ANGLED TRIANGULATIONS OF LINK COMPLEMENTS

A DISSERTATION
SUBMITTED TO THE DEPARTMENT OF MATHEMATICS
AND THE COMMITTEE ON GRADUATE STUDIES
OF STANFORD UNIVERSITY
IN PARTIAL FULFILLMENT OF THE REQUIREMENTS
FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

David Futer
July 2005

I certify that I have read this dissertation and that, in my opinion, it is fully adequate in scope and quality as a dissertation for the degree of Doctor of Philosophy.

Steven P. Kerckhoff
(Principal Advisor)

I certify that I have read this dissertation and that, in my opinion, it is fully adequate in scope and quality as a dissertation for the degree of Doctor of Philosophy.

Ralph L. Cohen

I certify that I have read this dissertation and that, in my opinion, it is fully adequate in scope and quality as a dissertation for the degree of Doctor of Philosophy.

Peter A. Storm

Approved for the University Committee on Graduate Studies.

Abstract

The goal of this thesis is to relate the projection diagram of a knot or link in S^3 to the geometry and topology of the link complement. We use the diagram of a link K to obtain a Dehn surgery description of K from a hyperbolic link L. The simple geometry of $S^3 \setminus L$ allows us to decompose it into ideal hyperbolic polyhedra, whose dihedral angles provide a lot of combinatorial information. One consequence of this approach is a mild condition on the original diagram that ensures K is hyperbolic and all its non-trivial Dehn fillings are hyperbolike. Another, closely related, consequence is a diagrammatic lower bound on the genus of K.

When K is an arborescent link, we use the correspondence between the link and a weighted tree to simplify the projection diagram into a particularly nice form. This simplified diagram then allows us to subdivide the link complement into hyperbolic polyhedra and tetrahedra whose dihedral angles fit together in a consistent fashion. An angled decomposition of this type implies that K is hyperbolic and provides a robust combinatorial framework for more detailed investigations into its geometry.

Acknowledgements

There are many people who have taught me, inspired me, and kept me sane during a six-year trek through graduate school.

Most importantly, I need to thank my advisor, Steve Kerckhoff. Steve has not only opened my eyes to the beautiful landscape of low-dimensional topology, but stood by with endless patience until its contours came into focus. Without his creative guidance, his support and encouragement, or a few well-timed kicks in the rear, I would never have needed to write the acknowledgements page of a thesis.

In my time at Stanford, I have been privileged to learn from several other mentors. Ralph Cohen has always been available with sage advice, while Pete Storm and Lenny Ng have fielded so many silly questions that one wonders how they've had time to work on their own projects. I have also been lucky to bounce ideas off many fellow students, including Josh Sabloff, Henry Segerman, Robert Lipshitz, and Jian He.

This thesis owes an especially large debt to Jessica Purcell. While we were reading papers and learning the terrain of research, she blazed a path through many conceptual thickets that would surely have ensnared me. In her own thesis research, Jessica developed most of the ideas about augmented links that underlie the first half of this document. Several of the figures in Chapters 1 and 3 are hers.

Finally, I must thank my parents for nurturing my interest in the world, for moving to the United States, and for enabling me to come this far. My muse, Simone Levine, has never ceased reminding me that life is wacky, unpredictable, and beautiful.

Contents

List of Tables

List of Figures

Chapter 1

Introduction

The goal of this thesis is to use a projection diagram of a knot or link to obtain geometric and topological information about the link complement. We provide a mild diagrammatic condition that implies the link complement has a complete hyperbolic structure, and a slightly stronger one that shows any non-trivial Dehn filling of the link has the topological properties of a hyperbolic manifold. For arborescent links, we also describe a concrete way to subdivide the link complement into ideal hyperbolic polyhedra whose dihedral angles fit together in a consistent fashion.

In the first four sections of this chapter, we give the relevant definitions about knots and their diagrams, state the main results of the thesis, and relate them to recent work in the field. The final section outlines the organization of the thesis.

1.1 Diagrams of hyperbolic links

Let K be a link in $S^3 = \mathbb{R}^3 \cup \{\infty\}$, projected onto a Euclidean plane. (In our usage, the term *link* refers to both knots and links of multiple components.) We can think of the *projection diagram* $D(K)$ as a 4-valent planar graph $G(K)$, with over-under crossing information at each vertex.

The theorems listed in the next two sections will work for all links whose diagrams are sufficiently complicated, where we estimate the complexity of $D(K)$ by counting *twist regions* and crossings.

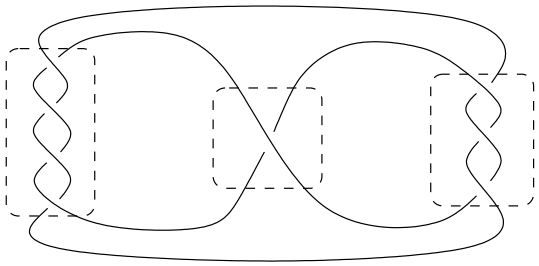

Figure 1.1: The above diagram has 3 twist regions, containing 2, $\frac{1}{2}$, and $1\frac{1}{2}$ twists, respectively.

Definition 1.1. A *bigon* is a contractible region in the complement of $G(K)$ that has two edges in its boundary. Following Lackenby [14, 16], we define a *twist region* of the link diagram to be a maximal string of bigons arranged end to end. A single crossing adjacent to no bigons is also a twist region.

We are also concerned with the amount of twisting that occurs in each twist region. We will count this either in terms of crossings or in terms of full twists, where a *full twist* of one strand about the other corresponds to two crossings. See Figure 1.1 for an illustration of these definitions.

Definition 1.2. A diagram $D(K)$ of a link $K \subset S^3$ is called *prime* if for any simple closed curve γ in the projection plane that intersects the graph $G(K)$ transversely in two points in the interior of edges, γ bounds a subdiagram containing no crossings of the original diagram. Note this ensures the diagram contains no monogons. See Figure 1.2.

Following Lackenby [16], we also require the diagram to be *twist-reduced*.

Definition 1.3. A link diagram $D(K)$ is *twist-reduced* if whenever a simple closed curve γ in the projection plane intersects the graph $G(K)$ transversely in four points in the interior of edges, with two points adjacent to one crossing and the other two points adjacent to another crossing, then γ bounds a subdiagram consisting of a (possibly empty) collection of bigons arranged in a row between these two crossings. See Figure 1.2.

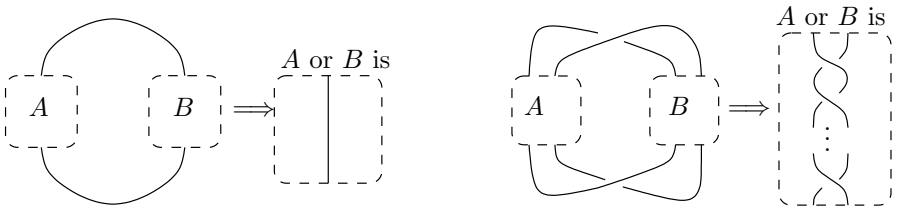

Figure 1.2: Left: A prime diagram; Right: A twist-reduced diagram.

Note that any diagram of a prime link K can be simplified into a prime, twist-reduced diagram: for if $D(K)$ is a diagram that fails to be prime, then all crossings on one side of a simple closed curve γ are extraneous and can be removed. Similarly, if $D(K)$ is not twist-reduced, then one can amalgamate the two twist regions adjacent to a curve γ into a single region, reducing the number of twist regions.

Following the work of Thurston [33], it is known that the complements of "most" knots and links admit a complete hyperbolic structure, that is, a complete Riemannian metric of constant sectional curvature -1. Specifically, Thurston proved that the only non-hyperbolic knots are *torus knots*, which follow along the surface of an unknotted torus, and *satellite knots*, which are contained in a tubular neighborhood of a simpler knot. A similar statement holds for links. Furthermore, the more complicated a link diagram is, the more likely it is that K is hyperbolic.

In practice, however, it is typically quite difficult to decide from a diagram $D(K)$ whether K is hyperbolic. One known special case is that of *alternating links*, which have projection diagrams where over- and under-crossings alternate. Menasco has used Thurston's work to prove that an alternating link K is hyperbolic if and only if a prime alternating diagram $D(K)$ has two or more twist regions [17].

The first main theorem of this thesis extends Menasco's result to all sufficiently complicated links.

Theorem 1.4. *Let $K \subset S^3$ be a link with a prime, twist-reduced diagram $D(K)$. If $D(K)$ has at least two twist regions and every twist region of $D(K)$ contains at least 6 crossings, then K is hyperbolic.*

Because there are torus and satellite links with arbitrarily many twist regions, some assumption on the number of crossings in each twist region is necessary. The hypothesis of 6 crossings per twist region is probably too strong, but 6 is a small enough number to make the theorem applicable in practice.

The methods used to prove Theorem 1.4 also provide a diagrammatic estimate on the genus of a hyperbolic link.

Definition 1.5. Let $K \subset S^3$ be a link. A *Seifert surface* for K is an orientable, incompressible surface whose boundary is K. We say that the *genus* of K is the smallest genus of any Seifert surface for K.

Theorem 1.6. *Let $K \subset S^3$ be a link of k components with a prime, twist-reduced diagram $D(K)$. If $D(K)$ has $t \geq 2$ twist regions and at least 6 crossings in each twist region, then*

$$\text{genus}(K) \geq \left\lceil 1 + \frac{t}{6} - \frac{k}{2} \right\rceil,$$

where $\lceil \cdot \rceil$ is the ceiling function that rounds up to the nearest integer.

Crowell [6] and Murasugi [21] have independently proved that the genus of an alternating link is equal to half the degree of its Alexander polynomial, and Gabai gave an algorithm to compute the genus of an arborescent link [10]. (See Section 1.4 for the definition of arborescent links.) The advantage of Theorem 1.6 is that it works for general links and, in fact, gives the exact value for certain families of links.

Theorems 1.4 and 1.6 are both proved using Dehn surgery techniques that have other interesting consequences.

1.2 Dehn surgery

Let M be a 3-manifold with torus boundary ∂M, for example a link complement. We can pick a *slope* s on ∂M – that is, an isotopy class of simple closed curves on ∂M. The manifold obtained by gluing a solid torus $S^1 \times D^2$ to ∂M in such a way that the slope s bounds a disk in the resulting manifold is called a *Dehn filling along the slope s*, or a *Dehn surgery along s*. More generally, if M is a 3-manifold with multiple

torus boundary components and along each component we have a slope s_i, we obtain a closed manifold by Dehn filling along these slopes.

Definition 1.7. For a link $K \subset S^3$, let the *exterior* $E(K)$ denote the complement of an open tubular neighborhood of K. On each boundary torus of $T \subset \partial E(K)$, we define a *meridian* μ to be the slope bounding a disk in S^3, and the *longitude* λ to be the slope of a curve that runs along a component of K and has trivial linking number with K. Then the pair $\langle \mu, \lambda \rangle$ gives a basis for the fundamental group of T.

Using the basis $\langle \mu, \lambda \rangle$ for the fundamental group of a torus T, slopes on T are parameterized by $\overline{\mathbb{Q}} = \mathbb{Q} \cup \{\infty\}$. Specifically, a slope corresponds to a/b if and only if the slope is equivalent to $a\mu + b\lambda$. For a knot exterior $E(K)$, Dehn filling along a meridian of K, i.e. $1/0$ filling, will always give S^3. This Dehn filling is called the *trivial filling*. All other Dehn fillings are *non-trivial*.

Dehn surgery on hyperbolic manifolds has been extensively studied in the last quarter-century. Thurston has shown that given a hyperbolic manifold M with torus boundary, all but finitely many choices of surgery slope on each component of ∂M yield a closed hyperbolic manifold [32]. More recently, Hodgson and Kerckhoff showed that if the surgery slope on each component of ∂M is longer than a given universal constant, then the resulting Dehn filled manifold is hyperbolic [13]. For hyperbolic manifolds whose boundary is a single torus, they prove that at most 60 slopes give non-hyperbolic manifolds, and a similar statement holds for multiple boundary components.

Using these results, Purcell was able to show that for all sufficiently complicated knots, every nontrivial Dehn filling is hyperbolic [28]. However, the required knots are quite complicated, needing over 140 crossings per twist region.

It turns out that we can rule out exceptional surgeries on much simpler links if we only require that the Dehn filled manifold satisfies the topological properties of hyperbolic manifolds.

Definition 1.8. A closed, orientable 3-manifold M is *hyperbolike* if

(1) M is irreducible and atoroidal,

(2) M is not Seifert fibered, and

(3) $\pi_1(M)$ is infinite and word-hyperbolic.

All hyperbolic manifolds are hyperbolike. Thurston's Geometrization Conjecture [33], whose proof was recently claimed by Perelman [24, 25], implies the converse.

Theorem 1.9. *Let K be a link in S^3 with a prime, twist-reduced diagram $D(K)$. Suppose that every twist region of $D(K)$ contains at least 6 crossings and each component of K passes through at least 7 twist regions (counted with multiplicity). Then*

(1) *any non-trivial Dehn filling of some but not all components of K is hyperbolic, and*

(2) *any non-trivial Dehn filling of all the components of K is hyperbolike.*

Jessica Purcell has independently proved Theorem 1.9 using geometric techniques, and the two proofs will appear in a joint paper [9].

Corollary 1.10. *Let K be a knot in S^3 with a prime, twist-reduced diagram $D(K)$. If $D(K)$ has at least 4 twist regions, and each twist region contains at least 6 crossings, then any non-trivial Dehn filling of K is hyperbolike.*

The corollary follows from Theorem 1.9 because if K is a knot, every twist region contains two strands of K. Thus in a diagram with 4 twist regions, K passes through a twist region 8 times.

In fact, an easy class of examples shows that the hypothesis of 4 twist regions in Corollary 1.10 is a sharp bound. Consider the family of *pretzel knots* with 3 twist regions, in which the twist regions all run in parallel. (See Figure 1.3.) When the numbers of crossings p, q, r in the three twist regions are all odd, the Seifert surface S is a punctured torus. (This can be seen by cutting S into two ideal triangles, glued together as in the figure.) Dehn filling along the boundary of S glues a disk into the puncture of the punctured torus, producing a torus that is incompressible by a theorem of Przytycki [27]. Thus the Dehn filled manifold is toroidal, so Corollary 1.10 assumes the smallest possible number of twist regions.

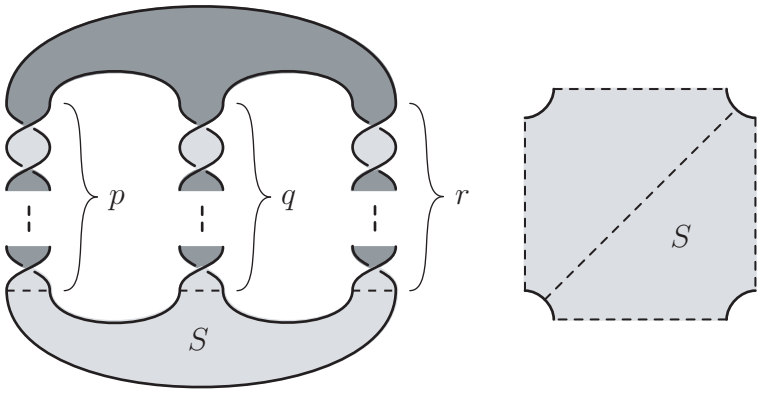

Figure 1.3: A pretzel knot with 3 twist regions. When p, q, r are odd, Dehn surgery along the boundary of the surface S produces an incompressible torus.

As for the requirement that each twist region contain at least 6 crossings, we know that some such requirement is necessary. It is known that there exist knots with non-trivial exceptional surgeries that have arbitrarily large volume, hence an arbitrarily high number of twist regions. These have been discovered by Eudave-Muñoz and Luecke [8], Eudave-Muñoz [7], as well as recently by Baker [2]. Thus a high number of twist regions alone is not enough to rule out exceptional surgeries. However, it seems difficult to tell if the requirement of six crossings per twist region is sharp.

Another advantage of Theorem 1.9 is that it gives information on Dehn fillings without requiring us to restrict our attention to a particular class of knots or links. This should be compared to other known results. If we restrict to alternating links, Lackenby has shown that all non-trivial Dehn surgeries on alternating knots with at least 9 twist regions are hyperbolike [14], as are surgeries on alternating links in which each component passes through 17 or more twist region. Wu proved that all non-trivial surgeries on a large class of arborescent knots are hyperbolic [35]. Theorem 1.9 applies to both of these classes of knots as well as non-alternating, non-arborescent knots and links.

Theorem 1.9 also gives a nice tool for understanding Dehn fillings on link complements as well as knot complements. Classifying Dehn fillings on links is often a more difficult problem than classifying fillings on knots, because of the way in which

the fillings of individual link components interact. For example, Wu proved that non-trivial surgeries on most arborescent links produce manifolds that are Haken and not small Seifert fibered [36], just as he did for knots [35], but he leaves open the question of whether these surgeries on links are atoroidal. On the other hand, our arguments apply equally well to both knots and links.

1.3 Angled polyhedra

The main tool used in the proofs of Theorems 1.4, 1.6, and 1.9 is the study of angled polyhedral decompositions. A 3-manifold with boundary $(M, \partial M)$ can always be subdivided into *ideal polyhedra*, whose vertices have been truncated and replaced by polygonal faces that lie on ∂M. In our context, we require that the polyhedra have dihedral angles corresponding to the angles of convex ideal polyhedra in the hyperbolic space $\mathbb{H}^3$. (By a theorem of Rivin [30], this amounts to checking a combinatorial condition on the dihedral angles.) We also require that the angles of neighboring polyhedra fit together in a consistent fashion, i.e., that the dihedral angles around each edge of M add up to 2π. If this is the case, we call the decomposition of M an *angled polyhedral decomposition*.

An angled decomposition of a manifold M does not necessarily give it a hyperbolic metric, because in gluing the polyhedra together we may have created singularities that amount to shearing along the edges. However, an angled polyhedral decomposition can be thought of as a piecewise hyperbolic structure on the manifold.

Angled polyhedral decompositions have many advantages. For starters, they are comparatively easy to obtain and deform. Because a hyperbolic structure on a finite-volume 3-manifold is unique by the Mostow-Prasad rigidity theorem [20, 26], any topological triangulation or polyhedral decomposition of a manifold corresponds to at most one geometric triangulation giving a hyperbolic metric. Finding the hyperbolic structure by gluing ideal polyhedra is also quite difficult, because it requires solving a system of non-linear, complex-valued equations.

On the other hand, an angle structure on a given ideal triangulation or polyhedral decomposition is the solution to a much more manageable system of linear equations and inequalities. Thus, if such a structure is found, it can be deformed to optimize some desired geometric property.

Angle structures also provide great control over the surfaces that can lie in a manifold. The work of Casson, expanded by Lackenby [14], defines a *combinatorial area* of surfaces intersecting angled polyhedra, and this notion of area has the same Gauss-Bonnet relationship with Euler characteristic as hyperbolic area. As a result, combinatorial area can be used to show that any 3-manifold with an angled polyhedral decomposition can, in fact, be given a hyperbolic structure (Theorem 2.8). Using combinatorial area, Lackenby has also found a way to define the *combinatorial length* $\ell(s)$ of a slope s on ∂M. (See Chapter 2 for the precise definition.)

Theorem 1.11 (Lackenby)**.** *Let M be a manifold with an angled polyhedral decomposition. Let $s_1, \ldots, s_n$ be a collection of slopes on ∂M, with one s_i on each component of ∂M. If $\ell(s_i) > 2\pi$ for each i, then*

(1) *Dehn filling M along the slopes $s_1, \ldots, s_n$ produces a hyperbolike manifold, and*

(2) *Dehn filling M along a proper subset of the s_i produces a hyperbolic manifold.*

Theorems 1.4, 1.6, and 1.9 are all proved using the same surgery description of the link K. In Chapter 3, we start with a prime, twist-reduced diagram $D(K)$ and use it to add a number of link components to K. The result is an *augmented link* J with the property that both K and its Dehn fillings can be obtained by surgery on J. The link exterior $E(J)$ subdivides into two particularly nice angled polyhedra. By studying the combinatorics of these polyhedra, we can estimate the length of surgery curves on $\partial E(J)$; applying Theorem 1.11 to these estimates gives Theorems 1.4 and 1.9. The combinatorial length of boundary slopes also bounds the Euler characteristic of surfaces in $E(K)$, yielding the genus estimate of Theorem 1.6.

For a large family of knots and links, called *arborescent links*, we can construct an angled decomposition for the original link, without any augmentation, obtaining more detailed information about the geometry of the link complement.

1.4 Arborescent links

Arborescent links were first defined by Conway [5] (who called them *algebraic* links) and studied in great detail by Montesinos [19], Gabai [10], and Wu [35, 36], among others. They are typically defined in terms of *tangles*.

Definition 1.12. A *tangle* is a pair (B, t), where B is a 3-ball and t is a compact 1-manifold, properly embedded in B, whose boundary consists of four points on ∂B. A *marked tangle* is a triple (B, t, D), where B and t are as above and $D \subset \partial B$ is a disk that contains exactly two endpoints of t. See Figure 1.4 for examples.

Definition 1.13. A *rational tangle* is a marked tangle (B, t, D), in which t consists of arcs γ_1 and γ_2 that are simultaneously parallel to ∂B.

A rational tangle can be visualized in a very concrete way by taking B to be a *solid pillowcase* (the 3-dimensional neighborhood of a unit square) and D to be the right or left half of ∂B. The arcs γ_1 and γ_2 start at the corners of the square and travel just inside ∂B with a fixed slope $s \in \mathbb{Q} \cup \{\infty\}$. Conway has proved that any rational tangle can be constructed in this way, and conversely, the slope s determines a rational tangle $R(s)$, uniquely up to isotopy fixing the four corners [5]. Furthermore, it follows from his work that every rational tangle (B, t, D), positioned in the solid pillowcase as above, admits an isotopy fixing ∂B that makes the diagram of t alternating. In effect, the isotopy implements the Euclidean algorithm to find a continued fraction expansion of the slope s whose integers all have the same sign as s. See Figure 1.5.

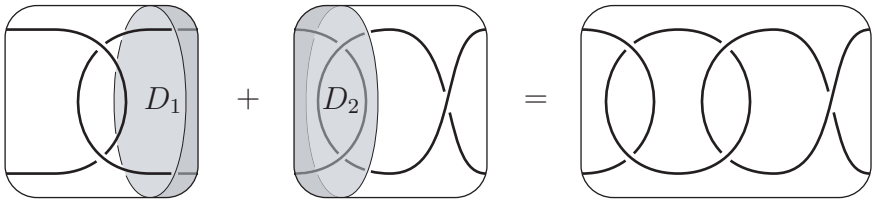

Figure 1.4: Two marked tangles, B_1 and B_2, can be glued along their marking disks D_1 and D_2 to form a *sum tangle* $B_1 + B_2$. In this figure, B_1 and B_2 are rational but their sum is not.

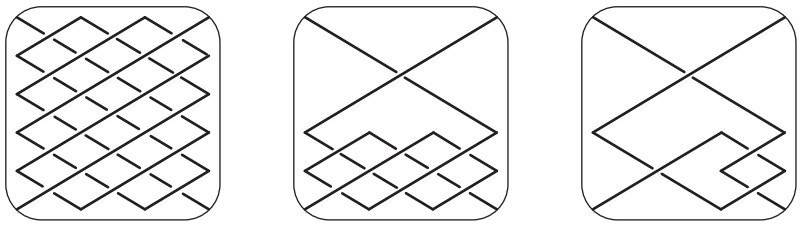

$$\frac{3}{5} \qquad = \frac{1}{1+\frac{2}{3}} \qquad = \frac{1}{1+\frac{1}{1+\frac{1}{2}}}$$

Figure 1.5: A rational tangle of slope 3/5, progressively simplified into an alternating projection.

Definition 1.14. Let B_1 and B_2 be two marked tangles. We can join them by gluing the marking disk D_1 to D_2. This forms the *sum tangle* $B_1 + B_2$, as in Figure 1.4. The sum is called *non-trivial* if neither B_1 nor B_2 is a rational tangle of slope 0 or ∞.

An *arborescent tangle* A is defined by the property that rational tangles are arborescent tangles, and any non-trivial sum of arborescent tangles is another arborescent tangle.

When we begin constructing an arborescent tangle by gluing two rational tangles B_1 and B_2, we can always position them as in Figure 1.4, with the marking disks in plain view. Later on, we may want to glue the tangle $A = B_1 + B_2$ to another marked tangle A', along a marking disk $D \subset \partial A$ that appears quite complicated. However, we can always place D in a simple position while preserving the diagram of B_1 and B_2, by isotoping ∂A outward along a braid. See Figure 1.6.

Definition 1.15. An *arborescent link* $K \subset S^3$ is obtained by gluing two arborescent tangles A and A' by some map identifying their entire boundaries. Again, the gluing map can be visualized concretely by connecting A to A' along a braid.

