology. My mathematical journey began under their kind and patient guidance. Looking back, it is difficult to overstate the impact they had on my decision to pursue mathematics and on the path that eventually led me here.

Finally, and above all, I wish to express my gratitude to my family. Mum, Dad, and Sepehrdokht, thank you for giving me the best family and the best home I could have asked for.

You were with me through every step of these past years. The thing I looked forward to most, at all times and more than anything else, was speaking with you. These years spent away from home have only made me realise how much I cherish you, and indeed that I cherish you above all else. Without you, none of this would have been possible. I hope that you are proud of me. More than that, I hope that I can someday return even a single drop of your immense love and sacrifice.

My *Aghajan* and my grandma Homa also deserve my heartfelt thanks for constantly keeping me in their prayers. I am likewise grateful to my dear aunt Atefe and uncle Masood for their constant expressions of love whenever we had the chance to speak. I miss all of you immensely and look forward to the day we meet again.

I also want to remember and thank my dear *Mamanbozorg*, grandpa Eidi, and uncle Kamal. It has been a long time since I last spoke with any of you, but I know that you are seeing this and, as always, showing me love beyond measure. I only wish that I could have told you about this myself.

I want to close by extending my gratitude to my broader family: the people of my country, Iran, together with our kin and friends across the Global South. In recent years, my country and its people have endured war and tragedy. It pains me that I was not there to share in their suffering, but I want to say that my heart will always be with them. I will always take pride in their beauty, resilience, and strength.

Abstract

This thesis studies structural patterns in knot invariants, with a particular focus on the Alexander polynomial and knot Floer homology. The guiding theme is that knot invariants often reveal additional structure when studied within special families of knots or under natural operations. The thesis develops this theme through two projects. The first concerns the Fox trapezoidal conjecture for alternating knots, while the second concerns twisting and stabilization phenomena in knot Floer homology.

The first part of the thesis studies the coefficients of the Alexander polynomial of alternating knots. Motivated by the recent proof of the Fox trapezoidal conjecture for special alternating knots, we extend these results beyond the special alternating setting. In particular, we prove the conjecture for alternating knots obtained as diagrammatic plumbings of special alternating knots. We also establish further coefficient inequalities for alternating 3-braids.

The second part of the thesis investigates stabilization in knot Floer homology for twist families of knots. Using bordered Floer homology, we introduce a combinatorial framework for studying stabilization phenomena under a non-zero linking condition. This yields general stabilization results for extremal knot Floer homology groups, the τ-invariant, and the total dimension of knot Floer homology. Together, these results uncover new structural patterns in Floer-theoretic knot invariants that were inaccessible through previous methods.

Statement of Originality

I declare that, to the best of my knowledge, the work presented in this thesis is original, except where otherwise indicated, and that no part of this material has been submitted for examination for any other degree.

Chapter 3 is based on joint work with András Juhász and Tamás Kálmán, as indicated in the thesis. My contributions to this work included leading the development of several key arguments, preparing the first full draft of the corresponding preprint, and participating in subsequent revisions and refinements through regular discussions with my co-authors.

The work presented in Chapter 4 is my own, except where otherwise indicated.

Declaration of generative AI use.

I confirm that I have read the University and MPLS Division guidance on generative AI and have complied with it in preparing this thesis.

I have not used generative AI to generate mathematical content, proofs, theorem statements, examples, figures, literature-review text, or any substantive part of the argument.

ChatGPT Edu was used only to assist with spelling, grammar, LaTeX syntax checking, and reference formatting. Its suggestions were used only for local corrections and did not generate or alter substantive mathematical content.

Contents

1 Introduction **1**

1.1 The Fox trapezoidal conjecture 2

1.2 Twisting and Stabilization 4

2 Background **8**

2.1 Knots and Knot Invariants 8

2.1.1 Classical Knot Invariants 8

2.1.2 Operations on Knots 12

2.1.3 Alternating Knots 19

2.2 Heegaard Floer Theory 29

2.2.1 Heegaard Floer Homology 30

2.2.2 Knot Floer Homology 35

2.3 Bordered Floer Theory 41

2.3.1 Working Guide to Bordered Floer Theory with Torus Boundary 42

2.3.2 Working Guide to Immersed Curve Invariants 56

2.3.3 Examples and Computations 64

3 The Fox Trapezoidal Conjecture **78**

3.1 A Survey of Known Results 81

3.2 Alternating Links with Large Twist Regions 86

3.3 Plumbings of Special Alternating Links 93

3.4 Alternating Three-Braids 107

4 Twisting and Stabilization in Knot Floer Homology **120**

4.1 Overview of the Proofs 127

4.2 Stabilization of Total Dimension 131

4.3 The Stabilization Framework 133

4.3.1 General structure of $\mathcal{X}_m^K$ 133

4.3.2 Gradings in $\mathcal{X}_m^K$ 139

4.3.3 Minimality and Homogeneity of Cycles 142
4.3.4 Length and Size of Minimal Cycles 144
4.3.5 Length and Size of Minimal Relations 148
4.3.6 Homology of $\mathcal{X}_m$: Algebra 151
4.3.7 Homology of $\mathcal{X}_m$: Combinatorics 158
4.4 Extremal Knot Floer Homology 161
4.4.1 Closed Balls $B_i^\bullet$ and Shift maps $R_\pm$ 161
4.4.2 Extremal knot Floer homologies fall in a closed ball 164
4.4.3 Proof of Theorem 4.3 165
4.5 The τ-Invariant 170
4.5.1 Reformulation of $\tau(K_m)$ 170
4.5.2 The set of elements with Maslov grading zero 171
4.5.3 Stabilization of maximizers of the Alexander grading 176
4.5.4 Proof of Theorem 4.4 183

Bibliography **184**

Chapter 1

Introduction

Oh, of course, this only classifies 0% of knots—but 0% is a start!

- John Conway

The starting point of knot theory was the problem of classifying and distinguishing knots. This became the main motivation behind the development of knot invariants. In a sense, the theory began with the crossing number, implicit in the earliest knot tabulations, and later expanded to include invariants such as the knot group, the Seifert genus, and the Alexander polynomial. Although these invariants arise from relatively basic ideas in algebraic and differential topology, they have proven to be remarkably rich, and many important questions about them remain open.

The theory of knot invariants has grown considerably since then. Nowadays, the study of knot invariants can be viewed as an area of research in its own right. Modern knot invariants, such as knot homologies, can be used not only to distinguish knots, but also to obtain important results in other areas of low-dimensional topology, such as the construction of exotic smooth structures in four-manifold topology.

Broadly speaking, research on knot invariants can be divided into several overlapping directions. One direction is the construction of new invariants. This includes both the introduction of entirely new knot invariants and the extraction of new invariants from existing ones. It also includes the construction of invariants of 3- and 4-manifolds from knot-theoretic data. Among the most influential examples in recent decades are the Jones polynomial, Khovanov homology, and knot Floer homology. A more recent example is the introduction of skein lasagna modules, which produce invariants of smooth four-manifolds from knot homologies.

A second direction is the application of these invariants to problems in knot theory and low-dimensional topology. Examples include genus detection, concordance questions, and the construction of exotic smooth structures on four-manifolds. Perhaps

the most prominent example of such a story in knot theory is the resolution of the Tait conjectures using the Jones polynomial.

A third direction is the study of the invariants themselves. This includes understanding their internal structure, the structural patterns they exhibit, the relationships between different invariants, and their behaviour under natural operations and within special families of knots. This is where most of the work in this thesis lies.

The three directions are, however, deeply intertwined, and progress in one often leads to developments in the others. Although the primary focus of this thesis is the third direction, several of the results discussed here also contribute to the first two.

This thesis revolves around two distinct projects. The first is concerned with a classical knot invariant, the Alexander polynomial, and the structural patterns it exhibits within a particularly well-behaved family of knots. The second is concerned with a modern invariant, knot Floer homology, and its behaviour under natural operations on knots. Although the techniques and motivations behind these projects are quite different, both fit within the broader theme of understanding the structure and behaviour of knot invariants. Moreover, the second project can also be viewed through the lens of categorification. In this context, this means extending results and phenomena from classical knot invariants to knot homologies.

To avoid unnecessary repetition, detailed introductions and historical discussions are given at the beginning of Chapters 3 and 4. Here we only provide a brief overview of the main ideas and results.

Chapter 2 provides the necessary background for the thesis. We review the classical knot invariants that will appear throughout our discussion, together with several natural operations on knots, including twisting, satellite constructions, and Murasugi sums. We also introduce alternating knots and their associated combinatorial structures. We then give a concise overview of Heegaard Floer theory, knot Floer homology, bordered Floer theory, and the immersed curve invariant. These topics provide the foundation for the two main projects of the thesis.

1.1 The Fox trapezoidal conjecture

The first part of this thesis concerns alternating knots. This is one of the most well-studied classes of knots. Many classical conjectures were first resolved in this setting. Moreover, several knot invariants that are difficult to compute in general admit particularly explicit descriptions for alternating knots. Nevertheless, many important questions about alternating knots remain open.

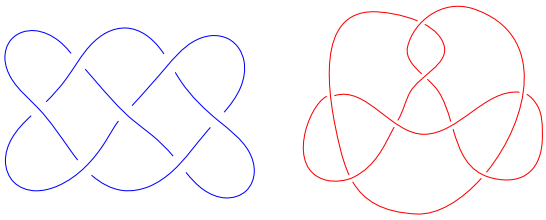


Figure 1.1: Left: an alternating knot, **Right:** a non-alternating knot

Among the oldest open problems in the area is the Fox trapezoidal conjecture, proposed by Fox [Fox61] in 1962. The conjecture states that the absolute values of the coefficients of the Alexander polynomial of an alternating knot form a *trapezoidal sequence*: strictly increasing, possibly constant, and then strictly decreasing.

Conjecture 1.1. *Let K be an alternating knot with $\Delta_K(t) = \sum_{j=0}^{2n} a_j t^j$. Then there exists an integer $r \in \{0, \dots, n\}$ such that*

$$|a_0| < |a_1| < \cdots < |a_{n-r}| = \cdots = |a_{n+r}| > |a_{n+r+1}| > \cdots > |a_{2n}|.$$

Stoimenow [Sto05] later strengthened this to the condition of *log-concavity*. Despite decades of work, these conjectures remain unresolved in full generality.

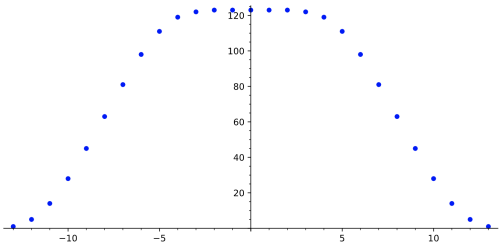


Figure 1.2: Trapezoidal behaviour of the coefficients of the Alexander polynomial of an alternating knot

The conjecture has been proved only for restricted classes, such as *two-bridge knots* [Har79], *arborescent alternating knots* [Mur85], and alternating knots of genus less than five [Sto16]. More recently, following Huh's breakthrough work on log-concavity, the conjecture has attracted renewed attention.

In 2023, Hafner, Mészáros, and Vidinas [HMV24] proved Fox's conjecture for *special alternating knots*. This is a particularly important subclass. Indeed, a theorem of Murasugi implies that every alternating knot decomposes as a *diagrammatic Murasugi sum* of special alternating knots. A diagrammatic Murasugi sum is a special case of the more general *Murasugi sum*, a natural cut-and-paste operation on knots.

The main goal of Chapter 3 is to extend these ideas beyond special alternating knots. This is done in two different directions. The first and main direction is to study the behaviour of the Alexander polynomial under certain diagrammatic Murasugi sums. This also ties the chapter to one of the broader themes of the thesis, namely the behaviour of knot invariants under natural operations on knots and links.

The main strategy is to find a decomposition formula expressing the Alexander polynomial of a diagrammatic Murasugi sum in terms of the Alexander polynomials of its summands. One can then try to run an inductive argument. However, such a formula appears to be out of reach in full generality. In the special case of diagrammatic plumbings it becomes possible to carry out such an argument. This leads to the following theorem.

Theorem 1.2 ([AJK24])**.** *The trapezoidal conjecture holds for any alternating knot that decomposes as a diagrammatic plumbing of special alternating knots.*

The second direction is more combinatorial in nature. It investigates whether the core ideas introduced by Hafner, Mészáros, and Vidinas for special alternating knots can be extended to non-special alternating knots. While a complete extension appears difficult, these methods can be adapted to alternating 3-braids, leading to the following partial result.

Theorem 1.3 ([AJK24])**.** *Let K be the alternating 3-braid obtained from the braid word $\sigma_1^{p_1}\sigma_2^{-q_1}\cdots\sigma_1^{p_n}\sigma_2^{-q_n}$. Let*

$$P = p_1 + \cdots + p_n \qquad \text{and} \qquad Q = q_1 + \cdots + q_n.$$

If $\Delta_K(t) = \sum_{j=0}^{2n} a_j t^j$, then

$$a_{j+1} > \frac{a_j}{n} \qquad \text{for} \qquad 0 \leq j \leq \max(P-2, Q-2).$$

The techniques used in both directions are primarily combinatorial in nature. They draw on classical knot theory, graph-theoretic interpretations of knot diagrams, and the interaction between these structures and the Alexander polynomial.

1.2 Twisting and Stabilization

The second part of this thesis concerns twisting, a natural operation on knots. Given a knot K and a disc D, one may cut K along D. After performing a full twist along the disc, one reglues to obtain a new knot. Iterating this construction produces a *twist family* $\{K_m\}_{m\in\mathbb{Z}}$, where the integer m records the number of twists.

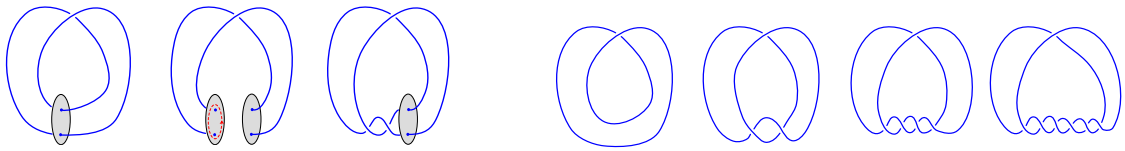


Figure 1.3: Left: twisting, **Right:** first four knots in the resulting twist family.

A central question is how knot invariants behave as the twisting parameter tends to infinity. In many situations one observes various forms of *stabilization*: invariants follow a predictable pattern, converging to a limit, becoming constant, or varying linearly. Such phenomena appear throughout low-dimensional topology, from hyperbolic geometry to quantum topology.

In the case of the Alexander polynomial, the classical work of Kidwell [Kid82] shows that the extremal coefficients of $\Delta_{K_m}(t)$ stabilize. More precisely, for each $n \in \mathbb{N}$, the first n non-trivial coefficients of $\Delta_{K_m}(t)$ become constant once m is sufficiently large.

It is natural to ask whether similar stabilization phenomena occur for more sophisticated knot invariants. Through the lens of *categorification*, the natural candidate is knot Floer homology.

Knot Floer homology was introduced by Ozsváth and Szabó [OS04b] building on the groundbreaking insights of Donaldson and Seiberg–Witten theory. Together with Khovanov homology, they have been the most influential developments in knot theory in the twenty-first century. Among its many remarkable properties are the following:

- It takes the form of a doubly graded vector space

$$\widehat{\mathrm{HFK}}(K) = \bigoplus_j \widehat{\mathrm{HFK}}(K, j), \qquad \widehat{\mathrm{HFK}}(K, j) = \bigoplus_i \widehat{\mathrm{HFK}}_i(K, j),$$

  where j denotes the Alexander grading and i the homological grading.
- It categorifies the Alexander polynomial:

$$\Delta_K(t) = \sum_j \chi\big(\widehat{\mathrm{HFK}}(K, j)\big) t^j.$$

- It detects the Seifert genus:

$$g(K) = \max\{j \mid \widehat{\mathrm{HFK}}(K, j) \neq 0\}.$$

- It gives rise to a numerical concordance invariant denoted by $\tau(K)$. This invariant yields a combinatorial proof of the Milnor conjecture, a landmark result previously obtained only through gauge theory.

	t^0	t^1	t^2	t^3	t^4	t^5	t^6	t^7	t^8	...
Δ_{K_0}	-1	3	-5	7	-5	3	-1	0	0	...
Δ_{K_1}	-1	4	-6	7	-6	4	-1	0	0	...
Δ_{K_2}	1	-1	-1	4	-6	7	-6	4	-1	...
Δ_{K_3}	1	-1	1	-1	-1	4	-6	7	-6	...
Δ_{K_4}	1	-1	1	-1	1	-1	-1	4	-6	...
Δ_{K_5}	1	-1	1	-1	1	-1	1	-1	-1	...
Δ_{K_6}	1	-1	1	-1	1	-1	1	-1	1	...
$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\ddots$

Figure 1.4: Extremal coefficients of Alexander polynomial $\Delta_{K_m}(t)$ become constant as $m \to \infty$

It is natural to ask whether the stabilization of the Alexander polynomial admits a categorified analogue. This was first observed in Hedden's work [Hed05, Hed09] on cabling and later extended by Lambert-Cole [LC18a] to two-strand twisting. A related stabilization phenomenon concerns the τ-invariant: Van Cott [VC10] and Roberts [Rob11] showed that $\tau(K_m)$ becomes a linear function of m as $m \to \infty$.

These results, however, relied on the *coherence assumption*, requiring all twisted strands to be oriented in the same direction. Moreover, the existing methods were inherently local in nature, capturing only limited portions of knot Floer homology and preventing a global analysis of $\widehat{\mathrm{HFK}}(K_m)$ across all Alexander gradings.

The main goal of Chapter 4 is to overcome these limitations. The key new ingredient is the introduction of bordered Floer homology into the twisting problem, leading to a unified proof strategy that subsumes all previous stabilization results in knot Floer homology and establishes a framework for further study of the subject.

The first step in this direction is to replace the strong coherence assumption with the much weaker *non-zero linking condition*

$$\mathrm{lk}(K, \partial D) \neq 0.$$

The resulting theory yields the most general stabilization statements to date for knot Floer homology, as summarized below.

Theorem 1.4 ([Aza25])**.** *Let $\{K_m\}$ be a twist family satisfying the non-zero linking condition. Then, for every $k \in \mathbb{N}$, the k extremal knot Floer homology groups stabilize for sufficiently large m, up to a fixed shift in homological grading. More precisely, there exists an integer $F \in \mathbb{Z}$ such that*

$$\widehat{\mathrm{HFK}}_i\big(K_m, g(K_m) - \ell\big) \cong \widehat{\mathrm{HFK}}_{i+F}\big(K_{m+1}, g(K_{m+1}) - \ell\big)$$

for all sufficiently large m and all $\ell = 0, 1, \ldots, k-1$.

Theorem 1.5 ([Aza25]). *Let $\{K_m\}$ be a twist family satisfying the non-zero linking condition. Then the τ-invariant is eventually linear in the twisting parameter. More precisely, there exist integers $T, t \in \mathbb{Z}$ such that*

$$\tau(K_m) = T \cdot m + t$$

for all sufficiently large m.

Theorem 1.6 ([Aza25]). *Let $\{K_m\}$ be a twist family satisfying the non-zero linking condition. Then there exist integers $D, d \in \mathbb{Z}$ such that, for all sufficiently large m,*

$$\dim \widehat{\mathrm{HFK}}(K_m) = D \cdot m + d.$$

As mentioned above, the key innovation of Chapter 4 is the introduction of *bordered Floer homology* into the twisting problem. Developed by Lipshitz, Ozsváth, and Thurston [LOT18], this framework extends Floer-theoretic invariants to 3-manifolds with boundary. This makes it particularly well suited to cut-and-paste techniques. The backbone of the theory is formed by $\mathcal{A}_\infty$-structures, and the resulting algebraic machinery is quite elaborate. However, the approach developed in this thesis translates the relevant structures into a combinatorial framework. As a result, many stabilization phenomena can ultimately be reduced to graph-theoretic and discrete-geometric arguments.

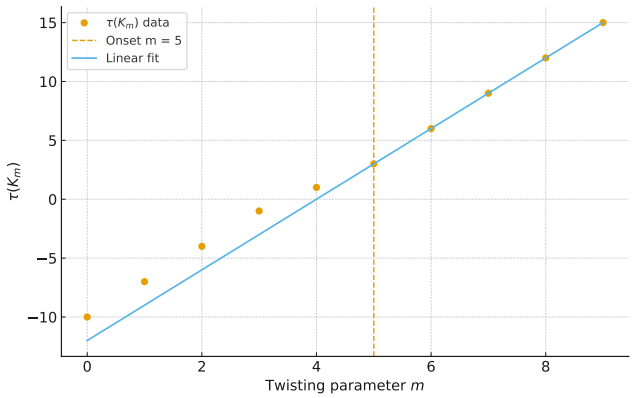


Figure 1.5: Stabilization of $\tau(K_m)$ to a linear function

Together, these advances provide the most general stabilization theorems to date and uncover new structural patterns in Floer-theoretic invariants that had previously remained out of reach.

Chapter 2

Background

2.1 Knots and Knot Invariants

2.1.1 Classical Knot Invariants

In this subsection we recall the classical knot invariants relevant to this thesis, and indicate how they are related to one another and what aspects of the topology of a knot they detect. For further details, see [Juh23].

Two fundamental examples of classical knot invariants are the *Seifert genus* and the *Alexander polynomial*, both arising from the topology of the knot exterior. Recall that the Seifert genus $g(K)$ is defined as the minimal genus of a connected, orientable surface $F \subset S^3$ with boundary $\partial F = K$. Such a surface is called a *Seifert surface* for K. The Seifert genus can be interpreted in terms of the genus function (i.e. Thurston norm) of the knot exterior $E_K := S^3 \setminus \mathrm{int}(N(K))$.

On the other hand, the Alexander polynomial admits several equivalent descriptions. From a topological perspective, it arises from the universal abelian cover of E_K. The deck transformation group of this cover is isomorphic to

$$\pi_1(E_K)/\pi_1(E_K)' \cong H_1(E_K) \cong \mathbb{Z},$$

and hence the cover is infinite cyclic. As a result, the first homology of this cover naturally carries the structure of a $\mathbb{Z}[t^{\pm 1}]$-module. The Alexander polynomial Δ_K is then defined from a presentation of this module. One such presentation can be obtained from a Seifert surface via the *Seifert form*, allowing Δ_K to be computed directly from a Seifert surface; see [Ale28, BZH14].

Remark 2.1. The Alexander polynomial is defined only up to multiplication by $\pm t^k$. In this manuscript, we use Δ_K to denote a normalization chosen so that $\Delta_K \in \mathbb{Z}[t]$

and has positive constant term. We also use the notation $\widetilde{\Delta}_K \in \mathbb{Z}[t^{\pm 1}]$ to denote the *symmetrized Alexander polynomial*.

The Seifert form description yields an easily computable lower bound for the Seifert genus, since the degree of Δ_K satisfies

$$\deg \Delta_K \leq 2g(K).$$

In this thesis we will primarily approach the Alexander polynomial from a combinatorial and diagrammatic perspective. This viewpoint originates in the work of Conway [Con70], who introduced *skein relations* and observed that the Alexander polynomial is determined by its behaviour under local modifications of knots. Although such relations already appear implicitly in Alexander's original work, Conway's formulation made their role explicit. More broadly, Conway's perspective of viewing knot invariants as functions on a "space" of knots, determined by their "local" behaviour, has had a deep and lasting influence on the study of modern knot invariants.

Building on these ideas, Kauffman [Kau83] gave the first explicit combinatorial description of the Alexander polynomial via his *state-sum formula*. In this formulation, the Alexander polynomial is expressed as a sum over combinatorial configurations, or states, associated to a knot diagram. We will recall this construction in the following.

Before doing so, we recall the definition of Tait graphs and fix notation that will be used throughout the discussion.

Definition 2.2. A knot diagram is said to be *decorated* if it has a marked edge, or equivalently, two marked adjacent regions. Fix a checkerboard colouring of the regions of a decorated link diagram. The black (resp. white) *Tait graph* has vertices given by the black (resp. white) regions, with an edge between two regions for each crossing they share. We denote these graphs by B and W, respectively. The graph B (resp. W) is rooted at the vertex corresponding to the marked region; see Figure 2.1. Note that B and W are planar duals.

Fix a decorated diagram D of the knot K and consider the associated Tait graphs B and W. Let $\mathcal{T}_B$ (resp. $\mathcal{T}_W$) denote the set of rooted spanning trees of B (resp. W), oriented away from the root. By planar duality, there is a natural one-to-one correspondence between $\mathcal{T}_B$ and $\mathcal{T}_W$. We denote by

$$\mathcal{D}_{(B,W)} := \{\, (T, T') \in \mathcal{T}_B \times \mathcal{T}_W \mid T \text{ and } T' \text{ are planar duals} \,\}$$

the set of dual pairs.

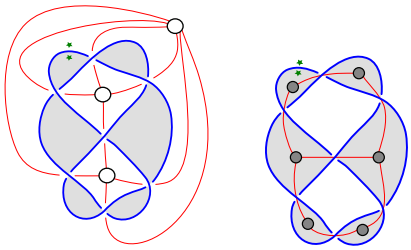


Figure 2.1: The white (left) and black (right) Tait graphs in a marked knot diagram.

Definition 2.3. An element $(T, T') \in \mathcal{D}_{(B,W)}$ is called a *Kauffman state* in D.

Kauffman's state-sum formula expresses the Alexander polynomial as a weighted sum over Kauffman states. The weights are assigned as follows. Each directed edge of a Tait graph passes through a crossing c and is directed towards one of the four corners around c. Figure 2.2 shows the weights assigned to each corner. The weight of a rooted tree is defined as the product of the weights assigned to its edges.

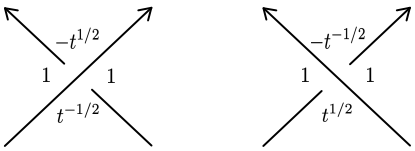


Figure 2.2: Kauffman weights.

Theorem 2.4 (Kauffman state-sum formula)**.** *The symmetrized Alexander polynomial is given by*

$$\widetilde{\Delta}_K(t) = \sum_{(T,T') \in \mathcal{D}_{(B,W)}} w(T)\, w(T').$$

Remark 2.5. Using the Kauffman state-sum formula, the skein relations can be recovered via a deletion–contraction argument.

Motivated by the Seifert genus, one can consider surfaces bounded by a knot in one higher dimension. Viewing S^3 as the boundary of D^4, we study smooth, properly embedded, orientable surfaces $F \subset D^4$ with $\partial F = K$. Such surfaces are called *slice surfaces*, and the minimal genus among them is the *slice genus* $g_4(K)$. An obvious example of a slice surface is obtained by pushing the interior of a Seifert surface into D^4, which shows that

$$g_4(K) \leq g(K).$$

Studying slice surfaces is considerably more subtle than studying Seifert surfaces. One approach is to place them in Morse position with respect to the radial height function on D^4. Decomposing the surface according to its critical points yields a sequence of links in S^3, called a *movie*, whose transitions consist of isotopies, *band attachments*, and *births* and *deaths*.

So far we have focused on invariants arising from the topology of knots and their exteriors. There are, however, invariants that originate purely from the combinatorics of knot diagrams and are equally fundamental. We briefly discuss some of these in the following.

The simplest measure of the complexity of a knot diagram is its number of crossings. For a diagram D, let $c(D)$ denote the number of crossings of D. The *crossing number* of a knot K is defined by

$$c(K) = \min\{\, c(D) \mid D \text{ is a diagram of } K \,\}.$$

A more subtle diagrammatic invariant is the *unknotting number*. Any knot diagram D can be transformed into a diagram of the unknot by a sequence of crossing changes. Let $u(D)$ denote the minimal number of crossing changes required to do so. The *unknotting number* of a knot K is defined by

$$u(K) = \min\{\, u(D) \mid D \text{ is a diagram of } K \,\}.$$

Remark 2.6. The unknotting number can also be defined sequentially by allowing crossing changes interspersed with isotopies (i.e. Reidemeister moves). For classical knots, this definition agrees with the one given above.

It is immediate from the definitions that $u(K) \leq c(K)$. Moreover, each crossing change can be realized by a pair of band attachments, and hence an unknotting sequence gives rise to a slice surface. In particular, this shows that

$$g_4(K) \leq u(K).$$

In a similar vein, we introduce another diagrammatic invariant related to genus. Given a diagram D, the Seifert algorithm produces a Seifert surface for the knot, called the *canonical Seifert surface* associated to D. Let $g_c(D)$ denote the genus of this surface. The *canonical genus* of a knot K is defined by

$$g_c(K) = \min\{\, g_c(D) \mid D \text{ is a diagram of } K \,\}.$$

By construction, one has $g(K) \leq g_c(K)$.

Remark 2.7. All of the inequalities discussed above can be strict. For example, the trefoil knot satisfies $u(K) < c(K)$, while the stevedore knot 6_1 is slice and has unknotting number 1, so that $g_4(K) < u(K)$ and $g_4(K) < g(K)$. There also exist knots for which $g(K) < g_c(K)$; for instance, the untwisted Whitehead double of the trefoil provides such an example [Mor86]. This can be detected using the *HOMFLY polynomial*, a more refined skein-theoretic invariant of knots.

Remark 2.8. Many of the invariants discussed above can be defined in essentially the same way for links. In particular, for an oriented link L, the Seifert genus $g(L)$, slice genus $g_4(L)$, canonical genus $g_c(L)$, and crossing number $c(L)$ admit direct analogues.

The situation for the Alexander polynomial is more subtle. For links, there are two notions of the Alexander polynomial: a multivariable version and a single-variable version, both of which admit algebraic interpretations. In this thesis, we will only consider the single-variable Alexander polynomial of links, which can be defined via the same Kauffman state-sum formula (up to multiplication by $\pm t^{\frac{k}{2}}$).

2.1.2 Operations on Knots

We now recall several operations on knots that will play a role in this thesis. These operations allow one to construct new knots from existing ones and to study how various invariants behave under such constructions.

We start with the most fundamental example, namely the connected sum. Given two knots K_1 and K_2, their connected sum $K_1 \# K_2$ is obtained by first choosing disjoint diagrams of the two knots. One then removes a small arc from each and joins the resulting endpoints by two parallel arcs disjoint from the rest of the diagram. This operation is well-defined up to isotopy. Topologically, this can be characterized by the existence of an embedded *separating 2-sphere* in S^3 intersecting $K_1 \# K_2$ transversely in exactly two points, such that each of the summands lies in one of the complementary components.

Recall that knots can be viewed as boundaries of surfaces. From this perspective, one can construct the connected sum $K_1 \# K_2$ by taking *boundary connected sums* of Seifert surfaces, gluing them along a boundary arc.

The Seifert surface perspective can be made precise and further generalized, leading to the notion of *Murasugi sum*. This is fundamentally a gluing operation on Seifert surfaces, and by passing to their boundaries, it induces an operation on links. The definition is given below.

Definition 2.9. Let L_1 and L_2 be links in S^3 with Seifert surfaces S_1 and S_2. Let $S^3 = B_1 \cup_{S^2} B_2$ be the standard decomposition of the 3-sphere into two 3-balls. Assume that $S_i \subset B_i$ for $i = 1, 2$, and that $S_1 \cap S_2 = D$ is a disk in S^2.

Suppose that D is the interior of a $2n$-gon, i.e., ∂D is decomposed into $2n$ arcs

$$a_1, b_1, \ldots, a_n, b_n,$$

labelled counterclockwise, such that for each $i \in \{1, \ldots, n\}$:

- $a_i \subset \partial S_1$ and is properly embedded in S_2,
- $b_i \subset \partial S_2$ and is properly embedded in S_1.

The compact oriented surface $S = S_1 \cup_D S_2$ is called the *Murasugi sum* of S_1 and S_2 along D. The link $L = \partial S$ is called the Murasugi sum of L_1 and L_2 along D; see Figures 2.3 and 2.4. The integer $2n$ is called the *gonality* of the Murasugi sum.

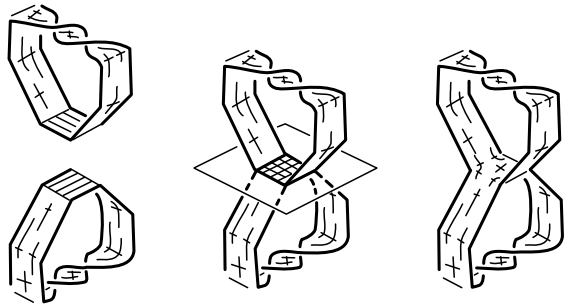


Figure 2.3: A 4-gon Murasugi sum; i.e., plumbing [Kaw96].

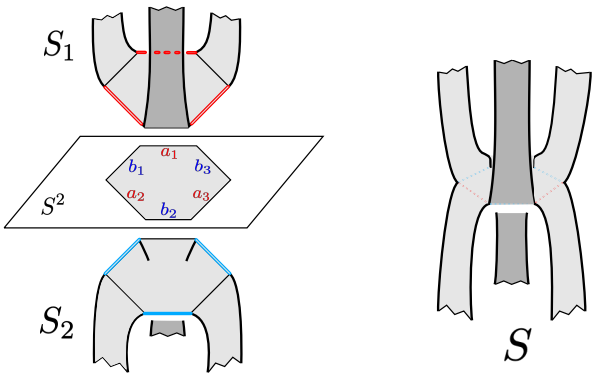


Figure 2.4: A 6-gon Murasugi sum [AH21].

Taking a 2-gon Murasugi sum is simply the connected sum operation. The 4-gon Murasugi sum operation is called *plumbing*.

We next consider satellite constructions, in which the ambient 3-manifold is modified via a cut-and-paste operation. A fundamental example is given by the satellite construction.

Let $C \subset S^3$ be a knot, called the *companion*, and let $P \subset S^1 \times D^2$ be a knot, called the *pattern*, which is not contained in a 3-ball and not isotopic to $S^1 \times \{0\}$. Let $N(C)$ be a tubular neighbourhood of C. Given a framing $n \in \mathbb{Z}$, choose a diffeomorphism

$$f : S^1 \times D^2 \to N(C)$$

realizing this framing. The resulting knot $f(P) \subset S^3$ is called the *n-twisted satellite* of C with pattern P, and is denoted by $C(P, n)$.

When the pattern P is a (p, q)-torus knot in $S^1 \times D^2$ and the framing is zero, the resulting satellite is called the (p, q)-*cable* of C, and is denoted by $C_{p,q}$.

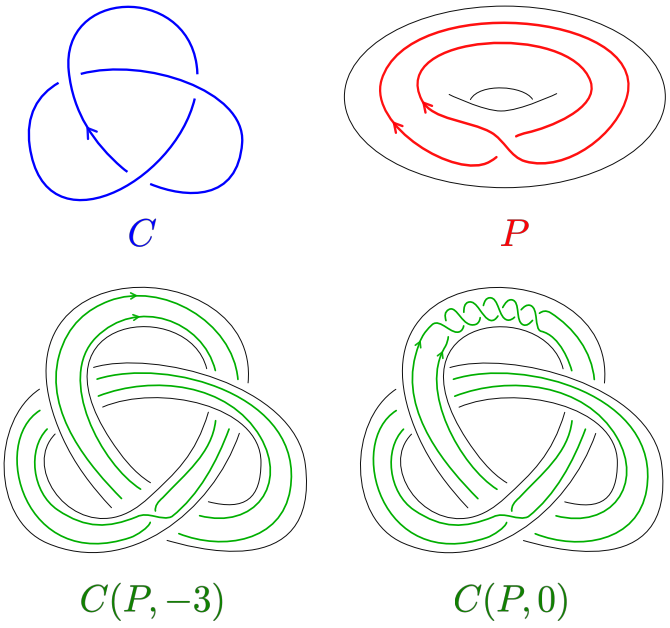


Figure 2.5: Two satellites of the trefoil knot with $(2, 1)$-torus knot pattern. The second example is the $(2, 1)$-cable of the trefoil.

Passing from framing n to framing $n + 1$ in a satellite construction is realized by precomposing the diffeomorphism f with a meridional Dehn twist of the solid torus. At the level of diagrams, this appears as twisting the strands of the satellite knot, as illustrated in Figure 2.5. This viewpoint admits a natural generalization, leading to the following definition.

Definition 2.10. Let $D \subset S^3$ be an embedded disk such that D meets the knot K transversely and $\partial D \cap K = \varnothing$. For each $n \in \mathbb{Z}$, let K_n denote the knot obtained by

cutting along D, rotating one side by $-2\pi n$, and regluing along D by the identity. The family $\{K_n\}_{n\in\mathbb{Z}}$ is called a *twist family*.

Construction 2.11. There are several equivalent ways to describe twisting. Let $c = \partial D$ be the unknot bounding the disk D. Then K can be viewed as a knot in the exterior E_c, which is an unknotted solid torus naturally identified with $S^1 \times D^2 \subset S^3$.

From this perspective, the knots K_n are obtained by applying the $(-n)$-th power of the meridional Dehn twist

$$\tau : S^1 \times D^2 \longrightarrow S^1 \times D^2,$$

that is,

$$K_n = \tau^{-n}(K) \subset E_c \subset S^3.$$

Alternatively, the knots K_n can be described via Dehn surgery, namely by performing $(1/n)$-surgery on c. Since c is an unknot, the resulting manifold $S^3_c(1/n)$ is diffeomorphic to S^3. Under this identification, the knot K is sent to K_n.

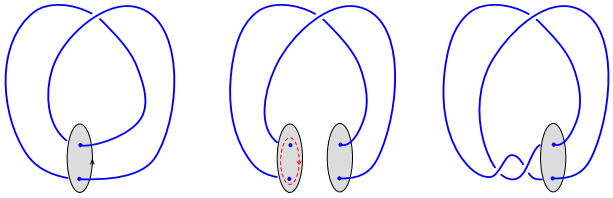


Figure 2.6: Twisting a knot along a disk D by a full -2π rotation.

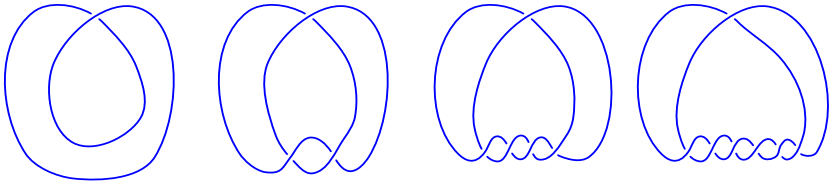


Figure 2.7: The knots K_0, K_1, K_2, K_3 in the resulting twist family.

To prepare for later results, we introduce some notation.

Notation 2.12. Let $\mathrm{wind}_c(K)$ and $\mathrm{wrap}_c(K)$ denote the winding number and wrapping number of K in the solid torus E_c, respectively. Recall that

$$\mathrm{wind}_c(K) = |\operatorname{lk}(K, c)|, \qquad \text{and} \qquad 0 \leq \mathrm{wind}_c(K) \leq \mathrm{wrap}_c(K).$$

For later use, let

$$E_{K\cup c} := \overline{S^3 \setminus N(K \cup c)}.$$

Let $\hat{D} := D \cap E_{K\cup c}$, which is a properly embedded planar surface in $E_{K\cup c}$. Let x denote the Thurston norm on

$$H_2(E_{K\cup c}, \partial E_{K\cup c}).$$

Then

$$\mathrm{wind}_c(K) - 1 \leq x([\hat{D}]) \leq \mathrm{wrap}_c(K) - 1.$$

Definition 2.13. We say that c links K *coherently* if

$$\mathrm{lk}(K, c) = \mathrm{wind}_c(K) = \mathrm{wrap}_c(K).$$

In this case, the associated twist family $\{K_m\}$ is called a *coherent twist family*.

We now discuss how the invariants introduced earlier behave under these operations.

Starting with connected sums, the Seifert genus is additive:

$$g(K_1 \# K_2) = g(K_1) + g(K_2),$$

and the Alexander polynomial is multiplicative:

$$\Delta_{K_1\#K_2}(t) = \Delta_{K_1}(t)\,\Delta_{K_2}(t).$$

When it comes to Murasugi sum, the situation is more complicated. In fact, Able and Hirasawa [AH21] show that any knot can be constructed via Murasugi sums from any two given knots. More precisely, they prove the following result:

Theorem 2.14 ([AH21])**.** *For any three knots $K_1, K_2, K_3 \subset S^3$, there exist Seifert surfaces F_1, F_2, F_3, such that F_3 is a Murasugi sum of F_1 and F_2.*

The key is to impose restrictions on the Seifert surfaces used in the Murasugi sum. In particular, when the surfaces are incompressible or of minimal genus, one obtains strong structural results. The first such results were established in the influential work of Gabai [Gab83].

Theorem 2.15 ([Gab83])**.** *The Murasugi sum of two incompressible (resp. minimal genus) Seifert surfaces is incompressible (resp. minimal genus). As a result, genus is additive under Murasugi sums along minimal genus Seifert surfaces.*

However, even under these restrictions, the behaviour of the Alexander polynomial under Murasugi sum does not admit a simple formula. More explicit results require additional hypotheses, which will be discussed in Chapter 3.

We now describe the behaviour of invariants under satellite constructions. To this end, we introduce some notation related to the pattern knot $P \subset S^1 \times D^2$. Similar to the twisting case, let $\mathrm{wind}(P)$ denote the winding number of P in the solid torus, i.e., the absolute value of its algebraic intersection with a meridian disk.

A *relative Seifert surface* for P is a compact orientable surface in $S^1 \times D^2$ whose interior is disjoint from P and whose boundary consists of P together with $\mathrm{wind}(P)$ disjoint, coherently oriented longitudes on $\partial(S^1 \times D^2)$. The *relative genus* $g(S^1 \times D^2, P)$ is defined as the minimal genus among all such surfaces.

Theorem 2.16 (Schubert [Sch53])**.** *Let $C \subset S^3$ be a nontrivial knot and let $P \subset S^1 \times D^2$ be a pattern. Then*

$$g(C(P,0)) = \mathrm{wind}(P) \cdot g(C) + g(S^1 \times D^2, P).$$

Remark 2.17. Schubert's formula can also be used to compute the genus of satellites with other framings. As discussed, changing the framing corresponds to replacing the pattern P by its n-twist $\tau^n(P)$, which alters the relative genus term accordingly.

We now describe the behaviour of the Alexander polynomial under satellite constructions.

Theorem 2.18 ([BZH14])**.** *Let $C \subset S^3$ be a knot and let $P \subset S^1 \times D^2$ be a pattern with winding number $w = \mathrm{wind}(P)$. Then,*

$$\Delta_{C(P,0)}(t) = \Delta_C(t^w) \cdot \Delta_P(t),$$

Here $\Delta_P(t)$ denotes the Alexander polynomial of P viewed as a knot in S^3 via the standard unknotted embedding of $S^1 \times D^2$ into S^3.

We now turn to the behaviour of invariants under twisting. In general, there is no simple formula describing invariants across an entire twist family $\{K_m\}_{m \in \mathbb{Z}}$. However, the asymptotic behaviour of these invariants as $m \to \infty$ is of particular interest.

In many cases, one observes a phenomenon of *stabilization*, whereby invariants eventually follow a predictable pattern, for example becoming constant or, more generally, varying linearly with respect to m. Although this phenomenon manifests differently for different invariants, it appears throughout low-dimensional topology, from hyperbolic geometry to quantum topology.

We begin by describing the stabilization phenomenon of the Seifert genus. In particular, the genus stabilizes to an affine linear function of m.

Theorem 2.19 ([BT19]). *Let $w = \mathrm{wind}_c(K)$. There exists a constant $G = G(K, c)$ such that*

$$2g(K_m) = x([\hat{D}])\, w\, m + 2G$$

for all sufficiently large m.

We now describe the stabilization of the Alexander polynomial. In this case, it is important to distinguish between whether $\mathrm{wind}_c(K)$ vanishes or not. If $\mathrm{wind}_c(K) = 0$, then the Alexander polynomial is an affine linear function of m for all $m \in \mathbb{Z}$.

The more interesting case is when $\mathrm{wind}_c(K) \neq 0$. In this case, we say that the twist family satisfies the *non-zero winding condition*. Here one observes a different form of stabilization, namely in the *extremal coefficients*: for each $n \in \mathbb{N}$, the first n non-trivial coefficients of $\Delta_{K_m}(t)$ become constant for sufficiently large m.

Theorem 2.20. *Let $\{K_m\}$ be a twist family of knots and suppose that $\mathrm{wind}_c(K) \neq 0$. Write*

$$\Delta_{K_m}(t) = \sum_{i \geq 0} a_{m,i}\, t^i.$$

Then for each $n \in \mathbb{N}$, there exists $M_n \in \mathbb{N}$, such that

$$a_{m,\ell} = a_{m+1,\ell}$$

for all $\ell = 0, 1, \ldots, n-1$ and all $m \geq M_n$.

An example of the stabilization phenomenon for the Alexander polynomial is shown in Figure 2.8. Note that the i-th row lists the coefficients of $\Delta_{K_i}(t)$.

Remark 2.21. In fact, one can give a general formula for Δ_{K_m} for all m in terms of the multivariable Alexander polynomial $\Delta_{K \cup c}$. This first appeared in the work of Kidwell [Kid82] and was later observed by Furusawa and Sakuma [FS83].

Theorem 2.20 can be deduced from this formula and has been regarded as a folklore result. It has been established in various special cases by Lambert-Cole [LC18a], Chen [Che22], and Boninger [Bon25]. Its most recent and general appearance is in the work of the present author [Aza25].

We will revisit the stabilization phenomenon in detail in Chapter 4.

	t^0	t^1	t^2	t^3	t^4	t^5	t^6	t^7	t^8	...
Δ_{K_0}	-1	3	-5	7	-5	3	-1	0	0	...
Δ_{K_1}	-1	4	-6	7	-6	4	-1	0	0	...
Δ_{K_2}	1	-1	-1	4	-6	7	-6	4	-1	...
Δ_{K_3}	1	-1	1	-1	-1	4	-6	7	-6	...
Δ_{K_4}	1	-1	1	-1	1	-1	-1	4	-6	...
Δ_{K_5}	1	-1	1	-1	1	-1	1	-1	-1	...
Δ_{K_6}	1	-1	1	-1	1	-1	1	-1	1	...
$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\ddots$

Figure 2.8: An illustration of the stabilization of the extremal coefficients of $\Delta_{K_m}(t)$ as described in Theorem 2.20.

2.1.3 Alternating Knots

Alternating knots form one of the most well-studied families in knot theory. Their definition places them at a natural intersection of topology and combinatorics, leading to rich structural properties.

Many invariants of alternating knots admit particularly nice descriptions, and even those that are typically difficult to compute become more accessible in this setting. Several classical conjectures in knot theory, such as the Tait conjectures and the Fox trapezoidal conjecture, are formulated for alternating knots. Moreover, important general problems, such as the cosmetic surgery conjecture, were first resolved in this class. Despite these advances, alternating knots continue to present interesting open questions and remain a central object of study in knot theory.

We begin with the definition of alternating diagrams and alternating knots.

Definition 2.22. A knot or link diagram is said to be *alternating* if the crossings alternate between over and under as one travels along each component of the diagram. A knot or link is *alternating* if it admits such a diagram.

Alternating diagrams also admit a useful characterization in terms of checkerboard colourings. We record this in the following proposition.

Proposition 2.23. *A non-split diagram D is alternating if and only if it admits a checkerboard colouring with the following local property: at each crossing, as one travels along the overstrand, the black region changes from the left side of the strand to the right; see Figure 2.9. We call such a colouring a standard checkerboard colouring of D.*



Figure 2.9: Standard colouring of a crossing in an alternating diagram.

This seemingly simple diagrammatic property has many important consequences. For example, the crossing number of an alternating link is particularly easy to determine: any reduced alternating diagram of a link L realizes its crossing number $c(L)$. This result was part of the celebrated Tait conjectures and was proved by Murasugi and Thistlethwaite, building on the discovery of the Jones polynomial.

Similarly, the computation of genus is particularly simple for alternating links. This fact follows from independent work by Crowell [Cro59], Murasugi [Mur58], and Gabai [Gab86], and is summarised in the following theorem.

Theorem 2.24. *Let L be an alternating link and let D be a reduced alternating diagram of L. Let $\mu(L)$ denote the number of components of L. Then*

$$g(L) = g_c(L) = \frac{1}{2}\big(\deg \Delta_L - \mu(L) + 1\big).$$

Moreover, the canonical Seifert surface associated to D has minimal genus.

Murasugi's work on the structure of canonical Seifert surfaces of alternating links led to a deeper understanding of alternating knot theory. In particular, it led to the introduction of an important subclass of *special alternating links*. We introduce this subclass and discuss its role in the theory throughout the rest of this section. We begin by defining *special diagrams*.

Definition 2.25. Let D be a diagram of a link L, consisting of crossings $c_1, \dots, c_n$ and arcs $e_1, \dots, e_{2n}$. Equivalently, D may be viewed as a 4-valent planar graph with vertices c_i, edges e_j, and regions $r_1, \dots, r_{n+2}$. Applying Seifert's algorithm to D, one obtains a collection of Seifert cycles $T_1, \dots, T_q$. A Seifert cycle T_i is called *type 1* if it bounds a region r_k of the diagram, and *type 2* otherwise. If all Seifert cycles are of type 1, then D is called a *special* diagram.

At first glance, it may not be clear why special diagrams are interesting. A first step toward understanding their role is to give some equivalent characterizations. We first recall the definition of Tait surfaces.

Definition 2.26. Let D be a diagram of a link L, equipped with a checkerboard colouring. The black (resp. white) Tait surface F_B (resp. F_W) is obtained by taking the black (resp. white) regions as disks in S^2 and attaching half-twisted bands at the crossings according to the diagram. These satisfy

$$\partial F_B = \partial F_W = L.$$

Note that the Tait surfaces are not necessarily orientable. In fact, the black (resp. white) Tait surface F_B (resp. F_W) is orientable if and only if the graph B (resp. W) is bipartite.

Now assume that a non-split diagram D is special. Note that no two Seifert cycles can bound regions of different colours. It follows that the Seifert surface obtained from Seifert's algorithm coincides with one of the two Tait surfaces, and hence one of the two Tait surfaces must be orientable. In fact, for knot diagrams, this necessary condition is also sufficient. This leads to the following proposition.

Proposition 2.27. *Let D be a diagram of a non-split link L. If D is special, then one of the two Tait surfaces associated to D is orientable, and the corresponding Tait graph is bipartite. Conversely, if L is a knot and one of the two Tait surfaces associated to D is orientable, then D is special.*

Proposition 2.27 shows why special diagrams are useful. However, special diagrams by themselves do not define a particularly restrictive family of links: it is a well-known fact that every link admits a special diagram. See Stoimenow [Sto16] and Hirasawa [Hir95] for algorithms that turn any diagram of a link into a special one; an example is shown in Figure 2.10.

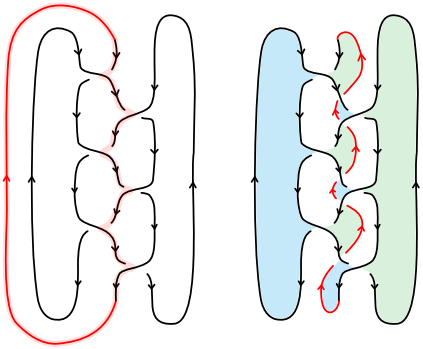


Figure 2.10: Non-special and special diagrams of a link. The left diagram is non-special as it contains a type 2 Seifert cycle, coloured yellow. One can turn the diagram into the special diagram on the right by moving the red segment.

If we require a diagram to be both alternating and special, we obtain an interesting family of links.

Definition 2.28. A knot or link is called *special alternating* if it admits a diagram that is both alternating and special.

Before discussing the structural role of special alternating links, we begin with a simple characterization. Recall that a link L is called *positive* (resp. *negative*) if it admits a diagram in which all crossings have positive (resp. negative) sign.

Proposition 2.29. *Let L be a non-split alternating link. Then L is special alternating if and only if L is positive or negative.*

Sketch of proof. Let D be a special alternating diagram. Suppose that there are two consecutive crossings of different signs. As illustrated in Figure 2.11, the boundaries of three distinct Seifert cycles appear in a neighbourhood of these crossings. These Seifert cycles cannot all bound regions of the diagram, and hence cannot all be of type 1. Therefore, all crossings in a special alternating diagram have the same sign.

For the converse, let L be a positive alternating link. By a result of Nakamura [Nak00], there exists an alternating diagram D of L in which all crossings have positive sign. Using Proposition 2.23, one sees that at each crossing the Seifert cycles run along the boundaries of the black regions in the standard checkerboard colouring (Figure 2.9). Since this holds at every crossing of D, all Seifert cycles bound black regions, and hence D is special. The negative case follows by taking mirror images. □

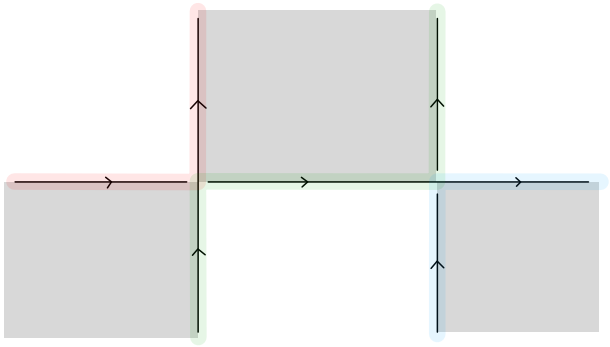


Figure 2.11: Two consecutive crossings of different signs and the Seifert cycles in their neighbourhood.

Proposition 2.29 shows, in particular, that not every alternating link is special alternating.

We now turn to Murasugi's work. We begin by describing *Murasugi's decomposition algorithm.*

Construction 2.30 (Murasugi's decomposition algorithm)**.** Let L be a link with connected diagram D, and use the notation of Definition 2.25. Let $C_1, \dots, C_m$ be the type 2 Seifert cycles in D. Each C_i is a simple closed curve in S^2 that divides S^2 into disks C_i^+ and C_i^-. Let C_i^+ be the disk containing ∞. For $\gamma_i \in \{+, -\}$, let $D(\gamma_1, \cdots, \gamma_m)$ be the closure of $C_1^{\gamma_1} \cap \cdots \cap C_m^{\gamma_m}$. Exactly $m+1$ of these sets will be nonempty, which we call *Seifert domains* and denote by $D_1, \dots, D_{m+1}$; see Figure 2.12.

For every $i \in \{1, \dots, n\}$, at least two of the four regions adjacent to crossing c_i belong to the same Seifert domain D_j for some $j \in \{1, \dots, m+1\}$. In this case, we say that c_i *belongs* to D_j. Each crossing belongs to exactly one Seifert domain.

For $j \in \{1, \dots, m+1\}$, let L_j be the link obtained by smoothing all the crossings along ∂D_j that do not belong to D_j. We call $\{L_1, \dots, L_{m+1}\}$ the *Murasugi decomposition* of the diagram D, and denote it by

$$L = L_1 * L_2 * \cdots * L_{m+1}.$$

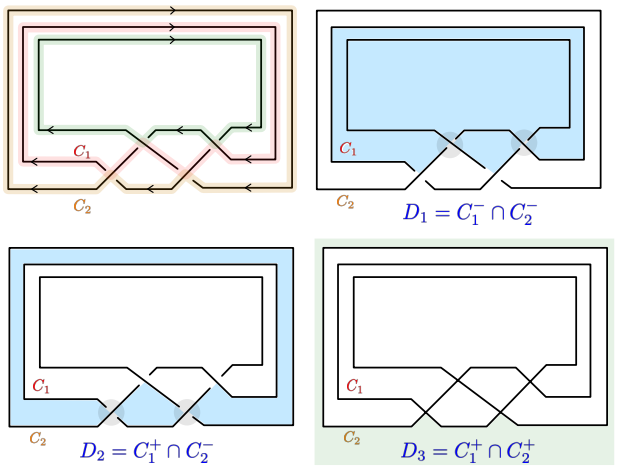


Figure 2.12: Seifert domains and their crossings. Note that none of the crossings belong to D_3.

The notation

$$L = L_1 * L_2 * \cdots * L_{m+1}$$

is slightly subtle. The construction above produces the links $L_1, \dots, L_{m+1}$ from the diagram D, but it does not immediately define a binary operation on links. More accurately, the decomposition should be understood as a gluing operation organised over a tree, analogous to plumbing trees. For our purposes, however, it is enough to describe the binary step from which the full decomposition is obtained inductively.

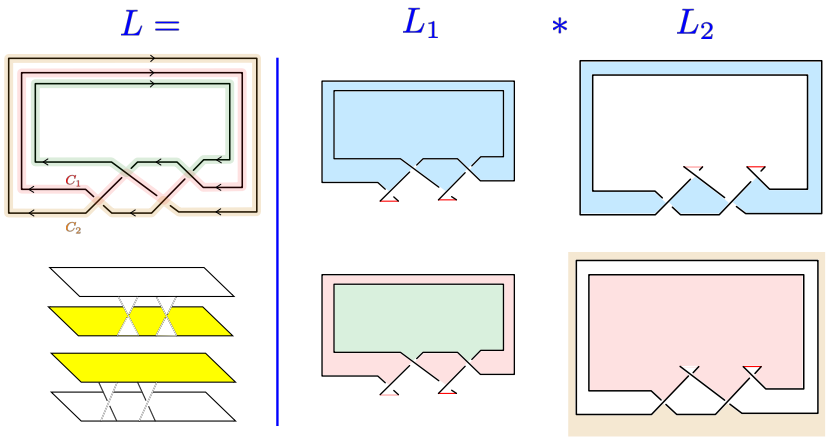


Figure 2.13: An example of a Murasugi decomposition. The first row shows two of the Seifert domains and the associated Murasugi summands. Note that the third summand L_3 is trivial and is hence not depicted. The leftmost picture in the second row shows the Murasugi decomposition as a Murasugi sum (related to Figure 2.14). Other figures in the second row depict Seifert surfaces of the Murasugi summands constructed by cutting the canonical Seifert surface of L.

Definition 2.31. Let D be a connected diagram of a link L, and let C be a type 2 Seifert cycle. Let C^+ and C^- be the two disks in S^2 bounded by C, with C^+ containing ∞. Let L^+ and L^- be the links obtained by restricting the diagram of L to C^+ and C^-, respectively, smoothing the crossings on C as in the Murasugi decomposition algorithm. We write

$$L = L^- * L^+$$

and call this decomposition a $*$-*product over the Seifert cycle* C.

Remark 2.32. With the convention of Definition 2.31, the notation

$$L = L_1 * L_2 * \cdots * L_{m+1}$$

is understood inductively. Each $*$-product occurs over a type 2 Seifert cycle. More precisely, if a type 2 Seifert cycle C separates the Seifert domains into

$$D_1, \ldots, D_i \subseteq C^- \quad \text{and} \quad D_{i+1}, \ldots, D_{m+1} \subseteq C^+,$$

then, by definition, $L = L^- * L^+$. Applying the Murasugi decomposition algorithm to L^- and L^+ gives

$$L^- = L_1 * \cdots * L_i \quad \text{and} \quad L^+ = L_{i+1} * \cdots * L_{m+1}.$$

Thus the full Murasugi decomposition is obtained by iterating $*$-products over type 2 Seifert cycles.

Remark 2.33. One can construct the links L^- and L^+ by adding vertical segments in the middle of the half-twisted bands used in Seifert's algorithm; see Figure 2.13. Cutting the canonical Seifert surface of L along these segments produces canonical Seifert surfaces for the links L^- and L^+. Conversely, gluing the canonical Seifert surfaces for L^- and L^+ recovers the canonical Seifert surface of L.

