DIAGRAMMATIC AND GEOMETRIC INVARIANTS OF HYPERBOLIC WEAKLY GENERALIZED ALTERNATING KNOTS

By

Brandon Bavier

A DISSERTATION

Submitted to
Michigan State University
in partial fulfillment of the requirements
for the degree of

Mathematics – Doctor of Philosophy

2021

ABSTRACT

DIAGRAMMATIC AND GEOMETRIC INVARIANTS OF HYPERBOLIC WEAKLY GENERALIZED ALTERNATING KNOTS

By

Brandon Bavier

We study the relationship between knot and link diagrams on surfaces and their invariants coming from hyperbolic geometry. A *link diagram on a surface*, denoted $\pi(L) \subset F$, is a way of projecting a link L in some manifold M onto some surface $F \subset M$. This generalizes the notion of a link diagram, and has been studied with a variety of conditions [1, 5, 12, 21, 23, 31]. We will work with the *weakly generalized alternating knots and links* of Howie and Purcell [23], which gives certain criteria to ensure the diagram is interesting. A *cusp* of a hyperbolic knot K is a neighborhood of K in $M \setminus K$, and the *cusp volume* is the Euclidean volume of a maximal cusp. We show that the cusp volume of a weakly generalized alternating knot, with some additional conditions, is bounded both above and below based on the twist number of $\pi(K)$ and $\chi(F)$. This is done by constructing a new essential surface for K that has nice properties, including having Euler characteristic based on the twist number as opposed to the crossing number. Our bound on cusp volume leads to interesting bounds on other geometric properties of K, including slope length and volumes of Dehn surgery. The *volume* of a hyperbolic link L in a manifold M is the hyperbolic volume of the complement $M \setminus L$. We can show that volume is also bounded below by the twist number of $\pi(L)$ and $\chi(F)$. We do this by generalizing the Jones polynomial to weakly generalized alternating links, and showing that there is a relation between this polynomial and the twist number, and this polynomial and the volume. Through the course of proving this bound, we also get relations between the guts of the checkerboard surfaces of L and this generalized Jones polynomial. In addition, if we are working inside a thickened surface, the twist number becomes an isotopy invariant of L.

ACKNOWLEDGEMENTS

I would like to thank my partner, Brad, for all the support and encouragement he has given me these past two years. He's kept me sane, happy, and motivated to keep moving forward.

I would also like to thank my advisor, Effie Kalfagianni, for her work and helping me in my research. Her advice and conversations on research and academia has been invaluable.

I would also like to thank my friends, Alex, Chris, Corleone, and Mark, for giving me something to look forward to every week for the past nine years.

The research done here was partially supported by the Dr. Williams L. Harkness Endowed Fellowship, and NSF Grants DMS-1404754, DMS-1708249, and DMS-2004155

TABLE OF CONTENTS

LIST OF FIGURES

CHAPTER 1

INTRODUCTION

Traditonally, knot theory is the study of embeddings of S^1 into S^3. A *knot*, then, is the image of this embedding, and a *link* a union of disjoint knots. Two knots (or links) K_1 and K_2 are said to be *equivalent* if there is a homeomorphism $f : S^3 \to S^3$ such that $f(K_1) = K_2$.

In addition to looking at links in S^3, we can also look at embeddings of S^1 into any manifold M. We call these *knots in* M, and disjoint unions of them *links in* M. When it is clear we are working in a general manifold, we will also just refer to these as knots and links.

When studying links, it is beneficial, and often easier, to work with a projection of the link onto S^2, while keeping track of crossing information. Such a projection is called a *link diagram*. We can generalize these in two different ways, both of which can occur simultaneously. First, we can look at projections of links in M onto S^2. Second, we can look at projections of L onto a surface F that are not S^2, such as the torus. In this second case, we call this projection a *link diagram on* F.

On S^2, we can use the Reidemeister moves, three local moves on a link diagram, to try to get one from one diagram to another. However, it can take a prohibitively large number of such moves to show two diagrams are equivalent. As such, it can be difficult, in general, to determine if they represent the same or different links. This is not made easier when we start dealing with diagrams on different surfaces. As such, we look to invariants to help distinguish when two diagrams represent two different links.

An *invariant* of a link is a way of associating some object to a link such that, if two links are equivalent, they are assigned the same object. A simple, but useful, invariant is the crossing number $\mathrm{cr}(K)$ of a knot K that counts the minimum number of crossings in a diagram of K. Other invariants include being alternating, which tells us if a knot has an alternating diagram, and the Jones polynomial, which associates a polynomial to a link.

Some link invariants, including the ones mentioned above, can be computed by looking

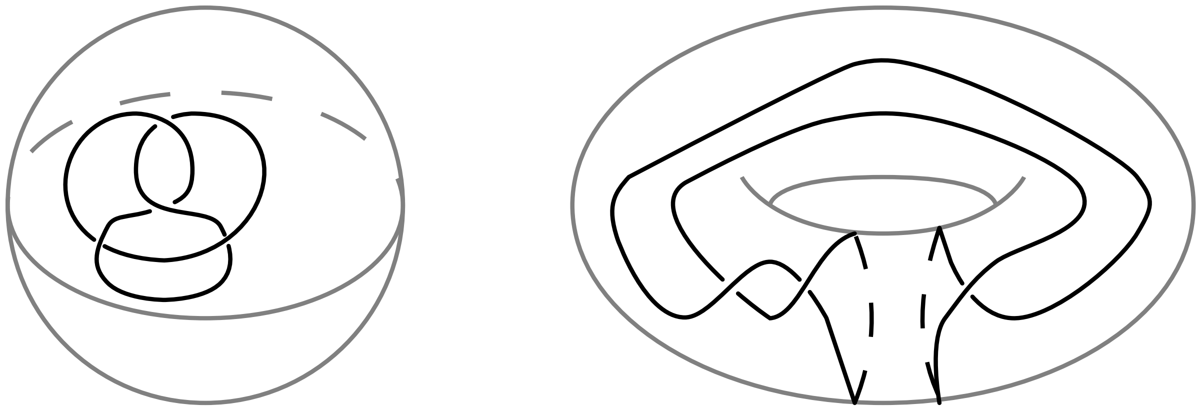

Figure 1.1. *(Left):* A diagram of the figure eight knot on S^2 *(Right):* A diagram of the unknot on the torus.

directly at a diagram of the link. These are called *diagrammatic invariants*. Another class of invairants comes from certain geometric properties of the link.

Given a link $L \subset M$, we can look at the link complement, $M \setminus L$. Many complements admit a complete finite-volume hyperbolic structure. That is, $M \setminus L$ admits a Riemannian metric of constant negative curvature. When a knot or link has such a complement in S^3, it is called a *hyperbolic knot or link*. Thurston showed that a knot is either a satellite knot, a torus knot, or a hyperbolic knot [33], so there are plenty of hyperbolic knots to study. Also, due to work of Mostow [30] and Prasad [32], these hyperbolic structures for link complements are unique up to isometry, and so invariants of the hyperbolic structure are also invariants of the link.

Two important *hyperbolic invariants* that we are interested in here are the volume of a link complement $M \setminus L$, called $\mathrm{vol}(L)$, and the cusp volume of a knot complement $M \setminus L$, called $\mathrm{CV}(K)$.

As diagrammtic invariants are computed based on just a diagram for a link, while hyperbolic invariants involve studying the geometric properties of the complement, diagrammatic invariants tend to be simpler to study. However, a natural question arises: is there a relation between diagrammatic invariants and hyperbolic invariants? For alternating knots and links, the answer is yes. Dasbach and Lin [15], and Lackenby with Agol and Dylan Thurston [25] showed that the volume of hyperbolic alternating links can be bounded both above and below by the twist number, a diagrammatic invariant, and Lackenby and Purcell [26] showed that such alternating knots have cusp volume bounded above and below by the twist number.

In this work, we will explore this relationship between these two hyperbolic invariants and diagrammatic invariants for a class of knots and links, called weakly generalized alternating links, that are either projected onto a surface that is not S^2 or are embedded in a manifold that is not S^3. In Chapter 2, we will introduce the necessary background material, including the exact definitions of the invariants we are working with, as well as the exact class of knots and links we will be examining. Then, in Chapter 3, we will construct an essential surface for this class of links. In Chapter 4, we will use this essential surface to show that cusp area for this class of links can be bounded based on a diagrammatic invariant. Finally, in Chapter 5, we will show that the volume of the link complement of this class of knots is also bounded by the same diagrammatic invariant. Sections 1.1, 1.2, and 1.3 will outline the major results of Chapters 3, 4, and 5, respectively.

1.1 Essential Surfaces of Weakly Generalized Alternating Knots

We begin our study of weakly generalized alternating knots by showing that they admit certain essential surfaces. By work of Howie and Rubinstein in S^3 [24] and Howie and Purcell in other manifolds [23], we know that these knots will admit at least two essential surfaces, in particular, the checkerboard surfaces. Our goal, then, is to construct two essential surfaces that we have more control over.

To construct the surfaces, we first augment our knot K by removing pairs of crossings from twist regions to get a new link K'. We can show that this augmented knot also remains weakly generalized alternating:

Proposition 3.3.1. *If K is a weakly generalized alternating knot with diagram $\pi(K)$ on a projection surface F in a 3-manifold Y, with $e(\pi(K), F) \geq 4$ and $r(\pi(K), F) > 4$, then K' is a weakly generalized alternating knot on the surface F in Y with $e(\pi(K'), F) \geq 4$ and $r(\pi(K'), F) > 4$.*

This means, then, that the checkerboard surfaces for our augmented knot must be essential.

We can pull these back to our original knot, K, and get what are called *twisted surfaces*, $S_{B,2}$ and $S_{R,2}$. Our goal, then, is to prove these are essential.

Theorem 3.1.1. *Let $f : S_{B,2} \to Y \setminus K$ be the immersion of our blue twisted surface $S_{B,2}$ into our knot complement $Y \setminus K$, where Y has no boundary components. Then this immersion is π_1-injective and boundary π_1-injective provided that $N_{tw} \geq 121$.*

We prove this by examining how a disk with boundary essential on one of our surfaces must intersect the the twisted surfaces. By visualizing these intersections as colored graphs, we show that these graphs must have simultaneously at least five adjacent bigons and less than five adjacent bigons.

Finally, we finish the chapter by modifying our proof of Theorem 3.1.1 to show that these twisted surfaces have the additional nice property:

Theorem 4.2.4. *Let S be the disjoint union of our twisted surfaces $S_{B,2} \sqcup S_{R,2}$. Suppose that two distinct essential arcs in the surface $S_{B,2}$ have homotopic images in $Y \setminus K$, but not $S_{B,2}$. Then the two arcs are homotopic in $S_{B,2}$ into the same subsurface associated with some twist region of K_2.*

Here, a *subsurface associated with some twist region* is a neighborhood of a twist region. This will allow us to limit how many essential arcs in either twisted surface are not homotopic in $Y \setminus K$ in the following chapter, based on the twist number of K. This contrasts to the case of checkerboard surfaces, where such essential arcs are bounded based on the crossing number, which grows much more quickly than the twist number.

1.2 Cusp Volumes of Weakly Generalized Alternating Knots

Given a finite volume hyperbolic manifold M, the ends for M will look like a half-open neighborhood of a torus, $T^2 \times [1, \infty)$. We call this end a *cusp*. If $p : \mathbb{H}^3 \to M$ is the covering map, then we can view the preimage of this cusp as a collection of horoballs. A cusp is called *maximal* if this preimage is tangent to itself.

Given a maximal cusp, we say the *cusp volume* is the Euclidean volume of the cusp, and the *cusp area* the euclidean area of the boundary of the cusp. We can use these, then, to examine the lengths of slopes of a knot, as well as study how Dehn surgery affects our manifold.

The cusp volume has been studied in a couple of settings for knots, although not as extensively as other geometric invariants such as the volume. Futer, Kalfagianni, and Purcell gave two-sided diagrammatic bounds for the cusp volume of closed 3-braids and 2-bridge knots [19]. Lackenby and Purcell were also able to find two sided bounds for the cusp volume of traditional alternating knots in S^3 based on the twist number.

Using the essential surfaces we construct in Chapter 3, we generalize the work of Lackenby and Purcell to get the following result:

Theorem 4.1.1. *Let Y be a closed, irreducible 3-manifold and $F \subset Y$ a closed surface with $\chi(F) \leq 0$ and such that $Y \setminus N(F)$ is atoroidal. Suppose that K is a weakly generalized alternating knot with a projection $\pi(K) \subset F$ that is weakly twist reduced, and we have $e(\pi(K), F) > 2$, and $r(\pi(K), F) > 4$. Then $Y \setminus K$ is hyperbolic, and we have*

$$\frac{5.8265514 \times 10^{-7}}{2\,(120 - \chi(F))^6}\,(\mathrm{tw}_\pi(K) - \chi(F)) \leq CV(K) \leq \frac{(360\mathrm{tw}_\pi(K) - 3\chi(F))^2}{\mathrm{tw}_\pi(K)},$$

where $CV(K)$ is the cusp volume of K and $\mathrm{tw}_\pi(K)$ the twist number of $\pi(K)$.

By some calculus, the cusp area is exactly twice the cusp volume, and so this bound also gives us a bound on the cusp area. Then, by using a result of Burton and Kalfagianni [11], we get the following bounds for the slopes of our knots:

Corollary 4.1.3. *Let Y be a closed, irreducible 3-manifold and $F \subset Y$ a closed surface with $\chi(F) < 0$ and such that $Y \setminus N(F)$ is atoroidal. Suppose that K is a weakly generalized alternating knot with a projection $\pi(K) \subset F$ that is weakly twist reduced, and we have $e(\pi(K), F) > 2$, and $r(\pi(K), F) > 4$. Let μ be the meridian and σ any non-meridian slope* *on maximal cusp of $Y \setminus K$. Then, we have*

$$\ell(\mu) \leq \frac{720\ (\mathrm{tw}_\pi(K) - \chi(F))}{\mathrm{tw}_\pi(K)}$$
$$\ell(\sigma) \geq \left(\frac{5.8265514 \times 10^{-7}}{720\ (120 - \chi(F))^6}\right)\ \mathrm{tw}_\pi(K),$$

where $\mathrm{tw}_\pi(K)$ *is the twist number of* $\pi(K)$.

Then, by combining with theorems from Howie and Purcell [23], Futer, Kalfagianni, and Purcell [18], and W. Thurston [33], we can get the following two sided bounds for the volume of manifolds obtained by Dehn surgery of highly twisted weakly generalized alternating knots:

Theorem 4.1.4. *Let K be a weakly generalized alternating knot on a surface $F \neq S^2$ in a closed manifold Y with conditions as in Theorem 4.1.1. Suppose that*

$$\mathrm{tw}_\pi(K) > \left(2.32928 \times 10^{10}\right)\ (120 - \chi(F))^6\,.$$

Then, any manifold M obtained from non-meridional surgery along K is hyperbolic and

$$\frac{v_8}{4}\ (\mathrm{tw}_\pi(K) - \chi(F)) \leq \frac{1}{2}\mathrm{vol}(Y \setminus K) \leq \mathrm{vol}(M) < \mathrm{vol}(Y \setminus K).$$

Here $v_8 \approx 3.66386$ *is the volume of a regular hyperbolic ideal octahedron.*

1.3 Twist Number, Jones-like polynomials, and the Volume

Given a link L in S^3, it is possible to manually construct the link complement from a diagram. However, it is impractical to do this for every hyperbolic link. Instead, we often look to generalize these constructions and find bounds for classes of links based on their diagrams.

Adams [3] was able to show that, if a hyperbolic knot has at least $c \geq 5$ crossings, then

$$\mathrm{vol}(S^3 \setminus K) \leq (c-5)v_{\mathrm{oct}} + 4v_{\mathrm{tet}},$$

where $v_{\mathrm{tet}} = 1.0149\ldots$ is the volume of a regular ideal tetrahedron and $v_{\mathrm{oct}} = 3.6638\ldots$ is the volume of a regular ideal octahedron. This was done by decomposing the complement into ideal tetrahedrons and octahedrons, and counting.

While the crossing number will give us an upper bound, it also tends to get large very quickly. As such, we also look towards other diagrammatic invariants to give us bounds on the volume. One of the more common ones is the twist region. By work of Lackenby, with an improvement on the upper bound by Agol and D. Thurston [25], we get the following bounds for alternating links:

$$\frac{v_{\text{oct}}}{2}(\text{tw}(L) - 2)\text{vol}(S^3 \setminus L) \leq 10v_{\text{tet}}(\text{tw}(L) - 1),$$

with the upper bound applying to non-alternating links, as well. The upper bound is found by looking at a polyhedral decomposition of an augmentation of L, while the lower bound is found by using a tool called the *guts of surfaces*.

There is also a way of relating the coefficients of the Jones polynomial to the volume. Described more completely in Chapter 2, the *Jones polynomial* is an invariant of knots that can be found based on the diagram of the knot. Dasbach and Lin [15] showed that, in S^3, the first two and last two coefficients of the Jones polynomial determine the twist number. Then we get the following bounds for the volume of alternating links

$$v_{\text{oct}}\ \max(|a_{m-1}|, |b_{n+1}|) \leq \text{vol}(S^3 \setminus L) \leq 10v_{\text{tet}}\ (|a_{m-1}| + |b_{n-1}| - 1)$$

where a_{m-1} is the second coefficient and b_{n-1} is the second to last coefficent of the Jones polynomial. The lower bound here is also found by examining the guts of surfaces.

We are able to generalize this lower bound of Dasbach and Lin for weakly generalized alternating links. We do this by first generalizing the Jones polynomial to these settings, and relating the twist number to it's coefficients:

Theorem 5.1.2. *Let M and F be as in Theorem 5.1.1 Suppose that a link L admits a weakly generalized alternating projection $\pi(L) \subset F$ that is twist-reduced and with all regions of $F \setminus \pi(L)$ disks. Then we have the following.*

1. $\chi(\text{guts}(M_A)) = -|a_{m-1}| + \frac{1}{2}\chi(M)$

2. $\chi(\text{guts}(M_B)) = -|b_{n+1}| + \frac{1}{2}\chi(M)$

3. $\mathrm{tw}_F(\pi(L)) = |a_{m-1}| + |b_{n+1}| + \chi(F)$.

Further, if $\pi(L)$ is a cellularly embedded, twist-reduced, reduced alternating diagram, we still get (3).

From there, we can use the relationship to the guts to get a lower bound for our volume:

Theorem 5.1.1. *Let M be an irreducible, compact 3-manifold with empty or incompressible boundary. Let $F \subset M$ be an incompressible, closed, orientable surface such that $M \setminus N(F)$ is atoroidal and $\partial - anannular$. Suppose that a link L admits a weakly generalized alternating projection $\pi(L) \subset F$ that is reduced, twist-reduced and with all regions of $F \setminus \pi(L)$ disks. Finally, suppose that $r(\pi(L), F) > 4$. Then L is hyperbolic and*

$$\mathrm{vol}(M \setminus L) \geq v_8 \max\{|a_{m-1}|, |b_{n+1}|\} - \frac{1}{2}\chi(\partial M),$$

where $v_8 = 3.66386\ldots$ is the volume of an ideal octahedron.

When we are working inside a thickened surface, we get some interesting consequences. In general, the twist number of a weakly generalized alternating link is not an isotopy invariant. That is, two different diagrams for such a link might have different twist numbers. However, as discussed by Boden, Karimi, and Sikora [9], the Jones-like polynomial we define is an invariant for oriented links in a thickened surface. As such, we get the following consequence:

Corollary 5.1.3. *Let L be a link in $F \times [-1, 1]$ that admits a weakly generalized alternating projection $\pi(L) \subset F$ that is reduced, twist-reduced and with all regions of $F \setminus \pi(L)$ disks. Then any two projections of Lhave the same twist number. That is, $\mathrm{tw}_F(\pi(L))$ is an isotopy invariant of L.*

CHAPTER 2

BACKGROUND

In this chapter, we will review the relevant background information that will be used through this work. First, in Section 2.1, we recall what a hyperbolic manifold is, as well as some important definitions and results. This will include the guts and cusps of a manifold. Next, in Section 2.2, we go over the concept of generalizing alternating knots; in particular, we give the definition of a weakly generalized alternating knot, and some of it's properties. Finally, in Section 2.3, we review diagrammatic properties of knot diagrams, as well as how some of these properties can be generalized to generalized knot diagrams.

2.1 Hyperbolic Geometry

In this section, we will quickly recall the definition of a hyperbolic manifold. Then, we will recall some results that make hyperbolic manifolds appealing to work with. Finally, we will introduce certain properties of hyperbolic manifolds, in particular the volume and the cusp area, that we will be working with.

2.1.1 Hyperbolic Manifolds

We will work with the upper half-space model of hyperbolic 3-space. That is, we take $\mathbb{H}^3 = \{(x+iy,t) \in \mathbb{C}\times\mathbb{R} : t > 0\}$, and give it the metric $ds^2 = \frac{dx^2+dy^2+dt^2}{t^2}$ The isometries of $\mathbb{H}^3$ correspond to elements of $\mathrm{PSL}(2,\mathbb{C}) = \mathrm{SL}(2,\mathbb{C})/\pm I$, where $\mathrm{SL}(2,C)$ is the group of 2×2 matrices of determinant 1 and I is the identity matrix.

Now let X be some manifold, and G a group acting on X. We say a manifold M has a (G,X)-*structure* on it if, for every point $x \in M$, there is a chart (U,ϕ), with $U \subset M$ a neighborhood of x and $\phi : U \to \phi(U) \subset X$ a homeomorphism. We also require charts to work nicely together; that is, if two charts (U,ϕ) and (V,ψ) overlap, then the *transition map* $\phi\circ\psi^{-1} : \psi(U\cap V) \to \phi(U\cap V)$ should also be an element of G. In particular, we say a manifold M has a *hyperbolic structure* if it has a $(\mathrm{PSL}(2,\mathbb{C}),\mathbb{H}^3)$-structure.

When a manifold has a hyperbolic structure, it also inherits a hyperbolic metric from $\mathbb{H}^3$, allowing us to ask questions such as what is the volume of a hyperbolic manifold. In this work, we will work exclusively with complete hyperbolic manifolds of finite volume.

An important theorem about the structures of a complete hyperbolic 3-manifold with finite volume comes from Mostow [30] and Prassad [32]; namely, there is only one hyperbolic structure, up to isometry:

Theorem 2.1.1 (Mostow-Prassad Rigidity)**.** *Suppose M_1 and M_2 are complete hyperbolic 3-manifolds with finite volume. Then any isomorphism of the fundamental groups $f : \pi_1(M_1) \to \pi_1(M_2)$ is realized by a unique isometry.*

This allows us to use the geometric properties of complete finite volume hyperbolic manifolds as topological invariants, as well.

2.1.2 Guts of a Manifold

As mentioned above, we can calculate $\mathrm{vol}(M)$, the volume of a hyperbolic manifold M. We can find a lower bound on $\mathrm{vol}(M)$ based on the *guts* of M.

Definition 2.1.2. Let M be a hyperbolic 3-manifold M, and $S \subset M$ some embedded surface. Then *M cut along S*, denoted $M\backslash\backslash S$, is the manifold obtained by taking the completion of $M \setminus S$.

Definition 2.1.3. Let M be a hyperbolic 3-manifold and $S \subset M$ a properly-embedded π_1-essential surface. Then $M\backslash\backslash S$ admits a JSJ-decomposition along essential tori and annuli into I-bundles, Seifert fibered solid tori, and the *guts of $M\backslash\backslash S$*.

The guts of $M\backslash\backslash S$, denoted $\mathrm{guts}(M)$, is exactly the parts of $M\backslash\backslash S$ that, after the decomposition, admit a hyperbolic metric.

The guts have been studied in how they relate to the volume of a hyperbolic manifold [6], as well as how they can, in knot theory, relate to the volume of hyperbolic alternating links [25], as well as the coefficients of the colored Jones polynomial [17]. In particular, we will be concerned with the following result of Agol, Storm, and Thurston [6, Corollary 2.2]:

Lemma 2.1.4. *Let M be a closed hyperbolic 3-manifold, and $S \subset M$ be an incompressible surface. Then*

$$\mathrm{vol}(M) \geq -v_8 \chi(\mathrm{guts}(M \backslash\backslash S)),$$

where $v_8 = 3.66396\ldots$ is the volume of a regular ideal octahedron.

If we choose our surface carefully, we will be able to calculate $\chi(\mathrm{guts}(M\backslash\backslash S))$ by looking at just the surface and ∂M.

2.1.3 Cusp Volume and Cusp Area

Let M be a finite-volume hyperbolic manifold with ∂M a collection of tori. Then the ends of M are of the form $T^2 \times [1, \infty)$. As M is hyperbolic, we have some covering map $\rho : \mathbb{H}^3 \to M$. Then if we look at the preimage of these ends, they look like horoballs in $\mathbb{H}^3$.

For each end, we can choose a horoball representative H_i. The image of this representative, $\rho(H_i) = C_i$ is called a *cusp* of M. In actuality, we get a family of such cusps as we shrink or expand H_i. In our case (a knot complement), there is only a single cusp — the place where we removed the knot. When we only have a single cusp, we can expand it as much as possible, until it becomes tangent to itself. In $\mathbb{H}^3$, this is expanding each horoball pre-image until they begin to become tangent to each other. Once we have done this, we get what is called the *maximal cusp* of M. In the case of a single cusp, this expansion is unique. We can, of course, generalize this to multi-cusped manifolds, however it is not quite as neat. Depending on the order we expand the cusps, we can get several maximal cusps, and each can have different properties.

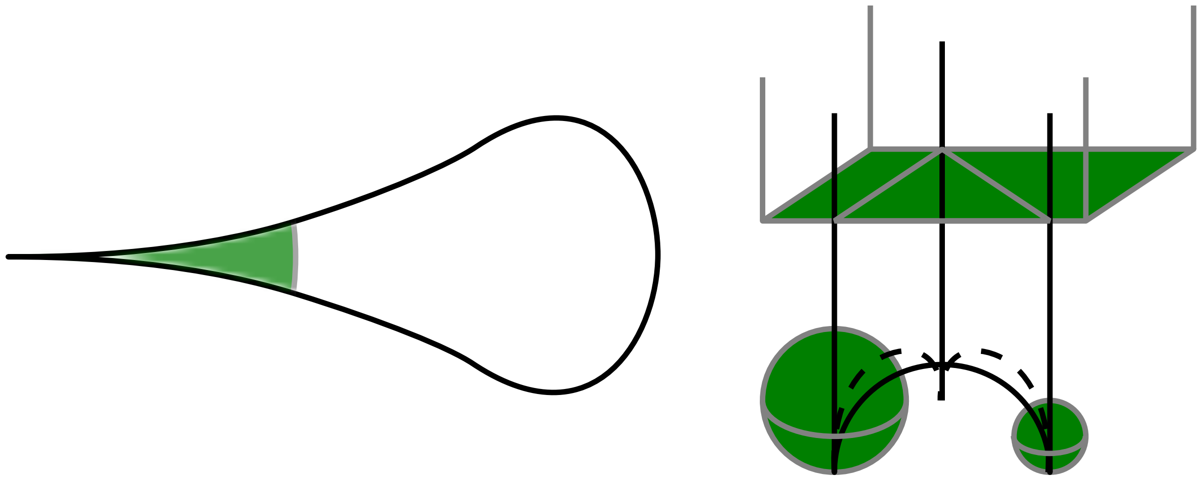

Figure 2.1. *(Left):* A manifold M with a single cusp. The green colored area is a neighborhood of the cusp. *(Right)* A preimage of M, with three different horoball representatives for the cusp. Note that the upper horoball extends to infinity.