In short, arborescent links are obtained by combining rational tangles along braids. Examples of arborescent links include *two-bridge links*, constructed from gluing only two rational tangles, and pretzel links like the one in Figure 1.3. (The knot in Figure 1.3 is constructed from rational tangles $R(\frac{1}{p})$, $R(\frac{1}{q})$, and $R(\frac{1}{r})$.) Another special case

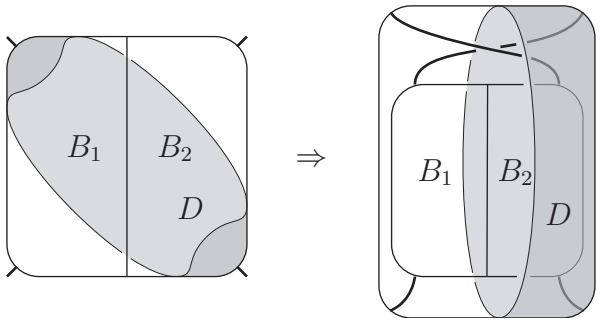

Figure 1.6: To place marking disk D in a simple position, we isotope the boundary of tangle $B_1 + B_2$ along a braid.

is the family of *Montesinos links*, which contains both two-bridge and pretzel links. Montesinos links are obtained by gluing several rational tangles cyclically along a band, as in Figure 1.7.

At any stage in the construction of an arborescent link, we also have some number of tangles $A_1, \ldots, A_n$ connected along a band. However, apart from Montesinos links, the tangles A_i could each consist of many rational tangles.

In Figure 1.7, the two crossings in the band can be inserted into any of the four rational tangles, changing its slope by 2. (Note, following Figure 1.5, that adding an integer to the slope twists the right or left pair of strands, which is equivalent to placing twists into a horizontal band.) Thus, when several rational tangles are connected along a band, their individual slopes are only fixed modulo $\mathbb{Z}$. To help avoid this ambiguity, we will choose slopes s with $|s| < 1$.

Definition 1.16. Let $R(s)$ be a rational tangle of slope s, with $|s| < 1$. We define the *length* $\ell(R)$ to be the number of crossings in the reduced alternating diagram of R. Equivalently, $\ell(R)$ is the sum of the absolute values of the integers in the shortest continued fraction expansion of s.

Observe that, for $s \in (0, 1)$, the continued fraction expansions of s and $s - 1$ have the same length. Thus $\ell(R(s)) = \ell(R(s-1))$, and we have an unambiguous notion of length for the rational tangles used to build an arborescent link.

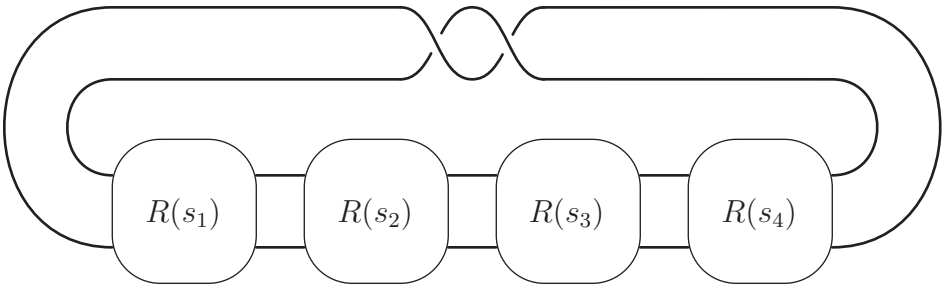

Figure 1.7: A Montesinos link is obtained by gluing several rational tangles cyclically along a band.

Definition 1.17. Suppose that, in an arborescent link K, arborescent tangles $A_1, \ldots, A_k$ are joined together cyclically along a band. For each i, we define ℓ_i to be the length $\ell(A_i)$ if A_i is a rational tangle, or ∞ otherwise. If

$$\frac{1}{l_{i-1}} + \frac{1}{l_i} + \frac{1}{l_{i+1}} \ < \ 1$$

for every triple of consecutive tangles around every band in K, we say the link K is *balanced*.

In fact, we will work with a definition of *balanced* that is more inclusive than the one above; see Definition 4.17.

Theorem 1.18. *Let K be a balanced arborescent link. Then there exists a decomposition of $S^3 \setminus K$ into angled hyperbolic polyhedra.*

Because angled polyhedral decompositions necessarily imply that a manifold admits a hyperbolic structure, we have

Corollary 1.19. *Every balanced arborescent link K is hyperbolic.*

The statement of Theorem 1.18 is sharp, in the following sense. For any triple of positive integers n_1, n_2, n_3 whose reciprocals sum to at least 1, there a Montesinos link K composed of three rational tangles R_1, R_2, R_3, such that $\ell(R_i) = n_i$ and K is not hyperbolic. For example, the pretzel knot $K(\frac{-1}{2}, \frac{1}{3}, \frac{1}{5})$ is actually a torus knot, and

the pretzel link $K(\frac{-1}{2}, \frac{1}{3}, \frac{1}{6})$ contains an incompressible torus in its complement. (By Menasco's theorem [17], non-hyperbolic examples of this sort cannot be alternating.)

Corollary 1.19 mostly replicates known results about hyperbolic arborescent links. Oertel has compiled a complete list of non-hyperbolic Montesinos links [23], and Wu's work on Dehn surgery indirectly proves that all non-Montesinos arborescent knots are hyperbolic [35]. In the remaining case of non-Montesinos links with multiple components, Corollary 1.19 may have some new content.

In any case, the primary interest of Theorem 1.18 involves the angled polyhedra themselves, because angled decompositions provide a way to glean concrete information about the hyperbolic structure of a manifold.

1.5 Organization of thesis

This thesis is organized as follows. Chapter 2 surveys the detailed definitions of angled polyhedra and the theory of normal surfaces that is used to prove results about combinatorial length and Dehn surgery. Chapter 3 then applies this theory to an augmented link constructed from a diagram $D(K)$, using combinatorial length to prove Theorems 1.4, 1.6, and 1.9.

The remaining chapters are devoted to arborescent links. Chapter 4 describes the correspondence between the diagram of an arborescent link and a *weighted tree* T, using the tree to simplify the diagram into a particularly nice form. Chapter 5 uses the simplified diagram and its correspondence to the tree to subdivide the link complement into ideal polyhedra. Chapter 6 studies the dihedral angles of these polyhedra, setting up the gluing equations sufficient for an angle structure and proving that for a balanced arborescent link K, the equations always have solutions.

Chapter 2

A Survey of Angled Triangulations

The goals of this chapter are to set up the definitions needed to study angled triangulations and to survey the theorems that make this point of view powerful. These theorems fall into two categories: combinatorial results using normal surfaces and analytic results using the technique of volume maximization.

This chapter is primarily a survey. None of the results listed here are original; they have either appeared in published work or are minor extensions of known results. The main ideas were developed by Marc Lackenby [14, 15] and Igor Rivin [29, 30].

2.1 Angled polyhedra

Definition 2.1. For the purposes of this thesis, a *polyhedron* is a 3-ball $\hat{P}$ with a specified graph Γ embedded in $\partial\hat{P}$, such that

(1) each vertex of Γ has valence at least 3,

(2) each edge of Γ has ends on distinct vertices, and

(3) each region of $\partial\hat{P} \setminus \Gamma$ is bounded by at least 3 edges.

$\hat{P}$ inherits vertices and edges from Γ, and the *faces* of $\hat{P}$ are regions of $\partial\hat{P} \setminus \Gamma$.

Definition 2.2. An *ideal polyhedron* P^{O} is a polyhedron $\hat{P}$ with the vertices removed. A *truncated ideal polyhedron* P is a polyhedron $\hat{P}$ with a closed regular neighborhood

of the vertices removed. Unless noted otherwise, the polyhedra mentioned in the sequel will be truncated ideal polyhedra.

A truncated ideal polyhedron P has two kinds of faces: *interior faces* that are truncated copies of the original faces of $\hat{P}$, and *boundary faces* that come from the truncated vertices of $\hat{P}$. P also has two kinds of edges: *interior edges* that come from the original truncated edges of $\hat{P}$, and *boundary edges* along the boundary faces. Note that we can recover P^{O} from P by removing the boundary faces.

The main objects of study in this thesis are three-dimensional manifolds with boundary (specifically, link complements) subdivided into truncated ideal polyhedra. These polyhedra are joined along their interior faces to form the manifold M, while the boundary faces fit together to tile ∂M. Our actual goal is to investigate the geometry of the open manifold $M \setminus \partial M$, whose corresponding subdivision does not include the boundary faces. However, it turns out that the asymptotic geometry of the truncated polyhedra – the angles of boundary faces, which correspond to dihedral angles on interior edges – carries a lot of information about the geometry of the interior.

Definition 2.3. Let P be a truncated ideal polyhedron, and let P^{O} be the ideal polyhedron obtained by removing the boundary faces of P. To each edge of P^{O} (and thus each interior edge of P), we assign an *internal angle* α_i and an *external angle* $\epsilon_i = \pi - \alpha_i$. We say that P or P^{O} is an *angled polyhedron* if

(1) $\alpha_i \in (0, \pi)$ (meaning the polyhedron is convex),

(2) around every ideal vertex of P^{O}, $\sum \epsilon_i = 2\pi$ (meaning each boundary face of P is a convex Euclidean polygon), and

(3) for any closed curve $\gamma \subset \partial P^{\mathrm{O}}$ that intersects each edge transversely at most once, does not lie wholly in a face, and does not merely encircle an ideal vertex, $\sum_\gamma \epsilon_i > 2\pi$.

Theorem 2.4 (Rivin [30])**.** *Every angled polyhedron P^{O} can be realized as a convex ideal polyhedron in $\mathbb{H}^3$ with the prescribed dihedral angles, uniquely up to isometry. Conversely, every convex ideal polyhedron in $\mathbb{H}^3$ is an angled polyhedron.*

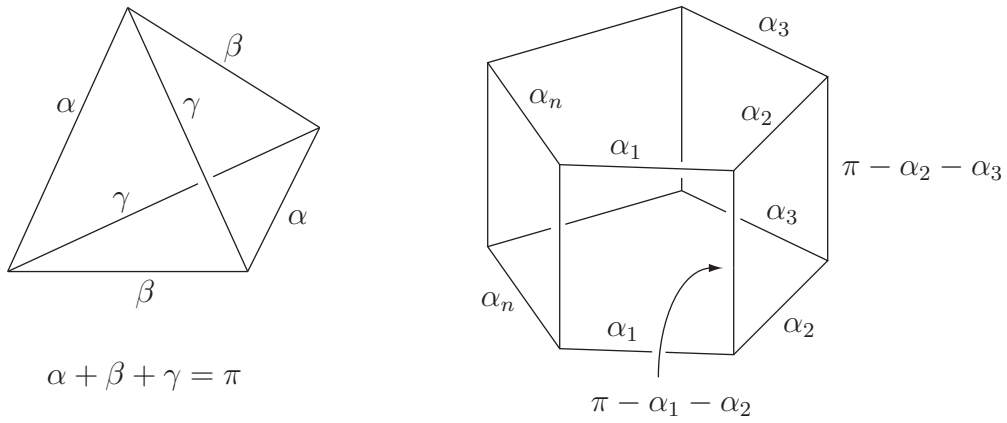

Figure 2.1: Left: in an ideal tetrahedron, opposite dihedral angles are equal. Right: dihedral angles in an ideal prism.

Although the conditions of Definition 2.3 seem complicated to check, in many cases they can be greatly simplified.

Example 2.5. The simplest type of an ideal polyhedron is a tetrahedron. Condition (2) says that the internal angles of the three edges that meet at each ideal vertex must add up to π. Since this holds at *each* ideal vertex, it follows that opposite edges must have equal dihedral angles. (See Figure 2.1(a).) It is easy to check that with such a choice of angles, the strict inequality of condition (3) is satisfied automatically.

Example 2.6. Let T be a polygon with $n \geq 3$ sides. An *ideal prism* is obtained by removing the vertices of $T \times [0, 1]$. We will call the polygons $T \times 1$ and $T \times 0$ the top and bottom faces, respectively. The prism will have n top edges, n bottom edges, and n *lateral* edges connecting the top face to the bottom face. (See Figure 2.1(b).)

Choose internal angles $\alpha_1, \ldots, \alpha_n$ for the top edges, and give the corresponding bottom edges the same labels. So long as $\alpha_i + \alpha_{i+1} < \pi$, we can choose internal angles $\theta_{i,i+1} = \pi - \alpha_i - \alpha_{i+1}$ for the lateral edges. This assignment of angles already satisfies conditions (1) and (2) of Definition 2.3. We will prove in Lemma 5.12 that condition (3) is satisfied whenever $\sum_{i=1}^{n} \alpha_i > \pi$ and $\alpha_i + \alpha_j < \pi$ for all $i \neq j$.

When we glue angled polyhedra together to form a manifold, we want to make sure that dihedral angles come together in a geometrically consistent way.

Definition 2.7. Let $(M, \partial M)$ be a compact 3-manifold with boundary. An *angled polyhedral decomposition* is a subdivision of M into (truncated) angled polyhedra, glued along their interior faces, such that $\sum \alpha_i = 2\pi$ around each interior edge. By removing the boundary faces, we also get a subdivision of $M \setminus \partial M$ into angled ideal polyhedra. If all the polyhedra are actually tetrahedra, the decomposition is called an *angled triangulation*.

Theorem 2.4 says that an angled decomposition constructs a manifold M out of hyperbolic pieces. In fact, Definition 2.7 is already restrictive enough to imply that M is hyperbolic, by a theorem of Lackenby. (Compare [14, Corollary 4.6] and the remark following [15, Lemma 4].)

Theorem 2.8 (Lackenby)**.** *Let M be an orientable 3-manifold with an angled polyhedral decomposition. Then ∂M is composed of tori, and the interior of M admits a complete, finite-volume hyperbolic structure.*

Its proof requires the techniques of normal surfaces.

2.2 Normal surfaces

The work of William Thurston in the late 1970s revealed deep connections between the geometry of 3-manifolds and the topology of surfaces in the manifold [33]. In particular, he proved the following important case of the Geometrization Conjecture:

Theorem 2.9 (Thurston's hyperbolization theorem)**.** *Let M be a compact, orientable 3-manifold with boundary consisting of tori. The interior of M admits a complete, finite-volume hyperbolic structure if and only if M contains no essential spheres, disks, tori, or annuli.*

When the manifold we are studying carries some additional structure, such as an angled polyhedral decomposition, it becomes natural to ask how surfaces in the manifold interact with that structure. The theory of normal surfaces, originally developed by Haken [12] and generalized and expanded in many directions, says that an embedded essential surface $(F, \partial F) \subset (M, \partial M)$ can be isotoped until its intersections with the polyhedra have a particularly nice, *normal* form.

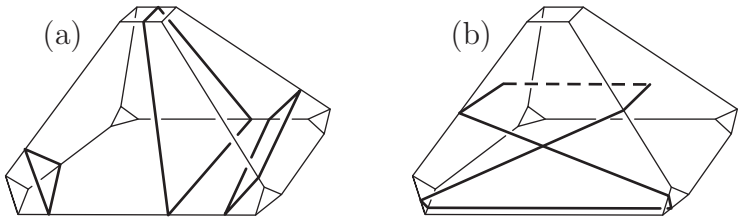

Figure 2.2: (a) Normal disks in a truncated polyhedron. (b) An admissible disk.

Definition 2.10. Let P be a truncated ideal polyhedron. A simple closed curve $\gamma \subset \partial P$ is called *normal* if

(1) γ is transverse to the edges of P,

(2) no arc of γ in a face of P has endpoints on the same edge, or on an interior edge and an adjacent boundary edge,

(3) γ doesn't lie entirely in a face of P,

(4) γ intersects each edge at most once, and

(5) γ intersects each boundary face at most once.

A disk in P bounded by a normal curve γ is called a *normal disk*. See Figure 2.2(a) for several examples.

Notation. To avoid confusion with longer arcs on ∂M, we will refer to the intersection between a curve γ and a face of P as a *segment*. (Thus an arc can consist of many segments.) Segments of γ lying in interior faces of P will be called *interior segments*, and the segments lying in boundary faces will be called *boundary segments*.

With a eye to Dehn surgery applications, Lackenby has extended the theory of normal surfaces to a more general class of surfaces that cannot be normalized [14]. These surfaces may not be embedded, and may even have boundary components in the interior of M.

Definition 2.11. Let P be a truncated ideal polyhedron. An immersed disk $D \subset P$ is called *admissible* if

(1) $\partial D \setminus \partial P$ is a (possibly empty) collection of disjointly embedded arcs with endpoints inside interior faces of P,

(2) $\partial D \cap \partial P$ is an immersed closed curve or an immersed collection of arcs,

(3) each segment of ∂D in a face of P is embedded,

(4) if $\partial D \cap \partial P$ is a closed curve, it satisfies conditions $(1) - (3)$ of Definition 2.10 of a normal curve, and

(5) each arc component of $\partial D \cap \partial P$ satisfies conditions $(1) - (2)$ of Definition 2.10.

An example is shown in Figure 2.2(b). We call an immersed surface $F \subset M$ an *admissible surface* if it intersects each polyhedron in a collection of admissible disks.

Angle structures on a polyhedral decomposition of M allow us to measure the complexity of surfaces using *combinatorial area*.

Definition 2.12. Let $D \subset P$ be an admissible disk in an angled polyhedron, with the boundary faces of P lying on ∂M. Let $E_1, \ldots, E_n$ be the interior edges crossed by ∂D (counted with multiplicity), and let $\epsilon_1, \ldots, \epsilon_n$ be the corresponding external angles. Then define the combinatorial area of D to be

$$a(D) = \sum_{i=1}^{n} \epsilon_i + \pi |\partial D \cap \partial M| - 2\pi + 3\pi |\partial D \setminus \partial P| \,.$$

For an admissible surface $F \subset M$, $a(F)$ is defined by summing the areas of its admissible disks.

For disks with $\partial D \subset \partial P$, this definition matches the formula for hyperbolic area. As a comparison, a polygon $T \subset \mathbb{H}^2$ with external angles ϵ_i has area $a(T) = \sum \epsilon_i - 2\pi$. (See, for example, Corollary 2.4.15 of [34].) Ideal vertices have internal angle 0 and thus add π to the area, just as each component of $\partial D \cap \partial M$ adds π to combinatorial area. As for the coefficient 3π per component of $\partial D \setminus \partial P$, it was chosen by Lackenby to make the combinatorial area of D automatically positive whenever ∂D passes through the interior of P.

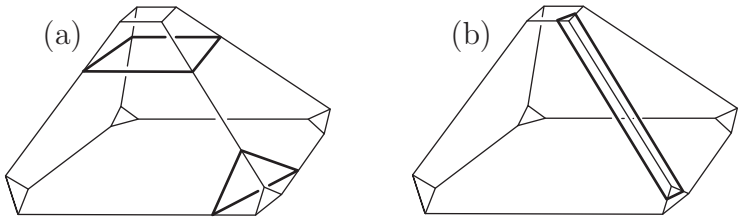

Figure 2.3: (a) Vertex links. (b) A boundary bigon.

In fact, there are only two types of admissible disks whose area is 0; both of them happen to be normal. The first is a *vertex link* cutting off a boundary face; its area is 0 by Condition (2) of Definition 2.3. The second is a *boundary bigon* cutting off an interior edge; it has area 0 because its boundary curve only picks up area from two boundary faces. They are shown in Figure 2.3.

Lemma 2.13. *Let $D \subset P$ be an admissible disk in an angled polyhedron. If D is not a vertex link or a boundary bigon, $a(D) > 0$.*

Proof. If D is normal, Lackenby [15, Lemma 4] proves that $a(D) > 0$ unless D is a vertex link or a bigon. If D is not normal, Lackenby [14, Lemma 4.2] proves that $a(D) > 0$. In both cases, the proofs rely on condition (2) of Definition 2.3 of an angled polyhedron and the observation that if ∂D self-intersects or crosses an edge multiple times, the area can actually be reduced by surgering the disk. □

Remark 2.14. In his study of admissible surfaces, Lackenby [14] worked with angled spines that are dual to the polyhedra. Our Definition 2.1 of a polyhedron is actually slightly more restrictive than his dual definition of a thickened spine, in that our polyhedra are not allowed to have bigon faces. This results in some simpler statements. Bigon faces of the (non-truncated) polyhedra would allow normal disks other than boundary bigons or vertex links to have zero area [14, Lemma 4.2]; in our scenario, every other admissible disk has strictly positive area.

The analogy between hyperbolic area and combinatorial area extends to the following combinatorial version of the Gauss-Bonnet Theorem [14, Proposition 4.3].

Proposition 2.15. *Let $F \subset M$ be an admissible surface in a manifold with an angled polyhedral decomposition. Let* $\text{Length}(\partial F \setminus \partial M)$ *be the number of arcs of intersection between $\partial F \setminus \partial M$ and the polyhedra. Then*

$$a(F) = -2\pi\chi(F) + 2\pi\text{Length}(\partial F \setminus \partial M)\,.$$

This development allows us to prove that only hyperbolic manifolds have angled polyhedral decompositions.

Proof of Theorem 2.8. Each component of ∂M is tiled by boundary faces of the polyhedra. Just inside each boundary face, a polyhedron has a normal disk of area 0. These vertex links glue up to form a closed, boundary-parallel normal surface F of area 0. By Proposition 2.15, $\chi(F) = 0$, and since M is orientable, F must be a torus. Thus ∂M consists of tori.

By Thurston's hyperbolization theorem, the interior of M is hyperbolic if and only if M contains no essential spheres, disks, tori, or annuli. In our situation, any such essential surface can be isotoped into normal form. An essential sphere or disk has positive Euler characteristic, hence negative area. Thus it cannot occur.

A normal torus $T \subset M$ has area 0 and thus, by Lemma 2.13, must be composed of normal disks of area 0. Since T has no boundary, these must all be vertex links, which glue up to form a boundary-parallel torus. Similarly, a normal annulus $A \subset M$ must be composed entirely of bigons, since a bigon cannot be glued to a vertex link. But a chain of bigons forms a tube around an edge of M, which is certainly not essential. Thus we can conclude that M is hyperbolic. □

2.3 Combinatorial length and Dehn surgery

Lackenby's crucial insight [14] is that one can use the combinatorial area of surfaces in a manifold M to define a combinatorial length of curves on ∂M, and that the length of surgery curves gives information about Dehn fillings of M.

Definition 2.16. Let P be an angled polyhedron, and let $D \subset P$ be an admissible disk that intersects at least one boundary face. Let γ be a segment of ∂D in a

boundary face of P. Then we define the *length of γ relative to D* to be

$$\ell(\gamma, D) = \frac{a(D)}{|\partial D \cap \partial M|}.$$

Definition 2.17. For a manifold M with an angled polyhedral decomposition, let γ be a (possibly non-closed) immersed arc in ∂M. We call γ a *simplicial arc* if

(1) γ is disjoint from the vertices of ∂M,

(2) the endpoints of γ (if any) lie on edges of ∂M,

(3) each segment of γ in a boundary face is embedded, and

(4) no segment of γ in a boundary face has endpoints on the same edge.

We can now define the combinatorial length of simplicial arcs on ∂M by considering all the possible *inward extensions* of the arc.

Definition 2.18. Let $\gamma \subset \partial M$ be a simplicial arc. Let $\gamma_1, \ldots, \gamma_n$ be the boundary segments that make up γ, ordered along a parametrization of γ. For each i, let D_i be an admissible disk in the corresponding polyhedron, whose boundary contains γ_i. Then $H = \cup_{i=1}^{n} D_i$ is called an *inward extension* of γ if

(1) ∂D_i agrees with ∂D_{i+1} on the shared face of their polyhedra, and

(2) if γ is closed, ∂D_n agrees with ∂D_1 on the common face.

We define the *combinatorial length* of γ to be

$$\ell(\gamma) = \inf \left\{ \sum_{i=1}^{n} \ell(\gamma_i, D_i) \right\},$$

where the infimum is taken over all inward extensions of γ.

Definition 2.19. Let s be a slope on a boundary component of M. Then define the combinatorial length of s to be

$$\ell(s) = \inf \{\ell(\gamma)\},$$

the infimum being taken over all closed simplicial curves $\gamma \subset \partial M$ that represent non-zero multiples of slope s.

This string of definitions gives us the following lemma [14, Proposition 4.8].

Lemma 2.20. *Let M be a manifold with an angled polyhedral decomposition, and let $F \subset M$ be an admissible surface. Let $C_1, \ldots, C_m$ be the components of $\partial F \cap \partial M$, each C_j representing a non-zero multiple of some slope $s_{i(j)}$. Then*

$$a(F) \geq \sum_{j=1}^{m} \ell(s_{i(j)}) \,.$$

Proof. The admissible disks of F bordering on each C_j form one inward extension of C_j. Definition 2.16 has us divide the area of each disk by the number of its intersections with ∂M, so we do not end up double-counting any area. □

As a consequence of Lemma 2.20, surfaces with long boundary have large combinatorial area, hence large genus. This yields the following powerful Dehn surgery result [14, Theorem 4.9]. (See Definition 1.8 for the meaning of *hyperbolike*.)