This gluing is in fact a Murasugi sum in the sense of Definition 2.9. More precisely, the $*$-product corresponds to a Murasugi sum along canonical Seifert surfaces and along Seifert disks.

We do not prove this fact in full generality, but illustrate it in the example of Figure 2.13; see Figure 2.14. Another example is shown in Figure 2.15.

Notation 2.34. The name $*$-product was introduced by Murasugi. We will use the term *diagrammatic Murasugi sum* interchangeably with $*$-product. In view of the previous remarks, this is a natural choice of terminology.

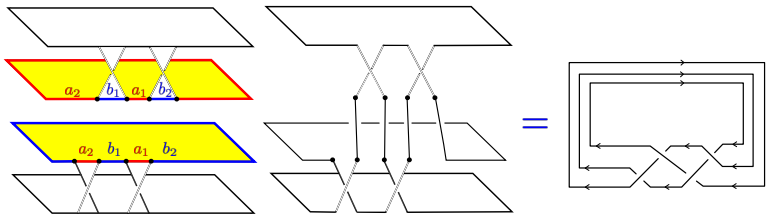


Figure 2.14: The $*$-product is a special case of Murasugi sum (shown in the example of Figure 2.13).

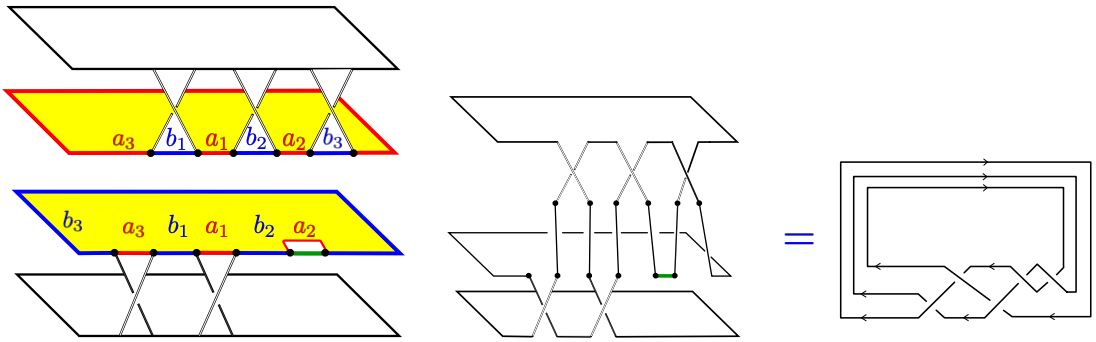


Figure 2.15: A second example showing that $*$-product is a special case of Murasugi sum. Note how the gluing disk of the Murasugi sum is slightly different from the one in Figure 2.14

We now discuss the relation between Murasugi's decomposition and special diagrams, together with its consequences for alternating links. Assume that we have a Murasugi decomposition

$$L = L_1 * L_2 * \cdots * L_{m+1}.$$

It follows naturally from Construction 2.30 that the diagrams obtained for the summands L_i are special, since the process exhausts all type 2 Seifert cycles.

Moreover, when the algorithm is applied to a reduced alternating diagram, one can show that the resulting diagrams are alternating as well, and hence the summands L_i are special alternating. This leads to the following celebrated result of Murasugi.

Theorem 2.35 (Murasugi [Mur65])**.** *Any alternating link can be decomposed, via the Murasugi decomposition, as a $*$-product of special alternating links. Furthermore, the genus is additive under this decomposition.*

In this sense, special alternating links serve as building blocks for alternating links. This perspective gives a more flexible view of alternating link theory, whose diagrams are often regarded as highly rigid. Furthermore, it provides a natural strategy for inductive arguments: one first proves results for special alternating links, and then extends them to arbitrary alternating links via the Murasugi decomposition. We will make essential use of this perspective in Chapter 3.

We conclude this discussion of diagrammatic Murasugi sums with a remark on their gonality.

Remark 2.36. We use the notation of Definition 2.31. Let the gonality of the Murasugi sum over the Seifert cycle C be $2s$. One can compute $2s$ directly from the diagram as follows.

Crossings on C are of two types: those belonging to C^- and those belonging to C^+ (shown in red and blue in Figure 2.16). Moving along C, one alternately encounters x_j crossings in C^- and y_j crossings in C^+, for $j = 1, \ldots, l$, where l is some positive integer. With this notation, we have

$$s = \sum_{j=1}^{l} x_j + \sum_{j=1}^{l} y_j - l. \tag{2.1}$$

The integer l appearing in Remark 2.36 will play an important role in our later discussion. This motivates the following definition, inspired by the notion of the length of an alternating 3-braid [AC21].

Definition 2.37. The integer l appearing in Remark 2.36 is called the *length* of the diagrammatic Murasugi sum.

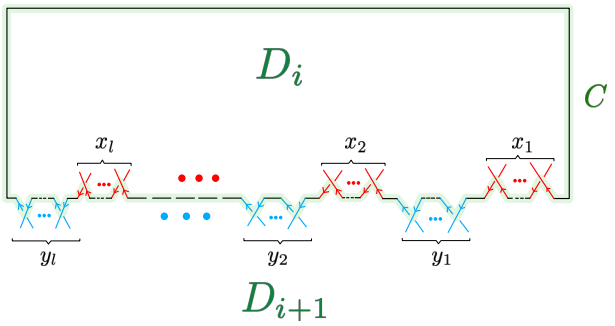


Figure 2.16: Local picture of a type 2 Seifert cycle C.

For the final part of this subsection, we return to the Alexander polynomial and discuss several of its special properties in the setting of alternating links.

Theorem 2.24 shows that the Alexander polynomial of an alternating link determines its genus. Beyond this, the polynomials themselves exhibit additional striking patterns. These were discovered in foundational works of Crowell and Murasugi, and are summarised in the following theorem.

Theorem 2.38 (Crowell–Murasugi)**.** *Let*

$$\Delta_L(t) = \sum_{i=0}^{n} a_i t^i$$

be the Alexander polynomial of a non-split alternating link L. Then the coefficients of $\Delta_L(t)$ alternate in sign and are all nonzero, i.e.

$$a_i a_{i+1} < 0$$

for all $i \in \{0, \dots, n-1\}$.

Recall from Remark 2.1 that we may assume $a_0 > 0$. It follows that the polynomial $\Delta_L(-t)$ has positive coefficients, while still encoding all the information of the Alexander polynomial.

Crowell presented a state-sum formula for the polynomial $\Delta_L(-t)$. As we will see later, it is useful to move between the perspectives of the Crowell and Kauffman state-sums. We first describe Crowell's formula below.

Construction 2.39 (Crowell graph)**.** Let D be an alternating diagram of a link L with m crossings $c_1, \dots, c_m$. We assume that $m \geq 2$.

Consider the underlying planar graph of the diagram D, which we denote by G_D and call the *Crowell graph* associated to D. The vertices of G_D are the crossings $c_1, \dots, c_m$, and the edges are the arcs of the diagram D.

We orient the edges of this graph such that the two under-crossing (resp. over-crossing) arcs at each crossing correspond to incoming (resp. outgoing) edges. Furthermore, we label the edges with weights t and 1 such that, as one traverses the over-strand at a crossing following the orientation of the link, the weight t (resp. 1) appears on the right (resp. left).

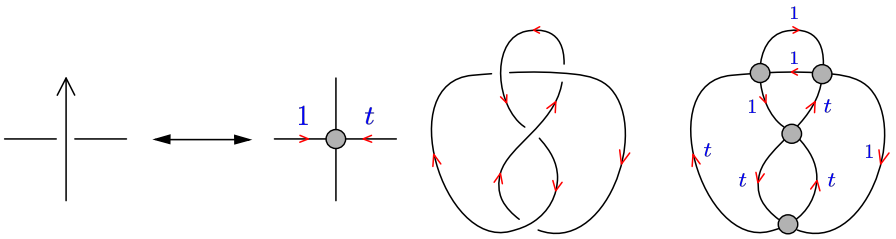


Figure 2.17: The Crowell graph associated to an alternating link diagram.

Choose a vertex c_k of G_D as the root. Let $\mathcal{T}(G_D, c_k)$ denote the set of all maximal oriented trees rooted at c_k. For a tree $T \in \mathcal{T}(G_D, c_k)$, define its weight $W(T)$ to be the product of the weights of all its edges.

Definition 2.40. The *Crowell polynomial* associated to (G_D, c_k) is

$$P_{G_D,c_k}(t) := \sum_{T \in \mathcal{T}(G_D,c_k)} W(T).$$

Theorem 2.41 (Crowell [Cro59])**.** *Let D be an alternating diagram of a non-split link L. Let c_k be a crossing in D. Then*

$$\Delta_L(-t) \doteq P_{G_D,c_k}(t),$$

where $\doteq$ denotes equality up to multiplication by t^m for some $m \in \mathbb{Z}$.

We now conclude our discussion of alternating links. We include a final remark, which hints at the deeper combinatorial structure underlying Crowell's polynomial.

Remark 2.42. More generally, the Crowell polynomial $P_{G,v}(t)$ can be defined for any weighted, oriented graph G rooted at a vertex v.

The non-zeroness phenomenon of Theorem 2.38 extends naturally to the general Crowell polynomial $P_{G,v}(t)$. This follows from the fact that $\mathcal{T}(G, v)$ forms the basis of a *greedoid*, and hence its adjacency graph is connected; see Corollary 5.7 of [BKL85]. In particular, for any two trees $T, T' \in \mathcal{T}(G, v)$, one can pass from T to T' through a sequence of edge exchanges while remaining inside $\mathcal{T}(G, v)$. Since each edge exchange changes the degree of the weight monomial by at most one, every intermediate monomial between $W(T)$ and $W(T')$ is realized by some element of $\mathcal{T}(G, v)$, and hence has nonzero coefficient.

2.2 Heegaard Floer Theory

The whole of human scholarship, from Socrates to Diderot, cannot compete with anything by Zoltán Sz. and Peter O.

- Adam Levine

Low-dimensional topology has been profoundly influenced by the introduction of geometric and analytic methods over the past decades. In dimension four, methods from gauge theory became particularly prominent through Donaldson's application of the *Yang–Mills equations* to the study of smooth four-manifolds. In dimension three, the same equations formed the basis of Floer's construction of what is now called *instanton Floer homology*. Seiberg and Witten later introduced a new set of equations that were easier to work with. Later, Kronheimer and Mrowka developed *monopole Floer homology* for 3-manifolds in this setting. A major difficulty in all of these theories is the computability of the resulting invariants, since their definitions ultimately arise from the study of nonlinear partial differential equations.

Work of Ozsváth and Szabó on understanding and computing the Seiberg–Witten invariant led to the introduction of *Heegaard Floer homology* at the turn of the 21st century. In their construction, gauge-theoretic methods were replaced by ideas from symplectic topology, most notably Floer's other celebrated construction, namely *Lagrangian Floer homology*. The result was a powerful package of invariants for knots, 3-manifolds, and 4-manifolds. In this sense, the theory also reflects another major motivation in the subject, namely Atiyah's *topological quantum field theory (TQFT)* picture and the *Atiyah–Floer conjecture*.

It is worth noting that later work of Kutluhan–Lee–Taubes and Colin–Ghiggini–Honda showed that Heegaard Floer homology and monopole Floer homology are equivalent, although the two theories are adapted to different aspects of 3-manifold topology. In particular, Heegaard Floer homology is based on Heegaard diagrams, possesses a strong combinatorial flavor, and is considerably more computable.

A detailed exploration of Heegaard Floer theory is beyond the scope of this thesis. In this section, we briefly recall the main definitions and properties of Heegaard Floer homology and its knot-theoretic counterpart, knot Floer homology. For a more detailed introduction, we refer the reader to the book of Juhász [Juh23], which serves as an excellent modern introduction to the subject.

2.2.1 Heegaard Floer Homology

Let M be a closed, connected, and oriented 3-manifold. A *Heegaard decomposition* of M consists of a closed, connected, and oriented surface $\Sigma \subset M$ that separates M into two handlebodies H_α and H_β.

Heegaard decompositions can be constructed combinatorially by considering a triangulation of M and letting H_α be a tubular neighborhood of the 1-skeleton. The construction that is more appropriate in our context, however, comes from Morse theory. Given a *self-indexing Morse function* $f\colon M \to [0,3]$, one can define the associated Heegaard surface by

$$\Sigma = f^{-1}\left(\frac{3}{2}\right).$$

The handlebodies H_α and H_β are then obtained from the sublevel and superlevel sets of f, respectively. Equivalently, the handlebody H_α (resp. H_β) corresponds to the union of the 0- and 1-handles (resp. the 2- and 3-handles) in the handle decomposition induced by f.

The data of a Heegaard splitting can be encoded combinatorially. Each of the handlebodies is represented by a collection of curves coming from the belt circles of its 1-handles. This leads to the definition of a *Heegaard diagram*.

Definition 2.43. A *Heegaard diagram* is a triple $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta})$, where Σ is a closed, connected, and oriented surface of genus g, and

$$\boldsymbol{\alpha} = \alpha_1 \cup \cdots \cup \alpha_{g+n} \quad \text{and} \quad \boldsymbol{\beta} = \beta_1 \cup \cdots \cup \beta_{g+n}$$

are collections of pairwise disjoint simple closed curves on Σ such that surgery on Σ along $\boldsymbol{\alpha}$ (resp. $\boldsymbol{\beta}$) yields a disjoint union of $n+1$ spheres.

As described above, a Heegaard decomposition of a 3-manifold M naturally determines a Heegaard diagram. Such a diagram is called a Heegaard diagram for M. Conversely, given a Heegaard diagram $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta})$ for a 3-manifold M, one can recover M by attaching 2-handles to $\Sigma \times [0,1]$ along the curves in $\boldsymbol{\alpha} \times \{0\}$ and $\boldsymbol{\beta} \times \{1\}$, and then capping off the resulting spherical boundary components with 3-balls.

This correspondence is far from one-to-one, as a given 3-manifold can be represented by many different Heegaard diagrams. For example, applying a diffeomorphism of the surface Σ to a Heegaard diagram does not change the diffeomorphism type of the associated 3-manifold.

We say that two Heegaard diagrams $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta})$ and $(\Sigma', \boldsymbol{\alpha}', \boldsymbol{\beta}')$ are *equivalent* if there exists an orientation-preserving diffeomorphism $f\colon \Sigma \to \Sigma'$ such that

$$f(\boldsymbol{\alpha}) = \boldsymbol{\alpha}' \quad \text{and} \quad f(\boldsymbol{\beta}) = \boldsymbol{\beta}'.$$

There are also other simple modifications of a Heegaard diagram that do not change the associated 3-manifold. The most important of these are called *Heegaard moves*. The moves consist of *handle slides*, *stabilization*, and *destabilization*. A handle slide is defined analogously to the corresponding notion for handle decompositions. Stabilization is obtained by taking the connected sum with the standard genus one Heegaard diagram of S^3, while destabilization is the inverse operation.

In fact, a combination of results of Singer and of Scharlemann–Thompson shows that these moves are sufficient to relate any two Heegaard diagrams of the same 3-manifold. This is summarised in the following theorem.

Theorem 2.44. *Let $(\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta})$ and $(\Sigma', \boldsymbol{\alpha}', \boldsymbol{\beta}')$ be Heegaard diagrams of the same 3-manifold M. Then, after applying a sequence of isotopies and Heegaard moves, the resulting diagrams become equivalent.*

For technical reasons arising in the construction of Heegaard Floer homology, one needs to equip Heegaard diagrams with additional basepoint data. This leads to the following notion of a *pointed Heegaard diagram*.

Definition 2.45. A *pointed Heegaard diagram* for a 3-manifold M is a pair $(\mathcal{H}, \boldsymbol{z})$, where $\mathcal{H} = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta})$ is a Heegaard diagram for M, and

$$\boldsymbol{z} = \{z_1, \dots, z_{n+1}\} \subseteq \Sigma \setminus (\boldsymbol{\alpha} \cup \boldsymbol{\beta})$$

is a collection of basepoints such that no two basepoints lie in the same connected component of $\Sigma \setminus (\boldsymbol{\alpha} \cup \boldsymbol{\beta})$.

The notions of equivalence and Heegaard moves extend naturally to pointed Heegaard diagrams.

We now describe the main setting of Heegaard Floer homology. Starting from a pointed Heegaard diagram $(\mathcal{H}, \boldsymbol{z})$, one can associate a symplectic manifold together with a pair of Lagrangian tori. The resulting Lagrangian Floer theory forms the foundation of Heegaard Floer homology.

Construction 2.46. Let $(\mathcal{H}, \boldsymbol{z})$ be a pointed Heegaard diagram.

- The associated *symmetric product* is defined by

$$\mathrm{Sym}^{g+n}(\Sigma) := \underbrace{\Sigma \times \cdots \times \Sigma}_{g+n \text{ times}} / S_{g+n},$$

  where the symmetric group S_{g+n} acts by permuting the factors.

- The collections of curves $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$ determine two tori inside $\mathrm{Sym}^{g+n}(\Sigma)$:

$$\mathbb{T}_\alpha := \alpha_1 \times \cdots \times \alpha_{g+n}, \quad \text{and} \quad \mathbb{T}_\beta := \beta_1 \times \cdots \times \beta_{g+n}.$$

- For each basepoint $z_i \in \boldsymbol{z}$, define the hypersurface

$$V_{z_i} := \{z_i\} \times \mathrm{Sym}^{g+n-1}(\Sigma) \subset \mathrm{Sym}^{g+n}(\Sigma).$$

It might not be immediately clear that $\mathrm{Sym}^{g+n}(\Sigma)$ is a manifold. However, the fundamental theorem of algebra implies that this is indeed the case. In fact, any complex structure j on Σ induces a natural complex structure on the symmetric product, denoted by $\mathrm{Sym}^{g+n}(j)$. Moreover, by a result of Perutz, one can choose compatible symplectic forms on $\mathrm{Sym}^{g+n}(\Sigma)$ with respect to which the tori $\mathbb{T}_\alpha$ and $\mathbb{T}_\beta$ are Lagrangian.

After a small perturbation if necessary, we may assume that the tori $\mathbb{T}_\alpha$ and $\mathbb{T}_\beta$ intersect transversely. The Heegaard Floer chain complex is generated by the intersection points in $\mathbb{T}_\alpha \cap \mathbb{T}_\beta$, which are called *Heegaard states*. These are described more explicitly in the following definition.

Definition 2.47. Let $(\mathcal{H}, \boldsymbol{z})$ be a pointed Heegaard diagram. The set of *Heegaard states* is

$$\mathfrak{G}(\mathcal{H}) := \mathbb{T}_\alpha \cap \mathbb{T}_\beta.$$

Equivalently, the elements of $\mathfrak{G}(\mathcal{H})$ are unordered tuples $\boldsymbol{x} = [x_1, \ldots, x_{g+n}]$ where each x_i is an intersection point between some α-curve and some β-curve, such that each α-curve and each β-curve contains exactly one point of $\boldsymbol{x}$.

The differential in Heegaard Floer homology is defined by counting pseudo-holomorphic Whitney disks connecting Heegaard states. This is the most technically challenging part of the construction. We only briefly recall some of the basic notions and definitions here; for further details, we refer the reader to the book of Juhász [Juh23]. The next two definitions introduce Whitney disks and the associated moduli spaces.

Definition 2.48. Given two Heegaard states $\boldsymbol{x}, \boldsymbol{y} \in \mathfrak{S}(\mathcal{H})$, a *topological Whitney disk* from $\boldsymbol{x}$ to $\boldsymbol{y}$ is a continuous map

$$u\colon D^2 \to \mathrm{Sym}^{g+n}(\Sigma)$$

such that

$$u(\partial D^2 \cap \{\Re(z) \geq 0\}) \subset \mathbb{T}_\alpha, \qquad u(\partial D^2 \cap \{\Re(z) \leq 0\}) \subset \mathbb{T}_\beta,$$

and

$$u(-i) = \boldsymbol{x}, \qquad u(i) = \boldsymbol{y}.$$

We denote by $\pi_2(\boldsymbol{x}, \boldsymbol{y})$ the set of homotopy classes of such Whitney disks.

Definition 2.49. Let $\boldsymbol{x}, \boldsymbol{y} \in \mathfrak{S}(\mathcal{H})$, and let $\phi \in \pi_2(\boldsymbol{x}, \boldsymbol{y})$ be a homotopy class. The moduli space $\mathcal{M}(\phi)$ of pseudo-holomorphic representatives of ϕ is a smooth manifold of dimension $\mu(\phi)$. The integer $\mu(\phi)$ is called the *Maslov index* of ϕ.

In what follows, we only consider homotopy classes for which this moduli space is nonempty. Abusing notation, we will often use ϕ both for a homotopy class and for a pseudo-holomorphic representative of that class.

The moduli space $\mathcal{M}(\phi)$ carries a natural $\mathbb{R}$-action, and hence one can define the *reduced moduli space*

$$\widetilde{\mathcal{M}}(\phi) := \mathcal{M}(\phi)/\mathbb{R}.$$

The reduced moduli space is then a smooth manifold of dimension $\mu(\phi) - 1$. A Whitney disk with Maslov index one is called *rigid*. In this case, the reduced moduli space is a discrete set.

Remark 2.50. The definition of the moduli spaces, and the fact that they are smooth manifolds, depend on the choice of suitable almost-complex structures. Since these analytic details will not play a role in this thesis, we omit them from the discussion.

Note that, up to this point, the basepoints have not played an essential role. The following definition introduces the main additional structure induced by the basepoints.

Definition 2.51. For each basepoint $z_i \in \boldsymbol{z}$, let $n_{z_i}(\phi)$ denote the intersection number of ϕ with the hypersurface V_{z_i}. We also define

$$n_{\boldsymbol{z}}(\phi) := \sum_i n_{z_i}(\phi).$$

Note that every intersection point of a pseudo-holomorphic curve with a complex hypersurface is positive. Hence $n_{z_i}(\phi) \geq 0$ for every $z_i \in \boldsymbol{z}$. In particular,

$$n_{\boldsymbol{z}}(\phi) = 0 \Leftrightarrow n_{z_i}(\phi) = 0 \quad \text{for every } z_i \in \boldsymbol{z}.$$

We are now ready to define Heegaard Floer homology.

Definition 2.52. Let $(\mathcal{H}, \boldsymbol{z})$ be a pointed Heegaard diagram for a 3-manifold M. The *Heegaard Floer chain complex* associated to $(\mathcal{H}, \boldsymbol{z})$ is the $\mathbb{F}_2$-vector space freely generated by the set of Heegaard states $\mathfrak{G}(\mathcal{H})$. Its differential is defined on generators by

$$\widehat{\partial}\,\boldsymbol{x} = \sum_{\boldsymbol{y} \in \mathfrak{G}(\mathcal{H})} \sum_{\substack{\phi \in \pi_2(\boldsymbol{x}, \boldsymbol{y}) \\ \mu(\phi) = 1 \\ n_{\boldsymbol{z}}(\phi) = 0}} \#\widetilde{\mathcal{M}}(\phi) \cdot \boldsymbol{y}.$$

We denote this chain complex by $\widehat{\mathrm{CF}}(\mathcal{H}, \boldsymbol{z})$, and its homology by $\widehat{\mathrm{HF}}(\mathcal{H}, \boldsymbol{z})$.

Ozsváth and Szabó proved that the chain homotopy type of $\widehat{\mathrm{CF}}(\mathcal{H}, \boldsymbol{z})$ is independent of the choice of pointed Heegaard diagram $(\mathcal{H}, \boldsymbol{z})$ for M. Equivalently, the construction is invariant under isotopies and Heegaard moves. Therefore, the chain homotopy type of $\widehat{\mathrm{CF}}(\mathcal{H}, \boldsymbol{z})$ is an invariant of the underlying 3-manifold M, and we denote it by $\widehat{\mathrm{CF}}(M)$. Its homology is denoted by $\widehat{\mathrm{HF}}(M)$.

Remark 2.53. In order for the differential $\widehat{\partial}$ to be well-defined, one needs to impose an additional technical condition on the Heegaard diagram, called *admissibility*. Roughly speaking, admissibility ensures the finiteness of the relevant counts of pseudo-holomorphic disks. Since admissibility will not play a direct role in this thesis, we omit the precise definition.

Example 2.54. Consider the standard genus one Heegaard diagram of S^3. In this case, the tori $\mathbb{T}_\alpha$ and $\mathbb{T}_\beta$ intersect in a single point. Hence the chain complex $\widehat{\mathrm{CF}}(S^3)$ is generated by a single Heegaard state and the differential vanishes. Therefore,

$$\widehat{\mathrm{HF}}(S^3) \cong \mathbb{F}_2.$$

Remark 2.55. Straight from the definition, Heegaard Floer homology is difficult to compute in practice. However, the situation simplifies significantly when the Heegaard surface Σ has genus one and $|\boldsymbol{\alpha}| = |\boldsymbol{\beta}| = 1$. In this case, rigid pseudo-holomorphic disks admit a concrete combinatorial description in terms of embedded bigons in the universal cover of the torus. This allows one to compute the Heegaard Floer homology of lens spaces explicitly.

This perspective becomes even more useful in the context of knot Floer homology and bordered Floer homology, as we shall see later in the thesis.

Heegaard Floer homology enjoys additional algebraic structure. In particular, it decomposes as a direct sum over Spin^c structures on the underlying 3-manifold M. Depending on the topology of M, these groups may also carry relative or absolute homological gradings. We will discuss these gradings further in the context of knot Floer homology in the next subsection.

2.2.2 Knot Floer Homology

Knot Floer homology, introduced independently by Ozsváth–Szabó and Rasmussen, is a refinement of Heegaard Floer homology. It is an invariant of knots in 3-manifolds. Together with the advent of Khovanov homology, these modern knot homologies reshaped knot theory in the 21st century. They recover classical numerical invariants through a phenomenon known as *categorification*, while also giving rise to entirely new invariants. Moreover, they enjoy strong *functorial* properties, making them particularly well suited for the study of *knot concordance* and *knotted surfaces*.

As we will see later in this subsection, knot Floer homology categorifies the Alexander polynomial and detects the Seifert genus. It also gives rise to the τ-invariant, which provides a lower bound for the slice genus. This led to new and eventually combinatorial proofs of the *Milnor conjecture* on the slice genus of torus knots, a landmark result previously obtained through gauge theory.

We now turn to the construction of knot Floer homology. Let L be a link in a closed, connected, and oriented 3-manifold M. The main idea is that the pair (M, L) can also be encoded by a Heegaard diagram equipped with additional basepoint data. This leads to the following notion.

Definition 2.56. A multi-pointed Heegaard diagram $(\mathcal{H}, \boldsymbol{z}, \boldsymbol{w})$ for a pair (M, L) consists of a Heegaard diagram $\mathcal{H} = (\Sigma, \boldsymbol{\alpha}, \boldsymbol{\beta})$ for M, together with two collections of basepoints

$$\boldsymbol{z} = \{z_1, \dots, z_{n+1}\}, \qquad \boldsymbol{w} = \{w_1, \dots, w_{n+1}\},$$

such that each connected component of $\Sigma \setminus \boldsymbol{\alpha}$ and $\Sigma \setminus \boldsymbol{\beta}$ contains exactly one z-basepoint and one w-basepoint.

The link L is recovered as follows. Choose disjoint collections of embedded arcs

$$\gamma_1, \dots, \gamma_{n+1} \subset \Sigma \setminus \boldsymbol{\alpha}, \qquad \delta_1, \dots, \delta_{n+1} \subset \Sigma \setminus \boldsymbol{\beta},$$

where each γ_i and δ_i connects z_i to w_i. After pushing the arcs γ_i into H_α and the arcs δ_i into H_β, the resulting union is the link L.

It is clear from Definition 2.56 that any multi-pointed Heegaard diagram determines a 3-manifold together with a link inside it. The converse is also true: any link $L \subset M$ admits a multi-pointed Heegaard diagram. In essence, this follows from placing L in *bridge position* with respect to the Heegaard surface Σ. A more precise proof is through Morse-theoretic arguments.

Similar to ordinary Heegaard diagrams, this correspondence is not one-to-one. The notions of equivalence and Heegaard moves extend naturally to this setting.

In the remainder of this thesis, we focus on the case of knots in 3-manifolds. We are now ready to define the knot Floer chain complex.

Definition 2.57. Let $(\mathcal{H}, \boldsymbol{z}, \boldsymbol{w})$ be a multi-pointed Heegaard diagram for a pair (M, K). The *knot Floer chain complex* associated to $(\mathcal{H}, \boldsymbol{z}, \boldsymbol{w})$ is the $\mathbb{F}_2$-vector space freely generated by the set of Heegaard states $\mathfrak{G}(\mathcal{H})$. Its differential is defined on generators by

$$\widehat{\partial}_K \boldsymbol{x} = \sum_{\boldsymbol{y} \in \mathfrak{G}(\mathcal{H})} \sum_{\substack{\phi \in \pi_2(\boldsymbol{x}, \boldsymbol{y}) \\ \mu(\phi) = 1 \\ n_{\boldsymbol{z}}(\phi) = n_{\boldsymbol{w}}(\phi) = 0}} \#\widetilde{\mathcal{M}}(\phi) \cdot \boldsymbol{y}.$$

We denote this chain complex by $\widehat{\mathrm{CFK}}(\mathcal{H}, \boldsymbol{z}, \boldsymbol{w})$, and its homology by $\widehat{\mathrm{HFK}}(\mathcal{H}, \boldsymbol{z}, \boldsymbol{w})$.

The chain homotopy type of $\widehat{\mathrm{CFK}}(\mathcal{H}, \boldsymbol{z}, \boldsymbol{w})$ is independent of the choice of multi-pointed Heegaard diagram $(\mathcal{H}, \boldsymbol{z}, \boldsymbol{w})$ for (M, K). Therefore, the chain homotopy type of $\widehat{\mathrm{CFK}}(\mathcal{H}, \boldsymbol{z}, \boldsymbol{w})$ is an invariant of the pair (M, K), and we denote it by $\widehat{\mathrm{CFK}}(M, K)$. Its homology is denoted by $\widehat{\mathrm{HFK}}(M, K)$.

From now on, we restrict to the case of knots in S^3, where knot Floer homology admits additional grading structures. To simplify notation, we will often suppress the ambient manifold S^3 from our notation.

Construction 2.58. Let $(\mathcal{H}, \boldsymbol{z}, \boldsymbol{w})$ be a multi-pointed Heegaard diagram for (S^3, K). There exist grading maps

$$M, A \colon \mathfrak{G}(\mathcal{H}) \to \mathbb{Z},$$

called the *Maslov grading* and the *Alexander grading*, respectively. These gradings satisfy the following: for any $\boldsymbol{x}, \boldsymbol{y} \in \mathfrak{G}(\mathcal{H})$ and any Whitney disk $\phi \in \pi_2(\boldsymbol{x}, \boldsymbol{y})$,

$$M(\boldsymbol{x}) - M(\boldsymbol{y}) = \mu(\phi) - 2n_{\boldsymbol{z}}(\phi), \tag{2.2}$$

$$A(\boldsymbol{x}) - A(\boldsymbol{y}) = n_{\boldsymbol{w}}(\phi) - n_{\boldsymbol{z}}(\phi). \tag{2.3}$$

These gradings endow the knot Floer chain complex with the structure of a bigraded vector space. We denote the homogeneous summand in Maslov grading i and Alexander grading j by $\widehat{\mathrm{CFK}}_i(K, j)$.

Furthermore, it is easy to see that the differential $\widehat{\partial}_K$ preserves the Alexander grading and decreases the Maslov grading by one. Hence the knot Floer homology inherits a bi-grading, denoted by $\widehat{\mathrm{HFK}}_i(K,j)$. We also use the notation

$$\widehat{\mathrm{HFK}}(K,j) = \bigoplus_{i\in\mathbb{Z}} \widehat{\mathrm{HFK}}_i(K,j).$$

Remark 2.59. When $\mathcal{H}$ is a Heegaard diagram for S^3, one can show that $\pi_2(\boldsymbol{x},\boldsymbol{y})$ is nonempty for any $\boldsymbol{x},\boldsymbol{y}\in\mathfrak{G}(\mathcal{H})$. Therefore, Equations (2.2) and (2.3) determine the Maslov and Alexander gradings up to an overall translation. One can further show that these definitions are independent of the choice of Whitney disk ϕ. Later, we will describe how these relative gradings are normalized to obtain absolute gradings.

We now describe how knot Floer homology captures classical invariants of knots. We begin with its relation to the Seifert genus, described in the following theorem. This was first proved by Ozsváth–Szabó [OS04a], and later Juhász [Juh06] gave a conceptually simpler proof using sutured Floer homology.

Theorem 2.60. *Let $K\subset S^3$ be a knot. Then*

$$g(K) = \max\left\{ j\in\mathbb{Z} \;\middle|\; \widehat{HFK}(K,j)\neq 0\right\}.$$

We now turn to the Alexander polynomial.

Theorem 2.61 (Ozsváth–Szabó [OS04b], Rasmussen [Ras03])**.** *Let $K\subset S^3$ be a knot. Then*

$$\sum_{i,j\in\mathbb{Z}} (-1)^i \dim \widehat{HFK}_i(K,j)\ \cdot\ t^j = \widetilde{\Delta}_K(t).$$

This means that knot Floer homology recovers the Alexander polynomial as its *graded Euler characteristic*. In other words, $\widehat{\mathrm{HFK}}$ categorifies $\widetilde{\Delta}_K$.

Furthermore, knot Floer homology satisfies a symmetry property generalizing the symmetry of the coefficients of the Alexander polynomial, as described in the following proposition.

Proposition 2.62 (Conjugation symmetry)**.** *For any $i,j\in\mathbb{Z}$, we have*

$$\widehat{HFK}_i(K,j) \cong \widehat{HFK}_{i-2j}(K,-j).$$

Beyond recovering classical invariants, knot Floer homology also gives rise to new concordance invariants. The most important example for our purposes is the τ-invariant, introduced by Ozsváth–Szabó. We briefly describe the construction following Juhász [Juh23].

The main idea is to observe that if $(\mathcal{H}, \boldsymbol{z}, \boldsymbol{w})$ is a multi-pointed Heegaard diagram for (S^3, K), then $(\mathcal{H}, \boldsymbol{z})$ is a pointed Heegaard diagram for S^3. Hence we can also consider the Heegaard Floer chain complex $\widehat{\mathrm{CF}}(\mathcal{H}, \boldsymbol{z})$. As vector spaces, the two chain complexes $\widehat{\mathrm{CFK}}(\mathcal{H}, \boldsymbol{z}, \boldsymbol{w})$ and $\widehat{\mathrm{CF}}(\mathcal{H}, \boldsymbol{z})$ are identical and generated by the set of Heegaard states $\mathfrak{G}(\mathcal{H})$. Therefore, the Maslov and Alexander gradings on $\widehat{\mathrm{CFK}}(\mathcal{H}, \boldsymbol{z}, \boldsymbol{w})$ can also be viewed as gradings on $\widehat{\mathrm{CF}}(\mathcal{H}, \boldsymbol{z})$.

Moreover, the differential $\widehat{\partial}$ on $\widehat{\mathrm{CF}}(\mathcal{H}, \boldsymbol{z})$ also decreases the Maslov grading by one. As a result, $\widehat{\mathrm{HF}}(\mathcal{H}, \boldsymbol{z}) = \widehat{\mathrm{HF}}(S^3)$ inherits a Maslov grading.

The main difference is that $\widehat{\partial}$ no longer preserves the Alexander grading. However, it still cannot increase it. More precisely,

$$\widehat{\partial} = \widehat{\partial}_K + \text{terms of lower Alexander grading.}$$

This means that $\widehat{\mathrm{CF}}(\mathcal{H}, \boldsymbol{z})$ admits an Alexander filtration as follows. Let

$$\mathcal{F}(K, j) \subseteq \widehat{\mathrm{CF}}(\mathcal{H}, \boldsymbol{z})$$

denote the subspace generated by those $\boldsymbol{x} \in \mathfrak{G}(\mathcal{H})$ satisfying $A(\boldsymbol{x}) \leq j$. Since the differential $\widehat{\partial}$ cannot increase the Alexander grading, each $\mathcal{F}(K, j)$ is preserved by $\widehat{\partial}$. As a result, $\mathcal{F}(K, j)$ inherits the structure of a chain complex. Hence we obtain a filtration of chain complexes

$$\cdots \subseteq \mathcal{F}(K, j-1) \subseteq \mathcal{F}(K, j) \subseteq \mathcal{F}(K, j+1) \subseteq \cdots \subseteq \widehat{\mathrm{CF}}(\mathcal{H}, \boldsymbol{z}).$$

Hence, the inclusions $\mathcal{F}(K, j) \hookrightarrow \widehat{\mathrm{CF}}(\mathcal{H}, \boldsymbol{z})$ induce maps on homology

$$H_*(\mathcal{F}(K, j)) \longrightarrow \widehat{\mathrm{HF}}(S^3).$$

We can now define the τ-invariant.

Definition 2.63. The τ-invariant of a knot $K \subset S^3$ is defined by

$$\tau(K) = \min\left\{ j \in \mathbb{Z} \;\middle|\; H_*(\mathcal{F}(K, j)) \to \widehat{\mathrm{HF}}(S^3) \text{ is nontrivial} \right\}.$$

The τ-invariant satisfies several remarkable properties. The most important for our purposes is that it provides a lower bound for the slice genus, which is a notoriously difficult invariant to compute. This is summarised in the following theorem.

Theorem 2.64 (Ozsváth–Szabó [OS03b])**.** *For any knot $K \subset S^3$,*

$$|\tau(K)| \leq g_4(K).$$

For the (p,q)-torus knot $T_{p,q}$, one can show that

$$\tau(T_{p,q}) = \frac{(p-1)(q-1)}{2}.$$

On the other hand, one can construct a Seifert surface for $T_{p,q}$ of genus $\frac{1}{2}(p-1)(q-1)$. It follows that

$$g_4(T_{p,q}) = g(T_{p,q}) = \frac{(p-1)(q-1)}{2}.$$

This gives a new proof of the celebrated Milnor conjecture on the slice genus of $T_{p,q}$.

Remark 2.65. Going back to Remark 2.59, we can now describe the normalization of the Maslov and Alexander gradings. The Maslov grading is normalized so that the unique generator of

$$\widehat{\mathrm{HF}}(\mathcal{H}, \boldsymbol{z}) = \widehat{\mathrm{HF}}(S^3) \cong \mathbb{F}_2$$

lies in Maslov grading zero. The Alexander grading is then normalized so that

$$|\{\boldsymbol{x} \in \mathfrak{G}(\mathcal{H}) \mid A(\boldsymbol{x}) = j\}| \equiv |\{\boldsymbol{x} \in \mathfrak{G}(\mathcal{H}) \mid A(\boldsymbol{x}) = -j\}| \pmod 2.$$

Equivalently, one can normalize the Alexander grading so that any of Theorem 2.61, Theorem 2.60, or Proposition 2.62 holds.

Remark 2.66. One must still show that the τ-invariant is independent of the choice of multi-pointed Heegaard diagram $(\mathcal{H}, \boldsymbol{z}, \boldsymbol{w})$. In fact, the more general invariance results of Ozsváth–Szabó show that the filtered chain homotopy type of

$$\left(\widehat{\mathrm{CF}}(\mathcal{H}, \boldsymbol{z}), \mathcal{F}(K, *)\right)$$

is an invariant of the knot K. In particular, the chain homotopy type of each filtered subcomplex $\mathcal{F}(K, j)$ is also independent of the choice of Heegaard diagram, which justifies the notation.

As we did for classical invariants, we briefly discuss the behaviour of knot Floer homology under knot operations.

Ozsváth–Szabó described the behaviour of knot Floer homology and the τ-invariant under connected sum. We summarise the relevant statements in Theorem 2.67.

Theorem 2.67. *For any knots $K_1, K_2 \subset S^3$, there is a bi-graded isomorphism*

$$\widehat{HFK}(K_1 \# K_2) \cong \widehat{HFK}(K_1) \otimes \widehat{HFK}(K_2).$$

Moreover,

$$\tau(K_1 \# K_2) = \tau(K_1) + \tau(K_2).$$

When it comes to Murasugi sums, the situation is considerably more subtle. Similar to the Alexander polynomial, there is no known general formula describing the behaviour of knot Floer homology under Murasugi sum. The most general results in this direction are due to Cheng, Hedden, and Sarkar, and concern the bottom-most Alexander grading. We recall this result in the following theorem.

Theorem 2.68 (Cheng–Hedden–Sarkar [CHS22])**.** *Let K be a Murasugi sum of knots K_1 and K_2, performed along minimal genus Seifert surfaces. Then,*

$$\widehat{HFK}(K, -g(K)) \cong \widehat{HFK}(K_1, -g(K_1)) \otimes \widehat{HFK}(K_2, -g(K_2)).$$

The behaviour of knot Floer homology under satellite operations is also considerably more complicated. The study of cables and satellites was among the earliest applications of knot Floer homology and has remained an active area of research. The advent of bordered Floer homology and immersed curve invariants has provided powerful tools for studying these operations and has reshaped the modern study of the subject. In particular, a special case of the gluing formula in bordered Floer homology provides a categorification of the satellite formula of Theorem 2.18.

The topic began with work of Hedden [Hed05], who computed the knot Floer homology of certain cable knots in terms of the filtered subcomplexes $\mathcal{F}(K, *)$. Later, Hanselman and Watson [HW23] essentially resolved the computation of knot Floer homology of cable knots in the setting of immersed curve invariants.

The resulting formulas are quite detailed and somewhat complicated, and hence we do not include them here. We revisit these ideas in a more general setting in Chapter 4, in the context of twisting and stabilization.

2.3 Bordered Floer Theory

As we mentioned in the previous section, one of the main challenges in Heegaard Floer theory is the computation of the invariants. Bordered Floer homology, introduced by Lipshitz, Ozsváth, and Thurston in 2008, was developed in part as a response to this problem. Today, bordered Floer homology is one of the most active areas of research in Heegaard Floer theory, and its influence has extended beyond the subject to areas such as Khovanov homology and quantum topology.

The main idea behind bordered Floer theory is a central one in 3-manifold topology: cut-and-paste methods. Bordered Floer theory extends Heegaard Floer homology to 3-manifolds with boundary. One can then decompose a 3-manifold M into simpler pieces where explicit computations are possible, and finally "glue" the resulting invariants back together to recover $\widehat{\mathrm{HF}}(M)$.

From the perspective of topological quantum field theory (TQFT), there is a natural way to set up such a theory. Roughly speaking, bordered Floer theory assigns:

- To a closed, connected, and oriented surface F equipped with additional marking data, a differential graded algebra $\mathcal{A}(F)$.
- To a 3-manifold Y with boundary identified with F, algebraic invariants in the form of modules over $\mathcal{A}(F)$.

More precisely, there are two such invariants:

$$\widehat{\mathrm{CFD}}(Y) \qquad \text{and} \qquad \widehat{\mathrm{CFA}}(Y),$$

called the type D and type A bordered invariants, respectively. The first is a left differential graded module over $\mathcal{A}(-F)$, while the second is a right $\mathcal{A}_\infty$-module over $\mathcal{A}(F)$. Both are invariants up to a notion of homotopy equivalence.

The next main piece of bordered Floer theory is the existence of gluing formulas. Roughly speaking, if a closed 3-manifold Y is decomposed into two pieces

$$Y = Y_1 \cup_F Y_2,$$

then the Heegaard Floer chain complex $\widehat{\mathrm{CF}}(Y)$ can be recovered from the bordered invariants $\widehat{\mathrm{CFA}}(Y_1)$ and $\widehat{\mathrm{CFD}}(Y_2)$ via the derived tensor product. More precisely, there is a chain homotopy equivalence

$$\widehat{\mathrm{CF}}(Y) \simeq \widehat{\mathrm{CFA}}(Y_1) \widetilde{\otimes}_{\mathcal{A}(F)} \widehat{\mathrm{CFD}}(Y_2).$$

The power of bordered Floer theory comes at the expense of considerable algebraic complexity. However, this difficulty was later addressed by Hanselman, Rasmussen, and Watson, who in 2016 gave a geometric interpretation of bordered Floer theory in terms of immersed curves. This viewpoint significantly simplifies many arguments and computations in the theory.

In the rest of this section, we briefly review the main constructions, definitions, and theorems from bordered Floer theory. We should emphasise that even a brief overview of the full theory would require substantial technical background and is beyond the scope of this thesis. Instead, we focus on the case of 3-manifolds with torus boundary, where many of the structures become considerably simpler while still capturing the main ideas of the theory. This is also the only setting that will be used in the later chapters.

2.3.1 Working Guide to Bordered Floer Theory with Torus Boundary

We begin by defining the main object of the theory: bordered 3-manifolds with torus boundary.

Definition 2.69. A *bordered 3-manifold with torus boundary* is a triple (M, α_1, α_2), where M is an orientable 3-manifold with torus boundary, and $\alpha_1, \alpha_2 \subset \partial M$ are oriented, essential, simple closed curves satisfying $|\alpha_1 \cap \alpha_2| = 1$. To define a bordered invariant of type A, we require the oriented intersection number $\alpha_1 \cdot \alpha_2$ to be $+1$. Conversely, to define a bordered invariant of type D, we require $\alpha_1 \cdot \alpha_2 = -1$.

As discussed above, bordered Floer theory starts by associating an algebra to the boundary surface. In the case of torus boundary, this algebra is called the *torus algebra*.

Definition 2.70. The algebra $\mathcal{A}$, called the *torus algebra*, is defined as a quotient of the path algebra of the quiver shown in Figure 2.18. As a vector space over $\mathbb{F}_2$, it has a basis consisting of two idempotent elements ι_0, ι_1, and six additional elements

$$\rho_1, \rho_2, \rho_3, \rho_{12}, \rho_{23}, \rho_{123}.$$

The non-trivial multiplications are given by

$$\iota_0\rho_1 = \rho_1 = \rho_1\iota_1, \qquad \iota_1\rho_2 = \rho_2 = \rho_2\iota_0, \qquad \iota_0\rho_3 = \rho_3 = \rho_3\iota_1,$$
$$\rho_{12} = \rho_1\rho_2, \qquad \rho_{23} = \rho_2\rho_3, \qquad \rho_{123} = \rho_1\rho_2\rho_3.$$

We also use the notation $\rho_\varnothing = \mathbb{1} = \iota_0 + \iota_1$ for the unit of $\mathcal{A}$. We denote the multiplication map of $\mathcal{A}$ by

$$\mu \colon \mathcal{A} \otimes \mathcal{A} \to \mathcal{A}.$$

Finally, the subring generated by the idempotents is denoted by $\mathcal{I}$.

Remark 2.71. Unless stated otherwise, all tensor products in this section are taken over $\mathcal{I}$.

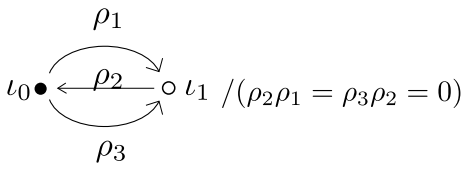


Figure 2.18: Quiver defining the torus algebra $\mathcal{A}$ [Che23].

We are now ready to define a type D structure. This can be viewed as a graded differential module over $\mathcal{A}$, although we use a slightly different formulation. We discuss the grading structures later in this section.

Definition 2.72. A *type D structure* over $\mathcal{A}$ is a unital left $\mathcal{I}$-module N equipped with an $\mathcal{I}$-linear map $\delta \colon N \to \mathcal{A} \otimes N$ satisfying the compatibility condition

$$(\mu \otimes \mathbb{I}_N) \circ (\mathbb{I}_\mathcal{A} \otimes \delta) \circ \delta = 0.$$

Equivalently, the following composition vanishes:

$$N \xrightarrow{\delta} \mathcal{A} \otimes N \xrightarrow{\mathbb{I}_\mathcal{A} \otimes \delta} \mathcal{A} \otimes \mathcal{A} \otimes N \xrightarrow{\mu \otimes \mathbb{I}_N} \mathcal{A} \otimes N.$$

For a type D structure N, define maps

$$\delta_k \colon N \to \mathcal{A}^{\otimes k} \otimes N$$

inductively by $\delta_0 = \mathbb{I}_N, \delta_1 = \delta$, and, for $k \geq 2$,

$$\delta_k = (\mathbb{I}^{\otimes(k-1)} \otimes \delta_1) \circ \delta_{k-1}.$$

A type D structure is called *bounded* if $\delta_k = 0$ for sufficiently large k. It is called *reduced* if there are no elements $x \in N$ such that $\delta(x)$ contains a summand of the form $\rho_\varnothing \otimes y$.

Remark 2.73. In this thesis, we work exclusively with bounded type D structures. This assumption is not particularly restrictive, since any type D structure is homotopy equivalent to a bounded one; see [HRW24].

Given a type D structure N, the tensor product $\mathcal{A} \otimes N$ forms a left differential module over $\mathcal{A}$. The left $\mathcal{A}$-action is induced by multiplication in $\mathcal{A}$. The differential $\partial_{\mathcal{A}\otimes N}$ is defined as the composition

$$\mathcal{A} \otimes N \xrightarrow{\mathbb{I}_{\mathcal{A}}\otimes\delta} \mathcal{A} \otimes \mathcal{A} \otimes N \xrightarrow{\mu\otimes\mathbb{I}_N} \mathcal{A} \otimes N.$$

The compatibility condition on δ implies that $\partial^2_{\mathcal{A}\otimes N} = 0$.

There is a convenient graph-theoretic description of type D structures. We work with decorated graphs whose vertices are labelled by $\{\bullet, \circ\}$, and whose directed edges are labelled by

$$\{\varnothing, 1, 2, 3, 12, 23, 123\}.$$

The labels must be compatible with the quiver defining the torus algebra. Thus the only possible edge types are

$$\bullet \xrightarrow{\varnothing} \bullet, \qquad \circ \xrightarrow{\varnothing} \circ, \qquad \bullet \xrightarrow{1} \circ, \qquad \bullet \xrightarrow{3} \circ,$$

$$\circ \xrightarrow{2} \bullet, \qquad \bullet \xrightarrow{12} \bullet, \qquad \circ \xrightarrow{23} \circ, \qquad \bullet \xrightarrow{123} \circ.$$

Construction 2.74. Let N be a type D structure. Since N is a left $\mathcal{I}$-module, it splits as $N = \iota_0 N \oplus \iota_1 N$. Choose a basis

$$B_N = B_{N_0} \sqcup B_{N_1},$$

where B_{N_0} is a basis for $\iota_0 N$ and B_{N_1} is a basis for $\iota_1 N$. The associated decorated graph Γ_N is defined as follows:

- Each element of B_{N_0} (resp. B_{N_1}) determines a vertex labelled by $\bullet$ (resp. $\circ$).
- Let $x \in B_N$ and expand $\delta(x)$ with respect to the basis B_N as follows:

$$\delta(x) = \sum_{j=1}^{k} \rho_{I_j} \otimes x_j.$$

Then for each summand we draw a directed edge $x \xrightarrow{I_j} x_j$.

An example of such a decorated graph is shown in Figure 2.19.

Remark 2.75. Higher-order maps δ_k are determined by directed paths in Γ_N. A type D structure N is bounded if and only if Γ_N contains no directed cycles. Similarly, N is reduced if and only if Γ_N has no edge labelled by $\varnothing$.

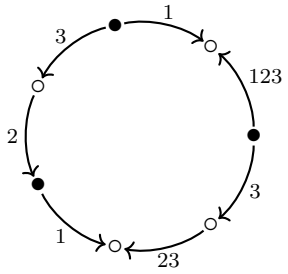


Figure 2.19: A decorated graph associated to a type D structure [HRW24].

We now define the corresponding notion of a type A structure. Unlike type D structures, these are $\mathcal{A}_\infty$-modules and are therefore slightly more complicated.

Definition 2.76. A *type A structure* over $\mathcal{A}$ is a right unital $\mathcal{I}$-module P equipped with a family of maps

$$m_{i+1}\colon P \otimes \mathcal{A}^{\otimes i} \to P$$

satisfying the $\mathcal{A}_\infty$ relations

$$\sum_{i=1}^{n} m_{n-i}\big(m_i(x \otimes a_1 \otimes \cdots \otimes a_{i-1}) \otimes \cdots \otimes a_{n-1}\big)$$

$$+ \sum_{i=1}^{n-2} m_{n-1}(x \otimes a_1 \otimes \cdots \otimes a_i a_{i+1} \otimes \cdots \otimes a_{n-1}) = 0$$

for all $x \in P$ and $a_1, \dots, a_{n-1} \in \mathcal{A}$. Furthermore,

$$m_2(x, \mathbb{1}) = x,$$

and

$$m_i(x, \cdots, \mathbb{1}, \cdots) = 0 \qquad \text{for } i \geq 3.$$

A type A structure is called *bounded* if $m_k = 0$ for sufficiently large k. It is called *reduced* if $m_1 = 0$.

We now turn to notions of homotopy equivalence for type D and type A structures. For type D structures, the notion of homotopy equivalence is equivalent to ordinary (chain) homotopy equivalence of the associated differential modules $\mathcal{A} \otimes N_1$ and $\mathcal{A} \otimes N_2$. We will not give the precise definition here. Instead, we focus on the corresponding notions for $\mathcal{A}_\infty$-modules, which play a more essential role in this thesis. We begin by defining $\mathcal{A}_\infty$ homomorphisms.

Definition 2.77. Let P and P' be right $\mathcal{A}_\infty$-modules over $\mathcal{A}$. An $\mathcal{A}_\infty$ *homomorphism* $f\colon P \to P'$ is a collection of maps

$$f_i\colon P \otimes \mathcal{A}^{\otimes(i-1)} \to P', \qquad i \geq 1,$$

satisfying the following compatibility conditions: for every $n \geq 1$, $x \in P$, and $a_1, \ldots, a_{n-1} \in \mathcal{A}$,

$$\begin{aligned} &\sum_{i+j=n+1} m'_i\big(f_j(x \otimes a_1 \otimes \cdots \otimes a_{j-1}) \otimes \cdots \otimes a_{n-1}\big) \\ &+ \sum_{i+j=n+1} f_i\big(m_j(x \otimes a_1 \otimes \cdots \otimes a_{j-1}) \otimes \cdots \otimes a_{n-1}\big) \\ &+ \sum_{l=1}^{n-2} f_{n-1}\big(x \otimes a_1 \otimes \cdots \otimes a_l a_{l+1} \otimes \cdots \otimes a_{n-1}\big) = 0. \end{aligned}$$

Furthermore, for all $i > 1$,

$$f_i(x \otimes a_1 \otimes \cdots \otimes \mathbb{1} \otimes \cdots \otimes a_{i-1}) = 0.$$

One can define a composition operation for $\mathcal{A}_\infty$ homomorphisms. Let

$$f\colon P \to P' \qquad \text{and} \qquad g\colon P' \to P''$$

be $\mathcal{A}_\infty$ homomorphisms. Their composition $g \circ f\colon P \to P''$ is defined by

$$(g \circ f)_n(x \otimes a_1 \otimes \cdots \otimes a_{n-1}) := \sum_{i+j=n+1} g_j\big(f_i(x \otimes a_1 \otimes \cdots \otimes a_{i-1}) \otimes \cdots \otimes a_{n-1}\big),$$

for all $n \geq 1$, $x \in P$, and $a_1, \ldots, a_{n-1} \in \mathcal{A}$.

The next step would be to define notions of homotopy, and in particular null-homotopy, for $\mathcal{A}_\infty$ homomorphisms.

Definition 2.78. An $\mathcal{A}_\infty$ homomorphism $f\colon P \to P'$ is called *null-homotopic* if there exists a collection of maps

$$H_i\colon P \otimes \mathcal{A}^{\otimes(i-1)} \to P'$$

such that the following hold:

1. For every $n \geq 1$, $x \in P$, and $a_1, \ldots, a_{n-1} \in \mathcal{A}$,

$$\begin{aligned} f_n(x \otimes a_1 \otimes \cdots \otimes a_{n-1}) = &\sum_{i+j=n+1} m'_i\big(H_j(x \otimes a_1 \otimes \cdots \otimes a_{j-1}) \otimes \cdots \otimes a_{n-1}\big) \\ &+ \sum_{i+j=n+1} H_i\big(m_j(x \otimes a_1 \otimes \cdots \otimes a_{j-1}) \otimes \cdots \otimes a_{n-1}\big) \\ &+ \sum_{l=1}^{n-2} H_{n-2}\big(x \otimes a_1 \otimes \cdots \otimes a_l a_{l+1} \otimes \cdots \otimes a_{n-1}\big). \end{aligned}$$

2. For every $i > 1$,

$$H_i(x \otimes a_1 \otimes \cdots \otimes \mathbb{1} \otimes \cdots \otimes a_{i-1}) = 0.$$

Now the setting is ready to define $\mathcal{A}_\infty$ homotopy equivalences.

Definition 2.79. Two $\mathcal{A}_\infty$ homomorphisms $f, g\colon P \to P'$ are called *homotopic* if $f - g$ is null-homotopic. An $\mathcal{A}_\infty$ homomorphism $h\colon P \to P'$ is called an $\mathcal{A}_\infty$ *homotopy equivalence* if there exists an $\mathcal{A}_\infty$ homomorphism $h'\colon P' \to P$ such that the compositions $h \circ h'$ and $h' \circ h$ are homotopic to the identity.

We now introduce our main bordered invariants. Recall from Definition 2.69 that there is a subtle sign convention in the definition of bordered 3-manifolds, depending on whether one considers type A or type D invariants. With this convention understood, we have the following theorem of Lipshitz, Ozsváth, and Thurston.

Theorem 2.80 ([LOT18])**.** *Let (M, α_1, α_2) be a bordered 3-manifold with torus boundary. Then:*

- *If $\alpha_1 \cdot \alpha_2 = -1$, then there exists a type D structure $\widehat{CFD}(M, \alpha_1, \alpha_2)$, well-defined up to homotopy equivalence.*
- *If $\alpha_1 \cdot \alpha_2 = +1$, then there exists a type A structure $\widehat{CFA}(M, \alpha_1, \alpha_2)$, well-defined up to $\mathcal{A}_\infty$ homotopy equivalence.*

Remark 2.81. Note that

$$\alpha_1 \cdot \alpha_2 = -\alpha_2 \cdot \alpha_1.$$

Consequently, if the bordered 3-manifold (M, α_1, α_2) can be used to define one of the bordered invariants, then (M, α_2, α_1) can be used to define the other. To simplify notation, we will occasionally suppress this change of boundary parametrization. In such cases, the intended parametrization will always be clear from the choice of bordered invariant under consideration.

Before turning to the gluing theorem, we take this opportunity to describe the graph-theoretic interpretation of type A structures. In general, one can give a graph-theoretic description of type A structures, although the graphs and edge labels become considerably more complicated. In the special case of the bordered invariant $\widehat{\mathrm{CFA}}(M, \alpha_1, \alpha_2)$, Hedden and Levine [HL16] used a certain duality between type A and type D structures to give a simplified graphical description of type A invariants. We briefly describe this construction below.