Now, once we have our choice of cusp, C, we can start analyzing it. As our manifold M has some (hyperbolic) geometric structure to it, C will naturally inherit a geometry from it. And because we know the pre-image of C is a horoball in $\mathbb{H}^3$, it must inherit a hyperbolic structure and it's boundary must be Euclidean.

Definition 2.1.5. The *volume* of a cusp C, $\mathrm{vol}(C)$, is the Euclidean volume of C. The *area* of a cusp, $\mathrm{area}(\partial C)$, the Euclidean area of ∂C. In the case that $M = Y \setminus K$ we will use $CV(K)$ to denote the volume of the maximal cusp of M.

A quick aside about why we look at the Euclidean area: We can always view the horoball preimage of our cusp as being a neighborhood at $\infty \in \mathbb{H}^3$. This amounts to looking at the set $\{(x, y, t) \in \mathbb{H}^3 \mid t \geq z\}$ for some $z > 0$. When we restrict to the boundary of our cusp, $\partial C = \{t = z\}$, we see that our hyperbolic metric becomes a scaled Euclidean metric. Thus it make sense to talk about the cusp in terms of Euclidean area and volume.

As we can view the boundary of a cusp as $T^2 \times 1$, by some simple calculus, we can see that $\frac{1}{2}\mathrm{area}(\partial C) = \mathrm{vol}(C)$. As the volume can be found from the area, we will often focus on just finding a single value, usually the area, and seeing what that gives us.

We have several methods to help us calculate the cusp area. First, if we have a triangulation of M, we can find the area by direct computation, as has been implemented in SnapPy [14]. When we work with fully augmented links, where all crossings of a knot have been removed and crossing circles have been added in, we can completely determine the geometry by a circle packing. In this case, we can use the circle packing to find the cusp area [20].

Definition 2.1.6. Let μ be the meridian of our knot, and λ the shortest longitude. Then the lengths of these curves, $\ell(\mu)$ and $\ell(\lambda)$, determine the similarity class of the Euclidean structure on ∂C. We call this similarity class the *cusp shape*.

Knowing the cusp shape can tell us about the area of the cusp, as we have $\text{Area}(\partial C) \leq \ell(\mu)\ell(\lambda)$.

Adams, Colestock, Fowler, Gillam, and Katerman found upper bounds of slope lengths and cusp shapes of a knot based on it's crossing number [2]. Burton and Kalfagianni found upper bounds coming from essential spanning surfaces of knots [11]. We need to recall their result [11, Theorem 4.1] as we will use it in Chapter 4.

Theorem 2.1.7 (Theorem 4.1 [11])**.** *Let K be a hyperbolic knot with maximal cusp C, in an irreducible 3-manifold Y. Suppose that S_1 and S_2 are essential spanning surfaces in $M = Y \setminus K$, and let $i(\partial S_1, \partial S_2) \neq 0$ denote the minimal intersection number of ∂S_1 and ∂S_2 in ∂C. Let $\ell(\mu)$ and $\ell(\lambda)$ denote the lengths of the meridian and the shortest longitude of K, respectively. Then:*

$$\ell(\mu) \leq \frac{6|\chi(S_1)| + 6|\chi(S_2)|}{i(\partial S_1, \partial S_2)} \qquad \ell(\lambda) \leq 3|\chi(S_1)| + 3|\chi(S_2)|,$$

and

$$\text{area}(\partial C) \leq \frac{18(|\chi(S_1)| + |\chi(S_2)|)^2}{i(\partial S_1, \partial S_2)}.$$

Note that the authors in [2] state their result for links in S^3. However, as noted by Burton and Kalfagianni in the paper, the results work for any manifold Y, as long as our knot is hyperbolic.

2.2 Weakly Generalized Alternating Knots

When generalizing knots to non-S^2 surfaces inside of manifolds that are not S^3, we will use the definitions of Howie and Purcell [23]. We review these definitions in this section.

Let Y be a compact, orientable, irreducible 3-manifold, with a closed, orientable surface F, which we call a *projection surface*. Given a knot $K \subset F \times I \subset Y$, we can get a projection $\pi(K)$ by flattening $F \times I$ down to just F, and keeping track of where the knot goes. We call $\pi(K)$ a *generalized diagram*. As we are not working with $F = S^2 \subset S^3$, we need to modify some of the usual definitions we have for knots. In the setting of $F = S^2$, one uses the fact that all curves in S^2 bound disks to define primeness. Because we do not have this property for general surfaces, we modify as below:

Definition 2.2.1. We say that $\pi(K) \subset F$ is *weakly prime* if, whenever $D \subset F$ is a disk with ∂D intersecting $\pi(K)$ transversely exactly twice, then either

- $F = S^2$, and either $\pi(K) \cap D$ is a single arc or $\pi(K) \cap (F \setminus D)$ is;
- F has positive genus, and $\pi(K) \cap D$ is a single arc.

In order to generalize the proofs that follow, we will want to make sure that our knots are sufficiently complex. To do this, we will introduce two conditions on the knot:

Definition 2.2.2. The *edge representativity* $e(\pi(K), F)$ is the minimum number of intersections between $\pi(K)$ and any essential curve on F. The *representativity* $r(\pi(K), F)$ is the minimum number of intersections between $\pi(K)$ and a compression disk of F, taken over all compression disks of F. If there are no essential curves, we set $e(\pi(K), F)$ to be ∞, while if there are no compression disks, we set $e(\pi(K), F)$ and $r(\pi(K), F)$ to be ∞.

While the definition of a weakly generalized alternating knot only uses the representativity, many of our proofs in Chapter 3 will also need the edge representativity to be high enough to deal with some edge cases. Now, though, we have enough to give a definition for a weakly generalized alternating knot.

Definition 2.2.3. Let $F \subset Y$ be a projection surface as above. Then the diagram $\pi(K)$ on F of a knot K is *reduced alternating* if

1. $\pi(K)$ is alternating on F,
2. $\pi(K)$ is weakly prime,
3. $\pi(K) \cap F \neq \emptyset$,
4. $\pi(K)$ has at least one crossing on F.

If, in addition, we also have

5. $\pi(K)$ is checkerboard colorable on F,
6. $r(\pi(K), F) \geq 4$

then $\pi(K)$ is a *weakly generalized alternating diagram*, and K is a *weakly generalized alternating knot*.

This definition generalizes to links on a single surface, as well, by taking K to be a link. While we will not consider it in this work, this definition can also be generalized to links with a projection onto a disjoint union of surfaces. This includes redefining what a projection surface F is, as well as requiring that each surface has at least one crossing projected onto it. For the complete definition in this more general case, see [23].

These knots have been studied in several other papers, as well as knots that satisfy subsets of this definition. These were introduced by Howie and Rubenstein [24] in S^3 as generalizations of Hayashi [22] and Ozawa [31], and generalized further to other manifolds Y by Howie and Purcell [23].

If we ignore the final two conditions of the above definition, and just look at reduced alternating, we get as a subset generalized alternating knots in S^3, as studied by Ozawa in [31], with the additional restrictions that $Y = S^3$ and $\pi(K)$ is strongly prime, which implies weakly prime. These knots are also checkerboard colorable, although the representativity is only at least 2 [31, Theorem 2.2]. If we don't require $\pi(K)$ to be weakly prime, and let F to be a Heegaard torus, we get toroidally alternating knots, as studied by Adams [5]. More broadly, this category also includes alternating knots on a Heegaard surface F, as studied by Hayashi [22]. This also fits the alternating projection of a knot $\pi(K)$ onto it's Turaev surface F which also has the property that $F \setminus \pi(K)$ are disks [16]. By contrast, for reduced alternating and weakly generalized alternating knots, F does not need to be a Heegaard surface, nor does $F \setminus \pi(K)$ need to be disks. In addition, the requirement that $r(\pi(K), F)$ be large enough also often guarantees that our diagram will be sufficiently complicated.

We will sometimes narrow our definition by requiring the representativity $r(\pi(K), F)$ to be strictly greater than 4, and adding in that the edge representativity $e(\pi(K), F)$ is at least 4. These restrictions will allow us to later rule out certain troublesome cases caused by working on a surface with genus. We will also define twist regions and twist reduced in this general case, following Howie and Purcell's definitions [23]:

Definition 2.2.4. Let $\pi(K)$ be a weakly generalized alternating diagram on a surface F. A *twist region* of $\pi(K)$ is either a string of bigon regions arranged vertex to vertex that is maximal in the sense that no larger string of bigons contains it, or a single crossing adjacent to no bigons.

Then two crossings are part of the same *twist region* if there exists a disk $D \subset F$, with ∂D intersecting our knot exactly four times - transversely at each of our crossings.

Definition 2.2.5. A diagram $\pi(K)$ on F is *weakly twist reduced* if every disk D in F, with ∂D meeting $\pi(K)$ in exactly two crossings, either contains only bigon faces of $F \setminus \pi(K)$ or $F \setminus D$ contains a disk D', where $\partial D'$ meets $\pi(K)$ in the same two crossings and bounds only bigon disks.

Note that this differs from the traditional definition of twist reduced in the following way. When $F = S^2$, we say a diagram is twist reduced if, whenever a simple closed curve intersects our diagram at exactly two crossings, that curve bounds a bigon region.

Finally, given a diagram $\pi(K)$ on F, we can ask about how many twist regions it might have:

Definition 2.2.6. The *twist number* of a weakly twist reduced diagram $\pi(K)$ on a surface F, labelled $\mathrm{tw}_F(\pi(K))$, is the number of twist regions in $\pi(K)$.

Note that, while the twist number is a diagram invariant for these generalized knots and links, it is not, in general, a link invariant. However, in the case of alternating links on S^2, by the use of flyping, we have a series of moves that will get us from any alternating diagram to another alternating diagram, and so the twist number is an invariant of alternating links [29]. Later, in Chapter 5, we will generalize another proof that the twist number is an invariant of alternating links [15] to show that, if we are working in a thickened surface, we can get a similar result for our generalized links.

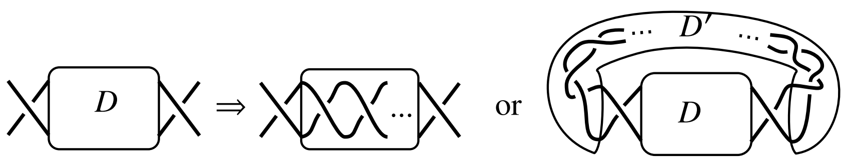

Figure 2.2. If a disk D has boundary intersecting our knot at exactly two crossings, then it either contains a twist region *(left)* or there is another disk D' intersecting at the same crossings containing a twist region *(right)*.

We will need the following result of Howie and Purcell that allows us to determine when a knot or link is hyperbolic based on from a generalized alternating projection and bounds the volume with diagrammatic quantities.

Theorem 2.2.7 (Theorem 1.1, [23]). *Let $\pi(L)$ be a weakly generalized alternating projection of a link L onto a generalized projection surface F in a 3-manifold Y. Suppose Y is compact, orientable, irreducible, and has empty boundary. Finally, suppose $Y \setminus N(F)$ is atoroidal. If F has genus at least one, the regions in the complement of $\pi(L)$ on F are disks, and the representativity $r(\pi(L), F) > 4$, then*

1. *$Y \setminus L$ is hyperbolic*

2. *$Y \setminus L$ admits two checkerboard surfaces that are esssential and quasifuchsian.*

3. *The hyperbolic volume of $Y \setminus L$ is bounded below by a function of the twist number of $\pi(L)$ and the Euler characteristic of F:*

$$\mathrm{vol}(Y \setminus L) \geq \frac{v_8}{2}(\mathrm{tw}_\pi(L) - \chi(F))$$

In order to use this theorem, we must show that our weakly generalized alternating diagrams have disk regions. We do this below for the cases we consider in Chapters 3 and 4, when the edge representativity is at least 4.

Lemma 2.2.8. *Let $\pi(K)$ be a weakly generalized alternating projection of K onto a generalized projection surface F. If $e(\pi(K), F) \geq 4$, then the regions of $F \setminus \pi(K)$ are all disks.*

Proof. Suppose not. Then one of the regions of $F \setminus \pi(K)$ either contains an annulus whose core is essential or meets itself at a crossing. If the region has an annulus, take γ to be the core. Then γ doesn't intersect $\pi(K)$, and so $e(\pi(K), F) = 0$, a contradiction. If a region meets itself at a crossing, let γ be the curve that meets this crossing and then connects back to itself in the region. Then γ intersects $\pi(K)$ exactly twice (at just the crossing), so $e(\pi(K), F) \leq 2$, a contradiction. Thus if $e(\pi(K), F) \geq 4$, then the regions of $F \setminus \pi(K)$ are all disks. □

2.3 Kauffman States and the Jones Polynomial

Finally, we will recall Kauffman states and the Jones polynomial for knots and links on $S^2 \subset S^3$. We will leave, however, the generalizations of these until Chapter 5.

A well-studied tool in the field of knot theory is the Jones polynomial. In particular, it has been used to study alternating knots and links, as in [15, 17]. Of particular note, Dasbach and Lin were able to prove that the twist number of alternating links can be recovered from the coefficients of the Jones polynomial, offering another proof of invariance.

We will review a construction of the Jones polynomial for links with diagrams on $S^2 \subset S^3$ using a function on link diagrams called the *Kauffman bracket function*.

Before we give the definition of the Kauffman bracket function, we introduce some additional definitions that will help us. First, given a crossing c in a link diagram, we define the A and B resolutions of c to be the two different smoothings of the crossing in Figure 2.3.

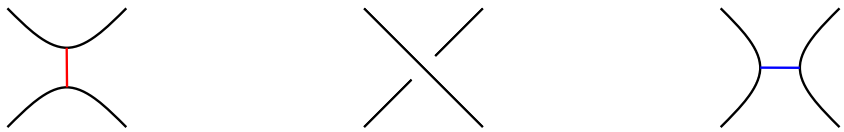

Figure 2.3. The A resolution *(left)* and the B resolution *(right)* of a crossing *(center)*.

Then a *Kauffman state* s for a diagram D is a choice of an A or B resolution at each crossing. For a state s, we can count how many A resolutions, $a(s)$, or B resolutions, $b(s)$, there are. In addition, we can ask how many *state circles*, or the number of disjoint simple closed curves, there are after we apply s to our diagram. We denote the number of state circles by $|s|$. We can now define the Kauffman bracket function.

Let $\mathcal{D}$ be the set of all unoriented link diagrams on S^2. Then the Kauffman bracket function,

$$\langle\ \rangle : \mathcal{D} \to \mathbb{Z}[A^{-1}, A]$$

is defined by the following relations

1. $\langle \diagup\!\!\!\!\diagdown \rangle = A\langle)(\rangle + A^{-1}\langle \asymp \rangle$
2. $\langle D \sqcup \bigcirc \rangle = (-A^2 - A^{-2})\langle D \rangle$
3. $\langle \bigcirc \rangle = 1.$

We can use Kauffman states to make this function more explicit:

$$\langle D \rangle = \sum_s A^{a(s)-b(s)}(-A^2 - A^{-2})^{|s|-1}.$$

We can then define the Jones polynomial of a link diagram to be

$$J(t) = \left((-1)^{w(D)} A^{-3w(D)} \langle D \rangle\right) |_{t=A^4},$$

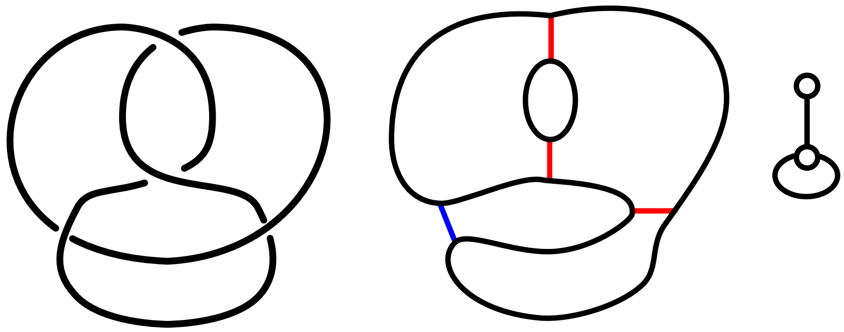

Figure 2.4. *(Left):* The figure eight knot. *(Center):* A Kauffman state for our knot. We've kept track of the type of resolutions by coloring the edges red for A resolutions and blue for B resolutions. *(Right):* The reduced graph for the state.

where $w(D)$ is the writhe of the diagram. Despite being defined using link diagrams, the Jones polynomial is an invariant of the link [28] itself.

Another useful tool coming from Kauffman states is the *state graph* and, by extension, the *reduced state graph*. We define the state graph for a Kauffman state s, G_s, to have vertices the state circles of s and edges connecting state circles that are adjacent across a crossing of the diagram. Then the reduced state graph, G_s', has the same vertices as G_s, but, if two edges of G_s are parallel, we remove one of them in G_s'.

These state graphs have been used to study information about the links. Dasbach and Lin showed that certain coefficients of the Jones polynomial can determine the twist number [15]. In particular, if L is an alternating link , a_{m-1} and b_{n+1} are the second and second last coefficients of the Jones polynomial for L, then the twist number is

$$\mathrm{tw}(L) = |a_{n-1}| + |b_{m+1}|.$$

As the Jones polynomial, and thus its coefficients, are invariants of the link, this offers another proof that the twist number is an invariant of the link.

CHAPTER 3

ESSENTIAL SURFACES

3.1 Introduction

This chapter is based on work that will be published in Algebraic & Geometric Topology [7].

In this chapter, we will construct two spanning surfaces for weakly generalized alternating knots, and then show that, under certain conditions, those surfaces are essential. This generalizes the work of Lackenby and Purcell [27] for alternating knots on the sphere. In addition to giving us new essential surfaces to work with, we will use these later in Chapter 4 to find bounds on cusp volume.

First, in Section 2, we will go through the construction of the surfaces, called $S_{B,2}$ and $S_{R,2}$, involving augmentations of our diagram. We will also go over some additional definitions we need here. Then, in Sections 3 and 4, we look at certain graphs that will allow us to study how these surfaces might intersect disks by studying colored graphs. Finally, in Section 5, we prove that, under certain conditions on the knot diagram, these spanning surfaces are essential. In particular, we will show the following theorem:

Theorem 3.1.1. *Let $f : S_{B,2} \to Y \setminus K$ be the immersion of our blue twisted surface $S_{B,2}$ into our knot complement $Y \setminus K$, where Y has no boundary components. Then this immersion is π_1-injective and boundary π_1-injective provided that $N_{tw} \geq 121$.*

Then, by a slight modification to our lemmas and proofs, we are also able to show the following theorem:

Theorem 3.1.2. *Let S be the disjoint union of our twisted surfaces $S_{B,2} \sqcup S_{R,2}$. Suppose that two distinct essential arcs in the surface $S_{B,2}$ have homotopic images in $Y \setminus K$, but not $S_{B,2}$. Then the two arcs are homotopic in $S_{B,2}$ into the same subsurface associated with some twist region of K_2.*

We will go over the exact definitions of a twisted surfaces and N_{tw} in Section 2. Then, in Section 3, we look at a specific twisted surface, by looking at a certain augmentation of our knot. In Sections 4 and 5, we study how a disk might intersect a twisted surface by introducing and studying colored graphs. Finally, in Section 6, we finish the proofs of both Theorems 4.1.1 and 4.2.4.

As some of the lemmas used for alternating knots on the sphere [27] apply to this more general case with no change, we will sometimes refer to these lemmas in this chapter. When we do, we will mention what the lemma tell us, and why we are able to use the lemma with no change to the proof.

3.2 Augmented Diagrams and Twisted Surfaces

In this section, we will construct the surfaces we are working with. Throughout, we will often omit the projection when talking about a diagram; that is, we will use K in place of $\pi(K)$.

First, let $\pi(K)$ be a weakly generalized alternating diagram on a surface F. As $\pi(K)$ is checkerboard colorable, we get two checkerboard surfaces, which we color blue and red and denote B and R, respectively.

Now, let N_{tw} be some number. We will eventually require that N_{tw} is at least 121, but for the initial construction, it can be any number we want. For each twist region of $\pi(K)$ with at least N_{tw} crossings, we add a crossing circle to get a link diagram, L.

Next, from each encircled twist region of L, we remove crossings in pairs until we have only one or two crossings remaining to get a diagram L_2. Finally, we remove the crossing circles from L_2 to get an augmented diagram K_2.

We will show in the next section that K_2 ends up being a weakly generalized alternating diagram, as well. In particular, this means that K_2 is checkerboard colorable, and so we get blue and red surfaces. In order to get our twisted surfaces, we will send these surfaces back to K by working backwards.

When we go back from K_2 to L_2, that is, when we put our crossing circles back in, each

.7

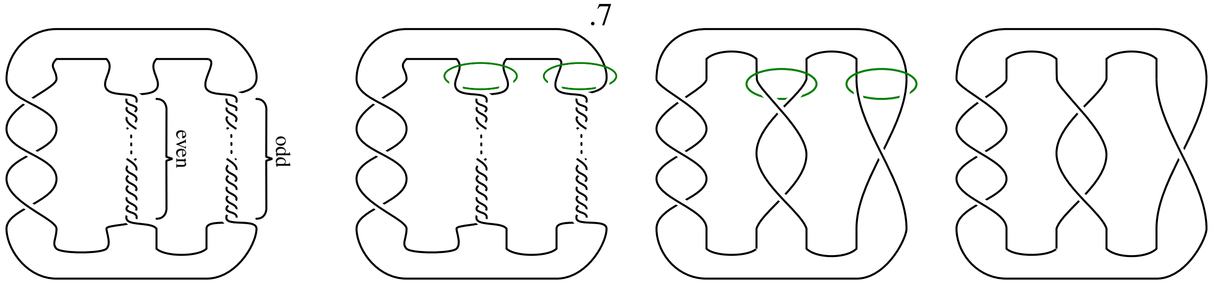

Figure 3.1. *(Left):* A knot diagram K. The center and right twist regions each have at least N_{tw} crossings with the center having an even number and the left an odd. *(Center Left:)* We adjoin circles to highly twisted regions to get L. *(Center Right:)* We remove crossing in pairs from encircled regions to get L_2. *(Right):* We remove the crossing circles to get K_2.

crossing circle will intersect either the red or the blue surface. Furthermore, when we look at $Y \setminus L_2$, each crossing circle will have a regular neighborhood that intersects either the red surface or the blue surface in a meridian disk. So then define the blue and red surfaces in the exterior of L_2 to be the blue and red surfaces of K_2 punctured by the crossing circles of L_2. We will denote these by B_2 and R_2, respectively.

Now we need to send B_2 and R_2 back to $Y \setminus L$. Adding back in the crossings is the same as twisting along the crossing circles. Most of B_2 and R_2 will twist to give us the usual checkerboard surface of L. Things will change, however, where the surfaces meet the crossing circle. If we removed $2n$ crossings from a twist region to get to L_2, we will need to twist a full n times around in order to get them back. That is, the meridian curves where our surface intersects our surfaces will go to $\pm 1/n$ curves on the boundary of the crossing circle, with sign chosen appropriately.

Finally, to get back to $Y \setminus K$, we need to fill in the crossing circles with a meridian Dehn filling. So in order to get our twisted surfaces, we will need to complete them in a way that makes sense with this filling. Note that each (meridional) cross-section of the crossing circle intersects the punctured surfaces $2n$ times on the boundary. We then connect opposite points by an interval running between them. When we do this for the whole crossing circle, we will either be attaching a single annulus or two Möbius bands, depending on the value of n. In either case, however, we now have two immersed surfaces, which we can still color blue and

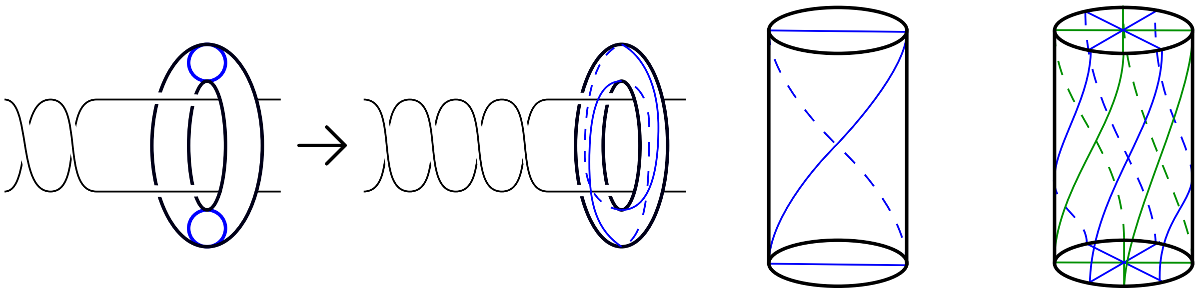

Figure 3.2. *(Left)* L_2 with the torus around the twist region. Blue circles mark where the torus punctures the blue surface. *(Center Left)* Twisting around the torus gives us back L. In this case, we are only adding in one full twist for simplicity. *(Center Right)* To complete the twisted surface, attach opposite sides of the torus with a band, as represented on top and bottom. Here this amounts to adding an Möbius band. *(Right)* When we have more twists, things can become more complicated. In addition, we might be adding two annuli instead.

red, and which we call *twisted checkerboard surfaces*. We will denote them $S_{B,2}$ and $S_{R,2}$ appropriately.

In order to prove that these twisted surfaces are essential, we will often only work with one surface at a time. As the construction of $S_{B,2}$ does not interfere with the construction of $S_{R,2}$, we can make our study of these surfaces simpler by looking at only one surface at a time. To do this, we modify our construction as follows.

When we start encircling twist regions with enough crossings, we only add crossing circles if they will intersect the blue checkerboard surfaces. We call this new link L_B. Then, when we remove the pairs of crossings from encircled regions, leaving only one or two crossings, we get a new link $L_{B,2}$. Finally, we remove the crossing circles to get $K_{B,2}$.

Now, the blue surface B_2 can also be embedded in the complement of $L_{B,2}$, while the red surface R_2 for $L_{B,2}$ is homeomorphic to the red checkerboard surface of K. This allows us, while studying our surfaces, to look at just one twisted surface at a time, say $S_{B,2}$, and R, the usual checkerboard surface for K.