Theorem 2.21 (Lackenby)**.** *Let M be a manifold with an angled polyhedral decomposition. Let $s_1, \ldots, s_n$ be a collection of slopes on ∂M, with one s_i on each component of ∂M. If $\ell(s_i) > 2\pi$ for each i, then the manifold obtained by Dehn filling M along the slopes $s_1, \ldots, s_n$ is hyperbolike.*

The machinery that Lackenby has developed also allows for an extension of his theorem to surgeries along only some components of ∂M.

Theorem 2.22. *Let M be a manifold with an angled polyhedral decomposition. Let $s_1, \ldots, s_m$ be a collection of slopes on some, but not all, of the boundary tori. If $\ell(s_i) > 2\pi$ for each i, then the manifold obtained by Dehn filling M along the slopes $s_1, \ldots, s_m$ is hyperbolic.*

Proof. By Thurston's Theorem 2.9, proving that the Dehn filled manifold is hyperbolic amounts to ruling out essential spheres, disks, tori, and annuli. Any such

surface F must intersect at least one of the solid tori added during the surgery process, because M is hyperbolic by Theorem 2.8. Thus F contains a punctured surface $G \subset M$, whose punctures (not counting the original boundary components of F) represent surgery slopes $s_{i(1)}, \ldots, s_{i(k)}$ of length greater than 2π. We can place G in normal form in the angled polyhedra and compute its combinatorial area. Then

$$\begin{aligned} a(G) &= -2\pi\chi(G), && \text{by Proposition 2.15} \\ &\leq 2\pi|\partial G \setminus \partial F|, && \text{given the choices of } F \\ &< \sum_{j=1}^{|\partial G \setminus \partial F|} \ell(s_{i(j)}), && \text{by assumption} \\ &\leq a(G), && \text{by Lemma 2.20,} \end{aligned}$$

obtaining a contradiction. □

2.4 Angle structures and volume maximization

We have seen in Section 2.2 that Thurston's hyperbolization theorem, together with the machinery of normal surfaces, implies that only hyperbolic manifolds can be decomposed into angled polyhedra. In this section, we describe a way to use angled triangulations to find the hyperbolic structure on a manifold in a much more concrete fashion, without resorting to Thurston's deep result.

Let M be an open manifold, the interior of a 3-manifold with torus boundary. The goal is to obtain M as a union of ideal tetrahedra, glued together in a way that extends the hyperbolic structure over all of M. Recall from Example 2.5 that the shapes of ideal hyperbolic tetrahedra are parametrized by a choice of three dihedral angles that add up to π. The faces of these tetrahedra – ideal triangles – are all isometric. Thus, given a topological triangulation of M into angled tetrahedra, there is a unique way to glue their geometric structures by isometries on the faces to give a hyperbolic structure on the complement of the edges.

In order to extend this hyperbolic structure over all of M, three conditions need to be satisfied:

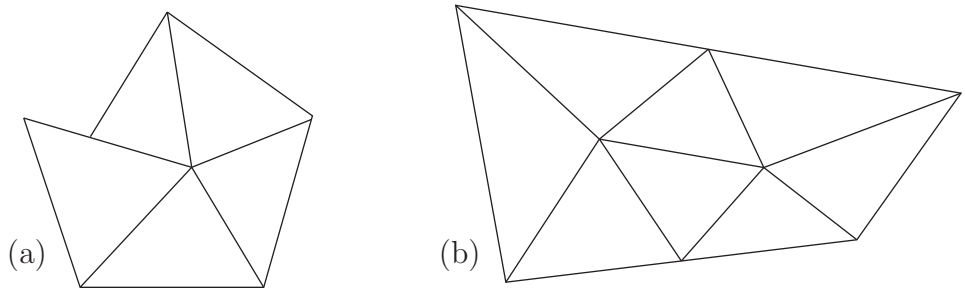

Figure 2.4: (a) A shearing singularity around an edge. (b) The fundamental domain at an incomplete cusp. The developing image of this torus would spiral in to a point.

(1) the dihedral angles at each edge must add up to 2π,

(2) there must not be any translational singularities (*shearing*) at the edges, and

(3) the resulting metric on M must be complete.

Condition (1) is familiar, and says that we have an angled triangulation. The other two conditions can be understood in terms of the geometry of the boundary triangles obtained when we truncate the tetrahedra. Condition (2) says that the side lengths of the triangles line up as we go around each edge of M. Condition (3) says that the boundary triangles at each cusp of M fit together to give a horospherical torus, whose developing map covers a horosphere in $\mathbb{H}^3$. See Figure 2.4 for examples of how these conditions can be violated.

Definition 2.23. Consider a topological ideal triangulation τ of a manifold M. An *angle structure* on τ is a choice of dihedral angles for the tetrahedra that makes τ into an angled triangulation. The set of all angle structures on τ is denoted $\mathcal{A}(\tau)$.

$\mathcal{A}(\tau)$ is the solution set of a system of linear equations (the three angles defining a tetrahedron sum to π; the angles around an edge sum to 2π) and linear inequalities (all angles are positive). Thus $\mathcal{A}(\tau)$ is a convex polytope, whose closure $\overline{\mathcal{A}(\tau)}$ is compact.

In the early 1990s, Rivin observed [29] that conditions (2) and (3) are closely related to critical points of the *volume function* $V : \overline{\mathcal{A}(\tau)} \to \mathbb{R}$. Every ideal hyperbolic

tetrahedron T with dihedral angles α, β, γ has an easily computable volume

$$\mathrm{Vol}(T_{\alpha\beta\gamma}) = \text{Л}(\alpha) + \text{Л}(\beta) + \text{Л}(\gamma)\,,$$

where

$$\text{Л}(x) = -\int_0^x \log|2\sin t|\,dt$$

is the Lobachevsky function. (See [18].) We can then associate a volume $V(p)$ to any angle structure $p \in \overline{\mathcal{A}(\tau)}$, by adding up the volumes of all the angled tetrahedra. Rivin proved that at any critical point of V, shears around the edges of M vanish [29], satisfying condition (2).

This can be combined with combinatorial properties of the triangulation due to Neumann and Zagier [22] to prove the following theorem [4].

Theorem 2.24 (Rivin, Chan–Hodgson)**.** *Suppose that V has a critical point at $p \in \mathcal{A}(\tau)$. Then the ideal tetrahedra whose shapes are described by p glue together to give a complete hyperbolic structure on M.*

It turns out that the volume function V is concave down; as a result, any critical point will be a maximum, and will be unique. V will always attain a maximum on the compact polytope $\overline{\mathcal{A}(\tau)}$, but this maximum could lie on the boundary. To find the hyperbolic structure on M using angled triangulations, it suffices to prove that the volume function is maximized in the interior.

In practice, the task of proving that the maximum of volume occurs in the interior of $\overline{\mathcal{A}(\tau)}$ seems challenging but doable. For example, François Guéritaud has carried out this program for all hyperbolic punctured torus bundles over the circle [11]. His methods translate directly to angled triangulations of two-bridge knot complements, and can be used to attack more complicated knots. The volume maximization process does indeed find geometric triangulations for at least a few special families of arborescent knots; the proof of this fact will be written up in future work.

Despite the promise of this approach, there are situations where the polytope of angle structures coming from some ideal triangulation τ attains the maximum volume on the boundary, and many more situations where the location of the maximum is

difficult to ascertain. In this case, Theorem 2.8 still tells us that the underlying manifold M is hyperbolic. In this case, it seems natural to ask how its hyperbolic volume compares to the volume associated to an angle structure.

Conjecture 2.25. *Let $p \in \mathcal{A}(\tau)$ be an angled triangulation of a manifold M. If p does not give the complete hyperbolic structure on M, then*

$$V(p) < \mathrm{Vol}(M)\,.$$

Chapter 3

Augmented Links and Dehn Surgery

In this chapter, we will use the tools of angled polyhedra and normal surfaces to prove Theorems 1.4 and 1.9. The proof strategy is slightly counterintuitive, in that we get information about Dehn fillings of a link K by first adding a number of extra components to the link. In a construction pioneered by Colin Adams [1] and expanded by Jessica Purcell [28], we will add *crossing circles* around the twist regions of $D(K)$ and remove crossings to obtain an augmented link J. Both K and its Dehn fillings can then be represented as Dehn fillings of J.

The payoff of this approach is that the link exterior $E(K)$ has a natural decomposition into angled polyhedra P_1 and P_2, whose combinatorics is fairly easy to understand. We will estimate the combinatorial length of surgery slopes on the tori of $\partial E(J)$, giving a way to rule out exceptional surgeries on both J and the original link K.

3.1 Constructing the augmented link

Let $D(K)$ be a prime, twist-reduced diagram of a link $K \subset S^3$. As described in the introduction, each twist region in $D(K)$ consists of two strands of K wrapping around each other. For each twist region R_i, add a simple closed curve C_i encircling

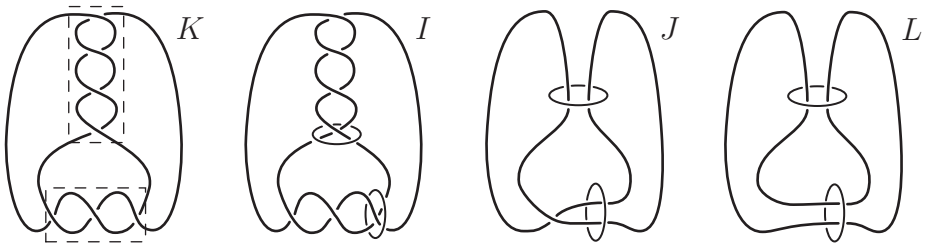

Figure 3.1: Left to right: The original knot with twist regions marked; the link I with crossing circles added; the homeomorphic link J; the flat augmented link L.

the twist region, known as a *crossing circle*. Let I be the resulting link.

Note that the manifold $E(I)$ is homeomorphic to the manifold $E(J)$, where J is the simpler link with all full twists removed at each twist region of I. We can recover the original link K from J by performing $1/n_i$ surgery on each C_i, $|n_i|$ being the number of full twists we removed. Furthermore, any Dehn filling of K can be viewed as a filling of J. We will spend the bulk of this chapter analyzing the geometry and combinatorics of $S^3 \setminus J$.

To assist us in this analysis, we will make J even simpler by removing all remaining single crossings from the twist regions. The resulting link L has two kinds of components: *knot strands* coming from K that lie flat in the projection plane, and crossing circles C_i perpendicular to the projection plane. We call L a *flat augmented link*. If some twist region R_i had an odd number of crossings, $E(L)$ is no longer homeomorphic to $E(J)$; indeed, J and L can have a different number of components. We will address this issue later, in Section 3.2. See Figure 3.1 for a visual summary of this construction.

To subdivide $E(L)$ into polyhedra, we follow the construction of Agol and Thurston [16, Appendix]. First, we slice S^3 along the projection plane, cutting it into two identical 3-balls. Since they are identical, we focus our attention on B_1, the ball above the projection plane. The decomposition of B_2 proceeds in the same way. Each crossing circle C_i bounds a disk D_i, half of which lies in B_1 and borders on three edges in the projection plane. We then further slice B_1 along each of these half-disks.

This allows us to pull apart the two sides of each half-disk and flatten them,

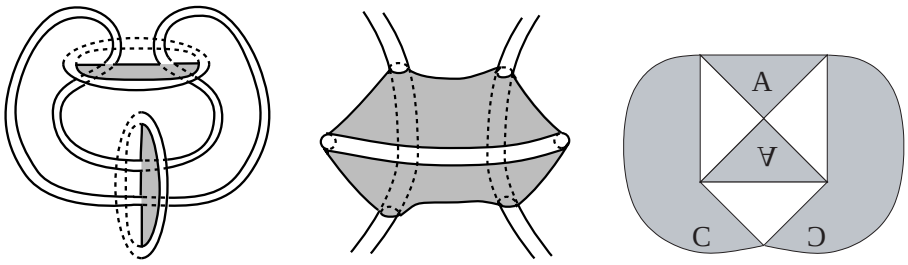

Figure 3.2: Decomposing $S^3 \setminus L$ into ideal polyhedra: First slice along the projection plane, then split remaining halves of two-punctured disks. Obtain polygon on right.

creating the planar diagram of a polyhedron. (See Figure 3.2.) This polyhedron will inherit one face from each region of the projection diagram and one face from each side of disk D_i. When we remove a closed regular neighborhood of L, we obtain a *truncated ideal polyhedron* $P_1 \subset E(L)$ (see Definition 2.2), whose *boundary faces* correspond to strands of L. The other ball B_2 becomes an identical polyhedron P_2.

P_1 and P_2 will each have six interior edges per crossing region, three from each side of the intersection between D_i and the projection plane. At each interior edge, a face coming from the projection place meets a face coming from D_i. This allows us to two-color the interior faces of P_1 and P_2 in a convenient fashion: the projection-plane faces will be white and the crossing-disk faces shaded, as in Figure 3.2.

To reconstruct $E(L)$ from P_1 and P_2, we first glue matching shaded faces in each P_j, and then glue the two polyhedra to each other along the white faces. Observe that in this gluing, the interior edges become 4-valent: each borders on two shaded faces (the two halves of D_i) and two white faces in the projection plane. In $\mathbb{R}^3$, we can position the crossing disks D_i perpendicular to the projection plane, creating dihedral angles of $\pi/2$ between adjacent faces. Conveniently, this feature carries over into hyperbolic geometry.

Theorem 3.1 (Purcell)**.** *Let $D(K)$ be a prime, twist-reduced diagram of a link K, with at least two twist regions. Let L be the flat augmented link obtained from $D(K)$. Then $E(L)$ is hyperbolic. Furthermore, the polyhedra P_1^o and P_2^o decomposing the interior of $E(L)$ are convex ideal polyhedra in $\mathbb{H}^3$, with totally geodesic faces that meet at right angles.*

Purcell proved this theorem by a direct construction [28, Theorem 6], using only Andreev's theorem about circle packings. We will give an alternate proof, shorter but less direct, that clarifies the relationship between our link K and the alternating links studied by Menasco [17].

Proof. Given the diagram $D(K)$, we can always reverse the crossings in some of the twist regions to obtain an alternating diagram $D(K')$. Furthermore, because the construction of the augmented link ignores the over-under crossing information, applying this construction to K' will yield the same link L. Thus every flat augmented link L is an example of what Adams calls an *augmented alternating link*. When L has 2 or more crossing circles, and thus K' has 2 or more twist regions, Menasco's theorem [17] implies $E(K')$ is hyperbolic. Then Adams' result on augmented alternating links [1] implies that every flat augmented link L is hyperbolic.

Additionally, note that there is an orientation-reversing involution of $S^3 \setminus L$ preserving L and our ideal polyhedra: namely, reflection through the projection plane. Every lift of this involution to the universal cover $\mathbb{H}^3$ is a reflection in a totally geodesic plane. Hence the polyhedra can be made totally geodesic in $\mathbb{H}^3$, with the shaded faces meeting the white faces at right angles. □

Remark 3.2. It is worth noting that the statement and proof of Theorem 3.1 do not assume that the original link K is hyperbolic. When $D(K)$ has at least two twist regions, it follows from Menasco's theorem that the alternating link K' is hyperbolic; we use this to bootstrap to a hyperbolic structure on $E(L)$. This will eventually be used to prove that K is hyperbolic (Theorem 1.4).

Recall that the truncated ideal polyhedra P_1 and P_2 have rectangular boundary faces on $\partial E(L)$. If we keep track of how these rectangles are glued along boundary edges in the gluing pattern of P_1 and P_2, we can construct a combinatorial picture of each cusp torus of $\partial E(L)$.

Lemma 3.3 (Purcell)**.** *The cusp tori of L are rectangular. For a crossing circle C_i, the cusp torus is composed of two boundary faces. For a knot strand K_j lying flat in the projection plane, the cusp torus is a $2\times n$ block of boundary faces, where n is the number of twist regions crossed by K_j (counted with multiplicity).*

Proof. Let $C_i \subset L$ be a crossing circle. As we saw in the construction of Section 3.1, P_1 intersects the tubular neighborhood of C_i in a rectangular boundary face F_1. The shaded faces on opposite sides of F_1 are glued to one another in the gluing pattern, since they glue to give half the disk D_i bounded by C_i. Thus an arc in F_1 connecting the two shaded faces represents a meridian of C_i.

The two white faces meeting F_1 are glued to corresponding faces of P_2, joining F_1 to the boundary rectangle F_2. Thus the cusp torus of C_i is tiled by F_1 and F_2, with the meridian and longitude as shown in Figure 3.3.

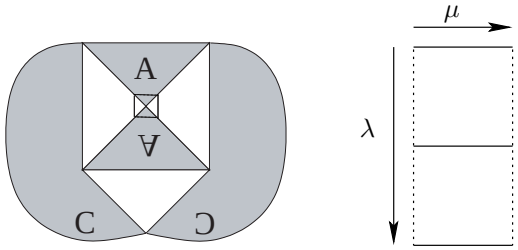

Figure 3.3: Left: Each crossing circle gives rise to one boundary face $F_1 \subset P_1$. Right: The cusp diagram corresponding to a crossing circle. Here μ is a meridian, and λ is a longitude.

For a knot strand $K_j \subset L$, P_1 will have one boundary face for each strand of K_j between adjacent crossing disks D_i. (See Figure 3.2.) These boundary rectangles are glued end to end along shaded faces coming from the D_i to complete a longitude of K_j. P_2 will give rise to an identical chain of rectangles, glued to the boundary rectangles of P_1 along the white faces of the projection diagram. Thus the cusp torus of K_j is tiled by a $2\times n$ block of rectangles, where n is the number of intersections between K_j and the crossing disks D_i, hence equal to the number of twist regions that K_j passes through, counted with multiplicity. See Figure 3.4. □

3.2 Half-twists and surgery slopes

Recall that to construct the flat augmented link L with its nice polyhedral decomposition, we took three steps, summarized in Figure 3.1. We added crossing circles to

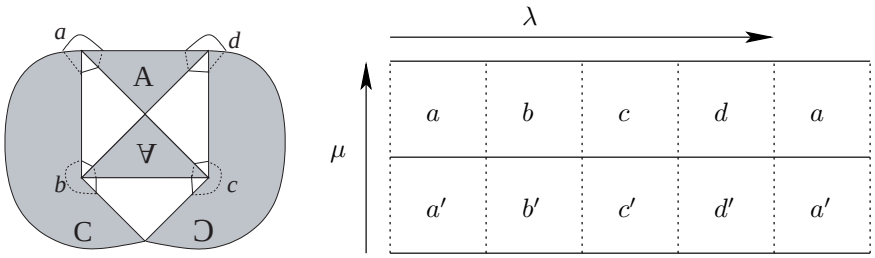

Figure 3.4: The cusp diagram for the knot strand cusp. Solid segments correspond to white faces, and dotted segments correspond to shaded faces.

K (obtaining link I); removed a whole number of twists per twist region (obtaining a homeomorphic link J); and then removed any remaining single crossings. Whereas any Dehn filling of K is a filling of J, the same is no longer true for L. Thus to obtain our results, we need to understand the combinatorics of the link J, with the half-twists re-inserted.

Conveniently, $E(J)$ can still be decomposed into the same polyhedra P_1 and P_2, only with a slightly modified gluing pattern. P_1 has one shaded face from each side of a crossing disk D_i; to construct $E(L)$, we glued those faces to each other. If instead we glue those shaded faces of P_1 to matching shaded faces of P_2, we effectively insert a half-twist along disk D_i and a single crossing into the projection diagram of L. We can do this wherever J has a single crossing. In particular, this simple rearrangement means that we have the following version of Theorem 3.1.

Theorem 3.4. *Let $D(K)$ be a prime, twist reduced diagram of a link K. Assume that $D(K)$ has at least two twist regions, with a_i crossings in twist region R_i. Let J be the augmented link constructed in Section 3.1, in which the number of crossings in region R_i is reduced to $a_i \mod 2$. Then*

(1) *$E(J)$ is hyperbolic,*

(2) *$E(J)$ subdivides into angled polyhedra P_1 and P_2 with dihedral angles $\pi/2$,*

(3) *K is the result of Dehn filling each crossing circle C_i of J along the surgery slope $1/s_i$, where we removed $2|s_i|$ crossings from R_i, and*

(4) *every Dehn filling of K is a filling of J.*

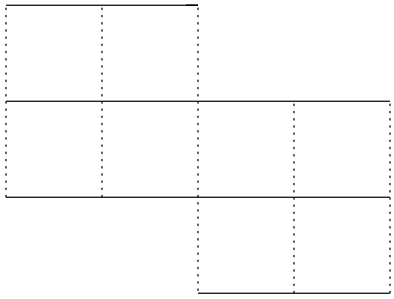

Figure 3.5: Cusp view: adding a half-twist to a flat augmented link shifts the gluing along the shaded faces.

Proof. The first two conclusions follow from Theorem 3.1 because $E(J)$ decomposes into the same polyhedra P_1 and P_2 as $E(L)$, with the same dihedral angles. By Theorem 2.4, P_1 and P_2 are angled polyhedra. The last two conclusions result from the process of constructing J, and are mentioned in Section 3.1. □

In the cusp diagrams of $\partial E(J)$, each half-twist in the transition from L to J will shift the gluing by one step along the shaded faces coming from D_i, as illustrated in Figure 3.5. Thus the neat rectangular pattern of Lemma 3.3 no longer holds. However, we can still make convenient statements about the universal cover $\widetilde{T}$ of each cusp torus of $S^3 \setminus J$.

Definition 3.5. Let T be a cusp torus of $\partial E(J)$, with universal cover $\widetilde{T} = \mathbb{R}^2$. Then $\widetilde{T}$ contains a rectangular lattice coming from white and shaded faces of P_1 and P_2. We construct a basis $\langle \boldsymbol{s}, \boldsymbol{w} \rangle$ of this $\mathbb{Z}^2$ lattice by letting $\boldsymbol{s}$ be a step parallel to a shaded face and $\boldsymbol{w}$ be a step parallel to a white face.

Lemma 3.6. *Let T be a cusp torus of $\partial E(J)$ and let $\langle \boldsymbol{s}, \boldsymbol{w} \rangle$ be the basis for the lattice on $\widetilde{T}$. In this basis, the fundamental domain of T appears as follows:*

(1) *If T comes from a crossing circle without a half-twist, it has meridian $\boldsymbol{w}$ and longitude $2\boldsymbol{s}$.*

(2) *If T comes from a crossing circle with a half-twist, it has meridian $\boldsymbol{w} \pm \boldsymbol{s}$ (depending on the direction of the twist) and longitude $2\boldsymbol{s}$.*

(3) *If T comes from a component K_j of the original link K, it has meridian $2\boldsymbol{s}$ and longitude $n\boldsymbol{w} + k\boldsymbol{s}$, where K_j runs through n twist regions with multiplicity and k is an undetermined integer.*

Proof. If J does not contain any half-twists, this is a restatement of Lemma 3.3. (See Figures 3.3 and 3.4.) Each half-twist along the crossing circle C_i shears the meridian of C_i by $\boldsymbol{s}$, as described above. It also shears by $\boldsymbol{s}$ the cusp of every component of the original link K passing through the crossing disk D_i.

Thus if K_j passes through a half-twist m times, the projection of the curve $n\boldsymbol{w}+m\boldsymbol{s}$ to T will be *some* longitude of K_i, in the sense of completing a basis of $\pi_1(T)$ along with $\mu = 2\boldsymbol{s}$. The *true* longitude, in the sense of having linking number 0 with K_j, is then some curve of the form $n\boldsymbol{w} + k\boldsymbol{s}$ for some integer k. □

The basis $\langle \boldsymbol{s}, \boldsymbol{w} \rangle$ also allows us to make precise statements about the surgery curves on $\partial E(J)$ that correspond to non-trivial surgeries on K.

Theorem 3.7. *Let $K = \cup_{j=1}^{m} K_j$ be a link in S^3 with a prime, twist reduced diagram $D(K)$. Suppose that $D(K)$ contains twist regions $R_1, \ldots, R_n$ ($n \geq 2$) and that twist region R_i contains a_i crossings. For each component K_j, let n_j be the number of twist regions crossed by K_j, counted with multiplicity; and let s_j be a non-trivial surgery slope on K_j.*

With this notation, the surgery on $S^3 \backslash K$ along slopes $s_1, \ldots, s_m$ can be represented as a surgery on J as follows:

(1) *On the (mostly) planar component of J corresponding to K_j, the surgery curve is $p_j n_j \boldsymbol{w} + q_j \boldsymbol{s}$, for some integers $p_j \neq 0$ and q_j.*

(2) *On the crossing circle C_i, the surgery curve is $\boldsymbol{w} \pm a_i \boldsymbol{s}$.*

Proof. By Lemma 3.6, K_j has meridian $2\boldsymbol{s}$ and a longitude of the form $n_j\boldsymbol{w} + k_j\boldsymbol{s}$. Since s_j is a non-trivial surgery slope, it must cover at least one longitude. In particular, the number of steps that a curve representing s_j takes along the white faces is a nonzero multiple of n_j.