Construction 2.82. Assume that $\widehat{\mathrm{CFD}}(M, \alpha_1, \alpha_2)$ is reduced, and let Γ_D be its associated decorated graph. Define a new decorated graph Γ_A by replacing edge labels according to the bijection

$$1 \leftrightarrow 3, \qquad 2 \leftrightarrow 2, \qquad 3 \leftrightarrow 1.$$

The vertices of Γ_A remain labelled by $\{\bullet, \circ\}$, while the edge labels now lie in

$$\{1, 2, 3, 32, 21, 321\}.$$

The graph Γ_A determines the type A structure $\widehat{\mathrm{CFA}}(M, \alpha_1, \alpha_2)$ as follows. The vertices form a basis for $\widehat{\mathrm{CFA}}(M, \alpha_1, \alpha_2)$ as a right $\mathcal{I}$-module. To compute the maps m_{n+1}, consider a directed path from a vertex x to a vertex y with edge labels $(J_1, \dots, J_l)$. Concatenate these labels to form the word $J_1 \cdots J_l$, and regroup this word as

$$I_1 \ \cdots \ I_n,$$

where n is minimal and each I_t belongs to $\{1, 2, 3, 32, 21, 321\}$. Then, the path contributes a y summand to

$$m_{n+1}(x \otimes \rho_{I_1} \otimes \cdots \otimes \rho_{I_n}).$$

Repeating this procedure for all paths determines the full $\mathcal{A}_\infty$-module structure. An example of such a decorated graph is shown in Figure 2.20.

Remark 2.83. Construction 2.82 is stated for reduced structures, and hence there are no edges labelled by $\varnothing$. However, given a decorated graph with edge labels

$$\{\varnothing, 1, 2, 3, 32, 21, 321\},$$

one can still read off a $\mathcal{A}_\infty$-module structure. In this case, the maps m_{n+1} for $n \geq 1$ are determined by directed paths which do not contain edges labelled by $\varnothing$. The edges labelled by $\varnothing$ are instead used to define the map m_1. Of course, in this more general setting the $\mathcal{A}_\infty$ relations impose additional conditions on the decorated graph. Since we will not need these conditions explicitly, we omit them from the discussion.

We now turn our attention to the gluing theorem. As mentioned at the beginning of this section, the gluing operation is reflected algebraically by a derived tensor product of bordered invariants. In practice, one works with a specific model for this derived tensor product, called the *box tensor product*, which we define below.

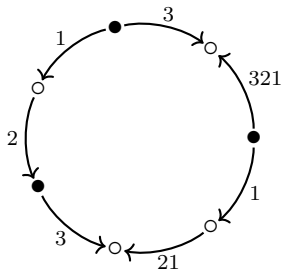


Figure 2.20: The graph Γ_A obtained from Γ_D in Figure 2.19 [HRW24].

Definition 2.84. Let P be a type A structure and N a type D structure. Assume that N is bounded. The *box tensor product* $P \boxtimes N$ is defined as follows. As a vector space over $\mathbb{F}_2$, it is given by

$$P \boxtimes N = P \otimes_{\mathcal{I}} N.$$

The differential on $P \boxtimes N$ is defined by

$$\partial^{\boxtimes}(x \otimes y) = \sum_{i=0}^{\infty} (m_{i+1} \otimes \mathbb{I}_N)(x \otimes \delta_i(y)). \tag{2.4}$$

We now state the gluing theorem of Lipshitz, Ozsváth, and Thurston.

Theorem 2.85 ([LOT18])**.** *Let (M, α_1, α_2) and $(M', \alpha_1', \alpha_2')$ be bordered 3-manifolds, and let $Y = M \cup_h M'$ be the closed 3-manifold obtained by gluing them along an orientation-reversing homeomorphism $h : \partial M' \to \partial M$ satisfying*

$$h(\alpha_1') = \alpha_1, \qquad h(\alpha_2') = \alpha_2.$$

If $\widehat{CFD}(M', \alpha_1', \alpha_2')$ is bounded, then

$$\widehat{CF}(Y) \simeq \widehat{CFA}(M, \alpha_1, \alpha_2) \boxtimes \widehat{CFD}(M', \alpha_1', \alpha_2').$$

Similar to Heegaard Floer homology, bordered Floer theory also produces knot invariants. Given a knot K in a bordered 3-manifold (M, α_1, α_2), one can define bordered invariants

$$\widehat{\mathrm{CFD}}(M, K) \qquad \text{and} \qquad \widehat{\mathrm{CFA}}(M, K),$$

which are invariants of the pair (M, K) up to homotopy equivalence.

Lipshitz, Ozsváth, and Thurston [LOT18] generalized the pairing theorem to bordered invariants of knots.

Theorem 2.86 ([LOT18])**.** *Let (M, α_1, α_2) and $(M', \alpha_1', \alpha_2')$ be bordered 3-manifolds, and let $Y = M \cup_h M'$ be as in Theorem 2.85. Let K be a knot in M. Then:*

1. If $\widehat{\mathrm{CFD}}(M', \alpha_1', \alpha_2')$ *is bounded, then*

$$\widehat{\mathrm{CFK}}(Y, K) \simeq \widehat{\mathrm{CFA}}(M, K) \boxtimes \widehat{\mathrm{CFD}}(M', \alpha_1', \alpha_2').$$

2. If $\widehat{\mathrm{CFD}}(M, K)$ *is bounded, then*

$$\widehat{\mathrm{CFK}}(Y, K) \simeq \widehat{\mathrm{CFA}}(M', \alpha_1', \alpha_2') \boxtimes \widehat{\mathrm{CFD}}(M, K).$$

Remark 2.87. Similar to Heegaard Floer homology, the construction of the bordered invariants of a bordered 3-manifold (M, α_1, α_2) is based on a *bordered Heegaard diagram* and involves counts of pseudo-holomorphic curves. These counts are used to define the maps δ and m_i appearing in the type D and type A structures. The general definitions are not required for this thesis. Later in this chapter, we describe the construction explicitly in the special case of genus one diagrams.

We now turn to the grading structures on bordered invariants. Type D and type A structures associated to bordered 3-manifolds with torus boundary carry additional grading data. These gradings take values in certain coset spaces of a non-commutative group G, defined below. At first sight, this grading structure may appear somewhat unusual. However, as we will see, in practice the resulting computations are quite manageable.

Definition 2.88. The *torus grading group* G has underlying set

$$\left\{ (m; i, j) \;\middle|\; m, i, j \in \frac{1}{2}\mathbb{Z}, \quad i + j \in \mathbb{Z}, \quad m \in \mathbb{Z} \iff i, j \in \mathbb{Z} \text{ and } i + j \in 2\mathbb{Z} \right\}.$$

The multiplication is given by

$$(m_1; i_1, j_1) \cdot (m_2; i_2, j_2) = \left(m_1 + m_2 + \begin{vmatrix} i_1 & j_1 \\ i_2 & j_2 \end{vmatrix} ; i_1 + i_2, j_1 + j_2 \right).$$

The component m is called the *Maslov component*, while the pair (i, j) is called the *Spin^c component*. The following central element of G is of particular importance:

$$\lambda = (1; 0, 0).$$

The torus algebra is graded by G. The grading of the basis elements is given by

$$\mathrm{gr}(\iota_0) = \mathrm{gr}(\iota_1) = (0; 0, 0),$$

$$\mathrm{gr}(\rho_1) = \left(-\frac{1}{2}; \frac{1}{2}, -\frac{1}{2} \right), \qquad \mathrm{gr}(\rho_2) = \left(-\frac{1}{2}; \frac{1}{2}, \frac{1}{2} \right), \qquad \mathrm{gr}(\rho_3) = \left(-\frac{1}{2}; -\frac{1}{2}, \frac{1}{2} \right),$$

$$\mathrm{gr}(\rho_{12}) = \mathrm{gr}(\rho_1)\,\mathrm{gr}(\rho_2), \qquad \mathrm{gr}(\rho_{23}) = \mathrm{gr}(\rho_2)\,\mathrm{gr}(\rho_3), \qquad \mathrm{gr}(\rho_{123}) = \mathrm{gr}(\rho_1)\,\mathrm{gr}(\rho_2)\,\mathrm{gr}(\rho_3).$$

We now describe how to compute the gradings on type D and type A structures. We denote the grading of a homogeneous element by $\mathrm{gr}(\cdot)$.

Let B denote a chosen basis for either a type D structure $\widehat{\mathrm{CFD}}(M, \alpha_1, \alpha_2)$ or a type A structure $\widehat{\mathrm{CFA}}(M, \alpha_1, \alpha_2)$, viewed as an $\mathcal{I}$-module. We require the basis elements to be homogeneous. In particular, the basis coming from Heegaard states in a bordered Heegaard diagram (see Remark 2.87) has this property.

One then chooses a reference generator $\mathbf{r}_0 \in B$ and assigns it grading

$$\mathrm{gr}(\mathbf{r}_0) = [(0; 0, 0)].$$

Associated to this choice is a subgroup $P_{\mathbf{r}_0} \subseteq G$. The grading takes values in the following coset spaces:

- For type D structures, in $G/P_{\mathbf{r}_0}$. Equivalently, the grading determines a map

$$\mathrm{gr}\colon B \to G/P_{\mathbf{r}_0}.$$

- For type A structures, in $P_{\mathbf{r}_0}\backslash G$. Equivalently, the grading determines a map

$$\mathrm{gr}\colon B \to P_{\mathbf{r}_0}\backslash G.$$

Remark 2.89. For the familiar reader, the subgroup $P_{\mathbf{r}_0}$ records the gradings of the periodic domains in a bordered Heegaard diagram. We will not need its precise definition here. As we will see in Remark 2.93, in many examples, and in particular in those relevant to this thesis, it can be computed directly from the description of the corresponding type A or type D structure.

Remark 2.90. The grading cosets themselves depend on the choice of reference generator $\mathbf{r}_0$. We will later describe how gradings associated to different choices of reference generator are related, and hence carry the same information. Apart from this, the grading is independent of the choice of homogeneous basis.

We now describe how to compute the grading of a type D structure. Strictly speaking, the grading is defined on the differential module $\mathcal{A} \otimes N$, where

$$N = \widehat{\mathrm{CFD}}(M, \alpha_1, \alpha_2).$$

One can consider a basis for $\mathcal{A} \otimes N$ consisting of tensors of the form $a \otimes x$, where $a \in \mathcal{A}$ and $x \in B$.

Construction 2.91 (Computing grading of a type D structure)**.** Recall that we have chosen a homogeneous basis B together with a reference generator $\mathbf{r}_0 \in B$. The grading of the remaining basis elements is determined by the rules

$$\mathrm{gr}(a \otimes x) = \mathrm{gr}(a)\,\mathrm{gr}(x),$$

$$\mathrm{gr}(\delta(x)) = \lambda^{-1}\,\mathrm{gr}(x).$$

In terms of the graph-theoretic description of a type D structure, these rules are summarised in Table 2.21.

Labelled edge	Grading change	Labelled edge	Grading change
$x \xrightarrow{1} y$	$\mathrm{gr}(x) = \left(\frac{1}{2}; \frac{1}{2}, -\frac{1}{2}\right) \cdot \mathrm{gr}(y)$	$x \xrightarrow{12} y$	$\mathrm{gr}(x) = \left(\frac{1}{2}; 1, 0\right) \cdot \mathrm{gr}(y)$
$x \xrightarrow{2} y$	$\mathrm{gr}(x) = \left(\frac{1}{2}; \frac{1}{2}, \frac{1}{2}\right) \cdot \mathrm{gr}(y)$	$x \xrightarrow{23} y$	$\mathrm{gr}(x) = \left(\frac{1}{2}; 0, 1\right) \cdot \mathrm{gr}(y)$
$x \xrightarrow{3} y$	$\mathrm{gr}(x) = \left(\frac{1}{2}; -\frac{1}{2}, \frac{1}{2}\right) \cdot \mathrm{gr}(y)$	$x \xrightarrow{123} y$	$\mathrm{gr}(x) = \left(\frac{1}{2}; \frac{1}{2}, \frac{1}{2}\right) \cdot \mathrm{gr}(y)$

Figure 2.21: Grading changes associated to edges in a decorated graph representing a type D structure [HRW22].

We now describe how to compute the grading of a type A structure. The procedure is analogous to that for type D structures.

Construction 2.92 (Computing grading of a type A structure)**.** Recall that we have chosen a homogeneous basis B together with a reference generator $\mathbf{r}_0 \in B$. The grading of the remaining basis elements is determined by the rule

$$\mathrm{gr}\big(m_{n+1}(x \otimes a_1 \otimes \cdots \otimes a_n)\big) = \lambda^{-1+n}\,\mathrm{gr}(x)\,\mathrm{gr}(a_1) \cdots \mathrm{gr}(a_n).$$

In terms of the graph-theoretic description of a type A structure, this rule is summarised in Table 2.22.

Labelled edge	Grading change	Labelled edge	Grading change
$x \xrightarrow{(3)} y$	$\mathrm{gr}(x) = \mathrm{gr}(y) \cdot \left(\frac{1}{2}; \frac{1}{2}, -\frac{1}{2}\right)$	$x \xrightarrow{(3,2)} y$	$\mathrm{gr}(x) = \mathrm{gr}(y) \cdot \left(\frac{1}{2}; 0, -1\right)$
$x \xrightarrow{(2)} y$	$\mathrm{gr}(x) = \mathrm{gr}(y) \cdot \left(\frac{1}{2}; -\frac{1}{2}, -\frac{1}{2}\right)$	$x \xrightarrow{(2,1)} y$	$\mathrm{gr}(x) = \mathrm{gr}(y) \cdot \left(\frac{1}{2}; -1, 0\right)$
$x \xrightarrow{(1)} y$	$\mathrm{gr}(x) = \mathrm{gr}(y) \cdot \left(\frac{1}{2}; -\frac{1}{2}, \frac{1}{2}\right)$	$x \xrightarrow{(3,2,1)} y$	$\mathrm{gr}(x) = \mathrm{gr}(y) \cdot \left(\frac{1}{2}; -\frac{1}{2}, -\frac{1}{2}\right)$

Figure 2.22: Grading changes associated to edges in a decorated graph representing a type A structure [HRW22].

Remark 2.93. In the applications considered in this thesis, the decorated graphs are connected. Hence, after choosing the reference generator $\mathbf{r}_0$, one can compute the grading of every vertex by following paths in the graph and applying the grading rules above. Different paths from $\mathbf{r}_0$ to the same vertex may give different representatives of the same grading coset; the ratios of these representatives give elements of the subgroup $P_{\mathbf{r}_0}$. In particular, when M is a rational homology solid torus, the subgroup $P_{\mathbf{r}_0}$ is cyclic, and in many examples it can be identified from cycles in the graph.

We now briefly discuss the grading in the knot case. For bordered invariants of knots, the torus grading group is enhanced by an additional Alexander component. More precisely, one works with

$$\widetilde{G} = G \times \mathbb{Z},$$

where the extra $\mathbb{Z}$ factor is called the *Alexander component*. The two main central elements of $\widetilde{G}$ are

$$\lambda = (1; 0, 0; 0), \qquad \mu = (0; 0, 0; 1).$$

By a slight abuse of notation, we use the same symbol λ as before.

The grading rules for type D and type A structures remain unchanged after replacing G by $\widetilde{G}$. The only additional information is the shift in the Alexander component, which cannot be recovered directly from the algebraic data. This shift is defined in the same way as the Alexander grading in knot Floer homology: it records the intersection number of the relevant pseudo-holomorphic curves with the additional basepoint used to encode the knot (see Remark 2.87). We will revisit this in detail in the genus one examples. For now, it suffices to know that the bordered knot invariants are graded in the following coset spaces:

- $\widehat{\mathrm{CFD}}(M, K)$ is graded in $\widetilde{G}/P_{\mathbf{r}_0}$.
- $\widehat{\mathrm{CFA}}(M, K)$ is graded in $P_{\mathbf{r}_0}\backslash\widetilde{G}$.

Similar to Heegaard Floer homology, the bordered grading contains information about the Maslov and Alexander gradings. This information becomes visible after passing to the box tensor product.

The gradings on type D and type A structures induce a grading on the box tensor product. More precisely:

- The box tensor product

$$\widehat{\mathrm{CFA}}(M, \alpha_1, \alpha_2) \boxtimes \widehat{\mathrm{CFD}}(M', \alpha_1', \alpha_2')$$

  is graded in the double coset space $P_{\mathbf{r}_0}\backslash G/P_{\mathbf{r}_0'}$.

- The box tensor product

$$\widehat{\mathrm{CFA}}(M,K)\boxtimes\widehat{\mathrm{CFD}}(M',\alpha_1',\alpha_2')$$

  is graded in the double coset space $P_{\mathbf{r}_0}\backslash\widetilde{G}/(P_{\mathbf{r}_0'};0)$.

The grading of a homogeneous tensor is given by

$$\mathrm{gr}(x\otimes y)=[\mathrm{gr}(x)\cdot\mathrm{gr}(y)]$$

in the first case, and by

$$\mathrm{gr}(x\otimes y)=[\mathrm{gr}(x)\cdot(\mathrm{gr}(y);0)]$$

in the second.

As in our discussion of Heegaard Floer homology, we restrict attention to the case where the resulting closed 3-manifold is S^3. In this setting, the bordered grading recovers the Maslov and Alexander gradings as follows.

Proposition 2.94. *Suppose that*

$$M\cup_h M'=S^3.$$

Let $x_1\otimes y_1$ and $x_2\otimes y_2$ be homogeneous elements of

$$\widehat{\mathrm{CFA}}(M,K)\boxtimes\widehat{\mathrm{CFD}}(M',\alpha_1',\alpha_2').$$

Then there exist unique integers h and a such that

$$\mathrm{gr}(x_1\otimes y_1)=\mathrm{gr}(x_2\otimes y_2)\lambda^h\mu^a, \tag{2.5}$$

where the equality is taken in the double coset space

$$P_{\mathbf{r}_0}\backslash\widetilde{G}/(P_{\mathbf{r}_0'};0).$$

Moreover,

$$h=M(x_1\otimes y_1)-M(x_2\otimes y_2),\qquad a=A(x_1\otimes y_1)-A(x_2\otimes y_2).$$

The analogous statement for

$$\widehat{\mathrm{CFA}}(M,\alpha_1,\alpha_2)\boxtimes\widehat{\mathrm{CFD}}(M',\alpha_1',\alpha_2')$$

recovers the relative Maslov grading in Heegaard Floer homology. In this case, one replaces $\widetilde{G}$ by G and omits the powers of μ from Equation (2.5).

As promised in Remark 2.90, we now explain how gradings associated to different choices of reference generator are related. This follows from a more general invariance theorem. We state only the type A version here, since it is the one that will be used later in the thesis.

Recall from Remark 2.87 that the construction of bordered invariants depends on a choice of bordered Heegaard diagram. Given such a diagram $\mathcal{H}$, we denote the associated type A structure by $\widehat{\mathrm{CFA}}(\mathcal{H})$. Similar to the admissibility condition discussed in Remark 2.53, bordered Heegaard diagrams must satisfy an additional condition called *provincial admissibility*.

We are now ready to state the invariance theorem.

Theorem 2.95 ([LOT18])**.** *Let $\mathcal{H}_1$ and $\mathcal{H}_2$ be provincially admissible bordered Heegaard diagrams for a bordered 3-manifold (M, α_1, α_2), and let $\mathbf{r}_1$ and $\mathbf{r}_2$ be reference generators for $\widehat{\mathrm{CFA}}(\mathcal{H}_1)$ and $\widehat{\mathrm{CFA}}(\mathcal{H}_2)$, respectively. Then there exist*

- *an isomorphism of G-sets*
$$\phi : P_{\mathbf{r}_1} \backslash G \longrightarrow P_{\mathbf{r}_2} \backslash G,$$
- *and an $\mathcal{A}_\infty$ homotopy equivalence*
$$\widehat{f} : \widehat{\mathrm{CFA}}(\mathcal{H}_1) \longrightarrow \widehat{\mathrm{CFA}}(\mathcal{H}_2),$$
given by maps
$$\widehat{f}_{i+1} : \widehat{\mathrm{CFA}}(\mathcal{H}_1) \otimes \mathcal{A}^{\otimes i} \longrightarrow \widehat{\mathrm{CFA}}(\mathcal{H}_2).$$

Moreover, for every homogeneous element
$$x \in \widehat{\mathrm{CFA}}(\mathcal{H}_1)$$
and every $a_1 \otimes \cdots \otimes a_l \in \mathcal{A}^{\otimes l}$, we have
$$\mathrm{gr}\big(\widehat{f}_{l+1}(x \otimes a_1 \otimes \cdots \otimes a_l)\big) = \phi(\mathrm{gr}(x)) \, \mathrm{gr}(a_1) \cdots \mathrm{gr}(a_l) \lambda^l.$$

We end this subsection with a final remark on invariance and $\mathcal{A}_\infty$ homotopy equivalence. The pairing theorem (Theorem 2.85) suggests that the chain homotopy type of the box tensor product should depend only on the homotopy types of the underlying type A and type D structures. The following proposition makes this precise.

Let $f : P \to P'$ be an $\mathcal{A}_\infty$ homomorphism and let N be a bounded type D structure. Then f induces a chain map

$$f \boxtimes \mathbb{I}_N : P \boxtimes N \to P' \boxtimes N$$

$$(f \boxtimes \mathbb{I}_N)(x \otimes y) = \sum_{k=0}^{\infty} (f_{k+1} \otimes \mathbb{I}_N)(x \otimes \delta_k(y)).$$

Proposition 2.96. *[[LOT18]] If f is an $\mathcal{A}_\infty$ homotopy equivalence, then $f \boxtimes \mathbb{I}_N$ is a chain homotopy equivalence.*

2.3.2 Working Guide to Immersed Curve Invariants

As discussed earlier, bordered Floer theory provides powerful computational tools, but at the expense of considerable algebraic complexity. In a breakthrough paper, Hanselman, Rasmussen, and Watson showed that the bordered invariants of 3-manifolds with torus boundary admit a geometric interpretation in terms of immersed curves in the punctured torus. This perspective greatly simplifies many arguments and computations.

In this subsection, we give a brief overview of this interpretation and the main results that we will use later in the thesis.

Definition 2.97. Let M be a compact, oriented 3-manifold with torus boundary, and fix a basepoint $z \in \partial M$. The *immersed curve invariant* of M, denoted by $\widehat{\mathrm{HF}}(M)$, consists of a collection of closed immersed curves $\{\gamma_1, \ldots, \gamma_n\}$ in the punctured torus $\partial M \setminus z$, each decorated with a local system (k_i, A_i), where k_i is a positive integer and A_i is an automorphism of a k_i-dimensional vector space over $\mathbb{F}_2$.
The invariant $\widehat{\mathrm{HF}}(M)$ is well-defined up to regular homotopy of the immersed curves and isomorphism of the local systems.

The immersed curve invariant can be constructed directly from a decorated graph representing a type D structure. We briefly describe the procedure.

Construction 2.98. Identify the punctured torus $\partial M \setminus z$ with the torus $T = \mathbb{R}^2/\mathbb{Z}^2$ punctured at the point

$$z = (1 - \varepsilon, 1 - \varepsilon).$$

Let α_1 and α_2 denote the images of the y- and x-axes in T, respectively. Consider the reduced type D structure

$$\widehat{\mathrm{CFD}}(M, \alpha_1, \alpha_2).$$

Let Γ_D be the associated decorated graph. Embed the vertices of Γ_D into T so that the $\bullet$ vertices lie at distinct points of α_1 and the $\circ$ vertices lie at distinct points of α_2. The edges are then embedded according to their labels, as illustrated in Figure 2.23.

Initially, this procedure may produce an immersed train track. A theorem of Hanselman, Rasmussen, and Watson [HRW24] states that, after choosing a suitable basis for the type D structure, this construction produces a collection of immersed curves. The resulting collection of immersed curves, equipped with the corresponding local systems, is the immersed curve invariant $\widehat{\mathrm{HF}}(M)$.

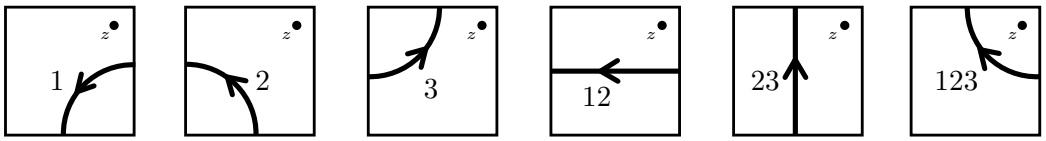


Figure 2.23: Embedding the edges of a decorated graph in the punctured torus [HRW24].

Hanselman, Rasmussen, and Watson [HRW24] proved an immersed-curve analogue of the bordered pairing theorem. Let

$$Y = M \cup_h M'$$

be as in Theorem 2.85. Let $\bar{h}$ denote the composition of h with the elliptic involution of ∂M.

Theorem 2.99 ([HRW24])**.** *Let $\boldsymbol{\gamma} = \widehat{HF}(M)$ and $\boldsymbol{\gamma}' = \widehat{HF}(M')$. Then*

$$\widehat{HF}(Y) \cong HF\big(\boldsymbol{\gamma}, \bar{h}(\boldsymbol{\gamma}')\big),$$

where the right-hand side denotes the Lagrangian intersection Floer homology of the immersed curve invariants.

We will not need the full definition of this version of Lagrangian intersection Floer homology and refer the reader to [HRW24] for details. Instead, we record the properties that will be used later in the thesis.

Notation 2.100. As in Theorem 2.99, we write

$$\boldsymbol{\gamma} = \boldsymbol{\gamma}_1 \cup \cdots \cup \boldsymbol{\gamma}_n$$

for a collection of decorated immersed curves. Each component is of the form

$$\boldsymbol{\gamma}_i = (\gamma_i, k_i, A_i),$$

where γ_i is an immersed curve and (k_i, A_i) is a local system.

The first property we will need is that Lagrangian intersection Floer homology splits over connected components.

Proposition 2.101. *[HRW24] Suppose that*

$$\boldsymbol{\gamma} = \boldsymbol{\gamma}_1 \cup \cdots \cup \boldsymbol{\gamma}_n, \qquad \bar{h}(\boldsymbol{\gamma}') = \boldsymbol{\gamma}'_1 \cup \cdots \cup \boldsymbol{\gamma}'_{n'}.$$

Then

$$HF(\boldsymbol{\gamma}, \bar{h}(\boldsymbol{\gamma}')) = \bigoplus_{\substack{1 \le i \le n \\ 1 \le j \le n'}} HF(\boldsymbol{\gamma}_i, \boldsymbol{\gamma}'_j).$$

Hence, it suffices to consider connected immersed curves equipped with local systems. We can simplify the situation further by restricting to primitive curves, in view of the following lemma.

Lemma 2.102 ([HRW24])**.** *Let $\boldsymbol{\gamma}_1 = (\gamma_1, k, A)$, where γ_1 is the n-fold cover of a primitive curve γ. Let*

$$\boldsymbol{\gamma} = (\gamma, nk, \widetilde{A}),$$

where $\widetilde{A}$ is the block matrix

$$\widetilde{A} = \left[\begin{array}{c|c|c|c} 0 & \cdots & 0 & A \\ \hline I_k & \cdots & 0 & 0 \\ \hline 0 & \ddots & 0 & 0 \\ \hline 0 & \cdots & I_k & 0 \end{array}\right].$$

Then, for any other decorated curve $\boldsymbol{\gamma}'$,

$$HF(\boldsymbol{\gamma}_1, \boldsymbol{\gamma}') \cong HF(\boldsymbol{\gamma}, \boldsymbol{\gamma}').$$

For primitive immersed curves, the dimension of the Lagrangian Floer homology can be computed from geometric intersection numbers, as stated in Theorem 2.105.

Recall that for essential immersed curves γ and γ', the *intersection number*

$$i(\gamma, \gamma')$$

is the minimal number of transverse intersection points between immersed curves $\widehat{\gamma}$ and $\widehat{\gamma}'$, where $\widehat{\gamma}$ is freely homotopic to γ and $\widehat{\gamma}'$ is freely homotopic to γ'.

Before stating Theorem 2.105, we introduce two technical notions.

Definition 2.103. Let $\gamma : S^1 \to T$ be a closed immersed curve in punctured torus T. Consider the composition

$$\gamma \circ p : \mathbb{R} \to T,$$

where $p : \mathbb{R} \to S^1$ is the universal covering map. Since

$$\widetilde{T} = \mathbb{R}^2 \setminus \mathbb{Z}^2$$

is the universal abelian cover of T, the map $\gamma \circ p$ admits a lift

$$\widetilde{\gamma} : \mathbb{R} \to \widetilde{T}.$$

We say that γ is *unobstructed* if $\widetilde{\gamma}$ is an embedding. By a theorem of Abouzaid [Abo06], this is equivalent to requiring that γ contain no immersed fishtails.

Definition 2.104. Two immersed curves $\gamma, \gamma' \subset T$ are called *commensurable* if there exists an immersed curve $\delta \subset T$ and positive integers n, n' such that γ is freely homotopic to δ^n and γ' is freely homotopic to $\delta^{n'}$.

Theorem 2.105 ([HRW24])**.** *Let*

$$\boldsymbol{\gamma} = (\gamma, k, A) \qquad \text{and} \qquad \boldsymbol{\gamma}' = (\gamma', k', A')$$

be primitive unobstructed immersed curves equipped with local systems. If γ and γ' are incommensurable, then

$$\dim HF(\boldsymbol{\gamma}, \boldsymbol{\gamma}') = kk' \cdot i(\gamma, \gamma').$$

To apply Theorem 2.105, one must compute the geometric intersection number $i(\gamma, \gamma')$. This requires choosing representatives of the free homotopy classes of γ and γ' that realize the minimal number of intersections. A theorem of Freedman, Hass, and Scott [FHS82] states that shortest-length geodesic representatives have this property. Motivated by this result, Hanselman, Rasmussen, and Watson [HRW24] introduced the notions of *pegboard diagrams* and *ε-geodesics*.

Intuitively, pegboard diagrams arise by viewing T as a flat Euclidean torus and imagining the puncture as a peg. One then places a rubber band in the free homotopy class determined by γ and allows it to contract until it reaches an equilibrium position.

To make the notion of a pegboard diagram precise, we first introduce a geometric model for the punctured torus. Let $\widehat{T}$ denote ∂M equipped with a flat metric g. Following Hanselman, Rasmussen, and Watson [HRW24], we construct a Riemannian manifold T_ε as follows.

The manifold T_ε consists of two pieces. The first is the complement

$$\widehat{T}_{2\varepsilon} = \widehat{T} \setminus B_{2\varepsilon}(z)$$

equipped with the flat metric g. The second is a surface of revolution attached along the boundary of $\widehat{T}_{2\varepsilon}$, as illustrated in Figure 2.24. The resulting manifold T_ε is homeomorphic to the punctured torus T and carries a complete Riemannian metric g_ε. Since T_ε embeds in $\widehat{T} \times \mathbb{R}$, there is a natural projection

$$p : T_\varepsilon \to \widehat{T}_\varepsilon \subset T,$$

which agrees with the identity on the flat region $\widehat{T}_{2\varepsilon}$.

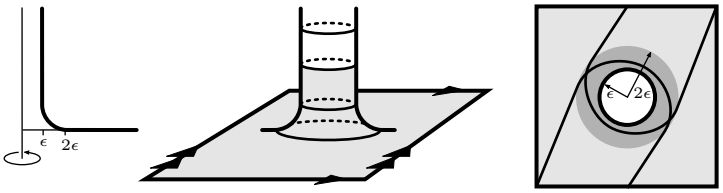


Figure 2.24: The manifold T_ϵ is a geometric model for T. The figure on the right shows the projection $p : T_\epsilon \to \widehat{T}_\epsilon$ [HRW24]

Definition 2.106. A curve $\gamma \subset T_\varepsilon$ is called an *ε-geodesic* if it is a geodesic with respect to the metric g_ε.

Hanselman, Rasmussen, and Watson [HRW24] proved that every nontrivial free homotopy class of loops in T admits an ε-geodesic representative. Moreover, this representative is either unique or entirely contained in the flat region of T_ε, in which case it is either boundary parallel or a Euclidean geodesic. Consequently, any two distinct ε-geodesics intersect minimally and transversally. As a result, the ε-geodesics can be used directly to compute geometric intersection numbers. They also provide a convenient way to construct pegboard diagrams.

Construction 2.107 (ε-pegboard diagrams)**.** Let $\gamma \subset T_\varepsilon$ be an ε-geodesic and consider its projection $p(\gamma) \subset T$. The projected curve consists of two parts. The first is the portion

$$p(\gamma) \cap \widehat{T}_{2\varepsilon},$$

which is a union of Euclidean geodesic segments. The second describes the behaviour of the curve near the puncture and is referred to as the collection of *geodesic corners*.

Modify the geodesic corners by a homotopy so that, instead of passing through the punctured neighbourhood, they follow a Euclidean circle centered at the puncture,

as illustrated in Figure 2.25. The resulting immersed curve is called the *ε-pegboard diagram* associated to γ.

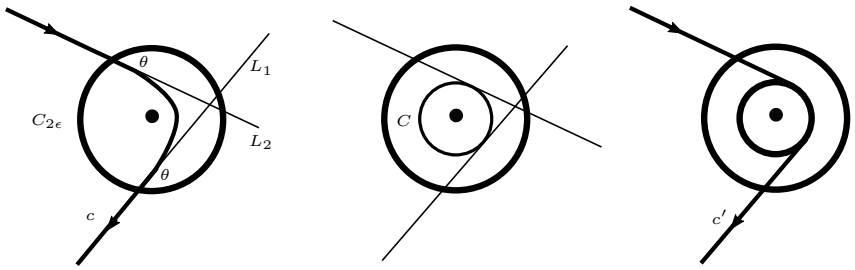


Figure 2.25: Turning a geodesic corner to a pegboard corner [HRW24]

Hanselman, Rasmussen, and Watson [HRW24] proved that if γ_1' and γ_2' are ε-pegboard diagrams associated to distinct ε-geodesics γ_1 and γ_2, then γ_1' and γ_2' are in minimal position. In particular,

$$|\gamma_1' \cap \gamma_2'| = i(\gamma_1', \gamma_2').$$

We will also need a degenerate version of a pegboard diagram. The general definition in [HRW24] is slightly more involved, but the following special case will suffice for our purposes.

Definition 2.108. Let γ_ε be the ε-pegboard diagram associated to a free homotopy class of curves. We say that γ_ε is *tight* if it has at least one corner, and *loose* otherwise.

Assume that γ_ε is tight. The *singular pegboard diagram* associated to γ_ε is obtained by taking the limit

$$\gamma_0 = \lim_{\varepsilon \to 0} \gamma_\varepsilon,$$

where the limit is taken in the Hausdorff metric. The limit γ_0 exists.

Since γ_ε consists of Euclidean geodesic segments outside the ball $B_{2\varepsilon}(z)$, the limiting curve γ_0 is a collection of Euclidean geodesics in $\widehat{T}$ meeting at the puncture z. The geodesics are completely determined by their slopes.

The limiting curve γ_0 does not retain the information contained in the corners of γ_ε. To retain this information, we record the corresponding *corner data*. Concretely, each c_i records how the corresponding corner of γ_ε winds around the puncture. We denote this data by $\mathbf{c} = (c_1, \ldots, c_n)$. The pair

$$\bar{\gamma} = (\gamma_0, \mathbf{c})$$

is called the *singular pegboard diagram* associated to γ.

We now arrive at the main advantage of singular pegboard diagrams: they provide a convenient way to compute geometric intersection numbers.

Theorem 2.109 ([HRW24])**.** *Let*

$$\bar{\gamma} = (\gamma_0, \mathbf{c}) \qquad \text{and} \qquad \bar{\gamma}' = (\gamma_0', \mathbf{c}')$$

be singular pegboard diagrams representing γ and γ', respectively. Suppose that no slope of γ_0 coincides with a slope of γ_0'. Then

$$i(\gamma, \gamma') = |\gamma_0 \cap \gamma_0'| + \sum_{(c_i, c_j')} i(c_i, c_j'),$$

where $i(c_i, c_j')$ denotes the local contribution to the intersection number determined by the corner data.

Remark 2.110. Suppose that one of the curves, say γ, is loose. Then the associated singular pegboard diagram has no corner data, and the second term in Theorem 2.109 vanishes. Consequently,

$$i(\gamma, \gamma') = |\gamma_0 \cap \gamma_0'|.$$

This is the only case that will be used later in the thesis.

So far, we have discussed the immersed curve invariant associated to a 3-manifold with torus boundary. Similar to bordered Floer theory, one can also associate immersed curve invariants to knots in manifolds with torus boundary.

The first such constructions were introduced by Chen [Che23] for $(1,1)$ patterns in the solid torus. More recently, Rasmussen introduced an immersed curve invariant $\mathrm{hfk}(P)$ associated to a knot $P \subset S^1 \times D^2$ in a talk at the Low-Dimensional Topology Workshop at the Erdős Center. Although a complete account of this construction has not yet appeared in the literature, the invariant has also been used by Baldwin, Hanselman, and Sivek [BHS26]. The properties relevant to this thesis will be described briefly below.

Construction 2.111 (The immersed curve invariant $\mathrm{hfk}(P)$)**.** The construction of $\mathrm{hfk}(P)$ parallels Construction 2.98. Starting from the decorated graph associated to

$$\widehat{\mathrm{CFD}}(S^1 \times D^2, P),$$

we embed the graph into the punctured torus T using the same local rules as before. As in the construction of $\widehat{HF}(M)$, one must also first choose a suitable basis for the type D structure.

The new feature in the knot setting is the possible appearance of *immersed arc* components. For type D structures associated to 3-manifolds with torus boundary, a condition called *extendability* ensures that the resulting immersed curve invariant consists entirely of closed curves. By contrast, the knot invariant $\widehat{\mathrm{CFD}}(S^1 \times D^2, P)$ need not be extendable. Consequently, after embedding the associated decorated graph, some connected components may appear as immersed segments

$$I\colon [0,1] \looparrowright T.$$

To obtain the invariant $\mathrm{hfk}(P)$, the endpoints of such segments are connected to the puncture. This produces immersed arcs

$$\widetilde{I}\colon \mathbb{R} \looparrowright T,$$

which form the non-compact components of $\mathrm{hfk}(P)$.

Rasmussen's immersed curve invariant satisfies a pairing theorem analogous to Theorem 2.99. Using the same gluing setup as in Theorem 2.85, let

$$Y = M \cup_h (S^1 \times D^2),$$

and let $\bar{h}$ denote the composition of h with the elliptic involution of ∂M.

Theorem 2.112. *For a knot $P \subset S^1 \times D^2$, let $K \subset Y$ denote its image under the gluing, and set*

$$\boldsymbol{\gamma} = \widehat{HF}(M), \qquad \boldsymbol{\gamma}' = \mathit{hfk}(P).$$

Then

$$\widehat{HFK}(Y,K) \cong HF\big(\boldsymbol{\gamma}, \bar{h}(\boldsymbol{\gamma}')\big).$$

Remark 2.113. The discussion of Lagrangian intersection Floer homology extends naturally to the knot invariant $\mathrm{hfk}(P)$. In particular, the decomposition over connected components (Proposition 2.101), the reduction to primitive curves (Lemma 2.102), the notions of unobstructed and commensurable curves, and the dimension formula of Theorem 2.105 all remain valid in this setting.

Likewise, the notions of ε-geodesics, pegboard diagrams, and singular pegboard diagrams extend to $\mathrm{hfk}(P)$. Consequently, Theorem 2.109 continues to hold for immersed curve invariants of knots.

2.3.3 Examples and Computations

We conclude this section with examples of the invariants introduced in previous subsections. These examples serve two purposes. First, they illustrate how genus-one bordered Heegaard diagrams can be used to compute type A, type D, and immersed curve invariants. Second, they provide the concrete input that will be used in later chapters.

In genus one, the analytic counts appearing in the definitions simplify considerably. More precisely, the relevant counts of pseudo-holomorphic curves reduce to combinatorial counts of embedded bigons in a suitable covering space. We recall the setup briefly before turning to examples.

Definition 2.114. A *pointed genus-one bordered Heegaard diagram* is a tuple

$$(\Sigma, \boldsymbol{\alpha}, \beta, z),$$

where Σ is a compact oriented surface of genus one with one boundary component, $\boldsymbol{\alpha} = (\alpha_1, \alpha_2)$ is a pair of disjoint properly embedded arcs in Σ whose endpoints alternate on $\partial\Sigma$, β is an embedded closed curve in the interior of Σ such that $\Sigma \setminus \beta$ is connected and β intersects $\boldsymbol{\alpha}$ transversely, and $z \in \partial\Sigma \setminus \partial\boldsymbol{\alpha}$ is a basepoint. The boundary arcs are labelled as in Figure 2.26.

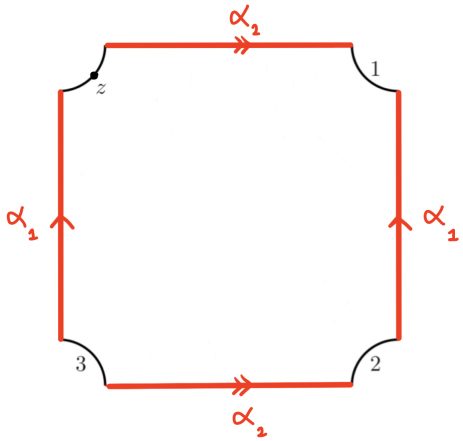


Figure 2.26: Boundary labelling convention for genus-one bordered Heegaard diagrams.

Notation 2.115. We use the symbol $I \in \{12, 23, 123\}$ to denote the arc obtained by concatenating the boundary arcs labelled by 1, 2, and 3 accordingly.

Such a diagram determines a bordered 3-manifold (M, α_1, α_2). The manifold M is obtained by attaching a 2-handle to $\Sigma \times [0,1]$ along $\beta \times \{1\}$ and then collapsing the disk bounded by $\partial\Sigma \times \{0\}$ so that the arcs α_1 and α_2 become closed curves on ∂M.

Adding a second basepoint

$$w \in \Sigma^\circ \setminus (\boldsymbol{\alpha} \cup \beta)$$

gives a *doubly-pointed genus-one bordered Heegaard diagram*. As in ordinary knot Floer homology, this additional basepoint specifies a knot in the bordered 3-manifold.

In genus one, the type A invariant admits a combinatorial description in terms of embedded bigons.

Construction 2.116. Let $(\Sigma, \boldsymbol{\alpha}, \beta, z)$ be a pointed genus-one bordered Heegaard diagram representing a bordered manifold (M, α_1, α_2).

The generators of $\widehat{\mathrm{CFA}}(\mathcal{H}, z)$ are the intersection points

$$\mathcal{G} = \beta \cap \boldsymbol{\alpha}.$$

The idempotent action is determined by

$$\widehat{\mathrm{CFA}}(\mathcal{H}, z) \cdot \iota_0 = \langle x \in \mathcal{G} \mid x \in \alpha_1 \cap \beta \rangle,$$

$$\widehat{\mathrm{CFA}}(\mathcal{H}, z) \cdot \iota_1 = \langle x \in \mathcal{G} \mid x \in \alpha_2 \cap \beta \rangle.$$

To describe the $\mathcal{A}_\infty$ operations, view

$$\Sigma = (S^1 \times S^1) \setminus \mathrm{int}(B),$$

where B is a disk. Let $\widetilde{\Sigma}$ be the covering space obtained from the universal cover $\mathbb{R}^2 \to S^1 \times S^1$ by removing the preimages of B.

The maps

$$m_{n+1} : \widehat{\mathrm{CFA}}(\mathcal{H}, z) \otimes \mathcal{A}^{\otimes n} \longrightarrow \widehat{\mathrm{CFA}}(\mathcal{H}, z)$$

are determined by counts of index-one embedded disks in $\widetilde{\Sigma}$. More precisely,

$$m_{n+1}(x, \rho_{i_1}, \dots, \rho_{i_n}) = \sum_{y \in \mathcal{G}} \#\mathcal{M}(x, y; \rho_{i_1}, \dots, \rho_{i_n})\, y,$$

where

$$i_j \in \{1, 2, 3, 12, 23, 123\}, \qquad j = 1, \dots, n,$$

and $\#\mathcal{M}(x, y; \rho_{i_1}, \dots, \rho_{i_n})$ is the mod two count of index-one embedded disks in $\widetilde{\Sigma}$ whose boundary, with the induced orientation, starts at a lift $\widetilde{x}$ of x, alternates between arcs on lifts of $\boldsymbol{\alpha}$ and arcs i_j on lifts of ∂B for $j = 1, \dots, n$, ends at a lift $\widetilde{y}$ of y, and finally follows an arc on a lift of β back to $\widetilde{x}$.

The disks counted by

$$\mathcal{M}(x, y; \rho_{i_1}, \dots, \rho_{i_n})$$

avoid all lifts of the basepoint z. In genus one, these moduli spaces are particularly simple and can often be identified combinatorially by counting embedded bigons. The index can be computed using the formula of [LOT18, Section 5.7], and in many examples it may also be determined using the Riemann mapping theorem.

The construction of $\widehat{\mathrm{CFA}}(\mathcal{H}, z, w)$ is identical, except that one replaces

$$\mathcal{M}(x, y; \rho_{i_1}, \dots, \rho_{i_n})$$

by the subspace

$$\mathcal{M}_w(x, y; \rho_{i_1}, \dots, \rho_{i_n})$$

consisting of disks that avoid all lifts of both basepoints z and w.

The type D invariant admits a similar description. The main difference is that the boundary labelling convention is changed as shown in Figure 2.27. We continue to use the notation introduced in Construction 2.116.

Construction 2.117. The generators of $\widehat{\mathrm{CFD}}(\mathcal{H}, z)$ are again the set of intersection points $\mathcal{G}$. The idempotent action is determined by

$$\widehat{\mathrm{CFD}}(\mathcal{H}, z) \cdot \iota_0 = \langle x \in \mathcal{G} \mid x \in \alpha_2 \cap \beta \rangle,$$

$$\widehat{\mathrm{CFD}}(\mathcal{H}, z) \cdot \iota_1 = \langle x \in \mathcal{G} \mid x \in \alpha_1 \cap \beta \rangle.$$

The differential

$$\delta : \widehat{\mathrm{CFD}}(\mathcal{H}, z) \longrightarrow \mathcal{A} \otimes \widehat{\mathrm{CFD}}(\mathcal{H}, z)$$

is determined by the same moduli spaces appearing in Construction 2.116:

$$\delta(x) = \sum_{y \in \mathcal{G}} \#\mathcal{M}(x, y; \rho_{i_1}, \dots, \rho_{i_n}) \, (\rho_{i_1} \cdots \rho_{i_n} \otimes y).$$

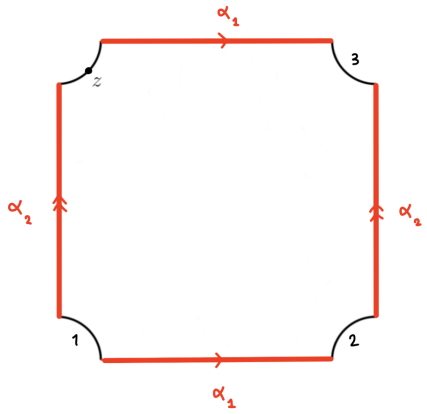


Figure 2.27: Labeling arcs on the boundary of the genus-one bordered Heegaard diagram used to define type D invariants

Remark 2.118. The $\mathcal{A}_\infty$ homotopy type of $\widehat{\mathrm{CFA}}(\mathcal{H}, z)$ and the homotopy type of $\widehat{\mathrm{CFD}}(\mathcal{H}, z)$ determine the bordered invariants

$$\widehat{\mathrm{CFA}}(M, \alpha_1, \alpha_2) \qquad \text{and} \qquad \widehat{\mathrm{CFD}}(M, \alpha_1, \alpha_2).$$

Likewise, the homotopy types of $\widehat{\mathrm{CFA}}(\mathcal{H}, z, w)$ and $\widehat{\mathrm{CFD}}(\mathcal{H}, z, w)$ determine the knot invariants

$$\widehat{\mathrm{CFA}}(M, K) \qquad \text{and} \qquad \widehat{\mathrm{CFD}}(M, K).$$

We now compute some examples.

Example 2.119. Let $Q \subset S^1 \times D^2$ be the Mazur pattern, shown in Figure 2.28. This example has also been studied by Petkova and Wong [PW20] and Levine [Lev14]. A doubly-pointed genus-one bordered Heegaard diagram $(\mathcal{H}_Q, z, w)$ for the pair $(S^1 \times D^2, Q)$ is shown in Figure 2.29.

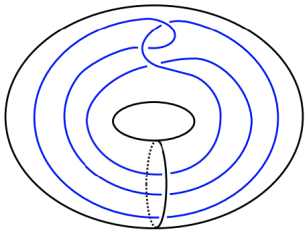


Figure 2.28: Mazur pattern in solid torus [PW20]

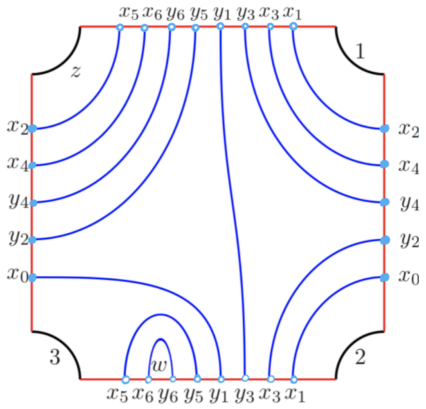


Figure 2.29: Doubly-pointed genus-one bordered Heegaard diagram for the Mazur pattern.

Hence

$$\widehat{\mathrm{CFA}}(\mathcal{H}_Q, z, w) = \langle x_0, \dots, x_6, y_1, \dots, y_6 \rangle_{\mathbb{F}_2}.$$

The idempotent actions are

$$\widehat{\mathrm{CFA}}(\mathcal{H}_Q, z, w) \cdot \iota_0 = \langle x_0, y_2, y_4, x_4, x_2 \rangle,$$

$$\widehat{\mathrm{CFA}}(\mathcal{H}_Q, z, w) \cdot \iota_1 = \langle x_5, x_6, y_6, y_5, y_1, y_3, x_3, x_1 \rangle.$$

The index-one disks contributing to the type A structure are shown in Figure 2.30. The nonzero $\mathcal{A}_\infty$ operations are

$$m_2(x_2 \otimes \rho_1) = x_1, \qquad m_2(x_4 \otimes \rho_1) = x_3, \qquad m_2(y_4 \otimes \rho_1) = y_3,$$

$$m_2(x_1 \otimes \rho_2) = x_0, \qquad m_2(x_3 \otimes \rho_2) = y_2, \qquad m_2(x_2 \otimes \rho_{12}) = x_0, \qquad m_2(x_4 \otimes \rho_{12}) = y_2,$$

$$m_3(y_3 \otimes \rho_2 \otimes \rho_1) = y_1, \qquad m_4(y_4 \otimes \rho_1 \otimes \rho_2 \otimes \rho_1) = y_1.$$

Figure 2.31 shows the corresponding decorated graph in the conventions of Hedden and Levine [HL16].

We next determine the grading structure. To do so, we first consider the type A structure

$$\widehat{\mathrm{CFA}}(\mathcal{H}_Q, z).$$

This requires allowing disks that pass through the basepoint w. The corresponding index-one disks are shown in Figure 2.33. Including these disks in the count yields the type A structure associated to the singly-pointed diagram $(\mathcal{H}_Q, z)$, shown in Figure 2.32.

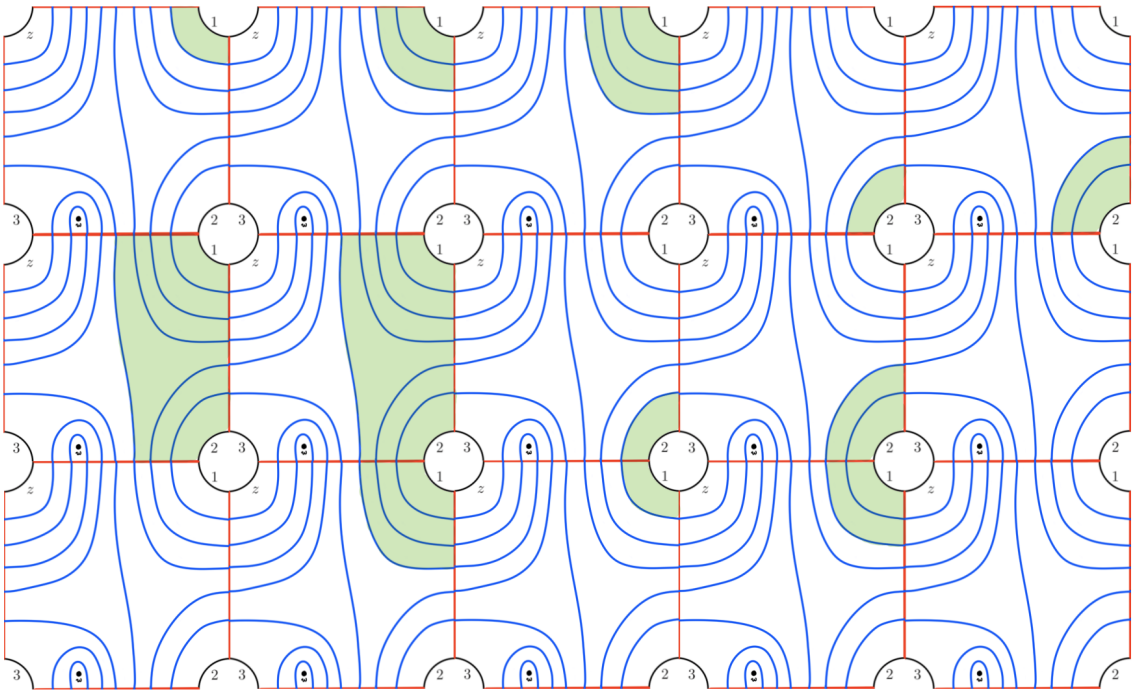


Figure 2.30: Index-one embedded disks in $\widetilde{\Sigma}$ avoiding all lifts of z and w.

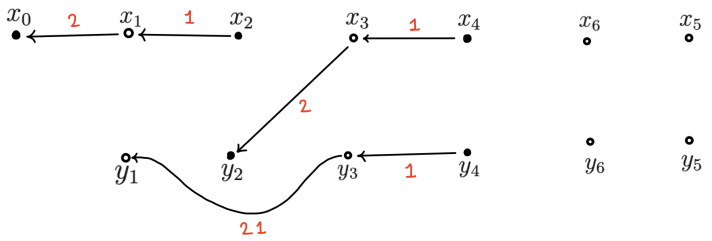


Figure 2.31: Decorated graph representing $\widehat{\mathrm{CFA}}(\mathcal{H}_Q, z, w)$.

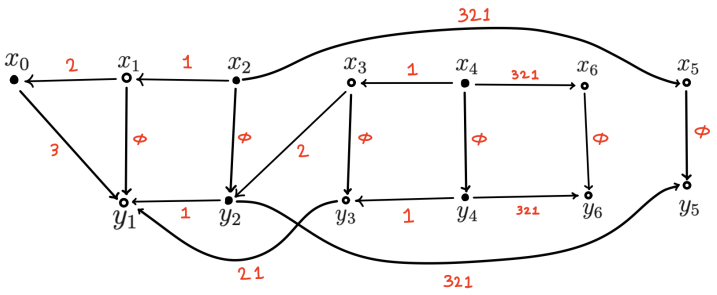


Figure 2.32: Decorated graph representing $\widehat{\mathrm{CFA}}(\mathcal{H}_Q, z)$.

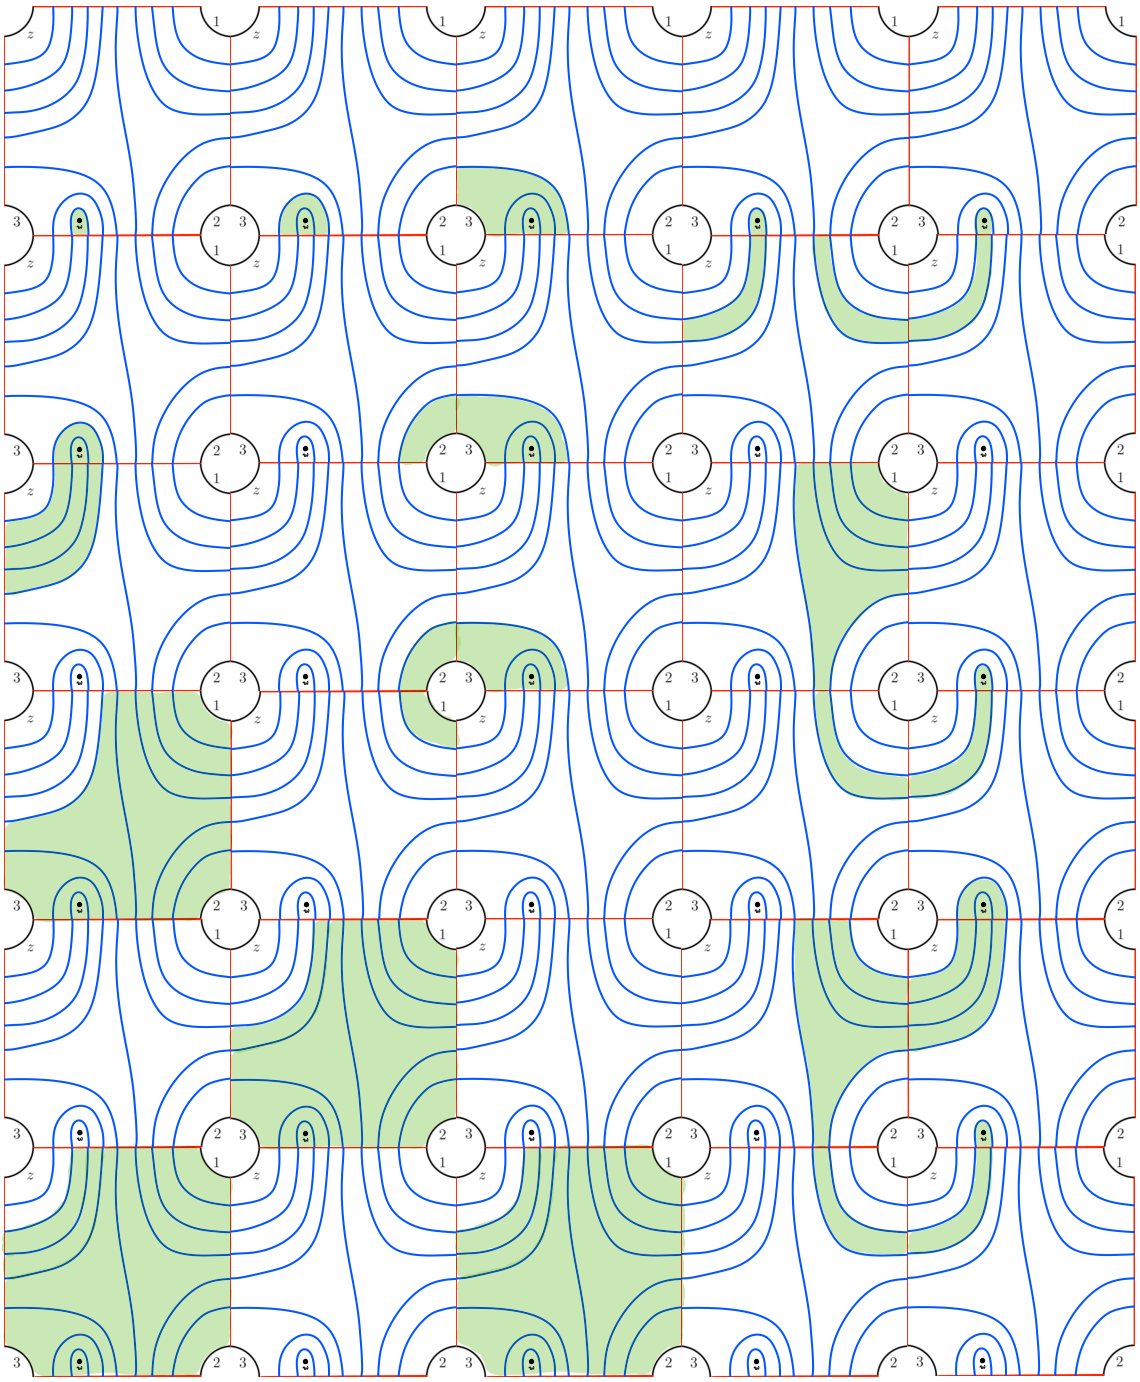


Figure 2.33: Index-one embedded disks in $\widetilde{\Sigma}$ avoiding lifts of z and passing through a lift of w.