3.3 Weakly Generalized Alternating Knots and Augmentation

In this section, our main goal is to show that, if we start with a weakly generalized alternating knot, with some additional properties, and then augment it by removing crossings, we will still have a weakly generalized alternating knot with those same properties.

Let K be a weakly generalized alternating knot with a closed, orientable projection surface F in a 3-manifold Y. Fix a (twist-reduced) diagram $\pi(K)$ on F with edge representativity, $e(\pi(K), F) \geq 4$, and representativity $r(\pi(K), F) > 4$.

Let L be the link obtained from K where we add a crossing circle to some number of twist regions. Also let L' the link where we remove, in pairs, all but one or two crossings from twist regions with a crossing circle in L, and K' the knot obtained from L' by removing the crossing circles. Our first goal is to prove the following which assures that the altered knot K' continues to be a weakly generalized alternating and has the same properties as K.

Proposition 3.3.1. *If K is a weakly generalized alternating knot with diagram $\pi(K)$ on a projection surface F in a 3-manifold Y, with $e(\pi(K), F) \geq 4$ and $r(\pi(K), F) > 4$, then K' is a weakly generalized alternating knot on the surface F in Y with $e(\pi(K'), F) \geq 4$ and $r(\pi(K'), F) > 4$.*

Proposition 3.3.1 allows us to look at the simpler augmented knot K' as opposed to the original K. The proof of the proposition requires a few lemmas that we will prove next.

Lemma 3.3.2. *$\pi(K')$ is reduced alternating on F.*

Proof. There are four conditions we need to show about $\pi(K')$ for it to be reduced alternating. First, as $\pi(K)$ is alternating on F, we must also have $\pi(K')$ alternating on F—we removed crossings in pairs, and so could not have introduced any non-alternating into our diagram. Next, we show that $\pi(K')$ is weakly prime, as stated in Definition 2.2.1. If it isn't, then we can find a disk $D \subset F$ such that ∂D intersects $\pi(K')$ exactly twice and $\pi(K') \cap D$ is not a single embedded arc. We can isotope D to be disjoint from crossing circles. When we put crossings back in to $\pi(K')$, we won't add any crossings to D—we only add crossings at crossing circles. But then we have a disk D whose boundary intersects $\pi(K)$ exactly twice with $\pi(K) \cap D$ not a single embedded arc. As $\pi(K)$ is reduced alternating, this is a contradiction.

Finally, we must show $\pi(K') \cap F \neq \emptyset$, and that there is at least one crossing on F. However, this follows from construction—$\pi(K')$ still lives on F, so the first condition is satisfied. For the second condition, as we never remove all crossings from a twist region, if $\pi(K')$ doesn't have a crossing on F, neither will $\pi(K)$, once again a contradiction. So then $\pi(K')$ must be reduced alternating on F. □

Lemma 3.3.3. *$\pi(K')$ is checkerboard colorable.*

Proof. As $\pi(K)$ is weakly generalized alternating, we know that $\pi(K)$ must be checkerboard colorable. We can use this coloring to obtain a coloring of $\pi(K')$. Because we remove all but one or two crossings from a region, we can obtain $\pi(K')$ from $\pi(K)$ by removing bigons from the same twist region. These bigons must have the same color, so removing them will not change the checkerboard colorability of the diagram. Using this coloring will then give us a coloring of $\pi(K')$. □

Lemma 3.3.4. *If $e(\pi(K), F) \geq 4$, then $e(\pi(K'), F) \geq 4$.*

Proof. Suppose not. Then we can find an essential curve ℓ in F that crosses our knot diagram, $\pi(K')$, a minimal number of times—zero, one, two, or three times. By isotopy, we can assume ℓ intersects $\pi(K')$ transversely away from crossings. First, we show that the number of intersections can't be odd. Because $\pi(K')$ is checkerboard colorable, we can look at the colored regions ℓ crosses through. Every time ℓ crosses $\pi(K')$, it must switch colored regions: either from red to blue or from blue to red. If ℓ intersects $\pi(K')$ an odd number of times, we would then have a curve that started in one color and ended in the other color. As it is a closed curve and our diagram is checkerboard colorable, this is impossible, so $e(\pi(K'), F)$ must be zero or two.

If $e(\pi(K'), F) = 0$, then ℓ is disjoint from $\pi(K')$. This means, in particular, that we can isotope ℓ to be disjoint from any twist region without changing the number of intersections. Each twist region in $\pi(K')$ has at least one crossing; if ℓ enters the twist region through one side, it must exit through the same side in order to not intersect $\pi(K')$. But then, as $\pi(K)$ differs from $\pi(K')$ only in twist regions, when we put the twists back in our diagram, we will have an essential curve that does not intersect $\pi(K)$, contradicting the fact that $e(\pi(K), F) = 4$.

If $e(\pi(K'), F) = 2$, then ℓ must intersect $\pi(K')$ exactly twice. First, suppose ℓ doesn't intersect a twist regions. Then, as before, when we go back to $\pi(K)$, we will have an essential curve intersecting it twice, a contradiction. So then ℓ must intersect $\pi(K')$ in a twist region. By the definition of $e(\pi(K'), F)$, we assume that ℓ intersects $\pi(K')$ transversely away from crossings. As such, in order for ℓ to intersect a twist region, it must intersect the twist region exactly twice, so we will focus on just this region. There are two options for how ℓ intersects the strands in this region: either ℓ intersects the same strand twice, or ℓ intersects both strands once. First, suppose it is the former. Then the arc of ℓ intersection the bigon of the twist region bounds a disk with the strand ℓ intersects. Using this disk, we can isotope ℓ away from the twist region, giving us an essential curve that intersects $\pi(K')$ zero times, a contradiction.

On the other hand, if ℓ intersects both strands of the twist region, we may isotope ℓ away from any crossing disks. Once we do this, when we put crossings back in to get $\pi(K)$, we will have an essential arc ℓ intersecting $\pi(K)$ exactly twice (as ℓ does not cross any crossing disks). This is a contradiction, and so we are done. □

We are now ready to give the proof of Proposition 3.3.1 .

Proof. All but the last part follows directly from the previous lemmas. Thus, all that we need to show is that $r(\pi(K'), F) > 4$. As any curve that bounds a compression disk must be essential, we can say, by Lemma 3.3.4, that we at least have $r(\pi(K'), F) \geq 4$. So suppose we can find some $\gamma \subset F$ such that γ bounds a compression disk and intersects $\pi(K')$ exactly four times. When we put the crossings back to get $\pi(K)$, we must have γ intersecting $\pi(K)$ more than 4 times, and so γ must be between a twist region. However, if γ doesn't intersect the knot strands associated to the twist region, then we have a region with zero crossings, contradicting our construction of K'. So at least two of the intersections of γ with $\pi(K')$ are part of the same twist region. But then, no matter how they intersect, we can then homotope γ to be outside of the twist region without introducing any more intersections (but possibly reducing them), as follows. Because we are working on a twist region, there must be a disk on F such that the

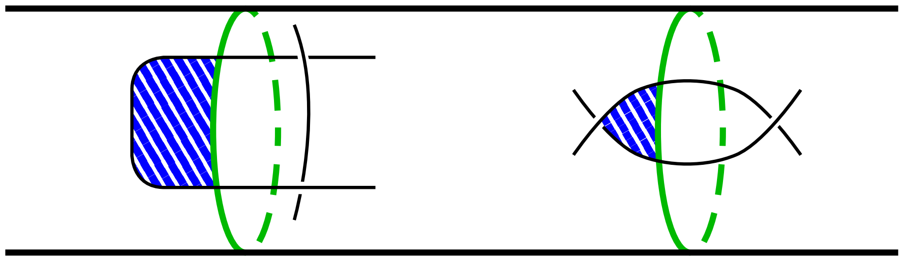

Figure 3.3. Two possible configurations to have $r(\pi(K'), F) \leq 4$. On the left, the blue disk can be used to homotope the knot away from the compression disk, reducing the number of intersections. On the right, homotopying along the blue disk will cause the compression disk to not intersect a twist region.

boundary passes through each of the crossings of our twist region. We can then use part of that disk to homotope away from the compression disk, as shown in Figure 3.3. But then, by putting crossings back in, we must have γ intersect $\pi(K)$ four or less times, a contradiction to $r(\pi(K), F) > 4$, and so we are done. □

Now we can apply Proposition 3.3.1 to our augmented knot $K_{B,2}$, and we get that $\pi(K_{B,2})$ must also be a weakly generalized alternating diagram on F with $r(\pi(K), F) > 4$ and $e(\pi(K), F) \geq 4$.

We will finish up this section by using Proposition 3.3.1 to prove two simple, but important, facts about $K_{B,2}$, as we now know some nice properties stay under augmentation. The first result is the following.

Lemma 3.3.5. *Label the regions of the complement of the diagram $\pi(K_{B,2})$ blue or red, depending on whether they meet the blue or red surface.*

(1) The blue regions on opposite sides of a crossing of $L_{B,2}$ cannot agree.

(2) The red regions on opposite sides of a crossing of $L_{B,2}$ cannot agree.

(3) The two blue regions that meet a single crossing circle of $L_{B,2}$ cannot agree.

Proof. To start, we assume that blue regions on the opposite side of a crossing do agree. Then we can create a simple closed curve on our surface F by drawing an arc in the blue region from one side of the crossing to the other, then drawing an arc across the red region near this crossing. Now we have a simple closed curve that intersects $\pi(K_{B,2})$ exactly twice. If this curve bounds a disk in F, then we contradict the fact that $\pi(K_{B,2})$ is weakly prime. If it doesn't bound a disk, then it is an essential curve intersecting our diagram exactly twice, contradicting $e(\pi(K_{B,2}), F) \geq 4$. A similar proof works to prove the second statement by switching the blue and red regions.

Now, to show the third statement, assume that two blue regions meeting a single crossing circle do agree. Then we can draw an arc entirely in the blue region from one side of the crossing circle to the other. We then finish it into a simple closed curve by attaching the arc of intersection of the crossing disk. As above, we then have a curve intersecting our diagram exactly twice, contradicting either $\pi(K_{B,2})$ being weakly prime or having edge representativity at least 4. □

This generalizes Lemma 3.9 of [27]. However, there is a fourth part in [27], stating that if distinct blue regions meeting a single crossing circle meet at the same crossing of $K_{B,2}$, then that crossing is associated with the crossing circle. This part does not carry through to the generalization of weakly generalized alternating knots. To prove part four the authors in [27] rely on a lemma which states that, in the case of usual alternating knots, K_2 $K_{B,2}$ are prime. In our case K_2 or $K_{B,2}$ are not prime, while this case can give us a simple closed curve that goes through both the crossing circle and the crossing, because it intersects our knot exactly 4 times. This curve it could very well be an essential curve, and thus tell us nothing about how the crossing relates to the crossing circle.

With $K_{B,2}$ defined in our general context, we need to introduce a new definition to generalize a diagram being blue twist reduced. As we already have a notion of weakly twist reduced, we will model our definition off of this.

Definition 3.3.6. A diagram of a link is *weakly blue twist reduced* if every disk D in F with ∂D meeting $\pi(L)$ in exactly two crossings, with sides on the blue checkerboard surface, either bounds a string of red bigons, or there is a disk D' in F meeting the diagram at the same crossings that bounds a string of red bigons.

Next, we will prove our augmented diagram is weakly blue twist reduced.

Lemma 3.3.7. *$\pi(K_{B,2})$ is weakly blue twist reduced.*

Proof. Suppose a disk D in F with ∂D meeting only the blue regions, and intersecting $\pi(K_{B,2})$ in exactly two places. We can isotope ∂D away from any crossing disk so that when we add back in crossings to get $\pi(K)$, ∂D remains disjoint from the red surface. Because $\pi(K)$ is weakly twist reduced, either D bounds a string of bigons, or there is another disk D' in $F - D$, with $\partial D'$ meeting the diagram in the same two crossings, that bounds a string of bigons. If it is the latter, then note that $\partial D'$ also only meets the blue region. In either case, these bigons must all be red bigons. Without loss of generality, suppose D contains the red bigons.

Because $\pi(K)$ is weakly twist reduced, all of the red bigons must come from the same twist region. When we remove some bigons to get back to $\pi(K_{B,2})$, we either don't touch any of the bigons inside D, in which case we are done, or we remove all but one or two crossings from the region. If this is the case, then we have to have removed all but two crossings - D intersects two distinct crossings, so there must be at least two distinct crossings left in the region. But then D contains a single red bigon, and we are done. □

3.4 Colored Graphs and Twisted Surfaces

In this section, we will prove a series of technical lemmas that we will need for the proof of Theorem 3.1.1. The lemmas aim to analyze how certain disks might intersect our twisted surfaces. The combinatorics of these intersections will be encoded by certain colored planar graphs. Before we can begin our analysis, we need to establish some notation and terminology.

3.4.1 Blue and red graphs

Recall that we need to prove that the immersed blue and red surfaces $S_{B,2}$ and $S_{R,2}$ are π_1-injective and boundary π_1-injective in $Y \setminus K$.

Working with both $S_{B,2}$ and $S_{R,2}$ simultaneously will be difficult. Instead, we will take advantage of the fact that the construction of one surface doesn't involve the other. That is, we could have constructed, say, $S_{B,2}$, the blue twisted surface, without ever having mentioned the red surface. The only part that would change is where we put the crossing circles—in order to not disturb the red surface, a crossing circle can only be added to a highly twisted region if it will intersect the blue surface. Thus in our proofs, we will often be working with the $S_{B,2}$ and R, where R is the usual red checkerboard surface for $\pi(K)$.

Recall the map $f : S_{B,2} \to Y \setminus K$ from the statement of Theorem 3.1.1. If f is not π_1−injective, then we have a disk $\phi : D \to Y \setminus K$ such that $\phi\mid_{\partial D}= f \circ l$, where l is an essential loop on $S_{B,2}$.

Let $\Gamma_B = \phi^{-1}(f(S_{B,2}))$ be the pre-image of $f(S_{B,2})$ under ϕ. By working with Γ_B, we can study the intersection $S_{B,2}$ intersects the disk $\phi(D)$. By Lemma 2.1 of [27], which also applies in our situation without changes, we have the following:

- Γ_B consists of a graph and a collection of simple closed curves.
- The vertices of the graph are points mapped to crossing circles of K.
- Each interior vertex of the graph has valence a multiple of $2n_j$, where $2n_j$ is the number of crossings removed from the twist region of $\pi(K)$ that corresponds to the twist region associated to the relevant crossing circle. Each boundary vertex has valence $n_j + 1$.

While the arguments from [27] are given in there for $Y = S^3$ and $F = S^2$, as long as our Y has no boundary components, many of the proofs require no modification. This is because, in many cases, the proofs are "local,'' and only reference a small neighborhood of where we are working. Instead of reproving everything, we will instead focus on the cases where the same proof doesn't work, requiring us to do something different, often either a completely different proof, or showing that the cases that occur when $Y \neq S^3$ or $F \neq S^2$ cannot happen.

To continue, recall the links L_B, $L_{B,2}$, and $K_{B,2}$ that we introduced in Section 2. We now have three surfaces to consider—B_2 and $S_{B,2}$ in $Y \setminus L_{B,2}$ and $Y \setminus K$, respectively; R_2, the red

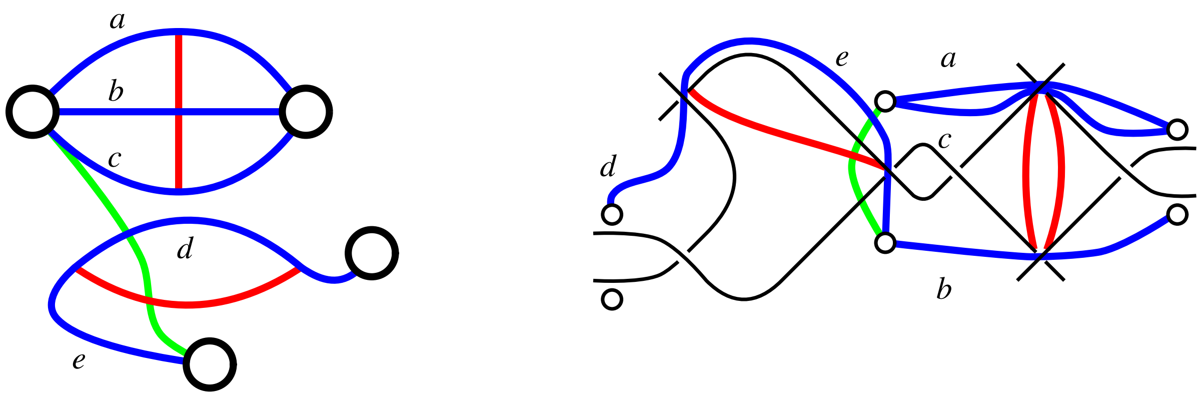

Figure 3.4. *(Left)* A possible configuration for Γ_{BRG}. *(Right)* How Γ_{BRG} would look when imposed on the diagram. The labeled blue edges in the graph map to the corresponding labeled arcs in the diagram.

checkerboard surface for $K_{B,2}$; and the crossing disks bounded by the crossing circles of $L_{B,2}$, which we color green. This also gives us three graphs to look at.

Definition 3.4.1. The *blue graph*, $\Gamma_B = \phi^{-1}(f(S_{B,2}))$, is defined as above. Next, if we remove the vertices of Γ_B (which map to crossing circles) from the disk D, we get an embedding of the punctured disk $\phi' : D' \to Y \setminus L_{B,2}$. Then the *blue-red-green graph* Γ_{BRG} is the pull-back of the union of all three surfaces—$S_{B,2}$, R_2, and the crossing disks— via ϕ'. Finally, by ignoring the edges and vertices of Γ_{BRG} coming from the green surface, we get the *blue-red graph*, Γ_{BR}. Note that $\Gamma_B \subset \Gamma_{BR} \subset \Gamma_{BRG}$.

We will be looking for what are called trivial bigons:

Definition 3.4.2. An edge of Γ_{BRG} is *trivial* if it is a blue arc disjoint from the red edges with endpoints on distinct vertices of Γ_B corresponding to the same crossing circle in L_B.

An important lemma relating to trivial arcs is the following:

Lemma 3.4.3 (Lemma 4.11 [27]). *If all but one of the edges of a region in Γ_B are trivial, then the remaining edge is also trivial*

In particular, this means if one edge of a blue bigon is trivial, then the other one must be as well. This will allow us to divide families of adjacent bigons into two cases, trivial and non-trivial. Then, as shown in Lemma 2.6 of [27], Γ_B must have more than $(N_{tw}/18) - 1$

adjacent non-trivial bigons. So, if we can show that there must be less than $5 \leq (121/18) - 1$ adjacent non-trivial bigons, then we get a contradiction. As our only assumption was that $S_{B,2}$ was not essential, we will have that our twisted surface must be essential.

Lemma 3.4.4 (Lemma 2.6, [27])**.** *The graph Γ_B must have more than $(R_{tw}/18) - 1$ adjacent non-trivial bigons, where R_{tw} is the minimum number of crossings removed from a twist region.*

Proof. The proof of Lemma 2..6 in [27] does not actually rely on either our knot or twisted surface, but is based purely on facts about the graph itself, and involves restricting to subdisks and subgraphs and removing trivial bigon families until we can get our result. As such, the proof and the lemma will continue to work in our setting, as well. □

To continue recall that we must prove that $f : S_{B,2} \to Y \setminus \mathrm{int}(N(K))$ in the statement of Theorem 3.1.1 is boundary π_1-injective: If not, then we can get a disk $\phi' : D' \to Y \setminus \mathrm{int}(K)$, where $\partial D'$ is a concatenation of two arcs, one mapped into $\partial N(K)$ and the other essential in $S_{B,2}$ [27, Lemma 2.2].

Definition 3.4.5. Define Γ'_B to be the pull-back of $f(S_{B,2}))$ via ϕ'. That is, $\Gamma'_B = \phi'^{-1}(f(S_{B,2}))$.

Once again the properties of Γ'_B remain the same if we replace $S^3 \setminus K$ from [27] with $Y \setminus K$. That is, the interior vertices of Γ'_B have valence $2n_j$, while exterior vertices have valence $n_j + 1$ (if they come from a crossing circle) or 1 (if the vertex maps to $\partial N(K)$). In particular, the following lemma from [27] works also in our setting.

Lemma 3.4.6 (Lemma 2.7 [27])**.** *The graph Γ'_B must have more than $(R_{tw}/18) - 1$ adjacent non-trivial bigons or more than $(R_{tw}/18) - 1$ adjacent triangles, where one edge of each triangle lies on $\phi'^{-1}(\partial N(K))$.*

There are a couple of important things to note about the last two lemmas. First, while R_{tw} is not necessarily equal to N_{tw}, we do know that $R_{tw} \geq 2\lceil N_{tw}/2 \rceil - 2$. With $N_{tw} \geq 121$, we get $R_{tw} \geq 120$, so there must be more than 5 adjacent non-trivial bigons in Γ_B. Second, our choice of $N_{tw} \geq 121$ is not optimal for these lemmas. In fact, the proof that $S_{B,2}$ is essential only needs $N_{tw} \geq 91$, so we have at least 5 adjacent bigons, or 5 adjacent triangles (we will prove this in the next section). The choice of 121 comes from a modification that will get us our final proof, Theorem 4.2.4.

3.4.2 Bigons in Γ_{BRG}

Here we will begin to study the structure and combinatorial properties of the graphs Γ_B, Γ_{BR} and Γ_{BRG} defined in the previous subsection. From the last subsection we know that these graphs have a certain number of adjacent non-trivial (blue) bigons, we will be looking at how the disk must intersect other surfaces near the blue surface. Our goal here is to study the possible configurations of these bigons and successively rule out several of then. By the end of this section, we will be left with only one outcome—the family of bigons must intersect the red surface at least twice.

Recall that Γ_B is the pull back of $S_{B,2}$ of a certain map $\phi : D \to Y \setminus K$ and that Γ_{BRG} is the pull-back of $S_{B,2}$, as well as the red surface R_2 and the green surface, of a map $\phi : D' \to Y \setminus L_{B,2}$ and that Γ_{BR} is obtained by deleting the green edges of Γ_{BRG}. Below we summarize some basic facts about these graphs that we will be using:

- As outside of the twisted parts, the twist surfaces are embedded, the edges of Γ_B (where the disk meets the blue surface) can only meet each other at the twisted components (where the crossing circles of $L_{B,2}$ intersected our blue surface). As such, the vertices of Γ_B are exactly the points that map to crossing circles of $L_{B,2}$. This means that, if two blue edges are adjacent to each other on a vertex, it must be that they are on opposite sides of the crossing circle (as they must meet on the twisted part).

- In addition to what we have for Γ_B, we also know that red edges can't intersect red edges, nor can green edges intersect green edges (as both red and green surfaces are embedded).

- Red edges can meet green edges (at the crossing disk). Likewise, green edges can meet blue edges at vertices of Γ_B (where the crossing disk meets $S_{B,2}$), and red edges can meet blue edges at crossings of the diagram. In addition, as D' crosses the projection surface when it meets a blue or red edge, regions of $D' \setminus \Gamma_{BR}$ are mapped above or below the projection surface, switching from one to the other when they meet one of these edges.

- Using a complexity-minimizing argument, where complexity is measured by the number of vertices in each graph, we get that there are no green edges, disjoint from blue, with both endpoints on red; and no green edges, disjoint from red, with both endpoints on a blue edge and one endpoint not a vertex. In addition, green edges cannot have both of their endpoints on the same blue vertex. In all three of these cases, we can use the arcs in question, in addition to the crossing disk, to find a homotopy that reduces how many (non-blue) vertices are in our graphs. So as long as we assume our disk's image, $\phi(D')$, and the induced graphs are minimal, these cannot happen.

As before, the lemmas are generalizations of the case when $F = S^2$ and $Y = S^3$. Most of the generalizations are relatively minor, with only a few complications brought in by the fact that we aren't working on a sphere anymore. For completeness sake, we will also talk about the proofs that still hold over from looking at the case of diagrams on S^2 [27], although in much broader and looser terms.

First, we will look at a potential blue-red bigon—a bigon in Γ_{BRG} with one side colored blue and one side colored red—and show that it cannot exist. We do this as follows:

Lemma 3.4.7. *The graph Γ_{BR} has no bigons with one blue side and one red, whether or not the bigon meets the green surface.*

Proof. If there is a bigon with only one red side and one blue side, then it must be mapped completely above or completely below the projection surface. Without loss of generality, we may assume it is mapped above. Then the union of the two sides of our bigon will form a simple closed curve γ meeting $\pi(K_{B,2})$ exactly twice with a crossing on either side. If this curve bounds a disk, we contradict $\pi(K_{B,2})$ being weakly prime. Because we are not working on the sphere, though, we are not guaranteed this. Instead, we must use the fact that the edge representativity (see Definition 2.2.2) of $\pi(K_{B,2})$ is at least four. If γ does not bound a disk, then we have an essential curve intersecting our knot transversely exactly twice, a contradiction of 3.3.4. So then γ must bound a disk with a crossing on either side, and so we get our desired contradiction. $\square$

Putting our observations above together, we get two important results. First, our disk D only meets the blue surface in “interesting” places—that is, near or through the crossing disks. Second, there are no monogons in Γ_B: any that existed couldn't meet either the red or green surfaces, as mentioned above, and so would have to connect two opposite sides of the crossing circle. However, by Lemma 3.3.5(3), this can't happen. While this later is not used immediately, we will use it later on, and so state it as a lemma:

Lemma 3.4.8 (Lemma 4.7 [27])**.** *The graph Γ_B has no monogons.*

As we want to work eventually with non-trivial bigons, it's important that we understand how trivial arcs and bigons work. As stated along with the definition, if one edge of a bigon is trivial, so is the other, giving us our trivial bigon families. In addition to this, we also find that trivial arcs have some restrictions to them—they must have at least one endpoint on the boundary of our disk, but cannot be a subset of the boundary. With that, we will set this aside to start looking at non-trivial bigons.

Lemma 3.4.9. *In the graph Γ_{BRG}, there are no non-trivial bigons with two blue sides, disjoint from red and green edges.*

Proof. Suppose we had such a non-trivial bigon. Because the bigon is non-trivial, the two vertices must correspond to separate crossing circles. When we put back in the crossings to get $\pi(K)$, we get a simple closed curve γ from the bigon. The blue edges of the bigon remain disjoint from any crossings, and connect across where the crossing disks met the projection surface. In particular, γ intersects $\pi(K)$ in exactly two crossings. We want to show that γ bounds a disk in our projection surface F—this will allow us to appeal to $\pi(K)$ being twist reduced and will give us a contradiction. So suppose γ doesn't bound a disk in F. Then the curve must be essential.

If γ bounds a compression disk, and $r(\pi(K), F) > 4$, then, after a small isotopy to take γ away from crossings, we get a curve bounding a compression disk that intersects our diagram only four times, a contradiction. So now we need to see what happens if γ doesn't bound a compression disk.