To prove conclusion (2), suppose first that a_i is even, so J has no half-twist at C_i. By Lemma 3.6, C_i has longitude $2\boldsymbol{s}$ and meridian $\boldsymbol{w}$. By Theorem 3.4, the surgery curve on C_i traverses $a_i/2$ longitudes and one meridian, proving the result.

Now, suppose that a_i is odd. Then in the construction of J, we have removed $2b_i = a_i - 1$ crossings; the remaining half-twist of J at C_i goes in the same direction as the twists of K. By Lemma 3.6, C_i has longitude $2\boldsymbol{s}$ and meridian $\boldsymbol{w} + \sigma_i \boldsymbol{s}$, for some $\sigma_i = \pm 1$. By Theorem 3.4, the surgery curve traverses $\sigma_i b_i$ longitudes (with the same σ_i) and one meridian. Thus, in the basis of $\langle \boldsymbol{s}, \boldsymbol{w} \rangle$, the surgery curve is

$$\begin{aligned} \mu + \sigma_i b_i \lambda &= (\boldsymbol{w} + \sigma_i \boldsymbol{s}) + \sigma_i b_i (2\boldsymbol{s}) \\ &= \boldsymbol{w} + \sigma_i (1 + 2b_i) \boldsymbol{s} \\ &= \boldsymbol{w} + \sigma_i a_i \boldsymbol{s} \,. \end{aligned}$$

□

3.3 Normal surfaces in the augmented link

Now that we understand exactly how the Dehn fillings of K are represented as fillings of J, we can estimate the lengths of surgery slopes by applying the normal surface theory of Sections 2.2–2.3 to the truncated ideal polyhedra P_1 and P_2. In fact, Theorem 3.4 tells us that P_1 and P_2 are examples of a special type of angled polyhedron, which we will call *rectangular-cusped*.

Definition 3.8. Let P be an angled polyhedron (see Definition 2.3) in which we have truncated the ideal vertices. We say that P is *rectangular-cusped* if

(1) each boundary face of P (each face of $P \cap \partial M$) meets 4 interior edges, and

(2) each interior edge is labeled with angle $\pi/2$.

Rectangular-cusped polyhedra have two convenient features. First, their interior faces can be two-colored, in a similar fashion to the white and shaded faces of P_1 and P_2. Around each rectangular boundary face, opposite interior faces have the same color. Second, making all dihedral angles equal to $\pi/2$ ensures that all combinatorial areas are multiples of $\pi/2$.

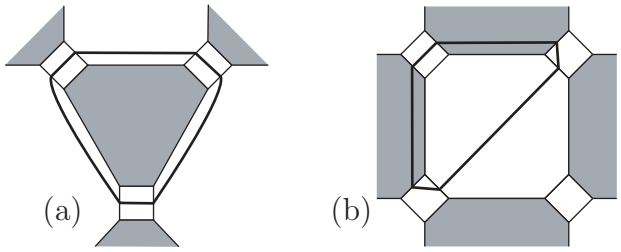

Figure 3.6: Two ideal triangles in a rectangular-cusped polyhedron.

In addition to the vertex links and boundary bigons of area 0 (see Figure 2.3), we need to define a third kind of special admissible disk.

Definition 3.9. Let P be a truncated ideal polyhedron. An admissible disk $D \subset P$ is called an *ideal triangle* if

(1) $\partial D \subset \partial P$,

(2) ∂D intersects the boundary faces of P exactly three times, and

(3) ∂D is disjoint from the interior edges of P.

Two examples are shown in Figure 3.6. Note that an ideal triangle D has area $a(D) = \pi$ and length $\ell(\gamma, D) = \pi/3$ for each segment γ of $\partial D \cap \partial M$.

Proposition 3.10. *Let $D \subset P$ be an admissible disk in a rectangular-cusped polyhedron, such that ∂D passes through at least one boundary face. Let $\gamma \subset \partial M$ be a boundary segment of ∂D. If D is not a bigon or an ideal triangle,*

$$\ell(\gamma, D) \geq \frac{\pi}{2}\,.$$

Proof. We consider different cases, conditioned on $n = |\partial D \cap \partial M|$. By Definition 2.16, $\ell(\gamma, D) = a(D)/n$.

Case 1: $n = 1$. For this case, we need to prove that $a(D) \geq \pi/2$. An admissible disk with one component of $\partial D \cap \partial M$ cannot be a vertex link or boundary bigon, so by Lemma 2.13, $a(D) > 0$. Since all areas in a rectangular-cusped polyhedron are multiples of $\pi/2$, $a(D) \geq \pi/2$.

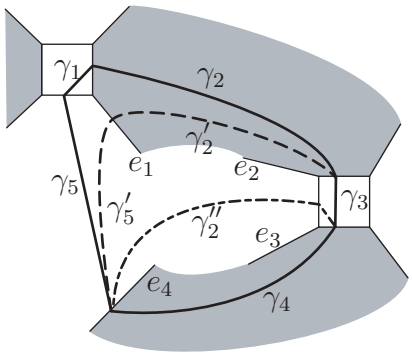

Figure 3.7: Schematic picture for Case 2 of Proposition 3.10.

Case 2: $n = 2$. For this case, we need to prove that $a(D) \geq \pi$. If $a(D) = 0$, D is a boundary bigon, excluded by the hypotheses. So we need to rule out the possibility that $a(D) = \pi/2$.

If such a disk were to occur, it would have to have $\partial D \subset \partial P$, and ∂D would have to intersect exactly one interior edge. Then ∂D passes through three interior faces, which cannot all have the same color because two of them share an edge. Thus a segment $\gamma_1 \subset \partial D$ in a boundary face must connect adjacent interior faces, for otherwise all three interior faces would have the same color. See Figure 3.7 for a schematic picture.

We can pull γ_1 off the boundary face and have it intersect interior edge e_1. This creates a new disk D' with one segment on ∂M and area 0, since this isotopy reduced the area by $\pi/2$. If D' were admissible, it would be a counterexample to Case 1. Thus $\partial D'$ must violate some condition of admissibility. The only way this can happen is if one of the new segments of $\partial D'$, γ_2' or γ_5', has both endpoints on the same edge, or on adjacent interior and boundary edges. But since D is admissible, e_1 and e_4 must be distinct edges, so γ_5' has endpoints on distinct edges.

Thus γ_2' connects adjacent interior and boundary edges, and so $e_1 = e_2$. We can then isotope γ_2' across this interior edge, creating a new disk D'' that has just one intersection with ∂M and one intersection with an interior edge. Since γ_4 and the new segment γ_2'' lie in adjacent faces of P, we have $e_3 = e_4$. Then γ_4 connects adjacent interior and boundary edges, contradicting the assumption that D was admissible.

Therefore, such a disk D does not exist.

Case 3: $n = 3$. For this case, we need to prove that $a(D) \geq 3\pi/2$. The three components of $\partial D \cap \partial M$ already ensure that $a(D) \geq \pi$. So if ∂D also intersects an interior edge or the interior of P, we have $a(D) \geq 3\pi/2$. Otherwise, D is an ideal triangle, excluded by the hypotheses.

Case 4: $n \geq 4$. For this case,

$$a(D) \geq n \cdot \pi - 2\pi \geq \frac{n}{2} \cdot \pi \,,$$

proving the lemma. □

Thus ideal triangles are the only admissible disks of nonzero area that contribute less than $\pi/2$ to combinatorial length. To obtain the best possible bounds on the length of surgery curves, we need to find out more about how these triangles fit into polyhedra P_1 and P_2 that decompose the link complement $E(J)$.

Lemma 3.11. *Let P be a truncated ideal polyhedron, and let $D \subset P$ be an ideal triangle. Then all the segments of ∂D lie in distinct faces of ∂P, and D is normal.*

Proof. ∂D consists of six segments, alternating between boundary and interior faces. Label them $\gamma_1, \ldots, \gamma_6$. If two of these segments (say, γ_1 and γ_3) lie in the same face of P, then a third segment (γ_2) must have both endpoints on the same edge, violating the definition of an admissible disk. Thus each γ_i lies in a different face, so ∂D is embedded. Since ∂D intersects each boundary face at most once and is disjoint from the interior edges altogether, D must be normal. □

For the rest of this chapter, we will work directly with the polyhedra P_1 and P_2, and the only manifolds we will consider are $E(J)$ and its Dehn fillings.

Definition 3.12. In polyhedra P_1 and P_2, we will classify ideal triangles into three types. A triangle of *type S* is one that is parallel to a shaded face, as in Figure 3.6(a). A triangle of *type W* is one that is parallel to a white face, as in Figure 3.6(b). An ideal triangle parallel to no face of its polyhedron will be of *type N*.

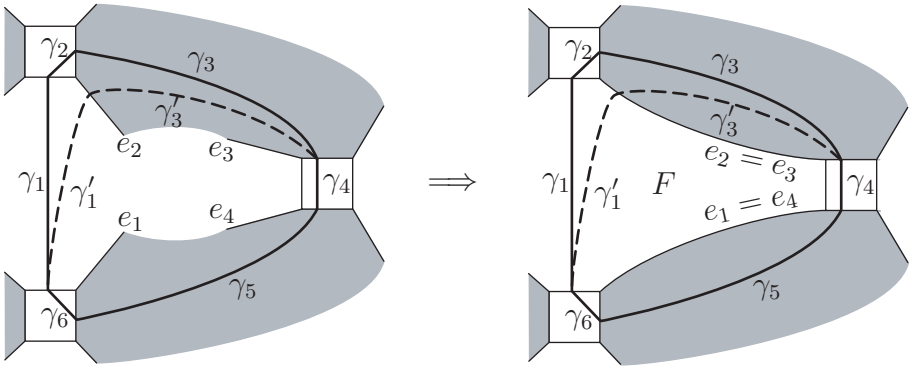

Figure 3.8: Schematic picture of an ideal triangle intersecting faces of both colors. The shading is generic, and might be reversed.

Lemma 3.13. *Let D be an ideal triangle in $P \in \{P_1, P_2\}$. Let $\gamma_1, \ldots, \gamma_6$ be the segments of ∂D. Then the following hold:*

(1) *If D is of type S or type W, then at least two of the γ_i are parallel to interior edges of P.*

(2) *If D is of type N, then no γ_i is parallel to an interior edge, and the three interior faces of P intersecting ∂D are all white faces.*

Proof. We consider two cases:

Case 1: ∂D intersects both white and shaded faces. Then D can be schematically represented by the left side of Figure 3.8. Label the stumps of interior edges $e_1, \ldots, e_4$, as in the figure; some of these are likely to be part of the same edge. Now, we can pull segment γ_2 off the boundary face and have ∂D intersect edge e_2 instead. This creates a disk D' of area $\pi/2$, which could a priori be normal. However, by Case 2 of Proposition 3.10, there are no normal disks that have two intersections with ∂M and area $\pi/2$. Thus D' fails some part of Definition 2.10.

Since D is normal by Lemma 3.11, the only way that D' can fail to be normal is if one of the new segments, γ_1' or γ_3', connects adjacent boundary and interior faces. If γ_1' violates normality, e_1 is the same edge as e_2. But then γ_3 and γ_5 must lie in the same face, contradicting Lemma 3.11.

If γ_3' violates normality by connecting adjacent boundary and interior faces, we can tighten $\partial D'$ by removing its intersection with $e_2 = e_3$. This creates a new disk D'' with area 0. Segment γ_5 and the isotopic image of γ_3 lie in distinct faces because they are on opposite sides of edge e_4. So D'' is normal, and thus a boundary bigon. Then we can conclude that $e_1 = e_4$, and the original disk D was parallel to face F, into which we have pulled γ_3 (see Figure 3.8, right). So D is of type S or W. Notice that both γ_3 and γ_5 are parallel to edges of F.

Case 2: All interior faces intersecting ∂D are the same color. If some segment γ_i is parallel to an interior edge, we can isotope ∂D across that edge, into a face of a different color, putting us in Case 1. Otherwise, if no γ_i is parallel to an interior edge, the three interior faces must all be white. (Shaded faces are all triangles, in which any arc connecting distinct ideal vertices is parallel to an edge.) By Lemma 3.11, the segments γ_i all lie in distinct faces, so since none of them is parallel to an edge, D cannot be parallel to a face. Thus D is of type N, and satisfies conclusion (2) of the Lemma. □

Corollary 3.14. *In an admissible surface in $E(J)$, an ideal triangle of type N cannot be glued to a bigon or a triangle of type S.*

Proof. Let F be a shaded face of P_1 or P_2, and D be a type S ideal triangle parallel to F. Since shaded faces are all triangles, every interior segment of ∂D is parallel to an interior edge of F, hence an edge of $E(J)$. Similarly, both interior segments on the boundary of a bigon are parallel to an edge of $E(J)$. On the other hand, by Lemma 3.13 the boundary of a type N ideal triangle does not have any segments parallel to interior edges. □

3.4 Progressive arcs and length estimates

We are now ready to estimate the combinatorial length of surgery slopes on $\partial E(J)$.

Definition 3.15. Let T be a torus of $\partial E(J)$. Recall that, by Definition 3.5, its universal cover $\widetilde{T}$ contains a lattice of shaded and white faces, generated by a basis $\langle \boldsymbol{s}, \boldsymbol{w} \rangle$. If T is a crossing circle cusp, we will say that the $\boldsymbol{w}$ direction is *meridional*

and the $\boldsymbol{s}$ direction is *longitudinal*. If T is a knot strand cusp, we will say that the $\boldsymbol{s}$ direction is *meridional* and the $\boldsymbol{w}$ direction is *longitudinal*. (By Lemma 3.6, the meridian and longitude of T are in fact aligned primarily in these directions.)

Thus if a segment γ spans opposite edges of a boundary face $B \subset \partial E(J)$, it makes sense to talk of γ lying in a meridional or longitudinal direction.

Definition 3.16. Let $P \in \{P_1, P_2\}$, and let $D \subset P$ be an admissible disk. Then D can intersect a boundary face $B \subset \partial E(J)$ in one of three types of segments: a *longitudinal segment*, connecting opposite edges of B in a longitudinal direction; a *meridional segment*, connecting opposite edges of B in a meridional direction; or a *diagonal segment*, connecting adjacent edges of B.

To estimate the combinatorial length of surgery slopes on $\partial E(J)$ representing a surgery slope, it helps to divide a curve into smaller pieces.

Definition 3.17. Let T be a torus of $\partial E(J)$, and let $\gamma \subset T$ be a non-closed simplicial arc (see Definition 2.17). Lift γ to an arc $\widetilde{\gamma} \subset \widetilde{T}$, and cut $\widetilde{T}$ into vertical strips along meridional faces in the lattice. We say that γ is a *progressive arc* if $\widetilde{\gamma}$ is contained entirely in one of these vertical strips, and the endpoints of $\widetilde{\gamma}$ lie on opposite sides of the strip.

In other words, a progressive arc on a crossing circle cusp has endpoints on consecutive white faces, and constitutes a step in the $\boldsymbol{s}$ direction. A progressive arc on a knot cusp has endpoints on consecutive shaded faces, and constitutes a step in the $\boldsymbol{w}$ direction. In either case, a progressive arc γ can consist of (a) a single longitudinal segment, (b) two diagonal segments connecting to different meridians, or (c) two diagonals with some number of meridional segments between them. These basic types are shown in Figure 3.9.

Lemma 3.18. *Let $\gamma \subset \partial E(J)$ be a progressive arc. Then $\ell(\gamma) \geq \pi/3$.*

Proof. Let H be an inward extension of γ (see Definition 2.18). For each admissible disk $D_i \subset H$ bordering on a segment $\gamma_i \subset \gamma$, $\ell(\gamma_i, D_i) = 0$ if and only if D_i is a boundary bigon. By Proposition 3.10, every other type of disk contributes at least

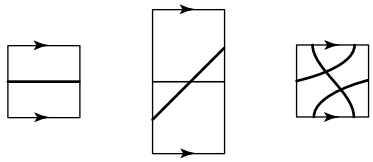

Figure 3.9: The three types of progressive arcs.

$\pi/3$ to $\ell(\gamma)$. So the only way to have $\ell(\gamma) < \pi/3$ is if H consists only of bigons. However, a string of bigons circles around a single edge of $E(J)$, which means that its intersection with a component of $\partial E(J)$ cannot be a progressive arc. □

Corollary 3.19. *Let T be a torus of $\partial E(J)$, and let s be a non-trivial surgery slope on T. If T comes from a crossing circle C_i, let n be the number of crossings in region R_i; if T comes from a component K_j of K, let n be the number of twist regions visited by K_j, counted with multiplicity. Then, in either case,*

$$\ell(s) \geq \frac{n\pi}{3}.$$

Proof. By Theorem 3.7, a surgery curve on a crossing circle corresponding to n crossings must cross at least n (white) meridional faces, and any surgery curve on a component of K passing through n twist regions with multiplicity must cross at least n (shaded) meridional faces. Specifically, we can say that they must each contain at least n progressive arcs. Thus the result follows from Lemma 3.18. □

For surgery curves on a crossing circle cusp, which by Theorem 3.7 look like $n\boldsymbol{s} \pm \boldsymbol{w}$ in the basis $\langle \boldsymbol{s}, \boldsymbol{w} \rangle$, we can obtain a slightly better estimate.

Proposition 3.20. *Let $s \subset \partial E(J)$ be a surgery slope on a crossing circle cusp that yields n crossings. Then we have the strict inequality*

$$\ell(s) > \frac{n\pi}{3}.$$

Proof. By Corollary 3.19, we must only rule out equality. Equality occurs when a simplicial curve c representing s contains exactly n progressive arcs, an inward

extension of c picks up length exactly $\pi/3$ per progressive arc, and any part of c not covered by progressive arcs contributes zero length. Consider such a curve.

If a progressive arc $\gamma \subset c$ has combinatorial length $\pi/3$, it must have an inward extension whose area comes from a single triangle D. D cannot be of type W, because white faces are meridional on a crossing circle cusp, and thus a triangle of this type, plus some bigons, cannot have their boundary segments add up to a progressive arc. Thus D must be a triangle of type S or type N.

Let H be an inward extension of c. We claim that if H contains a type-N triangle, then it consists entirely of type-N triangles. This is because by Corollary 3.14, a type-N triangle D cannot be glued to a type-S triangle or a bigon, and any other type of admissible disk glued to D would contribute extra area and bring the total length above $n\pi/3$. But if H consists entirely of type-N triangles, c consists entirely of longitudinal segments and never travels in the $\boldsymbol{w}$ direction. Thus we can conclude that H cannot contain any type-N triangles.

The only remaining possibility is that H consists entirely of type-S triangles and bigons. But in this case, all of H is parallel to a single shaded disk, and again c never traverses the lattice in the $\boldsymbol{w}$ direction. Thus the assumption that $\ell(c) = n\pi/3$ leads to a contradiction. □

We are now in a position to prove three of the theorems listed in the introduction.

Theorem 1.4. *Let $K \subset S^3$ be a link with a prime, twist-reduced diagram $D(K)$. If $D(K)$ has at least two twist regions and every twist region of $D(K)$ contains at least 6 crossings, then K is hyperbolic.*

Proof. The assumption that $D(K)$ has at least two twist regions ensures that all the constructions and results of Section 3.1 apply. Thus, by Theorem 3.4, K is obtained by Dehn surgery on the crossing circles of a hyperbolic link J. By Proposition 3.20, every surgery slope s_i on a crossing circle C_i has combinatorial length $\ell(s_i) > 2\pi$. Therefore, by Theorem 2.22, $E(K)$ is hyperbolic. □

Theorem 1.9. *Let K be a link in S^3 with a prime, twist-reduced diagram $D(K)$. Suppose that every twist region of $D(K)$ contains at least 6 crossings and each component of K passes through at least 7 twist regions (counted with multiplicity). Then*

(1) *any non-trivial Dehn filling of some but not all components of K is hyperbolic, and*

(2) *any non-trivial Dehn filling of all the components of K is hyperbolike.*

Proof. By Corollary 3.19, any non-trivial slope s on a component of K will have $\ell(s) > 2\pi$, and by Proposition 3.20, the same is true for surgery slopes on the crossing circles. Thus all surgery slopes on $\partial E(J)$ are sufficiently long. Conclusion (1) now follows by Theorem 2.22, and conclusion (2) by Theorem 2.21. □

Theorem 1.6. *Let $K \subset S^3$ be a link of k components with a prime, twist-reduced diagram $D(K)$. If $D(K)$ has $t \geq 2$ twist regions and at least 6 crossings in each twist region, then*

$$\text{genus}(K) \geq \left\lceil 1 + \frac{t}{6} - \frac{k}{2} \right\rceil,$$

where $\lceil \cdot \rceil$ is the ceiling function that rounds up to the nearest integer.

Proof. Let F be a Seifert surface for K, that is, an orientable incompressible surface whose boundary is K. Then F contains a punctured surface $G \subset E(J)$, where ∂G consists of curves $\gamma_1, \ldots, \gamma_k$ that run along K and curves $\gamma_{k+1}, \ldots \gamma_{k+n}$ along the crossing circles. We can place G in normal form in the polyhedra P_1 and P_2 and compute its combinatorial area. Observe that, by Corollary 3.19, the total length of $\gamma_1, \ldots, \gamma_k$ is at least $2t\pi/3$, because K passes through each twist region twice. By Proposition 3.20, $\ell(\gamma_i) > 2\pi$ for $i > k$. Thus we can compute that

$$\begin{aligned}
2\pi \cdot \text{genus}(G) &= 2\pi \left(1 - \frac{1}{2}\chi(G) - \frac{1}{2}(k+n)\right) \\
&= 2\pi + \frac{1}{2}a(G) - \pi k - \pi n \\
&\geq 2\pi + \frac{1}{2}\sum_{i=1}^{k} \ell(\gamma_i) - \pi k + \frac{1}{2}\sum_{i=k+1}^{k+n} \ell(\gamma_i) - \pi n \\
&\geq 2\pi + \frac{t\pi}{3} - \pi k \\
&= 2\pi \left(1 + \frac{t}{6} - \frac{k}{2}\right).
\end{aligned}$$

Since genus(F) = genus(G) is an integer, we are done. □

Observe that the inequality in the computation is an equality whenever G doesn't meet any crossing circles and consists of only type-W ideal triangles. This can happen when the twist regions of $D(K)$ always meet in threes and G lies in the projection plane. In this situation, Theorem 1.6 actually gives the exact genus of K.

Chapter 4

Diagrams of Arborescent Links

The remainder of this thesis is concerned with constructing an angled polyhedral decomposition for the complement of an arborescent link. The actual construction, which will be carried out in Chapters 5 and 6, relies in a concrete and fundamental way on a particularly nice diagram of the link. This chapter is devoted to proving that every arborescent link has such a diagram.

Following Gabai [10], we encode the diagrams of arborescent links using a combinatorial object called a *weighted tree*. In Section 4.1, we will define these trees and use them to construct a link projection. In Section 4.2, we will describe the correspondence between *branches* of a weighted tree and braids and rational tangles in the diagram. This allows us to show that the links constructed from a tree are the same arborescent links that we have defined in the Introduction. In Section 4.3, we will use the weighted tree to simplify the link diagram into a maximally alternating form. In this simplified diagram, every braid and every rational tangle has an alternating projection, as do many larger pieces of the link. Finally, in Section 4.4, we will turn the weighted tree into a new combinatorial object called an *expanded tree*, whose structure is particularly well suited for triangulating the link complement.

4.1 From weighted trees to arborescent diagrams

Definition 4.1. A *weighted tree* T is a finite, contractible 1-complex embedded in S^2. Let $v \in T$ be a vertex and e_i and e_{i+1} be edges emanating out of v, ordered counterclockwise. Each such ordered pair is assigned an integer $w(e_i, v, e_{i+1})$, called a *weight*. Thus every vertex of valence k will have k weights. (See Figure 4.1.)

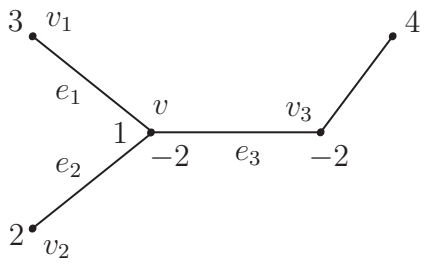

Figure 4.1: A non-degenerate weighted tree. Weights of 0 are suppressed.

Definition 4.2. A weighted tree T is called *degenerate* if some terminal vertex v has weight 0 or ± 1, or if a divalent vertex v' adjacent to a terminal vertex has weight 0. All other trees are called *non-degenerate.*

Definition 4.3. Consider a crossing in a link diagram, in which the four endpoints of the strands lie at the corners of a unit square. The crossing is called *positive* if the strand of slope $+1$ lies over the strand of slope -1, and *negative* otherwise.