We next determine the grading structure. Choose y_4 as the reference generator and set

$$\mathrm{gr}(y_4) = [(0;0,0;0)].$$

Using the grading rules of Table 2.22, one computes the gradings of the remaining generators:

$$\begin{aligned}
\mathrm{gr}(y_1) &= \left(-\frac{1}{2}; \frac{3}{2}, -\frac{1}{2}; 0\right), \qquad \mathrm{gr}(y_2) = (-2; 1, -1; 0),\\
\mathrm{gr}(y_3) &= \left(-\frac{1}{2}; \frac{1}{2}, -\frac{1}{2}; 0\right), \qquad \mathrm{gr}(y_5) = \left(-\frac{3}{2}; \frac{3}{2}, -\frac{1}{2}; -1\right),\\
\mathrm{gr}(y_6) &= \left(-\frac{7}{2}; \frac{1}{2}, -\frac{1}{2}; -2\right),
\end{aligned}$$

and

$$\begin{aligned}
\mathrm{gr}(x_0) &= \left(-\frac{1}{2}; 2, -1; 1\right), \qquad \mathrm{gr}(x_1) = \left(-\frac{3}{2}; \frac{3}{2}, -\frac{3}{2}; 1\right),\\
\mathrm{gr}(x_2) &= (-1; 1, -1; 1), \qquad \mathrm{gr}(x_3) = \left(-\frac{5}{2}; \frac{1}{2}, -\frac{3}{2}; 0\right),\\
\mathrm{gr}(x_4) &= \left(-\frac{5}{2}; 0, -1; 0\right), \qquad \mathrm{gr}(x_5) = \left(-\frac{1}{2}; \frac{3}{2}, -\frac{1}{2}; 0\right),\\
\mathrm{gr}(x_6) &= \left(-\frac{5}{2}; \frac{1}{2}, -\frac{1}{2}; -1\right).
\end{aligned}$$

Since the bordered manifold represented by $\mathcal{H}_Q$ is a solid torus, the subgroup $P(y_4)$ is cyclic. Applying the grading rules to the cycle in Figure 2.32 gives

$$[(0;0,0;0)] = \left[\left(-\frac{7}{2}; 0, -1; -1\right)\right]$$

in the grading coset. Hence

$$P(y_4) = \left\langle \left(-\frac{7}{2}; 0, -1; -1\right) \right\rangle.$$

We now compute the immersed curve invariant $\mathrm{hfk}(Q)$. A decorated graph representing

$$\widehat{\mathrm{CFD}}(\mathcal{H}_Q, z, w)$$

is obtained from Figure 2.31 by exchanging the labels 1 and 3. Embedding this graph in the punctured torus according to Construction 2.98 gives the collection of segments shown in Figure 2.34. Denote these segments by

$$I_1, \ldots, I_7.$$

In this example, four of the segments are constant, namely $I_4, \ldots, I_7$, with images x_6, x_6, y_5, y_6, respectively.

Since $\widehat{\mathrm{CFD}}(\mathcal{H}_Q, z, w)$ is not extendable, these segments do not close up to form immersed curves. Instead, the endpoints are connected to the puncture, producing immersed arc components. The resulting collection of arcs is precisely the immersed curve invariant $\mathrm{hfk}(Q)$ shown in Figure 2.35. In particular, $\mathrm{hfk}(Q)$ consists entirely of immersed arc components.

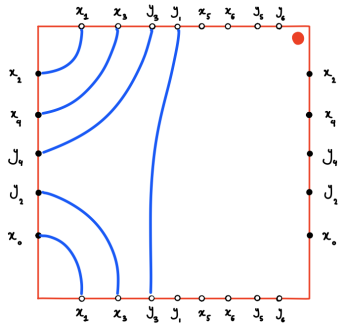


Figure 2.34: The decorated graph of $\widehat{\mathrm{CFD}}(\mathcal{H}_Q, z, w)$ embedded in the punctured torus.

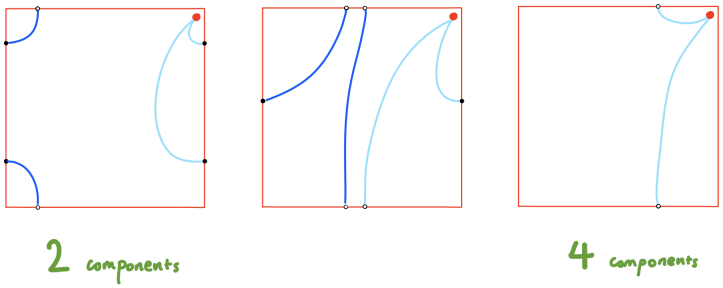


Figure 2.35: The immersed curve invariant $\mathrm{hfk}(Q)$ of the Mazur pattern.

Remark 2.120. The singly-pointed diagram $(\mathcal{H}_Q, z)$ represents a solid torus. This can be seen by isotoping the β curve to obtain the diagram $(\mathcal{H}_\infty, z)$ shown in Figure 2.36.

A direct computation gives

$$\widehat{\mathrm{CFA}}(\mathcal{H}_\infty, z) = \widehat{\mathrm{CFA}}(\mathcal{H}_\infty, z) \cdot \iota_0 = \langle x_0 \rangle,$$

and the only nontrivial higher operations are

$$m_{n+3}\big(x_0 \otimes \rho_3 \otimes \underbrace{\rho_{23} \otimes \cdots \otimes \rho_{23}}_{n} \otimes \rho_2\big) = x_0, \qquad n \geq 0.$$

The grading is determined by

$$\mathrm{gr}(x_0) = [(0;0,0)] \in \langle (-\tfrac{1}{2};0,1) \rangle \backslash G.$$

By Theorem 2.95, there exists an $\mathcal{A}_\infty$ homotopy equivalence

$$f : \widehat{\mathrm{CFA}}(\mathcal{H}_\infty, z) \longrightarrow \widehat{\mathrm{CFA}}(\mathcal{H}_Q, z).$$

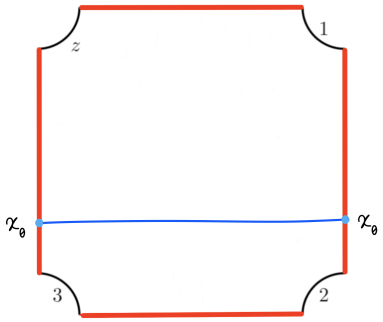


Figure 2.36: The diagram $(\mathcal{H}_\infty, z)$ obtained from $(\mathcal{H}_Q, z)$ by isotoping the β curve.

Example 2.121. Let $H'_{1/m}$ be a solid torus whose boundary is parametrized by curves α_1 and α_2. We choose the parametrization so that the meridian has slope $-m$ with respect to (α_1, α_2). Equivalently, the meridian represents $[\alpha_1] - m[\alpha_2]$ in homology.

We compute the type D invariant $\widehat{\mathrm{CFD}}(H'_{1/m}, \alpha_1, \alpha_2)$.

We first consider the case $m = 3$. Let $(\mathcal{H}'_{1/3}, z')$ be the genus-one bordered Heegaard diagram shown in Figure 2.37, representing the bordered manifold $(H'_{1/3}, \alpha_1, \alpha_2)$. The index-one disks contributing to the differential are shown in Figure 2.38.

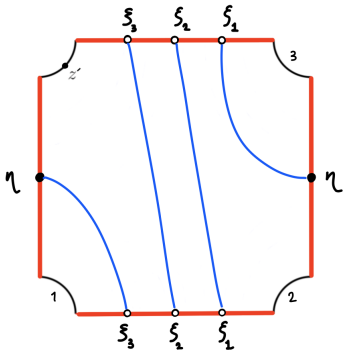


Figure 2.37: Genus-one bordered Heegaard diagram $(\mathcal{H}'_{1/3}, z')$ representing $(H'_{1/3}, \alpha_1, \alpha_2)$.

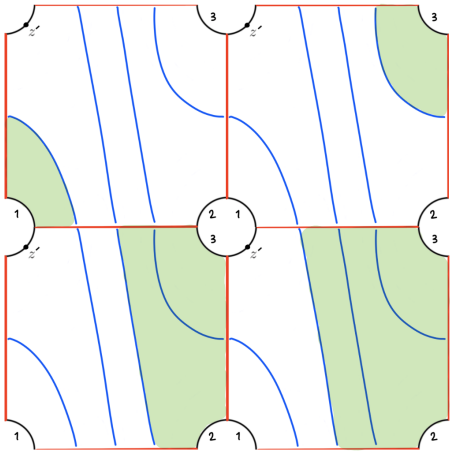


Figure 2.38: Index-one embedded disks in $\widetilde{\Sigma}$ contributing to the differential of $\widehat{\mathrm{CFD}}(\mathcal{H}'_{1/3}, z')$.

The computation gives

$$\widehat{\mathrm{CFD}}(\mathcal{H}'_{1/3}, z') = \langle \eta, \xi_1, \xi_2, \xi_3 \rangle_{\mathbb{F}_2},$$

with idempotent actions

$$\widehat{\mathrm{CFD}}(\mathcal{H}'_{1/3}, z') \cdot \iota_0 = \langle \eta \rangle, \qquad \widehat{\mathrm{CFD}}(\mathcal{H}'_{1/3}, z') \cdot \iota_1 = \langle \xi_1, \xi_2, \xi_3 \rangle.$$

The differential is given by

$$\delta(\eta) = \rho_3 \otimes \xi_1 + \rho_1 \otimes \xi_3,$$

$$\delta(\xi_i) = \rho_{23} \otimes \xi_{i+1}, \qquad i = 1, 2.$$

The corresponding decorated graph is shown in Figure 2.39.

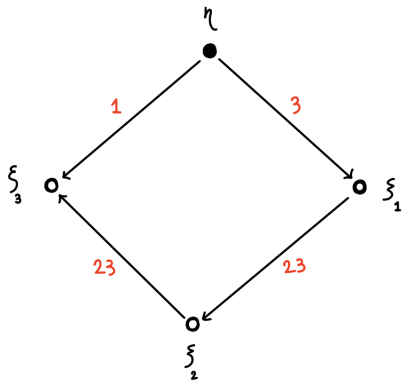


Figure 2.39: Decorated graph representing $\widehat{\mathrm{CFD}}(\mathcal{H}'_{1/3}, z')$.

The same computation extends to arbitrary $m \geq 1$. Starting from a curve parallel to α_1, the diagram $(\mathcal{H}'_{1/m}, z')$ is obtained by applying m Dehn twists in the α_2-direction to the β curve. Equivalently, β represents the meridian of $H'_{1/m}$ and intersects α_2 once and α_1 exactly m times.

The resulting type D structure is

$$\widehat{\mathrm{CFD}}(\mathcal{H}'_{1/m}, z') = \langle \eta, \xi_1, \dots, \xi_m \rangle_{\mathbb{F}_2},$$

with idempotent actions

$$\widehat{\mathrm{CFD}}(\mathcal{H}'_{1/m}, z') \cdot \iota_0 = \langle \eta \rangle, \qquad \widehat{\mathrm{CFD}}(\mathcal{H}'_{1/m}, z') \cdot \iota_1 = \langle \xi_1, \dots, \xi_m \rangle,$$

and differential

$$\delta(\eta) = \rho_3 \otimes \xi_1 + \rho_1 \otimes \xi_m,$$
$$\delta(\xi_i) = \rho_{23} \otimes \xi_{i+1}, \qquad i = 1, \dots, m-1.$$

The corresponding decorated graph is shown in Figure 2.40. Note that the underlying graph is a cycle consisting of $m+1$ vertices.

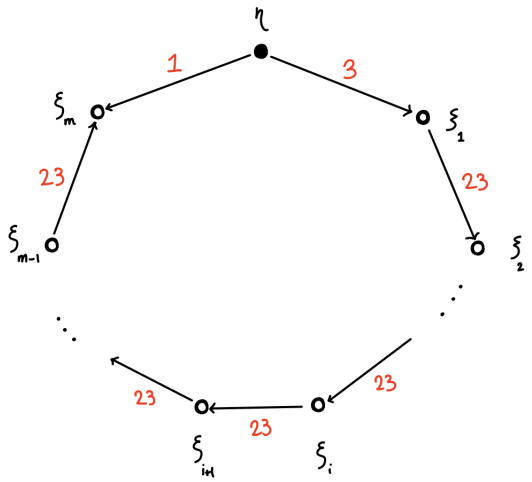


Figure 2.40: Decorated graph representing $\widehat{\mathrm{CFD}}(\mathcal{H}'_{1/m}, z')$.

We now compute the grading structure. Let gr_m denote the grading on $\widehat{\mathrm{CFD}}(\mathcal{H}'_{1/m}, z')$. Choose η as the reference generator, so

$$\mathrm{gr}_m(\eta) = [(0; 0, 0)].$$

Moving clockwise around the cycle in Figure 2.40 gives the representatives

$$\mathrm{gr}_m(\xi_i) = \left[\left(-\frac{1}{2}; \frac{1}{2}, -\frac{2i-1}{2}\right)\right], \qquad i = 1, \dots, m. \tag{2.6}$$

The cycle also determines the subgroup

$$P(\eta) = \left\langle \left(-\frac{m-1}{2}; 1, -m\right) \right\rangle,$$

so that gr_m takes values in the coset space

$$G/\left\langle \left(-\frac{m-1}{2}; 1, -m\right) \right\rangle.$$

The same cosets can also be represented by moving counterclockwise around the cycle. This gives

$$\mathrm{gr}_m(\xi_m) = \left[\left(-\frac{1}{2}; -\frac{1}{2}, \frac{1}{2}\right)\right],$$

and more generally

$$\mathrm{gr}_m(\xi_{m-i}) = \left[\left(\frac{2i-1}{2}; -\frac{1}{2}, \frac{2i+1}{2}\right)\right], \qquad i = 0, \dots, m-1. \tag{2.7}$$

These are different representatives of the same grading cosets. This second choice will be useful when comparing the complexes for consecutive values of m.

Finally, we identify the immersed curve invariant. Figure 2.41 shows the embedding of the decorated graph for

$$\widehat{\mathrm{CFD}}(\mathcal{H}'_{1/3}, z')$$

in the punctured torus. The resulting immersed curve has a loose representative given by a Euclidean geodesic lifting to a line of slope 3 in $\widetilde{T}$. More generally, $\widehat{HF}(H'_{1/m})$ is represented by a loose Euclidean geodesic lifting to a line of slope m in $\widetilde{T}$.

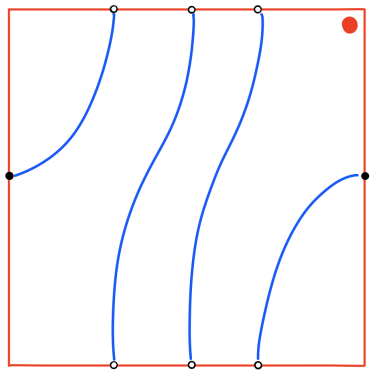


Figure 2.41: The decorated graph of $\widehat{\mathrm{CFD}}(\mathcal{H}'_{1/3}, z')$ embedded in the punctured torus.

Remark 2.122. Later, we compare the type D structures

$$\widehat{\mathrm{CFD}}(\mathcal{H}'_{1/m}, z') \qquad \text{and} \qquad \widehat{\mathrm{CFD}}(\mathcal{H}'_{1/(m+1)}, z')$$

as $m \to \infty$. To simplify notation, we use the same symbols for the corresponding generators:

$$\widehat{\mathrm{CFD}}(\mathcal{H}'_{1/m}, z') = \langle \eta, \xi_1, \dots, \xi_m \rangle, \qquad \widehat{\mathrm{CFD}}(\mathcal{H}'_{1/(m+1)}, z') = \langle \eta, \xi_1, \dots, \xi_{m+1} \rangle.$$

Definition 2.123. The representatives given by (2.6) are called the *clockwise representatives* of gr_m. Similarly, the representatives given by (2.7) are called the *counterclockwise representatives* of gr_m.

Combining the clockwise and counterclockwise representatives gives a convenient choice of representatives near the two ends of the cycle.

Corollary 2.124. *For any $k \in \mathbb{N}$ and any integer $m > 2k$, one may choose the clockwise representatives for $\xi_1, \dots, \xi_{k+1}$ and the counterclockwise representatives for $\xi_{m-k}, \dots, \xi_m$. With this choice, the gradings of*

$$\xi_{m-k}, \dots, \xi_m, \eta, \xi_1, \dots, \xi_{k+1}$$

admit representatives that are independent of m.

Chapter 3

The Fox Trapezoidal Conjecture

As discussed in Subsection 2.1.3, the invariants of alternating links often exhibit striking patterns. One such example is the behaviour of the coefficients of the Alexander polynomial. In particular, Theorem 2.38 showed that the coefficients of the Alexander polynomial of an alternating link alternate in sign. Consequently, it is natural to focus on the sequence of their absolute values.

This chapter is devoted to a much stronger conjectural pattern governing this sequence. In 1962, Fox conjectured that the absolute values of the coefficients of the Alexander polynomial of an alternating link form a trapezoidal sequence. The precise statement is given in Conjecture 3.1.

Conjecture 3.1 ([Fox61])**.** *Let* $(a_0, \dots, a_n)$ *be the sequence of absolute values of the coefficients of the Alexander polynomial* Δ_K *of an alternating link* K*, and let* $m = \left\lfloor \frac{n}{2} \right\rfloor$*. Then*

$$a_0 \leq a_1 \leq \cdots \leq a_m \qquad \text{and} \qquad a_m \geq a_{m+1} \geq \cdots \geq a_n.$$

Moreover, if $a_i = a_{i+1}$ *for some* $i < m$*, then*

$$a_i = a_{i+1} = \cdots = a_m.$$

A sequence satisfying these conditions is called trapezoidal.

An example of a trapezoidal sequence is shown in Figure 3.1. The graph has been normalized to be symmetric about the origin. Intuitively, the Fox trapezoidal conjecture predicts that the graph of the absolute values of the coefficients has a concave shape.

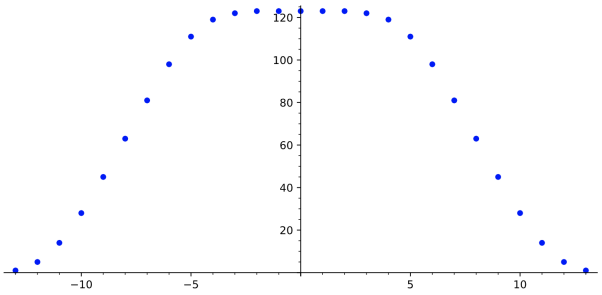


Figure 3.1: Graph of a trapezoidal sequence.

In 2005, Stoimenow proposed the following strengthening of Conjecture 3.1.

Conjecture 3.2 ([Sto05])**.** *Let* $(a_0, \dots, a_n)$ *be the sequence of absolute values of the coefficients of the Alexander polynomial* Δ_K *of an alternating link* K*. Then* $(a_0, \dots, a_n)$ *is* log-concave*; that is,*

$$a_{i-1}a_{i+1} \leq a_i^2$$

for every $i \in \{1, \dots, n-1\}$*.*

Log-concavity is a significantly stronger condition than trapezoidality. In particular, Conjecture 3.2 implies Conjecture 3.1. Hence, Conjecture 3.2 is often referred to as the *strong Fox conjecture*.

Despite decades of work, both conjectures remain open in full generality. Following Huh's breakthrough work on log-concavity, culminating in the proof of the Rota–Heron–Welsh conjecture, the Fox and Stoimenow conjectures have attracted renewed attention in recent years. In particular, Hafner, Mészáros, and Vidinas [HMV24] proved the Fox trapezoidal conjecture for special alternating links. Their proof builds on Kálmán's work [Kál13] on the combinatorics of bipartite plane graphs and relates the Alexander polynomial of special alternating links to *graphical matroids*, thereby bringing it within the scope of the *Lorentzian polynomial* machinery.

In this chapter, we discuss some of the most recent progress on these conjectures obtained in joint work with András Juhász and Tamás Kálmán [AJK24]. In particular, we prove the Fox trapezoidal conjecture for several new classes of alternating links and obtain further evidence for Stoimenow's stronger conjecture.

The chapter is organised as follows. In Section 3.1, we survey the main known results on the Fox and Stoimenow conjectures and briefly discuss the methods used to obtain them. Table 3.1 summarises the principal results discussed in Section 3.1.

Table 3.1: Known results on the Fox trapezoidal conjecture.

Known Results on the Fox Trapezoidal Conjecture		
Reference	**Class of alternating links**	**Method**
Parris [Par78]	Alternating pretzel knots	Explicit formula for Δ_K
Hartley [Har79]	Two-bridge links	Relating a formula for Δ_K to extended Schubert diagrams
Murasugi [Mur85]	Alternating arborescent links	Skein formula for plumbing trees and induction
Ozsváth–Szabó [OS03a]	Alternating knots with $g(K) \leq 2$	Computation of $\mathrm{HF}^+(S^3_0(K))$ using the large surgery formula
Jong[Jon10], **Stoimenow** [Sto16]	Alternating knots with $g(K) \leq 4$	Reduction to a finite list of genus generators using the $\overline{t_2'}$ move
Alrefai–Chbili [AC21]	Alternating three-braids of length less than 4	Formula for Δ_K using the Burau representation
Hafner–Mészáros–Vidinas [HMV24]	Special alternating links	Multivariable refinement of Δ_K and Lorentzian polynomial techniques

In Section 3.2, we study a class of alternating links with (linearly) large twist regions. Using Crowell's formula for the Alexander polynomial, we prove the trapezoidal conjecture for this class of alternating links. More generally, we show that the existence of a twist region with sufficiently many crossings forces the first inequalities predicted by the trapezoidal conjecture.

The next two sections extend the work of Hafner, Mészáros, and Vidinas in two different directions. In Section 3.3, we consider diagrammatic plumbings of special alternating links. We develop a formula for the Alexander polynomial of such plumbings. As an application, we prove the trapezoidal conjecture for this class.

Finally, in Section 3.4, we study alternating three-braids, which form one of the simplest classes of non-special alternating links. Motivated by the approach of Hafner, Mészáros, and Vidinas, we investigate a natural multivariable refinement of the Alexander polynomial. This refinement fails to be Lorentzian in general, preventing a direct extension of their argument to alternating three-braids. Nevertheless, we derive several inequalities between its coefficients and consequently obtain new restrictions on the coefficients of the Alexander polynomial.

3.1 A Survey of Known Results

Parris [Par78] proved the trapezoidal conjecture for alternating pretzel knots in his PhD thesis. His proof uses Fox calculus to derive an explicit formula for the Alexander polynomial of pretzel knots.

Hartley [Har79] proved the trapezoidal conjecture for two-bridge links $D(p,q)$. The Alexander polynomial $\Delta_{D(p,q)}$ can be computed using the following formula. Let p and q be coprime integers satisfying $0 < q < p$ with q odd. Then

$$\Delta_{D(p,q)}(t) = \sum_{0\le k\le p-1} (-1)^k\, t^{s_k}, \qquad s_k = \sum_{i=0}^{k} (-1)^{\left\lfloor \frac{iq}{p} \right\rfloor}. \tag{3.1}$$

Minkus [Min82] proved this formula using Fox calculus on a Schubert normal diagram of $D(p,q)$. Using similar ideas, Hartley derived another formula for $\Delta_{D(p,q)}$ that can be read directly from an extended diagram obtained by unwinding Schubert's normal diagram; see Figure 3.2. Note that all diagonal segments in the extended diagram have slope $\frac{q}{p}$.

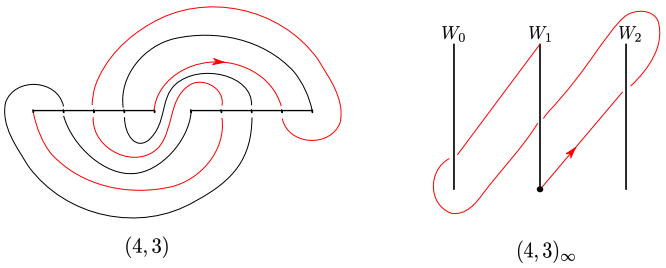


Figure 3.2: Hartley's extended diagram [Har79].

Proposition 3.3 ([Har79])**.** *Let α_i denote the number of segments joining the i-th and $(i+1)$-st vertical lines in the extended diagram of $D(p,q)$. Then*

$$\Delta_{D(p,q)} = \sum (-1)^i \alpha_i t^i.$$

Hartley used further combinatorial properties of the extended diagrams to construct an inductive argument proving the trapezoidal conjecture.

Remark 3.4. At first sight, Hartley's geometric interpretation of the coefficients of the Alexander polynomial appears somewhat mysterious. From the modern perspective of knot Floer homology and immersed curve invariants, however, such a relationship is perhaps less surprising. It is natural to ask whether Hartley's construction admits

a Floer-theoretic interpretation. To the best of the author's knowledge, no such connection has been formally established, and we leave this as an open question.

Murasugi [Mur85] proved the trapezoidal conjecture for alternating arborescent links. These are alternating links obtained by plumbing oriented twisted bands along a tree. Murasugi's proof uses skein relations on plumbing trees together with an inductive argument. Figure 3.3 illustrates the corresponding skein triple.

As depicted in Figure 3.3, changing the sign of a crossing decreases the number of half-twists in a band (equivalently, the weight of the corresponding vertex) by two, while smoothing a crossing produces a connected sum of the adjacent bands.

Furthermore, Murasugi derived a sufficient criterion for an algebraic link to be alternating. This criterion relies on relating the plumbing tree T to the Tait graph of the associated algebraic link $l(T)$. More precisely, Murasugi showed that if one augments T by an exterior vertex $\hat{v}$ and connects $\hat{v}$ to the vertices of T so that, in the augmented graph, the weight of each vertex is equal to the signed count of its incident edges, then $l(T)$ is isotopic to the link obtained by applying the median construction to the planar dual of $T \cup \{\hat{v}\}$; see Figure 3.4.

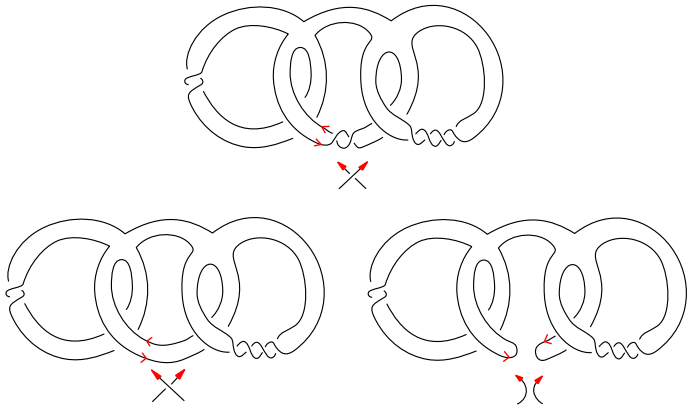


Figure 3.3: Skein relation on plumbing trees.

Ozsváth and Szabó [OS03a] used Heegaard Floer homology to prove the trapezoidal conjecture for alternating knots of genus two. Their starting point is the fact that alternating knots are *HF-thin*; namely,

$$\widehat{\mathrm{HFK}}(K,s) \cong \mathbb{Z}^{|a_s|}$$

supported in Maslov grading $s + \frac{\sigma}{2}$, where

$$\widetilde{\Delta}_K(t) = a_0 + \sum_{s=1}^{g} a_s(t^s + t^{-s}),$$

and $g(K)$ and $\sigma(K)$ denote the genus and signature of K, respectively.

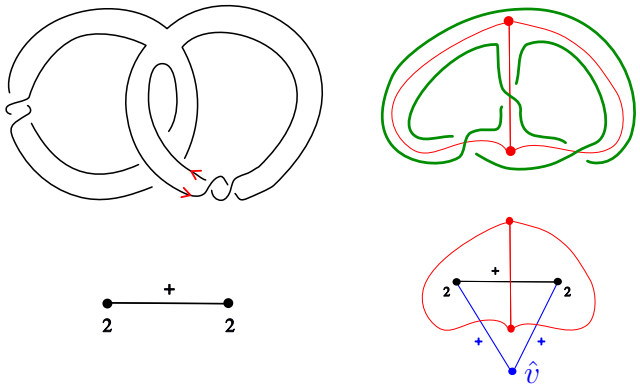


Figure 3.4: Murasugi's construction relating plumbing trees and Tait graphs; both the black and green diagrams represent the trefoil.

The proof relies on a Heegaard diagram whose Heegaard states are in one-to-one correspondence with Kauffman states. For alternating knots, this diagram satisfies

$$M(\boldsymbol{x}) - A(\boldsymbol{x}) = \frac{\sigma(K)}{2}$$

for every Heegaard state $\boldsymbol{x}$. As a consequence, all differentials in the hat version of the knot Floer chain complex vanish.

Ozsváth and Szabó went one step further and proved the following.

Proposition 3.5 ([OS03a])**.** *Let K be an alternating knot, oriented so that $\sigma = \sigma(K) \leq 0$. For all $s > 0$, let*

$$\delta(\sigma, s) := \max\left(0, \left\lceil \frac{|\sigma| - 2|s|}{4} \right\rceil\right)$$

and $t_s(K) := \sum_{j=1}^{\infty} j a_{|s|+j}$. Then we have a $\mathbb{Z}[U]$-module isomorphism

$$HF^+(S^3_0(K), s) \cong \mathbb{Z}^{b_s} \oplus \left(\mathbb{Z}[U]/U^{\delta(\sigma,s)}\right),$$

where $(-1)^{s+\frac{\sigma}{2}} b_s = \delta(\sigma, s) - t_s(K)$. Furthermore, for $s = 0$, we have

$$HF^+(S^3_0(K), 0) \cong \mathbb{Z}^{b_0} \oplus \mathcal{T}^+_{-1/2} \oplus \mathcal{T}^+_{-2\delta(\sigma,0)+\frac{1}{2}},$$

where $HF^+(S^3) = \mathcal{T}^+_0$ and the subscript denotes a shift in the $\mathbb{Q}$-grading.

Their proof uses the large surgery formula to compute $\mathrm{HF}^+(S^3_n(K), s)$ for sufficiently large n, and then applies the surgery exact sequence:

$$\cdots \to \mathrm{HF}^+(S^3) \to \mathrm{HF}^+(S^3_0(K), s) \to \mathrm{HF}^+(S^3_n(K), [s]) \to \cdots$$

A consequence of this computation is that $b_s \geq 0$. Using the inequality $b_0 \geq 0$ for alternating knots of genus two implies the trapezoidal conjecture.

Remark 3.6. One needs to combine the large surgery formula with some homological algebra to calculate $\mathrm{HF}^+(S^3_n(K), s)$. The following (Mayer–Vietoris) exact sequence of subcomplexes of $\mathrm{CFK}^\infty(S^3, K)$ is useful:

$$0 \to C\{i \geq 0 \text{ or } j \geq s\} \to C\{i \geq 0\} \oplus C\{j \geq s\} \to C\{i \geq 0 \text{ and } j \geq s\} \to 0$$

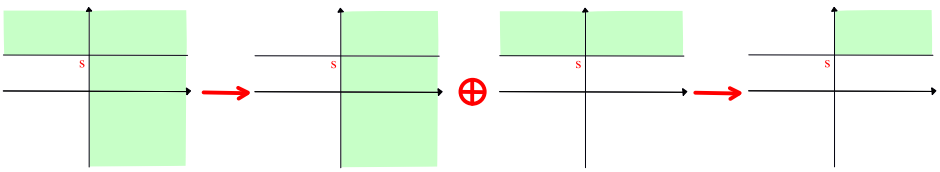


Figure 3.5: Exact sequence of subcomplexes of CFK^∞.

For an alternating knot, one can use $M(\boldsymbol{x}) - A(\boldsymbol{x}) = \frac{\sigma(K)}{2}$ to conclude that

$$M([\boldsymbol{x}, i, j]) = M(\boldsymbol{x}) + 2i = A(\boldsymbol{x}) + \frac{\sigma}{2} + 2i = (j - i) + \frac{\sigma}{2} + 2i = i + j + \frac{\sigma}{2}.$$

In particular, this means that $H_*(C\{i \geq 0 \text{ and } j \geq s\})$ is supported in degrees higher than $s + \frac{\sigma}{2}$. Combining this observation with the long exact sequence in homology associated to the above short exact sequence yields

$$\begin{aligned} H_{\leq s+\frac{\sigma}{2}-2}(C\{i \geq 0 \text{ or } j \geq s\}) &\cong H_{\leq s+\frac{\sigma}{2}-2}(C\{i \geq 0\}) \oplus H_{\leq s+\frac{\sigma}{2}-2}(C\{j \geq s\}) \\ &\cong \mathrm{HF}^+_{\leq s+\frac{\sigma}{2}-2}(S^3) \oplus \mathrm{HF}^+_{\leq -s+\frac{\sigma}{2}-2}(S^3). \end{aligned}$$

Recall that, based on the large surgery formula, we have

$$H_*(C\{i \geq 0 \text{ or } j \geq s\}) \cong \mathrm{HF}^+(S^3_n(K), [s])$$

for n sufficiently large.

Ni [Ni21] used these results to prove an extremal case of the inequalities predicted by the trapezoidal conjecture. More precisely, Ni showed that if $|a_g| = |a_{g-1}|$, then either K or its mirror is the torus knot $T_{2g+1,2}$.

Jong [Jon10] gave a combinatorial proof of the trapezoidal conjecture for alternating knots of genus at most two. The proof relies on Stoimenow's characterization of knot diagrams of canonical genus two [Sto08]. Using an analysis of Seifert graphs, Stoimenow showed that every alternating diagram of canonical genus two can be obtained from a finite list of diagrams, called *generators*, by applying flypes and $\overline{t_2'}$ moves.

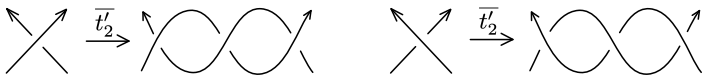


Figure 3.6: The move $\overline{t_2'}$; see Jong [Jon10].

Jong combined this characterization with Crowell's spanning tree formula for the Alexander polynomial of alternating links. By analyzing the effect of the move $\overline{t_2'}$ on Crowell's weighted graph, Jong established the trapezoidal conjecture through an inductive argument.

Stoimenow's characterization extends to any fixed genus. In principle, this reduces the trapezoidal conjecture for alternating knots of a given genus to a finite computation on the corresponding list of generators. Carrying out such a computation, however, is far from trivial. Using a number of computational techniques, **Stoimenow** [Sto16] proved the trapezoidal conjecture for all alternating knots of genus at most four.

Alrefai and Chbili [AC21] used the Burau representation to derive closed-form formulas for the Alexander polynomial of alternating three-braids of length at most three; that is, braid words of the form

$$\sigma_1^{p_1}\sigma_2^{-q_1}\sigma_1^{p_2}\sigma_2^{-q_2}\sigma_1^{p_3}\sigma_2^{-q_3}.$$

Using these formulas, they proved the trapezoidal conjecture for this class of alternating knots. Similar ideas were later used by **Alsukaiti and Chbili** [AC23] to prove the conjecture for three-braids of the form $(\sigma_1\sigma_2^{-1})^n$.

Banfield [Ban22] used Minkus's formula for the Alexander polynomial of two-bridge links to prove the strong Fox conjecture for this class. Their key observation is that Minkus's formula identifies the coefficients of the Alexander polynomial with the vertical slices of *folded Christoffel paths*.

The *Christoffel path* of slope $\frac{q}{p}$ is the lattice path in a $p \times q$ grid that stays below the diagonal and is closest to it. Here, "closest" means that there are no lattice points in the region bounded by the path and the diagonal. In this sense, Christoffel paths provide optimal lattice approximations to lines of rational slope; see Figure 3.7. Banfield used the combinatorics of *Christoffel words*, namely words in the alphabet $\{x, y\}$ obtained by recording the steps of a Christoffel path. Using a characterization of these words, they constructed an inductive proof of the strong Fox conjecture.

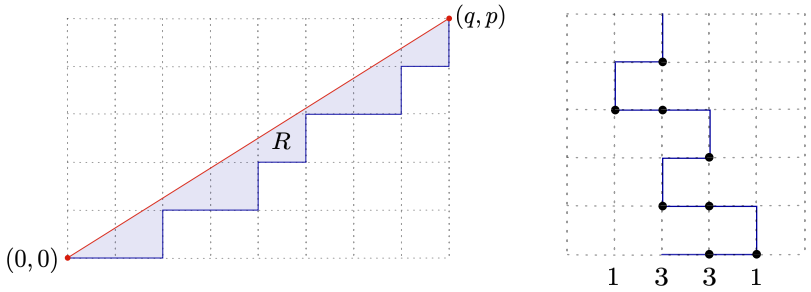


Figure 3.7: Christoffel paths (ordinary and folded); see Banfield [Ban22].

Hafner, Mészáros, and Vidinas [HMV24] proved the strong Fox conjecture for special alternating links. Their proof uses a multivariable refinement of the Alexander polynomial, which they show is Lorentzian. We discuss their work in Section 3.4.

3.2 Alternating Links with Large Twist Regions

The aim of this section is to prove the Fox trapezoidal conjecture for alternating links with (linearly) large twist regions; see Theorem 3.8. A neighborhood in a link diagram is called a *coherent twist region* if it is locally isomorphic to Figure 3.8. Such twist regions are obtained by twisting two parallel strands of a link.

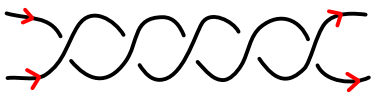


Figure 3.8: A coherent twist region.

Definition 3.7. The *maximal twist number* of an alternating link L, denoted by $\mathrm{MT}(L)$, is the maximum number of crossings contained in a coherent twist region among all reduced alternating diagrams of L.

Theorem 3.8. *Let L be an alternating link. Then the Fox trapezoidal conjecture holds for L if either of the following conditions is satisfied:*

1. *$MT(L)$ is odd and*
$$MT(L) - 2 \geq g(L) + \frac{|L|}{2},$$
2. *$MT(L)$ is even and*
$$MT(L) - 3 \geq g(L) + \frac{|L|}{2}.$$

The proof of Theorem 3.8 relies on the stabilization phenomenon for the coefficients of the Alexander polynomial under repeated twisting. In this section, we reprove this stabilization result for alternating links using Crowell's state-sum formula for the Alexander polynomial (Construction 2.39 and Theorem 2.41). This approach yields additional information about the stabilized coefficients, which will ultimately be used to prove Theorem 3.8.

Let D be a reduced alternating diagram of a link L satisfying $\mathrm{MT}(L) = n$, and let R be a coherent twist region with n crossings $c_1, \ldots, c_n$, labelled from left to right. For $k \geq 1$, let D_k be the diagram obtained by replacing R with a coherent twist region containing k crossings, and let L_k be the corresponding link. Let G_k be the Crowell graph associated to D_k, and let $v_1, \ldots, v_k$ denote the vertices corresponding to the crossings $c_1, \ldots, c_k$.

Consider the graph
$$G'_k = G_k \setminus \{v_{k-1}\}$$
for $k \geq 3$, where we also remove all edges adjacent to v_{k-1}; see Figure 3.9. To define G'_2, consider the two edges immediately to the left of v_1. Let
$$e_{li} = (v_{li}, v_1) \qquad \text{and} \qquad e_{lo} = (v_1, v_{lo})$$
be the incoming and outgoing edges, respectively. We define
$$G'_2 = G_2 \setminus \{v_1\} \cup \{e_{io}\},$$
where $e_{io} := (v_{li}, v_{lo})$. The weight of the new edge e_{io} is defined to be the product of the weights of e_{li} and e_{lo}.

Remark 3.9. The graph G'_k may be viewed as the unoriented smoothing of G_k at the vertex v_{k-1}. Let $L_k(c_{k-1}, ||)$ denote the 1-smoothing of L_k at the crossing c_{k-1}. Ignoring edge weights and orientations, the graph G'_k is isomorphic to the Crowell graph associated to $L_k(c_{k-1}, ||)$ after deleting two loop edges when $k \geq 3$ and one loop edge when $k = 2$. The two graphs, however, carry different edge weights and orientations; see Figure 3.9.

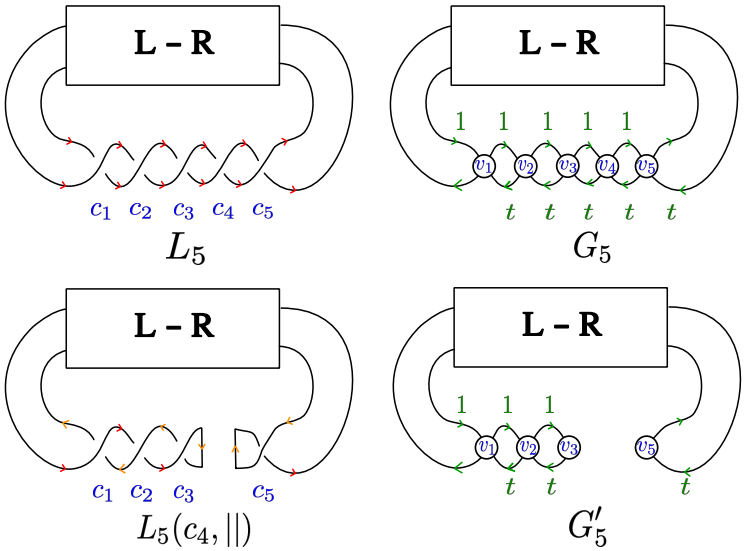


Figure 3.9: A twist region in a Crowell graph together with its unoriented smoothing.

The following lemma is the recursive relation at the heart of the stabilization phenomenon. It is an unoriented skein relation for the Crowell polynomial.

Lemma 3.10. *Let $P_k = P_{G_k,v_k}$ and $P'_k = P_{G'_k,v_k}$; i.e., the Crowell polynomials of the graphs G_k and G'_k rooted at v_k. Then*

$$P_k = tP_{k-1} + P'_k.$$

Proof. We decompose $\mathcal{T}(G_k, v_k)$ into two sets $\mathcal{T}_R$ and $\mathcal{T}_L$ of oriented spanning trees depending on the edge leading to v_{k-1}.

Let $\mathcal{T}_R$ be the set of oriented spanning trees containing the edge (v_k, v_{k-1}). By contracting this edge into the root vertex v_k and deleting the resulting loop edge, one obtains a one-to-one correspondence between this set and $\mathcal{T}(G_{k-1}, v_{k-1})$. The correspondence decreases the exponent of t in $W(T)$ by one; see Figure 3.10.

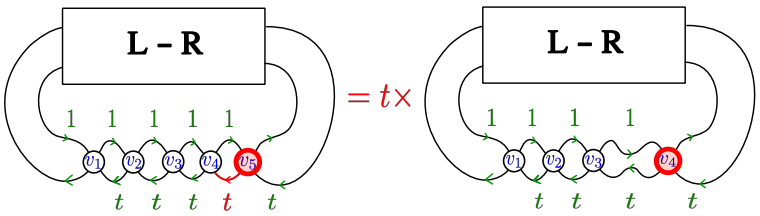


Figure 3.10: The set $\mathcal{T}_R$ is in bijection with $\mathcal{T}(G_{k-1}, v_{k-1})$. Trees containing the red edge belong to $\mathcal{T}_R$.

We write $\mathcal{T}_L$ for the set of oriented spanning trees containing the edge (v_{k-2}, v_{k-1}). These trees do not contain any other edge adjacent to v_{k-1}. As a result, $\mathcal{T}_L$ is in bijection with $\mathcal{T}(G'_k, v_k)$. This bijection preserves the weights; see Figure 3.11.

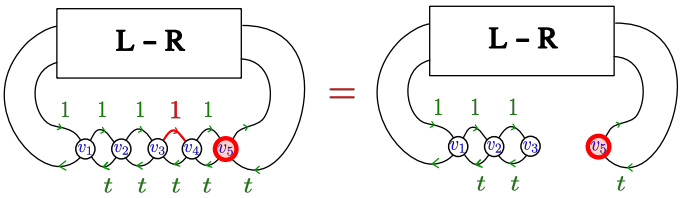


Figure 3.11: The set $\mathcal{T}_L$ is in bijection with $\mathcal{T}(G'_k, v_k)$. Trees containing the red edge are included in $\mathcal{T}_L$.

By the definition of P_k, we have

$$\begin{aligned} P_k &= \sum_{T \in \mathcal{T}(G_k, v_k)} W(T) \\ &= \sum_{T \in \mathcal{T}_R} W(T) + \sum_{T \in \mathcal{T}_L} W(T) \\ &= \sum_{T \in \mathcal{T}(G_{k-1}, v_{k-1})} tW(T) + \sum_{T \in \mathcal{T}(G'_k, v_k)} W(T) \\ &= tP_{k-1} + P'_k, \end{aligned}$$

as claimed. □

Before proceeding to the proof of Theorem 3.8, we introduce a couple of technical notions. Motivated by the symmetry of the coefficients of the Alexander polynomial, it is convenient to introduce the following one-sided version of trapezoidality.

Definition 3.11. A sequence of positive integers $(a_0, \ldots, a_m)$ is called *half-trapezoidal* if it is non-decreasing and, once it becomes constant, remains constant thereafter.

Remark 3.12. Note that, if $(a_0, \ldots, a_n)$ is a symmetric sequence of positive integers and $(a_0, \ldots, a_m)$ is half-trapezoidal for $m = \lfloor n/2 \rfloor$, then $(a_0, \ldots, a_n)$ is trapezoidal.

Finally, note that the Crowell polynomial need not have a nonzero constant term. Therefore, it is convenient to keep track of the smallest and largest exponents that occur in a polynomial.

Definition 3.13. Let $P(t) = \sum a_i t^i$ be a nonzero Laurent polynomial. We define

$$\operatorname{mindeg}(P) := \min\{\, i \mid a_i \neq 0 \,\}$$

$$\operatorname{maxdeg}(P) := \max\{\, i \mid a_i \neq 0 \,\}.$$

Remark 3.14. Whenever we refer to the sequence of coefficients of a Laurent polynomial, we mean the sequence

$$(a_{\text{mindeg}(P)}, a_{\text{mindeg}(P)+1}, \ldots, a_{\text{maxdeg}(P)}).$$

Notions such as trapezoidality can be extended to these sequences.

Proof of Theorem 3.8. Using the notation of Theorem 2.41 and Lemma 3.10, for $L = L_n$, we have

$$\Delta_L(-t) \doteq P_n.$$

Note that when L_n is alternating, L_1 is also alternating and hence $\Delta_{L_1}(-t) \doteq P_1$.

We prove Theorem 3.8 when n is odd. We are going to use the fact that $|L_n| = |L_1|$ for odd n. The same argument can be repeated for even n, but one needs to use L_2 instead of L_1. Hence, we will show that the sequence of coefficients of P_n is trapezoidal under the assumption

$$n - 2 \geq g(L) + \frac{|L|}{2}. \tag{3.2}$$

By applying Lemma 3.10 repeatedly, we deduce that

$$P_n = t^{n-1}P_1 + (P'_n + tP'_{n-1} + \cdots + t^{n-2}P'_2). \tag{3.3}$$

Note that the only difference between the graphs G'_k for different values of k is the length of the chain of vertices $v_1, \ldots, v_{k-2}$. Any tree in $\mathcal{T}(G'_k, v_k)$ must contain the edges $(v_1, v_2), \ldots, (v_{k-3}, v_{k-2})$.

Deleting the vertices $v_1, \ldots, v_{k-2}$ and the incident edges from each maximal rooted tree, and then replacing e_{li} with e_{io} whenever both e_{li} and e_{io} belong to the original tree, yields a one-to-one correspondence between $\mathcal{T}(G'_k, v_k)$ and $\mathcal{T}(G'_2, v_2)$; see Figure 3.12. This correspondence preserves weights, so $P'_k = P'_2$.

Combining this with Equation (3.3) gives

$$P_n = t^{n-1}P_1 + (1 + \cdots + t^{n-2})P'_2. \tag{3.4}$$

Let (a'_i) denote the sequence of coefficients of P'_2, and let (c'_i) denote the sequence of coefficients of

$$(1 + \cdots + t^{n-2})P'_2.$$

The first $n-1$ terms of (c'_i) are the cumulative sums of the sequence (a'_i); that is,

$$c'_{\text{mindeg}(P'_2)+j} = a'_{\text{mindeg}(P'_2)} + \cdots + a'_{\text{mindeg}(P'_2)+j}, \qquad j = 0, \ldots, n-2.$$

Note that the coefficients of Crowell polynomials are positive; see Remark 2.42. As a result, their cumulative sums form a half-trapezoidal sequence. Hence the first $n-1$ coefficients of $(1+\cdots+t^{n-2})P_2'$ form a half-trapezoidal sequence.

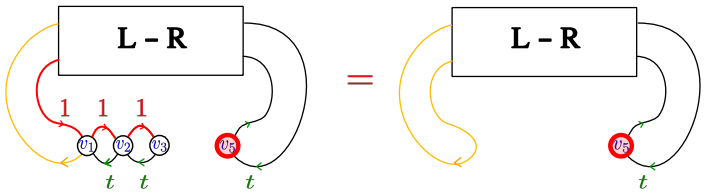


Figure 3.12: The bijection between $\mathcal{T}(G_k', v_k)$ and $\mathcal{T}(G_2', v_2)$. The oriented spanning trees in $\mathcal{T}(G_k', v_k)$ contain the red edges. The edges e_{lo} and e_{io} are shown in yellow.

We are going to show that the first $n-2$ coefficients of P_n agree with the first $n-2$ terms of (c_i') and hence are also half-trapezoidal. To do this, we separate the contribution of $(1+\cdots+t^{n-2})P_2'$ from the contribution of $t^{n-1}P_1$. More precisely, we prove that

$$n-1+\operatorname{mindeg}(P_1) = \operatorname{mindeg}(t^{n-1}P_1) = \operatorname{mindeg}(P_n)+n-1. \tag{3.5}$$

In fact, one can prove that $\operatorname{mindeg}(P_n) = \operatorname{mindeg}(P_1)$ for sufficiently large n by analysing spanning trees. However, to obtain the desired lower bound on n, we use a different argument. We show that

$$\operatorname{maxdeg}(P_n) = \operatorname{maxdeg}(t^{n-1}P_1) \tag{3.6}$$

$$\operatorname{Span}(P_n) - \operatorname{Span}(t^{n-1}P_1) = n-1. \tag{3.7}$$

Equations (3.6) and (3.7) imply Equation (3.5).

Note that any oriented spanning tree in $\mathcal{T}(G_2', v_2)$ is also a spanning tree in $\mathcal{T}(G_1, v_1)$, which means that all terms in P_2' also contribute to P_1. As a result,

$$\operatorname{maxdeg}(P_1) \geq \operatorname{maxdeg}(P_2'),$$

and hence $\operatorname{maxdeg}(t^{n-1}P_1) > \operatorname{maxdeg}((1+\cdots+t^{n-2})P_2')$.

Combining this with Equation (3.4) implies that the term of highest degree in P_n comes from $t^{n-1}P_1$. This proves Equation (3.6).

We compare the spans of P_1 and P_n by comparing the genera of L_1 and L_n, i.e.,

$$\operatorname{Span}(P_n) - \operatorname{Span}(P_1) = 2(g(L_n) - g(L_1)) + |L_n| - |L_1|.$$

Since n is odd, L_n and L_1 differ by a number of full twists and hence $|L_n| = |L_1|$. To compare their genera, we compare their Seifert graphs. The Seifert graph of L_n contains n parallel edges (equivalently, half-twisted bands) between two vertices (equivalently, Seifert cycles), whereas the Seifert graph of L_1 contains a single such edge; see Figure 3.13. Since L_n and L_1 are alternating, Theorem 2.24 implies that their genera are realised by Seifert's algorithm. It follows that

$$g(L_n) - g(L_1) = \frac{n-1}{2},$$

which proves Equation (3.7).

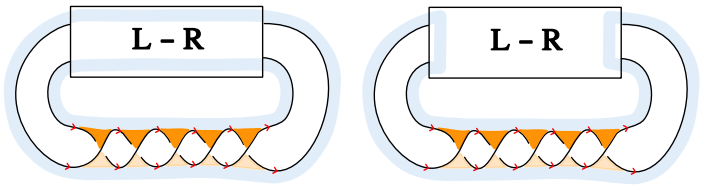


Figure 3.13: A twist region is a family of parallel bands between a pair of Seifert cycles (not necessarily distinct, as shown in the right figure).

We have proved that the first $n-2$ coefficients of P_n come from $(1 + \cdots + t^{n-2})P_2'$ and are therefore half-trapezoidal. Furthermore,

$$\operatorname{Span}(P_n) = 2g(L_n) + |L_n| - 1.$$

By assumption (3.2), all coefficients in the first half of the support of P_n are half-trapezoidal. By symmetry and Remark 3.12, it follows that the sequence of coefficients of P_n is trapezoidal. □

We conclude this section with a few corollaries of Theorem 3.8. The proof above immediately yields the following stronger statement.

Theorem 3.15. *Let L be an alternating link and let (a_i) be the sequence of absolute values of the coefficients of Δ_L. Then the sequence $(a_0, \ldots, a_{MT(L)-3})$ is half-trapezoidal.*

Remark 3.16. It is not always easy to verify the hypotheses of Theorem 3.8. However, the stronger condition

$$\operatorname{MT}(L) - 3 \geq \frac{c(L)}{2}$$

implies the hypotheses of Theorem 3.8. Since both $c(L)$ and $\operatorname{MT}(L)$ can be read directly from a reduced alternating diagram, this criterion is often easier to verify in practice. Roughly speaking, it means that more than half of the crossings of the diagram are concentrated in a single twist region.

3.3 Plumbings of Special Alternating Links

The goal of this section is to extend the following theorem.

Theorem 3.17 (Hafner–Mészáros–Vidinas [HMV24])**.** *The strong Fox conjecture holds for all special alternating links.*

In Subsection 2.1.3, we introduced diagrammatic Murasugi sums and Murasugi's decomposition of alternating links into special alternating links. In light of this decomposition, it is natural to investigate the behaviour of the Alexander polynomial under diagrammatic Murasugi sums, with the goal of extending Theorem 3.17 beyond the class of special alternating links.

However, as discussed in Subsection 2.1.2, no general formula is known for the Alexander polynomial of a Murasugi sum in terms of its summands. Nevertheless, for diagrammatic plumbings one can derive such a formula. This will allow us to prove the Fox trapezoidal conjecture for a new class of alternating links.

Theorem 3.18. *Let L be an alternating link with a Murasugi decomposition*

$$L = L_1 * \cdots * L_n,$$

where each L_i is special alternating and every $*$-product is a diagrammatic plumbing. Then the Fox trapezoidal conjecture holds for L.

We prove Theorem 3.18 by induction on the number of summands. The essential step is the case of a single diagrammatic plumbing, which we state below.

Theorem 3.19. *The Fox trapezoidal conjecture holds for any alternating link $L = L_1 * L_2$ that is a diagrammatic plumbing of two special alternating links.*

We start with a concrete diagrammatic description of a diagrammatic plumbing of two special alternating links. This description will be used repeatedly in the proof of Theorem 3.19.

Choose a reduced alternating diagram D of L. Since we only have two summands, there is only one type 2 Seifert cycle C, and the Seifert domains are the two disks it bounds in S^2. Using the notation of Remark 2.36 and Equation (2.1), assume first that $l = 1$. If $y_1 = 1$, the crossing connecting C to its exterior is reducible (also called nugatory), contradicting the assumption that D is reduced. Hence $y_1 \geq 2$. Similarly, $x_1 \geq 2$. By Equation (2.1), the gonality of the Murasugi sum is at least six, which contradicts the fact that we have a 4-gon Murasugi sum.

So $l \geq 2$, which, together with Equation (2.1), implies that $l = 2$ and $x_j = y_j = 1$ for $j \in \{1, 2\}$. Hence the diagram D looks locally as in Figure 3.14 near C.

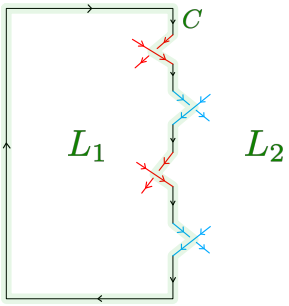


Figure 3.14: Plumbing in a Murasugi decomposition.

Note that Seifert cycles in a special diagram bound regions of the link diagram. As a result, one can describe the diagrammatic plumbing of two special alternating links as follows. Consider the special alternating links L_1 and L_2, and let R_i be a region of L_i for $i \in \{1, 2\}$ such that R_i has degree two as a vertex of the Tait graph of L_i (i.e., it has two crossings along its boundary). Isotope the diagram of L_1 so that R_1 becomes the unbounded region. Then place the diagram of L_1 inside the region R_2 of the diagram of L_2, identifying the boundaries of R_1 and R_2 according to Figure 3.14. See Figure 3.15 for an example.

Using this diagrammatic characterization, we can describe the effect of the operation on Tait graphs. Let B_{L_i} and W_{L_i} denote the black and white Tait graphs of the special alternating link L_i, respectively, for $i \in \{1, 2\}$. Without loss of generality, assume that R_1 and R_2 are both black regions in the respective colourings of the diagrams of L_1 and L_2. We also use R_1 and R_2 to denote the corresponding vertices of B_{L_1} and B_{L_2}, respectively. Furthermore, we can assume that the checkerboard colouring of L_2 extends to the colouring of L, and let B_L and W_L denote the black and white Tait graphs of L. These graphs can be described as follows. See Figures 3.16 and 3.17 for an example and some notation.

The graph B_L can be constructed from B_{L_2} by replacing the vertex R_2 with the graph W_{L_1}, as follows. In the link diagram of L_1, the region R_1 shares boundary arcs of the knot with exactly two white regions; denote these regions by w and w'. These correspond to vertices of W_{L_1}. Let b and b' be the two vertices of the black Tait graph B_{L_2} adjacent to R_2. To form B_L, we remove the vertex R_2 from B_{L_2} and add the two edges (b, w) and (b', w'); see Figure 3.17. We call this operation a *Tait blow-up of the region* R_2. The graph W_L is constructed in an analogous manner, with the roles of black and white regions reversed.