As γ bounds a bigon disjoint from red and green edges, it must bound a disk in a region in the pre-image of the blue surface in our disk. By our assumptions, γ can't bound a disk in F, nor can it bound a compressing disk for F. Therefore, the region it bounds must pass through our projection surface. But as our bigon is disjoint from red, this cannot happen (regions of the disk can only pass through the projection surface at blue or red edges, neither of which can be in our bigon). So then we get a contradiction of γ not bounding a compression disk.

Putting it all together, we see that we have a simple closed curve γ which must bound a disk on F. But then, as γ meets $\pi(K)$ in exactly two crossings and bounds a disk, because $\pi(K)$ is weakly twist reduced, we must have both of the crossings correspond to the same twist region, a contradiction of the bigon being non-trivial. □

So blue bigons must have some other colored edge intersecting them. A bigon intersected by another edge gives us two triangles, so it is natural to consider how adjacent triangles might act. As red cannot intersect red and green cannot intersect green, we find ourselves with four possible adjacent triangles, constructed as follows:

1. Blue-Red bigon intersected by green (this can't happen, as there are no blue-red bigons);

2. Red-Green bigon intersected by blue (we will prove this can't happen in Lemma 3.4.10 below);

3. Blue bigon intersected by green (this can't happen, as we can't have a blue-blue-green triangle by minimality, by [27] Lemma 4.16);

4. Blue bigon intersected by red.

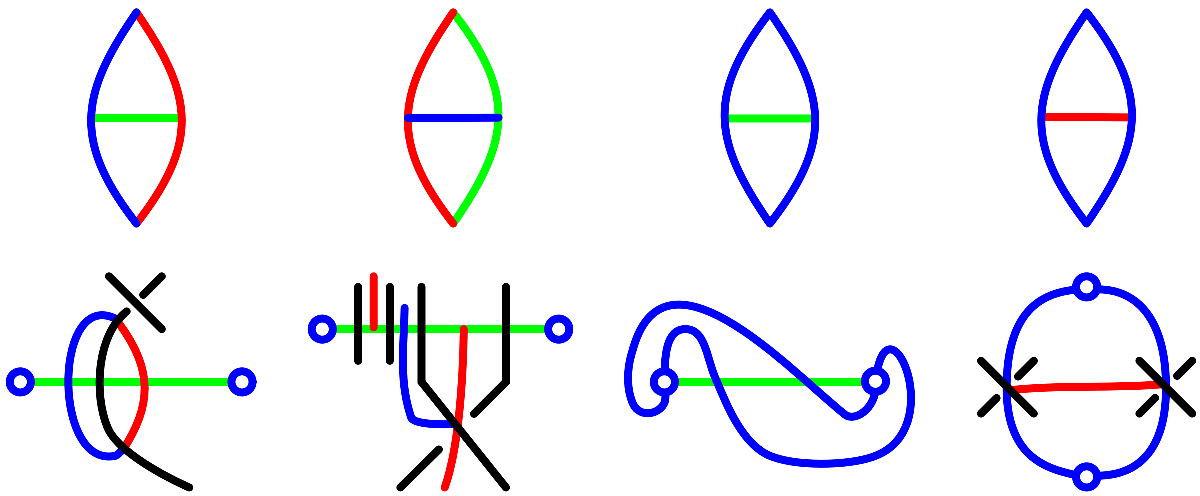

Figure 3.5. *(Top)* The four different types of adjacent triangles coming from a bigon. *(Bottom)* How the bigons might look in the diagram. Note that, in the red-green bigon case, the two red edges in the diagram are, in fact, the same edge, and must "wrap around" the diagram to create the bigon.

Lemma 3.4.10. *In Γ_{BRG}, there are no pairs of red-green-blue triangles adjacent across a blue edge. More generally, no pair of green edges can be added to Γ_{BR} in such a way that the result is a pair of triangles adjacent across a blue edge.*

Proof. We will prove the second statement. If such an edge could be added, we can find the innermost pair of triangular regions, so there are no additional red or blue edges in the triangles. Then one of these triangles must be mapped above the projection plane, and the other below. As there is only one place where red meets blue, there is only one crossing in the image of this diagram, so both red edges must meet this crossing, and meet no other crossing in our diagram.

Looking at the green edges, we see that both meet the same blue edge at the same point. This means that these green edges are related to the same crossing circle. As our two red edges meet these green edges, our red edges must both run back to the same crossing circle. So, we can create a simple closed curve with crossings on either side by following along the red edges, jumping from one to the next at either the crossing or the crossing circle. This curve intersects our knot exactly twice, so by the representativity and edge representativity of our knot, must bound a disk. However, this contradicts weak primality, so this cannot happen. ◻

The other two triangles we need to care about are blue-blue-green triangles (the third case listed above), and blue-blue-red triangles (the fourth case). In fact, with no change in proof from

Lemma 3.4.11. *The graph Γ_{BRG} contains no blue-blue-green triangles.*

Proof. Suppose we have a blue-blue-green triangle. We may assume we are working with the innermost green edge of such a triangle, and so it is disjoint from any other green edges. There are two cases to consider. First, the green edge can lie on the crossing circle of the vertex of the triangle, or second, the green edge can lie on a separate crossing circle.

If the vertex the blue edges meet at corresponds to the same crossing circle the green edge lies on, then we can construct two curves by following along one of the blue edges to the green edge, and then following the green edge back to the start of the blue edge. Each of these curves don't intersect the diagram, and so must bound disks. Then our entire triangle must lie in the union of these two disks along with a neighborhood of the green edge, which is also a disk. Then, by a similar argument as in [27], we may homotope the triangle through the crossing circle, thus removing a vertex from Γ_{BRG}. By our assumption of minimality, this gives us a contradiction.

If the vertex corresponds to a different crossing circle, we can get a curve that intersects our knot at exactly two crossings by following a blue edge down to the green edge, following the green edge to the other blue edge, then back up and across the vertex to get back to the start. We can modify this curve slightly to get a curve intersecting our knot at exactly two crossings, one associated to the vertex of the triangle and one associated to the green edge's crossing circle. Then, with a similar argument as in Lemma 3.4.9, we get that this curve must bound a disk. But then this means our two crossings must be part of the same twist region. However, as they are encircled by two different crossing circles, this is a contradiction. ◻

We will focus, then, on blue-blue-red triangles.

3.4.3 Blue-Blue-Red Triangles

Our ultimate goal is to show that there can only be so many adjacent blue bigons in Γ_B, contradicting Lemma 3.4.4 that says we need at least five. As a red edge must intersect any family of non-trivial blue bigons, we can look at blue-blue-red triangles constructed from the bigons and the red edge intersecting them.

One of the ways we will go about generalizing the results of [27] is by using a sort of "local" property—most of the proofs take place in a small subsection of the diagram, and so we don't necessarily need to know what's happening beyond that for a given proof. To take advantage of this for generalized surfaces, however, we will need to show that the subsections we care about are similar enough to subsections on a sphere. Namely, we will need to show that these small parts don't wrap around the surface, and instead lie on a disk. Once we have this, most of the proofs follow naturally. We will prove this, then, in two lemmas, dealing with two different cases.

First up, one of the more common cases we will need to deal with is where we have a triangle with two blue edges and a red edge. That such triangles induce a disk on F is split into two parts below, depending on how the red edge might pass through the crossing circle of the vertex. One important thing to note that will come up several times is that we don't actually need our triangle, or any of our future shapes, to bound a disk. In fact, looking at the case when the red edge passes through the crossing circle, this would be impossible. Instead, we just want it to lie within a disk, and hence induce a disk on F.

Before we begin, we introduce one piece of terminology. We say that a group of edges in a graph on F *induces a disk* if there is a disk on F with boundary containing those edges.

Lemma 3.4.12. *A blue-blue-red triangle in Γ_{BR} which, in the image on F, has none of its edges passing through the crossing circle of the blue vertex induces a disk on the projection surface F.*

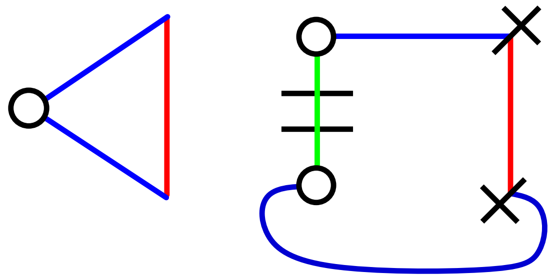

Figure 3.6. *(Left)* A blue-blue-red triangle as viewed on the graph Γ_{BR}. *(Right)* The same triangle as viewed on the knot diagram, $\pi(K_{B,2})$.

Proof. First, note that the triangle gives us a simple closed curve on F, as seen in Figure 3.6: let γ be the curve that follows our triangle along the edges, and joins the blue edges across the crossing disk associated to the vertex where the blue edges meet. By Lemma 4.18 in [27], the triangle meets two distinct crossings, and so then γ must also meet two distinct crossings. We want to show that γ bounds a disk on F. If it doesn't, there are two possibilities we will consider.

The first case: γ could bound a compression disk. If it does, then note when we put crossings back in, that γ intersects our knot in exactly four places: once across each of the crossings, and twice across the crossing disk of the vertex. But then we have a curve bounding a compression disk meeting our knot four times, contradicting the assumption that $r(\pi(K), F) > 4$.

So then if γ is essential and doesn't bound a compression disk, we can look at the triangular region γ bounds in the disk. Because γ bounds neither a compression disk nor a disk on F, the triangular region must pass through the projection surface. But then our region must have a red or blue edge in it, as it can only pass through the projection surface through an edge of Γ_{BR}. Then we don't have a blue-blue-red triangle, and so we are done. □

Lemma 3.4.13. *A blue-blue-red triangle in Γ_{BR} with the image of the red edge passing through the crossing circle of the vertex induces a disk on the projection surface F.*

Proof. As above, we know that the triangle has to meet two distinct crossings. But now, as our red edge must cross through the green edge, we are not guaranteed a single simple closed curve. Instead, we shall construct two separate disk , and then glue them together. Call our blue edges b and c and our red edge j. Let γ_1 be the curve that follows b, then j until we get to the crossing circle, and then the green edge back to b. Likewise, let γ_2 be the curve that follows c, then j until we get to the crossing circle, and then jumps back to c along the green edge.

First, each of these curves must intersect our diagram exactly twice - once at the crossings, and another time through the crossing disk. This allows us to immediately say that both γ_1 and γ_2 bound disks, as $e(\pi(K), F) \geq 4$. Next, note that γ_1 intersects γ_2 at exactly one point—where red meets green. This is because a red edge cannot intersect itself, and blue edges can only meet at vertices. But, as blue edges adjacent to the same vertex, they must map to different sides of the twist region, and so are disjoint. Thus the only point where γ_1 and γ_2 can meet is the one they share—where the crossing circle meets the red edge.

Now, we can finally construct our entire disk. Let D be the union of the disk bounded by γ_1, the disk bounded by γ_2, and a small neighborhood of the intersection of the crossing circle and red edge. As γ_1 doesn't interfere with γ_2 outside of this neighborhood, we do in fact have a disk as opposed to an annulus, and so are done. □

We can use these disks to construct even larger disks, with some caveats—we need to be careful that when we glue disks together, we aren't creating an annulus.

Lemma 3.4.14. *Families of blue-blue-red triangles, adjacent along blue edges, bound a disk.*

Proof. If we have a single blue-blue-red triangle, we are done by above. So we will suppose we have a group of them, and consider what happens when we want to add one more. As blue edges can only intersect at vertices, and red edges can never intersect, we only have to worry about two triangles—the one we are adding and the triangle that shares an edge with it. There are three cases we have to worry about, as shown in Figure 3.7. First, both triangles have their

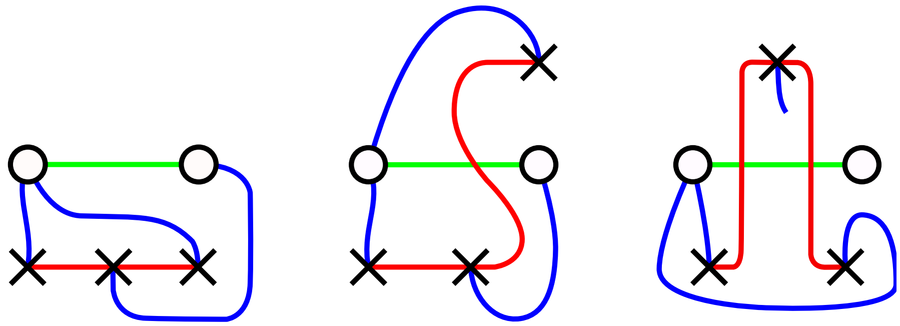

Figure 3.7. The three possible ways two blue-blue-red triangles can have their images glued together.

red edges not passing through the crossing circle. Then both triangles bound a disk, and share a single continuous segment—their common blue edge and the green edge—so we are done.

Second, one of the triangles doesn't pass through the crossing circle, but the other does. In this case, for the triangle that does pass through the crossing circle, we get two separate disks, say D_1 and D_2. The other triangle gives us a third disk, D_3. We will union these all together in the right order to get what we want. First, note that D_1 and D_3 agree only on a portion of the green edge, so gluing them together will still give us a disk. Then the boundary of D_2 only agrees with the boundary of this new disk in one continuous segment—the remainder of the green edge and the common blue edge. So we can add D_2 in and get our whole disk.

Finally, we could have both red edges pass through the crossing disk. Only one red region can pass through the a crossing disk, by definition. So for both red edges to pass through the crossing disk, they must be in the same region. But then at least one crossing, the one that meets both red edges, has opposite (red) sides agreeing, contradicting Lemma 3.3.5. So this case cannot happen, and we are done. □

Lemma 3.4.15. *A blue rectangle in Γ_{BR} with two vertices that correspond to the same crossing disk and a red diagonal has both crossings associated to the crossing circle of the vertices of the rectangle.*

Proof. Our image has to look like the one in Figure 3.8—the red edge tells us we need at least two crossings. In either case, look at the two curves a-i-f and b-i-e. Each of these curves intersects our knot transversely exactly twice. This means, then, as both the representativity

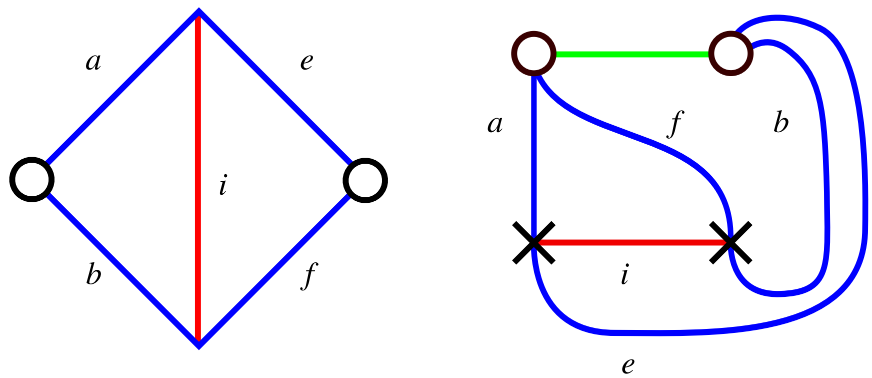

Figure 3.8. *(Left)* A blue rectangle with a red diagonal in the graph Γ_{BR}. *(Right)* A blue rectangle as viewed on the diagram, where both vertices represent the same crossing circle.

and the edge representativity are at least four, each of these curves must bound a disk. Now, though, we have two disks who agree at only one segment of the boundary. As such, we can glue the disks together without risking an annulus, and thus giving us another disk, this one with boundary a-f-e-b, and so we are done. □

To make sure we are on the right track, it is important to look at why we are studying these blue-blue-red triangles. We know that we have to have so many adjacent non-trivial blue bigons. As it turns out, given these blue bigons, we will eventually show that a red edge must pass through these bigons. If only one red passes through, we get a two collections of blue-blue-red triangles, with the families adjacent across the red edges. So it's important to see what happens here.

Before we begin studying these in earnest, there is an additional lemma that we will need here. In summary, Lemma 4.17 from [27] tells us that, if we have three adjacent blue-blue-red triangles that are adjacent at the vertex of the blue edges, two properties must hold. First, no green edge can intersect an interior blue edge (so, in figure 3.9, edges b or c). Second, if a green edge meets the interior red edge (so edge j), the green edge must run to the common vertex of the triangles. As we have families of blue-blue-red triangles bounding a disk by Lemma 3.4.14, the exact same proof as in [27] works. With this in mind, we can move on to looking at our triangles.

Lemma 3.4.16. *The graph Γ_{BR} cannot contain two pairs of three adjacent blue-blue-red triangles adjacent across the red edges.*

Proof. Suppose it did. Then we look at such a group of triangles, as pictured in Figure 3.9 (left). First, we need to show that the two vertices of the triangle group can't agree (that is, can't be related to the same crossing circle). If the vertices do agree, then the blue edges a, f, and c must be in the same region, and the blue edges e, b, and g must be in the same region. From there, we want to show that the endpoints at each of the edges are all associated with

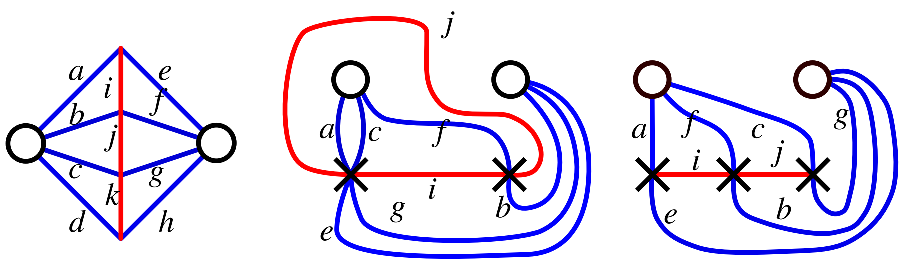

Figure 3.9. *(Left)* Two pairs of three adjacent blue-blue-red triangles. *(Center and Right)* The two possibilities if the vertices map to the same crossing circle, one with two crossings and one with three.

the crossing circle of the vertices. First, look at the a-e crossing and the b-f crossing. As a-e-f-b forms a blue rectangle with two vertices and a red edge down the diagonal, by Lemma 3.4.15, the image lies in a disk with boundary blue-blue-green. We can assume the boundary is a-e-green. But then, by definition, the crossing at a-e must be associated with the crossing disk. Then, by two applications of Lemma 3.4.15 (one with a-e-f-b and one with b-f-c-g), we get that the crossings of b-f and of c-g must be associated to the crossing circle as well.

As we are working with $L_{B,2}$, this does not automatically give us a contradiction—the diagram can still work if we have only two distinct crossings, such as in Figure 3.9. However, suppose there are only two distinct crossings for our triangle. Then we can get several curves that bound disks in F that intersect our diagram exactly twice: $c \cup i \cup f$, $e \cup i \cup b$, $f \cup a \cup j$, and $b \cup j \cup g$, as follows. To get the last curve, look at the rectangle b-f-c-g. It's disk must contain the j edge, so take the sub disk bounded by $b \cup j \cup g$. Then, for the curve $f \cup a \cup j$, note that a must be parallel to c, so use the sub disk bounded by $f \cup j \cup c$, and replace c by a.

In similar fashion, but working with the rectangle a-e-f-b and the edge i, we can get the other two curves. But then, as $\pi(K_{B,2})$ is weakly prime , we get that there are only two crossings on F.

We now have two possible options. First, $\pi(K_{B,2})$ can't be a single component knot, as any attempt to connect two crossings will result in either the unknot (which can't happen, as $e(\pi(K_{B,2}), F) \geq 4$) or a non-alternating knot on F. On the other hand, if $\pi(K_{B,2})$ is a link, by following the link around the outside, we can find a curve that completely encompasses the link. As $e(\pi(K_{B,2}), F) \geq 4$, this curve must bound a disk on F. But then we can find an essential curve that does not intersect our knot diagram at all, a contradiction. So there cannot be only two crossings for $\pi(K_{B,2})$ on F, and so the two vertices of our group of blue-blue-red triangles cannot agree.

Now we want to see how the red edges might intersect the crossing disks. First, the red edge j cannot intersect either crossing disk by Lemma 4.17 from [27], as if it did, we would have to have a green edge intersecting either b or c, contradicting Lemma 4.17 of [27]. Now, looking at i, we'll see that it can only intersect one of our crossing disks. Otherwise, if the crossing disk corresponding to e, f, and g intersects i, then we must have a green edge running from i to a. This means that f and a are in the same region. On the other hand, if the crossing disk corresponding to a, b, and c intersects i, we have a green edge running from i to e, telling us that b and e are in the same region. Then we can create two curves, one from a, b, and green, and the other from e, f, and green. Each intersects our knot transversely at the a-e crossing, but also intersects our knot at either the green edge corresponding to their respective crossing circles. But then we have a single crossing associated to two distinct crossing circles, which cannot happen. So then i can only intersect, at most, one of the crossing disks. Assume, then, it doesn't intersect the crossing disk associated to a, b, and c.

Next, we need to look at the crossings at the endpoints of a and e, which we'll call x, and at the endpoints of c and g, which we'll call y. We know that x must be distinct from the crossing at the endpoints of b and f, and y must be distinct from this third crossing as well. So we now want to show that x and y are distinct from each other. If they are not, then look at the curve $i \cup j$. This will give us a closed curve meeting both crossing disks, a contradiction to the crossing disk at a, b, and c not meeting either i or j.

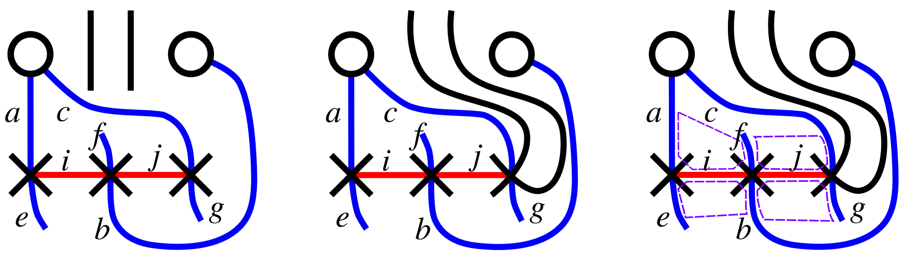

Figure 3.10. *(Left)* Focusing on the crossing circle of a, b, and c, our diagram must look something like this. *(Center)* Once we know that y is related to the crossing circle, we can draw in some extra information about our knot. *(Right)* The dotted lines are disk-bounding curves that will allow us to show x and z are also related to the crossing circle.

We want to show that this can't happen by showing the twist region associated to the shown crossing circle has three crossings—namely, the three drawn crossings. Note that edges g and e must lie in the same blue region, as they share a vertex (the crossing circle of e, f, and g) and are mapped to the same side of that vertex, but also g is bounded by $b \cup j \cup c \cup$ the green edge of the crossing disk, while e lies outside of it, so we should figure out how they meet. As g and c are at opposite sides of a crossing, by Lemma 3.3.5, g and e are not in the same region as c, so they can't cross over c. Because j is red, neither g nor e can meet it in the graph. So the only remaining options are that one of g or e meets the blue edge b or the crossing circle associated to a, b, and c. In either case, this will give us that g and e are in the same region as b.

Now we will show that the crossings x (at a-e), y (at c-g), and z (at h-f) are all associated to the same crossing circle. As we have already shown that each of these crossings are distinct, we will get our desired contradiction. First, we show that y is associated with the crossing circle of a, b, and c. To do this, we create a curve γ starting at y, following along c, using the crossing disk's green edge to get to b, and then switching over from b to g to get back to y. Because g is bounded by $b \cup j \cup c \cup$ the green edge, we can take γ to lie inside the disk on F bound by the triangle $b \cup j \cup c$ so that γ too bounds a disk on F. But now, when we put crossings back in to get K, we have a disk on F whose boundary intersects our knot at exactly two crossings, and so both must be related to the same crossing circle - namely, the crossing circle of a, b, and c. So then y is associated to this crossing circle.

In order to work with the crossings x and z, we will need to know a bit more about f—namely, that it is in the same region as a and c. Clearly, if f crosses a or c, then they must be in the same region as f. So then let's suppose it doesn't. Then note that the edge f is bounded by $a \cup c \cup j \cup i$. As f is blue, it can't cross the red edges i or j. Likewise, neither e nor g can cross i or j. So, in order for these three blue edges to form the triangles in our graph, the crossing disk of e, f, and g must enclose either i or j. By the same reason the crossing disk of a, b, and c can't meet j, the crossing disk of e, f, and g can't meet j either. So then the crossing disk must intersect i. But then, as shown in Lemma 4.17 of [27], we must have the green edge of the crossing circle intersect a. Now, with the green edge how it is, we must have f be in the same region as a and c.

Our last step is to show that x, y, and z connect to form bigons, thus are three distinct crossings in $\pi(K_{B,2})$ associated to the same crossing circle, a contradiction. To do this, note that we can form four closed curves intersecting our knot exactly twice: $i \cup \frac{1}{2}b \cup \frac{1}{2}e$; $i \cup \frac{1}{2}f \cup \frac{1}{2}a$; $j \cup \frac{1}{2}g \cup \frac{1}{2}b$; and $j \cup \frac{1}{2}c \cup \frac{1}{2}f$, as seen in Figure 3.10 (right). As we are assuming that $e(\pi(K_{B,2}), F) \geq 4$, each of these curves must bound a disk. But then, because $\pi(K_{B,2})$ is a weakly prime diagram, and because there are crossings outside of these disks, the intersection of each disk and $\pi(K_{B,2})$ must be a single embedded arc. But then, drawing the result, we see that these three crossings form bigons, and so must be associated to the same crossing circle. As we are working with $K_{B,2}$, and at most two crossings can be associated with a crossing circle, we're done. □

As mentioned earlier, our want to study adjacent families of blue-blue-red triangles came

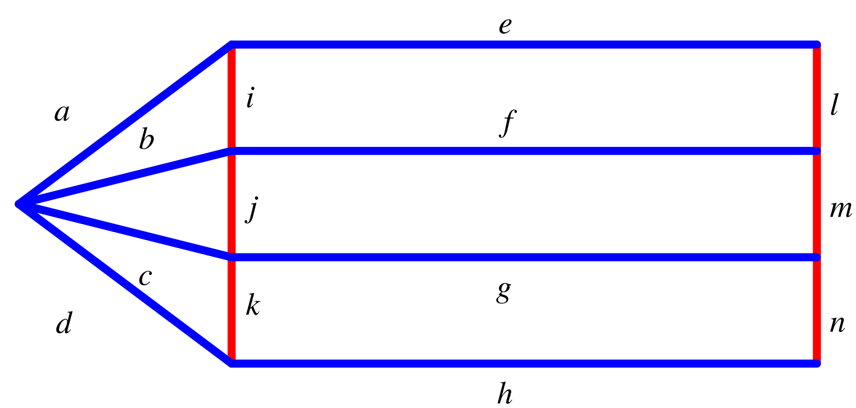

Figure 3.11. Three adjacent triangle-square pairs stacked on top of each other, with labelings for the edges.

from a red edge intersecting blue bigons. As this can't happen, we can't have just one red edge intersect our bigons. So, the next case to consider is two or more red edges intersecting the bigons.