Figure 4.2: Positive and negative crossings.

A weighted tree T encodes a link diagram $D_T(K)$, as follows. Start with any vertex v of T, and represent v by a horizontal band. Delete from the band the interiors of small 3-balls (*bubbles*) $B_1, \ldots, B_k$, with one B_i for each edge e_i incident to v. The band has a natural *core* direction c corresponding to a vector pointing counterclockwise around v and a *normal* direction n corresponding to a vector pointing toward v.

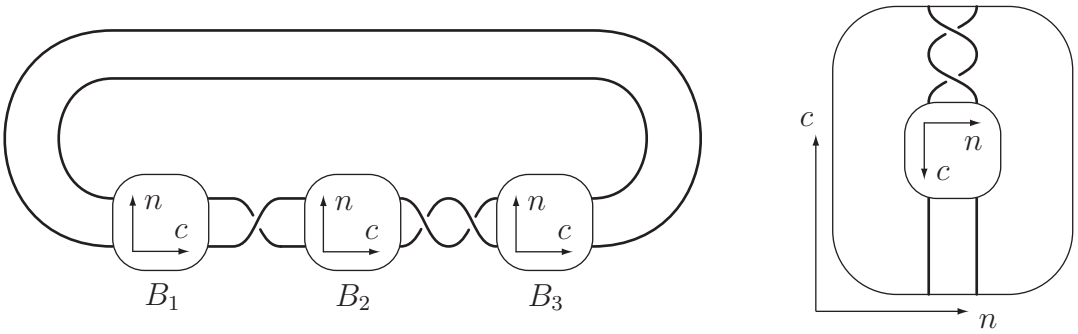

Figure 4.3: The bands of v and v_3 from Figure 4.1.

Between bubbles B_i and B_{i+1} (in that order along the core direction c), we place $w(e_i, v, e_{i+1})$ positive or negative crossings, as determined by the sign of the weight. See Figure 4.3.

In a similar way, each vertex v_i $(i = 1, \ldots, k)$ adjacent to v can be represented by a vertical band with the same number of bubbles as incident edges, and positive or negative crossings between the bubbles. Since we are performing the construction in S^3, it will be convenient to make the bubble corresponding to e_i *exterior*, i.e. containing the point at infinity. Now, the vertical band of v_i can be glued across its exterior bubble to the interior of bubble B_i. We perform the gluing in a way that aligns the core direction of one band with the normal direction of the other band.

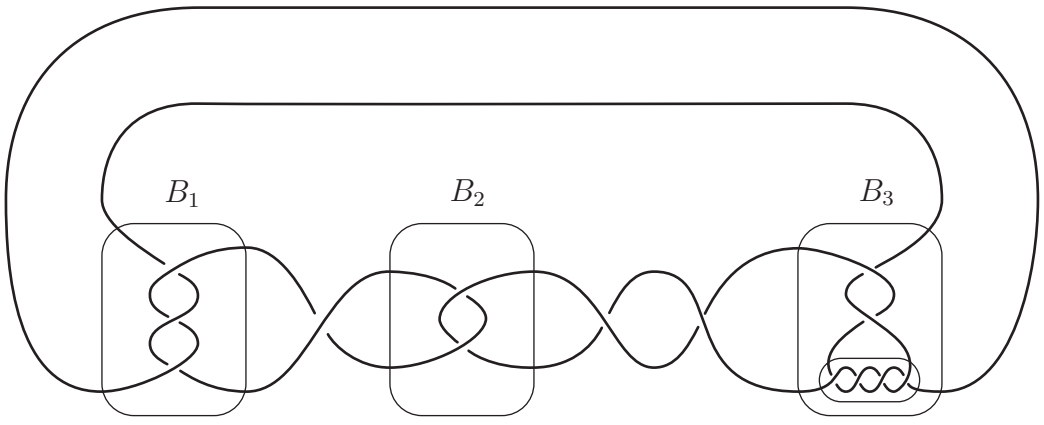

Figure 4.4: The projection diagram constructed from the tree of Figure 4.1.

We can iterate this process, with bands corresponding to vertices further from v glued into smaller and smaller bubbles. Each edge of T corresponds to a sphere along which the bubbles of two vertices are glued up. By the time we have traversed the entire tree, all the bubbles have been glued along these spheres, and we have a link projection as in Figure 4.4.

This construction contains one ambiguity: if we start the process from a different vertex, the alignment of horizontal and vertical bands could be reversed. (For example, this will be the case if we start from one of the vertices v_i adjacent to v.) Rotating the resulting link diagram would give Figure 4.4 with all the crossings reversed, because positive crossings become negative when they are rotated. However, the tree T does define a link projection uniquely up to isotopy and reflection.

Definition 4.4. An *arborescent link* is a link $K \subset S^3$ whose projection diagram can be constructed from a non-degenerate weighted tree T. The diagram $D_T(K)$ is called an *arborescent diagram*.

Remark 4.5. Although our construction of arborescent link diagrams is inspired by David Gabai's monograph [10], there are two key differences in notation. Firstly, Gabai uses a slightly different correspondence between weights and crossings. In our notation, an alternating diagram is represented by a tree whose weights all have the same sign; in his notation, the same diagram is represented by a tree in which neighboring vertices have weights with opposite signs.

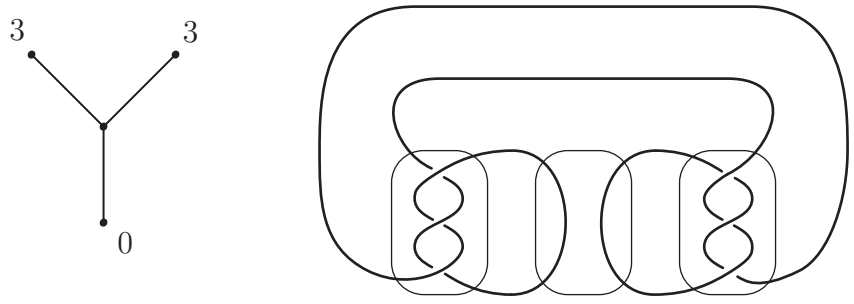

Figure 4.5: The connected sum of two trefoils, constructed from a degenerate tree.

Secondly, unlike Gabai, we restrict our definition of arborescent links to non-degenerate weighted trees. This restriction is necessary for our purposes, because degenerate trees allow the construction of composite and even disconnected links. For example, Figure 4.5 shows how to construct the connected sum of two trefoils from a degenerate tree. Split and composite links are definitely not hyperbolic, and many authors (for example, Wu [35, 36]) do not consider them to be arborescent. We will prove in Theorem 4.11 that this restrictive definition of arborescent links is equivalent to Definition 1.15, given in the introduction.

In the sequel, we will implicitly assume that weighted trees are non-degenerate.

We will simplify the tree T and its associated diagram $D_T(K)$ by a sequence of moves, the first two of which we can describe right now.

(1) *If a divalent vertex v has weight* 0*, collapse the vertex and the two edges.* The vertices v_1 and v_2 adjacent to v become identified to a single vertex, and the weights on each side of the pair of edges are added to form a single weight. As a result, some crossings may cancel. (See Figure 4.6.)

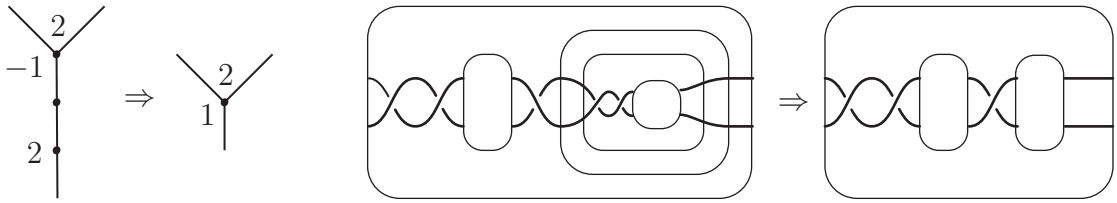

Figure 4.6: Collapsing a divalent vertex of weight 0 and the adjacent edges.

(2) *At each vertex v, move all (non-zero) weights to one spot.* Let B_i be a bubble on the band of v. A *flype* of the band that rotates the interior of B_i by 180° about the core axis will move a crossing from one side of B_i to the other. (See Figure 4.7.) In the tree T, this flype transfers a unit of weight from one side to the other side of the corresponding edge e_i, while also rotating the sub-tree at the far end of e_i. By repeating this procedure, we can collect all the weights at v at one spot while concatenating all the crossings on its band into a single twist region. In the process, some crossings may cancel.

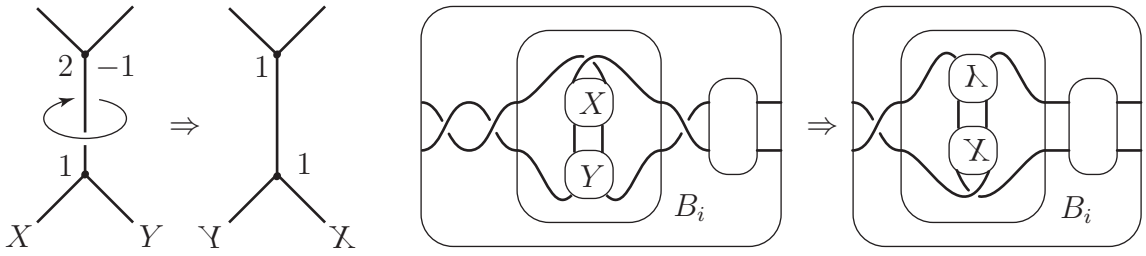

Figure 4.7: A *flype* of a band moves a crossing to the other side of a bubble. Repeated flypes combine the weights on a vertex to a single spot, possibly canceling crossings.

As a result of moves (1) and (2), each vertex of T with nonzero weight corresponds to a twist region in $D_T(K)$, where the number of crossings in the twist region is the absolute value of the weight. We will denote the weight of the vertex as simply $w(v)$, although we will continue to keep track of its location between two edges at v.

4.2 Branches, braids, and tangles

The correspondence between the tree and the arborescent diagram gives a way to subdivide an arborescent link into pieces larger than single twist regions.

Definition 4.6. A vertex v of a weighted tree T is called a *node* if it adjoins at least 3 edges.

Definition 4.7. Let $T \subset S^2$ be a weighed tree, and remove from T an open regular neighborhood of every node. The complement is a disjoint collection of linear subtrees, which may have truncated *stump edges* that connected to the nodes. Every such subtree containing at least one vertex is called a *branch*. A *terminal branch* is one that contains one terminal vertex of T and thus has only one stump edge, and an *interior branch* has two stump edges, one at each end. See Figure 4.8.

Notation. Let $v_1, \ldots, v_n$ be the vertices of a branch U, where v_1 is adjacent to a stump edge. If U is an interior branch, let S_0 and S_n be the spheres on the bands of v_1 and v_n, respectively, that correspond to stump edges at the two ends of U. If

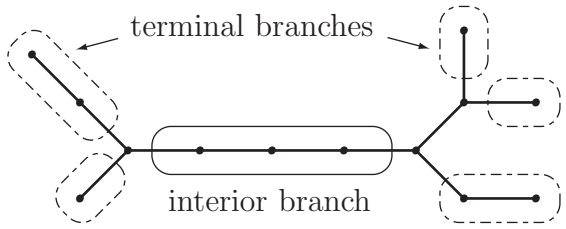

Figure 4.8: The branches of a weighted tree. (Weights are not shown.)

U is a terminal branch, let S_0 again be the sphere corresponding to the stump edge of U, and let S_n be a sphere on the band of v_n bounding a bubble that contains no crossings.

The branches of a tree correspond to braids and tangles in the diagram. Recall from Definition 1.13 that a rational tangle consists of a 3-ball B and two boundary-parallel arcs inside the ball.

Lemma 4.8. *Let T be a weighted tree and K be the corresponding link. Then*

(1) *Every interior branch $U \subset T$ defines a braid on 4 strings in $D_T(K)$, with boundary at the spheres S_0 and S_n. The complement of the braid is a product region between two 4-punctured spheres.*

(2) *Every terminal branch defines a rational tangle B, whose boundary is $\partial B = S_0$.*

Proof. Conclusion (1) follows from an easy induction on n, the number of vertices in the branch. If $n = 1$, then the band of v_1 looks like the right side of Figure 4.3. It is clear that the strands of K form a braid, whose complement is a product region between the two bubbles. Adding additional vertices and edges to the branch continues the braid (twisting on different strings), and the conclusion continues to hold.

Conclusion (2) follows from (1). On the band of v_n, the bubble B_n inside sphere S_n contains no crossings; thus the strings of K inside the bubble clearly form a rational tangle. Now, the crossings on the bands of $v_1, \ldots, v_n$ build a braid (hence a product region) between S_0 and S_n, preserving the property that the strands of K are boundary-parallel. Thus the sphere S_0 bounds a rational tangle. □

Remark 4.9. In Definition 4.7, we have omitted the case when the entire tree T is linear. Such a tree can be artificially split along some edge into two terminal branches. It then follows from Lemma 4.8 that the resulting link K is obtained by gluing two rational tangles – i.e., K is a *two-bridge link*.

One other special case is when T is a *star*, i.e., a tree with a single node. Following Lemma 4.8, we can conclude that the corresponding link K is obtained by gluing several rational tangles together in a cyclic fashion. Thus K is a *Montesinos link* (also called a a *star link*), as in Figure 1.7.

By studying branches, we can check that arborescent links as defined via weighted trees (Definition 4.4) are the same as the arborescent links defined in the Introduction.

Lemma 4.10. *Let A be an arborescent tangle. (See Definition 1.14.) Then a projection diagram of A can be represented by a non-degenerate weighted tree T that has been truncated along a single stump edge.*

Proof. We proceed by induction on the number of rational tangles used to construct A. If A is itself rational, then its projection can be simplified into a series of crossings along horizontal and vertical bands, as in Figure 1.5. Thus, by Lemma 4.8, A can be represented by a terminal branch U, which is a truncated tree satisfying the lemma. (U is non-degenerate because rational tangles of slope 0 or ∞ are ruled out by Definition 1.14, and rational tangles of slope ± 1 are never needed by the discussion preceding Definition 1.16.)

If A is not rational, then $A = A_1 + A_2$, a non-trivial sum of arborescent tangles. By the inductive hypothesis, some projection of each A_i can be represented by a truncated tree T_i. To visualize the sum, we would like to place each marking disk $D_i \subset \partial A_i$ into a simple position by a map of the 4-punctured sphere ∂A_i. It is well-known that the mapping class group of a 4-punctured sphere is generated by Dehn twists along horizontal and vertical curves [3, Theorem 4.5]; in other words, a sequence of twists along horizontal or vertical bands will move D_i into a simple position perpendicular to the projection plane. Thus we have added a (possibly empty) 4-string braid to the projection of A_i. By Lemma 4.8, this corresponds to adding a (possibly empty) interior branch U_i to the stump edge of T_i.

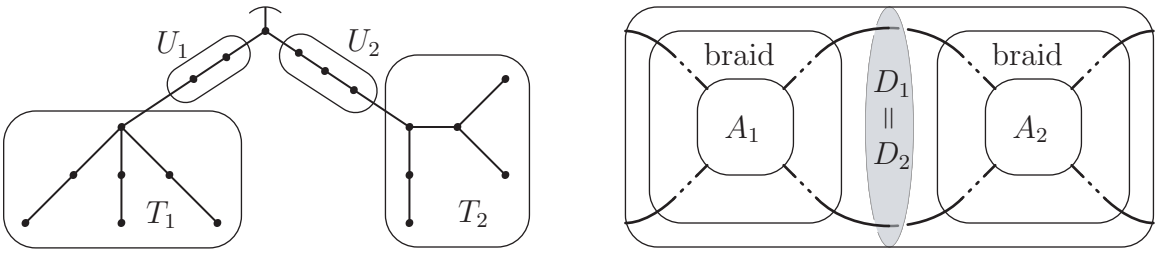

Figure 4.9: The inductive step in growing a tree to represent an arborescent tangle.

Now that D_1 and D_2 are clearly visible, we can glue A_1 to A_2 along these disks, as in Figure 1.4. In the language of this chapter, we are gluing A_1 and A_2 into bubbles along the same band, as in Figure 4.9. Thus, by joining $T_1 \cup U_1$ and $T_2 \cup U_2$ to the same node, we obtain a truncated tree representing A. □

Theorem 4.11. *Definitions 1.15 and 4.4 are equivalent. Every arborescent link K defined by a weighted tree T can be constructed from rational tangles; conversely, every link constructed from rational tangles can be represented by a weighted tree.*

Proof. Let K be an arborescent link as in Definition 4.4, constructed from a non-degenerate weighted tree T. By Lemma 4.8, every terminal branch of T corresponds to a rational tangle. Wherever two terminal branches U_1 and U_2 are adjacent at a node $v \in T$, the rational tangles B_1 and B_2 are connected along a disk perpendicular to the band of v, forming an arborescent tangle $B_1 + B_2$ as in Figure 1.4. (The restriction to non-degenerate trees rules out rational tangles of slope 0 or ∞, and thus the sum is non-trivial.) Continuing along T, we build a larger and larger arborescent tangle by gluing on rational tangles along braids. When we parse the last terminal branch of T, we glue the arborescent tangle to a rational tangle, forming an arborescent link that satisfies Definition 1.15.

For the converse, let K be an arborescent link as in Definition 1.15, obtained by gluing arborescent tangles A and A' along their entire boundaries. By Lemma 4.10, each of these tangles can be represented by a non-degenerate truncated tree. As above, the gluing map from ∂A to $\partial A'$ can be realized by a 4-string braid [3, Theorem 4.5]. Thus, by joining the two truncated trees along an interior branch representing the braid, we obtain a weighted tree T that represents the entire link K. □

4.3 Maximally alternating diagrams

Now that we have a correspondence between weighted trees and arborescent links as defined in the Introduction, we would like to simplify the link diagram in a way that will allow us to read off the length of each rational tangle. As a first step, we will bring the braids and rational tangles of $D_T(K)$ into an alternating projection.

Lemma 4.12. *Let $U \subset T$ be a branch. By flyping the bands of several vertices, as in Figure 4.7, we can ensure that the 4-braid corresponding to U has one strand that is free of crossings.*

Proof. In simplifying move (2), we have moved the weights of each vertex $v_i \subset U$ to one spot. Thus, in the band corresponding to v_i, only two strands pass through crossings. Proceeding from the outside in, from sphere S_0 toward S_n, we can flype the bands to ensure that, for example, every crossing in a horizontal band is on the left and every crossing in a vertical band is at the top. Then the bottom right strand of the braid will be free of crossings, as in Figure 4.10. □

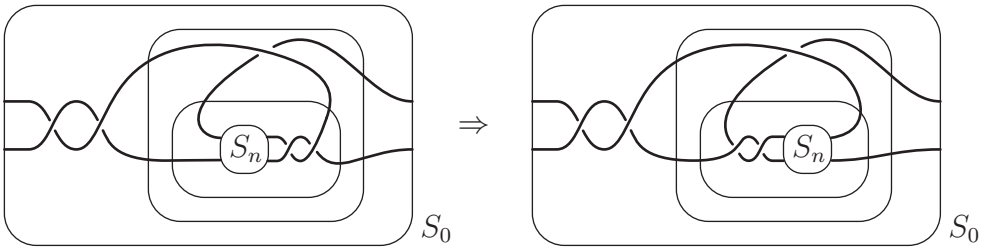

Figure 4.10: Every braid corresponding to a branch $U \subset T$ can be given a projection with one free strand. In the language of Theorem 4.13, the word of U is $\Omega = L^2R^{-1}L^2$.

Theorem 4.13. *Let $U \subset T$ be a branch containing vertices $v_1, \ldots, v_n$. Then the braid between spheres S_0 and S_n can be rearranged to have an alternating projection. This gives a new arborescent diagram $D_{T'}(K)$, corresponding to a tree T' in which U has been replaced by a branch U'.*

The relationship between T and T' is described at the end of the proof.

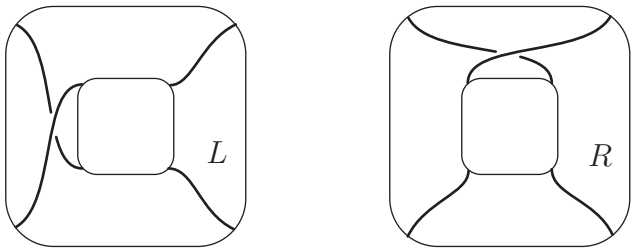

Figure 4.11: Braid moves L and R.

Proof. Assume, without loss of generality, that the braid of U has been positioned as in the right panel of Figure 4.10. That is: sphere S_0 lies outside the band of v_1, and the remaining bands are inserted into smaller and smaller bubbles; the band of v_1 is horizontal; and U has a free strand on the lower right. If a_i is the weight on vertex v_i, we may also assume that $a_1 > 0$.

We can represent a left-hand twist on the left side of a horizontal band by L and a right-hand twist on the top side of a vertical band by R. (See Figure 4.11; notice that both R and L encode positive crossings.) Then the braid between spheres S_0 and S_n is represented by the word

$$\Omega = \begin{cases} L^{a_1} R^{a_2} \cdots L^{a_n} & \text{if } n \text{ is odd,} \\ L^{a_1} R^{a_2} \cdots R^{a_n} & \text{if } n \text{ is even,} \end{cases}$$

read as we progress from the outside in. It is easy to check that the braid is alternating if and only if all the a_i have the same sign. Since we have assumed that $a_1 > 0$, our goal is to make all the a_i positive.

If the braid is not alternating, find the first sub-word of Ω of the form RL^{-1} or LR^{-1}. We will rotate the bubble B just inside this sub-word by a half-turn about a skewer of slope -1. (See Figure 4.12.) This has the effect of replacing RL^{-1} by L, or LR^{-1} by R, reducing the number of crossings in the braid by one. Inside bubble B, positive crossings at the top become negative crossings on the left, and vice versa. Thus the move replaces the rest of the word Ω according to the rule

$$L \leftrightarrow R^{-1}, \quad R \leftrightarrow L^{-1}.$$

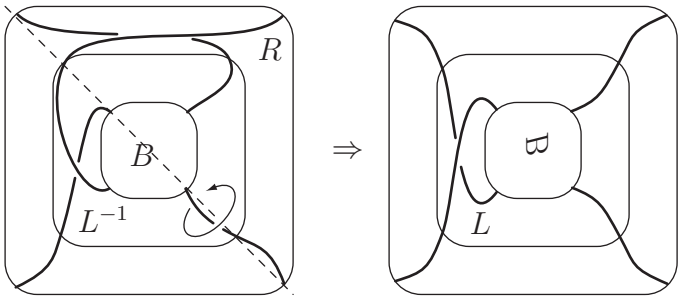

Figure 4.12: If a braid is not alternating, rotating about a line of slope -1 reduces the number of crossings.

Observe, as well, that this rotation preserves the property that the braid has a free strand on the lower right.

We can now continue scanning Ω for the next instance of RL^{-1} or LR^{-1}, and rotate the braid in the same fashion. Each iteration reduces the length of Ω by one letter, which corresponds to removing a crossing from the braid. Thus the process will terminate. In the end, we have replaced Ω by a word

$$\Omega' = \begin{cases} L^{b_1}R^{b_2}\cdots L^{b_m} & \text{if } m \text{ is odd,} \\ L^{b_1}R^{b_2}\cdots R^{b_m} & \text{if } m \text{ is even,} \end{cases}$$

where $b_2, \ldots, b_m$ are all positive and b_1 is non-negative. (The case $b_1 = 0$ occurs if Ω started with LR^{-1} and thus the first crossing of the modified braid is an R-crossing.) Since there are no sign changes, the modified braid is now alternating.

Let us explicitly describe the tree T' corresponding to the new link diagram. U has been replaced by a branch U', with vertex weights $b_1, \ldots, b_m$. The subtree of T that was adjacent to U at v_n may also be modified slightly, for each rotation interchanges horizontal and vertical bands. Therefore, if the total number of rotations is odd (i.e. if m has the opposite parity from n), all the vertex weights on the subtree inside sphere S_n will switch their sign. Because this switch occurs in unison, any braid that was previously made alternating will remain alternating. □

Remark 4.14. In case $b_1 = 0$, we will need to repeat simplifying move (1) to remove the redundant vertex v_1 of the new tree T'. (See Figure 4.6.)

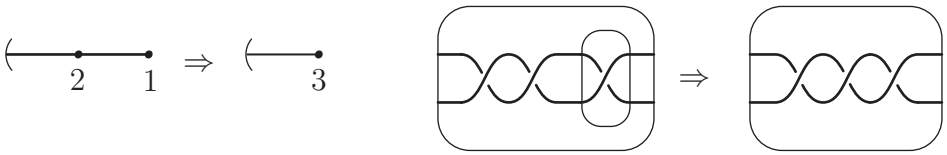

Figure 4.13: Adjusting the terminal vertex in a modified branch U'.