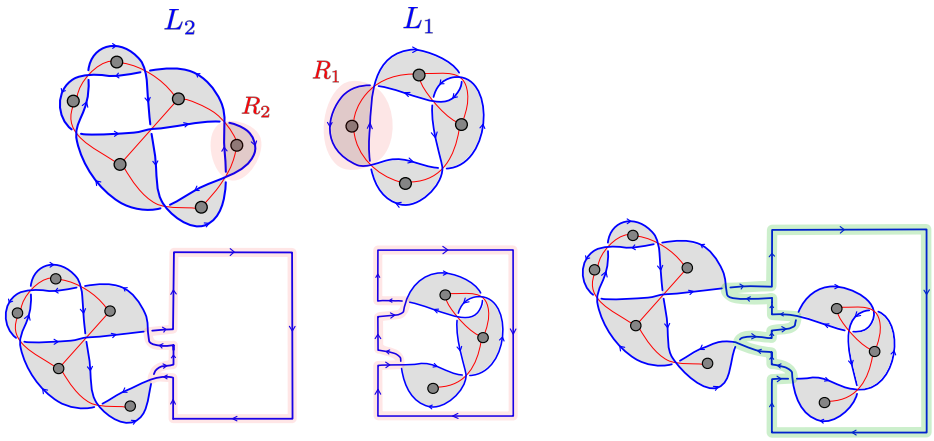


Figure 3.15: Example of diagrammatic plumbing. The links from left to right are L_2 and L_1. The red highlighted regions are R_2 and R_1. The green Seifert cycle is C, whose neighborhood is shown in Figure 3.14.

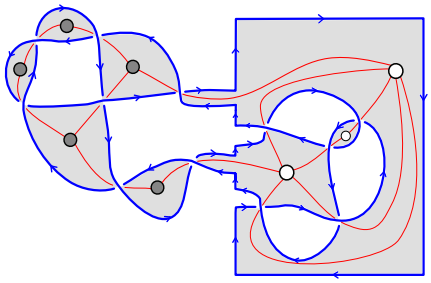


Figure 3.16: Example of Tait blow-up (Tait graph of a diagrammatic plumbing).

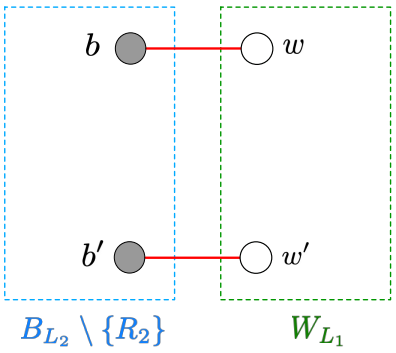


Figure 3.17: Schematic of the Tait graph of a plumbing.

We will use the notation introduced above throughout the remainder of this subsection. The technical crux of this subsection is an explicit formula for the Alexander polynomial of a diagrammatic plumbing. This is made precise in Proposition 3.20, which expresses $\widetilde{\Delta}_L$ in terms of the symmetrized Alexander polynomials of L_1, L_2, and

two auxiliary links $\widetilde{L}_1$ and $\widetilde{L}_2$ obtained by local modifications near the gluing region. Before stating this proposition, we introduce the notation used in the formula.

Let $\widetilde{L}_1$ (resp. $\widetilde{L}_2$) be the link obtained from L_1 (resp. L_2) by deleting the region R_1 (resp. R_2) and smoothing the two crossings along its boundary. See Figure 3.18 for an example. Note that this operation preserves the class of special alternating links, since deleting a region does not affect the orientability of the associated Tait surface (equivalently, deleting a vertex from the Tait graph does not change its bipartiteness).

Proposition 3.20. *Let the link $L = L_1 * L_2$ be the diagrammatic plumbing of the links L_1 and L_2. Then*

$$\widetilde{\Delta}_L = \widetilde{\Delta}_{L_1} \cdot \widetilde{\Delta}_{L_2} + \widetilde{\Delta}_{\widetilde{L}_1} \cdot \widetilde{\Delta}_{\widetilde{L}_2}.$$

The proof of Proposition 3.20 is postponed until after the proof of Theorem 3.19.

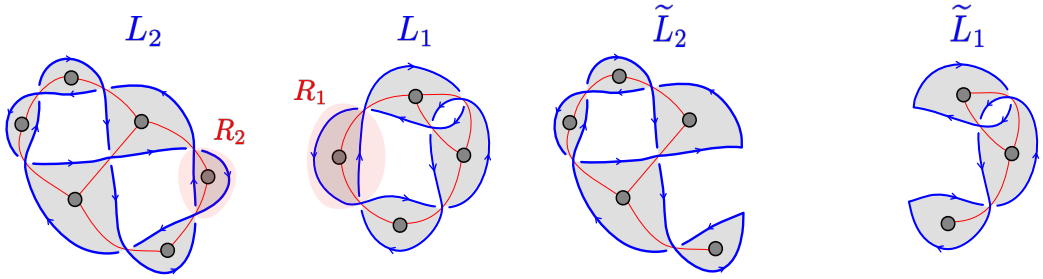


Figure 3.18: The auxiliary links $\widetilde{L}_1$ and $\widetilde{L}_2$. In this example, L_1 and L_2 are special alternating, and hence both auxiliary links are also special alternating.

Remark 3.21. One may notice that Proposition 3.20 does not assume that the summands are special alternating, or even alternating. Indeed, the notion of diagrammatic plumbing makes sense for arbitrary link diagrams. More generally, one can apply the Murasugi decomposition algorithm to any diagram, and a diagrammatic plumbing can be viewed as the reverse of a Murasugi decomposition over a type 2 Seifert cycle, whose local form is shown in Figure 3.14.

Furthermore, the description of diagrammatic plumbing in terms of Tait graphs given above extends readily to this more general setting. Consequently, the proof of Proposition 3.20 does not rely on any alternating or special alternating assumptions. These hypotheses only enter later.

The proof of Theorem 3.19 combines Proposition 3.20 with Theorem 3.17. Before carrying this out, we recall a basic result of Murasugi concerning the behaviour of trapezoidal polynomials under addition and multiplication. For convenience, we introduce the following terminology.

Notation 3.22. A Laurent polynomial $f(t) = \sum_{i=0}^{n} c_i t^i$ is called *trapezoidal* if it is symmetric, has nonnegative coefficients, and the sequence $(c_0, \ldots, c_n)$ is trapezoidal.

Proposition 3.23 (Murasugi [Mur85, Proposition 2.1])**.** *Let*

$$f(t) = \sum_{i=0}^{n} a_i t^i \qquad \text{and} \qquad g(t) = \sum_{j=0}^{m} b_j t^j$$

be trapezoidal polynomials.

1. *the product $f \cdot g$ is trapezoidal;*
2. *if $n = m + 2k$, $0 \leq k \leq 1$, then the sum $f + t^k \cdot g$ is trapezoidal.*

We now proceed with the proof of Theorem 3.19.

Proof of Theorem 3.19. We begin by recalling the formula of Proposition 3.20:

$$\tilde{\Delta}_L = \tilde{\Delta}_{L_1} \cdot \tilde{\Delta}_{L_2} + \tilde{\Delta}_{\widetilde{L}_1} \cdot \tilde{\Delta}_{\widetilde{L}_2}. \tag{3.8}$$

As one sees from the proof of Proposition 3.20, this formula comes from partitioning the Kauffman states into two classes. We claim that we can renormalize this equation to the following

$$\Delta_L = \Delta_{L_1} \cdot \Delta_{L_2} + t \cdot \Delta_{\widetilde{L}_1} \cdot \Delta_{\widetilde{L}_2}. \tag{3.9}$$

For this, we just need to compare the degrees of $\tilde{\Delta}_{L_1} \cdot \tilde{\Delta}_{L_2}$ and $\tilde{\Delta}_{\widetilde{L}_1} \cdot \tilde{\Delta}_{\widetilde{L}_2}$. More precisely, it suffices to show that

$$\deg\big(\tilde{\Delta}_{L_1} \cdot \tilde{\Delta}_{L_2}\big) = \deg\big(\tilde{\Delta}_{\widetilde{L}_1} \cdot \tilde{\Delta}_{\widetilde{L}_2}\big) + 1. \tag{3.10}$$

We determine the degrees of these polynomials using the corresponding Tait graphs. Note that the black Tait graph B_{L_1} (resp. B_{L_2}) gives rise, via the median construction, to a minimal genus Seifert surface for L_1 (resp. L_2). Consequently, the degrees of the symmetrized Alexander polynomials satisfy

$$\deg\big(\tilde{\Delta}_{L_1} \cdot \tilde{\Delta}_{L_2}\big) = \frac{1}{2}\big(\operatorname{rank}(H_1(B_{L_1})) + \operatorname{rank}(H_1(B_{L_2}))\big).$$

Similarly,

$$\deg\big(\tilde{\Delta}_{\widetilde{L}_1} \cdot \tilde{\Delta}_{\widetilde{L}_2}\big) = \frac{1}{2}\big(\operatorname{rank}(H_1(B_{\widetilde{L}_1})) + \operatorname{rank}(H_1(B_{\widetilde{L}_2}))\big).$$

For a graph B, we have

$$\operatorname{rank}(H_1(B)) = \#(\text{edges}) - \#(\text{vertices}) + \#(\text{components}).$$

This formula implies that

$$\operatorname{rank}\big(H_1(B_{\widetilde{L}_2})\big) = \operatorname{rank}\big(H_1(B_{L_2})\big) - 1.$$

Indeed, deleting the region R_2 does not increase the number of components of the black Tait graph, since the diagram is reduced and the degree of R_2 is two. Similarly, we have

$$\operatorname{rank}\big(H_1(B_{\widetilde{L}_1})\big) = \operatorname{rank}\big(H_1(B_{L_1})\big) - 1.$$

This proves equation (3.10). Equivalently, we have

$$\deg\big(\Delta_{L_1} \cdot \Delta_{L_2}\big) = \deg\big(\Delta_{\widetilde{L}_1} \cdot \Delta_{\widetilde{L}_2}\big) + 2.$$

We now use Proposition 3.23. Note that all links appearing in (3.9) are alternating. It is a classical fact that for an alternating link the coefficients of the Alexander polynomial have alternating signs. Consequently, the absolute values of the coefficients of $\Delta_L(t)$ form a trapezoidal sequence if and only if $\Delta_L(-t)$ is a trapezoidal polynomial. Furthermore, by Kauffman's clock theorem, for an alternating diagram the Kauffman states contributing to a fixed power of t have the same sign. Hence, applying the same grouping of Kauffman states as in the proof of Proposition 3.20, and replacing each contribution by its absolute value, we obtain

$$\Delta_L(-t) = \Delta_{L_1}(-t) \cdot \Delta_{L_2}(-t) + t \cdot \Delta_{\widetilde{L}_1}(-t) \cdot \Delta_{\widetilde{L}_2}(-t). \tag{3.11}$$

By hypothesis, L_1 and L_2 are special alternating, and, as noted above, the auxiliary links $\widetilde{L}_1$ and $\widetilde{L}_2$ are special alternating as well. Hence, by Theorem 3.17, the four polynomials

$$\Delta_{L_1}(-t), \quad \Delta_{L_2}(-t), \quad \Delta_{\widetilde{L}_1}(-t), \quad \Delta_{\widetilde{L}_2}(-t)$$

are trapezoidal in the sense defined above. Applying Proposition 3.23, we deduce that

$$\Delta_{L_1}(-t)\Delta_{L_2}(-t) \quad \text{and} \quad \Delta_{\widetilde{L}_1}(-t)\Delta_{\widetilde{L}_2}(-t)$$

are trapezoidal polynomials. The degree comparison above gives precisely the degree condition required for the second part of Proposition 3.23. Applying that proposition to (3.11), we conclude that $\Delta_L(-t)$ is trapezoidal. Therefore, the absolute values of the coefficients of $\Delta_L(t)$ form a trapezoidal sequence. This proves the theorem. □

We now use Theorem 3.19 to prove Theorem 3.18.

Proof. The proof proceeds by considering the innermost Seifert cycle C_m of type 2. The interior disk C_m^+ bounded by C_m is also a Seifert domain and contains a summand of the diagrammatic plumbing. Without loss of generality, assume this summand is L_n and define $L' = L_1 * \cdots * L_{n-1}$.

Note that a diagrammatic Murasugi decomposition always happens over a tree T. This is due to the inclusion hierarchy of type 2 Seifert cycles, which induces a partial order on the summands. In this Murasugi sum tree, L_n corresponds to a leaf vertex v. Without loss of generality, assume that the neighbour of v corresponds to L_{n-1}.

We use induction on the number of summands in the diagrammatic plumbing. Applying Proposition 3.20 to $L = L_n * L'$ gives

$$\widetilde{\Delta}_L = \widetilde{\Delta}_{L_n} \cdot \widetilde{\Delta}_{L'} + \widetilde{\Delta}_{\widetilde{L}_n} \cdot \widetilde{\Delta}_{\widetilde{L}'}.$$

The trapezoidal conjecture holds for L_n and $\widetilde{L}_n$, since both are special alternating. It also holds for L' by the induction hypothesis. Furthermore, $\widetilde{L}' = L_1 * \cdots * \widetilde{L}_{n-1}$. Indeed, the modification defining $\widetilde{L}'$ is local to the plumbing region between L_n and L_{n-1}. Hence the trapezoidal conjecture holds for $\widetilde{L}'$ as well. At this point, the remainder of the argument is identical to the proof of Theorem 3.19. □

We now prove Proposition 3.20 by analyzing the Kauffman state-sum for the Alexander polynomial of a diagrammatic plumbing.

Proof of Proposition 3.20. We work with the description of the black and white Tait graphs associated to the diagrammatic plumbing $L = L_1 * L_2$ given in Figures 3.16 and 3.17. Throughout the proof, we consider decorated link diagrams and the corresponding rooted Tait graphs. For each link diagram under consideration (namely L, L_1, L_2, and later $\widetilde{L}_1$ and $\widetilde{L}_2$), we always choose the unbounded region as one of the decorated regions and denote the corresponding root vertex of the Tait graph by y.

The choice of the second decorated region (and hence the second root vertex) depends on the link under consideration and on the case of the proof. For L, L_2, and $\widetilde{L}_2$, this root vertex will be the vertex b (as illustrated in Figure 3.17). For L_1, the second root vertex will be either w or w', depending on the case under consideration. The corresponding choice for $\widetilde{L}_1$ will be specified later in the proof, when it becomes relevant. The red stars in Figure 3.19 (and in the other figures appearing in this proof) indicate these root vertices.

Figure 3.19 also shows a pair of dual spanning trees $(T, T') \in \mathcal{T}_{B_L} \times \mathcal{T}_{W_L}$, coloured red and yellow, respectively. Note that the crossings in Figure 3.19 both belong to L_2; cf. Figure 3.16.

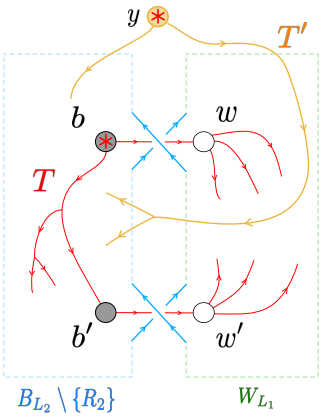


Figure 3.19: A pair of dual trees in $\mathcal{T}_{B_L} \times \mathcal{T}_{W_L}$.

Since T is a spanning tree rooted at b, it contains unique directed paths β from b to b' and ω from w to w'. Depending on β and ω, there are three distinct possibilities for a pair of dual spanning trees $(T, T') \in \mathcal{T}_{B_L} \times \mathcal{T}_{W_L}$; see Figure 3.20:

Case (i): β passes through w and then w'.

Case (ii): $\beta \subseteq B_{L_2}$ and $\omega \subseteq W_{L_1}$. This case has two subcases:

Case (ii).1: T contains the edge (b, w).

Case (ii).2: T contains the edge (b', w').

Case (iii): $\beta \subseteq B_{L_2}$ and T contains both (b, w) and (b', w').

For $* \in \{(i), (ii), (ii).1, (ii).2, (iii)\}$, we denote the set of (T, T') in Case $*$ by $\mathcal{D}^*_{(B_L, W_L)}$. By analyzing the possible options for the paths β and ω, we obtain the partition

$$\mathcal{D}_{(B_L, W_L)} = \mathcal{D}^{(i)}_{(B_L, W_L)} \sqcup \mathcal{D}^{(ii)}_{(B_L, W_L)} \sqcup \mathcal{D}^{(iii)}_{(B_L, W_L)}.$$

We now look at these cases one-by-one. In each case, we decompose the dual spanning trees T and T' into dual spanning trees of the Tait graphs of L_1 and L_2, and relate their weights to the weights of Kauffman states in $\mathcal{D}_{(B_{L_1}, W_{L_1})}$ and $\mathcal{D}_{(B_{L_2}, W_{L_2})}$.

Case (i): The edges of $T \cap W_{L_1}$ form a spanning tree T_1 of W_{L_1} rooted at w. The rest of the edges form a spanning tree T_2 of B_{L_2} after replacing the edges (b, w) and (w', b') with the edges (b, R_2) and (R_2, b'). There is an analogous decomposition of the planar dual T' of T into spanning trees T'_1 of B_{L_1} and T'_2 of W_{L_2}, with one small modification. Consider the vertex x of T' that lies between the summands L_1 and L_2, and between the two crossings between the two summands. The outgoing edges of x become the outgoing edges of y in T'_1. See Figure 3.21 for an illustration of the decomposition.

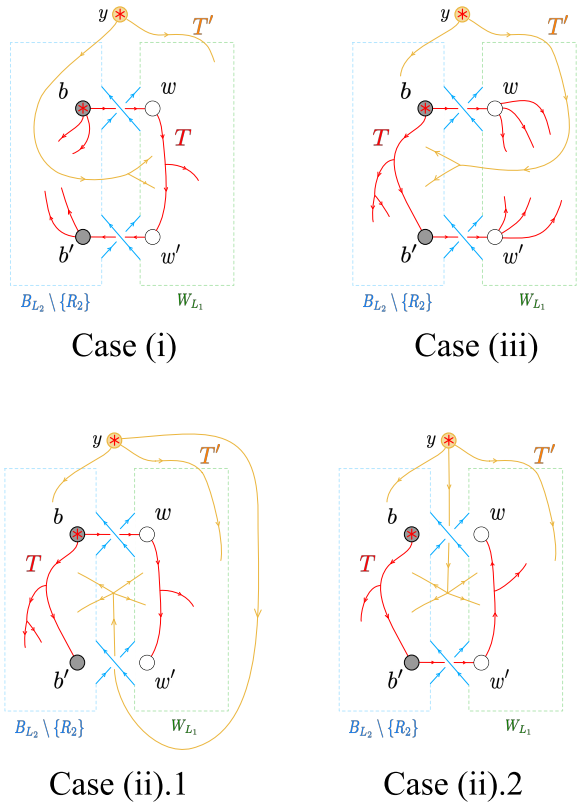


Figure 3.20: The possible cases for the paths β from b to b' and ω from w to w' in a rooted, directed spanning tree T of B_L, together with the dual spanning tree T' of W_L.

This decomposition gives us the map

$$\begin{aligned} \psi\colon \mathcal{D}^{(i)}_{(B_L,W_L)} &\hookrightarrow \mathcal{D}_{(B_{L_2},W_{L_2})} \times \mathcal{D}_{(W_{L_1},B_{L_1})}. \\ (T,T') &\mapsto \big((T_2,T_2'),(T_1,T_1')\big) \end{aligned} \tag{3.12}$$

If T_2 contains the edges (b,R_2) and (R_2,b'), then the construction can be reversed, so ψ is injective. Let

$$\mathcal{D}^{(i)}_{(B_{L_2},W_{L_2})} := \{(T_2,T_2') \in \mathcal{D}_{(B_{L_2},W_{L_2})} \,:\, (b,R_2),(R_2,b') \in T_2\}.$$

Then

$$\psi\colon \mathcal{D}^{(i)}_{(B_L,W_L)} \xrightarrow{\sim} \mathcal{D}^{(i)}_{(B_{L_2},W_{L_2})} \times \mathcal{D}_{(W_{L_1},B_{L_1})} \tag{3.13}$$

is a bijection. This map just changes the ends of some of the edges and does not change their position with respect to crossings, hence it preserves the weights of states: If $\psi(T,T') = \big((T_2,T_2'),(T_1,T_1')\big)$, then

$$w(T) = w(T_2)\cdot w(T_1) \;\text{ and }\; w(T') = w(T_2')\cdot w(T_1'). \tag{3.14}$$

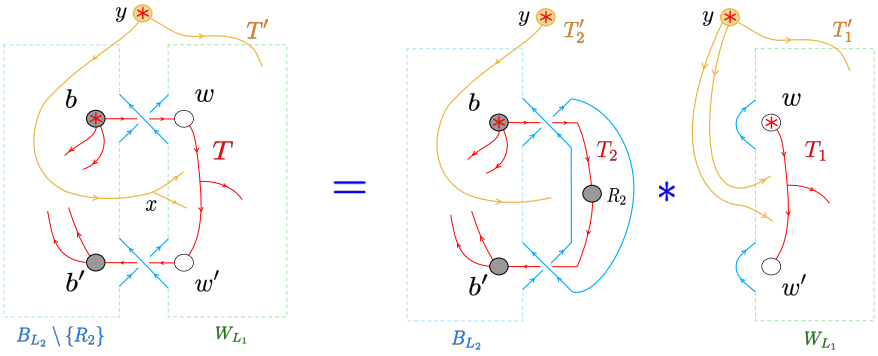


Figure 3.21: Decomposition of spanning trees in Case (i).

Case (ii): We decompose dual pairs of spanning trees in $\mathcal{D}^{(ii).1}_{B_L,W_L}$ as in Figures 3.22 and those in $\mathcal{D}^{(ii).2}_{B_L,W_L}$ as in Figure 3.23. Similarly to Case (i), we define the following sets of dual pairs of spanning trees:

$$\mathcal{D}^{(ii).1}_{(B_{L_2},W_{L_2})} := \{\, (T_2, T_2') \in \mathcal{D}_{(B_{L_2},W_{L_2})} \,:\, (b, R_2) \in T_2 \ ,\ (R_2, b') \notin T_2 \,\},$$

$$\mathcal{D}^{(ii).2}_{(B_{L_2},W_{L_2})} := \{\, (T_2, T_2') \in \mathcal{D}_{(B_{L_2},W_{L_2})} \,:\, (R_2, b') \in T_2 \ ,\ (b, R_2) \notin T_2 \,\}.$$

Decomposing dual pairs of spanning trees as

$$(T, T') \mapsto \big((T_2, T_2'), (T_1, T_1')\big)$$

gives bijections

$$\mathcal{D}^{(ii).1}_{(B_L,W_L)} \xrightarrow{\sim} \mathcal{D}^{(ii).1}_{(B_{L_2},W_{L_2})} \times \mathcal{D}_{(W_{L_1},B_{L_1})}, \tag{3.15}$$

$$\mathcal{D}^{(ii).2}_{(B_L,W_L)} \xrightarrow{\sim} \mathcal{D}^{(ii).2}_{(B_{L_2},W_{L_2})} \times \mathcal{D}_{(W_{L_1},B_{L_1})}. \tag{3.16}$$

These maps preserve the weights of the Kauffman states in the sense of equation (3.14).

A spanning tree of B_{L_2} rooted at b contains unique directed paths from b to b' and from b to R_2. There are three possibilities for these paths, giving the partition

$$\mathcal{D}_{(B_{L_2},W_{L_2})} = \mathcal{D}^{(i)}_{(B_{L_2},W_{L_2})} \sqcup \mathcal{D}^{(ii).1}_{(B_{L_2},W_{L_2})} \sqcup \mathcal{D}^{(ii).2}_{(B_{L_2},W_{L_2})}.$$

We can now compute the weighted sum of all the dual pairs in Cases (i) and (ii). Using the weight-preserving maps in equations (3.13), (3.15) and (3.16), we can write

$$\sum_{\mathcal{D}^{(i)}_{(B_L,W_L)} \,\sqcup\, \mathcal{D}^{(ii).1}_{(B_L,W_L)} \,\sqcup\, \mathcal{D}^{(ii).2}_{(B_L,W_L)}} w(T) \cdot w(T') =$$

$$\left(\sum_{\mathcal{D}^{(i)}_{(B_{L_2},W_{L_2})} \sqcup \mathcal{D}^{(ii).1}_{(B_{L_2},W_{L_2})} \sqcup \mathcal{D}^{(ii).2}_{(B_{L_2},W_{L_2})}} w(T_2)\cdot w(T_2') \right) \cdot$$

$$\left(\sum_{\mathcal{D}_{(W_{L_1},B_{L_1})}} w(T_1)\cdot w(T_1') \right) =$$

$$\left(\sum_{\mathcal{D}_{(B_{L_2},W_{L_2})}} w(T_2)\cdot w(T_2') \right) \left(\sum_{\mathcal{D}_{(W_{L_1},B_{L_1})}} w(T_1)\cdot w(T_1') \right) = \widetilde{\Delta}_{L_2}\cdot\widetilde{\Delta}_{L_1}.$$

The last equality comes from the Kauffman state-sum formula. Note that, although the root of the spanning trees of W_{L_1} is w in Cases (i) and (ii).1 and is w' in Case (ii).2, this does not affect the result of the state-sum.

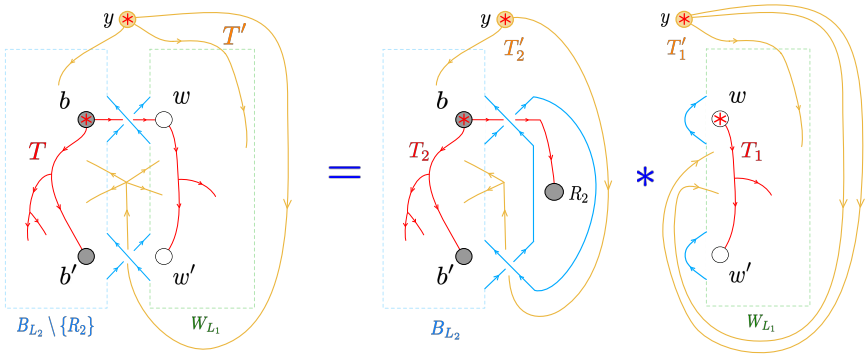


Figure 3.22: Decomposition of spanning trees in Case (ii).1.

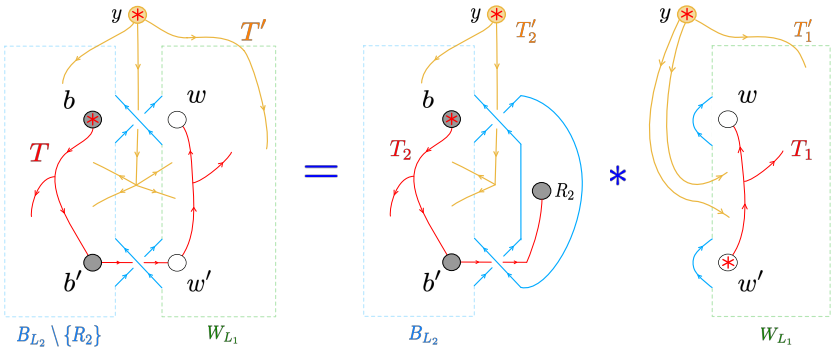


Figure 3.23: Decomposition of spanning trees in Case (ii).2.

Case (iii): It remains to show that the weighted sum of all Kauffman states in Case (iii) equals the product of the Alexander polynomials of the two auxiliary links.

Equivalently, our goal is to prove the formula

$$\sum_{\mathcal{D}^{(iii)}_{(B_L,W_L)}} w(T)\cdot w(T') \;=\; \widetilde{\Delta}_{\widetilde{L}_1}\cdot\widetilde{\Delta}_{\widetilde{L}_2}.$$

Recall that $\widetilde{L}_2$ was defined above as the result of deleting the region R_2 from L_2 and smoothing the two crossings on its boundary (see Figure 3.18); denote its black Tait graph by $B_{\widetilde{L}_2}$. Note that $B_{\widetilde{L}_2}$ sits inside B_L. In addition, we define $\widehat{L}_1$ to be the alternating link whose white Tait graph $W_{\widehat{L}_1}$ is obtained from W_{L_1} by contracting the two vertices w and w' to a single vertex $\{w,w'\}$. The link $\widehat{L}_1$ is constructed from this graph via the median construction.

For a spanning tree T of B_L occurring in Case (iii), set

$$\widetilde{T}_2 := T\cap B_{\widetilde{L}_2}.$$

By the definition of Case (iii), $\widetilde{T}_2$ is a rooted spanning tree of $B_{\widetilde{L}_2}$. The rest of the edges of T form a spanning forest of W_{L_1} with two components rooted at w and w'. By contracting w and w', this forest becomes a spanning tree $\widehat{T}_1$ of $W_{\widehat{L}_1}$ rooted at $\{w,w'\}$. This decomposition also extends to the planar duals T', $\widehat{T}'_1$, and $\widetilde{T}'_2$, with a small modification. The outgoing edges of x turn into the outgoing edges of y in $\widetilde{T}'_2$. See Figure 3.24.

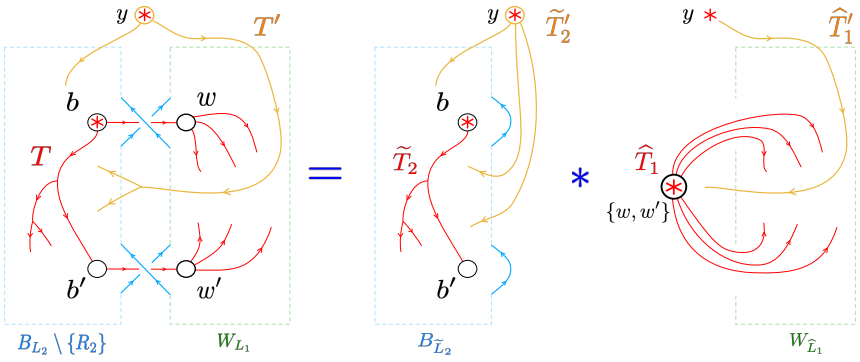


Figure 3.24: Decomposition of spanning trees in Case (iii).

Similarly to the previous cases, this construction is reversible and hence gives a bijection

$$\begin{aligned}\mathcal{D}^{(iii)}_{(B_L,W_L)} &\xrightarrow{\sim} \mathcal{D}_{(B_{\widetilde{L}_2},W_{\widetilde{L}_2})}\times\mathcal{D}_{(W_{\widehat{L}_1},B_{\widehat{L}_1})}.\\ (T,T') &\mapsto \big((\widetilde{T}_2,\widetilde{T}'_2),(\widehat{T}_1,\widehat{T}'_1)\big)\end{aligned} \tag{3.17}$$

This map is also weight-preserving in the sense of equation 3.14. This part of the argument is a bit different from the previous cases since this decomposition deletes

the edges (b, w) and (b', w'). Both of these edges have weight 1, and hence this does not affect the overall weight of the Kauffman states. This decomposition also deletes two of the edges of T'. These are the edges that pass over the other two crossings on the boundary of the type 2 Seifert cycle C. One can see that these two edges also have weight 1.

Using this weight-preservation, we can finally compute the sum of the weights of all the dual pairs of spanning trees in Case (iii) as follows:

$$\sum_{\mathcal{D}^{(iii)}_{(B_L, W_L)}} w(T) \cdot w(T') =$$

$$\left(\sum_{\mathcal{D}_{\left(B_{\widetilde{L}_2}, W_{\widetilde{L}_2}\right)}} w\big(\widetilde{T}_2\big) \cdot w\big(\widetilde{T}'_2\big) \right) \left(\sum_{\mathcal{D}_{\left(W_{\widehat{L}_1}, B_{\widehat{L}_1}\right)}} w\big(\widehat{T}_1\big) \cdot w\big(\widehat{T}'_1\big) \right) = \widetilde{\Delta}_{\widetilde{L}_2} \cdot \widetilde{\Delta}_{\widehat{L}_1}.$$

It remains to identify $\widehat{L}_1$ with the auxiliary link $\widetilde{L}_1$ defined earlier. Recall that $\widetilde{L}_1$ is obtained from L_1 by deleting the region R_1 and smoothing the two crossings on its boundary. These two crossings correspond to the two edges between w and w' in W_{L_1}. On the level of white Tait graphs, deleting R_1 first contracts w and w' to a single vertex $\{w, w'\}$. Under this contraction the two edges joining w and w' become loop edges. Smoothing the two crossings then removes these loop edges. Thus $W_{\widetilde{L}_1}$ and $W_{\widehat{L}_1}$ differ only by the presence of these two loop edges, which correspond to reducible crossings and do not change the isotopy type of the link. Hence $\widehat{L}_1$ is isotopic to $\widetilde{L}_1$, and in particular

$$\widetilde{\Delta}_{\widehat{L}_1} = \widetilde{\Delta}_{\widetilde{L}_1}.$$

Substituting this identification into the computation above shows that the contribution of Case (iii) equals $\widetilde{\Delta}_{\widetilde{L}_2} \cdot \widetilde{\Delta}_{\widetilde{L}_1}$. Combining this with the contributions from Cases (i) and (ii) yields the Alexander-plumbing formula

$$\widetilde{\Delta}_L = \widetilde{\Delta}_{L_1} \cdot \widetilde{\Delta}_{L_2} + \widetilde{\Delta}_{\widetilde{L}_1} \cdot \widetilde{\Delta}_{\widetilde{L}_2},$$

which proves Proposition 3.20. □

We conclude this section with a further extension of Theorem 3.18. The proof does not fundamentally rely on the Murasugi sums being plumbings, and can be adapted to a more general class of diagrammatic Murasugi sums. Using the notion of length introduced in Definition 2.37, one can formulate the following generalization.

Theorem 3.24. *Let L be an alternating link with a Murasugi decomposition*

$$L = L_1 * \cdots * L_n,$$

where each L_i is special alternating and every $*$-product in the decomposition has length less than three. Then the Fox trapezoidal conjecture holds for L.

Remark 3.25. The condition that a diagrammatic Murasugi sum has length less than three means that we allow x_1, x_2, y_1, and y_2 to be arbitrary positive integers, but there are no x_i or y_i with $i \geq 3$, using the notation of Figure 2.16.

The main modification concerns the definition of the auxiliary links $\widetilde{L}_1$ and $\widetilde{L}_2$ appearing in Proposition 3.20. In this more general setting, one constructs $\widetilde{L}_1$ (resp. $\widetilde{L}_2$) by gluing an unknotted band to L_1 (resp. L_2) inside the region R_1 (resp. R_2) and separating the two groups of twists along its boundary. Gluing these bands changes the Tait graph by a planar contraction of two vertices (or, dually, by splitting a vertex into two), which is precisely the operation needed to extend the proof of Theorem 3.19. See Figures 3.25, 3.26, and 3.27 for examples.

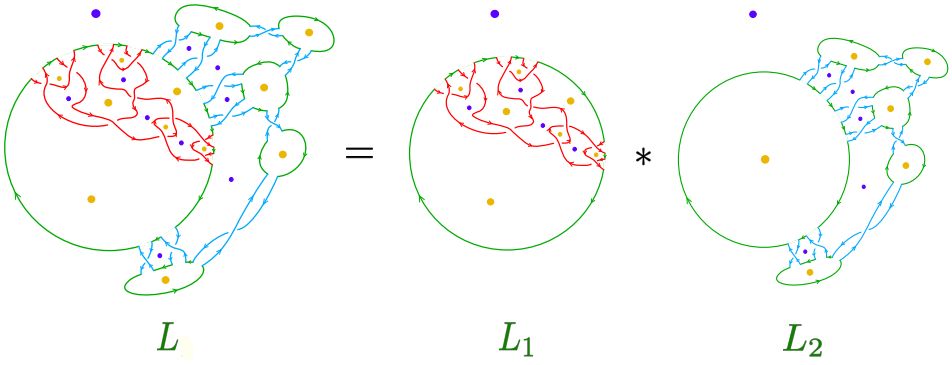


Figure 3.25: A diagrammatic Murasugi sum of length two.

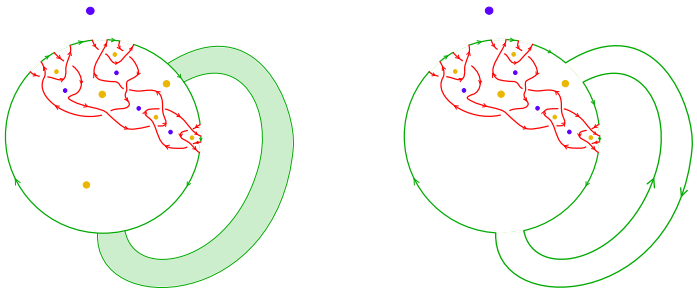


Figure 3.26: Construction of $\widetilde{L}_1$ in the example of Figure 3.25.

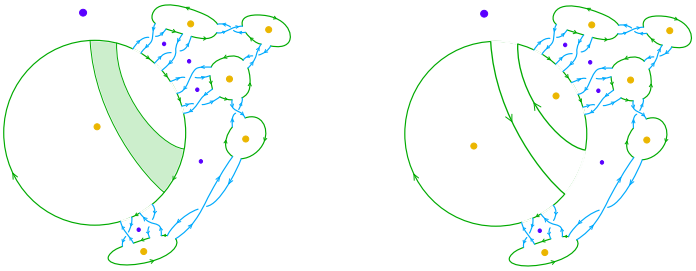


Figure 3.27: Construction of $\widetilde{L}_2$ in the example of Figure 3.25.

Note that this construction recovers the original definition in the plumbing case. Indeed, gluing a band to a bigon region creates two kinks that can be removed by Reidemeister 1 moves, and is therefore equivalent to smoothing the two crossings along the boundary.

3.4 Alternating Three-Braids

The goal of this subsection is to extend some of the ideas of Hafner, Mészáros, and Vidinas [HMV24] from special alternating links to alternating 3-braids. The class of alternating 3-braids is the simplest family of non-special alternating links. Indeed, the closure of an alternating 3-braid can be viewed as a diagrammatic Murasugi sum of two alternating torus links. On the other hand, the gonality of this Murasugi sum can be arbitrarily large. Thus alternating 3-braids already exhibit the full range of complexity arising from the gonality of a diagrammatic Murasugi sum.

As we will see, the resulting methods do not directly imply the Fox trapezoidal conjecture. Nevertheless, they lead to a collection of coefficient inequalities, the strongest of which is the following theorem.

Theorem 3.26. *Let L be the closure of the alternating 3-braid*

$$\sigma_1^{p_1}\sigma_2^{-q_1}\cdots\sigma_1^{p_n}\sigma_2^{-q_n},$$

and let (a_k) denote the sequence of absolute values of the coefficients of Δ_L. Then

$$a_{k+1} > \frac{a_k}{\min(n, k+1)}, \qquad k \in \{0, \dots, \max(P-2, Q-2)\},$$

where $P = p_1 + \cdots + p_n$ and $Q = q_1 + \cdots + q_n$.

We begin by recalling some of the key ideas of Hafner, Mészáros, and Vidinas [HMV24]. Let L be a special alternating link with diagram D, and let $B = B_L$

denote its bipartite Tait graph. Recall that the vertices of B correspond to the black regions of the diagram. We will use the same notation for a vertex of B and the corresponding black region of D. We start with a useful property of the Crowell graph G_D.

Proposition 3.27. *The edges of G_D are oriented so that the boundary of every black region $v \in B$ is directed clockwise. Furthermore, there is a 2-colouring*

$$V(B) = S \sqcup R$$

such that all edges on the boundary of a region in S have weight t, while all edges on the boundary of a region in R have weight 1.

Proof. The claim follows from a local analysis of the orientations and weights of the Crowell graph near a vertex of the bipartite Tait graph B; see Figures 3.28 and 3.29. □

Following Hafner, Mészáros, and Vidinas [HMV24], we define a multivariable refinement of the Alexander polynomial as follows.

Construction 3.28. Let $S = \{v_1, \ldots, v_n\}$ and $R = \{v_{n+1}, \ldots, v_{n+m}\}$. Let $\widetilde{G}_D$ be the oriented rooted graph obtained from the Crowell graph G_D by replacing the edge weights as follows. For each $i \in \{1, \ldots, n+m\}$, assign weight x_i to all edges along the boundary of the region corresponding to v_i.

Recall that $\mathcal{T}(\widetilde{G}_D, c_1)$ denotes the set of maximal oriented trees in $\widetilde{G}_D$ rooted at c_1. We define

$$P_{\widetilde{G}_D, c_1}(x_1, \ldots, x_{n+m}) = \sum_{T \in \mathcal{T}(\widetilde{G}_D, c_1)} W(T),$$

where $W(T)$ is the product of the weights of the edges of T; see Figure 3.30.

Note that $P_{\widetilde{G}_D, c_1}$ is clearly a refinement of the Crowell polynomial, since

$$P_{\widetilde{G}_D, c_1}(t, \ldots, t, 1, \ldots, 1) = P_{G_D, c_1}(t) = \Delta_L(-t).$$

The refinement $P_{\widetilde{G}_D, c_1}$ enjoys two remarkable properties. First, all of its nonzero coefficients are equal to 1. Second, its support satisfies a property known as *M-convexity*. Recall that the support $\mathrm{Supp}(P)$ of a multivariable polynomial P is the set of exponent vectors corresponding to monomials with nonzero coefficients.

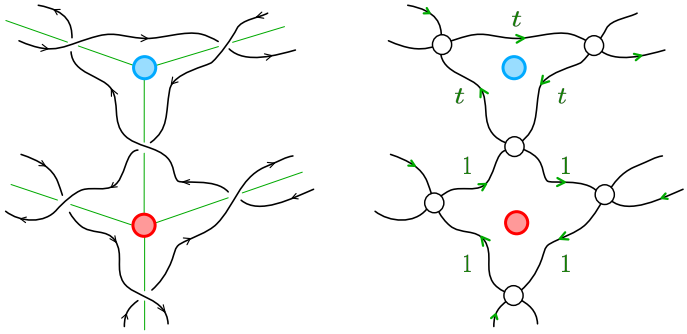


Figure 3.28: A special alternating link (left) and its Crowell graph (right). The blue and red vertices belong to S and R, respectively.

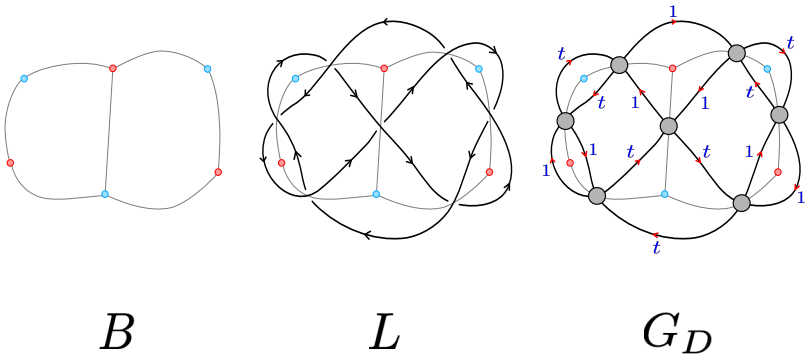
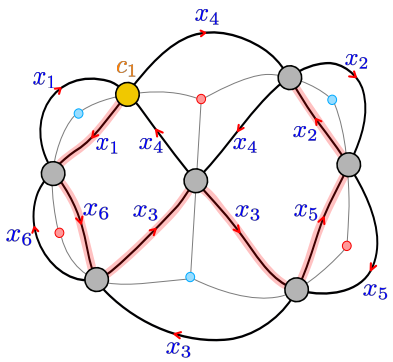


Figure 3.29: A bipartite Tait graph B, the corresponding special alternating link L, and the Crowell graph G_D; see [HMV24].



$$W(T) = x_1 x_2 x_3^2 x_5 x_6$$

Figure 3.30: Weights on the graph $\widetilde{G}_D$; see [HMV24].

Definition 3.29. A subset $J \subseteq \mathbb{N}^n$ is called *M-convex* if it satisfies the following exchange property. Whenever $\alpha, \beta \in J$ and $\alpha_i > \beta_i$ for some index i, there exists an index j such that $\alpha_j < \beta_j$ and

$$\alpha - e_i + e_j \in J, \qquad \beta - e_j + e_i \in J,$$

where $e_1, \ldots, e_n$ denote the standard basis vectors of $\mathbb{R}^n$.

Both properties follow from the following theorem of Kálmán [Kál13]; see also Hafner, Mészáros, and Vidinas [HMV24] for an alternative proof.

Lemma 3.30. *Let $C_1, \ldots, C_n, C_{n+1}, \ldots, C_{n+m}$ be cycles in $\widetilde{G}_D$, where C_i is the boundary of the region v_i. Let T be a spanning tree of the unoriented graph underlying $\widetilde{G}_D$, and let $a_i(T)$ denote the number of edges of T lying in C_i. Then there exists a unique oriented tree A rooted at c_1 such that*

$$a_i(A) = a_i(T)$$

for all $i \in \{1, \ldots, n+m\}$.

Lemma 3.30 implies that all nonzero coefficients of $P_{\widetilde{G}_D, c_1}$ are equal to 1. Moreover,

$$\operatorname{Supp}\big(P_{\widetilde{G}_D, c_1}\big) = \big\{(a_1(T), \ldots, a_{n+m}(T)) \ : \ T \in T(\widetilde{G}_D)\big\},$$

where $T(\widetilde{G}_D)$ denotes the set of spanning trees of the unoriented graph underlying $\widetilde{G}_D$. Since $T(\widetilde{G}_D)$ is the set of bases of the graphical matroid associated to $\widetilde{G}_D$, it is M-convex. Consequently, the support of $P_{\widetilde{G}_D, c_1}$ is M-convex.

Based on this, Hafner, Mészáros, and Vidinas conclude that $P_{\widetilde{G}_D, c_1}$ is *denormalized Lorentzian*, a generalization of log-concavity introduced by Brändén and Huh. As a consequence, they prove the strong Fox conjecture for special alternating links.

The main difficulty in extending this argument beyond the special alternating setting is that Proposition 3.27 no longer holds. Type 2 Seifert cycles still correspond to oriented cycles in the Crowell graph, but they contain edges with different weights. To construct a multivariable refinement of the Alexander polynomial, one therefore needs to use two variables for the edges of a type 2 Seifert cycle; see Figure 3.31.

Notation 3.31. For the remainder of this discussion, we will also use the term *Seifert cycle* for the corresponding oriented cycle in the Crowell graph.

We now carry out an analogous construction for alternating 3-braids.

Construction 3.32. Let D be the standard diagram of the closure of the alternating 3-braid

$$\sigma_1^{p_1}\sigma_2^{-q_1}\cdots\sigma_1^{p_n}\sigma_2^{-q_n},$$

where $p_i, q_i \in \mathbb{Z}_{>0}$ and σ_1, σ_2 are the standard generators of the braid group on three strands.

The diagram D is shown on the left of Figure 3.31 as a long knot, with the second strand closing at ∞. It contains two type 1 Seifert cycles, following the leftmost and rightmost strands. We denote these by L and R, respectively; they are shown in blue and green in Figure 3.31. The diagram also contains a type 2 Seifert cycle following the middle strand.

In the Crowell graph G_D, the left Seifert cycle is oriented counterclockwise and has weight t, while the right Seifert cycle is oriented clockwise and has weight 1. The middle strand is directed from top to bottom and contains two types of edges: the edges ending in L, denoted by EL and weighted by 1, and the edges ending in R, denoted by ER and weighted by t.

The refined Crowell graph $\widetilde{G}_D$ is shown on the right of Figure 3.31. We assign weights x_1, x_2, x_3, and x_4 to the edges of L, ER, EL, and R, respectively. To simplify notation, we will often use the colour-coding of Figure 3.31. Thus the variables x_1, x_2, x_3, x_4 will be referred to as the *blue*, *red*, *black*, and *green* weights, respectively.

The refined multivariable polynomial $P_{\widetilde{G}_D,c_1}(x_1, x_2, x_3, x_4)$ is defined by

$$P_{\widetilde{G}_D,c_1} = \sum_{T\in\mathcal{T}(\widetilde{G}_D,c_1)} W(T),$$

where $\mathcal{T}(\widetilde{G}_D, c_1)$ denotes the set of maximal oriented trees in $\widetilde{G}_D$ rooted at c_1 and $W(T)$ is the product of the weights of the edges of T.

Unlike the special alternating case, the coefficients of $P_{\widetilde{G}_D,c_1}$ need not all be equal to 1. Nevertheless, its support admits a particularly simple description.

Lemma 3.33. *Let $P = |V(L)|$ and $Q = |V(R)|$. The support of $P_{\widetilde{G}_D,c_1}$ is given by*

$$\big\{(a_1, a_2, a_3, a_4) \in \mathbb{Z}^4 : a_1, a_3, a_4 \geq 0,\ a_2 > 0,\ a_1 + a_3 = P - 1,\ a_2 + a_4 = Q\big\}.$$

Proof. Consider the braid diagram shown in Figure 3.31, with the standard height function in the plane. Let c_1 be the topmost vertex of L. Let $(a_1, a_2, a_3, a_4) \in \mathrm{Supp}\big(P_{\widetilde{G}_D,c_1}\big)$. Choose a tree T contributing to the corresponding monomial.

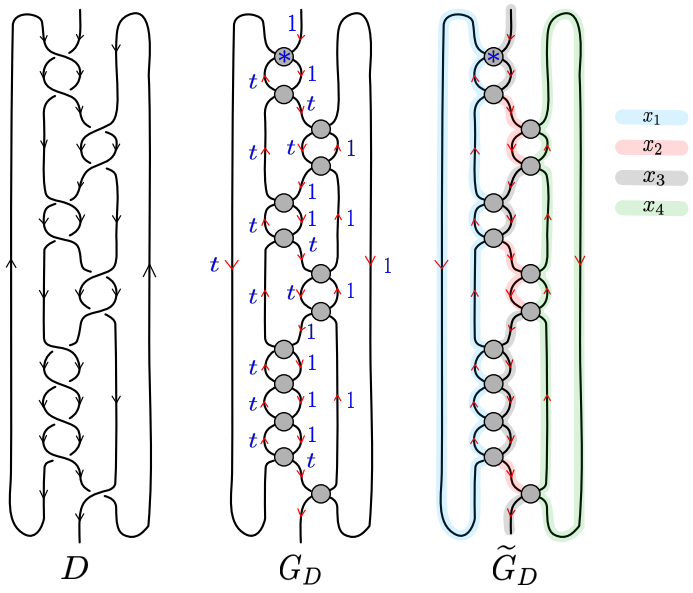


Figure 3.31: An alternating 3-braid (left), its Crowell graph (middle), and the refined Crowell graph (right). The distinguished root vertex c_1 is marked by a star.

Since T is rooted at c_1, every vertex other than c_1 is the terminal vertex of a unique edge of T. The edges of weight x_1 and x_3 are precisely the edges ending in L. Since L contains P vertices and c_1 is the root, it follows that $a_1+a_3 = P-1$. Similarly, the edges of weight x_2 and x_4 are precisely the edges ending in R, so $a_2 + a_4 = Q$. Moreover, T must contain at least one edge of type ER, and hence $a_2 > 0$.

Conversely, let $(a_1, a_2, a_3, a_4) \in \mathbb{Z}^4$ satisfy

$$a_1, a_3, a_4 \geq 0, \qquad a_2 > 0, \qquad a_1 + a_3 = P - 1, \qquad a_2 + a_4 = Q.$$

We construct a tree $T \in \mathcal{T}(\widetilde{G}_D, c_1)$ with weight

$$W(T) = x_1^{a_1} x_2^{a_2} x_3^{a_3} x_4^{a_4}.$$

Denote the topmost vertex of R by c_1'. Consider the following subsets of the edge set $E(\widetilde{G}_D)$:

$$E_1 = \text{the first } a_1 \text{ edges of } L \text{ starting at } c_1,$$
$$E_2 = \text{the topmost } a_2 \text{ red edges},$$
$$E_3 = \text{the topmost } a_3 \text{ black edges},$$
$$E_4 = \text{the first } a_4 \text{ edges of } R \text{ starting at } c_1'.$$

Let $E_T = E_1 \cup E_2 \cup E_3 \cup E_4$. We claim that the edges of E_T form the desired oriented rooted tree T. See the right-hand side of Figure 3.32 for the edges in E_T when $(a_1, a_2, a_3, a_4) = (3, 3, 4, 2)$, for the diagram in Figure 3.31.

Since $a_1 + a_3 = P - 1$ and $a_2 + a_4 = Q$, for any vertex c other than the root c_1, there is exactly one edge in E_T leading to c. Indeed, the edges in E_1 lead to the lowest a_1 vertices of L, and the edges in E_3 lead to the topmost a_3 vertices of L, excluding c_1. Similarly, edges in E_2 lead to the topmost a_2 vertices of R, and edges in E_4 lead to the lowest a_4 vertices of R. Since $a_2 > 0$, the edges E_T do not contain the cycle R.

By the above, for any vertex $v \neq c_1$, we can define the parent vertex $p(v)$ of v such that $(p(v), v) \in E_T$; i.e., as the source of the unique edge in E_T leading to v. Note that $p(v)$ is higher than v, unless possibly when $(p(v), v) \in E_1 \cup E_4$; i.e., when $(p(v), v)$ is an edge in one of the cycles L and R. For some $k \in \mathbb{N}$, the sequence

$$(p^k(v), \ldots, p(v), v)$$

is a path from c_1 to v in E_T. Hence E_T is the edge set of an oriented rooted tree T with the desired weight. □

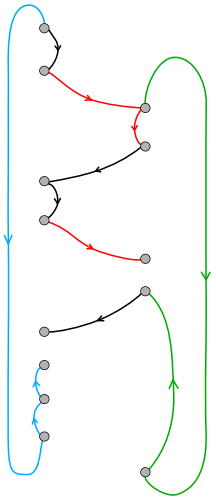


Figure 3.32: An oriented spanning tree T of the refined Crowell graph shown in Figure 3.31 for $(a_1, a_2, a_3, a_4) = (3, 3, 4, 2)$.

Notation 3.34. By Lemma 3.33, every monomial appearing in $P_{\widetilde{G}_D, c_1}$ has an exponent vector of the form $(P - 1 - i, Q - j, i, j)$. Accordingly, we write

$$P_{\widetilde{G}_D, c_1} = \sum_{\substack{0 \leq i \leq P-1 \\ 0 \leq j \leq Q-1}} A_{i,j}\, x_1^{P-1-i} x_2^{Q-j} x_3^i x_4^j \tag{3.18}$$

and refer to the matrix $A = (A_{i,j})$ as the *coefficient matrix* of the refined Crowell polynomial. The Crowell polynomial is recovered by setting

$$x_1 = x_2 = t, \qquad x_3 = x_4 = 1.$$

Thus

$$\Delta_L(-t) \doteq P_{G_D,c_1} = \sum_{0\le k\le P+Q-2} \left(\sum_{i+j=k} A_{i,j} \right) t^{P+Q-1-k}. \tag{3.19}$$

The remainder of this subsection is devoted to proving inequalities for the entries of the matrix A and translating them into inequalities for the coefficients of the Alexander polynomial. Our first goal is to prove the following theorem.

Theorem 3.35. *Let $A = (A_{i,j})$ be the coefficient matrix of the refined Crowell polynomial associated to the closure of the alternating 3-braid $\sigma_1^{p_1}\sigma_2^{-q_1}\cdots\sigma_1^{p_n}\sigma_2^{-q_n}$. Set $P = p_1 + \cdots + p_n$ and $Q = q_1 + \cdots + q_n$. Then*

$$nA_{i,j+1} \ge A_{i,j}, \qquad 0 \le i \le P-1, \quad 0 \le j \le Q-2.$$

Proof. Let $\mathcal{T}_{i,j}$ be the set of oriented rooted spanning trees of $\widetilde{G}_D$ with weight $x_1^{P-1-i}x_2^{Q-j}x_3^i x_4^j$; i.e., ones contributing to the coefficient $A_{i,j}$. Equivalently, $\mathcal{T}_{i,j}$ is the set of maximal oriented rooted trees in $\widetilde{G}_D$ with i black edges and j green edges.

We construct a map

$$\mathcal{L}\colon \mathcal{T}_{i,j} \to \mathcal{T}_{i,j+1}$$

which is at most n-to-1. This implies that

$$nA_{i,j+1} \ge A_{i,j}.$$

Recall that c_1 and c_1' were used to denote the highest vertices of L and R, respectively. Let $c_1, \ldots, c_P$ be the vertices of L, and $c_1', \ldots, c_Q'$ be the vertices of R, both labelled from top to bottom.

Let $T \in \mathcal{T}_{i,j}$ be an oriented, rooted spanning tree. We call c_i' a *maximal vertex* if the path in T from c_1 to c_i' only includes vertices in L, except for c_i'. Let the lowermost maximal vertex be c_m'. Consider the longest oriented path of green edges in T originating from c_m'. Note that the path might be of length zero. Assume that this path ends at c_k' and hence does not reach c_{k-1}'. Then there is a red edge in T leading to c_{k-1}', which we denote by e. The map $\mathcal{L}$ is defined by

$$\mathcal{L}(T) := T \setminus \{e\} \cup \{(c_k', c_{k-1}')\}.$$

See Figure 3.33 for an example, where the maximal vertices are c_1' and c_7', the lowest maximal vertex is $c_m' = c_7'$, and the green path goes from c_7' to $c_k' = c_6'$.

We first show that $\mathrm{Im}(\mathcal{L}) \subseteq \mathcal{T}_{i,j+1}$. By construction,

$$W(\mathcal{L}(T)) = W(T)x_4x_2^{-1}.$$

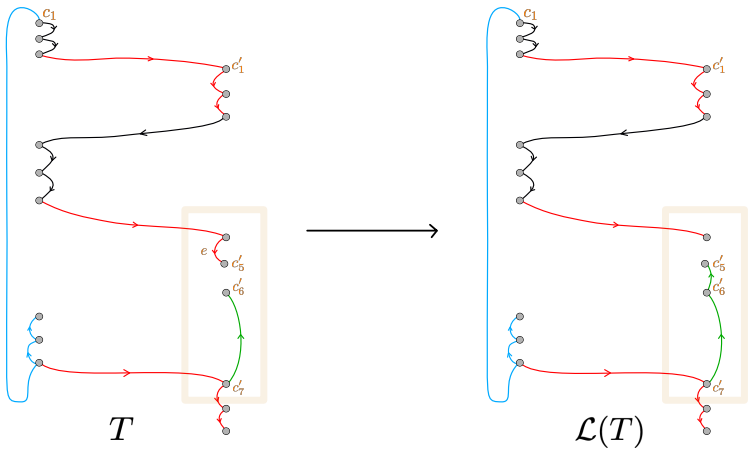


Figure 3.33: The map $\mathcal{L}\colon \mathcal{T}_{i,j} \to \mathcal{T}_{i,j+1}$.

Hence, it suffices to show that $\mathcal{L}(T)$ is an oriented rooted spanning tree.

The rooted spanning tree T induces a partial order $>_T$ on $V(\widetilde{G}_D)$ by declaring $v >_T w$ if and only if there is a path from v to w in T. For any vertex v, the subset

$$T_v = \{x \in V(\widetilde{G}_D) : v \geq_T x\}$$

forms an oriented subtree of T rooted at v. Conversely, any partial order on $V(\widetilde{G}_D)$ with c_1 as its maximum determines an oriented rooted spanning tree.

The construction of $\mathcal{L}(T)$ can be described as changing the parent of c'_{k-1} to c'_k. This produces an oriented rooted spanning tree if and only if

$$c'_k \notin T_{c'_{k-1}};$$

see Figure 3.34. This condition holds here. Indeed, if $c'_k \in T_{c'_{k-1}}$, then there would be an oriented path from c'_{k-1} to c'_k, contradicting the maximality of c'_m.