3.5 Triangles and Squares

To review where we are, we know that any non-trivial blue bigon families we have must be intersected by at least two (parallel) red edges. We can then break these bigons up into blue-blue-red triangles and red-blue-red-blue rectangles. The triangles come from the ends of the bigons, and must meet the vertex of the bigon, while the squares are from the interior, and have a red edge on either end. In this section, we will focus on adjacent triangles and squares, such as the three in Figure 3.11. Our goal now is to show that we can only have so many of these triangle-square pairs can be stacked on each other:

Lemma 3.5.1. *There cannot be five adjacent triangle-square pairs for the diagram* $\pi(L_{B,2})$.

Once we have proved this, we can combine it with Lemma 3.4.4, which tells us we must have at least five adjacent non-trivial blue bigons, to get a contradiction and show that $S_{B,2}$ is essential. To do this, we will first prove that these triangle-square pairs have to sit inside a disk on F, and hence can't have any topology to them. Next, we prove that, in a few specific cases, if two blue sides of a crossing meet the same crossing circle, then that crossing must be associated to that crossing circle. With these, we can then use the proof of [27, Section 5], which we will review below, with little modifications to prove Lemma 3.5.1.

Before we begin, we introduce another definition. We say that a crossing is *associated to a crossing circle* if the crossing is in the twist region the crossing circle encircles. There are two primary ways we can show that a crossing is associated to a crossing circle. First, we can show that a crossing x is part of the same twist region as crossing y, where we know that y is associated to a crossing circle. We do this in the typical way of showing two crossings belong to the same twist region, by finding a disk with boundary only intersecting the diagram at x and y. The second way we can show a crossing is associated to a crossing circle is by finding a disk with boundary intersecting the diagram exactly four times: twice transverse to the crossing in question, and twice where the crossing circle meets the diagram. That is, we find a disk whose boundary intersects the crossing and follows the crossing circle when it meets the surface F.

As before, where, when working with triangles, we saw that graphs induced disks on the projection surface, we will want to do something similar when working with triangle-square pairs, that is, when we glue one edge of a triangle to an edge of a square.

Lemma 3.5.2. *A blue-red rectangle with interior disjoint from the vertices of* Γ_B *induces a disk on the projection surface* F*. Furthermore, a triangle-square pair, with interior disjoint from the vertices of* Γ_B *induces a disk on the projection surface* F*.*

Proof. We will prove this by starting with the simplest case, a blue-red rectangle with interior disjoint from blue and red edges, and expand from there. So assume our rectangle has interior disjoint from blue and red edges. Then, as we have done previously, we can form a simple closed curve γ by taking the boundary of the rectangle, as below.

As red edges can't intersect red edges, neither of our two red edges can cross any additional crossings. Likewise, as blue edges can only meet blue edges at vertices, neither of our blue edges can meet additional crossings. So then γ intersects our knot exactly four times. As $r(\pi(K), F)$ is strictly greater than four, γ cannot bound a compression disk of F. On the other hand, if γ is an essential curve, the disk it bounds must pass through the projection surface F.

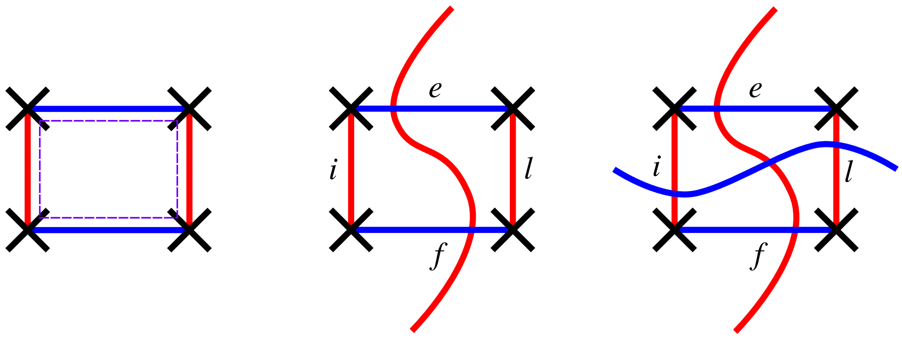

Figure 3.12. *(Left)* A blue-red rectangle disjoint from all other edges. The purple curve represents our γ. *(Center)* A blue-red rectangle that might intersect a red edge, but is disjoint from blue edges. *(Right)* A blue-red rectangle that might intersect any number of blue or red edges.

But such a disk can only pass through F at a red or blue edge, which our rectangle is disjoint from. So then γ must bound a disk on F.

Next, suppose our rectangle has interior disjoint from blue edges, but not necessarily from red edges. We first want to show that our red edges must run entirely through the rectangle, without any crossings in the interior. First, if any of the red edges meets a crossing, it must then either meet a blue edge at that crossing or meet no other edges at that crossing (as red edges can't meet red edges, and can only meet green at crossing circles). If it meets a blue edge, as our interior is disjoint from blue edges, that means the crossing must be on the boundary blue edge, and we are good. If it meets no other edges, then our crossing is in the interior. But then, we can go from one side of the red edge to the other by going around the crossing, with this path never intersecting our graph. But then, as the projection of the disk switches from above to below the projection plane at blue and red edges, we have a single region of our disk that must be both above and below this plane, which cannot happen. So any crossings our red edge meets must be on the boundary, and we may proceed.

Now we will proceed by induction. We already know how to proceed if we have no red edges. Now suppose we can get a disk if our rectangle has at most $n-1$ red edges in the interior. Label the blue boundary edges e, f and the red boundary edges i, l. Given a rectangle with n red edges in the interior, pick one of them and label it α. As red edges cannot intersect red edges, this red edge must be parallel to all of the other red edges, including those on the boundary. Then we can come up with two smaller blue-red rectangles, each with less than n red edges in the interior, by splitting along α. The boundaries of these rectangles will be $\gamma_1 = i \cup \frac{1}{2}e \cup \alpha \cup \frac{1}{2}f$ and $\gamma_2 = \alpha \cup \frac{1}{2}e \cup l \cup \frac{1}{2}f$. Each of these curves, by induction, will give us disks on F. To get a disk on F bounded by $i \cup e \cup l \cup f$, we can glue our two smaller disks along their common edge α, and we are done with this case.

Finally, we deal with the broad case of a rectangle that is only disjoint from the vertices of Γ_B. As Γ_B is a connected graph, and blue edges only meet blue edges at vertices/crossing circles, our blue edges must run all the way from one boundary (red) edge to the other. As it turns out, this doesn't actually matter for our proof - we just need the blue edges to run from one red edge to the next. But if the blue edges didn't run all the way from one end to the other, we could find a region where our disk was both above and below the projection plane, a contradiction.

Once again, we proceed by induction. If our rectangle has no blue edges in the interior, we can use what we've shown above to find a disk. Now suppose we can get a disk if our rectangle has at most $n-1$ blue edges, and suppose our rectangle has n blue interior edges. Pick a blue edge and call it β. As blue edges can only meet blue edges at vertices, β must run parallel to both the boundary edges e and f, as well as the other interior blue edges. So we can construct two smaller blue-red rectangles with boundary $\gamma_1 = e \cup \frac{1}{2}l \cup \beta \cup \frac{1}{2}i$ and $\gamma_2 = \beta \cup \frac{1}{2}l \cup f \cup \frac{1}{2}i$, with each rectangle having less than n blue edges. By induction, we can then find disks on F with boundaries γ_1 and γ_2. To get a disk with boundary $e \cup i \cup f \cup l$, we can glue the disks for γ_1 and γ_2 along β, and we are done.

Finally, we will show that a triangle-square pair with no interior Γ_B vertices induces a disk on F. We can split the triangle-square pair into a blue-blue-red triangle and a blue-red rectangle, whose boundaries share a red edge. The triangle bounds a disk by 3.4.14, while the square bounds a disk by above. So we can form a disk for the triangle-square pair by gluing these two disks together along their common red side. $\square$

Now we will examine what three adjacent triangle-square pairs might look like. As we now know that the triangle-square pairs must all lie on a disc, and thus there is no genus involved, we don't have to worry about possibilities beyond the four drawn in Figure 17 of [27]. Likewise, many of the lemmas proven for $F = S^2$ will also hold for a general projection surface. We only have to be careful when we are referencing [27, Lemma 3.2(4)], as this part of the lemma relied on the whole projection surface as opposed to just a small (local) portion of it. So now we will prove [27, Lemma 3.2(4)] for the two cases in a triangle-square pair.

Before we begin, there is something that must be noted. When the lemmas analogue to the ones we prove are applied in [27], they are to families of adjacent bigons. As such, the interior of these bigons will not have any vertices—if not, then any interior vertex must have a blue edge extending to either vertex of the bigon, breaking the fact that we have adjacent bigons. So the condition that the triangle-square pairs are disjoint from vertices holds for the ones we care about, allowing us to apply the following two lemmas in place of [27, Lemma 3.2(4)]:

Lemma 3.5.3. *Suppose that, in a triangle-square pair with interior disjoint from vertices, two blue edges, x and y, meet such that the other endpoint of x is the vertex of the triangle-square pair, and y is in the same region as the crossing circle associated to the vertex. Then the crossing corresponding to this endpoint must be associated to the crossing circle.*

Proof. First, note that, by Lemma 3.5.2, our triangle-square pair bounds a disk on F. Let z be a blue edge adjacent to x that also has an endpoint on the vertex, and w the red edge from x to z. Then, as x and z must be in distinct blue regions, and x and y must be in distinct blue region, we must have that y and z are in the same blue region.

There is a simple closed curve connecting the crossing at the endpoint of x and y and the crossing circle. Let γ be the curve that follows x from the crossing to the crossing circle, goes through the crossing circle to z, and then jumps from z to y to make it back to the crossing. We claim that γ bounds a disk. If γ is contained entirely within the disk corresponding to the triangle-square pair, then we are done. Otherwise, at some point, we would have to leave this disk. By construction, the only time we could have done this is the jump from z to y.

That would mean that, although z and y are in the same blue region, we have to leave the triangle-square pair to realize this.

While this could be a problem, we are going to show that we can add to our triangle-square pair disk, and get a disk that includes all of γ. Let α be the simple closed curve that follows the jump from z to y, then the red edge w back to z, and then follows z back to where α started.

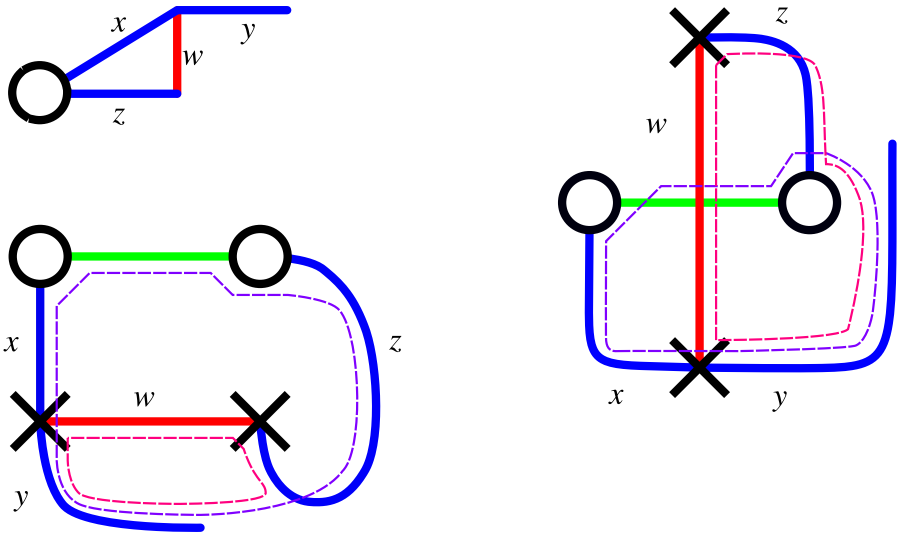

Figure 3.13. *(Upper Left)* The graph of the case we are considering. *(Lower Left and Right)* The possible diagram configurations of this case. The purple curve following the crossing disk is γ, while the pink curve following w is α.

Then α must intersect our knot in exactly two places, and so bounds a disk in F. Further, α is either entirely contained in the disk of the triangle square pair (as in the lower left picture of Figure 3.13), or can be split into two arcs, one of which is entirely contained within the disk of the triangle square pair (as in the right picture of Figure 3.13). This will allow us to glue the α disk to the triangle-square pair disk, and still get a disk. This, then, will give us γ living inside a disk, and so then it must bound a disk. Then we have a disk with boundary intersecting the crossing and with part of the boundary running along where the crossing disk meets the surface F. This means that the crossing must be part of the twist region the crossing circle encloses, and so the crossing is associated to the twist region, and we are done. □

Lemma 3.5.4. *Suppose that, in a group of triangle-square pairs with interior disjoint from vertices, two blue edges, x and y, both have an endpoint at the vertex, and in the image of the* *knot, have their other endpoint on opposite sides of the same crossing. Then that crossing is associated to the crossing circle of the vertex.*

Proof. As with above, we can use Lemma 3.5.2 to get a disk that the triangle-square pair lives on. In particular, both x and y live on the disk, as does the green edge associated with the crossing circle. Also, as x and y are on opposite sides of a crossing, they must also be on opposite sides of the crossing circle associated to the vertex. So we will create a simple closed curve that intersects our knot twice at the crossing circle and twice at the crossing of x and y by taking γ to be the curve that follows x from the crossing to the crossing circle, then the green edge across, then y back down to the crossing. As x and y live in distinct blue regions, they cannot intersect except at the crossing, so γ is simple. And, as γ lives in a disk, it must also then bound a disk. Then we have γ intersecting our diagram twice at the crossing at twice at the crossing circle and bounding a disk, and so the crossing must be associated to the crossing circle, and we are done. □

We may now prove Lemma 3.5.1. The proof does not change much from it's version of [27], but we outline the method here.

Proof. The proofs of Lemmas 5.1 through 5.5 of [27] all work as intended, where we use the above two lemmas in place of [27, Lemma 3.2 (4)]. To summarize, these lemmas tell us how three adjacent triangle-square pairs must be placed in our diagram. Of particular note, the last three of these, Lemmas 5.3 through 5.5, state that, of the four red-blue crossings of the triangles, three of them must correspond to the vertex of the crossing circle, and they cannot all be adjacent to each other. That is, either the top two crossings and the bottom crossing are associated to the vertex, or the bottom two and the top crossing are.

Now, by starting with three adjacent triangle-square pairs, we can start adding additional triangle-square pairs, and keep track of if the crossings are associated to the vertex of the crossing circle or not. Suppose we have three adjacent triangle square pairs, and the top two red-blue crossings of the triangles are associated ot the vertex, as is the bottom crossing. To add a fourth triangle-square pair, we must add it to the bottom, and the crossing we add must also be associated to the crossing. If we try to add it to the top, then the three top triangle-square pairs would have either three adjacent crossings all associated to the vertex, or only two crossings associated to the vertex, both contradicting Lemmas 5.3 through 5.5 of [27]. Next, when we try to add a fifth triangle-square pair, we must add it to the bottom agagin. However, no matter if the added crossing is associated to the vertex or not, the botttom three pairs will either have only two associted crossings or three adjacent associated crossings, both contradictions. So then we can't have five adjacent triangle-square pairs, and so we are done. □

We may use Lemma 3.5.1 to get the following result.

Proposition 3.5.5. *The graph Γ_B cannot contain five or more adjacent non-trivial bigons.*

Proof. Suppose Γ_B contains more than five adjacent non-trivial bigons. Then, by Lemma 3.4.9, the bigons cannot be disjoint from green and red. If the bigons are disjoint red, then they must intersect a green edge. However, as we cannot have blue-green bigons or blue-blue-green triangles ([27, Lemmas 4.4, 4.16]), so the green edge must go from one vertex to another. If the green edge meets only one vertex, we have a green monogon, which cannot happen ([27, Lemma 4.16] says we can remove such monogons), while if it meets two distinct vertices of our bigon, then the blue edges must be trivial by definition, a contradiction. So the blue bigons must meet the red surface.

As there are no blue-red bigons (Lemma 3.4.7), any red edge must run straight through all five bigons, giving us a collection of blue-blue-red triangles. By Lemma 3.4.16, we can't have three or more adjacent such triangles, so another red edge must run parallel, giving us adjacent triangles and squares instead. However, by Lemma 3.5.1, we cannot have five adjacent triangle-square pairs, so we get a contradiction, and thus Γ_B cannot have five adjacent non-trivial bigons. □

3.6 Completing the Proof of the Theorems

We are now ready to prove Theorem 4.1.1. Once we do, we will then introduce some additional lemmas to prove Theorem 4.2.4

We begin by showing that $f : S_{B,2} \to Y \setminus K$ is π_1-injective:

Theorem 3.6.1. *Let $f : S_{B,2} \to Y \setminus K$ be the immersion of $S_{B,2}$ into $Y \setminus K$. Then this immersion is π_1-injective, provided $N_{tw} \geq 121$.*

Proof. If not, then we can get a disk $\phi : D \to Y \setminus K$, such that $\phi \mid_{\partial D} = f \circ \ell$, for some essential loop ℓ, and graph $\Gamma_B = \phi^{-1}(f(S_{B,2}))$, with vertices corresponding to crossing circles. These vertices have valence $2n_j$ if in the interior, or $n_j + 1$ in the boundary, where $2n_j$ is the number of crossings removed from the twist region when constructing $S_{B,2}$. In particular, $2n_j \geq 120$. By Lemma 3.4.4, as the minimum number of crossings removed is $R_{tw} \geq 2\lceil N_{tw}/2 \rceil - 2 \geq 120$, we must then have more than five adjacent non-trivial bigons. This, however, contradicts Lemma 3.5.5 above, and so f must be π_1-injective. □

Likewise, the proof of boundary-π_1-injectivity can be proved as in [27]. As before, instead of reproving every lemma we need, we will instead prove [27, Lemma 3.2(4)] in the specific case we need, and replace the uses of Lemma 3.2(4) with this new one:

Lemma 3.6.2. *Suppose a triangle with two blue edges and one edge on $\partial N(K)$ does not meet the red edge. Then the two blue edges must be on opposite sides of a crossing, and that crossing must be associated to the crossing circle of the vertex.*

Proof. As the two blue edges of the triangle meet at the same vertex, they must map opposite sides of the crossing circle. However, as the third side is the knot strand, and the triangle does not meet red, the strand must run through a single crossing, proving the first part. To see that the crossing is associated to the crossing circle, we can construct a closed curve intersecting the knot exactly four times: Take one of the blue edges from the crossing circle down to the crossing, then follow the other blue edge up to the crossing circle, and pass over to the start through the green edge. We can view the interior of this disk as being the interior of our triangle. We know immediately that this curve cannot bound a compressing disk—$r(\pi(K_{B,2}), F) > 4$. So the other option is that it is essential. However, if this curve is essential and does not bound a compression disk, the disk, and thus the interior of the triangle, must pass through the projection surface. It can only do this where red and blue meet. As the triangle is disjoint from red, this cannot happen, so the curve must bound a disk on the projection surface. But then we have a disk whose boundary passes only through the crossing in question and the crossing circle, so the crossing is related to the crossing circle, and we are done. □

From here, the remaining lemmas follow nicely. To summarize, first, by looking at how blue-blue-$\partial N(K)$ triangles can be placed in the diagram, we get that, if we have three adjacent triangles in Γ_B, they must meet the red surface. This is the lemma that references [27, Lemma 3.2(4)], and so we can use the above lemma in place of it. Further, by careful examination, [27, Lemma 6.2] tells us that three adjacent such triangles must meet the red surface at least twice. This, then, gives us blue-red triangle square pairs, and we can use Lemma 3.5.1 to say that we can't have five adjacent triangle-square pairs, and thus, can't have five adjacent blue-blue-$\partial N(K)$ triangles.

Now, we can prove the ∂-π_1-injective:

Theorem 3.6.3. *Let $f : S_{B,2} \to Y \setminus \mathrm{int}(N(K))$ be the immersion of $S_{B,2}$ into $Y \setminus \mathrm{int}(N(K))$. Then this immersion is ∂-π_1-injective, provided $N_{tw} \geq 121$.*

Proof. If not, then we can get a disk $\phi : D \to Y \setminus \mathrm{int}(N(K))$, such that ∂D is a concatenation of two arcs, one mapped into $\partial N(K)$, and the other an essential arc in $S_{B,2}$ [27, Lemma 2.2]. By Lemma 3.4.6, as the minimum number of crossings removed is $R_{tw} \geq 2\lceil N_{tw}/2\rceil - 2 \geq 120$, we must have more than five adjacent non-trivial bigons or five adjacent triangles with an edge on $\phi^{-1}(\partial N(K))$. However, as mentioned above, by [27, Lemma 6.3], we cannot have five adjacent such triangles. Also, by Lemma 3.5.5, we cannot have five adjacent non-trivial bigons. This gives us a contradiction, so f must be ∂-π_1-injective. □

Putting Theorems 3.6.1 and 3.6.3 together, we get:

Theorem 3.1.1. *Let $f : S_{B,2} \to Y \setminus K$ be the immersion of our blue twisted surface $S_{B,2}$ into our knot complement $Y \setminus K$, where Y has no boundary components. Then this immersion is π_1-injective and boundary π_1-injective provided that $N_{tw} \geq 121$.*

Proof. We first prove that $f : S_{B,2} \to Y \setminus K$ is π_1-injective. If not, then, by [27, Lemma 2.1] (and the remarks at the beginning of section 5 about why this carries over to our general case), we get a map of a disk $\phi : D \to Y \setminus K$ with $\phi|_{\partial D} = f \circ \ell$ for some essential loop ℓ in $S_{B,2}$, with $\Gamma_B = \phi^{-1}(f(S_{B,2}))$ a collection of embedded closed curves and an embedded graph in D. Each vertex in the interior of D has valence a non-zero multiple of $2n_j > 121$, where $2n_j$ is the number of crossings removed, and each vertex in the exterior has valence $n_j + 1 > 60$. Then, as we are working with a graph with exactly the same properties as the graph in Section 2 of [27], we can get that Γ_B has more than $(N_{tw}/18) - 1 > 5$ adjacent non-trivial bigons. □

3.6.1 Properties of Twisted Surfaces

Now that we have proven that our twisted surfaces are essential surfaces in $Y \setminus K$, we can get a bit more out of them as well. Ultimately, we want to show the following:

Theorem 4.2.4. *Let S be the disjoint union of our twisted surfaces $S_{B,2} \sqcup S_{R,2}$. Suppose that two distinct essential arcs in the surface $S_{B,2}$ have homotopic images in $Y \setminus K$, but not $S_{B,2}$. Then the two arcs are homotopic in $S_{B,2}$ into the same subsurface associated with some twist region of K_2.*

This theorem is proven in a series of lemmas, all but one of which require little to no alteration in statement and proof. We will go through each of the lemmas, and give a sketch of why they still hold for our case. Recall that *the subsurface associated to a twist region of K_2* is the intersection of a surface with a regular neighborhood of the twist region.

The first step in our proof of Theorem 4.2.4 is to get a disk to work on.

Lemma 3.6.4. *Suppose homotopically distinct essential arcs a_1 and a_2 in $S_{B,2}$ map by $f : S_{B,2} \to Y \setminus K$ to homotopic arcs e_1 and e_2 in $Y \setminus K$. Then there is a map of a disk $\phi : D \to Y \setminus \mathrm{int}(N(K))$ with ∂D expressed as four arcs, with opposite arcs mapping by ϕ to e_1 and e_2, and the other two arcs mapping to $\partial N(K)$. Moreover, $\Gamma_B = \phi^{-1}(f(S_{B,2}))$ forms a graph on D whose edges have endpoints either at vertices where $\phi(D)$ meets a crossing circle, or on $\phi^{-1}(\partial N(K))$ on ∂D. Each vertex has valence either a multiple of $2n$ (if in the interior of D), $n + 1$ (if in the interior of an arc on ∂D that maps to $S_{B,2}$), or 1 (if on an arc that maps to $\partial N(K)$), where $2n$ is the number of crossings removed from the twist region.*

Proof. As e_1 and e_2 are homotopic, we immediately get a disk $\phi : D \to Y \setminus \mathrm{int}(N(K))$ with the correct edges. We then modify this disk to get the remaining properties.

First, we lift e_1 and e_2 to the orientable double cover of $S_{B,2}$ By pushing e_1 and e_2 in the transverse direction, and using the fact that $\partial N(K)$ is transversely orientable, we can get the map of ϕ on a neighborhood of ∂D. Then we can extend ϕ over the rest of D, making it transverse to all crossing circles and to $f(S_{B,2})$.

Now look at $\Gamma_B = \phi^{-1}(f(S_{B,2}))$. As $S_{B,2}$ is embedded in $Y \setminus K$ except as crossing circles, Γ_B must consist of embedded closed curves (where $S_{B,2}$ intersects $\phi(D)$ away from crossing circles or $\partial N(K)$), arcs with endpoints on vertices (where $S_{B,2}$ intersects a crossing circle), and arcs with endpoints on $\partial N(K)$.

As a vertex of Γ_B in the interior of D comes from where $\phi(D)$ intersects a crossing circle, the valence is equal to the number of crossings removed, by construction of the twisted surface $S_{B,2}$. If we have a vertex on an arc in ∂D, it maps either to $S_{B,2}$ or $\partial N(K)$. If it maps to $S_{B,2}$, then a neighborhood of the arc must map to half a meridian disk of a crossing circle, and so has valence $n + 1$, where we removed $2n$ crossings from the corresponding twist region in our construction of $S_{B,2}$. If it maps to $\partial N(K)$, however, because the arc must be transverse to $S_{B,2}$, the vertex must have valence 1. □

We can also take our disk to be minimal, in the sense that it is a disk that gives us the fewest number of vertices of Γ_B (the graph coming from how the disk intersects our blue surface) and Γ_{BR} (the graph coming from how the disk intersects both our blue and red surfaces). Next, by examining graphs in a disk that satisfy the properties above, we get an analogue to Lemma 3.4.6:

Lemma 3.6.5. *Suppose Γ_B contains at least one blue vertex. Then either Γ_B has more than $(R_{tw}/24) - 1$ adjacent non-trivial bigons, or more than $(R_{tw}/24) - 1$ adjacent triangles, each with one arc on $\phi^{-1}(\partial N(K))$.*

The proof of this lemma is a graph theoretic proof, and so does not rely on the ambient space Y or the twisted surfaces $S_{B,2}$ or $S_{R,2}$.

Proof (Sketch). First, suppose we have a graph Γ on $I \times I$ with no monogons and only valence one vertices on $I \times \partial I$, with interior vertices with valence at least R and vertices on $\partial I \times (I - \partial I)$ have valence at least $R/2 + 1$. Then Γ must have more than $R/8 - 1$ adjacent bigons or more than $R/8 - 1$ adjacent triangles with an edge on $I \times \partial I$.