If U is a terminal braid, a tiny bit of fussing may also be required to ensure that the new tree T' is non-degenerate. That is, if $b_m = 1$, we need to move the final crossing to the band of v_{m-1}, as in Figure 4.13. Since $b_{m-1} \geq 1$ by construction, one move of this sort will suffice.

Remark 4.15. When U is a terminal branch, Conway's classification of rational tangles by slope [5] implies

$$\cfrac{1}{b_1 + \cfrac{1}{b_2 + \cfrac{1}{\ddots + \cfrac{1}{b_m}}}} = \cfrac{1}{a_1 + \cfrac{1}{a_2 + \cfrac{1}{\ddots + \cfrac{1}{a_n}}}} = s,$$

where s is the *slope* of the tangle. (See Figure 1.5.) The b_i give its shortest continued fraction expansion. If we define the *length* of U to be $\ell(U) = \sum_i |b_i|$, the length of a terminal branch will agree with the length of its rational tangle (see Definition 1.16).

Once we have simplified the branches of a weighted tree T, ensuring that they correspond to alternating braids, we can read off the length of each rational tangle in terms of weights on the tree. The next step is to look at the nodes.

Definition 4.16. Let v be a node of a weighted tree T. We say that v is an *alternating node* if the weights of v and every vertex adjacent to v all have the same sign. (For the purpose of this definition, 0 has the same sign as any positive or negative number.)

The braids and rational tangles that meet at the band of an alternating node v combine to form an alternating sub-diagram of $D_T(K)$.

Definition 4.17. Let $D_T(K)$ be an arborescent diagram represented by a weighted tree T. Suppose that in every branch of T, all vertex weights are non-zero and have the same sign. Let $v \in T$ be a node adjacent to edges $e_1, \ldots, e_k$. If e_i is part of a terminal branch B_i, we define $\ell_i = \ell(B_i)$ to be the number of crossings in the corresponding tangle; otherwise, define $\ell_i = \infty$. Suppose that, for each i,

$$1 > \begin{cases} \dfrac{1}{l_i} + \dfrac{1}{l_{i+1}} & \text{if } v \text{ is an alternating node,} \\ \dfrac{1}{l_{i-1}} + \dfrac{1}{l_i} + \dfrac{1}{l_{i+1}} & \text{if } v \text{ is a non-alternating node.} \end{cases}$$

If this condition is satisfied for every node of T, the diagram $D_T(K)$ is called *balanced*.

For alternating nodes, Definition 4.17 is more inclusive than Definition 1.17. Theorem 1.18, which guarantees that the complement of a balanced link has an angled polyhedral decomposition, will hold with this looser definition. In fact, as the next theorem shows, we can turn many nodes of T into alternating nodes.

Theorem 4.18. *Let $D_T(K)$ be a diagram of an arborescent link. We can modify the diagram until the corresponding weighted tree T satisfies the following conditions:*

(1) *In every branch of T, all vertex weights are non-zero and have the same sign.*

(2) *Every node of T is either alternating or has weight 0 (or both).*

Furthermore, this alteration preserves the length of each terminal branch. Thus a balanced diagram will remain balanced.

Proof. In the simplification process of Section 4.1, we removed all divalent vertices of weight 0. This, combined with Theorem 4.13, guarantees condition (1). It remains to prove that we can satisfy condition (2) without breaking (1).

Suppose that some node $v \in T$ fails condition (2): it is non-alternating and (without loss of generality) $w(v) > 0$. We will use the property that v is not alternating to simplify the diagram. At each step, we will switch the weight of a neighboring vertex v_i (and, if v_i is part of a branch, on the entire branch) from negative to positive,

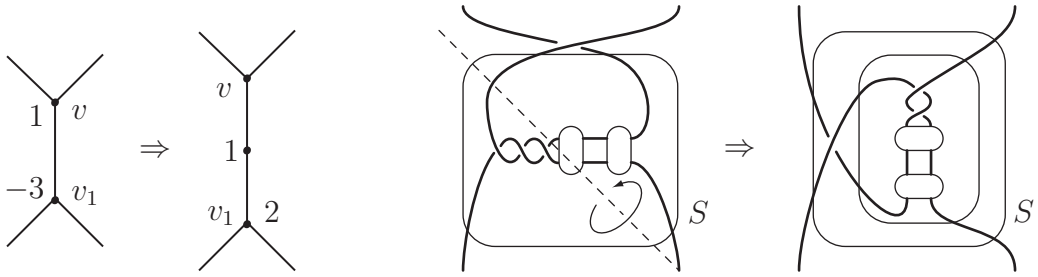

Figure 4.14: When adjacent nodes have opposite weights, rotating about a diagonal axis removes a crossing from $D_T(K)$.

reduce $w(v)$ by one, and remove a crossing from $D_T(K)$. This procedure will terminate when one of two things happens: either all adjacent vertices have non-negative weights, or $w(v)$ is reduced all the way to zero. In either case, v will satisfy condition (2), and we can move on to another node.

Suppose, therefore, that $w(v) > 0$ but $w(v_1) < 0$ for a neighboring vertex v_1. Suppose, without loss of generality, that the band of v is vertical and the band of v_1 horizontal. By performing flypes on each band, as in Figure 4.10, we can also assume that the positive crossings at v and the negative crossings at v_1 are both located next to the sphere S that separates the two bands. In other words, we have the situation depicted in Figure 4.12, represented by a word RL^{-1}.

If the vertex v_1 is part of a branch U, we proceed exactly as in the proof of Theorem 4.13. That is: we rotate the bubble just interior to the first negative L-crossing about a diagonal skewer, removing a crossing from the diagram and replacing RL^{-1} by L. Because all the vertices of U start out with negative weights, there are no more sign changes in its word Ω. Thus, after the single diagonal flip, U is replaced by a branch U' with all positive weights, and its length remains unchanged. (The crossing removed during the flip came from the band of v, not from the braid.)

If the vertex v_1 is not part of a branch, it must be another node. Although we have not yet addressed this situation, the procedure is the same. In the diagram, we rotate the band of v_1 about a diagonal skewer, removing a crossing. The alteration to the tree T is a bit more complex. The node v_1 now corresponds to a vertical band, and is now two steps away from v. In between v and v_1 is a new divalent vertex with

weight 1, corresponding to the single L-crossing created during the flip. This new vertex forms a new, very short, interior branch. The weight $w(v)$ is reduced by one, while the weight $w(v_1)$ switches sign and is then reduced by one. (On the sub-tree that starts at v_1, all the weights will also switch signs, just as in Theorem 4.13). The net effect is that $D_T(K)$ has one less crossing, at the expense of $w(v)$, and one more vertex adjacent to v has positive weight. See Figure 4.14.

Therefore, whether or not v_1 is part of a branch, we remove a crossing from the band of v, as claimed. We can repeat this step as many times as needed, continuing to remove crossings from $D_T(K)$ until every node satisfies condition (2). □

4.4 Unloaded nodes and expanded trees

For the constructions of the next chapter, we will need the bands corresponding to nodes of T to be free of crossings. Thus for each node $v \in T$ with $w(v) \neq 0$ (which, by Theorem 4.18, must be an alternating node) we will "unload" the weight $w(v)$ into one of the branches meeting at v, in a partial undoing of simplifying move (1). See Figure 4.15 for a graphical description of this procedure.

In unloading the weight of a node v, we can typically choose from multiple branches. We make the choice as follows. If every edge emanating from v connects to another node, we unload the weight $w(v)$ by creating a new branch between v and v', leaving the diagram unchanged as in Figure 4.15. If v is adjacent to one or more interior branches, we move $w(v)$ into any one of those branches. The remaining

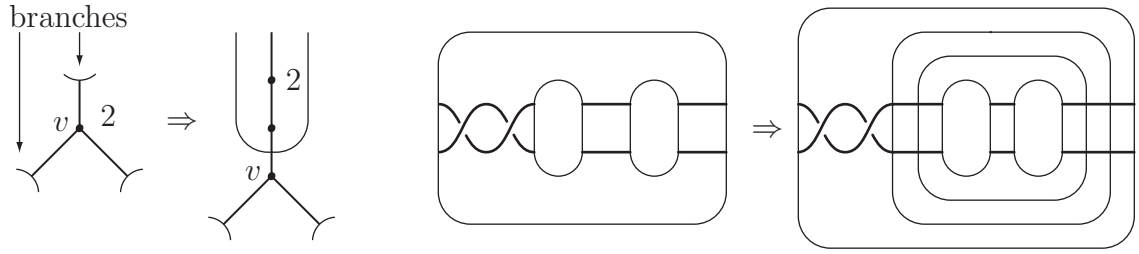

Figure 4.15: We move the weight of a node v into a neighboring branch, creating a "dummy vertex" in between. As a result, the band of v has no crossings.

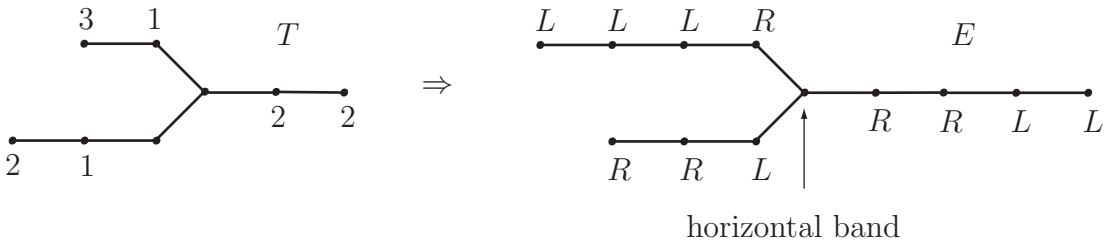

Figure 4.16: Turning a weighted tree T into an expanded tree E.

case is when every edge at v is part of a terminal branch, i.e., K is a Montesinos link. In this case, we move $w(v)$ into the shortest terminal branch.

In each case, the crossings on the band of v either add to an alternating braid, or start a new alternating braid. Thus every alternating node of T remains alternating.

Definition 4.19. Let T be a weighted tree in which every node has zero weight. We will use T to construct an *expanded tree* E, as follows. Every node of T defines a node of E, of the same valence. Every branch $U \subset T$ also defines a branch $B \subset E$, which has one vertex for each crossing in the braid of U. (See Figure 4.16.)

We think of a node $v \in E$ as *horizontal* or *vertical* according as the corresponding band of $D_T(K)$ is horizontal or vertical. Vertices of E corresponding to positive or negative crossings in a vertical band are labeled by R or R^{-1}, respectively. Similarly, vertices of E corresponding to crossings in a horizontal band are labeled L or L^{-1}.

Note that, just as with weighted trees, the nodes of E correspond to bands with no crossings. Just as with weighted trees, the interior and terminal branches of an expanded tree E correspond to braids and rational tangles, respectively.

If we concatenate the vertex labels of a branch $B \subset E$ into a single word in L and R, we get exactly the word Ω that was described in the proof of Theorem 4.13. The notion of length for rational tangles also extends naturally to branches of E. Thus the *length* $\ell(B)$ of a branch B is equal to the number of vertices in B, the number of letters in Ω, and the number of crossings in the corresponding braid.

Definition 4.20. Let v be a node of an expanded tree E. We say that v is *alternating* if every branch vertex adjacent to v is labeled with the same positive or negative power

of L or R. (The node in Figure 4.16 is an example.) Alternating nodes of E come from alternating nodes of T.

Definition 4.21. Let e be an edge of an expanded tree E, connecting a node v to a branch vertex v_1. We call e an *exceptional* edge if the band of v is vertical and v_1 is labeled $R^{\pm 1}$, or if the band of v is horizontal and v_1 is labeled $L^{\pm 1}$. A node $v \in E$ adjacent to an exceptional edge is called an *exceptional node*.

Exceptional edges are created when the weight of an alternating node is unloaded into a branch. In Figure 4.16, the lower left branch ends in an exceptional edge.

We can now summarize the results of this chapter in terms of expanded trees.

Theorem 4.22. *Every arborescent link $K \subset S^3$ can be represented by an expanded tree E such that*

(1) *in any branch $B \subset E$, the vertex labels are all positive or all negative,*

(2) *every exceptional node $v \in E$ is alternating,*

(3) *every exceptional node $v \in E$ is adjacent to one exceptional edge, and*

(4) *if K is not a Montesinos link, all exceptional edges belong to interior branches.*

Proof. Conclusion (1) is a restatement of Theorem 4.13. Conclusion (2) follows from Theorem 4.18, because the only nodes of T whose weights are pushed into a branch are alternating nodes. Conclusions (3) and (4) are a consequence of the choices we make in unloading the weight of a node: all the crossings on the band of a node are pushed into a single branch, which is interior when K is not a Montesinos link. □

Remark 4.23. It is worth pointing out that some information is lost in constructing an expanded tree. Unlike weighted trees, where the location of weights at a vertex encodes the location of crossings in a band, the vertex labels L or R do not preserve this information. As a result, an expanded tree is not enough to reconstruct a link, because it does not distinguish between certain mutants. However, we will see in Chapter 5 that the vertex labels of E carry just enough data to aid us in subdividing the link complement. In Chapter 6, we will see that when E represents a balanced diagram, its edges correspond to parameters for angle structures.

Chapter 5

Arborescent Link Complements

The last chapter described a strong correspondence between a weighted tree T and an arborescent diagram $D_T(K)$. We can now use the pieces of T to construct a topological decomposition of the link complement. Recall, in particular, that in Definition 4.7, we subdivided a weighted tree into *terminal branches* and *interior branches* that come together at *nodes* (vertices of valence at least 3). The corresponding subdivision of the expanded tree E allows us to cut up an arborescent link complement into three types of pieces:

(1) *Product regions.* Each branch of an expanded tree E describes a braid in $D_T(K)$. The complement of such a braid is a product region that we will subdivide into layers of tetrahedra. The word Ω labeling the branch controls the combinatorics of the subdivision.

(2) *Clasps*, corresponding to the last two vertices of each terminal branch in E. A clasp will actually not need any 3-dimensional pieces, and will be constructed by folding a pleated surface in a particular fashion.

(3) *Prisms*, as in Example 2.6. Each band corresponding to a node of E will give two prisms in the link complement.

This chapter is devoted to describing these pieces in some detail.

5.1 Product regions

Recall, from the last chapter, the correspondence between branches of an expanded tree E and product regions in the link complement. We will subdivide each such product region into layers of tetrahedra. To describe this subdivision, we need a better understanding of triangulations of a 4-punctured sphere.

Definition 5.1. Let S be a 4-times punctured sphere. An *arc pair* $\gamma \subset S$ consists of two disjoint, properly embedded arcs γ_1 and γ_2, such that γ_1 connects two punctures of S and γ_2 connects the remaining two punctures of S. A *slope* on S is an isotopy class of arc pairs.

To visualize slopes, it helps to picture S as a pillowcase surrounding the unit square in $\mathbb{R}^2$, with punctures at the corners. Any arc pair $\gamma \subset S$ can then be straightened so that its intersections with the front of the pillow have a well-defined Euclidean slope. Thus slopes on S are in 1-to-1 correspondence with elements of $\overline{\mathbb{Q}} = \mathbb{Q} \cup \{\infty\}$.

It is easy to check that at most 3 slopes can be disjoint on a 4-punctured sphere; such a choice of 3 disjoint arc pairs subdivides S into ideal triangles. (See Figure 5.1.) Ideal triangulations of S can be neatly represented by the *Farey complex* $\mathcal{F}$, shown in Figure 5.2. Vertices of $\mathcal{F}$ correspond to slopes, edges of $\mathcal{F}$ to disjoint slopes, and triangles to ideal triangulations of S.

Now, let B be an interior branch of E. By Theorem 4.13, we may assume that E defines an alternating braid, whose complement is a product region between spheres S_0 and S_c (where $c = \ell(B)$ equals the number of vertices in B and the number of

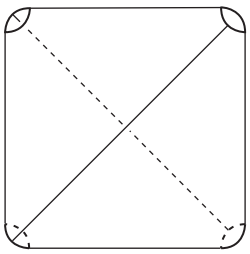

Figure 5.1: Arcs of slope 0, 1, and ∞ give an ideal triangulation of a 4-punctured sphere S.

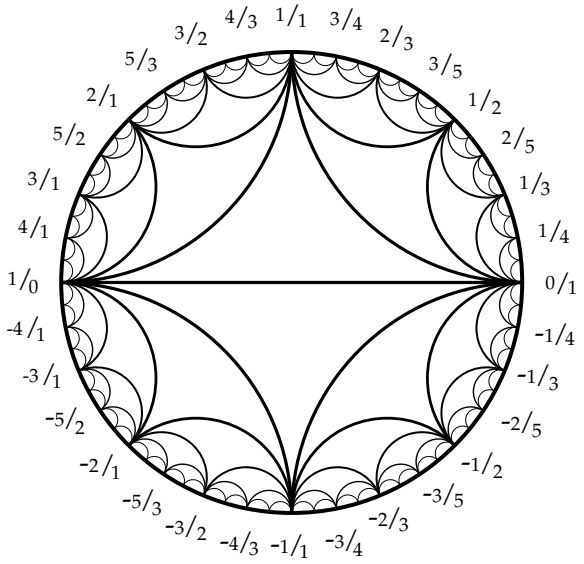

Figure 5.2: The Farey graph $\mathcal{F}$ of a 4-punctured sphere (graphic by Allen Hatcher).

crossings in the braid). Sphere S_0, exterior to the braid, intersects the projection plane in horizontal and vertical edges whose slopes we will identify with 0 and ∞, respectively. Once we have fixed this identification, we can say that horizontal and vertical edges on S_c have slope $\varphi(0)$ and $\varphi(\infty)$, where $\varphi : \overline{\mathbb{Q}} \to \overline{\mathbb{Q}}$ is the map on slopes induced by isotopy through the entire product region.

Notice that slopes 0 and ∞ are disjoint, and thus 0 and ∞ are connected by an edge $b \subset \mathcal{F}$. Similarly, $\varphi(0)$ and $\varphi(\infty)$ are connected by an edge $t \subset \mathcal{F}$.

Lemma 5.2. *Let Ω be the word in R and L obtained by concatenating the vertex labels of B. Then the map $\varphi : \overline{\mathbb{Q}} \to \overline{\mathbb{Q}}$ can be identified with Ω. In particular, each letter Ω_i of Ω corresponds to a Farey triangle τ_i between b and t. If $\Omega_i = R$, the path from b to t takes a right turn across τ_i; if $\Omega_i = L$, the path from b to t takes a left turn across τ_i.*

Proof. Recall that a vertex label of L corresponds to a positive crossing in a horizontal band, and a vertex label of R corresponds to a positive crossing in a vertical band. To prove this lemma, it suffices to identify the action of L and R on the Farey graph $\mathcal{F}$. By sliding horizontal and vertical arcs along the braid from the inner sphere (closer

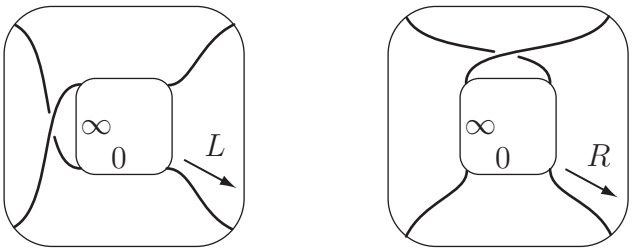

Figure 5.3: Braid move L takes slopes 0 and ∞ to 1 and ∞, respectively. R takes slopes 0 and ∞ to 0 and 1, respectively.

to S_c) to the outer sphere (closer to S_0), we can see in Figure 5.3 that

$$L(0) = 1,\ L(\infty) = \infty, \quad R(0) = 0,\ R(\infty) = 1.$$

This determines their action on all slopes: L acts on the Farey graph by a left turn fixing ∞, and R acts by a right turn fixing 0.

If we read the word Ω from right to left, the actions of L and R are combined (in the usual sense of function composition) to give the map φ from slopes on S_c to slopes on S_0. On the other hand, if we read Ω from left to right, it tells us the sequence of turns across Farey triangles that gets from b to t. That is, if the i^{th} letter of Ω is $\Omega_i = L$, we take a left turn across τ_i; if $\Omega_i = R$, we take a right turn across τ_i.

The key consequence of the braid of B being alternating is that the path described by Ω is *non-backtracking*: in taking right and left turns while going forward from b to t, we only cross each triangle τ_i once. Thus the number of these triangles is exactly $c = \ell(B)$, the number of crossings in the braid. □

Example 5.3. If $\Omega = R^2L^2$, two right turns followed by two left turns take b to the edge t connecting slopes $\varphi(0) = \frac{2}{5}$ and $\varphi(\infty) = \frac{1}{2}$. (See Figure 5.4.)

Let $b = e_0, e_1, \ldots, e_c = t$ be the sequence of Farey edges interpolating between slopes on S_0 and slopes on S_c. Each e_i, $1 \leq i \leq c-1$, can be thought of as a diagonal exchange between τ_i and τ_{i+1}. We will use these diagonal exchanges to construct tetrahedra.

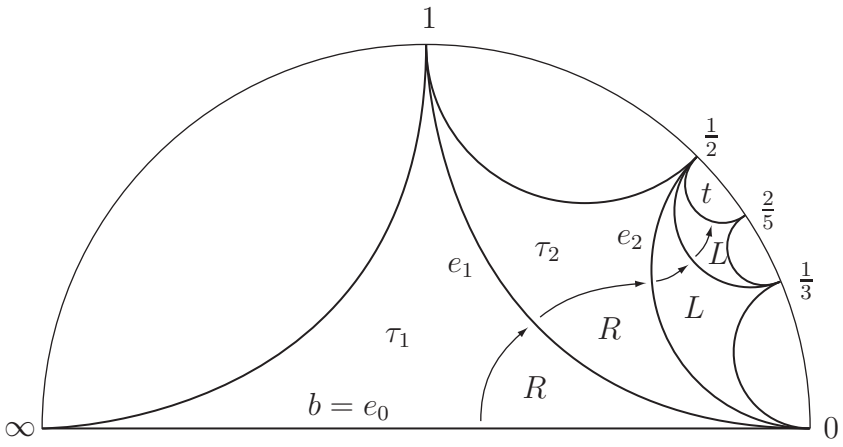

Figure 5.4: When $\varphi = \Omega = R^2L^2$, we can get from b to t with two right turns followed by two left turns.

Definition 5.4. For each e_i, $1 \leq i \leq c-1$, construct the product region $S\times[0,1]$, where S is a 4-punctured sphere. Give the bottom surface $S\times 0$ the triangulation τ_i, and the top surface $S\times 1$ the triangulation τ_{i+1}. Now, for each of the two shared slopes corresponding to endpoints of e_i, pick an arc pair γ of that slope and collapse $\gamma\times[0,1]$ to $\gamma\times 0$. The result is a union of two tetrahedra, denoted $\Delta_i = \Delta(e_i)$, connected across two arc pairs whose slopes are the endpoints of e_i. The boundary of this layer of tetrahedra consists of two *pleated surfaces*, S_i on the bottom and S_{i+1} on top, each homotopic to S. (See Figure 5.5.)

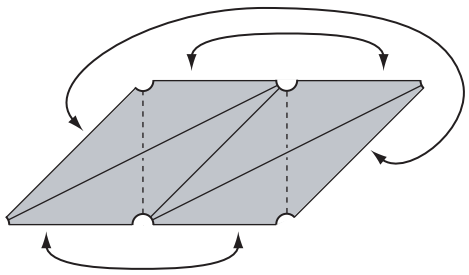

Figure 5.5: The layer of tetrahedra corresponding to the Farey edge e_1 from 0 to 1 can be thought of as exchanging diagonals of slope ∞ for diagonals with slope $1/2$.

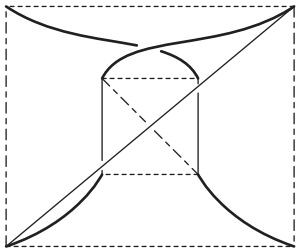

Figure 5.6: The pleated surface corresponding to a crossing in a braid can be seen as living either right before or right after the crossing. In this figure, arcs with the same dashing pattern have the same slope.

We can glue the top of tetrahedron layer Δ_i to the bottom of Δ_{i+1}, since their triangulations come from the same Farey triangle τ_{i+1}. In this way, we obtain a stack of tetrahedron layers, progressively interchanging diagonals. Since the bottom of Δ_1 is identified with S_0 and the top of Δ_{c-1} with S_c, we have a triangulation of the product region. Note that this stack of tetrahedra is homeomorphic to $S\times[0,1]$ if and only if b and t have disjoint endpoints. If b and t share an endpoint, all the triangulations share the corresponding slope, and the entire stack of tetrahedra is pinched along the corresponding arc pair.

Remark 5.5. The correspondence of Lemma 5.2 identifies crossings in the braid with pleated surfaces between S_0 and S_c. The pleated surface S_i can be concretely visualized as living near the i^{th} crossing in the braid; the edges of its triangulation are the ones that look vertical and horizontal in the projection plane immediately before and after the crossing. Figure 5.6 shows two views of the same pleated surface.