We now prove that $\mathcal{L}$ is at most n-to-1; i.e., we show that, for any $S \in \mathcal{T}_{i,j+1}$, we have $|\mathcal{L}^{-1}(S)| \leq n$. To find the elements of $\mathcal{L}^{-1}(S)$, first notice that the lowermost maximal vertices of T and $\mathcal{L}(T)$ are the same, since the construction of $\mathcal{L}(T)$ does not affect the path from c_1 to c'_m, and just turns c'_{k-1} from a possibly maximal vertex to a non-maximal one.

Using this fact, consider the lowermost maximal vertex of S and denote it by c'_m. Denote by J the largest path of green edges in S starting from c'_m. By construction of $\mathcal{L}$, we know that, if $T \in \mathcal{L}^{-1}(S)$, then S can be constructed by taking the longest path of green edges in T starting from c'_m and adding an edge at the end. The trees T and S only differ in two edges, one green and one red. Denote the one extra red edge of T by $e_T = (x, y)$. Based on the explanation above, y must be in J.

To construct T from S, we delete the green edge leading to y and replace it with e_T. As a result, x must be outside of J, otherwise the green edge from (y, x) in S is also in T and forms a cycle with e_T. This restricts the options for e_T: it is either one of the edges starting in L and ending in R, or it is an edge leading to the final leaf of J. So there are at most n options for e_T, hence $|\mathcal{L}^{-1}(S)| \leq n$. See Figure 3.35 for an example. Therefore $\mathcal{L}$ is at most n-to-1, which proves the theorem. □

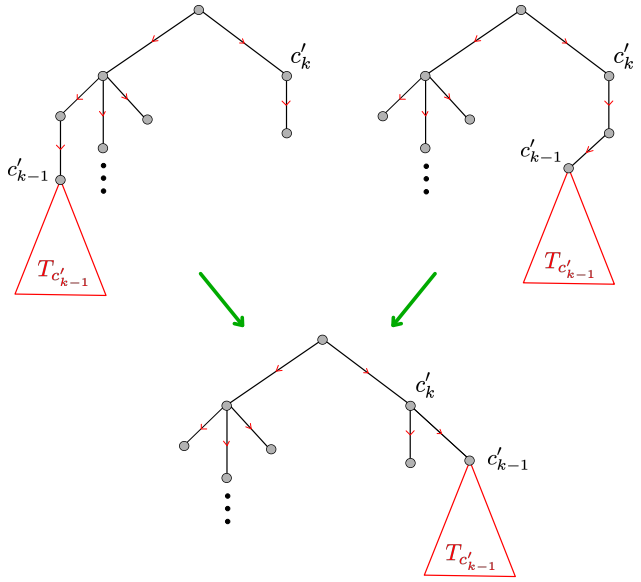


Figure 3.34: Two different examples of the *change of parent operation* in a rooted tree. In both examples, $c'_k \notin T_{c'_{k-1}}$ and the result of the operation is still an oriented rooted tree.

Combining Theorem 3.35 and Equation (3.19) immediately yields inequalities for the coefficients of the Alexander polynomial. Furthermore, by taking mirrors and turning the diagram upside down, one can interchange the roles of σ_1 and σ_2, obtaining the same inequality with P and Q exchanged. This gives us the following corollary.

Corollary 3.36. *Let (a_k) be the sequence of absolute values of the coefficients of the Alexander polynomial of the alternating 3-braid $\sigma_1^{p_1}\sigma_2^{-q_1}\cdots\sigma_1^{p_n}\sigma_2^{-q_n}$. Then*

$$a_{k+1} \geq \frac{a_k}{n} + A_{0,k+1} > \frac{a_k}{n}, \qquad 0 \leq k \leq \max(P-2, Q-2).$$

Furthermore, the estimate on the number of possible choices for the edge e_T can be sharpened. This yields the stronger inequality stated in Theorem 3.26.

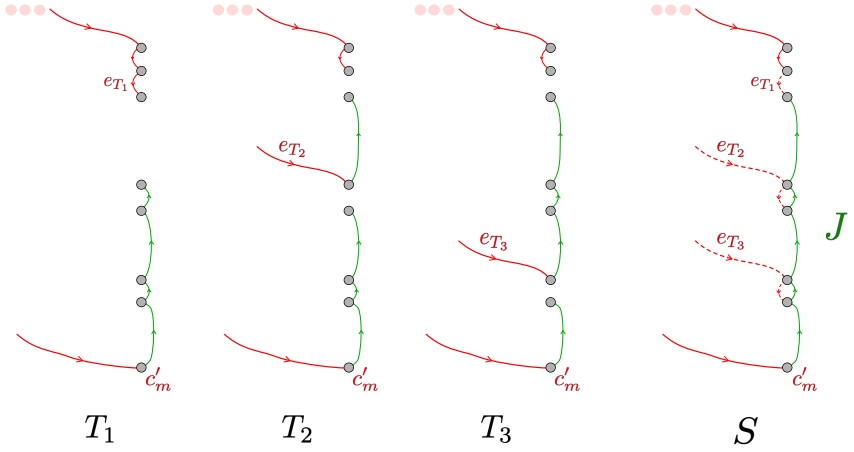


Figure 3.35: Construction of the inverse image $\mathcal{L}^{-1}(S)$. The possible options for e_T come from the red edges leading to J from outside. One can see that using a red edge with both ends in J does not work.

Proof of Theorem 3.26. The argument is identical to the proof of Theorem 3.35, except that we improve the estimate on the size of the fibres of $\mathcal{L}$. As discussed in the proof of Theorem 3.35, the cardinality of $\mathcal{L}^{-1}(S)$ is bounded above by the number of possible choices for the edge e_T, namely a red edge leading to the path J. Thus

$$|\mathcal{L}^{-1}(S)| \leq |J|.$$

Each edge of J contributes weight 1 to $W(T)$. Consequently, if $W(S) = t^{\mathrm{pw}(S)}$, then

$$|J| < P + Q - 2 - \mathrm{pw}(S).$$

The remainder of the proof is identical to that of Theorem 3.35. □

One might ask whether this argument extends beyond alternating 3-braids. The key ingredient in the proof of Theorem 3.35 is the structure of the Crowell graph near the Seifert cycle R. In fact, the argument only uses that R is a Seifert cycle of type 1.

The local shape of an alternating diagram near a type 1 Seifert cycle is shown at the top of Figure 3.36. The diagram consists of n twist regions along the boundary of the cycle. The edges are coloured analogously to the neighborhood of the Seifert cycle R in the Crowell graph of an alternating 3-braid. Using this colouring, the tree-manipulation argument from the proof of Theorem 3.35 can be repeated in this more general setting, leading to the following theorem.

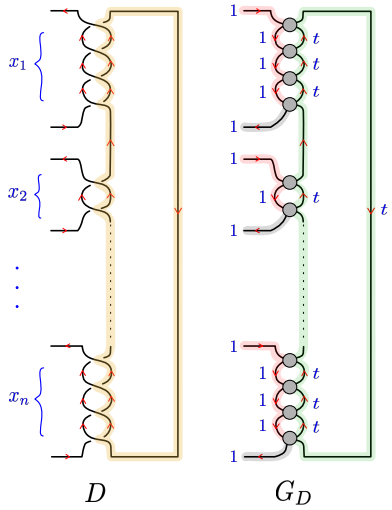


Figure 3.36: Left: a neighborhood of a type 1 Seifert cycle in an alternating diagram. Right: the corresponding portion of the Crowell graph, together with the colouring used in the proof of Theorem 3.35.

Theorem 3.37. *Let D be an alternating diagram of a link L, and suppose that D contains a type 1 Seifert cycle C surrounded by n twist regions with $x_1, \dots, x_n$ crossings. Let $N = x_1 + \cdots + x_n$. Choose a root crossing c disjoint from C and write*

$$P_{G_D,c}(t) = \sum_{k \geq 0} a_k t^k.$$

Then

$$a_{i+1} \geq \frac{a_i}{n}, \qquad 0 \leq i \leq N-2.$$

We conclude this subsection with a brief remark on the possibility of extending the argument of Hafner, Mészáros, and Vidinas directly to alternating 3-braids. Although the coefficients of the refined Crowell polynomial $P_{\widetilde{G}_D,c_1}$ are no longer all equal to 1, one might still hope that the polynomial is denormalized Lorentzian. Unfortunately, this is not true in general, as stated in the following proposition.

Proposition 3.38. *Let $P_{\widetilde{G}_D,c_1}$ be the refined Crowell polynomial of the closure of $B = \sigma_1^{p_1}\sigma_2^{-q_1} \cdots \sigma_1^{p_n}\sigma_2^{-q_n}$. Then $P_{\widetilde{G}_D,c_1}$ is not denormalized Lorentzian unless the braid is a connected sum of two torus links; equivalently, $B = \sigma_1^{p_1}\sigma_2^{-q_1}$.*

We do not prove Proposition 3.38 here and instead refer to our paper [AJK24]. The proof proceeds by analyzing the definition of denormalized Lorentzian polynomials for polynomials of the form given in Equation (3.18). This leads to restrictive conditions on the coefficient matrix $A = (A_{i,j})$. Proposition 3.39 records these conditions. A

further tree-counting argument then shows that these conditions cannot be satisfied by the refined Crowell polynomial of an alternating 3-braid; see Proposition 3.40.

Proposition 3.39. *For a $P \times Q$ matrix A with positive entries, the polynomial*

$$f = \sum_{\substack{0 \le i \le P-1 \\ 0 \le j \le Q-1}} A_{i,j}\ x_1^{P-1-i} x_2^{Q-j} x_3^i x_4^j$$

is denormalized Lorentzian if and only if there exist log-concave sequences $(z_0, \ldots, z_{P-1})$ and $(y_0, \ldots, y_{Q-1})$ such that

$$A_{i,j} = z_i y_j.$$

Proposition 3.40. *The extremal entries of the matrix $A = [A_{i,j}]$ of coefficients of the refined Crowell polynomial of the closure of an alternating 3-braid $B = \sigma_1^{p_1}\sigma_2^{-q_1} \cdots \sigma_1^{p_n}\sigma_2^{-q_n}$ satisfy the following:*

- $A_{i,0} = 1$ *for* $i \in \{0, \ldots, P-1\}$,
- $A_{P-1,j} = 1$ *for* $j \in \{0, \ldots, Q-1\}$, *and*
- $A_{0,j} \ge n$ *for* $j \in \{1, \ldots, Q-1\}$.

Chapter 4

Twisting and Stabilization in Knot Floer Homology

Any time you have a framing, take it to the limit

- Matthew Hedden

In Section 2.1.2, we introduced the twisting operation and twist families of knots. We also discussed the *stabilization phenomenon*, namely the behaviour of knot invariants as the number of twists grows to infinity. In particular, we recalled stabilization results for the Seifert genus (Theorem 2.19) and the Alexander polynomial (Theorem 2.20). The latter was revisited in Section 3.2 in the context of alternating links.

In this chapter, we study the stabilization phenomenon for knot Floer homology. The first motivation comes from categorification. Since knot Floer homology categorifies the Alexander polynomial, it is natural to ask whether the stabilization results for the Alexander polynomial discussed above admit a categorified analogue. Moreover, knot Floer homology detects the Seifert genus, so a stabilization theorem for knot Floer homology also provides a natural route toward genus stabilization. From this point of view, it is interesting to seek a Floer-theoretic proof of the genus stabilization theorem of Baker and Taylor, whose original proof uses three-dimensional methods involving the Thurston norm.

Second, understanding stabilization in knot Floer homology can be useful for studying the stabilization of more subtle knot invariants. For example, the asymptotic behaviour of the slice genus and the unknotting number in twist families has been studied by Baker–Motegi [BM19] and Baker–Miyazawa–Motegi [BMM25]. Since knot Floer homology provides lower bounds and structural information for several such invariants, its stabilization provides a natural framework for approaching such questions.

The question is also intrinsic to knot Floer homology. Stabilization phenomena already appear in the study of cables and satellites. In particular, taking the framing to be sufficiently large plays an essential role in Hedden's work on knot Floer homology of cables, and related *large-framing arguments* appear in several other contexts. More recently, Alishahi, Gorsky, and Liu [AGL25] used stabilization of knot Floer homology to define a coloured version of knot Floer homology, in analogy with coloured invariants and coloured Khovanov-type theories in quantum topology. This is part of a broader and active interaction between quantum topology and Heegaard Floer theory, with ideas moving in both directions.

Finally, stabilization results can be used to build and analyse new families of examples. For instance, Lambert-Cole [LC18a] used stabilization under twisting to construct an infinite family of prime positive mutant knots with isomorphic knot Floer homology groups. This illustrates that stabilization is not only a structural phenomenon, but also a tool for testing the strength and limitations of knot Floer homology.

The main results of this chapter concern three different stabilization phenomena in knot Floer homology. We prove stabilization of the total dimension of knot Floer homology, stabilization of the extremal knot Floer homology groups, and stabilization of the τ-invariant. These results are summarised below.

Throughout this chapter, we use the notation of Definition 2.10, Definition 2.13, Notation 2.12, and the descriptions of twisting from Construction 2.11. Thus, we always work with a twist family $\{K_m\}$ obtained from a knot K and a twisting unknot c, and we often regard K as a knot in the solid torus E_c.

The first result describes the asymptotic behaviour of the total dimension of knot Floer homology in a twist family. It shows that, when $\mathrm{wind}_c(K) \neq 0$, this dimension eventually becomes an affine linear function of m. This phenomenon does not seem to have been explicitly isolated in the literature in this generality. In the special case where c links K coherently and $\mathrm{wind}_c(K) = 2$, however, it follows from Lambert-Cole's stabilization theorem for knot Floer homology [LC18a].

Theorem 4.1. *Let $\{K_m\}$ be a twist family of knots. Assume that $\mathrm{wind}_c(K) \neq 0$. Then there exist integers $D, d \in \mathbb{Z}$ such that*

$$\dim\big(\widehat{HFK}(K_m)\big) = D \cdot m - d$$

for sufficiently large m.

The second result concerns stabilization of the extremal knot Floer homology groups. This can be viewed as an analogue of the stabilization of the extremal coefficients of the Alexander polynomial: as $m \to \infty$, the extremal groups stabilize, up to a fixed shift in the Maslov grading.

Stabilization phenomena of this kind first appeared in the study of cables. For example, for fixed $p > 1$, the family of cables $K_{p,pm+1}$ can be viewed as a coherent twist family, and Hedden's work [Hed05] on knot Floer homology of cables implies stabilization in the extremal Alexander gradings. Related stabilization phenomena also appear in Hedden's work [Hed07] on twisted Whitehead doubles.

The first stabilization theorem for a general twist family was due to Lambert-Cole [LC18a] in the case of a coherent twist family with $\mathrm{wind}_c(K) = 2$. The author [Aza23] later showed that the same argument extends to coherent twist families of higher winding number, although it fails to address absolute Maslov gradings.

More recently, Alishahi, Gorsky, and Liu [AGL25] studied analogous stabilization phenomena in the broader setting of the "full" version of link Floer homology.

Theorem 4.3 replaces the coherence assumption with the weaker condition that $\mathrm{wind}_c(K) \neq 0$, and identifies the resulting shift in Maslov grading. In particular, it generalizes all previous stabilization results for the hat version of knot Floer homology.

To state Theorem 4.3, we introduce an auxiliary invariant of a knot in a solid torus.

Definition 4.2. Using the notation of Notation 2.12, define

$$F_K := \frac{1}{4}\big(\mathrm{lk}(K,c) - x([\hat{D}]) + 1\big)\big(\mathrm{lk}(K,c) - x([\hat{D}]) - 1\big).$$

Note that if c links K coherently, then

$$x([\hat{D}]) = \mathrm{lk}(K,c) - 1,$$

and hence $F_K = 0$.

Theorem 4.3. *Let $\{K_m\}$ be a twist family of knots and assume that*

$$\mathrm{wind}_c(K) \neq 0.$$

For every $\ell \in \mathbb{N}$, there exists $N \in \mathbb{N}$ such that for all $m \geq N$ and $0 \leq j \leq \ell - 1$, there exists a graded isomorphism

$$\widehat{HFK}(K_m, -g(K_m) + j) \cong \widehat{HFK}(K_{m+1}, -g(K_{m+1}) + j)[F_K],$$

where $[F_K]$ denotes a downward shift in the Maslov grading by F_K.

The final result concerns the behaviour of the τ-invariant under twisting. Similar to the Seifert genus and the total dimension of knot Floer homology, the τ-invariant also stabilizes to an affine linear function.

The first such linearity results appeared implicitly in Hedden's work on $K_{p,pm+1}$ cables [Hed09]. Van Cott [VC10] later proved this result for coherent twist families. Roberts [Rob11] subsequently extended it to twist families arising from general twisted satellites. However, in this general setting, Roberts only proved that the τ-invariant remains within a uniformly bounded distance of a linear function.

Theorem 4.4 establishes the corresponding statement for arbitrary twist families with nonzero winding number.

Theorem 4.4. *Let $\{K_m\}$ be a twist family of knots and assume that* $\mathrm{wind}_c(K) \neq 0$. *Then there exist integers $T, t \in \mathbb{Z}$ such that*

$$\tau(K_m) = T \cdot m + t$$

for sufficiently large m.

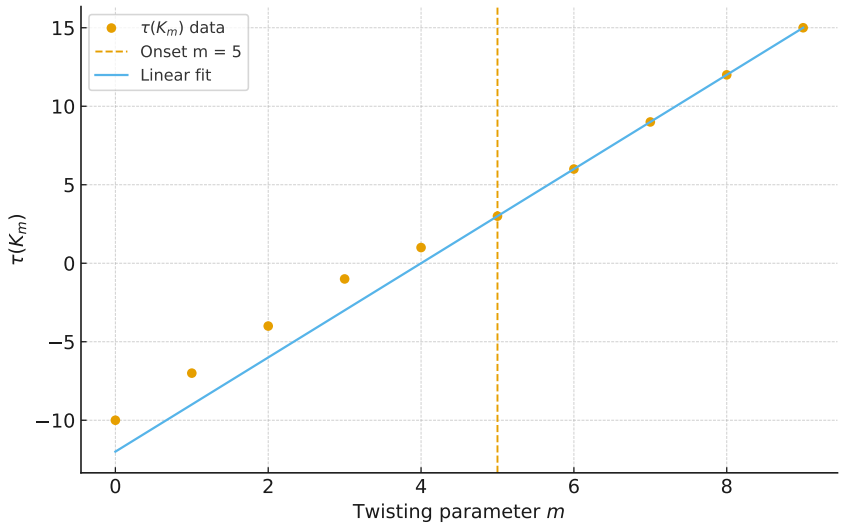


Figure 4.1: Stabilization of $\tau(K_m)$ to an affine linear function.

We will not discuss the results mentioned above in detail, nor restate them formally, as Theorems 4.1, 4.3, and 4.4 subsume the stabilization statements relevant to this chapter. Nevertheless, it is worth briefly comparing the methods used in the existing literature with those employed here.

In terms of methodology, the previous approaches can be roughly divided into the following two categories:

1. Hedden [Hed05], Lambert-Cole [LC18a], and Alishahi–Gorsky–Liu [AGL25] study the local effect of twisting on a multipointed Heegaard diagram representing the knot. Using a suitable Heegaard diagram, the twisting disk D can be isotoped to a small disk on the Heegaard surface, so that twisting corresponds to applying powers of the Dehn twist $\tau_{\partial D}$ to the $\boldsymbol{\beta}$ curve collection.

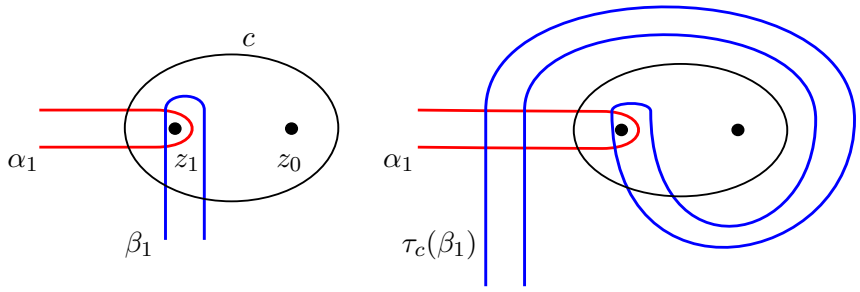


Figure 4.2: A local region of a multipointed Heegaard diagram before (left) and after (right) twisting. The diagram on the right is obtained from the diagram on the left by applying a Dehn twist along the curve c to the $\boldsymbol{\beta}$ curves [LC18b].

The rest of the argument consists of analyzing the Heegaard states and their relative gradings in a neighbourhood of D. Roughly speaking, they show that the Heegaard states contributing to the extremal Alexander gradings are forced to behave in a highly constrained way near D. This local control determines the differentials between such states and ultimately leads to the stabilization results.

The diagrammatic approach has two main limitations:

(I) All existing results obtained by this method concern coherent twist families.

(II) The approach is fundamentally local and therefore primarily detects information in the extremal Alexander gradings. As a consequence, it does not provide effective access to global features of knot Floer homology, namely phenomena that depend on the behaviour across many Alexander gradings rather than only the extremal ones. In particular, it does not naturally address quantities such as the total dimension, absolute Maslov gradings, or the τ-invariant.

Remark 4.5. There are ways to overcome some aspects of the Maslov grading problem. Lambert-Cole uses the skein exact triangle together with an inductive argument to determine the grading shift in the two-strand case. Alishahi, Gorsky, and Liu use a more modern version of this idea via cobordism maps in link Floer homology. Nevertheless, these techniques still do not appear to provide access to stabilization phenomena involving other global features of knot Floer homology.

2. Van Cott [VC10] and Roberts [Rob11] study the τ-invariant by building clever knot cobordisms and leveraging the inequality between τ and the slice genus.

Unlike the diagrammatic approach discussed above, this method is not restricted to coherent twist families. On the other hand, it does not directly imply that the τ-invariant eventually becomes an affine linear function. However, it applies in greater generality and does not require a large number of twists.

Our approach: The approach taken in this chapter is fundamentally different. We interpret twisting as a gluing operation and analyse it using bordered Floer homology and the immersed curve invariant. To the author's knowledge, this is the first systematic use of bordered Floer theory in the study of stabilization phenomena for knot invariants. This perspective has several advantages:

- It provides a unified combinatorial framework for studying stabilization in knot Floer homology. The same machinery is used to prove Theorems 4.1, 4.3, and 4.4, and is flexible enough to address further stabilization questions.
- It relaxes the coherence assumption appearing in previous work to the weaker condition that $\mathrm{wind}_c(K) \neq 0$.
- It is inherently a global method. The construction produces explicit representatives for both $\widehat{\mathrm{CF}}(S^3, K_m)$ and $\widehat{\mathrm{CF}}(S^3)$ on the same underlying vector space. It naturally determines the differentials $\widehat{\partial}_{K_m}$ and $\widehat{\partial}$.
- This additional information gives access to the absolute Maslov grading, the knot filtration, and consequently to invariants such as the τ-invariant.
- The resulting constructions are sufficiently explicit to allow effective computations and the analysis of concrete examples.

The chapter is organised as follows. Since the proofs of the main results are long and technically involved, we begin in Section 4.1 with an overview of the arguments. This section is intended to give the reader a clearer sense of the main ideas and techniques, while also serving as a road map for the remainder of the chapter.

Section 4.2 applies the immersed curve invariant to prove Theorem 4.1.

In Section 4.3, we use the pairing theorem in bordered Floer homology to analyse the box tensor product that computes the knot Floer homology of the knots K_m. This section contains the combinatorial framework used throughout the chapter. Section 4.4 applies this framework to study the extremal Alexander gradings and prove Theorem 4.3. Finally, Section 4.5 uses the same framework to prove Theorem 4.4.

We conclude the introduction with a few additional remarks. As discussed above, the stabilization framework developed in this chapter can be applied to a broader range of questions concerning stabilization phenomena in knot Floer homology. Some of these directions remain work in progress, while others are developed in greater detail in our paper [Aza25]. We briefly record two such results here.

The first concerns the knot Floer thickness, denoted by $\mathrm{th}(K)$. Using the same methods, one obtains the following stabilization theorem.

Theorem 4.6. *Let $\{K_m\}$ be a twist family of knots. Assume that $\mathrm{wind}_c(K) \neq 0$. Then there exist integers $W, w \in \mathbb{Z}$ such that for sufficiently large m,*

$$\mathrm{th}(K_m) = W \cdot m + w.$$

Theorem 4.7. *Let $\{K_m\}$ be a twist family. Then there exists an integer $e \geq 0$ such that for all sufficiently large m, there is a finite subset*

$$E_m \subset \mathbb{Z} \qquad \text{with} \qquad |E_m| = e,$$

such that for every $j \notin E_m$, there is an isomorphism

$$\Psi_j : \widehat{\mathrm{HFK}}(K_m, j) \xrightarrow{\sim} \widehat{\mathrm{HFK}}(K_m, j + \mathrm{wind}_c(K)),$$

preserving the homological grading modulo 2.

Remark 4.8. This phenomenon of near-periodicity, namely periodic behaviour with only a bounded number of exceptions, appears to be new in the context of knot Floer homology. It reveals a form of global regularity in Floer-theoretic data. Moreover, it can be viewed as a categorification of the periodicity phenomenon underlying Theorem 2.20.

4.1 Overview of the Proofs

Throughout the chapter, we assume $\mathrm{lk}(K, c) < 0$ (see Assumption 4.32).

Recall from Construction 2.11 that the knot K_m can be obtained from K by performing $\frac{1}{m}$-surgery on the unknot c. The main strategy is to reinterpret this description as a gluing of bordered three manifolds.

Construction 4.9. Consider the bordered manifolds

$$(E_c, \mu_c, -\lambda_c) \qquad \text{and} \qquad (H'_{1/m}, \mu_c, -\lambda_c),$$

where $E_c = S^3 \setminus N(c)$ is the solid torus containing the knot K. The curve $\lambda_c \subset \partial E_c$ is a Seifert longitude of c, which serves as a meridian of E_c, while $\mu_c \subset \partial E_c$ is a meridian of c, which serves as a longitude of E_c. The manifold $H'_{1/m}$ is a solid torus whose meridian represents the class $[\mu_c] + m[\lambda_c]$.

Proposition 4.10. *The manifold obtained by gluing*

$$(E_c, \mu_c, -\lambda_c) \qquad \textit{and} \qquad (H'_{1/m}, \mu_c, -\lambda_c)$$

is S^3. Moreover, under this identification, the knot $K \subset E_c$ is sent to the knot K_m.

In Section 4.2, we combine Proposition 4.10 with the computation of the immersed curve invariant of $H'_{1/m}$ (Example 2.121) and the pairing theorem for immersed curves (Theorem 2.99) to prove Theorem 4.1.

The bordered Floer approach begins by applying the pairing theorems of Chapter 2 to the decomposition of Proposition 4.10. Let $(\mathcal{H}_K, z, w)$ be a doubly-pointed bordered Heegaard diagram for (E_c, K) and let $(\mathcal{H}'_{1/m}, z')$ be the bordered Heegaard diagram for $H'_{1/m}$ used in Example 2.121.

Notation 4.11. For convenience, we introduce the following notation:

$$\mathcal{K} := \widehat{\mathrm{CFA}}(\mathcal{H}_K, z, w), \qquad \mathcal{S} := \widehat{\mathrm{CFA}}(\mathcal{H}_K, z), \qquad \mathcal{D}_m := \widehat{\mathrm{CFD}}(\mathcal{H}'_{1/m}, z'),$$

$$\mathcal{X}_m^K := \mathcal{K} \boxtimes \mathcal{D}_m, \qquad \text{and} \qquad \mathcal{X}_m := \mathcal{S} \boxtimes \mathcal{D}_m.$$

By Theorems 2.85 and 2.86, there are chain homotopy equivalences

$$\widehat{\mathrm{CFK}}(S^3, K_m) \simeq \mathcal{X}_m^K, \qquad \widehat{\mathrm{CF}}(S^3) \simeq \mathcal{X}_m. \tag{4.1}$$

The remainder of the chapter is devoted to a simultaneous analysis of these two box tensor products and their gradings.

Notation 4.12. We denote the absolute Maslov grading on both $\mathcal{X}_m^K$ and $\mathcal{X}_m$ by h_m. We denote the Alexander grading on $\mathcal{X}_m^K$ by a_m.

In Subsection 4.3.1, we construct a basis $\boldsymbol{C}^m$ for both $\mathcal{X}_m^K$ and $\mathcal{X}_m$. This basis admits a natural circular arrangement, illustrated in Figure 4.3, and decomposes into a distinguished *black box* $C^\bullet$ together with a collection of *white boxes* $C_1^\circ, \dots, C_m^\circ$ (see Definition 4.15 and Construction 4.16 for precise statements).

We also define two inclusion maps in Construction 4.20,

$$\Phi_m, \Phi'_m : \boldsymbol{C}^m \to \boldsymbol{C}^{m+1},$$

which endow the family $\{\boldsymbol{C}^m\}_{m\geq 0}$ with two nested structures (see Figures 4.4 and 4.5). This framework greatly simplifies many of our subsequent discussions on stabilization and is deeply reflected in the notation used throughout the chapter; see Remark 4.21.

We also describe the differentials in the box tensor products, as detailed in Proposition 4.22 and illustrated in Figure 4.8. Throughout the chapter, we adopt a graph-theoretic presentation for all chain complexes (see Notation 4.17).

In Subsection 4.3.2, we analyse the relative gradings of elements in the basis $\boldsymbol{C}^m$ and their behaviour as m varies. Moreover, we compute the grading differences between elements in $\boldsymbol{C}^m \setminus C^\bullet$ and their natural shifts.

In Subsection 4.3.3, we begin analyzing elements of $\mathcal{X}_m^K$ by viewing them as subsets of the basis $\boldsymbol{C}^m$. Working over $\mathbb{F}$, there is a natural one-to-one correspondence

$$I \subseteq \boldsymbol{C}^m \qquad \longleftrightarrow \qquad S_I = \sum_{\nu \in I} \nu$$

(see Notation 4.17). This perspective allows us to define a *minimality* condition, which can be viewed as a refinement of homogeneity and plays a key role in our later structural arguments.

In Subsection 4.3.4, we use the circular arrangement of $\boldsymbol{C}^m$ to define a natural metric on the set of basis elements. This allows us to introduce the notion of *length* for subsets of $\boldsymbol{C}^m$, which can be viewed as a refinement of diameter. Intuitively, the length measures how spread out a subset is with respect to the circular ordering.

We then establish a uniform upper bound on the length of subsets corresponding to minimal cycles, providing a key constraint for the stability arguments that follow.

In Subsection 4.3.5, we study *relations* between cycles, namely equalities of the form

$$\partial S_R = S_{I_1} + \cdots + S_{I_r},$$

where each S_{I_j} is a minimal cycle. We introduce a minimality condition for relations and establish upper bounds on the length of minimal relations, analogous to the bounds obtained for minimal cycles.

The previous two subsections provide a highly controlled description of the homology. The homology is determined by spaces of cycles and relations. As we will see, these spaces are generated by minimal cycles and minimal relations, whose size and length are uniformly bounded. Thus, the study of homology can be reduced to the study of a well-controlled collection of minimal objects and ultimately leads to the stabilization results in extremal gradings. The main remaining difficulty is the behaviour of the absolute Maslov grading, which requires additional analysis.

The next two subsections focus on the complex

$$\mathcal{X}_m = \mathcal{S} \boxtimes \mathcal{D}_m.$$

Note that the homology is equal to $\widehat{\mathrm{HF}}(S^3) = \mathbb{F}_2$. Hence our goal is to understand representatives of its unique non-trivial homology class in terms of the basis $\boldsymbol{C}^m$.

In Remark 2.120, we introduced the standard bordered Heegaard diagram $(\mathcal{H}_\infty, z)$ for the solid torus $(E_c, \mu_c, -\lambda_c)$ and computed the corresponding type A structure.

Notation 4.13. We denote $\mathcal{S}_\infty = \widehat{\mathrm{CFA}}(\mathcal{H}_\infty, z)$.

By Theorem 2.95, there exists an $\mathcal{A}_\infty$ homotopy equivalence

$$f : \mathcal{S}_\infty \longrightarrow \mathcal{S}.$$

In Subsection 4.3.6, we use this homotopy equivalence to study representatives of the unique non-trivial homology class. This analysis identifies distinguished stabilized elements of Maslov grading zero, which later serve as reference points for determining absolute Maslov gradings, particularly in the proof of Theorem 4.3.

In Subsection 4.3.7, we introduce combinatorial notions for comparing gradings of basis elements appearing in representatives of the non-trivial homology class. Rather than working directly with the gradings, we focus on information that can be extracted from the graph-theoretic description of the complex. While the resulting notions are elementary, they provide a convenient language that simplifies several arguments in Section 4.5.

Section 4.4 is devoted to the proof of Theorem 4.3, which is completed in Subsection 4.4.3. In Subsection 4.4.1, we introduce the closed metric balls around the black box, denoted by $B^\bullet_t$, together with the shift maps $R_\pm$ (see Figure 4.12). These balls capture the stable portion of the chain complex (illustrated in Figure 4.13) and form the foundation for nearly all stabilization arguments.

In Subsection 4.4.2, we show that the extremal knot Floer homologies are supported inside a fixed closed ball by analyzing the action of the shift maps on differentials and null-homologous classes. Finally, in Subsection 4.4.3, we prove that these extremal homologies are completely determined by the data contained in that fixed ball, namely the differentials and grading double coset representatives of its elements. This completes the proof of Theorem 4.3. In Section 4.5, we prove Theorem 4.4. We begin by reformulating the τ-invariant as

$$\tau(K_m) = \min_{S_I \in \mathcal{HZ}_m} \left(\max_{\nu \in I} a_m(\nu) \right),$$

where $\mathcal{HZ}_m$ denotes the set of homogeneous representatives of the unique non-trivial homology class of $\mathcal{X}_m$.

To compute $\tau(K_m)$, we introduce a set $\mathcal{P}_m$ consisting of basis elements that realize the maximal Alexander grading in some cycle of $\mathcal{HZ}_m$. The main goal is to show that $\mathcal{P}_m$ stabilizes inside a fixed closed ball, as stated in Lemma 4.81. The proof of Theorem 4.4 is then reduced to this stabilization statement.

In Subsection 4.5.2, we study the Maslov grading zero elements of $\boldsymbol{C}^m$ using the grading machinery developed in Subsection 4.3.2. Although this material could be placed in Section 4.3, we include it here because its proofs are closely tied to the stabilization arguments used throughout the remainder of the section.

Subsection 4.5.3 is devoted to the proof of Lemma 4.81. We first show that any element that can potentially realize the maximal Alexander grading in a cycle of $\mathcal{HZ}_m$ must lie in a fixed closed ball $B^\bullet_{L_\tau}$ (Proposition 4.89 and Corollary 4.90).

Next, we use the distinguished stabilized element from Lemma 4.57 to compare Alexander gradings inside this ball. This allows us to further restrict the set of possible maximizers and ultimately identify the stabilized subset that determines $\mathcal{P}_m$.

Finally, we prove that this collection of subsets of potential maximizers stabilizes as m grows. This is done in Lemmas 4.92 and 4.94. To this end, we decompose the closed ball $B^\bullet_{L_\tau}$ into its positive and negative half-balls (as defined in Subsection 4.4.1). The stabilization over the positive half-ball follows readily from the results in Subsection 4.3.6, especially Equation 4.7. In contrast, stabilization over the negative half-ball requires the combinatorial tools developed in Subsection 4.3.7.

4.2 Stabilization of Total Dimension

In this section, we prove Theorem 4.1 using immersed curve invariants.

Proof of Theorem 4.1. By Proposition 4.10, the pair (S^3, K_m) is obtained by gluing

$$(E_c, K) \qquad \text{and} \qquad (H'_{1/m}, \mu_c, -\lambda_c).$$

Let

$$\boldsymbol{\gamma} := \mathrm{hfk}(K) \qquad \text{and} \qquad \boldsymbol{l}_m := \widehat{\mathrm{HF}}(H'_{1/m}).$$

By Theorem 2.112,

$$\widehat{\mathrm{HFK}}(S^3, K_m) \cong \mathrm{HF}(\boldsymbol{l}_m, \bar{h}(\boldsymbol{\gamma})).$$

By Example 2.121, the immersed curve invariant $\boldsymbol{l}_m$ consists of a single loose Euclidean geodesic of slope m, which we denote by l_m. Write

$$\bar{h}(\boldsymbol{\gamma}) = \boldsymbol{\gamma}_1 \cup \cdots \cup \boldsymbol{\gamma}_n,$$

where

$$\boldsymbol{\gamma}_j = (\gamma_j, k_j, A_j) \qquad (j = 1, \ldots, n).$$

Using Lemma 2.102, we may assume that each γ_j is primitive. Since Floer homology splits over connected components, we have

$$\mathrm{HF}(\boldsymbol{l}_m, \bar{h}(\boldsymbol{\gamma})) = \bigoplus_{j=1}^{n} \mathrm{HF}(\boldsymbol{l}_m, \boldsymbol{\gamma}_j).$$

By Theorem 2.105,

$$\dim \mathrm{HF}(\boldsymbol{l}_m, \boldsymbol{\gamma}_j) = k_j \cdot i(l_m, \gamma_j).$$

Hence

$$\dim\big(\widehat{\mathrm{HFK}}(S^3, K_m)\big) = \sum_{j=1}^{n} k_j \cdot i(l_m, \gamma_j).$$

Thus it remains to understand the behaviour of the geometric intersection numbers $i(l_m, \gamma_j)$ as $m \to \infty$.

We use $(\widehat{\gamma}_j, \boldsymbol{c}^j)$ to denote the singular pegboard diagram of γ_j. We now apply Theorem 2.109. Since $\widehat{\gamma}_j$ consists of finitely many geodesic segments, it has only finitely many slopes. Therefore, for sufficiently large m, none of these slopes coincide with the slope of l_m. Moreover, the curve l_m is loose, i.e. it has no corners. Hence Theorem 2.109 gives

$$i(l_m, \gamma_j) = |l_m \cap \widehat{\gamma}_j|.$$

Recall that the punctured torus T is obtained from the marked torus $\widehat{T}$ by removing the basepoint z. As discussed earlier, $\widehat{\gamma}_j$ can be viewed as a collection of Euclidean geodesics in $\widehat{T}$ passing through z. We denote these geodesics by $L_1^j, \dots, L_{t_j}^j$. Then

$$|l_m \cap \widehat{\gamma}_j| = \sum_{i=1}^{t_j} |l_m \cap L_i^j|.$$

Suppose that L_i^j lifts to a line of slope p_i^j/q_i^j in the universal cover of $\widehat{T}$. We can assume that $q_i^j > 0$. Since l_m lifts to a line of slope m, we have

$$|l_m \cap L_i^j| = \left|\det \begin{pmatrix} 1 & m \\ q_i^j & p_i^j \end{pmatrix}\right| = |p_i^j - mq_i^j|.$$

For sufficiently large m, this becomes

$$|l_m \cap L_i^j| = mq_i^j - p_i^j.$$

Therefore, for sufficiently large m,

$$\begin{aligned} \dim\big(\widehat{\mathrm{HFK}}(S^3, K_m)\big) &= \sum_{j=1}^{n} k_j\, i(l_m, \gamma_j) \\ &= \sum_{j=1}^{n} k_j \left(\sum_{i=1}^{t_j} (mq_i^j - p_i^j) \right) \\ &= m \underbrace{\left(\sum_{j=1}^{n} \sum_{i=1}^{t_j} k_j q_i^j \right)}_{D} - \underbrace{\left(\sum_{j=1}^{n} \sum_{i=1}^{t_j} k_j p_i^j \right)}_{d}. \end{aligned}$$

This completes the proof. □

Remark 4.14. One might hope to use immersed curve invariants to prove stronger stabilization results. The main obstacle is the determination of absolute Maslov gradings.

To access these gradings, one needs an explicit chain complex carrying both differentials $\widehat{\partial}_{K_m}$ and $\widehat{\partial}$. The immersed curve invariant $\mathrm{hfk}(K)$ does not contain enough information to recover this structure. To phrase this in terms of Example 2.119, one needs access to both type A structures $\widehat{\mathrm{CFA}}(\mathcal{H}_Q, z, w)$ and $\widehat{\mathrm{CFA}}(\mathcal{H}_Q, z)$, just as in the computation of the bordered grading structure. The immersed curve invariant $\mathrm{hfk}(Q)$, however, only retains the information encoded by the former.

In the special case of twist families arising from cabling, there are ways to work around this by computing the immersed curve invariants of the knot complements. See Hanselman–Watson [HW23].

4.3 The Stabilization Framework

4.3.1 General structure of $\mathcal{X}_m^K$

Using the notation of Construction 4.9 and Notation 4.11, we now begin the analysis of the box tensor products $\mathcal{X}_m^K$ and $\mathcal{X}_m$.

Consider a basis for $\mathcal{S}$ (and hence for $\mathcal{K}$) as an $\mathcal{I}$-module:

$$B_K = B_K^0 \sqcup B_K^1.$$

We can assume that the basis comes from a doubly-pointed bordered Heegaard diagram. Example 2.121 gives a basis of $\mathcal{D}_m$,

$$B'_m := \{\eta, \xi_1, \dots, \xi_m\}.$$

Definition 4.15. Define $\boldsymbol{C}^m = C^\bullet \sqcup C_1^\circ \sqcup \cdots \sqcup C_m^\circ$, where

$$C^\bullet = \{x^0 \otimes \eta \mid x^0 \in B_K^0\},$$

$$C_i^\circ = \{x^1 \otimes \xi_i \mid x^1 \in B_K^1\}, \qquad i = 1, \dots, m.$$

Note that $\boldsymbol{C}^m$ is the standard basis of both $\mathcal{X}_m^K$ and $\mathcal{X}_m$.

Construction 4.16. We arrange the elements of $\boldsymbol{C}^m$ around a circle as follows. Considering the vertices of the decorated graph in Figure 2.40, we turn each vertex into a box. The box on top, associated to the vertex η, contains all elements of $C^\bullet$, labelled by elements of B_K^0. The box associated to ξ_i, for $i = 1, \dots, m$, contains the elements of C_i°, labelled by elements of B_K^1.

By abuse of notation, we also use $C^\bullet, C_1^\circ, \dots, C_m^\circ$ to denote the corresponding boxes. We refer to $C^\bullet$ as the *black box*, and to $C_1^\circ, \dots, C_m^\circ$ as the *white boxes*. A schematic of this arrangement is shown in Figure 4.3.

We give the black and white boxes the structure of a cyclic graph by taking the boxes as vertices and connecting neighbouring boxes in the arrangement of Figure 4.3. This is precisely the underlying graph of Figure 2.40. We denote this graph by $\boldsymbol{GC^m}$.

Note that $\mathcal{X}_m^K$ and $\mathcal{X}_m$ have the same underlying vector space and differ only in their differentials. Since we work over $\mathbb{F}_2$, linear algebra reduces to set combinatorics. This motivates the following notation.

Notation 4.17. We identify subsets of $\boldsymbol{C}^m$ with elements of $\mathcal{X}_m^K$ (and $\mathcal{X}_m$) via

$$I \subseteq \boldsymbol{C}^m \qquad \longleftrightarrow \qquad S_I := \sum_{\nu \in I} \nu.$$

We also represent chain complexes by directed graphs. For $\nu, \nu' \in \boldsymbol{C}^m$, we draw a directed edge from ν to ν' if and only if ν' appears as a summand of $\partial\nu$.

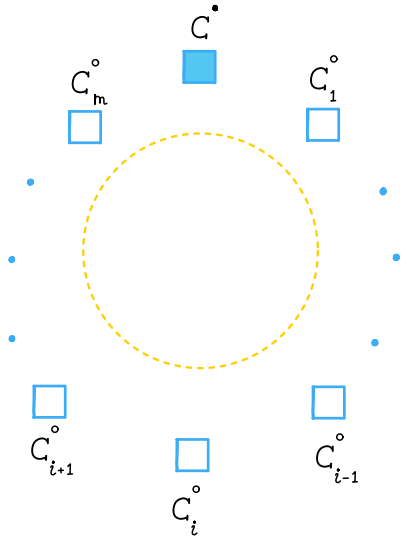


Figure 4.3: The circular arrangement of the boxes in $\boldsymbol{C}^m$.

Example 4.18. An explicit example is illustrated in Figures 4.6 and 4.7 for $m = 3$ and the Mazur pattern Q, using the computations of $\mathcal{K}$ and $\mathcal{S}$ from Example 2.119. The differentials are shown as green arrows.

Remark 4.19. The homology of $\mathcal{X}_3$ is $\widehat{\mathrm{HF}}(S^3) \cong \mathbb{F}_2$. In the example of Figures 4.6 and 4.7, it is generated by

$$x_0 \otimes \eta + x_1 \otimes \xi_1 + x_1 \otimes \xi_2 + x_1 \otimes \xi_3.$$

The circular arrangement of $\boldsymbol{C}^m$ suggests a natural relationship between the bases $\boldsymbol{C}^m$ and $\boldsymbol{C}^{m+1}$. To formalize this idea, we introduce the notions of *inclusion maps* and *nested sequence structures*, which will be used throughout the chapter.

Construction 4.20. We define the following two inclusions:

$$\Phi'_m : \boldsymbol{C}^m \hookrightarrow \boldsymbol{C}^{m+1}, \quad \begin{cases} \Phi'_m(x^0 \otimes \eta) := x^0 \otimes \eta & \text{for all } x^0 \otimes \eta \in C^\bullet, \\ \Phi'_m(x^1 \otimes \xi_i) := x^1 \otimes \xi_i & \text{for all } x^1 \otimes \xi_i \in C^\circ_i,\ i \in \{1, \dots, m\}. \end{cases} \tag{4.2}$$

$$\Phi_m : \boldsymbol{C}^m \hookrightarrow \boldsymbol{C}^{m+1}, \quad \begin{cases} \Phi_m(x^0 \otimes \eta) := x^0 \otimes \eta & \text{for all } x^0 \otimes \eta \in C^\bullet, \\ \Phi_m(x^1 \otimes \xi_i) := x^1 \otimes \xi_i & \text{for all } x^1 \otimes \xi_i \in C^\circ_i,\ i \le \frac{m}{2}, \\ \Phi_m(x^1 \otimes \xi_i) := x^1 \otimes \xi_{i+1} & \text{for all } x^1 \otimes \xi_i \in C^\circ_i,\ i > \frac{m}{2}. \end{cases} \tag{4.3}$$

Intuitively, the map Φ'_m corresponds to constructing $\boldsymbol{C}^{m+1}$ by adding a new white box after the existing white boxes. On the other hand, the map Φ_m corresponds to constructing $\boldsymbol{C}^{m+1}$ by inserting a new white box in the middle of the existing white boxes. These two constructions are illustrated in Figures 4.4 and 4.5.

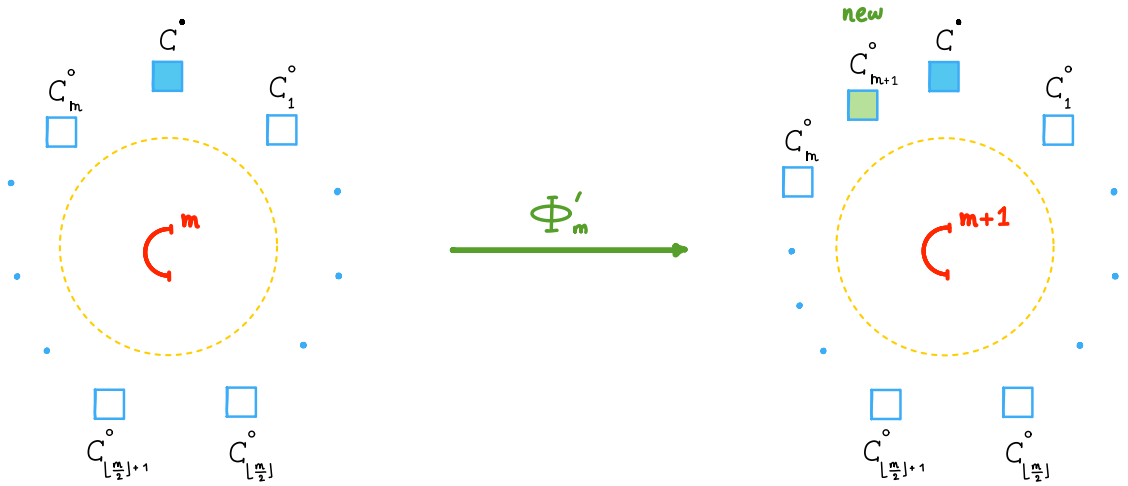


Figure 4.4: The inclusion Φ'_m

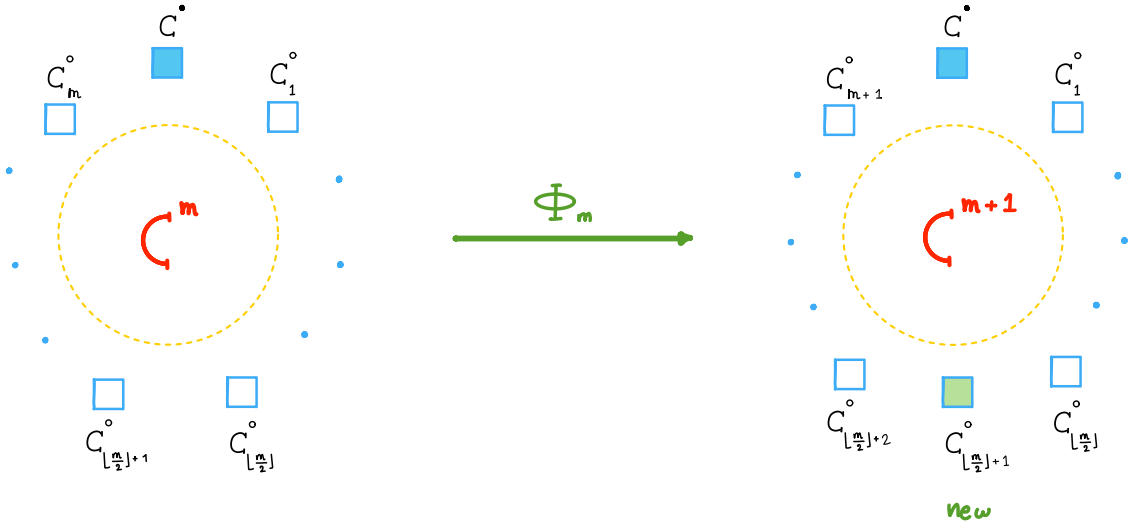


Figure 4.5: The inclusion Φ_m

Remark 4.21. The nested structures given by Φ_m and Φ'_m are reflected in our notation. Throughout the chapter, we often use the same symbols to denote corresponding basis elements and boxes in $\boldsymbol{C}^m$ for different values of m. The intended identification will

always be determined by one of the two nested structures above. Although this introduces a mild ambiguity, the relevant choice of inclusion map will be clear from the context in each instance.

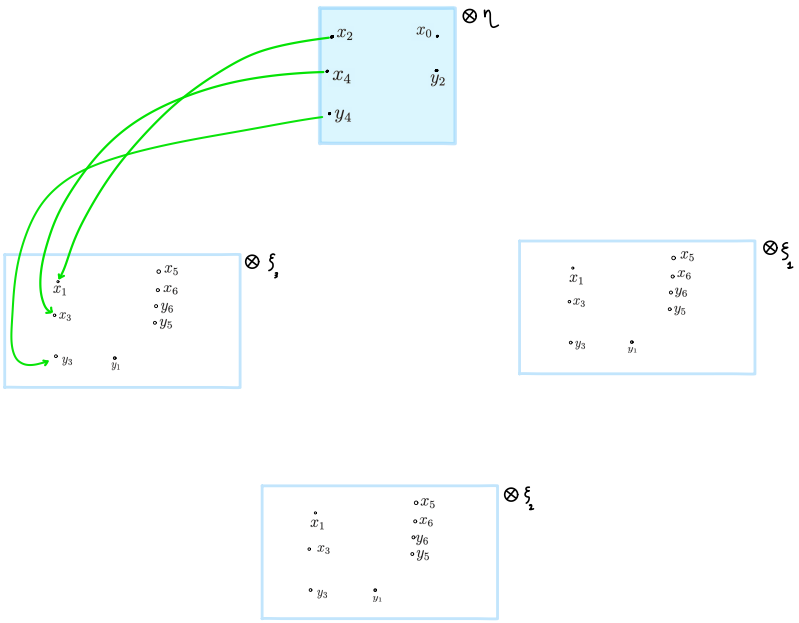


Figure 4.6: The complex $\widehat{\mathrm{CFA}}(\mathcal{H}_Q, z, w) \boxtimes \widehat{\mathrm{CFD}}(\mathcal{H}'_{\frac{1}{3}}, z')$

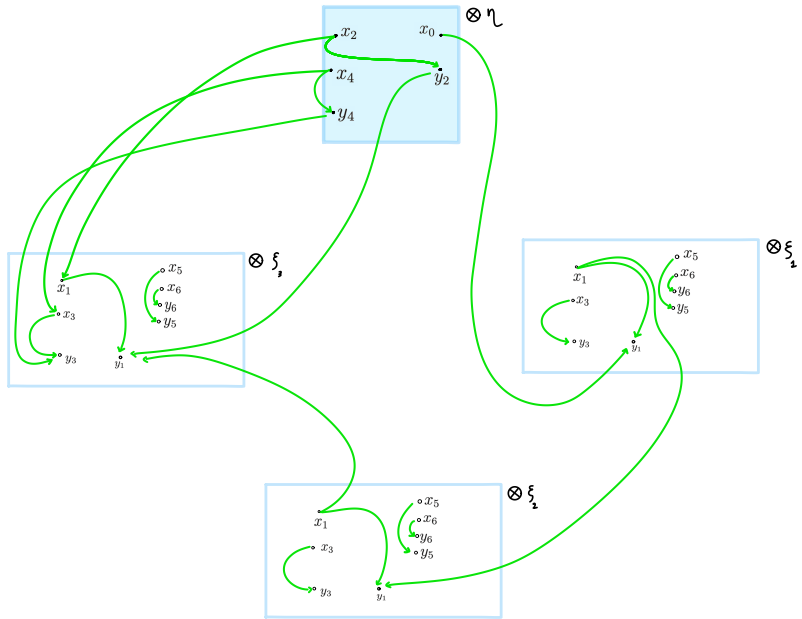


Figure 4.7: The complex $\widehat{\mathrm{CFA}}(\mathcal{H}_Q, z) \boxtimes \widehat{\mathrm{CFD}}(\mathcal{H}'_{\frac{1}{3}}, z')$

We now begin analyzing the differentials.

Proposition 4.22. *There are four types of directed edges in the graph-theoretic descriptions of the chain complexes $\mathcal{X}_m$ and $\mathcal{X}_m^K$, with respect to the standard basis $\boldsymbol{C}^m$:*

1. *The directed edges inside $C^\bullet$. For any $x^0, y^0 \in B_K^0$ such that y^0 appears in $m_1(x^0)$, we have a directed edge*

$$x^0 \otimes \eta \longrightarrow y^0 \otimes \eta.$$

2. *The directed edges originating from $C^\bullet$ and moving in the clockwise direction. For any $x^0 \in B_K^0$ and $x^1 \in B_K^1$ such that x^1 appears in*

$$m_{n+2}(x^0 \otimes \rho_3 \otimes \underbrace{\rho_{23} \otimes \cdots \otimes \rho_{23}}_{n})$$

for some $n \geq 0$, we have a directed edge

$$x^0 \otimes \eta \longrightarrow x^1 \otimes \xi_{n+1}.$$

3. *The directed edges originating from $C^\bullet$ and moving in the counterclockwise direction. For any $x^0 \in B_K^0$ and $x^1 \in B_K^1$ such that x^1 appears in $m_2(x^0 \otimes \rho_1)$, we have a directed edge*

$$x^0 \otimes \eta \longrightarrow x^1 \otimes \xi_m.$$

4. *The directed edges between white boxes, possibly moving in the clockwise direction. For any $x^1, y^1 \in B_K^1$ such that y^1 appears in*

$$m_{n+1}(x^1 \otimes \underbrace{\rho_{23} \otimes \cdots \otimes \rho_{23}}_{n})$$

for some $n \geq 0$, we have a family of directed edges

$$x^1 \otimes \xi_i \longrightarrow y^1 \otimes \xi_{i+n}$$

for $1 \leq i \leq m - n$.

Proof. This follows directly from the definition of the differential on the box tensor product in Equation (2.4) together with the description of the type D structure $\mathcal{D}_m$ from Example 2.121. □

Figure 4.8 depicts Proposition 4.22.

Remark 4.23. The type (4) edges are invariant under circular shift maps. This observation plays an important role in many of the stabilization arguments.

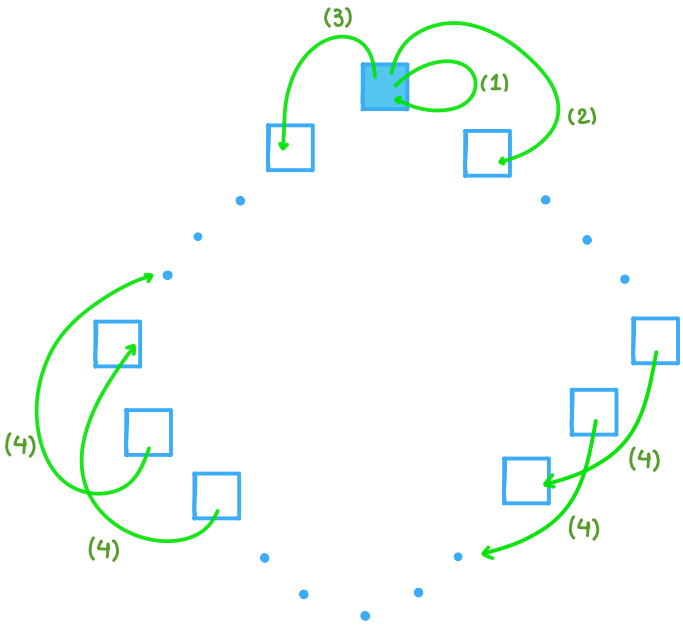


Figure 4.8: The decomposition of edges into the four types

Notation 4.24. Let ν, ν' be elements of the basis $\boldsymbol{C}^m$, viewed as vertices in the graph-theoretic description. We write

$$\nu \xrightarrow{[i]} \nu'$$

for an edge of type (i) from ν to ν'.

For $i \in \{1, 2, 3, 4\}$, let $E_i(\nu, \nu')$ denote the set of type (i) edges from ν to ν'. Similarly, let $E_i(\nu, \cdot)$ (resp. $E_i(\cdot, \nu)$) denote the set of type (i) edges starting (resp. ending) at ν. We then have

$$\partial(\nu) = \sum_{i=1}^{4} \sum_{\nu' \in \boldsymbol{C}^m} |E_i(\nu, \nu')| \cdot \nu' \pmod 2.$$

We also define

$$\partial_i(\nu) := \sum_{\nu' \in \boldsymbol{C}^m} |E_i(\nu, \nu')| \cdot \nu' \pmod 2,$$

so that

$$\partial = \partial_1 + \partial_2 + \partial_3 + \partial_4.$$

Note that the maps ∂_i are not necessarily differentials.

4.3.2 Gradings in $\mathcal{X}_m^K$

We now study the relative gradings of elements of the standard basis $\boldsymbol{C}^m$. We begin with a description of the subgroup $P(\boldsymbol{r}_0)$ associated with the grading of $\mathcal{K}$.