Next, we can modify our graph Γ_B coming from Lemma 3.6.4 by looking at subdisks, doubling along edges, and collapsing certain bigons and triangles so that we satisfy the above statement, with $R = R_{tw}$ is the minimal number of crossings removed from a twist region. When we collapsed bigons or triangles, we made it so at most three bigons or triangles in Γ_B are associated to a bigon or triangle in this new graph. As such, we get $(R_{tw}/24) - 1$ adjacent non-trivial bigons in Γ_B, or more than $(R_{tw}/24) - 1$ adjacent triangles, and are done. □

Using this, we can prove:

Lemma 3.6.6 (Lemma 7.6 [27])**.** *If $N_{tw} \geq 121$, then the graph Γ_B contains no blue vertices. That is, $\phi(D)$ meets no crossing circles.*

Proof. As shown in Lemma 3.4.8, Γ_B has no monogons. If Γ_B has any trivial arcs in ∂D, then that arc must bound a disk E that is a subset of $S_{B,2}$. By sliding ∂D along E, we can remove such an arc, so we may assume that Γ_B has no trivial arcs in ∂D. Then, if there is a blue vertex in Γ_B, we can use Lemma 3.6.5, and get more than $((N_{tw} - 1)/24) - 1$ adjacent non-trivial bigons or adjacent triangles with an arc on $\phi^{-1}(\partial N(K))$. When $N_{tw} \geq 121$, we get at least 5 such bigons or triangles. By Lemma 3.5.5, we cannot have five such bigons, and, as discussed before, by [27, Lemma 6.3], we cannot have five adjacent triangles, and so Γ_B cannot have any blue vertices. □

So now we can start looking at how exactly the surfaces $S_{B,2}$ and R_2 (the red surface arising from K_2) intersects our disk. First, with no blue vertices, Γ_B consists only of blue arcs with end points on the north and south edges of the disk, corresponding to $\partial N(K)$. If both of the edges are on the same edge, by Theorem 3.6.3, because the disk is $\partial\pi_1$-injective, the blue edge is trivial and we can find a disk that doesn't meet that edge. So blue edges must run from north to south. Likewise, the red edges of Γ_{BR} are only arcs that go from one distinct edge to another, from Lemma 3.4.7 (no blue red bigons), and the fact that the red checkerboard surface R_2 is essential.

Now we need to show that the red arcs run along opposite edges, either east to west (e_1 to e_2), or north to south ($\partial N(K)$ to $\partial N(K)$):

Lemma 3.6.7. *The graph Γ_{BR} consists of red and blue arcs with endpoints on opposite sides of D (north-south or east-west). Further, blue arcs run north to south.*

Proof. In the case of blue edges, we are done, as we know the arcs run from north to south as discussed above. For red arcs, if they don't have endpoints on opposite sides, then they must form a triangle with one endpoint on either the north or south side, and the other on the east or west side. This gives us a blue-red-$\partial N(K)$ triangle. The blue and red arcs must meet at a crossing of our arc, so, assuming the triangle region maps above the projection plane, there are two ways this could happen. Either the red edge runs from one crossing (where it meets the blue arc) to near another crossing, or runs from one crossing to $\partial N(K)$, with the blue arc on the opposite side of this meeting.

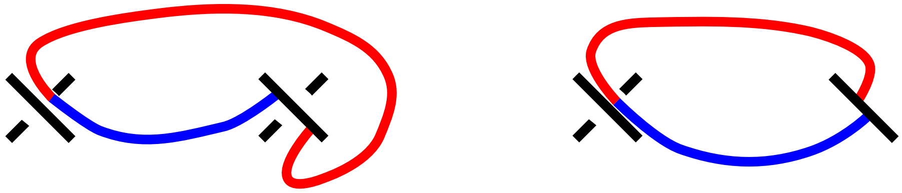

Figure 3.14. Two ways a triangle with one red on red, one edge on blue, and one edge on the knot can sit in the diagram.

Before we can use primality arguments to show neither of the cases can happen, we need to show that we can get a disk from this triangle. By joining the blue and red arcs together at the endpoints that don't meet, either by sliding along K or connecting across a crossing, we get a closed curve that intersects our knot diagram twice. As $e(\pi(K), F) \geq 4$, this curve cannot be essential, and so must bound a disk on F.

From here, the remainder of the proof follows as in [27]. If we are in the left case of Figure 3.14, connecting the endpoints will give us a blue-red bigon. But then, as in the proof of Lemma 3.4.7, we can get a contradiction to primality. On the other hand, if we are in the other case, because $K_{B,2}$ is weakly prime, there must be no crossings in the interior of the closed curve we constructed. As such, both the red and blue arcs must be homotopic to a portion of $\partial N(K)$. We can then use a homotopy to slide our disk D away from where the blue and red arc intersect. This will give us one less vertex on D, and so contradicts the minimality of D. In either situation, we get a contradiction, so the red arcs must run either north-south or east-west. □

As red arcs cannot intersect red arcs, we must also have that all red arcs run in the same direction—either all north-south or all east west. We consider the first case now:

Lemma 3.6.8. *If the graph Γ_{BR} cuts D into a subrectangle with two opposite sides mapped to $\partial N(K)$, one side on blue, one side on red, and the interior disjoint from blue and red, then the blue and red sides of that rectangle are homotopic to the same crossing arc, and the homotopies can be taken to lie entirely in $S_{B,2}$ and R_2.*

Proof. First, note that such a rectangle, with interior disjoint from both blue and red, must be mapped either completely above or below the projection surface. Suppose it is above. Then there are three cases to consider. First, the blue and red arcs could straddle over crossings at both ends. Second, one set of endpoints could straddle an over crossing, while the other set lies on opposite sides of a strand. Third, both sets of endpoints lie on opposite ends of strands.

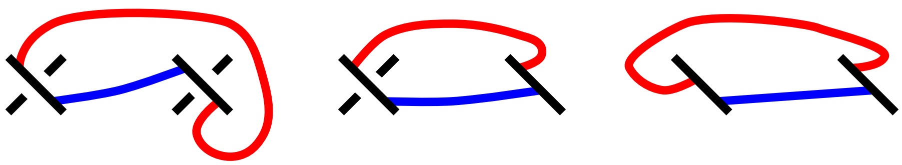

Figure 3.15. The three ways a red-blue rectangle can sit in our diagram

Note that in all three of these cases, by connecting the red and blue arcs, we get a curve that intersects our knot exactly twice and so, as before, must bound a disk on F. The first case, with two over crossings, as in Lemma 3.4.7 (there are no blue-red bigons), violates the fact that our diagram is Weakly Prime. The third case, with two strands, can be homotoped away from as in Lemma 3.6.7. That leaves us with just our second case. In this case, both the blue and red arc are homotopic to the same crossing arc (the over crossing they straddle), and the homotopies lie entirely in their corresponding surfaces, either $S_{B,2}$ or R_2, so we are done. □

From here, with no change to proof except to use $e(\pi(K), F) \geq 4$ to show the curves in question bound a disk, by carefully examining how rectangles with two colored sides (either both blue, both red, or blue and red) and two $\partial N(K)$ sides can give us homotopies, we get the following result:

Lemma 3.6.9. *If the graph Γ_{BR} consists of disjoint red and blue arcs on D, all running north to south, then e_1 and e_2 are each homotopic in the blue surface to arcs in the same subsurface associated with a twist region of K_2.*

The e_1 and e_2 here are the two homotopic arcs in $Y \setminus K$ from Lemma 3.6.4.

Proof (Sketch). Our arcs split our disk into several sub-rectangles, with east and west sides colored either red or blue, and the north and south side on $\partial N(K)$. If both colored sides are blue, the subrectangle must be mapped to a single side of the projection surface, and the north and south sides must run over crossings (or else we could homotope D away to remove intersections). This means our subrectangle gives us a simple closed curve that meets our diagram at two crossings with the subrectangle on only one side of the projection surface, and thus gives us a disk. Because $K_{B,2}$ is weakly blue-twist reduced, this disk bounds a collection of red bigons, and so we can isotope the blue arcs to lie in the same subsurface. In a similar fashion, having to be a bit more careful as we don't know $K_{B,2}$ is weakly red-twist reduced, we get the same result for when both sides are red.

If the two sides are colored differently (one red, one blue), then Lemma 3.6.8 allows us to homotope the two sides into the same subsurface. Putting it all together, two successive arcs are homotopic into the same subsurface, in either $S_{B,2}$ or R_2, as appropriate. Expanding this out, we get that e_1 and e_2 must be homotopic in the blue surface into the same subsurface. □

For the last lemma, we need an analogue to the polyhedron decomposition of alternating knots. We will use the chunk decomposition of Howie and Purcell [23]:

Definition 3.6.10. A *chunk* C is a compact, oriented, irreducible 3-manifold with boundary ∂C containing an embedded non-empty graph with all vertices having valence at least 3. The graph separates ∂C into regions called *faces*. If the region is disjoint from the graph, it is called an *exterior face*, otherwise, it is an *interior face*.

A *truncated chunk* is a chunk where a regular neighborhood of each vertex of the edge graph has been removed. This leaves a *boundary face* surrounded by *boundary edges*.

A *chunk decomposition* for a 3-manifold M is a decomposition of M into chunks, such that M is obtained by gluing chunks by homeomorphisms of non-exterior, non-boundary faces with edges mapping to edges homomorphically.

Howie and Purcell showed in Proposition 3.1 in the same paper that weakly generalized alternating knots admit a chunk decomposition:

Theorem 3.6.11. *[Proposition 3.1 [23]] Let Y be a compact, orientable, irreducible 3-manifold with no boundary components, containing a generalized projection surface F. Let K be a weakly generalized alternating knot with a diagram $\pi(K)$ on F. Then $Y \setminus K$ can be decomposed into pieces such that:*

1. *Pieces are homeomorphic to components of $Y \setminus N(F)$, except each piece has a finite set of points removed from $\partial(Y \setminus N(F))$, namely the ideal vertices below.*

2. *On each copy of F, there is an embedded graph with vertices, edges, and regions identified with the diagram graph $\pi(K)$. All vertices are ideal and 4-valent.*

3. *To obtain $Y \setminus K$, glue pieces as follows. Each region of $F \setminus \pi(L)$ is glued to the corresponding region on the opposite copy of F by a homeomorphism that is the identity composed with a rotation along the boundary. The rotation takes an edge of the boundary to the nearest edge in the direction of that boundary component's orientation.*

4. *Edges correspond to crossing arcs, and are glued in fours. At each ideal vertex, two opposite edges are glued together.*

With this in mind, we can generalize the following lemma, initially proven by Lackenby [25]. Before we begin, note that an *essential square* is a simple closed curve, intersecting $\pi(K)$ exactly four times, thus dividing the curve into four sides, each of which intersects a different edge of the diagram. As we will be working in Γ_B, and essential squares are blue-red rectangles, we can use Lemma 3.5.2 to assume that these squares must bound a disk in F.

Lemma 3.6.12. *Let C be a chunk in a chunk decomposition for $Y \setminus K$, where K is a weakly generalized alternating knot on F with $e(\pi(K), F) \geq 4$. Then, if S and T are essential squares, each bounding a disk on F and isotoped to minimize $|S \cap T|$, they intersect either zero or two times. If they intersect two times, then the points of intersection lie in distinct regions of the diagram with the same color.*

Proof. We can color the faces of our chunk blue or red, depending on how it meets the checkerboard coloring of $F \setminus K$. An essential square intersects the knot projection at four distinct points, and so the square must have two blue edges and two red edges, each opposite each other. As S and T have been isotoped to have minimal intersection, there are at most four points of intersection, one at each edge. If there are four points of intersection, then S and T are isotopic to each other, and can be pushed off to have no intersection. In addition, they cannot intersect an odd number of times, as both bound disks on F. If they do, look at how T as a curve intersects S. It must cross the boundary of S exactly three times. As T bounds a disk on F, however, this is impossible. So S and T intersect either zero or two times.

Suppose they intersect exactly twice, and one of those intersection points happens on red, and the other on blue. Because S and T must intersect the knot where blue meets red, if we look at the arcs of $S \setminus T$ and $T \setminus S$, they must intersect our knot either once or three times. Look at the two arcs that intersect the knot once, each, and glue them together to get a curve intersecting our knot exactly twice. As $e(\pi(K), F) \geq 4$, this curve must bound a disk on F. Then, as K is weakly prime, because the curve intersects the knot exactly twice and bounds a disk, the disk must intersect our knot in a single arc with no crossings. We can then use this arc to homotope S and T away from each other, contradicting the fact that our intersection was minimal. So, if S and T intersect, they must intersect in distinct regions of the same color. □

As a consequence to this, we get the following result:

Lemma 3.6.13. *Let C be a chunk in a chunk decomposition for $Y \setminus K$, where K is a weakly generalized alternating knot on F with $e(\pi(K), F) \geq 4$. Let S and T be essential squares in P, each bounding a disk on F and moved by normal isotopy to minimize $|S \cap T|$. Suppose S and T pass through the same red face W, and that edges $S \cap W$ and $T \cap W$ differ by a single rotation of W. Then exactly one of the following two conclusions holds:*

1. *Each of S and T cuts off a single ideal vertex in W, and S and T are disjoint.*

2. *Neither S nor T cuts off a single vertex in W. The two essential squares intersect in W and in another face W'.*

Proof. First, if S cuts off a single ideal vertex, then so must T, and so do not intersect in W. By Lemma 3.6.12, then, if S and T intersect, it must be in a blue face. As $e(\pi(K), F) \geq 4$, W can only meet a given blue face once (otherwise we could find an essential curve on F intersecting our knot exactly twice). Then, if S and T do intersect in a blue face (and thus in two blue faces), then $S \cap W$ and $T \cap W$ must be parallel to each other. We could then isotope S and T away from each other in the blue faces so that S and T are disjoint.

If S does not cut off a single vertex, then neither does T. As $S \cap W$ and $T \cap W$ differ by a single rotation, they must intersect in W. By Lemma 3.6.12, then, they must intersect in two red faces, W and W', and so we are done. □

We can use this to get our final needed lemma. There is one last definition to introduce.

Definition 3.6.14. An *essential product disk* for the blue checkerboard surface B of a knot K is an essential disk properly embedded in $S^3 \setminus \text{int}(N(B))$, with boundary a rectangle with two oppostie sides on $N(B)$ and two opposite sides on $\partial N(K)$.

Then note that the disk of Lemma 3.6.4 is an essential product disk, so if we can prove the following lemma for essential product disks in general, it will also apply to the specific case we are concerned with.

Lemma 3.6.15. *Let e_1 and e_2 denote the boundary arcs on the blue surface in an essential product disk for the blue checkerboard surface of $K_{B,2}$. Then there is a subsurface associated with a twist region of the diagram of K_2, and arcs a_1 and a_2 in that subsurface, such that e_1 is homotopic in $S_{B,2}$ to a_1, and e_2 is homotopic in $S_{B,2}$ to a_2.*

Proof. Let E be the essential product disk. If it is disjoint from the red surface, then the blue edges of E are homotopic to arcs in a subsurface associated to one twist region of $K_{B,2}$. This follows from [27, Lemma 7.10]. To summarize, if we have a rectangle made up of two blue arcs and two edges on $\partial N(K)$, and interior disjoint from red or blue, the rectangle must be mapped to a single side of the projection surface. If the $\partial N(K)$ edges don't run over a crossing, then we could homotope the rectangle away from the projection surface, and thus Γ_{BR} was not minimal, a contradiction. So then the edges on $\partial N(K)$ must run over crossings, and thus we can get a curve that intersects the diagram at exactly two crossings. Then this curve must bound a series of red bigons, and so we may isotope the blue arcs into a neighborhood of these bigons, and are done.

If E does meet the red surface, then, by Lemma 3.6.7, either the red edges all run north-south ($\partial N(K)$ to $\partial N(K)$) or east-west (from blue edge to blue edge). In the former case, Lemma 3.6.8 gives us our result. So we may assume we are in the latter case, and red edges run from blue edge to blue edge, cutting the rectangle into blue-red-blue-red sub-rectangles.

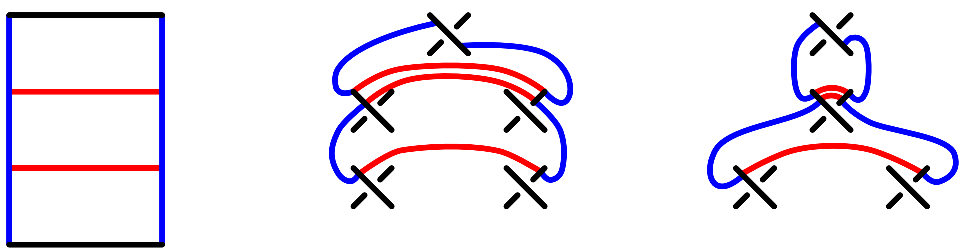

Figure 3.16. *(Left)* An essential product disk with blue sides e_1 and e_2 intersected horizontally by red edges. *(Right)* Two sub-rectangles whose shared red edge meets two separate crossings. *(Center)* Two sub-rectangles whose share red edge meets a single crossing. We want to prove that this is what happens for all adjacent sub-rectangles.

By Theorem 3.6.11, we get a chunk decomposition for our knot, with pieces homeomorphic to components of $Y \setminus N(F)$, except for ideal vertices corresponding to crossings. Each face of the chunk can be colored either red or blue, depending on what color it meets on F. As F is a connected surface, $Y \setminus N(F)$ has either one or two components. Each of our sub-rectangles that make up E must be contained entirely in a single chunk. Further, if two rectangles glue to each other along a red edge, the rectangles are either not in the same chunk, if we have two chunks, or meet two separate red faces that glue to each other, if there is only one chunk. We proceed based on whether we have only one chunk or two.

If we only have a single chunk, then look at the first two rectangles in E, E_1 the north most one that meets $\partial N(K)$, and E_2, the one below that which glues to the red edge, r of E_1. Looking at the side of E_1 that meets $\partial N(K)$, it must run through an ideal vertex of the chunk decomposition, and so we can push it off into the red face so that it cuts off a single vertex. Consider the other side of E_1 now, which glues by a clockwise or counterclockwise turn to E_2. While the red side of E_1 and the red side of E_2 are not on the same face, we can create a new rectangle $\bar{E_2}$ by copying r in E_2 into the same face of r in E_1, after rotating by the corresponding turn, and then copying the rest of E_2 over. By Lemma 3.6.12, if the two copies of r (for E_1 and $\bar{E_2}$) intersect, then E_1 and $\bar{E_2}$ must also intersect in the other red face. However, as the other edge of E_1 cuts of an ideal vertex, we can homotope it so it does not intersect the other edge of $\bar{E_2}$. This means, then, that E_1 and $\bar{E_2}$ cannot intersect in a red face. By Lemma 3.6.13, then, both r in E_1 and in $\bar{E_2}$ must cut off a single vertex, and so r in E_2 must also cut off a single vertex. Continuing by induction, each rectangle making up E must have sides in the red faces cutting off a single vertex.

Now suppose we have two chunks. Once again, we start by looking at the first two rectangles of E, E_1 and E_2, which glue by a single turn along a red edge r. The first step is to get both of these rectangles into the same chunk. Like before, we can homotope the north end of E_1 so that it cuts off a single vertex. As E_1 and E_2 are in separate chunks, we can impose E_2 into the same chunk as E_1 by copying r in E_2 into the chunk containing E_1 under a clockwise turn, and continuing for the rest of E_2. We call this new rectangle $\bar{E_2}$. Now, with both rectangles in the same chunk, we can use the same argument as in the single chunk case, and get that r of both E_1 and E_2 must cut off a single vertex, and, by induction, each red edge of E must cut off a single ideal vertex.

In either case, each rectangle of E maps to a region of the diagram meeting $\pi(K_{B,2})$ exactly four times, adjacent to two crossings (one for each red edge). As $\pi(K_{B,2})$ is weakly blue twist reduced by Lemma 3.3.7, this region must bound a string of red bigons, and so the boundaries all lie in the same twist region. By applying this to all the rectangles in E, we get that the two blue edges must lie in a neighborhood of the same twist region, and so are done. □

Now Theorem 4.2.4 follows as a consequence to several of the above lemmas.

Proof of Theorem 4.2.4. Call our two arcs e_1 and e_2. Then, by Lemma 3.6.4, we get a disk $\phi : D \to Y \setminus K$ and a graph Γ_B. Then, by Lemma 3.6.6, there are no blue vertices in our graph. Next, Lemma 3.6.7 tells us that blue edges must run north to south, while red edges either run all north to south or all east to west. If they all run north to south, then Lemma 3.6.9 tells us e_1 and e_2 are homotopic into the same subsurface. If the red edges all run east to west, then Lemma 3.6.15 gives us the same result, and we are done. $\square$

CHAPTER 4

CUSP AREA

4.1 Introduction

This chapter is based on work that will be published in Algebraic & Geometric Topology [7].

The aim of this chapter is to show that weakly generalized alternating knots, under certain conditions, have two-sided bounds for the cusp area based on the twist number of a diagram and $\chi(F)$, thus generalizing Lackenby and Purcell's results for alternating knots on $S^2 \subset S^3$ [26].

Theorem 4.1.1. *Let Y be a closed, irreducible 3-manifold and $F \subset Y$ a closed surface with $\chi(F) \leq 0$ and such that $Y \setminus N(F)$ is atoroidal. Suppose that K is a weakly generalized alternating knot with a projection $\pi(K) \subset F$ that is weakly twist reduced, and we have $e(\pi(K), F) > 2$, and $r(\pi(K), F) > 4$. Then $Y \setminus K$ is hyperbolic, and we have*

$$\frac{5.8265514 \times 10^{-7}}{2\,(120 - \chi(F))^6}\,(\mathrm{tw}_\pi(K) - \chi(F)) \leq CV(K) \leq \frac{(360\mathrm{tw}_\pi(K) - 3\chi(F))^2}{\mathrm{tw}_\pi(K)},$$

where $CV(K)$ is the cusp volume of K and $\mathrm{tw}_\pi(K)$ the twist number of $\pi(K)$.

This result, then, allows us to investigate certain properties of the knot by looking at the cusp, as we talk about next.

Given a hyperbolic knot complement $M = Y \setminus K$ there is a well defined notion of its maximal cusp C. The boundary of C, denoted by ∂C, inherits a Euclidean structure from the hyperbolic metric. Given an essential simple closed curve $\sigma \subset \partial C$ (a.k.a. a slope), the length of σ is the Euclidean length of the unique geodesic in the homotopy class of σ. We well denote this length by $\ell(\sigma)$.

Definition 4.1.2. If K is a knot in a 3-manifold Y, take a regular torus neighborhood of K. The meridian m is the curve that bounds a disk in Y that intersects K exactly once.

Using Theorem 4.1.1 and a result of Burton and Kalfagianni [11], we can obtain the following corollary:

Corollary 4.1.3. *Let Y be a closed, irreducible 3-manifold and $F \subset Y$ a closed surface with $\chi(F) < 0$ and such that $Y \setminus N(F)$ is atoroidal. Suppose that K is a weakly generalized alternating knot with a projection $\pi(K) \subset F$ that is weakly twist reduced, and we have $e(\pi(K), F) > 2$, and $r(\pi(K), F) > 4$. Let μ be the meridian and σ any non-meridian slope on maximal cusp of $Y \setminus K$. Then, we have*

$$\ell(\mu) \leq \frac{720\ (\mathrm{tw}_\pi(K) - \chi(F))}{\mathrm{tw}_\pi(K)}$$
$$\ell(\sigma) \geq \left(\frac{5.8265514 \times 10^{-7}}{720\ (120 - \chi(F))^6}\right)\ \mathrm{tw}_\pi(K),$$

where $\mathrm{tw}_\pi(K)$ is the twist number of $\pi(K)$.

In addition, we are able to use Theorem 4.1.1 with a theorem of Howie and Purcell and a theorem of Futer, Kalfagianni, and Prucell to prove the following result:

Theorem 4.1.4. *Let K be a weakly generalized alternating knot on a surface $F \neq S^2$ in a closed manifold Y with conditions as in Theorem 4.1.1. Suppose that*

$$\mathrm{tw}_\pi(K) > \left(2.32928 \times 10^{10}\right)\ (120 - \chi(F))^6\,.$$

Then, any manifold M obtained from non-meridional surgery along K is hyperbolic and

$$\frac{v_8}{4}\,(\mathrm{tw}_\pi(K) - \chi(F)) \leq \frac{1}{2}\mathrm{vol}(Y \setminus K) \leq \mathrm{vol}(M) < \mathrm{vol}(Y \setminus K).$$

Here $v_8 \approx 3.66386$ is the volume of a regular hyperbolic ideal octahedron.

We will proceed as follows. First, in Section 2, we will prove the theorems and lemmas we will need to show our main result. Then, in Section 3, we will put these all together to prove our main result, Theorem 4.1.1. Finally, in Section 4, we will go through the proofs of Corollary 4.1.3 and Theorem 4.1.4.

4.2 Proofs for Calculations

Many of the statements below will hold as long as we have an essential spanning surface for our knot, including our checkerboard surfaces. Occasionally, we will specialize to the essential twisted surfaces introduced in the previous chapter. This will allow us use numbers specific to the twisted surfaces that will give us better results.

For the lower bound we will adapt the approach of Lackenby and Purcell [26] to our setting. We note that some of the results we reference from [26], and used in there for knot complements in S^3, apply immediately to knot complements in any closed Y. We will restate them and we will review them and their proofs briefly below.

Before we start the proofs, there is an important definition that we will use throughout. Most geodesics in a hyperbolic manifold are bi-infinite, but we will want to be able to talk about their lengths. As such, we use the following definition to get around this complication:

Definition 4.2.1. Let γ be a (bi-)infinite geodesic in a hyperbolic surface R, with both ends in a horoball neighborhood of the cusps of R, H. Then $\gamma \cap H$ is a collection of closed intervals (if γ enters and exits H) along with two helf-open intervals (the ends). The *length of γ with respect to H* is the hyperbolic length of γ minus the two half-open intervals.

If α is an arc with boundary on ∂R, we first find a geodesic γ such that α is homotopic to γ in R, and define the length of α with respect to H to be the length of γ with respect to H.

The next result concerns an estimate of geodesic arcs on hyperbolic surfaces. We will apply the result specifically to our twisted surfaces, but it works in more general settings.

Theorem 4.2.2 (Theorem 2.7 [26])**.** *Let S be a (possibly disconnected) finite-area hyperbolic surface. Let H_S be an embedded horoball neighborhood of the cusps of S. Let $k = \mathrm{Area}(H_S)/\mathrm{Area}(S)$ and let $d > 0$. Then there is a collection of at least*

$$\frac{(ke^d - 1)\pi}{(e^d - 1)(\sinh(d) + 2\pi)}|\chi(S)|$$

embedded disjoint bi-infinite geodesic arcs, each with both ends in H_S, and each having length at most $2d$ with respect to H_S.