Remark 5.6. It follows from the last remark that the Farey triangles τ_1 and τ_c force a choice of diagonal for spheres S_0 and S_c, respectively. That is, S_0 will receive a diagonal of slope 1 or -1 depending on whether the letters of Ω have positive or negative powers. Similarly, the triangulation on S_c, which already contains two slopes coming from the vertices of t, will receive a third slope from the third vertex of τ_c.

Remark 5.7. If the braid has only one crossing, i.e. $c = 1$, Definition 5.4 results in no tetrahedra at all. In this scenario, we glue S_0 directly to S_1. (In Figure 5.6, the

outer sphere would be S_0 and the inner sphere S_1.) Notice that, just as in Remark 5.6, the direction of the single crossing still forces a choice of diagonals, completing a triangulation of the identified sphere.

5.2 Clasps

We have just described how to triangulate the product region corresponding to an interior branch of E. For terminal branches corresponding to rational tangles, the procedure is very similar. We once again have a sphere S_0, exterior to the tangle, and a sphere S_c, which bounds a bubble containing no crossings. These spheres again give rise to Farey edges b and t. The intermediate triangles $\tau_1, \ldots, \tau_c$ define a sequence of triangulations of a 4-punctured sphere, giving rise to layered stack of tetrahedra.

The one important difference is that in constructing the triangulation corresponding to a terminal branch, we stop at triangle τ_{c-1}. That is, the final layer of tetrahedra is Δ_{c-2}, whose top is the pleated surface S_{c-1}.

Definition 5.8. On the pleated surface $S_{\text{cl}} = S_{c-1}$, let the *peripheral edges* be the two edges whose slope is the Farey vertex of τ_{c-1} opposite edge e_{c-1}.

Following Remark 5.5, we can picture S_{cl} as lying near the penultimate crossing of the braid of B. Its peripheral edges are the two edges that would be replaced if we were to glue on the tetrahedron layer Δ_{c-1}: that is, the peripheral edges run parallel to the band containing the last pair of crossings. (See Figure 5.7.)

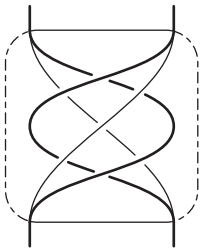

Figure 5.7: The peripheral edges of pleated surface S_{cl} are dashed. Notice that the remaining four edges are all isotopic.

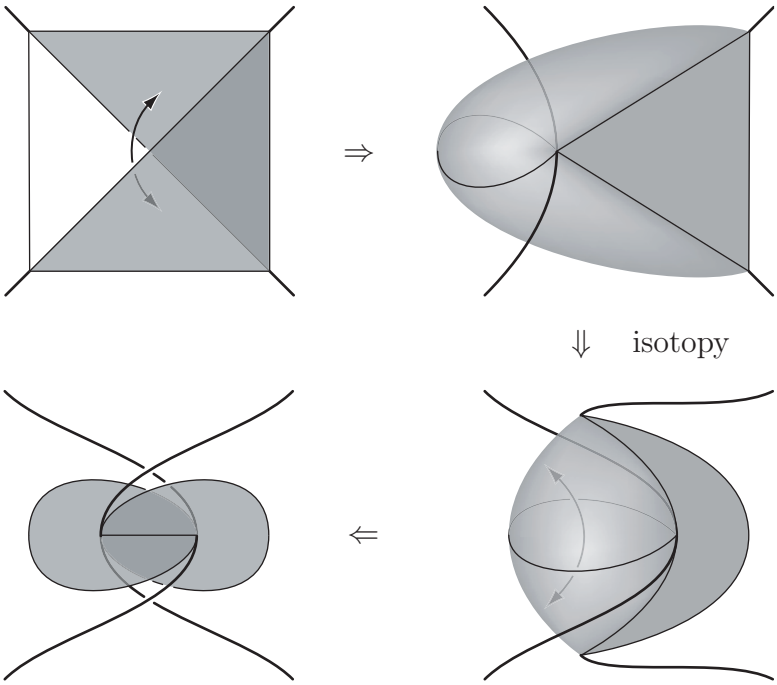

Figure 5.8: Folding the pleated surface S_{cl}.

Instead of gluing another tetrahedron layer to S_{cl}, we will identify its faces in pairs by folding along the two peripheral edges. Figure 5.8 shows that this creates exactly the desired effect of connecting the strands of the braid in pairs while twisting along the band, producing the last pair of crossings.

The four non-peripheral edges of S_{cl} get identified to a single edge, called the *core* of the rational tangle. The two faces that remain after folding meet at the core of the tangle, in the manner of a *clasp*. Outside the clasp, we have a standard product region, subdivided into layers of tetrahedra all the way out to sphere S_0.

Remark 5.9. Note that in triangulating a rational tangle with c crossings, we have used $c-2$ pairs of tetrahedra. Thus if a rational tangle contains only two crossings, no tetrahedra are needed at all: we simply identify the 4-punctured sphere S_0 with S_{cl}, and fold as above. (The terminal branches of length 2 that correspond to these tangles will be called *short*.) Because our weighted trees are non-degenerate (see Definition 4.2), we do not need to deal with rational tangles containing a single crossing.

Remark 5.10. This way of triangulating rational tangles is not new. Sakuma and Weeks [31] use exactly this method to construct the (presumably canonical) triangulations of two-bridge link complements. To get a two-bridge link, one needs to cap off both boundary spheres of a 4-string braid, corresponding to clasps with the triangulations of τ_2 on one end and τ_{c-1} on the other end.

In constructing layered triangulations from the branches of an expanded tree E, we have made use of a number of objects: the Farey complex $\mathcal{F}$, the word Ω, and of course the arborescent diagram $D_T(K)$. Their correspondence to the pieces of the link complement is summarized in Table 5.9.

Edge of expanded tree E	Non-node vertex of E	Node of E
Farey edge e_i	Farey triangle τ_i	
Tetrahedron layer Δ_i	Pleated surface S_i	Pair of prisms
Pleating angle w_i	Letter $\Omega_i = L$	
	Crossing in $D_T(K)$	

Table 5.9: The three correspondence classes of objects used to construct a polyhedral decomposition of the link complement.

5.3 Prisms

Let v be a node of an expanded tree, i.e. a vertex of valence $k \geq 3$. Recall from the last chapter that v corresponds to a band in the projection diagram, intersecting bubbles $B_1, \ldots, B_k$. (See Figure 4.3.) In Section 4.4, we have ensured that the band contains no crossings.

We will subdivide the complement of the band and the bubbles into two ideal polyhedra, as follows. First, cut the picture along the projection plane. This yields two identical pieces, each homeomorphic to a 3-ball; we focus our attention on the ball in front of the projection plane. Next, contract each rectangular piece of the ribbon to a segment. This identifies each pair of isotopic edges on adjacent bubbles to a single edge, and shrinks strands of K to ideal vertices. (See Figure 5.10.)

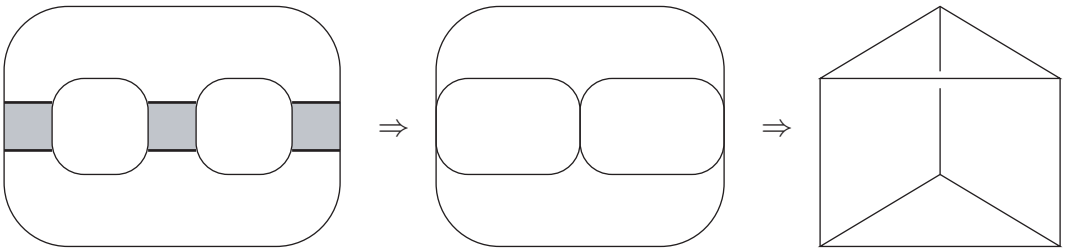

Figure 5.10: Constructing a prism from the band of a node.

The result is an ideal polyhedron P that has rectangular faces $F_1, \ldots, F_k$ (one for the front side of each bubble) and two k-gon faces (one for each region of the projection plane in the complement of the band). Thus P is a *prism*, as defined in Example 2.6. (In [32, Section 6.8], Thurston calls the this type of polyhedron a *drum*.) We will call the k-gon faces the *top* and *bottom* faces of P, and the rectangles bordering on bubbles the *lateral* faces of P. The ideal polyhedron behind the projection plane is an identical prism P', whose lateral faces border on the back sides of the bubbles.

Prisms P and P' are glued to each other along the top and bottom faces, which came from regions of the projection plane. They also share the same lateral edges, but the interiors of their lateral faces are disjoint. Matching lateral rectangles of P and P' are glued along their edges to form 4-times punctured spheres $S_1, \ldots, S_k$ bounding bubbles $B_1, \ldots, B_k$. These spheres connect the pair of prisms to other pieces of the link complement.

Recall that the band of v has one bubble B_i for each adjacent edge e_i. If e_i connects to another node, sphere S_i connects $P \cup P'$ to another pair of prisms. If e_i is the stump edge at the end of a branch, S_i is a pleated surface bounding a product region. In the latter case, the triangulation of the product region will actually force a choice of diagonals on the rectangular faces $F_i \subset P$ and $F_i' \subset P'$, by Remark 5.6.

Our goal in the next chapter will be to give convex angles for the polyhedral decomposition of the link complement. In particular, the individual pieces will need to be angled polyhedra, as in Definition 2.3. The dihedral angles of a single tetrahedron are easy to understand and parametrize; for prisms, it takes a little work to ensure that a given choice of dihedral angles produces a convex hyperbolic polyhedron.

Definition 5.11. A polyhedron P that is obtained from a prism by adding diagonals to some lateral faces is called a *bent-faced prism*. If no diagonals have been added, we may call P a *flat-faced prism* for emphasis.

If P is a flat-faced prism, we will assign dihedral angles as in Figure 2.1. That is, the top and bottom edges of every lateral face F_i are labeled with the same internal angle $\alpha_i \in (0, \pi)$. So long as $\alpha_i + \alpha_{i+1} < \pi$, we can choose internal angles $\theta_{i,i+1} = \pi - \alpha_i - \alpha_{i+1}$ for the lateral edges.

Lemma 5.12. *Let P be a flat-faced prism, with angles assigned as above. Then P is an angled polyhedron if and only if*

(1) $\alpha_i + \alpha_j < \pi$ *for all* $i \neq j$ *and*

(2) $\sum_{i=1}^{k} \alpha_i > \pi$.

Proof. We have ensured by construction that all dihedral angles are positive, and that the external angles around each ideal vertex sum to 2π. Thus P will be an angled polyhedron if and only if it satisfies the last condition of Definition 2.3:

($*$) for any normal curve $\gamma \subset \partial P$ that does not encircle a single vertex, $\sum_{\gamma} \epsilon_i > 2\pi$.

Consider such a curve γ, and let m be the number of intersections between γ and the top and bottom faces of P.

Case 0: $m = 0$. If γ is disjoint from the top and bottom faces, it must be a belt curve, as in Figure 5.11(a). Then γ picks up the external angles from all the lateral edges. For each lateral edge, the external angle is $\pi - \theta_{i,i+1} = \alpha_i + \alpha_{i+1}$, so the angle sum is greater than 2π precisely if and only if $\sum_{i=1}^{n} \alpha_i > \pi$.

Case 1: $m = 1$. Suppose, without loss of generality, that γ intersects the top face, entering it at the edge e_1 and exiting at edge e_j. Since γ is disjoint from the bottom face, it must complete its loop through the lateral faces $F_1, \ldots, F_j$. Thus γ encircles two or more vertices of the top face, as well as the edges $e_2, \ldots, e_{j-1}$ that connect them. (See Figure 5.11(b).)

Consider the first vertex $v_{1,2}$ inside γ, and notice that γ intersects every edge into that vertex except e_2. Because the external angles around $v_{1,2}$ must sum to 2π, and

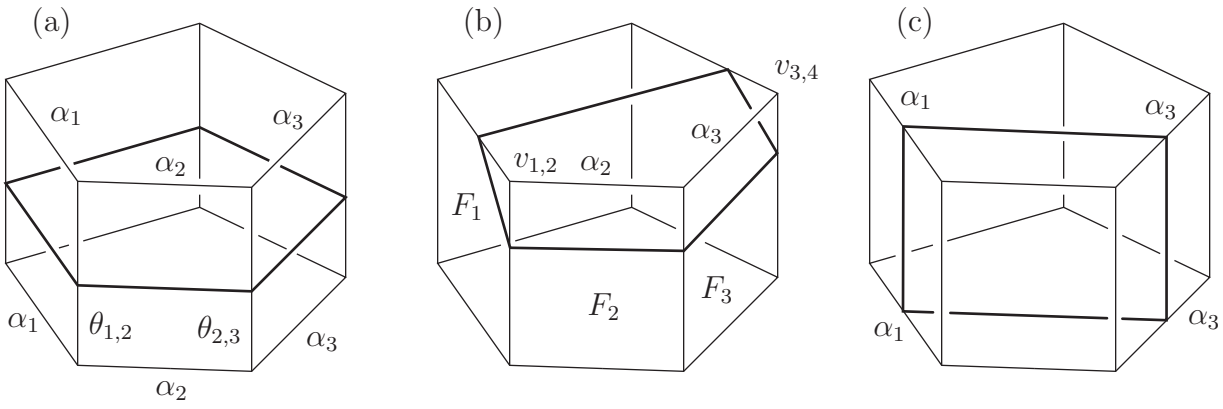

Figure 5.11: Normal curves in an angled prism.

$\alpha_2 < \pi < 2\pi - \alpha_2$, pulling γ to the other side of the vertex will reduce its exterior angle sum. We can continue shortening γ in this fashion until it encircles the last vertex $v_{j-1,j}$, around which the external angles sum to 2π by construction. Thus the original angle sum of γ was greater than 2π, and $(*)$ is satisfied automatically.

Case 2: $m \geq 2$. Let us remove from γ its intersections with the top and bottom faces. This leaves arcs $\gamma_1, \ldots, \gamma_m$, ordered along an orientation of γ. If any γ_i runs from the top face back to the top face (or bottom to bottom), then its exterior angle sum is already at least 2π, by Case 1. Thus we may assume that each of the γ_i runs from the top to the bottom (or vice versa).

We claim that if γ_1 ends in lateral face F_i, its external angle sum is at least $2\pi - 2\alpha_i$. If γ_1 only passes through F_i, this expression gives the exact value. But if γ_1 starts in some other face, it will intersect all but one edge around some vertex of that face. Thus pulling the arc γ_1 into F_i will reduce the external angle sum at each step, as in Case 1.

Now, consider the next arc γ_2. Because γ intersects each edge of P at most once, γ_2 starts in lateral face F_j, with $j \neq i$. By the same argument as above, the external angle sum of γ_2 is at least $2\pi - 2\alpha_j$. Because we have assumed that $\alpha_i + \alpha_j < \pi$, these sums add up to more than 2π.

Conversely, for any choice of $i \neq j$, a curve γ as in Figure 5.11(c) will have its external angle sum equal to $4\pi - 2\alpha_i - 2\alpha_j$, so condition (1) is necessary. □

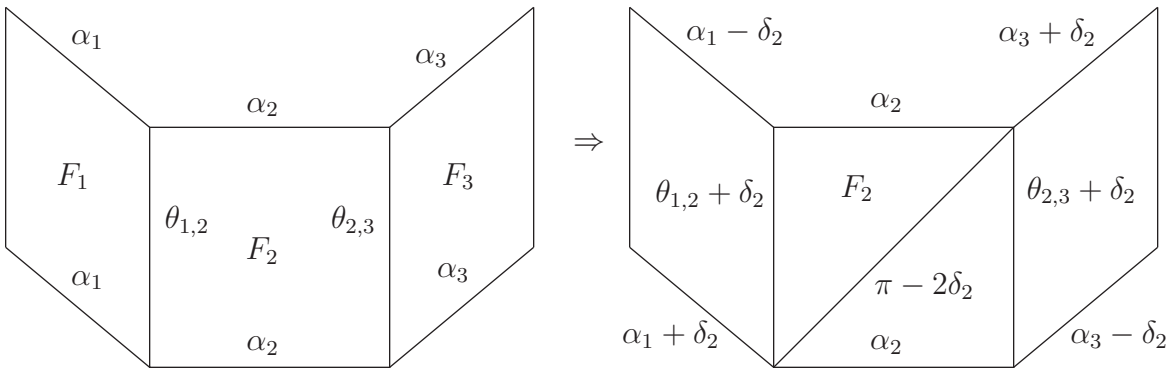

Figure 5.12: Bending a lateral face of a prism along a diagonal changes the angles of neighboring edges.

Now that we understand the dihedral angles of flat-faced prisms, we can bend the diagonals one by one. Suppose a prism P has been assigned dihedral angles, as above. We will choose a diagonal for lateral face F_i and give it dihedral angle $\pi - 2\delta_i$, where $\delta_i \in [0, \pi/2)$. We will then adjust the angles on the lateral edges of F_i, as well as the top and bottom edges of F_{i-1} and F_{i+1}. Edges sharing a vertex with the new diagonal increase their angle by δ_i, and ones that do not share a vertex with the diagonal decrease their angle by δ_i. (See Figure 5.12.)

We can then repeat this procedure, bending some other lateral face F_j along a diagonal. Each time, the change of angles is local: the only edges whose dihedral angles change are the lateral edges of F_j, the top and bottom edges of $F_{j\pm1}$, and (of course) the new diagonal. When we have bent all the faces as needed, the top edge of a face F_i will have a dihedral angle of the form $\alpha_i \pm \delta_{i-1} \pm \delta_{i+1}$, counting indices mod k. (The signs of $\pm\delta_{i-1}$ and $\pm\delta_{i+1}$ will be the different if and only if the corresponding branches have vertex labels with the same positive or negative powers of R and L, or equivalently, if and only if the diagonals of F_{i-1} and F_{i+1} have the same slope on the prism.) Note also that flat-faced prisms correspond to choosing $\delta_i = 0$ for all i.

Remark 5.13. By Remark 5.6, the diagonals of $F_i \subset P$ and $F_i' \subset P'$ have the same slope. However, relative to the top and bottom of those prisms, they are oriented in opposite directions, as in Figure 5.1. (In other words, bent-faced prisms P and P'

are symmetric by a rotation around the core of the band, which interchanges top and bottom edges.) Thus if the top edge of F_i has dihedral angle $\alpha_i \pm \delta_{i-1} \pm \delta_{i+1}$, the top edge of F_i' has dihedral angle $\alpha_i \mp \delta_{i-1} \mp \delta_{i+1}$.

This apparent complication has the fortuitous consequence that the top edges of F_i and F_i', which are glued to the same edge in a product region, have a total dihedral angle of $2\alpha_i$ (and similarly with the bottom edges).

Theorem 5.14. *Let P be a bent-faced prism, with angles labeled as above. Then P is an angled polyhedron if and only if*

(1) $(\alpha_i - \delta_i) + (\alpha_j - \delta_j) < \pi$ *for all* $i \neq j$,

(2) $\sum_{i=1}^{k} \alpha_i > \pi$,

(3) $(\alpha_i - \delta_i) + (\alpha_{i+1} - \delta_{i+1}) > 0$,

(4) $|\delta_{i-1} - \delta_{i+1}| < \alpha_i < \pi - |\delta_{i-1} - \delta_{i+1}|$, *and*

(5) $\delta_{i-1} + \delta_{i+1} < \alpha_i < \pi - (\delta_{i-1} + \delta_{i+1})$ *for those i where the diagonals of F_{i-1} and F_{i+1} have opposite slopes.*

Note that when P corresponds to an alternating node in the expanded tree E (see Definition 4.20), all diagonals of P have the same slope, and condition (5) is unnecessary.

Proof. For the "if" direction, let us first check that all the dihedral angles of P are in the range $(0, \pi)$. The top and bottom edges of P have dihedral angles of the form $\alpha_i \pm \delta_{i-1} \pm \delta_{i+1}$. If the diagonals of F_{i-1} and F_{i+1} have the same slope, the signs of δ_{i-1} and δ_{i+1} offset each other, and condition (4) suffices to place the angles in the right range. If the diagonals of F_{i-1} and F_{i+1} have opposite slopes, condition (5) suffices. For lateral edges, the dihedral angle $\pi - \alpha_i - \alpha_{i+1} + \delta_i + \delta_{i+1}$ is positive by condition (1) and less than π by condition (3). The only remaining edges are the new diagonals, whose angles are in the right range because we have chosen $\delta_i \in [0, \pi/2)$. (If $\delta_i = 0$, F_i will be a totally geodesic quadrilateral, whose diagonal is only a bookkeeping device.)

In bending lateral faces along diagonals, we have preserved the property that external angles around each vertex of P sum to 2π. Thus it remains to check that

$(*)$ for any normal curve $\gamma \subset \partial P$ that does not encircle a single vertex, $\sum_\gamma \epsilon_i > 2\pi$.

The proof of this statement is identical to the argument in Lemma 5.12. In Case 0, the exterior angle sum of a belt curve γ is still $\sum_{i=1}^{k} 2\alpha_i$. In Case 1, we can still tighten γ, reducing its angle sum at each step, until it encircles a single vertex. In Case 2, we can once again cut γ into arcs $\gamma_1, \ldots, \gamma_m$. The same tightening argument shows that the angle sum of γ_1 is at least $2\pi + 2\delta_i - 2\alpha_i$ (the minimum value being attained when γ_1 only passes through F_i) and the angle sum of γ_2 is at least $2\pi + 2\delta_j - 2\alpha_j$, for some $j \neq i$. Thus

$$\sum_\gamma \epsilon_i \;\geq\; 4\pi + (2\delta_i - 2\alpha_i) + (2\delta_j - 2\alpha_j) \;>\; 2\pi,$$

by condition (1), and P is an angled polyhedron.

For the "only if" direction, observe from the beginning of the proof that conditions $(3)-(5)$ are necessary to make the edges of P have dihedral angles in the range $(0, \pi)$. Meanwhile, conditions (1) and (2) are necessary for the same reasons as in Lemma 5.12, illustrated by the normal curves in Figure 5.11. $\square$

Chapter 6

Angle Structures for Arborescent Links

In the previous chapter, we have constructed an ideal polyhedral decomposition for an arborescent link K, consisting of tetrahedra and prisms. In this chapter, we will show how to assign convex angles to the ideal polyhedra. The construction will succeed whenever the arborescent diagram $D_T(K)$ is *balanced*, i.e. so long as neighboring rational tangles have enough crossings between them. (See Definition 4.17 for the precise statement.) This will prove that all balanced arborescent links are hyperbolic.

The proof proceeds as follows. In Section 6.1, we study the cusp triangulation in the product regions of a link complement. Using the cusp combinatorics, we can parametrize the angle structures of each product region in terms of *pleating angles* assigned to edges of E. In Section 6.2, we investigate how the angles on product regions interact with clasps and prisms, setting up a system of local gluing equations sufficient for an angle structure on the entire link complement. Finally, in Section 6.3 we find angles that satisfy the gluing equations, giving us an angle structure.

6.1 Cusp combinatorics of the product regions

Our exposition in this section borrows heavily from François Guéritaud's description of punctured torus bundles [11].

Let $B \subset E$ be a branch of an expanded tree, with vertices labeled by letters $L^{\pm 1}$ or $R^{\pm 1}$ of a word Ω. In the layered triangulation of the corresponding product region, each layer Δ_i consists of two tetrahedra, T_i and T_i', as in Figure 5.5. It is clear from the figure that each tetrahedron has exactly one ideal vertex at each puncture of S, i.e. at each strand of the 4-string braid. Since the combinatorics of the four strands is identical, let us focus on a single puncture of the 4-punctured sphere.

The tetrahedron layer Δ_i intersects the neighborhood of a puncture in two *boundary triangles*, one from a truncated vertex of T_i and one from T_i'. These boundary triangles meet at two vertices that come from shared edges of T_i and T_i'. (This completes a loop, corresponding to the meridian of a component of K.) Following Guéritaud, let us call the non-shared vertices the *apices* of the two boundary triangles. The two triangles are aligned so that one apex points up and the other one points down (where, as before, we identify *up* as the direction of increasing indices).

The top and bottom pleated surfaces that form the boundary of Δ_i can be seen in Figure 6.1 as zigzag lines of three segments each. Since the layers of tetrahedra are stacked along these pleated surfaces, we must stack the pairs of boundary triangles in a way that respects the combinatorics of the braid.

In the resulting cusp triangulation, each vertex comes from an edge shared by a

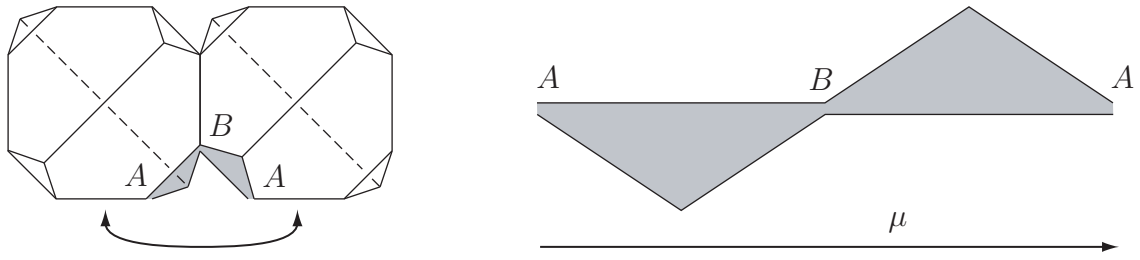

Figure 6.1: The truncated vertices of a layer of two tetrahedra intersect a puncture in two conjoined boundary triangles, pointing in different directions.