Lemma 4.25. *For any choice of reference generator $\boldsymbol{r}_0 \in \mathcal{K}$, the associated subgroup $P(\boldsymbol{r}_0)$ is generated by an element $(M; 0, 1; \omega) \in \widetilde{G}$, where $M \in \frac{\mathbb{Z}}{2}$ and $\omega = \mathrm{lk}(K, c)$.*

Proof. Since E_c is a solid torus, $P(\boldsymbol{r}_0)$ is cyclic. As explained in [HRW22], the Spin^c component of the generator of $P(\boldsymbol{r}_0)$ is determined by the topology of the bordered 3-manifold. In particular, it agrees with the corresponding computation for the Mazur pattern in Example 2.119, and hence is of the form $(0, \pm 1)$.

To elaborate further, as mentioned in Chapter 2, the subgroup $P(\boldsymbol{r}_0)$ is determined by the gradings of periodic domains. As proved in [LOT18, Lemma 4.20], periodic domains in $\mathcal{H}_K$ are in one-to-one correspondence with $H_2(E_c, \partial E_c)$. In fact, for any periodic domain B, one can construct a properly embedded surface T_B.

Since E_c is a solid torus, $H_2(E_c, \partial E_c) \cong \mathbb{Z}$ is generated by a meridian disk D. Let B be a periodic domain such that $[T_B] = [D]$. The *Spin*$^{\mathbb{C}}$ component of the corresponding element of $P(\boldsymbol{r}_0)$ is determined by the expansion of $[\partial D] \in H_1(\partial E_c)$ in terms of the parametrizing curves. Since

$$[\partial D] = 0 \cdot [\mu_c] - 1 \cdot [-\lambda_c],$$

the *Spin*$^{\mathbb{C}}$ component of the generator is $(0, -1)$.

The Alexander component of the grading is defined by $n_w(B) - n_z(B)$. Hence,

$$n_w(B) - n_z(B) = -(K \cdot T_B) = -\mathrm{lk}(K, c) = -\omega.$$

Therefore,

$$P(\boldsymbol{r}_0) = \langle (-M; 0, -1; -\omega) \rangle$$

for some $M \in \frac{\mathbb{Z}}{2}$. The result follows. □

To study relative gradings, we first fix notation for the bordered grading structures introduced in Chapter 2.

Notation 4.26. We denote the grading functions on $\mathcal{K}$ and $\mathcal{D}_m$ by

$$\mathrm{gr}_K : B_K \to P(\boldsymbol{r}_0) \backslash \widetilde{G}, \qquad \mathrm{gr}_m : B'_m \to G/P(\eta).$$

These induce a grading function on $\mathcal{X}_m^K$:

$$\mathrm{gr}_{K_m} : \boldsymbol{C}^m \to P(\boldsymbol{r}_0) \backslash \widetilde{G}/(P(\eta); 0).$$

Similarly, we denote by gr_S and gr_{S_m} the corresponding grading functions on $\mathcal{S}$ and $\mathcal{X}_m$. Note that gr_S (resp. gr_{S_m}) is obtained from gr_K (resp. gr_{K_m}) by forgetting the Alexander component.

We examine the behaviour of gr_{K_m} and the relative Maslov and Alexander gradings as $m \to \infty$ in Lemmas 4.28 and 4.29. Recall from Notation 4.12 that h_m and a_m denote the Maslov and Alexander gradings on $\mathcal{X}_m^K$.

Notation 4.27. Let

$$p_K := (M; 0, 1; \omega), \qquad p_m := (-\frac{m-1}{2}; 1, -m; 0).$$

By Lemma 4.25, p_K generates $P(\boldsymbol{r}_0)$. As shown in Example 2.121, p_m generates $P(\eta)$.

Lemma 4.28. *Let $n \in \mathbb{N}$. Consider two elements $\nu_1, \nu_2 \in \boldsymbol{C}^n$ such that, for every $m \geq n$, the double cosets $\mathrm{gr}_{K_m}(\nu_1)$ and $\mathrm{gr}_{K_m}(\nu_2)$ have representatives which do not depend on m. In other words, there exist elements*

$$g_i = (\alpha_i; \beta_i, \gamma_i; \delta_i) \in \widetilde{G}, \qquad i \in \{1, 2\},$$

such that, for every $m \geq n$, $\mathrm{gr}_{K_m}(\nu_i)$ admits g_i as a representative. Then

$$\begin{aligned} h_m(\nu_1) - h_m(\nu_2) = {} & ((\beta_2 - \beta_1)(M - \tfrac{1}{2} - \beta_1 - \beta_2)) \cdot m \\ & + (\gamma_2 - \gamma_1)M + (\alpha_1 - \alpha_2) + \tfrac{1}{2}(\beta_2 - \beta_1) \\ & - (\beta_2 - \beta_1)\gamma_1 - (\gamma_2 - \gamma_1)\beta_2, \\ a_m(\nu_1) - a_m(\nu_2) = {} & ((\beta_2 - \beta_1)\omega) \cdot m + (\delta_1 - \delta_2) - (\gamma_1 - \gamma_2)\omega. \end{aligned}$$

Proof. To compute the relative Maslov and Alexander gradings, we need to find a representative of the double coset $[g_1]$ whose Spin^c component is equal to (β_2, γ_2). Such a representative is

$$p_K^{\gamma_2 - \gamma_1 + (\beta_2 - \beta_1)m} \cdot g_1 \cdot p_m^{\beta_2 - \beta_1}.$$

Computing the Maslov and Alexander components of this representative yields the stated formulas. □

Lemma 4.29. *Let $x^1 \in B_K^1$, and let $(a; b, c; d) \in \widetilde{G}$ be a representative of the grading coset $\mathrm{gr}_K(x^1)$. Let i, j be positive integers. Then for any $m \geq \max\{i, j\}$ we have*

$$\begin{aligned} h_m(x^1 \otimes \xi_j) - h_m(x^1 \otimes \xi_i) &= (j - i)\Big(M - 2b - \frac{1}{2}\Big), \\ a_m(x^1 \otimes \xi_j) - a_m(x^1 \otimes \xi_i) &= (j - i)\omega. \end{aligned}$$

Proof. By the definition of the induced grading on the box tensor product,

$$\begin{aligned}\mathrm{gr}_{K_m}(x^1 \otimes \xi_i) &= [\mathrm{gr}_K(x^1) \cdot \mathrm{gr}_m(\xi_i)] \\ &= [(a - \tfrac{1}{2} - \tfrac{2i-1}{2}b - \tfrac{c}{2};\, b + \tfrac{1}{2},\, c - \tfrac{2i-1}{2};\, d)].\end{aligned}$$

Similarly,

$$\mathrm{gr}_{K_m}(x^1 \otimes \xi_j) = [(a - \tfrac{1}{2} - \tfrac{2j-1}{2}b - \tfrac{c}{2};\, b + \tfrac{1}{2},\, c - \tfrac{2j-1}{2};\, d)].$$

The $Spin^{\mathbb{C}}$ components of these representatives have the same first entry, and we can find representatives with equal $Spin^{\mathbb{C}}$ components by multiplying by p_K^{i-j}. Thus

$$\mathrm{gr}_{K_m}(x^1 \otimes \xi_i) = \left[p_K^{i-j} \cdot (a - \tfrac{1}{2} - \tfrac{2i-1}{2}b - \tfrac{c}{2};\, b + \tfrac{1}{2},\, c - \tfrac{2i-1}{2};\, d)\right].$$

Computing the Maslov and Alexander components of this representative yields the stated formulas. □

As seen in Lemma 4.29, the second component of a representative of the grading coset $\mathrm{gr}_K(x^1)$ plays a role in the computation of relative Maslov gradings. The following lemma records a useful fact about this component.

Lemma 4.30. *For $j \in \{1, 2\}$, let $x_j \in B_K^1$, and let*

$$(a_j; b_j, c_j; d_j) \in \widetilde{G}$$

be a representative of the grading coset $\mathrm{gr}_K(x_j)$. Then $b_1 - b_2 \in \mathbb{Z}$.

Proof. We use the graph-theoretic description of type A invariants. Let Γ_A be the decorated graph representing $\mathcal{K}$. For any (not necessarily directed) path P in Γ_A connecting x_1 and x_2, let P_l denote the subset of edges of P labelled by

$$l \in \{\varnothing, 1, 2, 3, 32, 21, 321\}.$$

By the grading rules of Table 2.22, it suffices to show that

$$\frac{|P_1| + |P_2| + |P_3| + |P_{321}|}{2} \in \mathbb{Z}.$$

Note that P_2 is the subset of edges in P going from a $\circ$-labelled vertex to a $\bullet$-labelled vertex. On the other hand, $P_1 \cup P_3 \cup P_{321}$ is the subset of edges in P going from a $\bullet$-labelled vertex to a $\circ$-labelled vertex. Since both x_1 and x_2 belong to B_K^1, they are $\circ$-labelled. Therefore,

$$|P_2| = |P_1| + |P_3| + |P_{321}|.$$

It follows that

$$\frac{|P_1| + |P_2| + |P_3| + |P_{321}|}{2} = \frac{2|P_2|}{2} = |P_2| \in \mathbb{Z},$$

which completes the proof. □

A similar argument gives the following corollary.

Corollary 4.31. *Let $x_1 \in B^1_K$ and $x_0 \in B^0_K$. Let*

$$(a_j; b_j, c_j; d_j) \in \widetilde{G}$$

be a representative of the grading coset $\mathrm{gr}_K(x_j)$. Then $b_1 - b_0 + \frac{1}{2} \in \mathbb{Z}$.

4.3.3 Minimality and Homogeneity of Cycles

We now begin the study of the homology of $\mathcal{X}^K_m$. Before proceeding, we fix a standing assumption that will be used throughout the remainder of the chapter.

Assumption 4.32. Without loss of generality, we assume that $\mathrm{lk}(K, c) = \omega < 0$. Indeed, reversing the orientation of K changes the sign of $\mathrm{lk}(K, c)$ while leaving the knot Floer homology groups unchanged. The assumption that $\mathrm{lk}(K, c) \neq 0$ is the essential part, and it also appears in the statements of our main results.

Recall that, by Notation 4.17, elements of $\mathcal{X}^K_m$ are identified with subsets of $\boldsymbol{C}^m$.

Definition 4.33. Let $I \subseteq \boldsymbol{C}^m$. We say that S_I is *homogeneous* if all elements $\nu \in I$ have the same Maslov and Alexander gradings. Equivalently, all elements of I have the same gr_{K_m} value.

Definition 4.34. Let S_I be a cycle. We say that S_I is *minimal* if it is minimal among cycles with respect to the inclusion partial order on $2^{\boldsymbol{C}^m}$. Equivalently, for every non-empty subset $I' \subsetneq I$, we have $\partial S_{I'} \neq 0$.

Remark 4.35. Similarly, we define homogeneity and minimality for homology classes. A homology class is called *homogeneous* (resp. *minimal*) if it admits a homogeneous (resp. minimal) cycle representative.

We extend these properties to a (linear) basis of the homology as well.

Definition 4.36. Let

$$\boldsymbol{S}^m = \{S_{I_1}, \dots, S_{I_{d_m}}\}$$

be a collection of cycles. By abuse of notation, we call $\boldsymbol{S}^m$ a *basis for the homology* if the corresponding homology classes form a basis of the homology vector space.

Definition 4.37. Let $\boldsymbol{S}^m = \{S_{I_1}, \dots, S_{I_{d_m}}\}$ be a basis for the homology. We say that $\boldsymbol{S}^m$ is *homogeneous* if each S_{I_j} is homogeneous.

Definition 4.38. Let $\boldsymbol{S}^m = \{S_{I_1}, \ldots, S_{I_{d_m}}\}$ be a basis for the homology. We say that $\boldsymbol{S}^m$ is *minimal* if it is minimal among all homology bases with respect to the inclusion partial order. Equivalently, for every $j \in \{1, \ldots, d_m\}$ and every subset $I'_j \subsetneq I_j$:

$$\boldsymbol{S}^m \setminus \{S_{I_j}\} \cup \{S_{I'_j}\}$$

is not a basis for the homology.

Remark 4.39. Unless otherwise stated, we will always assume that homology bases are chosen to be minimal. As we will see in the next lemma, every minimal basis is homogeneous.

Lemma 4.40. *Let $\boldsymbol{S}^m = \{S_{I_1}, \ldots, S_{I_{d_m}}\}$ be a minimal basis for homology. Then each S_{I_j} is a minimal cycle. Consequently, every minimal basis is homogeneous.*

Proof. We prove the statement by contradiction. Assume there exist k and a subset $I'_k \subsetneq I_k$ such that $S_{I'_k}$ is a cycle. Let

$$I''_k := I_k \setminus I'_k.$$

Note that $S_{I_k} = S_{I'_k} + S_{I''_k}$, and hence

$$\partial S_{I_k} = \partial S_{I''_k} = 0.$$

The cycle $S_{I'_k}$ determines a homology class $[S_{I'_k}]$, which can be written as a linear combination of the basis elements. Therefore, there exists $J \subseteq \{1, \ldots, d_m\}$ such that

$$[S_{I'_k}] = \sum_{j \in J} [S_{I_j}].$$

We now consider two cases.

1. If $k \notin J$, then

$$[S_{I_k}] - [S_{I''_k}] = [S_{I'_k}] = \sum_{j \in J} [S_{I_j}].$$

   It follows that

$$\boldsymbol{S}^m \setminus \{S_{I_k}\} \cup \{S_{I''_k}\}$$

   is again a basis for homology, contradicting the minimality of $\boldsymbol{S}^m$.

2. If $k \in J$, then

$$[S_{I'_k}] - [S_{I_k}] = \sum_{j \in J \setminus \{k\}} [S_{I_j}].$$

   It follows that

$$\boldsymbol{S}^m \setminus \{S_{I_k}\} \cup \{S_{I'_k}\}$$

   is a basis for homology, again contradicting the minimality of $\boldsymbol{S}^m$.

This contradiction shows that every S_{I_j} is minimal. Finally, note that a minimal cycle must be homogeneous, since otherwise any nonzero homogeneous component would be a cycle supported on a proper subset, contradicting minimality. This completes the proof. □

4.3.4 Length and Size of Minimal Cycles

We now turn to the study of minimal cycles. We begin by introducing two constants associated to the type A structure $\mathcal{K}$.

Notation 4.41. Let $N_A := |B_K|$. Let L_A be the smallest integer such that

$$m_{n+1} : \mathcal{K} \otimes \mathcal{A}^n \to \mathcal{K}$$

vanishes for all $n \geq L_A$.

Remark 4.42. For L_A to be well-defined, we need $\mathcal{K}$ to be bounded (as defined in Definition 2.76). Similar to the convention for type D structures (see Remark 2.73), we may always assume that the type A structures $\mathcal{K}$ and $\mathcal{S}$ are bounded. This assumption is not restrictive, since every type A structure is homotopy equivalent to a bounded one.

We now introduce a distance function and a partial order on elements of $\boldsymbol{C}^m$ leading to the introduction of length for subsets of $\boldsymbol{C}^m$.

Notation 4.43. The cyclic graph $\boldsymbol{GC^m}$ induces a metric on the set of boxes. We denote the corresponding distance function on $\boldsymbol{C}^m$ by dist. In other words, the distance between two elements of $\boldsymbol{C}^m$ is defined to be the distance between the boxes containing them. We also use dist_H to denote the associated Hausdorff distance on subsets of $\boldsymbol{C}^m$.

Notation 4.44. Ordering the boxes $C^\bullet, C_1^\circ, \ldots, C_m^\circ$ clockwise starting from the black box induces a partial order on $\boldsymbol{C}^m$. We call this the *clockwise order* and denote it by $\prec$.

Definition 4.45. Let $I \subseteq \boldsymbol{C}^m$, and let

$$\nu_0, \ldots, \nu_{n-1}$$

be the elements of I ordered with respect to the clockwise order. Throughout, the indices are taken modulo n. We define the *length* of S_I by

$$\mathrm{len}(S_I) = \min\Big\{ \sum_{i=0}^{n-2} \mathrm{dist}(\nu_{t+i}, \nu_{t+i+1}) \;\Big|\; t \in \{0, \ldots, n-1\} \Big\}. \tag{4.4}$$

Note that $\mathrm{diam}(I) \leq \mathrm{len}(S_I)$.

The following series of lemmas forms the core of this subsection and establishes the fundamental constraints satisfied by minimal cycles. More precisely, Lemma 4.46 bounds the size of a homogeneous element. Lemma 4.47 relates the differentials to the clockwise order and the distance function. Lemma 4.48 bounds the gaps between consecutive elements of a minimal cycle. Finally, Lemma 4.49 gives a uniform upper bound on the length of minimal cycles.

Lemma 4.46. *Assume that S_I is homogeneous. Then $|I| \leq N_A$.*

Proof. We prove this by contradiction. Assume that $|I| > N_A$. Then

$$\left| I \cap \left(\bigcup_{i=1}^{m} C_i^{\circ} \right) \right| > N_A - |C^{\bullet}| = |B_K^1|.$$

Therefore, there exist $x^1 \in B_K^1$ and integers $1 \leq j_1 < j_2 \leq m$ such that

$$x^1 \otimes \xi_{j_1}, \quad x^1 \otimes \xi_{j_2} \in I.$$

By Lemma 4.29,

$$a_m(x^1 \otimes \xi_{j_2}) - a_m(x^1 \otimes \xi_{j_1}) = (j_2 - j_1) \cdot \omega \neq 0.$$

Here, recall from Lemma 4.25 that $\omega = \mathrm{lk}(K, c)$. This contradicts the homogeneity assumption. □

Lemma 4.47. *Assume that there is a directed edge $e = (\nu, \nu')$ in the graph-theoretic description of $\mathcal{X}_m^K$. Then*

$$\nu \preceq \nu' \quad \textit{and} \quad \textit{dist}(\nu, \nu') < L_A.$$

We call $\textit{dist}(\nu, \nu')$ the length *of the edge e.*

Proof. This follows from Proposition 4.22. All edges that do not originate from $C^{\bullet}$ move in the clockwise direction, and they neither lead to nor pass over $C^{\bullet}$. Therefore, $\nu \preceq \nu'$. Furthermore, by the definition of L_A, the directed edge (ν, ν') arises from a map m_n with $n \leq L_A$. Hence $\mathrm{dist}(\nu, \nu') \leq L_A - 1 < L_A$. □

Lemma 4.48. *Assume that S_I is a minimal cycle. Order the elements of I with respect to the clockwise order and denote them by*

$$\nu_0, \ldots, \nu_{n-1}.$$

Then there exists at most one $i \in \{0, \ldots, n-1\}$ such that $\textit{dist}(\nu_i, \nu_{i+1}) > L_A$.

Proof. We prove this by contradiction. Assume that there exist

$$0 \leq j_1 < j_2 \leq n-1$$

such that $\operatorname{dist}(\nu_{j_l}, \nu_{j_l+1}) > L_A$ for $l = 1, 2$.

Now consider the set

$$I' = \{\nu_{j_1+1}, \ldots, \nu_{j_2}\}.$$

This is clearly an interval in the ordered set I. By the definition of the clockwise order and the assumption on $\operatorname{dist}(\nu_{j_1}, \nu_{j_1+1})$, we have $\nu_{j_1+1} \in C_l^\circ$ for some $l > L_A$. Hence

$$I' \cap C^\bullet = \varnothing.$$

Let I'' denote the complement of I' in I. Since $I = I' \sqcup I''$, we have

$$S_I = S_{I'} + S_{I''}$$

and therefore $\partial S_I = \partial S_{I'} + \partial S_{I''}$. Let

$$\partial S_{I'} = S_{J'}, \qquad \partial S_{I''} = S_{J''}.$$

Since S_I is a cycle,

$$S_{J'} + S_{J''} = \partial S_I = 0.$$

We claim that $J' \cap J'' = \varnothing$. If so, there can be no cancellations in the sum $S_{J'} + S_{J''}$, and hence

$$S_{J'} + S_{J''} = S_{J' \sqcup J''}.$$

Therefore $J' = J'' = \varnothing$, so both $S_{I'}$ and $S_{I''}$ are cycles. This contradicts the minimality of S_I.

To prove that $J' \cap J'' = \varnothing$, we show that edges originating from I' and I'' cannot have the same endpoint. Decompose

$$I'' = I_1'' \sqcup I_2'',$$

where I_1'' consists of the elements of I'' appearing before I' in the clockwise order, and $I_2'' := I'' \setminus I_1''$. Note that both I_1'' and I_2'' are intervals.

First consider I_2''. Suppose that there are edges e' and e'' originating from I' and I_2'', respectively, with the same endpoint. By Lemma 4.47, both edges move in the clockwise direction. Consequently, e' must pass over the gap between the last element of I', namely ν_{j_2}, and the first element of I_2'', namely ν_{j_2+1}. This contradicts Lemma 4.47, since the length of e' would be greater than $\operatorname{dist}(\nu_{j_2}, \nu_{j_2+1}) > L_A$.

The same argument applies to clockwise edges originating from I_1''. It remains to consider counterclockwise edges originating from $I_1'' \cap C^\bullet$. By Proposition 4.22, such an edge ends in C_m° and has length one. Therefore the gap between the endpoint of such an edge and the last element of I' has size at least

$$\operatorname{dist}(\nu_{j_2}, \nu_{j_2+1}) - 1,$$

which is still greater than the maximal edge length allowed by Lemma 4.47. Hence these edges also cannot share an endpoint with an edge originating from I'.

This proves that $J' \cap J'' = \varnothing$ and completes the proof. □

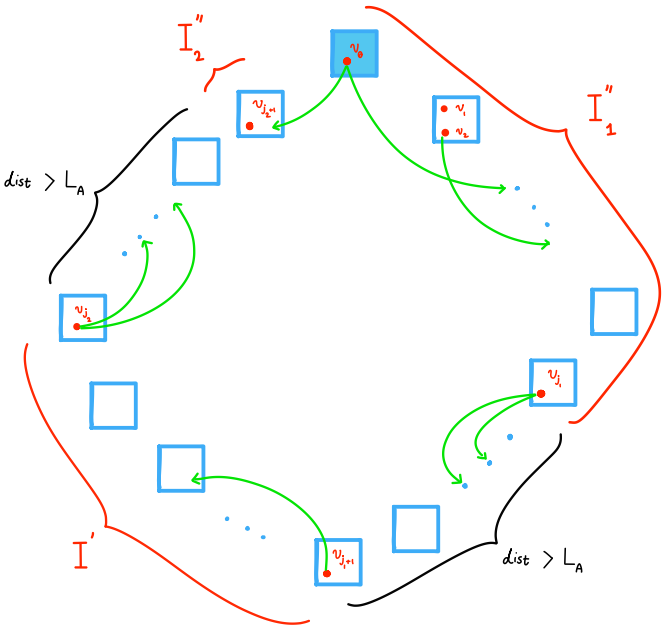


Figure 4.9: Illustration of the proof of Lemma 4.48

Lemma 4.49. *Assume that S_I is a minimal cycle. Then $\operatorname{len}(S_I) \leq N_A L_A$.*

Proof. Order the elements of I as $\nu_0, \ldots, \nu_{n-1}$ with respect to the clockwise order. Choose $t \in \{0, \ldots, n-1\}$ such that

$$\operatorname{dist}(\nu_t, \nu_{t+1}) = \max\{\operatorname{dist}(\nu_j, \nu_{j+1}) \mid 0 \leq j \leq n-1\}.$$

Then

$$\operatorname{len}(S_I) \leq \sum_{i=0}^{n-2} \operatorname{dist}(\nu_{t+1+i}, \nu_{t+2+i}) \leq N_A L_A,$$

where the final inequality follows from Lemmas 4.46 and 4.48. □

4.3.5 Length and Size of Minimal Relations

To understand the homology, we also need to examine some of the properties of boundaries. To this end, we define the space of relations.

Definition 4.50. Let $I_1, \dots, I_r \subseteq \boldsymbol{C}^m$, and assume that $S_{I_1}, \dots, S_{I_r}$ are minimal cycles. A *relation* between these cycles is a subset $R \subseteq \boldsymbol{C}^m$ such that

$$\partial S_R = S_{I_1} + \cdots + S_{I_r}.$$

Definition 4.51. A relation is called *homogeneous* if all elements of $R \cup I_1 \cup \cdots \cup I_r$ have the same grading. It is called *minimal* if there do not exist a non-empty subset $R' \subsetneq R$ and a subset $T \subseteq \{1, \dots, r\}$ such that

$$\partial S_{R'} = \sum_{t \in T} S_{I_t}.$$

The following lemmas establish for minimal relations the same type of control obtained previously for minimal cycles.

Lemma 4.52. *Assume that we have a minimal relation*

$$\partial S_R = S_{I_1} + \cdots + S_{I_r}.$$

Order the elements of R with respect to the clockwise order and denote them by $\nu_0, \dots, \nu_{n-1}$. Then there exists at most one $i \in \{0, \dots, n-1\}$ such that

$$dist(\nu_i, \nu_{i+1}) > (N_A + 1)L_A.$$

Proof. The proof is similar to Lemma 4.48, and we use the same notation. Suppose, for a contradiction, that there exist $0 \leq j_1 < j_2 \leq n-1$ such that

$$\operatorname{dist}(\nu_{j_l}, \nu_{j_l+1}) > (N_A + 1)L_A$$

for $l = 1, 2$. Consider the subsets

$$R' = \{\nu_{j_1+1}, \dots, \nu_{j_2}\}, \qquad R'' = R \setminus R'.$$

Let

$$\partial S_{R'} = S_{J'}, \qquad \partial S_{R''} = S_{J''}.$$

As in the proof of Lemma 4.48, we have $J' \cap J'' = \varnothing$. Moreover, the same argument gives

$$\operatorname{dist}_H(J', J'') \geq N_A L_A.$$

Now consider the two gaps between J' and J''. We treat the case where neither gap contains $C^\bullet$; the other case is analogous, and the bound can only improve since differentials neither lead to nor pass over $C^\bullet$.

Let

$$E = \{\kappa_0, \kappa_1, \dots, \kappa_l\} \qquad \text{and} \qquad E' = \{\kappa'_0, \kappa'_1, \dots, \kappa'_{l'}\}$$

be the ordered sets of elements of $I_1 \cup \dots \cup I_r$ lying inside or at the endpoints of the first and second gaps, respectively. Since J' and J'' are subsets of $I_1 \cup \dots \cup I_r$, both E and E' are non-empty.

We claim that at least one of l and l' is strictly greater than N_A. Otherwise, if $l, l' \leq N_A$, then the inequality $\operatorname{dist}_H(J', J'') \geq N_A L_A$ implies, by the pigeonhole principle, that there exist indices

$$0 \leq o \leq l-1, \qquad 0 \leq o' \leq l'-1$$

such that

$$\operatorname{dist}(\kappa_o, \kappa_{o+1}) > L_A, \qquad \operatorname{dist}(\kappa'_{o'}, \kappa'_{o'+1}) > L_A.$$

By Lemma 4.48, for each $t \in \{1, \dots, r\}$, all elements of I_t lie on one side of the gap (κ_o, κ_{o+1}). The same holds for the gap $(\kappa'_{o'}, \kappa'_{o'+1})$. Define

$$T' = \left\{ t \in \{1, \dots, r\} \;\middle|\; \text{all elements of } I_t \text{ lie after } (\kappa_o, \kappa_{o+1}) \text{ and before } (\kappa'_{o'}, \kappa'_{o'+1}) \right\},$$

and let $T'' = \{1, \dots, r\} \setminus T'$. Then

$$\left(\bigcup_{t' \in T'} I_{t'} \right) \cap J'' = \varnothing, \qquad \left(\bigcup_{t'' \in T''} I_{t''} \right) \cap J' = \varnothing.$$

Hence

$$\partial S_{R'} = S_{J'} = \sum_{t' \in T'} S_{I_{t'}}, \qquad \partial S_{R''} = S_{J''} = \sum_{t'' \in T''} S_{I_{t''}}.$$

This contradicts the minimality of the relation.

Therefore at least one of l and l' is strictly greater than N_A. But then one of the sets E or E' contains more than N_A elements of the homogeneous relation, contradicting Lemma 4.46. This completes the proof. □

Corollary 4.53. *There exists an integer L_B, independent of m, such that for every minimal relation*

$$\partial S_R = S_{I_1} + \dots + S_{I_r},$$

we have

$$\mathit{len}(S_R) \leq L_B.$$

Proof. Set

$$L_B := N_A(N_A+1)L_A.$$

Since every minimal relation is homogeneous, Lemma 4.46 implies that $|R| \leq N_A$. The result now follows from Lemma 4.52, exactly as in the proof of Lemma 4.49. □

We have already established a bound for $\mathrm{len}(S_R)$, and hence also for $\mathrm{diam}(R)$. However, a relation also involves the supports of the minimal cycles $S_{I_1}, \ldots, S_{I_r}$. Using the same type of gap argument, one obtains the following analogous bound.

Corollary 4.54. *Consider a minimal relation $\partial S_R = S_{I_1} + \cdots + S_{I_r}$. Then*

$$\mathit{diam}\left(\bigcup_{1\leq t\leq r} I_t\right) \leq N_A(N_A+1)L_A = L_B.$$

Consequently,

$$\mathit{diam}\left(R \cup \bigcup_{1\leq t\leq r} I_t\right) \leq 2L_B + L_A.$$

We do not provide the full proof of Corollary 4.54. The argument follows the same strategy as the proof of Corollary 4.53: one analyses large gaps between consecutive elements of $\bigcup_{1\leq t\leq r} I_t$. The minimality of the relation shows that such gaps must be filled by elements of R. Lemma 4.52, together with the bound $|R| \leq N_A$, then gives a uniform bound on the size of such gaps. A direct implementation of this argument gives

$$\mathrm{diam}\left(\bigcup_{1\leq t\leq r} I_t\right) \leq N_A^2(N_A+1)L_A = N_A L_B.$$

However, this estimate can be sharpened to the bound stated above.

The second inequality follows from the first one, Corollary 4.53, and Lemma 4.47: since

$$\partial S_R = S_{I_1} + \cdots + S_{I_r},$$

there is at least one edge from an element of R to an element of $\bigcup_{1\leq t\leq r} I_t$, and such an edge has length less than L_A.

Remark 4.55. Lemma 4.40 shows that the space of cycles is generated by minimal cycles. A similar observation holds for relations.

Combining these two facts shows that the space of boundaries is generated by the boundaries of minimal relations. Hence, minimal cycles and minimal relations completely determine the homology.

4.3.6 Homology of $\mathcal{X}_m$: Algebra

As indicated in the overview, we begin by considering the type A structure

$$\mathcal{S}_\infty = \widehat{\mathrm{CFA}}(\mathcal{H}_\infty, z),$$

described in Remark 2.120. This module is generated by a single element x_0 and is depicted in Figure 4.10. Moreover, the grading coset $\mathrm{gr}(x_0)$ has representative $(0;0,0)$, and the subgroup $P(x_0)$ is generated by $(-\frac{1}{2};0,1)$.



Figure 4.10: Decorated graph representing $\mathcal{S}_\infty$

As in Subsection 4.3.2, let $\mathbf{r}_0$ be the reference generator in $\mathcal{S}$. By Theorem 2.95, there exist an isomorphism of G-sets

$$\phi : P(x_0)\backslash G \longrightarrow P(\mathbf{r}_0)\backslash G$$

and an $\mathcal{A}_\infty$ homotopy equivalence

$$\widehat{f} : \mathcal{S}_\infty \longrightarrow \mathcal{S}.$$

As discussed in Proposition 2.96, $\widehat{f}$ induces a chain homotopy equivalence

$$\widehat{f} \boxtimes \mathbb{I}_m : \mathcal{S}_\infty \boxtimes \mathcal{D}_m \longrightarrow \mathcal{X}_m.$$

Note that the chain complex $\mathcal{S}_\infty \boxtimes \mathcal{D}_m$ is generated by a single element $x_0 \otimes \eta$ and has trivial differential. Therefore,

$$(\widehat{f} \boxtimes \mathbb{I}_m)(x_0 \otimes \eta)$$

represents the unique non-trivial homology class in $\mathcal{X}_m$. The following lemma studies the properties of this representative.

Lemma 4.56. *Let $\widehat{f}$ be the $\mathcal{A}_\infty$ homotopy equivalence defined above. Consider a subset $\widehat{I}_m \subseteq \boldsymbol{C}^m$ such that*

$$(\widehat{f} \boxtimes \mathbb{I}_m)(x_0 \otimes \eta) = S_{\widehat{I}_m}.$$

Then $S_{\widehat{I}_m}$ is homogeneous with respect to gr_{S_m}. Moreover, every element of $\widehat{I}_m$ has grading double coset admitting a representative $(\alpha;\beta,\gamma)$ satisfying

$$2\beta = M + \frac{1}{2}.$$

Proof. By the box-tensor formula,

$$(\widehat{f} \boxtimes \mathbb{I}_m)(x_0 \otimes \eta) = \sum_{k\geq 0} (\widehat{f}_{k+1} \otimes \mathbb{I}_m)\big(x_0 \otimes \delta_k(\eta)\big).$$

Consider an element $y \otimes \sigma \in \widehat{I}_m$. Then for some $k \geq 0$ we have

$$y \text{ is a summand of } \widehat{f}_{k+1}(x_0 \otimes a_1 \otimes \cdots \otimes a_k),$$

and

$$a_1 \otimes \cdots \otimes a_k \otimes \sigma \text{ is a summand of } \delta_k(\eta).$$

By Theorem 2.95, we obtain

$$\mathrm{gr}_S(y) = \phi(\mathrm{gr}(x_0))\ \mathrm{gr}(a_1) \cdots \mathrm{gr}(a_k)\lambda^k.$$

Using the grading rule for type D structures, we also have

$$\mathrm{gr}_m(\sigma) = \mathrm{gr}(a_k)^{-1} \cdots \mathrm{gr}(a_1)^{-1}\lambda^{-k}\, \mathrm{gr}_m(\eta).$$

Recall that here $\lambda = (1; 0, 0)$. Therefore,

$$\mathrm{gr}_{S_m}(y \otimes \sigma) = \mathrm{gr}_S(y)\ \mathrm{gr}_m(\sigma) = \phi(\mathrm{gr}(x_0))\ \mathrm{gr}_m(\eta).$$

In particular, all elements of $\widehat{I}_m$ have the same gr_{S_m}-coset, so $S_{\widehat{I}_m}$ is homogeneous.

For the second statement, write

$$\phi(\mathrm{gr}(x_0)) = [(\alpha; \beta, \gamma)].$$

Since $(-\frac{1}{2}; 0, 1) \in P(x_0)$, we have

$$\mathrm{gr}(x_0) \cdot (-\tfrac{1}{2}; 0, 1) = \mathrm{gr}(x_0).$$

Applying ϕ and using that it is an isomorphism of G-sets, we get

$$\phi(\mathrm{gr}(x_0)) \cdot (-\tfrac{1}{2}; 0, 1) = \phi(\mathrm{gr}(x_0)).$$

Equivalently,

$$(\alpha; \beta, \gamma) \cdot (-\tfrac{1}{2}; 0, 1) \cdot (-\alpha; -\beta; -\gamma) \in P(\mathbf{r}_0).$$

A direct computation gives

$$(2\beta - \tfrac{1}{2}; 0, 1) \in P(\mathbf{r}_0).$$

By Lemma 4.25, $P(\mathbf{r}_0)$ is generated by $(M; 0, 1)$. Hence, $(2\beta - \frac{1}{2}; 0, 1)$ must be a power of $(M; 0, 1)$. Comparing the third components, this power must be 1, and therefore

$$2\beta - \tfrac{1}{2} = M.$$

Finally, since $\mathrm{gr}_m(\eta) = (0; 0, 0)$, every element of $\widehat{I}_m$ has grading double coset represented by $(\alpha; \beta, \gamma)$ with $2\beta = M + \frac{1}{2}$. □

The previous lemma used the additional grading information provided by Theorem 2.95. The next lemma is intentionally phrased for arbitrary $\mathcal{A}_\infty$ homotopy equivalences. It shows that every resulting representative contains a distinguished element whose position is independent of m. In particular, this element is stabilized with respect to either of the nested sequence structures arising from Construction 4.20.

Lemma 4.57. *Consider an arbitrary $\mathcal{A}_\infty$ homotopy equivalence*

$$f : \mathcal{S}_\infty \longrightarrow \mathcal{S},$$

where $\mathcal{S}_\infty$ is defined in Notation 4.13. Let $I_m \subseteq \boldsymbol{C}^m$ be such that

$$(f \boxtimes \mathbb{I}_m)(x_0 \otimes \eta) = S_{I_m}.$$

Then there exists an element $\nu_0 \in C^\bullet$ such that $\nu_0 \in I_m$ for every $m \in \mathbb{N}$.

Proof. Define $J_k \subseteq B_K$ for each integer $k \geq 0$ by

$$f_{k+1}\big(x_0 \otimes \underbrace{\rho_3 \otimes \rho_{23} \otimes \cdots \otimes \rho_{23}}_{k}\big) = \sum_{y \in J_k} y.$$

Also, let $J' \subseteq B_K$ be defined by

$$f_2(x_0 \otimes \rho_1) = \sum_{y \in J'} y.$$

Then for any $m \in \mathbb{N}$,

$$f \boxtimes \mathbb{I}_m(x_0 \otimes \eta) = \sum_{y \in J_0} y \otimes \eta + \sum_{k=1}^{m} \sum_{y \in J_k} y \otimes \xi_k + \sum_{y \in J'} y \otimes \xi_m.$$

It therefore suffices to prove that $J_0 \neq \varnothing$, or equivalently that $f_1(x_0) \neq 0$. Assume, for a contradiction, that $f_1(x_0) = 0$.

Since f is an $\mathcal{A}_\infty$ homotopy equivalence, there exists an $\mathcal{A}_\infty$ homomorphism

$$g : \mathcal{S} \longrightarrow \mathcal{S}_\infty$$

such that $g \circ f$ is homotopic to the identity. Hence $g \circ f - \mathbb{I}_{\mathcal{S}_\infty}$ is a null-homotopic $\mathcal{A}_\infty$ homomorphism. By the definition of a null-homotopic $\mathcal{A}_\infty$ homomorphism, there exists a map

$$H_1 : \mathcal{S}_\infty \longrightarrow \mathcal{S}_\infty$$

such that

$$(g \circ f)_1(x_0) - x_0 = m_1(H_1(x_0)) + H_1(m_1(x_0)).$$

Since m_1 vanishes on $\mathcal{S}_\infty$ (see Remark 2.120), the right-hand side is zero. On the other hand,

$$(g \circ f)_1(x_0) - x_0 = g_1(f_1(x_0)) - x_0 = -x_0,$$

where the last equality follows from the assumption that $f_1(x_0) = 0$. This contradiction proves that $f_1(x_0) \neq 0$, and hence $J_0 \neq \varnothing$.

Now, choose an element $y_0 \in J_0$ and set

$$\nu_0 := y_0 \otimes \eta.$$

By the above expansion, $\nu_0 \in I_m \cap C^\bullet$ for every $m \in \mathbb{N}$. This completes the proof. □

So far, we have used $\mathcal{A}_\infty$ homotopy equivalences to construct representatives of the unique non-trivial homology class in $\mathcal{X}_m$. The following proposition shows that, conversely, every such representative arises from an $\mathcal{A}_\infty$ homotopy equivalence.

Proposition 4.58. *Let $m_0 \in \mathbb{N}$ and let $S_I \in \mathcal{X}_{m_0}$ be a cycle representing the unique non-trivial homology class. Then there exists an $\mathcal{A}_\infty$ homotopy equivalence*

$$f_{S_I} : \mathcal{S}_\infty \longrightarrow \mathcal{S}$$

such that

$$(f_{S_I} \boxtimes \mathbb{I}_{m_0})(x_0 \otimes \eta) = S_I.$$

Sketch of proof. Fix an arbitrary $\mathcal{A}_\infty$ homotopy equivalence $f' : \mathcal{S}_\infty \longrightarrow \mathcal{S}$. Let $I' \subseteq \boldsymbol{C}^{m_0}$ be such that

$$(f' \boxtimes \mathbb{I}_{m_0})(x_0 \otimes \eta) = S_{I'}.$$

Since both S_I and $S_{I'}$ represent the unique non-trivial homology class of $\mathcal{X}_{m_0}$, there exists a subset $R \subseteq \boldsymbol{C}^{m_0}$ such that

$$\partial S_R = S_I - S_{I'}.$$

We construct a null-homotopic $\mathcal{A}_\infty$ homomorphism $g : \mathcal{S}_\infty \longrightarrow \mathcal{S}$ such that

$$(g \boxtimes \mathbb{I}_{m_0})(x_0 \otimes \eta) = S_I - S_{I'}.$$

Setting $f_{S_I} = f' + g$, we will have the desired result. To this end, consider the collection of maps

$$H_{n+1} : \mathcal{S}_\infty \otimes \mathcal{A}^{\otimes n} \longrightarrow \mathcal{S}$$

defined to be zero except in the following cases:

$$H_1(x_0) = S_{R \cap C^\bullet},$$

and, for $n \geq 1$,

$$H_{n+1}(x_0 \otimes \rho_3 \otimes \underbrace{\rho_{23} \otimes \cdots \otimes \rho_{23}}_{n-1}) = S_{R \cap C_n^\circ}.$$

One can verify that these maps satisfy the required $\mathcal{A}_\infty$ relations. Consequently, they define a null-homotopic homomorphism g with the desired property. □

As a corollary, we can now extend the grading constraint from Lemma 4.56 to the stabilized elements obtained in Lemma 4.57.

Corollary 4.59. *Use the notation of Lemma 4.57, and let*

$$\nu_0 \in C^\bullet \cap I_m$$

be the stabilized element obtained there. Then, for every $m \in \mathbb{N}$, the grading $\mathrm{gr}_{S_m}(\nu_0)$ admits a representative $(\alpha_0; \beta_0, \gamma_0)$ such that

$$2\beta_0 = M + \frac{1}{2}.$$

Proof. We also use the notation of Lemma 4.56. If

$$I_m \cap \widehat{I}_m \cap C^\bullet \neq \varnothing,$$

then the result follows immediately from Lemma 4.56. Hence assume that this intersection is empty. Since both S_{I_m} and $S_{\widehat{I}_m}$ represent the unique non-trivial homology class in $H_*(\mathcal{X}_m)$, there exists a relation

$$\partial S_{R_m} = S_{I_m} - S_{\widehat{I}_m}.$$

Using the graph-theoretic description of $\mathcal{X}_m$, we claim that there exist elements

$$\nu_{R_m} \in R_m \cap C^\bullet, \qquad \nu_0 \in I_m \cap C^\bullet, \qquad \widehat{\nu}_0 \in \widehat{I}_m \cap C^\bullet$$

and directed edges

$$\nu_{R_m} \longrightarrow \nu_0, \qquad \nu_{R_m} \longrightarrow \widehat{\nu}_0.$$

Indeed, otherwise consider the subset

$$R'_m := \{\nu \in R_m \cap C^\bullet \mid \text{there exists } \nu' \in I_m \cap C^\bullet \text{ and an edge } \nu \to \nu'\}.$$

Then

$$\partial S_{R'_m} = S_{I_m \cap C^\bullet} + S_{R''_m}$$

for some subset $R''_m \subseteq \boldsymbol{C}^m \setminus C^\bullet$. Therefore

$$S_{I_m} - \partial S_{R'_m} = S_{I_m \setminus C^\bullet} - S_{R''_m}$$

is a cycle supported outside the black box. Moreover,

$$[S_{I_m} - \partial S_{R'_m}] = [S_{I_m}],$$

so it still represents the unique non-trivial homology class. This contradicts the combination of Lemma 4.57 and Proposition 4.58. This proves the claim.

It remains to compare gradings. More generally, we can show that every edge between elements of $C^\bullet$ preserves the Spin^c component of the grading. By Proposition 4.22, any such edge is of type (1), and hence comes from the m_1 map on $\mathcal{S}$. Therefore, by Construction 2.92, the two endpoints have grading double cosets admitting representatives with the same Spin^c component.

Applying this to the two edges above, we see that ν_0 and $\widehat{\nu}_0$ (in combination with Lemma 4.56) completes the proof. □

Another consequence of Proposition 4.58 is that any generating cycle at a fixed level m_0 can be extended to a family of generating cycles across all values of m. This family is close to being nested: it decomposes into one part that stabilizes under the inclusions Φ_m and another part that evolves under the inclusions Φ'_m, as described below.

Construction 4.60. Let $m_0 \in \mathbb{Z}_{>0}$, and let $S_I \in \mathcal{X}_{m_0}$ be a cycle representing the unique non-trivial homology class. Consider the $\mathcal{A}_\infty$ homotopy equivalence f_{S_I} introduced in Proposition 4.58. For every $m \geq m_0$, define a cycle $S_{I_m} \in \mathcal{X}_m$ by

$$(f_{S_I} \boxtimes \mathbb{I}_m)(x_0 \otimes \eta) = S_{I_m}.$$

In particular, $S_{I_{m_0}} = S_I$. We call the family

$$\{S_{I_m}\}_{m \geq m_0}$$

the *growing family* associated to S_I.

Set $f = f_{S_I}$. Define subsets

$$J', J_0, J_1, J_2, \cdots \subseteq B_K$$

as in the proof of Lemma 4.57. Then

$$S_{I_m} = \sum_{y\in J_0} y\otimes\eta + \sum_{k=1}^{m}\sum_{y\in J_k} y\otimes\xi_k + \sum_{y\in J'} y\otimes\xi_m. \tag{4.5}$$

Motivated by this expansion, we decompose S_{I_m} as

$$S_{I''_m} := \sum_{y\in J_0} y\otimes\eta + \sum_{k=1}^{m}\sum_{y\in J_k} y\otimes\xi_k, \qquad S_{I'_m} := \sum_{y\in J'} y\otimes\xi_m. \tag{4.6}$$

Thus

$$S_{I_m} = S_{I''_m} + S_{I'_m}.$$

The two summands interact differently with the inclusions

$$\Phi_m, \Phi'_m : \boldsymbol{C}^m \hookrightarrow \boldsymbol{C}^{m+1}.$$

Indeed,

$$S_{I''_{m+1}} = S_{\Phi'_m(I''_m)} + \sum_{y\in J_{m+1}} y\otimes\xi_{m+1}, \qquad S_{I'_{m+1}} = S_{\Phi_m(I'_m)}. \tag{4.7}$$

Equivalently, if $NI''_{m+1}\subseteq C^\circ_{m+1}$ is defined by

$$S_{NI''_{m+1}} := \sum_{y\in J_{m+1}} y\otimes\xi_{m+1},$$

then

$$I''_{m+1} = \Phi'_m(I''_m)\sqcup NI''_{m+1}, \qquad I'_m = \Phi_m(I'_m).$$

We refer to $S_{I''_m}$ as the *nested component* of S_{I_m} and to $S_{I'_m}$ as the *stable component*. Similarly, we call $S_{NI''_{m+1}}$ the *new contribution* to the nested component when passing from m to $m+1$.

We conclude this subsection by establishing an algebraic relation between the stable component $S_{I'_m}$ and the black-box contribution $S_{I^\bullet}$ defined below. The relation also involves the differential component ∂_3 introduced in Notation 4.24.

Proposition 4.61. *Let*

$$S_{I^\bullet} := S_{I_m\cap C^\bullet} = \sum_{y\in J_0} y\otimes\eta.$$

Then

$$\partial_3(S_{I^\bullet}) + \partial(S_{I'_m}) = 0.$$

Proof. By Proposition 4.22, the desired identity is equivalent to

$$\sum_{y \in J_0} m_2(y \otimes \rho_1) \otimes \xi_m + \sum_{y \in J'} m_1(y) \otimes \xi_m = 0. \tag{4.8}$$

Recall that

$$f_1(x_0) = \sum_{y \in J_0} y, \qquad f_2(x_0 \otimes \rho_1) = \sum_{y \in J'} y. \tag{4.9}$$

The $\mathcal{A}_\infty$-homomorphism relations give

$$m_2(f_1(x_0) \otimes \rho_1) + m_1(f_2(x_0 \otimes \rho_1)) + f_2(m_1(x_0) \otimes \rho_1) + f_1(m_2(x_0 \otimes \rho_1)) = 0.$$

Since $m_1(x_0) = 0$ and $m_2(x_0 \otimes \rho_1) = 0$ in $\mathcal{S}_\infty$, we obtain

$$m_2(f_1(x_0) \otimes \rho_1) + m_1(f_2(x_0 \otimes \rho_1)) = 0.$$

Substituting Equation 4.9 yields Equation 4.8. □

4.3.7 Homology of $\mathcal{X}_m$: Combinatorics

In this subsection, we introduce several combinatorial notions associated to the graph-theoretic description of $\mathcal{X}_m$. These notions provide a convenient way to compare gradings of basis elements. We start by defining an alternating path between elements of a subset $I \subseteq \boldsymbol{C}^m$.

Definition 4.62. Let $I \subseteq \boldsymbol{C}^m$. An *alternating path* between two elements $\nu, \nu' \in I$ consists of a sequence $\nu_0, \ldots, \nu_k$ of elements of I satisfying the following conditions:

- $\nu_0 = \nu$ and $\nu_k = \nu'$.
- For every $i \in \{1, \ldots, k\}$, there exists an element $\nu_i' \in \boldsymbol{C}^m$ and directed edges

$$\nu_i \longrightarrow \nu_i', \qquad \nu_{i-1} \longrightarrow \nu_i'.$$

An alternating path is called *good* if none of the above edges are of type (3) in the terminology of Proposition 4.22.

Figure 4.11 illustrates an alternating path.

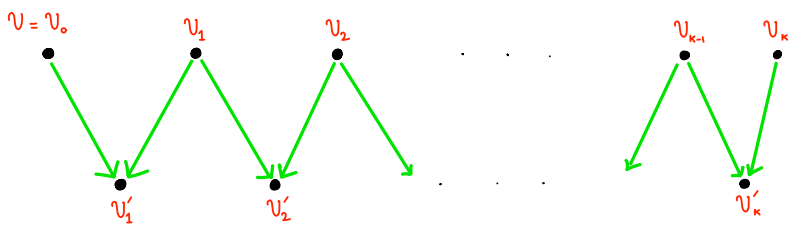


Figure 4.11: An alternating path

Definition 4.63. Let $I \subseteq \boldsymbol{C}^m$. We define an equivalence relation $\sim$ on I by

$$\nu \sim \nu' \quad \Longleftrightarrow \quad \text{there exists a good alternating path between } \nu \text{ and } \nu'.$$

The relevance of this equivalence relation comes from the following simple observation.

Lemma 4.64. *For any two elements $\nu, \nu' \in I$, we have*

$$\nu \sim \nu' \quad \Longrightarrow \quad h_m(\nu) = h_m(\nu').$$

Proof. The differential in $\mathcal{X}_m$ decreases the Maslov grading by one. The statement therefore follows immediately by induction on the length of a good alternating path connecting ν and ν'. □

The following proposition describes the structure of the equivalence classes of $\sim$ inside the nested component of a growing family of generators. This fact will play an important role in Section 4.5. Let

$$S_{I_m} = S_{I''_m} + S_{I'_m}$$

be the decomposition of a growing family of generators, as in Construction 4.60.

Proposition 4.65. *Consider the equivalence relation $\sim$ on the subset $I''_m \subseteq \boldsymbol{C}^m$. Let $\nu \in I''_m$ and set*

$$T = [\nu]_\sim \subseteq I''_m.$$

Then at least one of the following holds:

- $T \cap C^\bullet \neq \varnothing$;
- S_T *is a nullhomologous cycle.*

Furthermore, if S_{I_m} is minimal, then the second alternative can occur only if

$$T \cap I'_m \neq \varnothing.$$

Proof. By Proposition 4.61, we have

$$(\partial_1 + \partial_2 + \partial_4)(S_{I''_m}) = 0.$$

By definition of the equivalence relation, if

$$[\nu]_\sim \neq [\nu']_\sim,$$

then no element of $\boldsymbol{C}^m$ can occur simultaneously in

$$(\partial_1 + \partial_2 + \partial_4)(S_{[\nu]_\sim})$$

and

$$(\partial_1 + \partial_2 + \partial_4)(S_{[\nu']_\sim}).$$

In other words, cancellations between different equivalence classes are impossible. Hence,

$$(\partial_1 + \partial_2 + \partial_4)(S_T) = 0.$$

Now assume that

$$T \cap C^\bullet = \varnothing.$$

By Proposition 4.22, type (3) edges originate in the black box. Hence

$$\partial_3(S_T) = 0,$$

and therefore

$$\partial(S_T) = (\partial_1 + \partial_2 + \partial_4)(S_T) = 0.$$

Thus S_T is a cycle.

Since S_T is supported outside the black box, Lemma 4.57 together with Proposition 4.58 implies that S_T is nullhomologous. This proves the first statement.

For the second statement, assume that S_{I_m} is minimal and that

$$T \cap I'_m = \varnothing.$$

Then no element of S_T is cancelled in the decomposition

$$S_{I_m} = S_{I''_m} + S_{I'_m},$$

and hence

$$T \subseteq I_m.$$

Since S_T is a nullhomologous cycle, this contradicts the minimality of S_{I_m}. Therefore the second alternative can occur only when

$$T \cap I'_m \neq \varnothing.$$

□

4.4 Extremal Knot Floer Homology

In this section, all references to homology, cycles, boundaries, and grading-preserving differentials are understood to be with respect to the knot Floer complex $\mathcal{X}_m^K$. We use the framework developed in Section 4.3 to prove Theorem 4.3.

4.4.1 Closed Balls $B_t^\bullet$ and Shift maps $R_\pm$

Notation 4.66. For a non-negative integer t, let $B_t^\bullet$ denote the closed metric ball of radius t centered at the black box $C^\bullet$. Explicitly,

$$B_t^\bullet = \{C_{m-t+1}^\circ, \dots, C_m^\circ, C^\bullet, C_1^\circ, \dots, C_t^\circ\}.$$

By abuse of notation, we also use $B_t^\bullet$ to denote the set of all basis elements contained in these boxes. We further define the positive and negative half-balls by

$$B_{[0,t]}^\bullet := \{C^\bullet, C_1^\circ, \dots, C_t^\circ\}, \qquad B_{[-t,0]}^\bullet := \{C_{m-t+1}^\circ, \dots, C_m^\circ, C^\bullet\}.$$

Notation 4.67. We define the shift maps

$$R_+ : \boldsymbol{C}^m \setminus \{C^\bullet, C_m^\circ\} \longrightarrow \boldsymbol{C}^m, \qquad R_- : \boldsymbol{C}^m \setminus \{C^\bullet, C_1^\circ\} \longrightarrow \boldsymbol{C}^m$$

by

$$R_+(x^1 \otimes \xi_i) = x^1 \otimes \xi_{i+1}, \qquad R_-(x^1 \otimes \xi_i) = x^1 \otimes \xi_{i-1}.$$

Since these maps are injective, we also regard them as maps on subsets of $\boldsymbol{C}^m$.

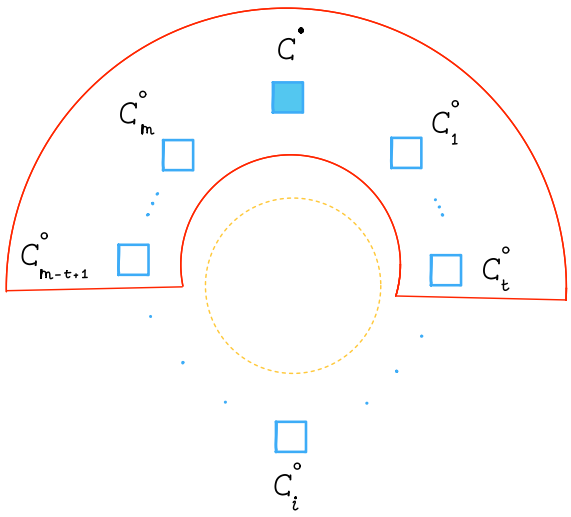


Figure 4.12: The closed metric ball $B_t^\bullet$

The following proposition shows that, away from the black box, the differential is compatible with the shift maps. Recall definition of L_A from Notation 4.41.

Proposition 4.68. *Let $I \subseteq \boldsymbol{C}^m \setminus B^\bullet_{L_A+1}$. If $\partial S_I = S_J$, then*

$$\partial S_{R_\pm(I)} = S_{R_\pm(J)}.$$

Proof. By Lemma 4.47, every edge has length at most L_A. Since $I \subseteq \boldsymbol{C}^m \setminus B^\bullet_{L_A+1}$, it follows that $J \subseteq \boldsymbol{C}^m \setminus B^\bullet_1$. In particular, $R_\pm(J)$ is well-defined.

By Proposition 4.22, all differentials originating outside $B^\bullet_{L_A+1}$ are of type (4). These differentials are preserved by the shift maps, and therefore

$$\partial S_{R_\pm(I)} = S_{R_\pm(J)}.$$

□

The shift maps not only preserve differentials, but also preserve non-trivial homology classes sufficiently far from the black box. Recall the constant L_B from Corollary 4.53.

Proposition 4.69. *Let*

$$I \subseteq \boldsymbol{C}^m \setminus B^\bullet_{L_B+2L_A+1}$$

and assume that S_I is a minimal cycle representing a non-trivial homology class. Then $[S_{R_\pm(I)}] \neq 0$.

Proof. We prove the statement for R_+; the proof for R_- is identical. By Proposition 4.68, the element $S_{R_+(I)}$ is a cycle. Moreover, Proposition 4.68 and the minimality of S_I imply that $S_{R_+(I)}$ is also minimal.

Assume, for a contradiction, that $[S_{R_+(I)}] = 0$. Since $S_{R_+(I)}$ is minimal, there exists a minimal relation

$$\partial S_T = S_{R_+(I)}.$$

By Corollary 4.53,

$$\operatorname{diam}(T) \leq \operatorname{len}(S_T) \leq L_B.$$

Furthermore, Lemma 4.47 gives

$$\operatorname{dist}_H(T, R_+(I)) < L_A.$$

Combining these inequalities with the assumption

$$I \subseteq \boldsymbol{C}^m \setminus B^\bullet_{L_B+2L_A+1},$$

we conclude that

$$T \subseteq \boldsymbol{C}^m \setminus B^\bullet_{L_A+1}.$$

Hence $R_-(T)$ is well-defined. Applying Proposition 4.68 again, we obtain

$$\partial S_{R_-(T)} = S_{R_-(R_+(I))} = S_I.$$

This contradicts the non-triviality assumption on $[S_I]$ and finishes the proof. □

We conclude this subsection with a remark on the stabilization of closed balls and half-balls under the nested sequence structures introduced in Construction 4.20.

Remark 4.70. Let L be a fixed positive integer. For sufficiently large m, the closed ball $B_L^\bullet \subseteq \boldsymbol{C}^m$ stabilizes under the inclusions

$$\Phi_m : \boldsymbol{C}^m \hookrightarrow \boldsymbol{C}^{m+1}.$$

More precisely, once $m > 2L$, the map Φ_m restricts to a bijection on $B_L^\bullet$. Hence, for all sufficiently large m, we may regard $B_L^\bullet$ as a fixed subset of the complexes $\boldsymbol{C}^m$.

The same statement holds for the half-balls $B_{[0,L]}^\bullet$ and $B_{[-L,0]}^\bullet$. Furthermore, the positive half-ball $B_{[0,L]}^\bullet$ also stabilizes under the second nested sequence structure induced by the inclusions Φ'_m once $m > L$.