We do not need to modify this proof for our more general setting, as it does not rely on the ambient manifold, and instead only involves the surface. This is done by looking at how much "area" disjoint geodesics take up of the cusp, and then using this to find a lower bound for the area of H_S. We can then rearrange the equation to get our result. For a complete proof, refer to [26].

The next lemma that holds without change deals with estimating the cusp area of a manifold, based on essential arcs of bounded length:

Lemma 4.2.3. *Suppose that a one-cusped hyperbolic 3-manifold M contains at least p homotopically distinct essential arcs, each with length at most L measured with respect to the maximal cusp H of M. Then the cusp area* $\mathrm{Area}(\partial H)$ *is at least $p\sqrt{3}e^{-2L}$.*

Proof (Sketch). The proof is similar to that of Lemma 2.8 in [26]: Each arc in our collection has a geodesic representative in M, which then lifts to two geodesics in the universal cover $\mathbb{H}^3$. Look at the "shadows" of these geodesics on the lift of the maximal cusp ∂H (that is, a neighborhood of where the geodesic intersects ∂H). For the maximal cusp to be embedded, each of these neighborhoods must be disjoint, and, because the length is at most L, the diameter of this neighborhood must be at least e^{-L}. This means that the total area of the disks must be $p\pi e^{-2L}/2$. Finally, combining with a disk packing argument, we get the area of the cusp must be at least $p\sqrt{3}e^{-2L}$. □

While these two lemmas suggest that we should use them one after the other, we must be careful. The cusp in Theorem 4.2.2 is a two-dimensional cusp, while the cusp in Lemma 4.2.3 is three-dimensional. In addition, the geodesic arcs of Theorem 4.2.2 live on a surface, while the essential arcs of Lemma 4.2.3 are in the whole manifold. To get from one to the other, we will need the following theorem, proved in the previous chapter:

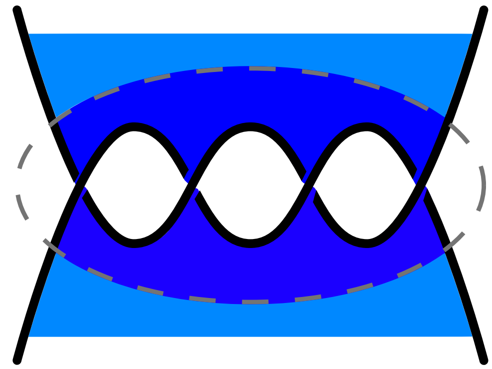

Figure 4.1. The subsurface associated to the twist region. When it comes from checkerboard surfaces, the subsurface can be colored red and blue, with one component (blue in the picture) being two disks attached with Möbius bands at each crossing, and the other being a twisted disk.

Theorem 4.2.4. *Let S be the disjoint union of our twisted surfaces $S_{B,2} \sqcup S_{R,2}$. Suppose that two distinct essential arcs in the surface $S_{B,2}$ have homotopic images in $Y \setminus K$, but not $S_{B,2}$. Then the two arcs are homotopic in $S_{B,2}$ into the same subsurface associated with some twist region of K_2.*

Here, a *subsurface associated to a twist region* is the intersection of one of the checkerboard surfaces with a regular neighborhood of the twist region. See Figure 4.1.

As a corollary to Theorem 4.2.4 we get the following:

Corollary 4.2.5. *Suppose $N \geq 121$. Let C be a collection of disjoint embedded essential non-parallel arcs in $S_{B,2} \sqcup S_{R,2}$ which are all homotopic in $Y \setminus K$. Then the number of arcs in C is at most $2(3N-2)$.*

Proof. By Theorem 4.2.4, all such arcs in C that lie in $S_{B,2}$ can be homotoped in $S_{B,2}$ to lie in a subsurface associated with a twist region of K_2. This homotopy keeps theses arcs disjoint, essential, and non-parallel. Because of this, we can use the fact that the number of disjoint, essential, and non-parallel arcs in a subsurface is at most $3c-2$, where c is the number of crossings in the twist region associated to that subsurface. Then, as we are working in K_2, $c \leq N$, so there are at most $3N-2$ such arcs in the blue twisted surface. By switching out blue for red, we get that there are also at most $3N-2$ such arcs in $S_{R,2}$. Adding these together, we get that there are at most $2(3N-2)$ such arcs total. □

From Corollary 4.2.5 we can then limit how many disjoint arcs we get from Theorem 4.2.2.

In addition to these, [26, Proposition 3.5] gives us an embedded cusp in our twisted surfaces with a lower bound on area, which we will use with Theorem 4.2.2:

Lemma 4.2.6 (Proposition 3.5 [26])**.** *If S is the disjoint union of our twisted surfaces with an immersion $f : S \to Y \setminus K$, and H is the maximal cusp for $Y \setminus K$, then there is an embedded cusp H_S for S contained in $f^{-1}(H)$ with area at least $2^{5/4}\mathrm{tw}_\pi(K)$, as long as K is neither the figure-eight knot or the 5_2 knot.*

We can find such a cusp by pleating S and giving it a hyperbolic structure, and then letting the embedded cusp be the union of the cusps of the two twisted surfaces. The lower bound on area then comes from carefully examining the boundaries of our twisted surfaces. By resolving any intersections between the boundaries, we will get a number of curves with length the same as the union of the two boundaries. The exact number is based on the intersection number of our surfaces; in our case, $S_{B,2}$ and $S_{R,2}$ intersect $2\mathrm{cr}(K_2) \geq \mathrm{tw}_\pi(K)$ times. Then, as long as K isn't the figure-eight knot or the 5_2 knot, by Adams [4], we get that each curve must have length at least $2^{5/4}$. Putting this together, we then get the desired lower bound on area for our lemma.

Finally, we will need the following lemma computing the Euler characteristic of the disjoint union of our twisted surfaces:

Lemma 4.2.7. *Suppose K is a weakly generalized alternating Knot on a projection surface F. Let $\mathrm{tw}_\mathrm{N}(K)$ denote the number of twist regions of $\pi(K)$ with at least N crossings, and let $\mathrm{cr}(K_2)$ be the number of crossings of the diagram of $\pi(K_2)$, the knot where we remove all but 1 or 2 crossings from twist regions with more than N crossings in even pairs. Then the Euler characteristics of the blue and red twisted surfaces satisfy*

$$|\chi(S_{R,2}) + \chi(S_{B,2})| = \mathrm{cr}(K_2) + 2\mathrm{tw}_\mathrm{N}(K) - \chi(F).$$

Proof. Note that the union of the two (non-twisted) checkerboard surfaces in K_2 will have each crossing arc of K_2 appear twice, once in the red surface, and once in the blue surface. Then, by splitting the crossing arcs of K_2 into two homotopic arcs and attaching the blue regions to one arc and the red regions to the other, we get that the Euler characteristic for the disjoint union of these non-twisted surfaces is $\mathrm{cr}(K_2) - \chi(F)$. To obtain R_2 and B_2, the punctured surfaces, we remove two disks from each crossing circle, and get $|\chi(B_2) + \chi(R_2)| = \mathrm{cr}(K_2) + 2\mathrm{tw}_N(K) - \chi(F)$. Then, finally, to obtain the twisted surfaces, we connect these punctures by annuli or Möbius bands. This doesn't change the Euler characteristic, and so we get our desired result. □

4.3 Calculation of Cusp Area Bounds

We are now ready to give the proof of Theorem 4.1.1.

Proof. We will begin with the upper bound, using Theorem 2.1.7 (from [11]) to almost immediately get our result. We know that $|\chi(S)| = \mathrm{cr}(K_2) + 2\mathrm{tw}_N(K) - \chi(F) \leq 120\mathrm{tw}_\pi(K) - \chi(F)$. Then, note that, by construction, our twisted surfaces must intersect at least once for every twist region in our projection, so $i(\partial S_{B,2}, \partial S_{R,2}) \geq \mathrm{tw}_\pi(K)$. Finally, as $\chi(S_{R,2})$ and $\chi(S_{B,2})$ are both negative, we get

$$|\chi(S_{B,2})| + |\chi(S_{R,2})| = |\chi(S_{B,2}) + \chi(S_{R,2})| \leq 120\mathrm{tw}_\pi(K) - \chi(F).$$

So, combining with Theorem 2.1.7, we get:

$$\mathrm{Area}(\partial C) \leq \frac{18(|\chi(S_{B,2})| + |\chi(S_{R,2})|)^2}{i(\partial S_{B,2}, \partial S_{R,2})} \leq \frac{18(120\mathrm{tw}_\pi(K) - \chi(F))^2}{\mathrm{tw}_\pi(K)},$$

which finishes the proof of the upper bound.

To proceed with the proof of the lower bound define

$$A(F) = \frac{18^{1/4}}{5776\pi^4} \; \frac{1}{2^{3/4}\pi(120 - \chi(F))^6 + 4(120 - \chi(F))^5}.$$

To prove the lower bound of $CV(K)$ in Theorem 4.1.1 we will show that

$$A(F)\ (\mathrm{tw}_\pi(K) - \chi(F)) \leq \mathrm{area}(\partial C), \;\; \textit{and} \;\; A(F) \geq \frac{5.8265514 \times 10^{-7}}{(120 - \chi(F))^6}.$$

Since $2CV(K) = \text{area}(\partial C)$, we obtain the desired bound.

Let S be the disjoint union of the twisted blue and red surfaces, $S_{B,2} \sqcup S_{R,2}$, with an immersion f into $Y \setminus K$. Then let H be the maximal cusp of $Y \setminus K$, and H_S the embedded cusp for S contained in $f^{-1}(H)$ with area at least $2^{5/4}\text{tw}_\pi(K)$ from Lemma 4.2.6. Note that $\text{cr}(K_2)$ is less than or equal to 120 times the number of twist regions with less than 121 twists. Thus $\text{cr}(K_2) + 2\text{tw}_N(K) \leq 120\text{tw}_\pi(K)$. Then, by Lemma 4.2.7, we have

$$\text{Area}(S) = 2\pi(\text{cr}(K_2) + 2\text{tw}_N(K) - \chi(F)) \leq 2\pi(120\text{tw}_\pi(K) - \chi(F)).$$

Recall that

$$k = \text{Area}(H_S)/\text{Area}(S) \geq \frac{\sqrt[4]{2}}{\pi} \frac{\text{tw}_\pi(K)}{120\text{tw}_\pi(K) - \chi(F)}.$$

The inequality comes from combining our upper bound on $\text{Area}(S)$ and the lower bound on $\text{Area}(H_S)$ of $2^{5/4}\text{tw}_\pi(K)$ from Lemma 4.2.6. Then, let $d = \log(2/k)$, so $ke^d = 2$, $\sinh(d) = \frac{1}{k} - \frac{k}{4}$, and $e^d = \frac{2}{k}$. Note that $\frac{1}{k} \leq 2^{-1/4}\pi(120 - \chi(F))$, so we can say the following:

$$\begin{aligned}
\frac{2}{k} - 1 &< 2^{3/4}\pi\left(120 - \frac{\chi(F)}{\text{tw}_\pi(K)}\right) = 2^{3/4}\pi O(F,K) \\
\sinh(d) + 2\pi &\leq 2^{-1/4}\left(120 - \frac{\chi(F)}{\text{tw}_\pi(K)}\right) - \frac{k}{4} + 2\pi \\
&< 2^{-1/4}O(F,K) + 2\pi \\
\left(\frac{2}{k} - 1\right)(\sinh(d) + 2\pi) &< 2^{-1/4}\pi(2^{3/4}\pi O(F,K)^2 + 4O(F,K)),
\end{aligned}$$

where $O(F,K) = 120 - \frac{\chi(F)}{\text{tw}_\pi(K)}$. As a side note, while we will leave in $O(F,K)$ for our calculations, once we choose a projection surface, we immediately get an upper and lower bound: $O(F,K)$ is between 120 and $120 - \chi(F)$.

Putting this together, by Theorem 4.2.2, we have at least

$$\frac{(ke^d - 1)\pi}{(e^d - 1)(\sinh(d) + 2\pi)}|\chi(S)| > \frac{2^{1/4}}{2^{3/4}\pi O(F,K)^2 + 4O(F,K)}|\chi(S)|$$

embedded disjoint bi-infinite geodesics, with both ends in H_S, of length at most $2d$. Now we can map these arcs into $Y \setminus K$. As S is essential, each arc remains essential under this mapping. Furthermore, the length of these arcs can only decrease, so the upper bound of length $2d$ remains. Some of these arcs may be homotopic now in $Y \setminus K$, however, by Corollary 4.2.5, there are at most $2(3N-2) = 722$ such arcs. So, dividing out, we have at least

$$\frac{2^{1/4}}{722(2^{3/4}\pi O(F,K)^2 + 4O(F,K))}|\chi(S)|$$

disjoint embedded geodesics with both ends in H and length at most $2d$ with respect to H.

Now we can use Lemma 4.2.3 to say that, if the above number is p,

$$\mathrm{Area}(\partial H) \geq \sqrt{3}e^{-2(2d)}.$$

Before we put all the math together, to make the numbers look nicer, we will find an upper bound for e^{-4d}:

$$\begin{aligned} e^{4d} &= \left(\frac{2}{k}\right)^4 \\ &\leq \left(2^{3/4}\pi O(F,K)\right)^4 = 8\pi^4 O(F,K)^4 \\ e^{-4d} &\geq \frac{1}{8\pi^4 O(F,K)^4} \end{aligned}$$

Now, combining this, with all of our other bounds, we get that our cusp volume satisfies

$$\begin{aligned} \mathrm{Area}(\partial H) &\geq \frac{2^{1/4}\sqrt{3}|\chi(S)|}{722(2^{3/4}\pi O(F,K)^2 + 4O(F,K))}e^{-4d} \\ &\geq \frac{18^{1/4}}{722(2^{3/4}\pi O(F,K)^2 + 4O(F,K))} \; \frac{1}{8\pi^4 O(F,K)^4} \left(\mathrm{tw}_\pi(K) - \chi(F)\right). \end{aligned}$$

Because [26] gives us a lower bound for the case when $F = S^2$, we will focus on the other cases. When $\chi(F) \neq 2$, $O(F,K) \leq 120 - \chi(F)$, so we can replace $O(F,K)$ in our equation with $120 - \chi(F)$, and we get a lower bound of $A(F)(\mathrm{tw}_\pi(K) - \chi(F))$, and so are done with the lower bound, and our proof. □

Remark 4.3.1. *In the course of proving the lower bound in Theorem 4.1.1, we proved that we can replace $A(F)$ with a function $E(F,K)$ depending on the Euler characteristic of F and and*

$\mathrm{tw}_\pi(K)$. *In particular, we can replace* $A(F)$ *with the following function for a more precise lower bound:*

$$E(F,K) = \frac{18^{1/4}}{5776\pi^4}\ \frac{1}{2^{3/4}\pi O(F,K)^6 + 4O(F,K)^5} \quad where \quad O(F,K) = 120 - \frac{\chi(F)}{\mathrm{tw}_\pi(K)}.$$

We can get a better lower bound if we replace $A(F)$ *with* $E(F,K)$*. Below are some value for* $E(F,K)$ *and* $A(F)$ *for different values of* $\chi(F)$ *and* $\mathrm{tw}_\pi(K)$*, showing that, while as* $\chi(F)$ *decreases, both our bounds get worse, as* $\mathrm{tw}_\pi(K)$ *increases, the bound* $E(F,K)$ *starts to improve.*

$\chi(F)$	$A(F)$	$E(F,K)$ *when* $\mathrm{tw}_\pi(K) = 10$	$E(F,K)$ *when* $\mathrm{tw}_\pi(K) = 100$
0	2.3059×10^{-19}	2.3059×10^{-19}	2.3059×10^{-19}
-2	2.0884×10^{-19}	2.2830×10^{-19}	2.3036×10^{-19}
-4	1.8945×10^{-19}	2.2604×10^{-19}	2.3013×10^{-19}
-100	6.0903×10^{-21}	1.4272×10^{-19}	2.1941×10^{-19}

4.4 Additional Results

Now we can apply Theorem 4.1.1 to obtain estimates of slope lengths on the maximal cusp of K. We discuss an upper bound on the length on meridian curves for K and lower length bounds for non-meridian slopes.

Corollary 4.1.3. *Let* Y *be a closed, irreducible 3-manifold and* $F \subset Y$ *a closed surface with* $\chi(F) < 0$ *and such that* $Y \setminus N(F)$ *is atoroidal. Suppose that* K *is a weakly generalized alternating knot with a projection* $\pi(K) \subset F$ *that is weakly twist reduced, and we have* $e(\pi(K), F) > 2$*, and* $r(\pi(K), F) > 4$*. Let* μ *be the meridian and* σ *any non-meridian slope on maximal cusp of* $Y \setminus K$*. Then, we have*

$$\ell(\mu) \leq \frac{720\ (\mathrm{tw}_\pi(K) - \chi(F))}{\mathrm{tw}_\pi(K)}$$

$$\ell(\sigma) \geq \left(\frac{5.8265514 \times 10^{-7}}{720\ (120 - \chi(F))^6}\right)\ \mathrm{tw}_\pi(K),$$

where $\mathrm{tw}_\pi(K)$ *is the twist number of* $\pi(K)$*.*

Proof. As shown in the proof of Theorem 4.1.1 above, if $S_{B,2}$ and $S_{R,2}$ are our two twisted essential surfaces, then:

$$|\chi(S_{B,2})| + |\chi(S_{R,2})| \leq 120\text{tw}_\pi(K) - \chi(F) \leq 120(\text{tw}_\pi(K) - \chi(F)),$$

where the right inequality comes from the fact that $\chi(F) \leq 0$, so $-\chi(F) \leq -120\chi(F)$. Also, $i(\partial S_{B,2}, \partial S_{R,2}) \geq \text{tw}_\pi(K)$. Using these two results with Burton and Kalfagianni [11] on upper bounds for slope lengths, we get our desired upper bound on $\ell(\mu)$:

$$\begin{aligned}\ell(\mu) &\leq \frac{6|\chi(S_{B,2})| + |\chi(S_{R,2})|}{i(\partial S_{B,2}, \partial S_{R,2})} \\ &\leq \frac{6(120\text{tw}_\pi(K) - \chi(F))}{\text{tw}_\pi(K)} \\ &\leq \frac{720(\text{tw}_\pi(K) - \chi(F))}{\text{tw}_\pi(K)}.\end{aligned}$$

Next, Theorem 4.1.1 tells us that $A(F)(\text{tw}_\pi(K) - \chi(F)) \leq \text{Area}(\partial C)$. Then, as in [13],

$$\ell(\mu)\ell(\sigma) \geq \text{Area}(\partial C)\Delta(\mu, \sigma).$$

Then, using our bound for Area(∂C), and that μ and σ must intersect at least once, we get

$$\ell(\mu)\ell(\sigma) \geq A(F)(\text{tw}_\pi(K) - \chi(F)).$$

Next, we use our upper bound for $\ell(\mu)$ to get:

$$\frac{720(\text{tw}_\pi(K) - \chi(F))}{\text{tw}_\pi(K)}\ell(\sigma) \geq \ell(\mu)\ell(\sigma) \geq A(F)(\text{tw}_\pi(K) - \chi(F)).$$

Finally, we rearrange:

$$\begin{aligned}\ell(\sigma) &\geq A(F)\text{tw}_\pi(K) - \chi(F)\frac{\text{tw}_\pi(K)}{720(\text{tw}_\pi(K) - \chi(F))} \\ &= A(F)\frac{\text{tw}_\pi(K)}{720},\end{aligned}$$

giving the desired inequality. □

For our final result, recall that *Dehn surgery* along a knot $K \subset Y$ involves removing a torus neighborhood of the knot from Y and replacing it with a torus glued along a slope σ of the knot to get a manifold M. An important result related to such surgery is as follows.

Theorem 4.4.1 (Theorem 1.1 [18]). *Let M be a complete, finite-volume hyperbolic manifold with a single cusp C. Let σ be a slope on ∂C with length $\ell(\sigma)$ greater than 2π, and let $M(\sigma)$ denote the 3-manifold obtained by Dehn filling M along σ. Then, $M(\sigma)$ is hyperbolic and we have*

$$\mathrm{vol}(M(\sigma)) \geq \left(1 - \left(\frac{2\pi}{\ell(\sigma)}\right)^2\right)^{3/2} \mathrm{vol}(M).$$

This, combined with Theorem 2.2.7 from [23], gives us our other result.

Theorem 4.1.4. *Let K be a weakly generalized alternating knot on a surface $F \neq S^2$ in a closed manifold Y with conditions as in Theorem 4.1.1. Suppose that*

$$\mathrm{tw}_\pi(K) > \left(2.32928 \times 10^{10}\right) \, (120 - \chi(F))^6 \, .$$

Then, any manifold M obtained from non-meridional surgery along K is hyperbolic and

$$\frac{v_8}{4} \left(\mathrm{tw}_\pi(K) - \chi(F)\right) \leq \frac{1}{2}\mathrm{vol}(Y \setminus K) \leq \mathrm{vol}(M) < \mathrm{vol}(Y \setminus K).$$

Here $v_8 \approx 3.66386$ is the volume of a regular hyperbolic ideal octahedron.

Proof. As K is a weakly generalized alternating knot and, by Lemma 2.2.8, has disk regions, we get the leftmost inequality from Theorem 2.2.7. Then, by Corollary 4.1.3, as we've taken the number of twist regions to be high enough, any non-meridional slope must have length strictly greater than 2π. This allows us to use Theorem 4.4.1 from [18] to get our center inequality. Finally, we get $\mathrm{vol}(M) < \mathrm{vol}(Y \setminus K)$ from the fact that volume must decrease under surgery.w □

CHAPTER 5

JONES POLYNOMIAL

5.1 Introduction

Parts of this chapter is based on joint work with Efstratia Kalfagianni that will be published in Fundamenta Mathematicae [8].

The goal of this chapter is to show that the hyperbolic volume of certain weakly generalized alternating links on a closed surface F in a compact, irreducible 3-manifold M is bounded below in terms of a Kauffman bracket function defined on link diagrams on F. Further, for $M = F \times [-1, 1]$, this function leads to a Jones polynomial link invariant and coefficients of it provide 2-sided linear bounds on the volume of hyperbolic alternating links.

As a corollary, we obtain that the twist-number of a twist-reduced, weakly generalized alternating link projection with representativity $r(\pi(L), F) > 4$ and checkerboard disk regions, is an invariant of the link in $F \times [-1,\ 1]$. This generalizes work of Dasbach and Lin [15] and Futer, Kalfagianni and Purcell [17, 18], who have obtained similar results for families of links in S^3.

Let M be an irreducible, compact 3-manifold with or without boundary, and let $F \subset M$ be a closed, orientable, embedded surface. Then, if we have a link $L \subset F \times [-1, 1]$, we can get a projection under the map $\pi : F \times [-1, 1] \to F = F \times \{0\}$. We can then define, for connected F, a Kauffman bracket function, and construct a family of Jones-like polynomials

$$J_X(\pi(L)) = a_{m,X}t^m + a_{m-1,X}t^{m-1} + \cdots + b_{n+1,X}t^{n+1} + b_{n,X}t^n,$$

where X is a collection of simple disjoint closed curves on F. We will give the precise definitions for these polynomials in the next section. We can consider the sums of these coefficients,

$$a_k = \sum_X a_{k,X} \qquad b_k = \sum_X b_{k,X},$$

where we sum over all possible collections of simple disjoint closed curves on F. Then we are able to say the following:

Theorem 5.1.1. *Let M be an irreducible, compact 3-manifold with empty or incompressible boundary. Let $F \subset M$ be an incompressible, closed, orientable surface such that $M \setminus N(F)$ is atoroidal and $\partial - anannular$. Suppose that a link L admits a weakly generalized alternating projection $\pi(L) \subset F$ that is reduced, twist-reduced and with all regions of $F \setminus \pi(L)$ disks. Finally, suppose that $r(\pi(L), F) > 4$. Then L is hyperbolic and*

$$\mathrm{vol}(M \setminus L) \geq v_8 \max\{|a_{m-1}|, |b_{n+1}|\} - \frac{1}{2}\chi(\partial M),$$

where $v_8 = 3.66386\ldots$ is the volume of an ideal octahedron.

Now, given an alternating link projection $\pi(L)$ as in the statement of Theorem 5.1.1, let S_A and S_B denote the two checkerboard surfaces corresponding to $\pi(L)$, and let M_A and M_B denote the manifolds obtained by cutting $M \setminus L$ along S_A and S_B, respectively. We can show the following.

Theorem 5.1.2. *Let M and F be as in Theorem 5.1.1 Suppose that a link L admits a weakly generalized alternating projection $\pi(L) \subset F$ that is twist-reduced and with all regions of $F \setminus \pi(L)$ disks. Then we have the following.*

1. $\chi(\mathrm{guts}(M_A)) = -|a_{m-1}| + \frac{1}{2}\chi(M)$
2. $\chi(\mathrm{guts}(M_B)) = -|b_{n+1}| + \frac{1}{2}\chi(M)$
3. $\mathrm{tw}_F(\pi(L)) = |a_{m-1}| + |b_{n+1}| + \chi(F)$.

Further, if $\pi(L)$ is a cellularly embedded, twist-reduced, reduced alternating diagram, we still get (3).

Theorem 5.1.1 follows directly from 5.1.2. We can also obtain the following corollary for links in thickened surfaces:

Corollary 5.1.3. *Let L be a link in $F \times [-1, 1]$ that admits a weakly generalized alternating projection $\pi(L) \subset F$ that is reduced, twist-reduced and with all regions of $F \setminus \pi(L)$ disks. Then any two projections of Lhave the same twist number. That is,* $\mathrm{tw}_F(\pi(L))$ *is an isotopy invariant of L.*

The rest of this chapter is outlined as follows. First, in Section 2, we introduce the exact functions and polynomials we will be working with. Then, in Section 3, we will show that there is a relation between certain coefficients of our polynomials and certain reduced state graphs. Finally, in Section 4, we will prove Theorem 5.1.2, and show how Theorem 5.1.1 follows.

5.2 Background

Our definition of a Kauffman bracket function for links on general surfaces will follow, in many regards, the definition of the Kauffman bracket function for links on S^2, as defined in Chapter 2. Let M be a 3-manifold, and let $F \subset M$ be an orientable surface. Given a link $L \subset F \times [-1, 1] \subset M$, we can consider it's image under the projection $\pi : F \times [-1, 1] \to F = F \times \{0\}$.

Let $\mathcal{D}(F)$ denote the set of all (unoriented) link diagrams on F, taken up isotopy on F. Also let X_F denote the set of all collections of disjoint simple closed curves (a.k.a multi-curves) on F. We define a Kauffman bracket function

$$\langle \ \rangle : \mathcal{D}(F) \longrightarrow \mathbb{Z}[A, A^{-1}, (A^2 - A^{-2})^{-1}]X_F,$$

by the following skein relations.