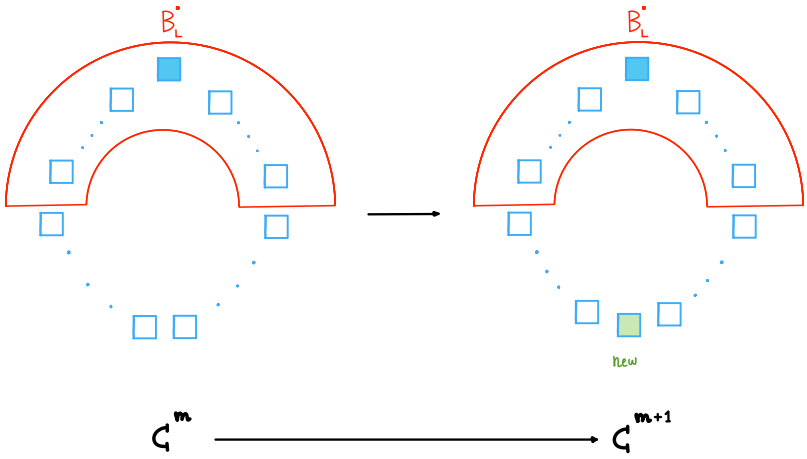


Figure 4.13: The nested sequence structure induced by the inclusions Φ_m. For $m > 2L$, the closed ball $B_L^\bullet$ is unaffected by the addition of a new white box.

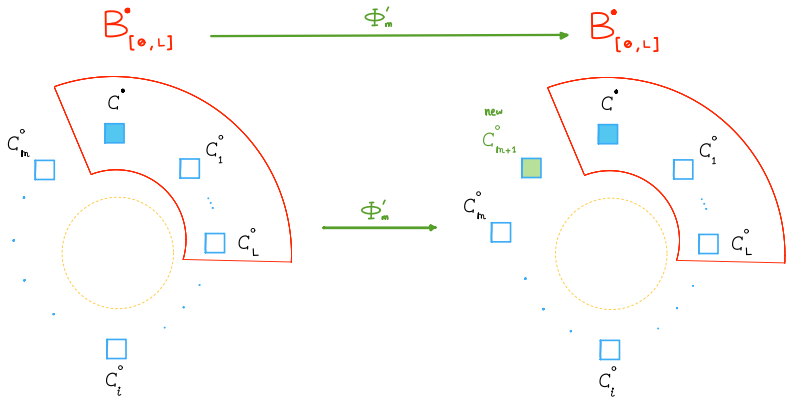


Figure 4.14: The half-ball $B_{[0,L]}^\bullet$ and its stabilization under the inclusions Φ'_m.

4.4.2 Extremal knot Floer homologies fall in a closed ball

Recall Assumption 4.32. Throughout this section, we assume that $\mathrm{lk}(K, c) = \omega < 0$. As discussed previously, the pairing theorem gives

$$\widehat{\mathrm{CFK}}(S^3, K_m) \simeq \mathcal{X}_m^K.$$

It helps to consider a notation for extremal knot Floer homologies.

Notation 4.71. For $\ell \in \mathbb{N}$, let

$$\mathcal{E}_m^\ell := \bigoplus_{j=0}^{\ell-1} \widehat{\mathrm{HFK}}(K_m, -g(K_m) + j) \subseteq \widehat{\mathrm{HFK}}(K_m) = \mathrm{H}_*(X_m^K).$$

The following constant will be used throughout this subsection.

Notation 4.72. We define $L_S := L_B + N_A L_A + 2L_A + 1$.

The next lemma shows that representatives of the extremal knot Floer homologies are supported inside a fixed closed ball.

Lemma 4.73. *Fix $\ell \in \mathbb{N}$. For all sufficiently large m, let S_I be a minimal cycle such that*

$$[S_I] \in \mathcal{E}_m^\ell \qquad \text{and} \qquad [S_I] \neq 0.$$

Then $I \subseteq B^\bullet_{L_S+\ell-1}$.

Proof. Assume, for a contradiction, that there exists $\nu \in I$ such that

$$\mathrm{dist}_H(\{\nu\}, C^\bullet) \geq L_S + \ell.$$

By Lemma 4.49, $\mathrm{diam}(I) \leq \mathrm{len}(S_I) \leq N_A L_A$. Hence

$$\mathrm{dist}_H(I, C^\bullet) \geq L_S + \ell - N_A L_A,$$

and therefore

$$I \subseteq \boldsymbol{C}^m \setminus B^\bullet_{L_B+2L_A+\ell+1}.$$

Consequently, the iterated shifts

$$R_+(I),\ R_+^2(I),\ \ldots,\ R_+^\ell(I)$$

are all contained in $\boldsymbol{C}^m \setminus B^\bullet_{L_B+2L_A+1}$. By Propositions 4.68 and 4.69, the classes

$$[S_{R_+(I)}],\ [S_{R_+^2(I)}],\ \ldots,\ [S_{R_+^\ell(I)}]$$

are all non-trivial. Since m is sufficiently large, we may apply Lemma 4.29 to obtain

$$a_m(S_{R^i_+(I)}) = a_m(S_I) + i\omega \qquad \text{for } 1 \le i \le \ell.$$

Since $[S_I] \in \mathcal{E}^\ell_m$, we have

$$a_m(S_I) \le -g(K_m) + \ell - 1.$$

Therefore

$$a_m(S_{R^\ell_+(I)}) \le -g(K_m) + \ell - 1 + \ell\omega.$$

Since $\omega \le -1$, it follows that

$$a_m(S_{R^\ell_+(I)}) < -g(K_m).$$

This contradicts the fact that knot Floer homology vanishes in Alexander gradings below $-g(K_m)$. □

Lemma 4.73 shows that the extremal knot Floer homologies are represented by cycles supported in a fixed closed ball. Combined with the stabilization properties of Remark 4.70, this reduces the proof of Theorem 4.3 to understanding the stabilized data contained in that ball.

4.4.3 Proof of Theorem 4.3

Lemma 4.73 motivates the following definitions.

Notation 4.74. Let $\mathcal{C}_{L_S+\ell-1}$ denote the set of cycles $S_I \in \mathcal{X}^K_m$ such that $I \subseteq B^\bullet_{L_S+\ell-1}$. Let

$$\mathcal{H}_{L_S+\ell-1} \subseteq \mathrm{H}_*(\mathcal{X}^K_m)$$

denote the subspace represented by cycles in $\mathcal{C}_{L_S+\ell-1}$.

The notation introduced above suppresses the dependence on m. The following lemma shows that this is justified for sufficiently large m.

Lemma 4.75. *Assume that*

$$m > 2L, \qquad L := (L_S + \ell - 1) + L_B + 2L_A.$$

Then the vector spaces

$$\mathcal{C}_{L_S+\ell-1} \qquad \text{and} \qquad \mathcal{H}_{L_S+\ell-1}$$

stabilize under the nested sequence structure induced by the inclusions $\{\Phi_m\}$.

Proof. Since $m > 2L$, Remark 4.70 allows us to identify the balls

$$B^\bullet_{L_S+\ell-1} \subseteq B^\bullet_L \subseteq \boldsymbol{C}^m$$

across all sufficiently large values of m.

By Lemma 4.47, the differential of every element of $B^\bullet_{L_S+\ell-1}$ is supported inside

$$B^\bullet_{L_S+\ell-1+L_A} \subseteq B^\bullet_L.$$

Furthermore, Proposition 4.22 implies that these differentials are independent of m. Hence the space $\mathcal{C}_{L_S+\ell-1}$ stabilizes.

Now consider a minimal relation

$$\partial S_R = S_{I_1} + \cdots + S_{I_r},$$

where

$$S_{I_1}, \ldots, S_{I_r} \in \mathcal{C}_{L_S+\ell-1}.$$

By Corollary 4.53 and Lemma 4.47,

$$R \subseteq B^\bullet_{(L_S+\ell-1)+L_B+L_A} \subseteq B^\bullet_{L-L_A}.$$

Therefore the differentials of all elements involved in such relations are also contained in the stabilized ball $B^\bullet_L$. Consequently, the space of minimal relations among cycles in $\mathcal{C}_{L_S+\ell-1}$ stabilizes.

By Remark 4.55, minimal cycles and minimal relations determine the homology. Since both stabilize, it follows that the homology space $\mathcal{H}_{L_S+\ell-1}$ also stabilizes. □

We now examine the relative Alexander gradings of homogeneous classes in $\mathcal{H}_{L_S+\ell-1}$. By Corollary 2.124, every element of the stabilized ball $B^\bullet_L$ has a grading double coset admitting a representative independent of m. Consequently, Lemma 4.28 implies that the relative Alexander grading between any two homogeneous classes in $\mathcal{H}_{L_S+\ell-1}$ is a linear function of m.

Therefore, for sufficiently large m, the signs of all relative Alexander gradings stabilize. In particular, we may choose a homogeneous class

$$H_{\min} \in \mathcal{H}_{L_S+\ell-1}$$

such that

$$a_m(H) - a_m(H_{\min}) \geq 0$$

for every homogeneous class $H \in \mathcal{H}_{L_S+\ell-1}$.

Notation 4.76. Let

$$\mathcal{H}^{\ell}_{L_S+\ell-1} \subseteq \mathcal{H}_{L_S+\ell-1}$$

denote the subspace generated by homogeneous classes H satisfying

$$a_m(H) - a_m(H_{\min}) \leq \ell - 1$$

for sufficiently large m.

Again by Lemma 4.28, we know that

$$a_m(H) - a_m(H_{\min})$$

stabilizes to a linear function of m. Hence, $\mathcal{H}^{\ell}_{L_S+\ell-1}$ consists precisely of those homogeneous classes H for which this linear function has slope zero and whose eventual value is less than ℓ. Consequently, $\mathcal{H}^{\ell}_{L_S+\ell-1}$ is independent of m for sufficiently large m, and therefore defines a stabilized subspace. Moreover, the relative Alexander grading on $\mathcal{H}^{\ell}_{L_S+\ell-1}$ also stabilizes. In fact we can prove the following corollary.

Corollary 4.77. *For sufficiently large m, we have*

$$\mathcal{E}^{\ell}_m \cong \mathcal{H}^{\ell}_{L_S+\ell-1}.$$

Proof. By Lemma 4.73, every non-trivial class in $\mathcal{E}^{\ell}_m$ admits a representative supported in $B^{\bullet}_{L_S+\ell-1}$. Hence

$$\mathcal{E}^{\ell}_m \subseteq \mathcal{H}_{L_S+\ell-1}.$$

By the definition of $H_{\min}$ and $\mathcal{H}^{\ell}_{L_S+\ell-1}$, it follows that

$$a_m(H_{\min}) = -g(K_m),$$

and that the classes in $\mathcal{H}^{\ell}_{L_S+\ell-1}$ are precisely those whose Alexander gradings lie in

$$\{-g(K_m), \ldots, -g(K_m) + \ell - 1\}.$$

Therefore $\mathcal{E}^{\ell}_m \cong \mathcal{H}^{\ell}_{L_S+\ell-1}$. □

We now examine the absolute Maslov grading. Let

$$\mathrm{gr}_{K_m}(H_{\min}) = [(\alpha_{\min}; \beta_{\min}, \gamma_{\min}; \delta_{\min})].$$

For any homogeneous class $H \in \mathcal{H}^{\ell}_{L_S+\ell-1}$ with

$$\mathrm{gr}_{K_m}(H) = [(\alpha_H; \beta_H, \gamma_H; \delta_H)],$$

we have

$$\beta_H = \beta_{\min}. \tag{4.10}$$

Indeed, by Lemma 4.28, the quantity $(\beta_{\min} - \beta_H)\omega$ is the stabilized slope of the relative Alexander grading $a_m(H) - a_m(H_{\min})$. Here, we use the assumption that $\omega \neq 0$.

To determine the absolute Maslov grading, we need a reference element. By Lemma 4.56, there exists a fixed element $\nu_0 \in B^\bullet$ of absolute Maslov grading zero. Furthermore, combining Lemma 4.56 and Corollary 2.124, we obtain a fixed element

$$(\alpha_0; \beta_0, \gamma_0; \delta_0) \in \widetilde{G}$$

such that $2\beta_0 = M + \frac{1}{2}$, and

$$gr_{K_m}(\nu_0) = [(\alpha_0; \beta_0, \gamma_0; \delta_0)].$$

Since $h_m(\nu_0) = 0$, Lemma 4.28 implies that

$$h_m(H_{\min}) = h_m(H_{\min}) - h_m(\nu_0) = (\beta_0 - \beta_{\min}) \left(M - \frac{1}{2} - \beta_{\min} - \beta_0 \right) m + \text{cte}.$$

It therefore remains to show that

$$(\beta_0 - \beta_{\min}) \left(M - \frac{1}{2} - \beta_{\min} - \beta_0 \right) = F_K.$$

Together with Equation (4.10), this will complete the proof.

Define

$$f(x) := (\beta_0 - x) \left(M - \frac{1}{2} - x - \beta_0 \right).$$

The preceding arguments apply equally well to the extremal summands of highest Alexander grading,

$$\bigoplus_{j=0}^{k-1} \widehat{HFK}(K_m, g(K_m) - j).$$

Accordingly, let $H_{\max}$ be a homogeneous class of maximal relative Alexander grading in $\mathcal{H}_{L_S+k-1}$. After increasing m if necessary, we may assume that

$$gr_{K_m}(H_{\max}) = [(\alpha_{\max}; \beta_{\max}, \gamma_{\max}; \delta_{\max})]$$

for all sufficiently large m. As above,

$$a_m(H_{\max}) = g(K_m),$$

and

$$h_m(H_{\max}) = f(\beta_{\max}) \cdot m + \text{cte}.$$

Here and throughout, cte denotes a constant independent of m. Consequently,

$$\widehat{HFK}(K_m, -g(K_m)) \cong \widehat{HFK}(K_{m+1}, -g(K_{m+1}))[f(\beta_{\min})],$$

and

$$\widehat{HFK}(K_m, g(K_m)) \cong \widehat{HFK}(K_{m+1}, g(K_{m+1}))[f(\beta_{\max})].$$

On the other hand, knot Floer homology satisfies the symmetry

$$\widehat{HFK}(K_m, g(K_m))[2g(K_m)] \cong \widehat{HFK}(K_m, -g(K_m)),$$

and

$$\widehat{HFK}(K_{m+1}, g(K_{m+1}))[2g(K_{m+1})] \cong \widehat{HFK}(K_{m+1}, -g(K_{m+1})).$$

These four isomorphisms fit into the commutative diagram

$$\begin{array}{ccc} \widehat{HFK}(K_{m+1}, g(K_{m+1})) & \xrightarrow{[f(\beta_{\max})]} & \widehat{HFK}(K_m, g(K_m)) \\ \downarrow{\scriptstyle [2g(K_{m+1})]} & & \downarrow{\scriptstyle [2g(K_m)]} \\ \widehat{HFK}(K_{m+1}, -g(K_{m+1})) & \xrightarrow{[f(\beta_{\min})]} & \widehat{HFK}(K_m, -g(K_m)) \end{array}$$

Comparing the total grading shifts around the square yields

$$-f(\beta_{\max}) - 2g(K_m) + f(\beta_{\min}) = -2g(K_{m+1}),$$

and hence

$$f(\beta_{\max}) - f(\beta_{\min}) = 2\big(g(K_{m+1}) - g(K_m)\big).$$

Equivalently,

$$(\beta_{\min} - \beta_{\max})\left(M - \frac{1}{2} - \beta_{\max} - \beta_{\min}\right) = 2\big(g(K_{m+1}) - g(K_m)\big). \tag{4.11}$$

Furthermore,

$$a_m(H_{\max}) - a_m(H_{\min}) = 2g(K_m).$$

Lemma 4.28 therefore gives

$$2g(K_m) = \big((\beta_{\min} - \beta_{\max})\omega\big)m + \text{cte}. \tag{4.12}$$

Combining Equations (4.11), (4.12), and Theorem 2.19, we obtain

$$\begin{cases} \beta_{\max} + \beta_{\min} = M - \omega - \dfrac{1}{2}, \\ \beta_{\min} - \beta_{\max} = -x([\hat{D}]). \end{cases}$$

Hence,

$$2\beta_{\min} = M - \omega - \frac{1}{2} - x([\hat{D}]).$$

This completes the proof.

4.5 The τ-Invariant

4.5.1 Reformulation of $\tau(K_m)$

Recall the definition of the τ-invariant from Definition 2.63. We now reformulate $\tau(K_m)$ in terms of the Alexander grading a_m on the basis $\boldsymbol{C}^m$.

Notation 4.78. Define

$$\mathcal{HS}_m = \Big\{ S_I \in \mathcal{X}_m \;\Big|\; \partial S_I = 0,\ [S_I] \neq 0 \Big\},$$

$$\mathcal{HZ}_m = \Big\{ S_I \in \mathcal{HS}_m \;\Big|\; S_I \text{ is homogeneous} \Big\}.$$

Proposition 4.79. *For every m,*

$$\tau(K_m) = \min_{S_I \in \mathcal{HS}_m} \Big(\max_{\nu \in I} a_m(\nu) \Big).$$

Proof. The family $\mathcal{HS}_m$ consists of all cycles whose homology classes generate

$$H_*(\mathcal{X}_m) \cong \widehat{\mathrm{HF}}(S^3).$$

Hence, by Definition 2.63,

$$\tau(K_m) = \min \Big\{ i \in \mathbb{Z} \;\Big|\; \mathcal{F}(K_m, i) \cap \mathcal{HS}_m \neq \varnothing \Big\}.$$

Since $S_I \in \mathcal{F}(K_m, i)$ if and only if $a_m(\nu) \leq i$ for every $\nu \in I$, the result follows. □

Since every homology class admits a homogeneous representative, for every $S_I \in \mathcal{HS}_m$ there exists $S_T \in \mathcal{HZ}_m$ such that $[S_T] = [S_I]$. In fact, T consists precisely of those elements of I with Maslov grading zero.

Since $T \subseteq I$, we have

$$\max_{\nu \in T} a_m(\nu) \leq \max_{\nu \in I} a_m(\nu).$$

Therefore, Proposition 4.79 implies

$$\tau(K_m) = \min_{S_I \in \mathcal{HZ}_m} \Big(\max_{\nu \in I} a_m(\nu) \Big). \tag{4.13}$$

Equation (4.13) shows that the computation of $\tau(K_m)$ reduces to understanding the maximal Alexander grading attained by elements of cycles in $\mathcal{HZ}_m$.

Notation 4.80. For a subset $I \subseteq \boldsymbol{C}^m$, let $V_{\max}(I)$ denote the first element of maximal Alexander grading with respect to the clockwise ordering. For a cycle S_I, write

$$V_{\max}(S_I) := V_{\max}(I).$$

Define

$$\mathcal{P}_m = \{V_{\max}(S_I) \mid S_I \in \mathcal{HZ}_m\}.$$

Equation (4.13) can then be rewritten as

$$\tau(K_m) = \min_{\nu \in \mathcal{P}_m} a_m(\nu). \tag{4.14}$$

The proof of Theorem 4.4 is reduced to understanding the stabilization of the sets $\mathcal{P}_m$. The key step is the following result.

Lemma 4.81. *For sufficiently large m,*

$$\Phi_m(\mathcal{P}_m) = \mathcal{P}_{m+1}.$$

Moreover, there exists an integer L_τ such that, for all sufficiently large m, the set $\mathcal{P}_m$ is a fixed subset of the ball $B^\bullet_{L_\tau}$.

Before proving Lemma 4.81, we need a detailed analysis of the Maslov grading zero elements of $\boldsymbol{C}^m$.

4.5.2 The set of elements with Maslov grading zero

Notation 4.82. Let

$$GZ_m = \{\nu \in \boldsymbol{C}^m \mid h_m(\nu) = 0\}.$$

The goal of this subsection is to describe the structure of the sets GZ_m as $m \to \infty$.

We begin with some consequences of the grading formulas established in Subsection 4.3.2. Some of theses results were already used implicitly in Section 4.4. We now isolate and prove these statements in their natural level of generality.

Recall from Corollary 2.124 that, for every fixed ball $B^\bullet_t$ and sufficiently large m, the grading double coset $\mathrm{gr}_{K_m}(\nu)$ of each element $\nu \in B^\bullet_t$ admits a fixed representative. We will use these representatives throughout the remainder of this subsection.

Corollary 4.83. *Let $t > 0$. For sufficiently large m, let $\nu_1, \nu_2 \in B^\bullet_t$, and let*

$$gr_{K_m}(\nu_i) = [(\alpha_i; \beta_i, \gamma_i; \delta_i)], \qquad i \in \{1, 2\}.$$

Then

$$a_m(\nu_1) - a_m(\nu_2) \text{ is constant} \iff \beta_1 = \beta_2, \tag{4.15}$$

$$\beta_2 > \beta_1 \implies a_m(\nu_2) > a_m(\nu_1), \tag{4.16}$$

and

$$(\beta_2 - \beta_1)\Big(M - \frac{1}{2} - \beta_1 - \beta_2\Big) > 0 \implies h_m(\nu_1) > h_m(\nu_2). \tag{4.17}$$

Proof. By Lemma 4.28, the relative Alexander and Maslov gradings are affine linear functions of m. Consequently, for sufficiently large m, the sign of each relative grading is determined by its slope unless the slope vanishes, in which case the relative grading is constant. The stated conclusions follow directly from the formulas of Lemma 4.28. □

Recall from Corollary 4.59 that the distinguished Maslov grading zero element $\nu_0 \in C^\bullet$ has grading representative

$$\mathrm{gr}_{K_m}(\nu_0) = [(\alpha_0; \beta_0, \gamma_0; \delta_0)],$$

where

$$\beta_0 = \frac{M}{2} + \frac{1}{4}.$$

Corollary 4.84. *Let $t > 0$. For sufficiently large m, let $\nu \in B_t^\bullet$, with representative*

$$\mathrm{gr}_{K_m}(\nu) = [(\alpha; \beta, \gamma; \delta)].$$

If $h_m(\nu) = 0$, then

$$\beta \in \{\beta_0, \beta_0 - 1\}.$$

Moreover, for sufficiently large m, the set $GZ_m \cap B_t^\bullet$ is independent of m.

Proof. We compare ν with the distinguished element ν_0. By Lemma 4.28, the relative Maslov grading

$$h_m(\nu) - h_m(\nu_0)$$

is an affine linear function of m. For sufficiently large m, the equality $h_m(\nu) = 0$ holds if and only if this affine linear function is identically zero.

Using the formula of Lemma 4.28, this is equivalent to

$$(\beta_0 - \beta)\Big(M - \frac{1}{2} - \beta - \beta_0\Big) = 0,$$

together with

$$(\gamma_0 - \gamma)M + (\alpha - \alpha_0) + \frac{1}{2}(\beta_0 - \beta) - (\beta_0 - \beta)\gamma - (\gamma_0 - \gamma)\beta_0 = 0.$$

The first equation implies

$$\beta = \beta_0 \qquad \text{or} \qquad \beta = M - \frac{1}{2} - \beta_0 = \beta_0 - 1.$$

Since the above conditions do not depend on m, the set $GZ_m \cap B_t^\bullet$ stabilizes for sufficiently large m. □

We now extend Corollary 4.84 from the stabilized sets

$$GZ_m \cap B_t^\bullet$$

to the entire set GZ_m.

Recall that the grading gr_{K_m} is induced from the gradings gr_K and gr_m. Throughout the remainder of this subsection, we will make use of both the clockwise and counterclockwise representatives of gr_m from Definition 2.123.

We are now ready to prove the following proposition.

Proposition 4.85. *Assume that $m \gg 0$. Let $\nu \in GZ_m$, and let $(\alpha; \beta, \gamma; \delta)$ be the representative of $\mathrm{gr}_{K_m}(\nu)$ coming from the counterclockwise representatives of gr_m. Then*

$$\beta \in \{\beta_0, \beta_0 - 1, \beta_0 - 2\}.$$

Proof. If $\nu \in C^\bullet$, then the statement follows from Corollary 4.84. Hence we may assume $\nu = x \otimes \xi_i$ for some $x \in B_K$ and $i \in \{1, \dots, m\}$.

Let $\nu' = x \otimes \xi_m \in C_m^\circ$, and let $(\alpha'; \beta', \gamma'; \delta')$ be the representative of $\mathrm{gr}_{K_m}(\nu')$ coming from the counterclockwise representatives of gr_m. By construction,

$$\text{if } \mathrm{gr}_K(x) = [(a; b, c; d)], \qquad \text{then} \qquad \beta' = b - \frac{1}{2}.$$

Note that this representative is independent of m. Moreover, $\beta = \beta'$, by the definition of the counterclockwise representatives of gr_m.

Let $\nu_0 \in C^\bullet$ be the element from Lemma 4.57. By Lemma 4.28,

$$h_m(\nu') = h_m(\nu') - h_m(\nu_0) = (\beta_0 - \beta')\big((\beta_0 - 1) - \beta'\big)m + F, \tag{4.18}$$

where F is a fixed integer.

On the other hand, by Lemma 4.29,

$$h_m(x \otimes \xi_j) - h_m(\nu') = h_m(x \otimes \xi_j) - h_m(x \otimes \xi_m) = (j - m)\Big(M - 2b - \frac{1}{2}\Big),$$

and hence

$$h_m(x \otimes \xi_j) - h_m(\nu') = (j - m)(2\beta_0 - 2\beta' - 2) \qquad \text{for } j \in \{1, \dots, m\}. \tag{4.19}$$

Equations (4.18) and (4.19) show that $h_m(\nu')$ is a linear function of m with slope

$$(\beta_0 - \beta')\big((\beta_0 - 1) - \beta'\big),$$

while the set

$$\big\{h_m(x \otimes \xi_j) - h_m(\nu') \mid 1 \leq j \leq m\big\}$$

forms an arithmetic progression of length m and common difference

$$2\beta_0 - 2\beta' - 2.$$

Since $m \gg 0$, the existence of an index $i \in \{1, \ldots, m\}$ satisfying $h_m(x \otimes \xi_i) = 0$ forces

$$\big|2\beta_0 - 2\beta' - 2\big| \geq \big|(\beta_0 - \beta')\big((\beta_0 - 1) - \beta'\big)\big|.$$

A direct calculation shows that this can occur only when $|\beta_0 - \beta'| \leq 2$.

It remains to rule out the possibilities $\beta' \in \{\beta_0 + 1, \beta_0 + 2\}$. In this case, Corollary 4.83, and in particular Equation (4.17) implies that $h_m(\nu') > 0$. Moreover,

$$h_m(x \otimes \xi_j) - h_m(\nu') = -4(j - m) \qquad \text{or} \qquad -6(j - m),$$

and therefore

$$h_m(x \otimes \xi_j) \geq h_m(\nu') > 0$$

for every $j \in \{1, \ldots, m\}$. This contradicts the assumption that $\nu = x \otimes \xi_i \in GZ_m$.

Hence $\beta' \in \{\beta_0, \beta_0 - 1, \beta_0 - 2\}$. Since $\beta = \beta'$, the proposition follows. □

Proposition 4.85 induces the following decomposition of GZ_m.

Notation 4.86. For $i \in \{0, 1, 2\}$, let GZ_m^i denote the subset of elements $\nu \in GZ_m$ whose counterclockwise representative $(\alpha; \beta, \gamma; \delta)$ satisfies $\beta = \beta_0 - i$. Thus

$$GZ_m = GZ_m^0 \sqcup GZ_m^1 \sqcup GZ_m^2.$$

The proof of Proposition 4.85 contains additional information about this decomposition, which we record in the following corollary.

Corollary 4.87. *Assume that $m \gg 0$. Then there exists an integer $L_{GZ} > 0$ such that*

$$GZ_m^0 \subseteq B^\bullet_{[-L_{GZ},0]}, \qquad \Phi_m(GZ_m^0) = GZ_{m+1}^0,$$
$$GZ_m^2 \subseteq B^\bullet_{[0,L_{GZ}]}, \qquad \Phi_m(GZ_m^2) = GZ_{m+1}^2.$$

Moreover, there exists a subset $ZB \subseteq B_K$ such that

$$GZ_m^1 \setminus C^\bullet = \{x \otimes \xi_i \mid x \in ZB,\ 1 \leq i \leq m\}.$$

Proof. We first consider GZ_m^0. Let

$$\nu \in GZ_m^0 \setminus C^\bullet, \qquad \nu = x \otimes \xi_i.$$

Let $\nu' = x \otimes \xi_m$, and let $(\alpha'; \beta', \gamma'; \delta')$ be the representative of $\mathrm{gr}_{K_m}(\nu')$ coming from the counterclockwise representatives of gr_m. Since $\nu \in GZ_m^0$, we have $\beta' = \beta_0$.

Rewriting Equations (4.18) and (4.19), we obtain

$$h_m(\nu') = F, \qquad h_m(x \otimes \xi_j) - h_m(\nu') = -2(j - m)$$

for every $j \in \{1, \dots, m\}$. Consequently,

$$h_m(\nu) = 0 \iff m - \frac{F}{2} \in \mathbb{Z}_+ \quad \text{and} \quad i = m - \frac{F}{2}.$$

In particular,

$$\nu \in B^\bullet_{[-F/2,0]}.$$

Taking the maximum over all possible choices of $\nu' \in C_m^\circ$ gives the desired statement for GZ_m^0. A similar argument yields the corresponding statement for GZ_m^2.

Finally, consider GZ_m^1. Proceeding as above, we now have $\beta' = \beta_0 - 1$. Hence

$$h_m(x \otimes \xi_j) - h_m(\nu') = 0 \qquad \text{for all } j \in \{1, \dots, m\}.$$

Therefore,

$$h_m(x \otimes \xi_i) = 0 \Longrightarrow h_m(x \otimes \xi_j) = 0$$

for every $j \in \{1, \dots, m\}$, which proves the final statement. □

Remark 4.88. At first glance, Corollaries 4.84 and 4.87 may appear contradictory. The reason is that they use different choices of representatives for the grading double cosets.

In Corollary 4.84, we use the fixed representatives provided by Corollary 2.124. These coincide with the counterclockwise representatives on the negative half-ball, but differ on the positive half-ball. More precisely, passing from the counterclockwise representatives to the fixed representatives increases the second grading component by 1 on the positive half-ball (apart from the elements of $C^\bullet$).

Since every element of GZ_m^2 lies in the positive half-ball, its fixed representative has second grading component

$$(\beta_0 - 2) + 1 = \beta_0 - 1.$$

This is consistent with Corollary 4.84.

4.5.3 Stabilization of maximizers of the Alexander grading

In this subsection, we prove Lemma 4.81. We begin by showing that the set $\mathcal{P}_m$ remains within a bounded distance of the black box $C^\bullet$.

Proposition 4.89. *There exists an integer $L_\tau > 0$ such that*

$$\mathcal{P}_m \subseteq B^\bullet_{L_\tau}$$

for all sufficiently large m.

Proof. Let $S_I \in \mathcal{HZ}_m$. Since S_I is homogeneous, $I \subseteq GZ_m$. By Corollary 4.87, the subsets GZ^0_m and GZ^2_m lie within a bounded distance of $C^\bullet$. Hence it suffices to bound the distance of potential maximizers lying in GZ^1_m.

Using the notation of Corollary 4.87, let

$$x \in ZB, \qquad \nu' = x \otimes \xi_m \in GZ^1_m.$$

Let $\nu_0 \in I$ be the distinguished element from Lemma 4.57 and Corollary 4.59. By Lemmas 4.28 and 4.29,

$$a_m(\nu_0) - a_m(\nu') = -\omega m + F'_x, \tag{4.20}$$

and

$$a_m(x \otimes \xi_j) - a_m(x \otimes \xi_m) = (j-m)\omega, \qquad j \in \{1, \dots, m\}, \tag{4.21}$$

where F'_x depends only on x.

Combining (4.20) and (4.21) gives

$$a_m(x \otimes \xi_j) - a_m(\nu_0) = j\omega - F'_x.$$

Since $\omega < 0$,

$$j \geq \left|\frac{F'_x}{\omega}\right| \Longrightarrow a_m(x \otimes \xi_j) \leq a_m(\nu_0),$$

and therefore

$$x \otimes \xi_j \neq V_{\max}(S_I).$$

Thus only the first $\left|\frac{F'_x}{\omega}\right|$ elements of the string

$$x \otimes \xi_1, \dots, x \otimes \xi_m$$

can occur as maximizers. Taking the maximum over all $x \in ZB$, we obtain the result whenever

$$L_\tau \geq \max_{x \in ZB} \left|\frac{F'_x}{\omega}\right| \qquad \text{and} \qquad L_\tau \geq L_{GZ}.$$

□

Corollary 4.90 follows immediately from Proposition 4.89 and Corollary 4.87.

Corollary 4.90. *For sufficiently large m, let $S_I \in \mathcal{HZ}_m$. Then*

$$V_{\max}(S_I) \in (I \cap B^{\bullet}_{[0,L_\tau]}) \cup (I \cap GZ^0_m).$$

To prove Lemma 4.81, we prove a stronger statement in the form of Lemmas 4.92 and 4.94. These lemmas give stabilization results for the collections of subsets of closed balls $B^{\bullet}_L$ that appear in homogeneous generators, i.e. elements of $\mathcal{H}Z_m$.

Notation 4.91. For $L > 0$, define

$$\mathcal{HI}^+_{m,L} = \{I \cap B^{\bullet}_{[0,L]} \mid S_I \in \mathcal{HZ}_m\}.$$

Lemma 4.92. *For sufficiently large m, the family $\mathcal{HI}^+_{m,L}$ stabilizes under Φ_m.*

Notation 4.93. Define

$$\mathcal{HI}^-_m = \{I \cap GZ^0_m \mid S_I \in \mathcal{HZ}_m\}.$$

Lemma 4.94. *For sufficiently large m, the family $\mathcal{HI}^-_m$ stabilizes under Φ_m.*

We begin with the proof of Lemma 4.92. The proof is based on Construction 4.60. The key observation is that the behaviour of a homology generator inside the positive half-ball is completely determined by the stable part of its nested component.

Proof. First, note that

$$\mathcal{HI}^+_{m,L} = \left\{ I \cap B^{\bullet}_{[0,L]} \cap GZ_m \;\middle|\; S_I \in \mathcal{HS}_m \right\}. \tag{4.22}$$

Fix $m_0 \in \mathbb{N}$ and let $S_I \in \mathcal{HS}_{m_0}$. By Proposition 4.58, there exists an $\mathcal{A}_\infty$-homotopy equivalence f whose associated generator is S_I. Using Construction 4.60, let

$$S_{I_m} \in \mathcal{HS}_m$$

be the associated growing family. Write

$$S_{I_m} = S_{I''_m} + S_{I'_m},$$

where $S_{I''_m}$ and $S_{I'_m}$ are the nested and stable components. By Equation 4.7,

$$I''_{m+1} = \Phi'_m(I''_m) \sqcup NI''_{m+1}, \qquad I'_{m+1} = \Phi_m(I'_m).$$

For sufficiently large m, the subsets $I'_m,\ I'_{m+1},\ NI''_{m+1}$ lie outside the positive half-ball $B^{\bullet}_{[0,L]}$. Hence

$$I_m \cap B^{\bullet}_{[0,L]} = I''_m \cap B^{\bullet}_{[0,L]},$$

and similarly

$$I_{m+1} \cap B^{\bullet}_{[0,L]} = I''_{m+1} \cap B^{\bullet}_{[0,L]}.$$

Moreover, for sufficiently large m, the inclusions Φ_m and Φ'_m agree on $B^{\bullet}_{[0,L]}$. Therefore

$$\Phi_m(I_m \cap B^{\bullet}_{[0,L]}) = \Phi_m(I''_m \cap B^{\bullet}_{[0,L]}) = \Phi'_m(I''_m) \cap B^{\bullet}_{[0,L]}$$

$$= I''_{m+1} \cap B^{\bullet}_{[0,L]} = I_{m+1} \cap B^{\bullet}_{[0,L]}.$$

Thus $I_m \cap B^{\bullet}_{[0,L]}$ stabilizes for sufficiently large m. In fact, for $m \geq L+1$, this subset is completely determined by the maps $f_1, \ldots, f_{L+1}$, and hence is independent of m.

Running this argument over all $\mathcal{A}_\infty$-homotopy equivalences

$$f : \mathcal{S}_\infty \longrightarrow \mathcal{S},$$

we conclude that the family

$$\left\{ I \cap B^{\bullet}_{[0,L]} \;\middle|\; S_I \in \mathcal{HS}_m \right\}$$

stabilizes for sufficiently large m.

Finally, $GZ_m \cap B^{\bullet}_{[0,L]}$ stabilizes by Corollary 4.84. The lemma now follows from Equation (4.22). □

Before proving Lemma 4.94, we establish a compatibility between the equivalence relation introduced in Subsection 4.3.7 and the decomposition of GZ_m from Proposition 4.85.

Proposition 4.95. *Let $S_I \in \mathcal{H}Z_m$, and let $\sim$ be the equivalence relation on I introduced in Subsection 4.3.7. If $\nu, \nu' \in I \setminus C^{\bullet}$ satisfy $[\nu]_\sim = [\nu']_\sim$, then ν and ν' belong to the same subset GZ^i_m.*

Proof. Let

$$\nu = \nu_0, \nu_1, \ldots, \nu_k = \nu'$$

be a good alternating path. For each $i \in \{1, \ldots, k\}$, let μ_i denote the common target appearing in the definition of a good alternating path, so that

$$\nu_i \longrightarrow \mu_i, \qquad \nu_{i-1} \longrightarrow \mu_i,$$

and neither edge is of the third type. Let

$$(\alpha_i; \beta_i, \gamma_i; \delta_i) \qquad \text{and} \qquad (\alpha_i'; \beta_i', \gamma_i'; \delta_i')$$

be the counterclockwise representatives of $gr_{K_m}(\nu_i)$ and $gr_{K_m}(\mu_i)$, respectively.

We first compare the second grading component across a single edge. Suppose

$$\nu_i \longrightarrow \mu_i$$

is an edge of the fourth type. Then there exist elements $x_i, x_i' \in B_K$ and integers $j, j' \in \{1, \ldots, m\}$ such that

$$\nu_i = x_i \otimes \xi_j, \qquad \mu_i = x_i' \otimes \xi_{j'}, \qquad \text{and}$$

$$x_i' \text{ appears in } m_{j'-j+1}\big(x_i \otimes \underbrace{\rho_{23} \otimes \cdots \otimes \rho_{23}}_{j'-j}\big).$$

Using Construction 2.92, we obtain

$$gr_K(x_i') = \lambda^{j'-j+1}\, gr_K(x_i)\, \underbrace{gr(\rho_{23}) \cdots gr(\rho_{23})}_{j'-j}.$$

Hence, if

$$gr_K(x_i) = [(a_i; b_i, c_i; d_i)] \quad \text{and} \quad gr_K(x_i') = [(a_i'; b_i', c_i'; d_i')]$$

then $b_i = b_i'$.

Using the description of the counterclockwise representatives of gr_m, we have

$$\beta_i = b_i - \frac{1}{2}, \qquad \beta_i' = b_i' - \frac{1}{2},$$

and therefore $\beta_i = \beta_i'$. Thus edges of type (4) preserve the second grading component. The same argument applies to edges of the first type.

Now suppose $\nu_i \longrightarrow \mu_i$ is an edge of the second type. In this case

$$\nu_i = x_i \otimes \eta,$$

and the same calculation gives

$$b_i = b_i' + \frac{1}{2}.$$

Since

$$\beta_i = b_i, \qquad \beta_i' = b_i' - \frac{1}{2},$$

we obtain $\beta_i = \beta_i' + 1$.

We now return to the alternating path. The only configurations that change the second grading component are

$$\nu_i \xrightarrow{[2]} \mu_i, \qquad \nu_{i-1} \xrightarrow{[4]} \mu_i,$$

and

$$\nu_i \xrightarrow{[4]} \mu_i, \qquad \nu_{i-1} \xrightarrow{[2]} \mu_i.$$

In the first case, the alternating path enters $C^\bullet$, while in the second case it exits $C^\bullet$. Since both ν and ν' lie outside $C^\bullet$, every entrance is paired with a corresponding exit. The resulting changes in the second grading component therefore cancel.

Consequently,

$$\beta_0 = \beta_k,$$

which completes the proof. □

Before proving Lemma 4.94, we introduce two families of subsets that will be used throughout the proof.

Notation 4.96. Let $\mathcal{I}'_m$ denote the family of all subsets $I'_m \subseteq \boldsymbol{C}^m$ arising as the stable component of a growing family. Equivalently,

$$\mathcal{I}'_m := \Big\{ f_2(x_0 \otimes \rho_1) \otimes \xi_m \;\Big|\; f : \mathcal{S}_\infty \longrightarrow \mathcal{S} \text{ is an } \mathcal{A}_\infty\text{-homotopy equivalence} \Big\}.$$

Notation 4.97. Define

$$\mathcal{NI}^-_m := \Big\{ T \subseteq GZ^0_m \setminus C^\bullet \;\Big|\; \partial S_T = 0 \Big\}.$$

We are now ready to prove Lemma 4.94.

Proof. First, note that

$$\mathcal{HI}^-_m = \Big\{ I \cap GZ^0_m \;\Big|\; S_I \in \mathcal{HS}_m \Big\}. \tag{4.23}$$

Fix $m_0 \in \mathbb{N}$ and let $S_I \in \mathcal{HS}_{m_0}$. By Proposition 4.58, there exists an $\mathcal{A}_\infty$-homotopy equivalence f whose associated generator is S_I. Using Construction 4.60, let

$$S_{I_m} \in \mathcal{HS}_m$$

be the associated growing family. Write

$$S_{I_m} = S_{I''_m} + S_{I'_m},$$

where $S_{I''_m}$ and $S_{I'_m}$ are the nested and stable components, respectively.

Since we work over $\mathbb{F}_2$, we have

$$I_m = I''_m \ominus I'_m,$$

where $\ominus$ denotes symmetric difference. Consequently,

$$I_m \cap GZ^0_m = (I''_m \cap GZ^0_m) \ominus (I'_m \cap GZ^0_m). \tag{4.24}$$

The stable component I'_m belongs to the family $\mathcal{I}'_m$, which stabilizes under the inclusions Φ_m. Moreover,

$$I_m \cap C^\bullet = I''_m \cap C^\bullet$$

depends only on $f_1(x_0)$, and therefore stabilizes as well. Thus it remains to study

$$(I''_m \cap GZ^0_m) \setminus C^\bullet.$$

To do so, we use the equivalence relation $\sim$ introduced in Subsection 4.3.7. It is important to note that, throughout the remainder of the proof, we consider $\sim$ as a relation on the subset I''_m (as in Proposition 4.65). By Proposition 4.95,

$$\nu \in (I''_m \cap GZ^0_m) \setminus C^\bullet \Longrightarrow [\nu]_\sim \subseteq I''_m \cap GZ^0_m.$$

Therefore $I''_m \cap GZ^0_m$ is a disjoint union of equivalence classes.

The next observation is crucial:

$$\forall\, \nu \in (I''_m \cap GZ^0_m) \setminus C^\bullet, \qquad [\nu]_\sim \cap C^\bullet = \varnothing \tag{4.25}$$

for sufficiently large m.

Indeed, by Corollary 4.87,

$$GZ^0_m \subseteq B^\bullet_{[-L_{GZ},0]}.$$

For sufficiently large m, the only edges connecting $C^\bullet$ to $B^\bullet_{[-L_{GZ},0]} \setminus C^\bullet$ are edges of type (3), since the lengths of all edges are uniformly bounded by Lemma 4.47. Hence, no good alternating path that starts in $C^\bullet$ and remains inside $B^\bullet_{[-L_{GZ},0]}$ can leave $C^\bullet$. This proves the claim.

The previous observation allows us to isolate a large family of null-homologous cycles inside GZ^0_m. Let

$$T = (I''_m \cap GZ^0_m) \setminus C^\bullet.$$

Combining Equation (4.25) with Proposition 4.65, we obtain

$$\partial S_T = 0 \Longrightarrow S_T \text{ is a null-homologous cycle supported in } GZ^0_m. \tag{4.26}$$

Equivalently,

$$T \in \mathcal{NI}_m^-.$$

On the other hand, every element of $\mathcal{NI}_m^-$ corresponds to a null-homologous cycle that may be added to a homogeneous generator without changing the represented homology class. Therefore, the family of subsets of the form

$$(I_m'' \cap GZ_m^0) \setminus C^\bullet$$

is precisely the family $\mathcal{NI}_m^-$.

Since GZ_m^0 stabilizes by Corollary 4.87, and the differentials inside $B^\bullet_{[-L_{GZ},0]}$ stabilize by Proposition 4.22, the family $\mathcal{NI}_m^-$ stabilizes as well.

We now combine the stabilized contribution of I_m' with the null-homologous cycles contained in GZ_m^0. Using Equation (4.24) and the discussion above, we obtain

$$\mathcal{HI}_m^- = \Big\{ T \ominus I_m' \;\Big|\; T \in \mathcal{NI}_m^-,\ I_m' \in \mathcal{I}_m' \Big\}.$$

Since both families $\mathcal{NI}_m^-$ and $\mathcal{I}_m'$ stabilize under the inclusions Φ_m, the same is true for $\mathcal{HI}_m^-$. This proves the lemma. □

We are now ready to prove the main result of this subsection, Lemma 4.81. Motivated by Corollary 4.90, we introduce the following family of subsets.

Notation 4.98. Define

$$\mathcal{PI}_m := \Big\{ (I \cap B^\bullet_{[0,L_\tau]}) \cup (I \cap GZ_m^0) \;\Big|\; S_I \in \mathcal{H}Z_m \Big\}.$$

Proof. By Corollary 4.90,

$$\mathcal{P}_m = \big\{ V_{\max}(J) \mid J \in \mathcal{PI}_m \big\}.$$

Thus it suffices to show that the family $\mathcal{PI}_m$ stabilizes under the inclusions Φ_m, and that for each stabilized $J \in \mathcal{PI}_m$, the maximizer $V_{\max}(J)$ is eventually independent of m. Now, for

$$J = (I \cap B^\bullet_{[0,L_\tau]}) \cup (I \cap GZ_m^0), \qquad S_I \in \mathcal{H}Z_m,$$

the positive-half-ball part $I \cap B^\bullet_{[0,L_\tau]}$ is controlled by Lemma 4.92, while the negative part $I \cap GZ_m^0$ is controlled by Lemma 4.94. Hence the family $\mathcal{PI}_m$ stabilizes for sufficiently large m as a family of subsets of $B^\bullet_{L_\tau}$.

Moreover, by Corollary 4.83, the sign of the relative Alexander grading between any pair of elements of $B^\bullet_{L_\tau}$ stabilizes for sufficiently large m. Therefore, for each stabilized $J \in \mathcal{PI}_m$, the element $V_{\max}(J)$ is eventually independent of m. Equivalently, the family $\mathcal{P}_m$ stabilizes for sufficiently large m.

This proves Lemma 4.81. □

4.5.4 Proof of Theorem 4.4

In this final subsection of Section 4.5, we finally prove Theorem 4.4.

Proof. By Equation (4.14),

$$\tau(K_m) = \min_{\nu \in \mathcal{P}_m} a_m(\nu).$$

By Lemma 4.81, the family $\mathcal{P}_m$ stabilizes for sufficiently large m. Hence we may choose a fixed element $\nu_{\min} \in \mathcal{P}_m$ such that

$$a_m(\nu_{\min}) = \min_{\nu \in \mathcal{P}_m} a_m(\nu)$$

for all sufficiently large m.

It therefore suffices to show that $a_m(\nu_{\min})$ is eventually a linear function of m. Let

$$L = \max\{L_S + L_B + 2L_A,\ L_\tau\}.$$

By the proof of Theorem 4.3, there exists a fixed subset

$$I \subseteq B^\bullet_L$$

such that S_I is a cycle in $\mathcal{X}^K_m$ representing a non-trivial homology class in Alexander grading $-g(K_m)$ for all sufficiently large m. Choose $\nu \in I$ so that

$$a_m(\nu) = -g(K_m)$$

for all sufficiently large m.

By Lemma 4.28, the relative Alexander grading

$$a_m(\nu_{\min}) - a_m(\nu)$$

is a linear function of m for sufficiently large m. By Theorem 2.19, $g(K_m)$ is eventually an affine linear function of m. Therefore

$$\tau(K_m) = a_m(\nu_{\min}) = \big(a_m(\nu_{\min}) - a_m(\nu)\big) - g(K_m)$$

is eventually an affine linear function of m. This proves Theorem 4.4. □

Bibliography

[Abo06] Mohammed Abouzaid. On the Fukaya categories of higher genus surfaces. *Advances in Mathematics*, 217:1192–1235, 2006.

[AC21] Marwa E. Alrefai and Nafaa Chbili. On Fox's trapezoidal conjecture for closed 3-braids. *Internat. J. Math.*, 32(1):Paper No. 2150002, 11, 2021.

[AC23] Mark E. AlSukaiti and Nafaa Chbili. Alexander and Jones polynomials of weaving 3-braid links and Whitney rank polynomials of Lucas lattice. *preprint, arXiv:2303.11398*, 2023.

[AGL25] Akram Alishahi, Eugene Gorsky, and Beibei Liu. Colored knot Floer homology: structures and examples, 2025.

[AH21] Jared Able and Mikami Hirasawa. Construction and decomposition of knots as Murasugi sums of Seifert surfaces. *preprint, arXiv:2112.00297*, 2021.

[AJK24] Soheil Azarpendar, András Juhász, and Tamás Kálmán. On Fox's trapezoidal conjecture. *preprint, arXiv:2406.08662*, 2024.

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

[Aza23] Soheil Azarpendar. Twisting and stabilization in knot Floer homology. *preprint, arXiv:2312.01501*, 2023.

[Aza25] Soheil Azarpendar. Twisting, stabilization and bordered Floer homology. *preprint, arXiv:2507.15144*, 2025.

[Ban22] Ian M. Banfield. Christoffel words and the strong Fox conjecture for two-bridge knots. *preprint, arXiv:2212.04561*, 2022.

[BHS26] John A. Baldwin, Jonathan Hanselman, and Steven Sivek. Ribbon concordance and fibered predecessors, ii: the general case. *preprint, arXiv:2602.21109*, 2026.

[BKL85] Anders Björner, Bernhard Korte, and László Lovász. Homotopy properties of greedoids. *Adv. in Appl. Math.*, 6(4):447–494, 1985.

[BM19] Kenneth L. Baker and Kimihiko Motegi. Seifert vs. slice genera of knots in twist families and a characterization of braid axes. *Proceedings of the London Mathematical Society*, 119(6):1493–1530, 2019.

[BMM25] Kenneth L. Baker, Yasuyuki Miyazawa, and Kimihiko Motegi. Asymptotic behavior of unknotting numbers of links in a twist family. *Topology Appl.*, 367:109350, 2025.

[Bon25] Joe Boninger. Twisted knots and the perturbed Alexander invariant. *New York J. Math.*, 31:126–166, 2025.

[BT19] Kenneth L. Baker and Scott A. Taylor. Dehn filling and the Thurston norm. *J. Differential Geom.*, 112(3):391–409, 2019.

[BZH14] Gerhard Burde, Heiner Zieschang, and Michael Heusener. *Knots*, volume 5 of *De Gruyter Studies in Mathematics*. De Gruyter, Berlin, extended edition, 2014.

[Che22] Daren Chen. Twistings and the Alexander polynomial, 2022.

[Che23] Wenzhao Chen. Knot Floer homology of satellite knots with $(1, 1)$-patterns. *Selecta Math. (N.S.)*, 29(2):Paper No. 22, 63, 2023.

[CHS22] Zhechi Cheng, Matthew Hedden, and Sucharit Sarkar. Murasugi sum and extremal knot Floer homology. *Quantum Topology*, 2022.

[Con70] J. H. Conway. An enumeration of knots and links, and some of their algebraic properties. In *Computational Problems in Abstract Algebra (Proc. Conf., Oxford, 1967)*, pages 329–358. Pergamon, Oxford-New York-Toronto, Ont., 1970.

[Cro59] Richard H. Crowell. Genus of alternating link types. *Annals of Mathematics*, 69:258, 1959.

[FHS82] Michael H. Freedman, Joel Hass, and Peter Scott. Closed geodesics on surfaces. *Bulletin of The London Mathematical Society*, 14:385–391, 1982.

[Fox61] R. H. Fox. Some problems in knot theory. In *Topology of 3-manifolds and related topics (Proc. The Univ. of Georgia Institute, 1961)*, pages 168–176. Prentice-Hall, Inc., Englewood Cliffs, NJ, 1961.

[FS83] Satoshi Furusawa and Makoto Sakuma. Dehn surgery on symmetric knots. *Math. Sem. Notes Kobe Univ.*, 11(2):179–198, 1983.

[Gab83] David Gabai. The Murasugi sum is a natural geometric operation. In *Low Dimensional Topology*, volume 20 of *Contemporary Mathematics*, pages 131–143. American Mathematical Society, Providence, RI, 1983.

[Gab86] David Gabai. Genera of the alternating links. *Duke Mathematical Journal*, 53:677–681, 1986.

[Har79] R. I. Hartley. On two-bridged knot polynomials. *Journal of the Australian Mathematical Society*, 28(2):241–249, 1979.

[Hed05] Matthew Hedden. On knot Floer homology and cabling. *Algebraic & Geometric Topology*, 5(3):1197–1222, 2005.

[Hed07] Matthew Hedden. Knot Floer homology of Whitehead doubles. *Geometry & Topology*, 11(4):2277–2338, 2007.

[Hed09] Matthew Hedden. On Knot Floer Homology and Cabling: 2. *International Mathematics Research Notices*, 2009(12):2248–2274, 01 2009.

[Hir95] Mikami Hirasawa. The flat genus of links. *Kobe journal of mathematics*, 12(2):155–159, 12 1995.

[HL16] Matthew Hedden and Adam Simon Levine. Splicing knot complements and bordered Floer homology. *J. Reine Angew. Math.*, 720:129–154, 2016.

[HMV24] Elena S. Hafner, Karola Mészáros, and Alexander Vidinas. Log-concavity of the Alexander polynomial. *Int. Math. Res. Not. IMRN*, pages 10273–10284, 2024.

[HRW22] Jonathan Hanselman, Jacob Rasmussen, and Liam Watson. Heegaard Floer homology for manifolds with torus boundary: properties and examples. *Proc. Lond. Math. Soc. (3)*, 125(4):879–967, 2022.

[HRW24] Jonathan Hanselman, Jacob Rasmussen, and Liam Watson. Bordered Floer homology for manifolds with torus boundary via immersed curves. *J. Amer. Math. Soc.*, 37(2):391–498, 2024.

[HW23] Jonathan Hanselman and Liam Watson. Cabling in terms of immersed curves. *Geom. Topol.*, 27(3):925–952, 2023.

[Jon10] In Dae Jong. Alexander polynomials of alternating knots of genus two II. *J. Knot Theory Ramifications*, 19(8):1075–1092, 2010.

[Juh06] András Juhász. Floer homology and surface decompositions. *Geometry & Topology*, 12:299–350, 2006.

[Juh23] András Juhász. *Differential and low-dimensional topology*, volume 104 of *London Mathematical Society Student Texts*. Cambridge University Press, Cambridge, 2023.

[Kál13] Tamás Kálmán. A version of Tutte's polynomial for hypergraphs. *Advances in Mathematics*, 244:823–873, 2013.

[Kau83] Louis H. Kauffman. *Formal knot theory*, volume 30 of *Mathematical Notes*. Princeton University Press, Princeton, NJ, 1983.

[Kaw96] Akio Kawauchi. *A survey of knot theory*. Birkhäuser Verlag, Basel, 1996. Translated and revised from the 1990 Japanese original by the author.

[Kid82] Mark E. Kidwell. Relations between the Alexander polynomial and summit power of a closed braid. *Math. Sem. Notes Kobe Univ.*, 10(2):387–409, 1982.

[LC18a] Peter Lambert-Cole. Twisting, mutation and knot Floer homology. *Quantum Topol.*, 9(4):749–774, 2018.

[LC18b] Peter Lambert-Cole. Twisting, mutation and knot Floer homology. *Quantum Topol.*, 9(4):749–774, 2018.

[Lev14] Adam Simon Levine. Nonsurjective satellite operators and piecewise-linear concordance. *Forum of Mathematics, Sigma*, 4, 2014.

[LOT18] Robert Lipshitz, Peter S. Ozsváth, and Dylan P. Thurston. Bordered Heegaard Floer homology. *Mem. Amer. Math. Soc.*, 254(1216):viii+279, 2018.

[Min82] Jerome Minkus. The branched cyclic coverings of 2 bridge knots and links. *Mem. Amer. Math. Soc.*, 35(255):iv+68, 1982.

[Mor86] H. R. Morton. Seifert circles and knot polynomials. *Math. Proc. Cambridge Philos. Soc.*, 99(1):107–109, 1986.

[Mur58] Kunio Murasugi. On the genus of the alternating knot. I, II. *Journal of the Mathematical Society of Japan*, 10(1), 1958.

[Mur65] Kunio Murasugi. On a certain numerical invariant of link types. *Transactions of the American Mathematical Society*, 117:387–422, 1965.

[Mur85] Kunio Murasugi. On the Alexander polynomial of alternating algebraic knots. *J. Austral. Math. Soc. Ser. A*, 39(3):317–333, 1985.

[Nak00] Takuji Nakamura. Positive alternating links are positively alternating. *Journal of Knot Theory and Its Ramifications*, 9(1):107–112, 2000.

[Ni21] Yi Ni. A characterization of $T_{2g+1,2}$ among alternating knots. *Acta Math. Sin. (Engl. Ser.)*, 37(12):1841–1846, 2021.

[OS03a] Peter S. Ozsváth and Zoltán Szabó. Heegaard Floer homology and alternating knots. *Geometry & Topology*, 7:225–254, 2003.

[OS03b] Peter S. Ozsváth and Zoltán Szabó. Knot Floer homology and the four-ball genus. *Geometry & Topology*, 7(2):615–639, 2003.

[OS04a] Peter S. Ozsváth and Zoltán Szabó. Holomorphic disks and genus bounds. *Geometry & Topology*, 8(1):311–334, 2004.

[OS04b] Peter S. Ozsváth and Zoltán Szabó. Holomorphic disks and knot invariants. *Advances in mathematics (New York. 1965)*, 186(1):58–116, 2004.

[Par78] Richard L. Parris. *Pretzel knots*. PhD thesis, Princeton University, 1978.

[PW20] Ina Petkova and Biji Wong. Twisted Mazur pattern satellite knots and bordered Floer theory. *Michigan Mathematical Journal*, 2020.

[Ras03] Jacob Andrew Rasmussen. Floer homology and knot complements, 2003.

[Rob11] Lawrence P. Roberts. Extending Van Cott's bounds for the τ and s-invariants of a satellite knot. *Journal of Knot Theory and Its Ramifications*, 20:1237–1245, 2011.

[Sch53] Horst Schubert. Knoten und vollringe. *Acta Mathematica*, 90:131–286, 1953.

[Sto05] Alexander Stoimenow. Newton-like polynomials of links. *Enseign. Math. (2)*, 51(3-4):211–230, 2005.

[Sto08] A. Stoimenow. Knots of (canonical) genus two. *Fund. Math.*, 200(1):1–67, 2008.

[Sto16] Alexander Stoimenow. *Diagram genus, generators, and applications*. Monographs and Research Notes in Mathematics. CRC Press, Boca Raton, FL, 2016.

[VC10] Cornelia A. Van Cott. Ozsváth–szabó and Rasmussen invariants of cable knots. *Algebraic & Geometric Topology*, 10(2):825–836, 04 2010.