1. $\langle$ 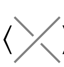 $\rangle = A\langle$ 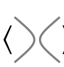 $\rangle + A^{-1}\langle$ 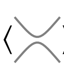 $\rangle$

2. $\langle L \sqcup$  $\rangle = (-A^2 - A^{-2})\langle L\rangle$

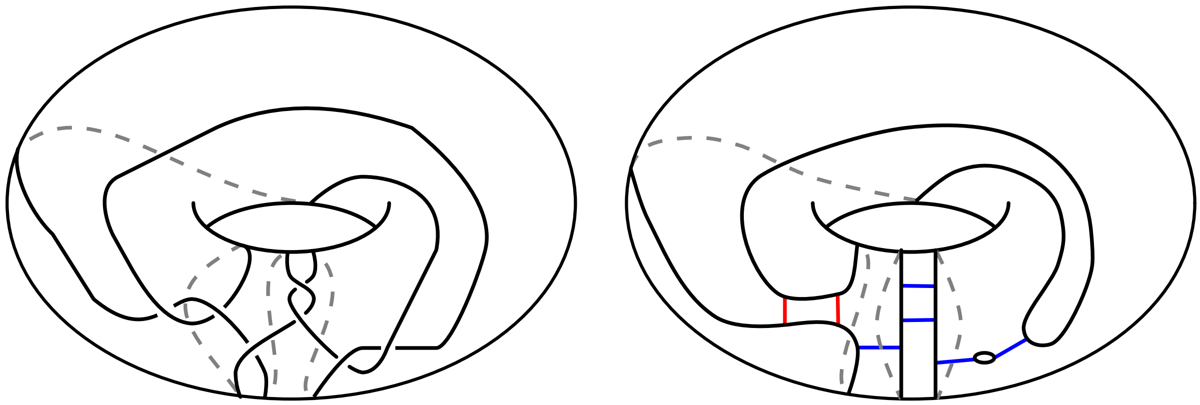

Figure 5.1. A Kauffman State *(right)* for an alternating knot on the torus *(left)*. The color of the edges represent the type of resolution at each crossing. Note that we now have states that don't bound disks on T^2.

3. $\langle \bigcirc \rangle = -A^2 - A^{-2}$

As we are working on a surface that may have essential curves, we also require that the unknots in the above relations bound disks on F. We can then describe $\langle D \rangle$ as a function in $\mathbb{Z}[A, A^{-1}]X_F$, as follows.

Recall that, given a Kauffman state s, $a(S)$ is the number of A-resolutions, $b(S)$ the number of B-resolutions, and $|s|$ the number of state circles. We will need to make a distinction between contractible and non-contractible state circles, so we say s_t is the collection of contractible state circles, s_{nt} the collection of non-contractible state circles, and $|s|_t$ the number of contractible state circles.

Given a link diagram $D = \pi(L)$ in $\mathcal{D}(F)$, we define

$$\langle D \rangle_0 = \sum_{\{s | s_{nt} = \emptyset\}} A^{a(s)-b(s)} (-A^2 - A^{-2})^{|s|_t} \tag{5.2.1}$$

$$\langle D \rangle_X = \sum_{\{s | s_{nt} = X\}} A^{a(s)-b(s)} (-A^2 - A^{-2})^{|s|_t} \tag{5.2.2}$$

where $X \in X_F$ is a collection of disjoint simple closed curves on F. Then

$$\langle D \rangle = \langle D \rangle_0 + \sum_{X \in X_F} \langle D \rangle_X X.$$

We can then define the Jones-like polynomials as follows:

$$J_0(\pi(L)) = \left((-1)^{w(D)}A^{-3w(D)}\langle\pi(L)\rangle_0\right)\Big|_{t=A^4}$$
$$J_X(\pi(L)) = \left((-1)^{w(D)}A^{-3w(D)}\langle\pi(L)\rangle_X\right)\Big|_{t=A^4}.$$

We can write these out with coefficients:

$$J_0(\pi(L)) = a_{m,0}t^m + a_{m-1,0}t^{m-1} + \cdots + b_{n+1,0}t^{n+1} + b_{n,0}t^n$$
$$J_X(\pi(L)) = a_{m,X}t^m + a_{m-1,X}t^{m-1} + \cdots + b_{n+1,X}t^{n+1} + b_{n,0}t^n.$$

As $J_0(\pi(L))$ and $J_X(\pi(L))$ might have different highest or lowest degrees, we take m to be the highest degree over all possible $J_X(\pi(L))$ and $J_0(\pi(L))$, and n to be the lowest. Then, if any polynomial has a lower maximum degree, we take all higher coefficients to be 0.

Note that, in general, these are not link invariants, but instead diagram invariants. As discussed in [8, 10],these polynomials can be promoted to isotopy invariants in the special case that $M = F \times [-1, 1]$.

5.2.1 Reduced Graphs

Recall that a Kauffman state s gave us two different graphs, the state graph G_s and the reduced graph G'_s. When working with generalized link diagrams, we will also get similar graphs, however a slight modification of definition is necessary. We let G_s be the the graph with a vertex for each state circle, and edges connecting state circles that are adjacent across crossings of $\pi(L)$. Then the reduced graph G'_s is obtained from G_s by, whenever two edges of G_s bound a disk on F, we remove one of them. Then we let e'_s denote the number of edges of G'_s. This will mean that two edges adjacent to the same vertices that form an essential curve on F will both still be in G'_s.

In particular, we will be concerned with the reduced graphs for the all A state s_A and the all B state s_B. We call these G'_A and G'_B, respectively.

5.3 Coefficients and Graphs

As a note before we begin, many of our proofs will work with a more general category of links: ones which are cellularly embedded and with no removable nugatory crossings (a crossing is removable nugatory if there is a disk on F that intersects the diagram at only the crossing). These are called *reduced* diagrams Boden, Karimi, and Sikora [9].

As we will use some results for reduced diagrams, it's important to show that weakly generalized alternating link diagrams are also reduced diagrams. To see this, note that a weakly generalized alternating link must be cellularly embedded, as the complement has disk regions, and there can be no nugatory crossings, as if there were, we would have a disk violating weakly prime.

We begin now by examining the coefficients for weakly generalized alternating diagrams.

Lemma 5.3.1. *Suppose $\pi(L)$ is twist-reduced, reduced alternating diagram on a closed orientable surface F, with $F \setminus \pi(L)$ disks. Suppose also that $\pi(L)$ is checkerboard colorable. Then $|a_m| = |b_n| = 1$.*

Proof. As the regions of $F \setminus \pi(L)$ are disks, the all A state, s_A, has all state circles compressible disks (in particular, we can view s_A as the collection of all regions of $F \setminus \pi(L)$ colored a single color). Then s_A's contributions to $J(\pi(L))$ will have highest degree $(c + 2|s_A|_t)$, and the highest term contributed will be ± 1.

Now suppose we have a state s_k with exactly k B-resolutions, and a state s_{k-1} which differs from s_k by exactly one resolution change from an A-resolution to a B-resolution. This resolution change can, at most, increase the number of contractible state circles by 1. Look, then, at any state that has exactly one B resolution, s_1. As s_A has only contractible state circles, there are exactly three ways a resolution change from s_A to a state s_1 can act.

1. We could merge two contractible circles into a single contractible circle. Then $|s_1|_t = |s_A|_t - 1$.

2. We could split a single contractible circle into two noncontractible circles. Then $|s_1|_t = |s_A|_t - 1$.

3. We could split a single contractible circle into two contractible circles. Then $|s_1|_t = |s_A|_t + 1$.

Note that the highest degree s_1 can contribute to a polynomial is $(c + 2|s_1|_t - 2)$. The only way s_1 could contribute to the highest degree of $J(\pi(L))$ is if the last case, (3), happens. However, note that for this to happen, the resolution change must occur at a crossing that meets a single state circle of s_A twice. Then we can construct a loop γ that intersects the diagram at exactly this crossing, and then follows along the boundary of the state circle. As there are no nugatory crossings, this must mean that γ is a non-separating curve. In particular, γ must be essential. After the resolution change, we will get two state circles that must have essential boundary, and thus be noncontractible. But then we must be in the second case, (2), and have a contradiction. So no s_1 can contribute to the highest degree.

Finally, look at any state s_k with exactly k B resolutions. This must have highest degree $(c + 2|s_k|_t - 2k)$. In order to contribute to the highest degree, we must have $|s_k|_t = |s_A|_t + k$. However, as each resolution change can, at most, introduce one contractible state circle, and we know that the change from s_A to s_1 does not, we must have s_k not contribute to the highest degree. So then $a_m = \pm 1$, and we are done. □

Before we begin, we introduce one last definition that we will use in the later half of the proof for the second coeffient.

Definition 5.3.2. Let $\pi(L) \subset F$ be a reduced alternating diagram, and s a Kauffman state, with reduced graph G'_s. We say two edges, ℓ and ℓ' in G'_s are *parallel* if they are adjacent to the same vertices (or, equivalently, connect the same state circles) and, when placed on F with the state circles, there is a disk in F with boundary containing ℓ, ℓ', and a portion of the state circles they connect.

Lemma 5.3.3. *Suppose $\pi(L)$ is a twist-reduced, reduced alternating diagram on a closed orientable surface F, with $F \setminus \pi(L)$ disks. Suppose also that $\pi(L)$ is checkerboard colorable. Then*

1. $|a_{m-1}| = e'_A - |s_A|_t$.
2. $|b_{n+1}| = e'_A - |s_B|_t$.

Proof. First, note that s_A contributes to the second highest degree. In particular, it contributes $(-1)^{|s_A|_t}|s_A|_t$. Next, we also know that any state with exactly one B resolution, s_1, will contribute to the second highest degree. We know this because we have one less contractible circle than s_A, by work above, one less A resolution, and one more B resolution; then we can use either Equation 5.2.1 or 5.2.2. Now suppose we have a sequence of states $s_A, s_1, \cdots, s_k$, each differing from the previous state by a resolution change from an A-resolution to a B-resolution. Then, for s_k to contribute, we must have each resolution change after the first introducing a new contractible state circle. We know that s_1 either has all contractible state circles, or has exactly two non-contractible state circles that are parallel to each other. Focus first on the states s_1 and s_2. Our goal is to show that s_2 will only contribute to $|a_{m-1}|$ if the edge changed from s_1 is parallel to the edge changed from s_A.

There are three possible ways a resolution change can act and still have s_2 contribute to the second highest degree:

1. We split a single contractible circle into two contractible circles.
2. We merge two noncontractible circles into a single contractible circle.
3. We split a single noncontractible circle into a noncontractible circle and a contractible circle.

We focus on each of these cases individually. In the case of (1), where we split a contractible circle into two contractible circles and don't change the pattern of noncontractible state circles, the edge changed, ℓ_2 would have to, in s_1, be adjacent to a single state circle. If this edge is parallel to the edge changed from s_A to s_1, ℓ_1, we are fine. Otherwise, we must have both ℓ_2 and ℓ_1 appear in G'_A, and so ℓ_1 and ℓ_2, with a portion of some state circles, forms an essential curve on F. Then s_2 will introduce a noncontractible circle, and so we cannot be in the first case.

In the case of (2), where we merge two noncontractible circles into a single contractible circle, denote the edge changed in the resolution change from s_A to s_1 as ℓ_1, and the edge changed from s_1 to s_2 as ℓ_2. We know that the two noncontractible circles must be parallel, and so, to merge them to get a contractible circle, we must have ℓ_1 parallel to ℓ_2. Look at ℓ_2 in s_A, and construct a new state s'_1, which has all A-resolutions except at ℓ_2. In Figure 5.2, the leftmost state is s_1, the center s_A, and the right s'_1.

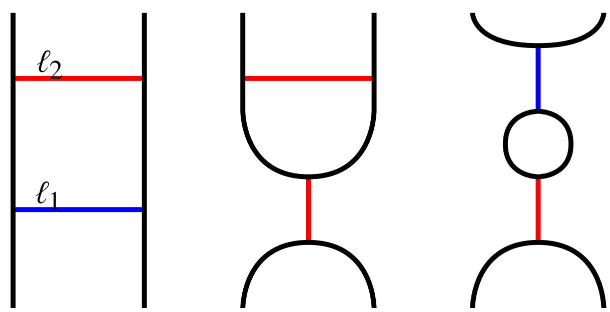

Figure 5.2. *(Left)* The state s_1 gotten from s_A by merging two noncontractible circles into a single contractible circle; *(Center)* The all A state, s_A; *(Right)* The modified state, s'_1, with exactly one A resolution.

Note that s'_1 must, then, introduce a new contractible state circle from s_A. In addition, it will have the same types of non-contractible state circles as s_A. Then, by Boden, Karimi, and Sikora [9, Theorem 11], any state of a reduced alternating diagram (and thus a weakly generalized alternating diagram) with one B resolution must have, at most, the same number of contractible state circles as s_A, so we get a contradiction.

Finally, in the case of (3), we split a noncontractible state into a contractible circle and a noncontractible circle. Label the edges as before: ℓ_1 is the edge first changed in the resolution change from s_A to s_1, and ℓ_2 the edge changed from s_1 to s_2. There are two possibilities for ℓ_2: either it is parallel to ℓ_1 in s_A, or it is not. If the two edges are not parallel, then ℓ_2 must have it's endpoints both to the same side of ℓ_1, and cut off a contractible circle disjoint from

ℓ_1. Then we can construct a new state, s_1' that has all A resolutions except for at ℓ_2, where we have a B resolution. As ℓ_2 cuts off a contractible circle disjoint from ℓ_1, s_1' must have more contractible circles than s_A, a contradiction. So then ℓ_1 and ℓ_2 must be parallel in s_A.

Note that, in order for s_2 to contribute to the second highest degree, the second edge must be parallel to the first. If it isn't, it either does not introduce a new contractible circle (like in the first case), or we can find a state with 1 B resolution that has more contractible circles than s_A. We can do this same process for any state s_k with k B resolutions, and so all states that contribute to $|a_{m-1}|$ must have B resolutions at only parallel edges.

Also note, if any two edges are parallel, they must be part of the same twist region. To see this, as parallel edges bound a disk on F, we get a disk whose boundary intersects the link diagram at exactly two crossings. As the diagram is twist reduced, this must mean both of these crossings, and thus the edges, belong to the same twist region.

So then we get a contribution to $|a_{m-1}|$ only when we make resolution changes to edges that, in s_A, are all parallel to each other. As we are asking about $|a_{m-1}|$, we will leave off the essential curve variable X during our calculations. Each such state will contribute $(-1)^{|s_A|-2}(-1)^j$, where j is the number of changed edges, and there are $\binom{k}{j}$ possible states with exactly j resolution changes, where k is the total number of parallel edges. Using the binomial theorem, we get that each family of parallel edges must contribute $(-1)^{|s_A|-1}$ to the second highest coefficient. As each family of parallel edges corresponds to a single edge in G_A', and factoring in the contribution to the second coefficient from s_A, we get that

$$\begin{aligned}|a_{m-1}| &= |(-1)^{|s_A|_t-1}e_A' + (-1)^{|s_A|_t}(|s_A|_t)| \\ &= |(-1)^{|s_A|_t-1}(e_A' - |s_A|_t)| \\ &= e_A' - |s_A|_t\end{aligned}$$

and so we are done with part (1) of the theorem.

To prove part (2), we can apply the same argument to the diagram D^*.

□

5.4 Twist Number, Volume, and Coefficients

Finally, in this section, we will prove Theorems 5.1.1 and 5.1.2 using Lemmas 5.3.1 and 5.3.3. We will start by proving the following relation between the coefficients, the twist number, and the Euler characteristic of F:

Theorem 5.4.1. *Suppose that $\pi(L)$ is a twist-reduced, reduced alternating link diagram with twist number $t_F(\pi(L))$ on a closed orientable surface $F \subset M$ of genus at least 1. Suppose also that the regions of $F \setminus \pi(L)$ are all disks. Then*

$$|a_{m-1}| + |b_{n+1}| = t_F(\pi(L)) - \chi(F)$$

Proof. First, note that we have

$$t_F(\pi(L)) = c - (c - e'_A) - (c - e'_B) = e'_A + e'_B - c,$$

where c is the number of crossings of $\pi(L)$. We can see this as follows: we call a twist region an A (or B) twist region if the edges in G_A (or G_B) are all parallel. Then G'_A (or G'_B) will have exactly one edge for each A (or B) twist region, as well as edges for each crossing in a B (or A) twist region. Note that $c - e'_A$ counts exactly the number of edges in G_A that aren't in G'_A. We can think of this as counting each edge that is an A twist region, except for one for each such region. Likewise, $c - e'_B$ counts the number of edges in a B twist region, except one for each such region. Then $(c - e'_A) + (c - e'_B) = c - t_F(\pi(L))$.

Next, as the regions of $F \setminus \pi(L)$ are disks, we also have

$$|s_A| + |s_B| = c + 2 - 2g(F) = c + \chi(F).$$

It is important to note that, as $F \setminus \pi(L)$ are all disks, then $|s_A| = |s_A|_t$ and $|s_B| = |s_B|_t$, as there are only contractible circles. Then we get

$$\begin{aligned}|a_{m-1}| + |b_{n+1}| &= e'_A - |s_A|_t + e'_B - |s_B|_t \\ &= e'_A + e'_B - |s_A| - |s_B| \\ &= t_F(\pi(L)) + c - c - \chi(F) \\ &= t_F(\pi(L)) - \chi(F),\end{aligned}$$

and we are done. □

We can now use this to prove Theorem 5.1.2:

Theorem 5.1.2. *Let M and F be as in Theorem 5.1.1 Suppose that a link L admits a weakly generalized alternating projection $\pi(L) \subset F$ that is twist-reduced and with all regions of $F \setminus \pi(L)$ disks. Then we have the following.*

1. $\chi(\mathrm{guts}(M_A)) = -|a_{m-1}| + \frac{1}{2}\chi(M)$
2. $\chi(\mathrm{guts}(M_B)) = -|b_{n+1}| + \frac{1}{2}\chi(M)$
3. $\mathrm{tw}_F(\pi(L)) = |a_{m-1}| + |b_{n+1}| + \chi(F)$.

Further, if $\pi(L)$ is a cellularly embedded, twist-reduced, reduced alternating diagram, we still get (3).

Proof. We have already proved the twist number result above, in Theorem **??**, as well as the last last statement about reduced alternating diagrams. So we now focus on the other two statements.

Let S_A and S_B be the checkerboard surfaces of $\pi(L)$, $X = M \setminus L$, and $M_A = X \backslash\backslash S_A$ and $M_B = X \backslash\backslash S_B$. As $\pi(L)$ is a weakly generalized alternating diagram, [23] gives us

$$\chi(\mathrm{guts}(M_A)) = \chi(F) + \frac{1}{2}\chi(\partial M) - |s'_B|, \tag{5.4.1}$$

where $|s'_B|$ is the number of non-bigon B regions of S_A. Then note that

$$\chi(F) = |s_A| - e'_A + |s'_B|.$$

In particular, by rearranging, we get

$$\chi(F) - |s'_B| = |s_A| - e'_A$$

Finally, by Lemma 5.3.3, we get

$$\chi(F) - |s'_B| = -|a_{m-1}|. \tag{5.4.2}$$

Then by combining Equations 5.4.1 and 5.4.2, we immediately get our result. The proof for guts(M_B) follows similarly, by swapping A and B.

$\square$

Theorem 5.1.1. *Let M be an irreducible, compact 3-manifold with empty or incompressible boundary. Let $F \subset M$ be an incompressible, closed, orientable surface such that $M \setminus N(F)$ is atoroidal and $\partial - anannular$. Suppose that a link L admits a weakly generalized alternating projection $\pi(L) \subset F$ that is reduced, twist-reduced and with all regions of $F \setminus \pi(L)$ disks. Finally, suppose that $r(\pi(L), F) > 4$. Then L is hyperbolic and*

$$\mathrm{vol}(M \setminus L) \geq v_8 \max\{|a_{m-1}|, |b_{n+1}|\} - \frac{1}{2}\chi(\partial M),$$

where $v_8 = 3.66386\ldots$ is the volume of an ideal octahedron.

Proof. Let S_A and S_B be the checkerboard surfaces of $\pi(L)$, $X = M \setminus L$ the link complement, and $M_A = X\backslash\backslash S_A$ and $M_B = X\backslash\backslash S_B$. Then, by [23]Theorem 9.1, we have

$$\mathrm{vol}(M) \geq -v_8\chi(\mathrm{guts}(M_A)) \qquad \text{and} \qquad \mathrm{vol}(M) \geq -v_8\chi(\mathrm{guts}(M_B)).$$

Then, by Theorem 5.1.2, we immediately get our result, and are done. $\square$

BIBLIOGRAPHY

BIBLIOGRAPHY

[1] C. Adams, C. Albors-Riera, B. Haddock, Z. Li, D. Nishida, B. Reinoso, and L. Wang, *Hyperbolicity of links in thickened surfaces*, Topology Appl. **256** (2019), 262–278. MR 3916014

[2] C. Adams, A. Colestock, J. Fowler, W. Gillam, and E. Katerman, *Cusp size bounds from singular surfaces in hyperbolic 3-manifolds*, Trans. Amer. Math. Soc. **358** (2006), no. 2, 727–741. MR 2177038

[3] Colin Adams, *Triple crossing number of knots and links*, J. Knot Theory Ramifications **22** (2013), no. 2, 1350006, 17. MR 3037297

[4] ______, *Waist size for cusps in hyperbolic 3-manifolds II*, Geom. Dedicata **203** (2019), 53–66. MR 4027583

[5] Colin C. Adams, *Toroidally alternating knots and links*, Topology **33** (1994), no. 2, 353–369. MR 1273788

[6] I. Agol, P. A. Storm, and W. P. Thurston, *Lower bounds on volumes of hyperbolic Haken 3-manifolds*, J. Amer. Math. Soc. **20** (2007), no. 4, 1053–1077, With an appendix by Nathan Dunfield. MR 2328715

[7] Brandon Bavier, *Cusp volumes of alternating knots on surfaces*, 2020.

[8] Brandon Bavier and Efstratia Kalfagianni, *Guts, volume and skein modules of 3-manifolds*, 2020, To be Published in Fundamenta Mathematicae.

[9] H. U. Boden, H. Karimi, and A. S. Sikora, *Adequate links in thickened surfaces and the generalized Tait conjectures*, arXiv e-prints (2020), arXiv:2008.09895.

[10] Hans U. Boden and Homayun Karimi, *The jones-krushkal polynomial and minimal diagrams of surface links*, 2019.

[11] Stephan D. Burton and Efstratia Kalfagianni, *Geometric estimates from spanning surfaces*, Bull. Lond. Math. Soc. **49** (2017), no. 4, 694–708. MR 3725490

[12] Abhijit Champanerkar, Ilya Kofman, and Jessica S. Purcell, *Geometry of biperiodic alternating links*, J. Lond. Math. Soc. (2) **99** (2019), no. 3, 807–830. MR 3977891

[13] Daryl Cooper and Marc Lackenby, *Dehn surgery and negatively curved* 3*-manifolds*, J. Differential Geom. **50** (1998), no. 3, 591–624. MR 1690741

[14] Marc Culler, Nathan M. Dunfield, Matthias Goerner, and Jeffrey R. Weeks, *SnapPy, a computer program for studying the geometry and topology of* 3*-manifolds*, Available at http://snappy.computop.org (2019).

[15] Oliver Dasbach and Xiao-Song Lin, *A volumish theorem for the jones polynomial of alternating knots*, Pacific Journal of Mathematics - PAC J MATH **231** (2007), 279–291.

[16] Oliver T. Dasbach, David Futer, Efstratia Kalfagianni, Xiao-Song Lin, and Neal W. Stoltzfus, *The Jones polynomial and graphs on surfaces*, J. Combin. Theory Ser. B **98** (2008), no. 2, 384–399. MR 2389605

[17] David Futer, Efstratia Kalfagianni, and Jessica Purcell, *Guts of surfaces and the colored Jones polynomial*, Lecture Notes in Mathematics, vol. 2069, Springer, Heidelberg, 2013. MR 3024600

[18] David Futer, Efstratia Kalfagianni, and Jessica S. Purcell, *Dehn filling, volume, and the Jones polynomial*, J. Differential Geom. **78** (2008), no. 3, 429–464. MR 2396249

[19] ———, *Cusp Areas of Farey Manifolds and Applications to Knot Theory*, International Mathematics Research Notices **2010** (2010), no. 23, 4434–4497.

[20] David Futer and Jessica S. Purcell, *Links with no exceptional surgeries*, Comment. Math. Helv. **82** (2007), no. 3, 629–664. MR 2314056

[21] Chuichiro Hayashi, *Links with alternating diagrams on closed surfaces of positive genus*, Mathematical Proceedings of the Cambridge Philosophical Society **117** (1995), 113 – 128.

[22] ———, *Links with alternating diagrams on closed surfaces of positive genus*, Math. Proc. Cambridge Philos. Soc. **117** (1995), no. 1, 113–128. MR 1297898

[23] Joshua A. Howie and Jessica S. Purcell, *Geometry of alternating links on surfaces*, Trans. Amer. Math. Soc. **373** (2020), no. 4, 2349–2397. MR 4069222

[24] Joshua A. Howie and J. Hyam Rubinstein, *Weakly generalised alternating links*, (2017), In Preparation.

[25] Marc Lackenby, *The volume of hyperbolic alternating link complements*, Proc. London Math. Soc. (3) **88** (2004), no. 1, 204–224, With an appendix by Ian Agol and Dylan Thurston. MR 2018964

[26] Marc Lackenby and Jessica S. Purcell, *Cusp volumes of alternating knots*, Geom. Topol. **20** (2016), no. 4, 2053–2078. MR 3548463

[27] ———, *Essential twisted surfaces in alternating link complements*, Algebr. Geom. Topol. **16** (2016), no. 6, 3209–3270. MR 3584257

[28] W. B. R. Lickorish, *An introduction to knot theory*, Graduate Texts in Mathematics, vol. 175, Springer-Verlag, New York, 1997. MR 1472978

[29] W. Menasco and M. Thistlethwaite, *The classification of alternating links*, Ann. of Math. (2) **138** (1993), no. 1, 113–171. MR 1230928

[30] George D. Mostow, *Quasi-conformal mappings in n-space and the rigidity of hyperbolic space forms*, Publications Mathématiques de l'IHÉS **34** (1968), 53–104 (en). MR 38 #4679

[31] Makoto Ozawa, *Non-triviality of generalized alternating knots*, J. Knot Theory Ramifications **15** (2006), no. 3, 351–360. MR 2217501

[32] Gopal Prasad, *Strong rigidity of Q-rank* 1 *lattices*, Invent. Math. **21** (1973), 255–286. MR 385005

[33] William P. Thurston, *Geometry and topology of three-manifolds*, Princeton University Math Department Notes, 1979.