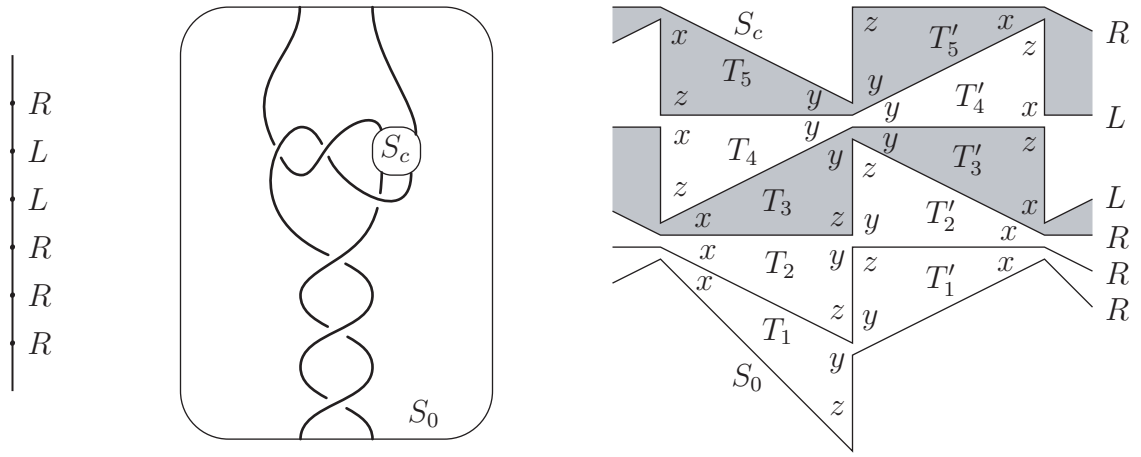

Figure 6.2: Three views: a branch $B \subset E$ labeled with the word $\Omega = R^3L^2R$; the corresponding braid of K; and the cusp triangulation of the braid exterior.

few consecutive layers of tetrahedra. The slope of this edge appears on the Farey graph as a vertex shared by several consecutive triangles $\tau_i, \ldots, \tau_{i+k}$, corresponding to a subword of Ω of the form RL^*R or LR^*L, where $* \geq 0$. Each Farey triangle labeled with one of the letters of this subword contributes a different pleated surface, whose zigzag intersection with the cusp visits the given vertex.

The resulting cusp triangulation, corresponding to the word $\Omega = R^3L^2R$, is shown in Figure 6.2. To emphasize the layered structure of the triangulation, the cusp triangles are shown "opened up," separating neighboring pleated surfaces.

Definition 6.1. Recall, from Table 5.9, that we think of layers of tetrahedra as positioned between consecutive letters of Ω. If the letters Ω_i, and Ω_{i+1} on either side of Δ_i are different, we say that i is a *hinge index* and the tetrahedra of Δ_i are *hinge tetrahedra*. In Figure 6.2, the hinge layers are shaded.

Between consecutive hinge layers in a cusp triangulation, a series of triangles shares a single vertex (and the corresponding layers of tetrahedra share an edge of the same slope). We say that these boundary triangles form a *fan*. Each fan in the cusp triangulation corresponds to a twist region in the braid, where the shared edge is isotopic to a vertical arc at each crossing. Each fan also corresponds to a vertex of

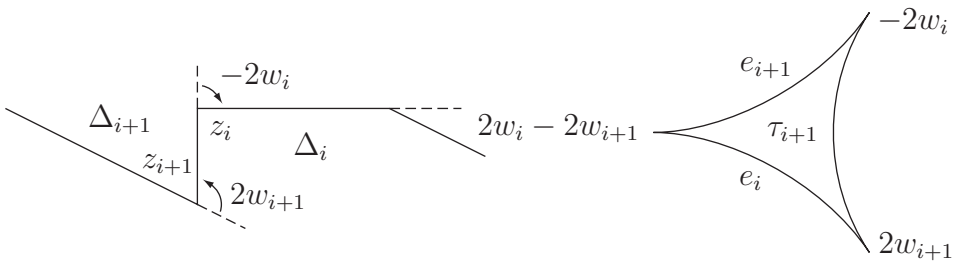

Figure 6.3: Pleating angles, seen on the cusp (left) and in the Farey graph (right).

the original weighted tree T.

Definition 6.2. We label the internal angles of boundary triangles as follows. The apex of each triangle in the i^{th} layer is labeled z_i; continuing clockwise from the apex, the other two angles are x_i and y_i. (The subscripts are omitted in Figure 6.2.) Following Example 2.5, note that these angles give dihedral angles of tetrahedra T_i and T_i' as long as they are non-negative and $x_i + y_i + z_i = \pi$. Note as well that we have chosen to make the two tetrahedra isometric; in the future, we will not distinguish between them.

We will parametrize the dihedral angles of tetrahedra in terms of *pleating angles* assigned to edges of E.

Definition 6.3. Let $e_0, \ldots, e_c$ be the edges of a branch B. To each e_i, we assign a *pleating angle* $w_i \in [0, \pi/2]$. By Table 5.9, we can identify e_i with the corresponding Farey edge. Then in the triangle $\tau_i \subset \mathcal{F}$, we label the vertex opposite e_i by $2w_i$; in the triangle $\tau_{i+1} \subset \mathcal{F}$, we label the vertex opposite e_i by $-2w_i$. This gives each of the Farey triangles $\tau_1, \ldots \tau_c$ two vertex angles. The third vertex of τ_i, shared by e_i and e_{i+1}, is labeled $2w_i - 2w_{i+1}$. (See Figure 6.3.)

The parameters w_i are called pleating angles because a consecutive pair w_i, w_{i+1} defines the (oriented) external angles of the pleated surface S_{i+1}. The slope of each edge of S_{i+1} is given by one of the vertices of τ_{i+1}, and the external angle at that edge is given by the corresponding vertex label. (We have picked signs so that a pleating angle is positive if a surface is convex toward the inside, i.e. toward higher indices,

and negative otherwise.) Note that the three vertex labels of τ_{i+1} add up to 0, and thus a meridian around a puncture of S_{i+1} has angle holonomy 0.

For $1 \leq i \leq c-1$, the edge e_i corresponds to a tetrahedron layer Δ_i. By locating the pleated surface on either side of Δ_i in Figure 6.2, we can see that the tetrahedron angle z_i has to equal $\pi - 2w_i$. In fact, Guéritaud has made the clever observation that all the angles of tetrahedron layer Δ_i can be computed from pleating angles, via the following table.

Ω_i, Ω_{i+1}	$L \quad L$	$R \quad R$	$L \quad R$	$R \quad L$
x_i	$w_{i-1} + w_{i+1}$	$2w_i - w_{i-1} - w_{i+1}$	$w_i + w_{i-1} - w_{i+1}$	$w_i - w_{i-1} + w_{i+1}$
y_i	$2w_i - w_{i-1} - w_{i+1}$	$w_{i-1} + w_{i+1}$	$w_i - w_{i-1} + w_{i+1}$	$w_i + w_{i-1} - w_{i+1}$
z_i	$\pi - 2w_i$	$\pi - 2w_i$	$\pi - 2w_i$	$\pi - 2w_i$

Table 6.4: Pleating angles determine the dihedral angles of Δ_i.

Definition 6.4. For $1 \leq i \leq c-1$, i.e. for all indices corresponding to tetrahedra, we will require that the pleating angles satisfy the following conditions:

(Range) $0 < w_i < \pi/2$.

(Concavity) If i is not a hinge index, $2w_i > w_{i-1} + w_{i+1}$.

(Hinge) If i is a hinge index, $w_i > |w_{i+1} - w_{i-1}|$.

See Figure 6.9 for a graphical interpretation of these conditions.

Lemma 6.5 (Guéritaud)**.** *Suppose that the pleating angles w_i satisfy the range, concavity, and hinge conditions, and set the tetrahedron angles as in Table 6.4. Then*

(1) *for each Δ_i, the angles x_i, y_i, z_i are positive and add up to π, and*

(2) *the dihedral angles around any edge interior to the product region add up to 2π.*

Proof. Conclusion (1) is evident from Table 6.4. Conclusion (2) also follows by a quick computation from the table, keeping track of the change in pleating angle as we "scan" each vertex on the cusp by pleated surfaces. □

6.2 The pleated boundary of product regions

The choice of pleating angles in a product region is somewhat constrained by the geometry of the boundary spheres S_0 and S_c.

Lemma 6.6. *Let $B \subset E$ be a terminal branch of length c, with edges labeled by pleating angles $w_0, \ldots, w_{c-1}$. These labels w_i give an angled triangulation for the rational tangle corresponding to B, so long as*

(1) *$w_1, \ldots, w_{c-2}$ satisfy the range, concavity, and hinge conditions, and*

(2) *$w_{c-1} = \frac{\pi}{2}$.*

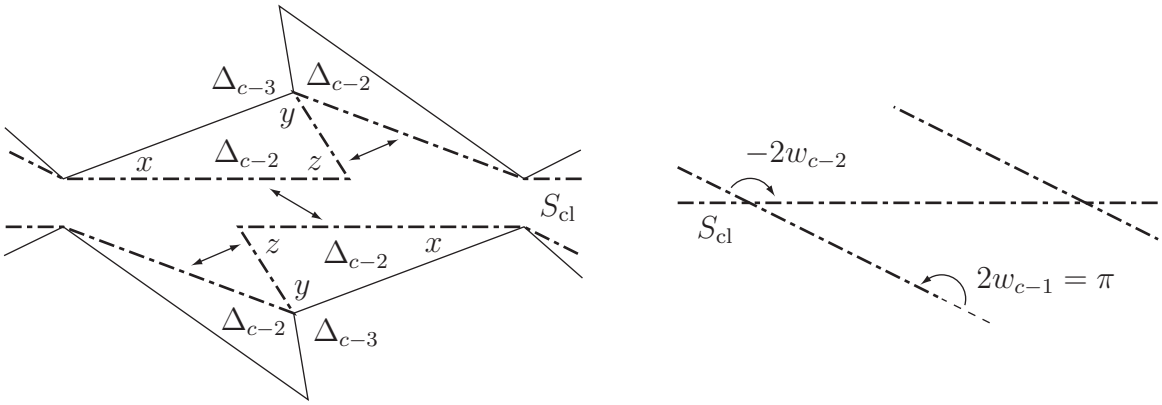

Figure 6.5: Cusp view: folding pleated surface S_{cl} requires a pleating angle $w_{c-1} = \frac{\pi}{2}$.

Proof. Recall from Section 5.2 that the layered triangulation of a rational tangle consists of a product region exactly as above, ending at the pleated surface $S_{\text{cl}} = S_{c-1}$ that corresponds to Farey triangle τ_{c-1}. Thus Lemma 6.5 takes care of the tetrahedron angles inside the product region.

On the cusp torus, the folding pattern of S_{cl} brings together the layered triangulations of two different strands of K. We can see the surface S_{cl} making a hairpin turn at the peripheral edges, as in Figure 6.5. By Definitions 5.8 and 6.3, the pleating angle at the peripheral edges is labeled $2w_{c-1}$, and the hairpin turn along those edges

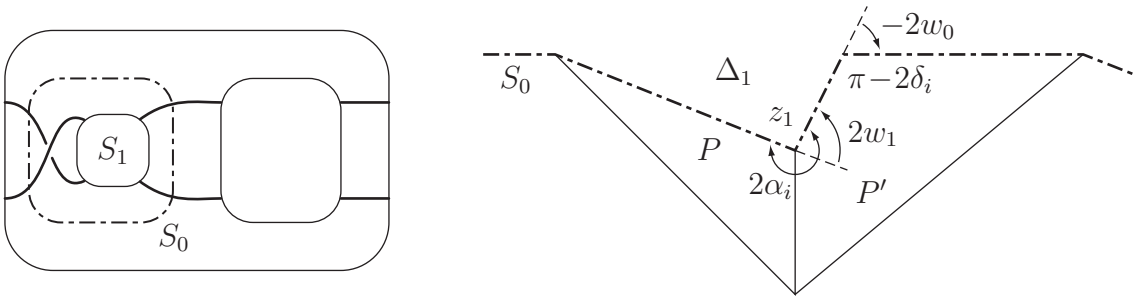

Figure 6.6: Gluing a product region to a pair of prisms when the edge e_0 is exceptional. Left: sphere S_0 as seen in the link diagram. Right: the cusp combinatorics.

requires an external angle of π. Thus setting $w_{c-1} = \pi/2$ allows us to construct the clasp in a geometrically consistent fashion. □

When a branch B connects to a node $v \in E$, and thus the product region connects to a pair of prisms P and P', the pleating angles on the boundary of the product region determine the dihedral angles along matching faces $F_i \subset P$ and $F_i' \subset P'$. This can happen in one of two ways, depending on whether the edge e_0 of B adjacent to v is exceptional (see Definition 4.21).

Lemma 6.7. *Let S_0 be a 4-punctured sphere bounding the product region of B, which is glued to faces $F_i \cup F_i'$ in the pair of prisms. Let the dihedral angles on F_i and F_i' be labeled as in Section 5.3. Then the pleating angles on S_0 determine these dihedral angles as follows:*

(1) *If e_0 is exceptional, $\alpha_i = \frac{\pi}{2} + w_1$ and $\delta_i = w_0$.*

(2) *If e_0 is not exceptional, $\alpha_i = \frac{\pi}{2} + w_0 - w_1$ and $\delta_i = w_0$.*

Proof. Let v_1 be the first vertex of B, across e_0 from v. If e_0 is exceptional, the crossing corresponding to v_1 is on a band parallel to the band of v. In this case, the edge of S_0 isotopic to a vertical arc at the crossing is glued to a lateral edge of $F_i \cup F_i'$. (See Figure 6.6.) Thus the top and bottom edges of $F_i \cup F_i'$ are glued to the edges of S_0 that are not shared with S_1, i.e., the edges whose slope is opposite e_1 in the Farey graph. By Definition 6.3, the pleating angle on those edges is $2w_1$, and by Remark

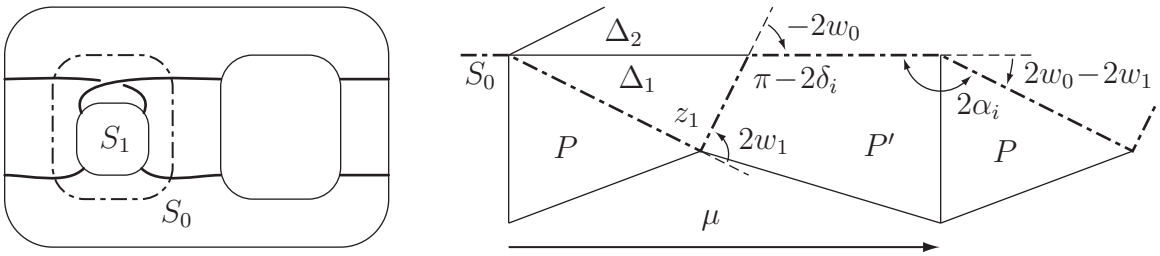

Figure 6.7: Gluing a product region to a pair of prisms when the edge e_0 is not exceptional. Left: sphere S_0 as seen in the link diagram. Right: the cusp combinatorics.

5.13, the total internal angle of $F_i \cup F_i'$ at those edges is $2\alpha_i$. Thus we can compute that $2\alpha_i = \pi + 2w_1$.

If e_0 is not exceptional, the crossing corresponding to v_1 is on a band perpendicular to the band of v. In this case, the vertical edge at the crossing and its twin of the same slope are glued to the top or bottom edges of F_i and F_i'. (See Figure 6.7.) In the cusp picture, we can see the crossing edge as a vertex of the first fan of Ω; thus the pleating angle of S_0 at this edge (and its twin of the same slope) is $2w_0 - 2w_1$. The signs work out so that $2w_0 - 2w_1$ is negative when $2\alpha_i < \pi$. We can therefore compute that $2\alpha_i = \pi + 2w_0 - 2w_1$.

In either case, the diagonals of $F_i \cup F_i'$ are glued to edges of S_0 whose pleating angle is $-2w_0$. Thus $\pi - 2\delta_i = \pi - 2w_0$, and $\delta_i = w_0$. □

Remark 6.8. Lemma 6.7 describes the dihedral angles of the prisms glued to boundary sphere S_0 of a product region. If the other boundary sphere S_c is also glued to a pair of prisms, the dihedral angles of those prisms will of course have the same description, with e_0 replaced by e_c, v_1 by v_c, w_0 by w_c, and w_1 by w_{c-1}.

We have now described the interaction between the dihedral angles at all 4-punctured spheres connecting pieces of the link complement, with one exception: places where a prism pair $P \cup P'$ is glued directly to another prism pair $Q \cup Q'$ rather than to a product region. In the expanded tree E, this corresponds to two nodes separated by a single edge, without a branch in between.

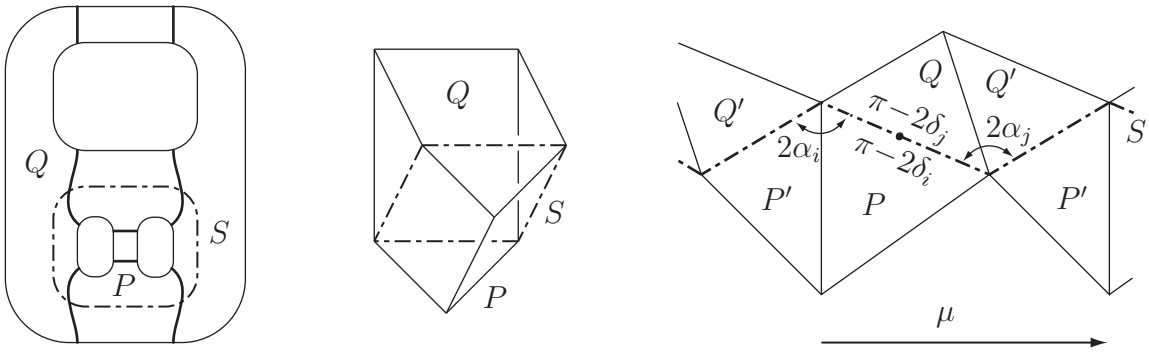

Figure 6.8: Gluing prisms $P \cup P'$ directly to prisms $Q \cup Q'$. Left to right: the bands of v_1 and v_2; the gluing by a 90° twist; the cusp combinatorics showing dihedral angles.

Lemma 6.9. *Let P and P' be prisms corresponding to a node $v_1 \in E$, and Q and Q' be prisms corresponding to an adjacent node v_2. Then*

(1) *the lateral faces $F_i \cup F_i'$ of $P \cup P'$ are glued to lateral faces $G_j \cup G_j'$ of $Q \cup Q'$ by a 90° twist,*

(2) *these faces not bent along diagonals ($\delta_i = \delta_j = 0$), and*

(3) *the total internal angle at the top and bottom edges of F_i and F_i' matches the total internal angle at the top and bottom edges of G_j and G_j' (i.e., $\alpha_i = \alpha_j$).*

Proof. In converting a weighted tree T into an expanded tree E, we have expanded the branches of T, but did not touch the edges that connect two nodes. Thus v_1 and v_2 must have come from adjacent nodes of T, which have perpendicular bands. Suppose without loss of generality that the band of v_1 is horizontal, as in Figure 5.10. Then the top and bottom edges of faces $F_i \cup F_i'$ of $P \cup P'$, which are horizontal in the projection diagram, are glued to edges that span the band of v_2, i.e. to lateral edges of $G_j \cup G_j'$. Thus the prisms are glued by a 90° twist, as in Figure 6.8.

If the faces $F_i \cup F_i'$ are bent along diagonals, P and P' must be convex at those diagonals, implying that Q and Q' are concave. Thus the only way to keep all four prisms convex if $\delta_i = \delta_j = 0$. As a result, the 4-punctured sphere S along which we perform the gluing is only bent along vertical and horizontal edges. As with 4-punctured spheres in a product region, the signed pleating angles on S must sum to

0. Since those pleating angles are equal to $2\alpha_i - \pi$ at the top/bottom edges of $F_i \cup F_i'$ and $\pi - 2\alpha_j$ at the top/bottom edges of $G_j \cup G_j'$, it follows that $\alpha_i = \alpha_j$. □

Remark 6.10. Since $\alpha_i = \alpha_j$, we can define both angles via a pleating angle assigned to the edge e connecting v_1 to v_2.

We can combine the results of this section into the following statement.

Theorem 6.11. *Let K be an arborescent link represented by an expanded tree E. To each edge $e \subset E$, we assign a pleating angle $w(e) \in [0, \pi/2]$. These pleating angles determine dihedral angles of the tetrahedra and prisms of $S^3 \setminus K$ as follows:*

- *For a tetrahedron layer $\Delta(e)$ corresponding to an interior edge of a branch, the dihedral angles are defined via Table 6.4.*
- *For a prism pair corresponding to a node $v \in E$, let $e_0^1, \dots e_0^k$ be the edges adjacent to v. If e_0^i is part of a branch B_i, let e_1^i be the next edge of the branch. Now,*
  - *if e_0^i is part of a branch and not exceptional, $\alpha_i = \frac{\pi}{2} + w_0^i - w_1^i$ and $\delta_i = w_0^i$,*
  - *if e_0^i is part of a branch and is exceptional, $\alpha_i = \frac{\pi}{2} + w_1^i$ and $\delta_i = w_0^i$, and*
  - *if e_0^i connects to another node, $\alpha_i = \frac{\pi}{2} - w_0^i$ and $\delta_i = 0$.*

This assignment of angles gives the polyhedra of $S^3 \setminus K$ an angle structure as long as

(1) *every terminal edge of E is labeled $\frac{\pi}{2}$,*

(2) *if e is an interior edge of a branch, the pleating angle $w(e)$ satisfies the range, concavity, and hinge conditions of Definition 6.4, and*

(3) *the angles of each prism satisfy the conditions of Theorem 5.14.*

Proof. See Lemmas 6.5, 6.6, 6.7, and 6.9, and Theorem 5.14. □

6.3 Existence of angle structures

The goal of this section is to assign a pleating angle to each edge of E in a way that satisfies the conditions of Theorem 6.11, completing the proof of Theorem 1.18.

We start with the branches of E. For each branch, we will think of pleating angles via a graph in the plane, as in Figure 6.9. More concretely, we think of pleating angle w_i as lying at the point (i, w_i).

Lemma 6.12. *Let $B \subset E$ be a terminal branch of length $c \geq 3$, with edges labeled by pleating angles $w_0, \ldots, w_{c-1}$. Chose values of $w_0, w_1 \in (0, \frac{\pi}{2})$ such that*

$$m \;:=\; \frac{\pi/2 - w_1}{c-2} \;<\; \min(w_0,\, w_1 - w_0).$$

Given any such choice, there exist pleating angles $w_2, \ldots, w_{c-2}$ that give the rational tangle of B an angle structure.

Following Remark 5.9, we note that if B is short (i.e., $c = 2$), the rational tangle of B contains no tetrahedra at all.

Proof. The proof is a construction in coordinate geometry. We construct a line ℓ connecting points $(1, w_1)$ and $(c-1, \frac{\pi}{2})$; the value m defined in the statement of the

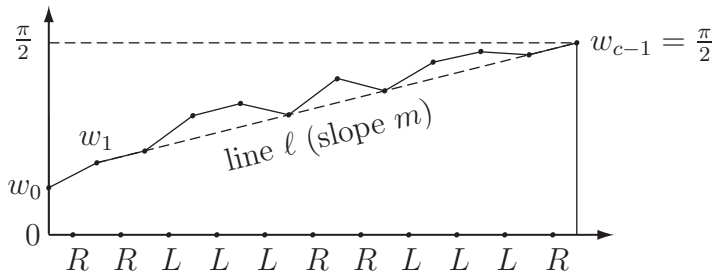

Figure 6.9: In finding pleating angles for a terminal branch, we place hinge indices on the line ℓ and connect them by segments of parabolas. When the parabolas are sufficiently close to ℓ, all the conditions are satisfied.

lemma is the slope of ℓ. For every hinge index $i \in [2, c-2]$, we will place the pleating angle w_i at the point of ℓ whose x-coordinate is i. We will then connect the pleating angles of hinge indices by segments of parabolas that lie above ℓ but below $\frac{\pi}{2}$. See Figure 6.9.

This construction defines an angle $w_i \in (0, \frac{\pi}{2})$ for $i = 2, \ldots, c-2$. Thus the range condition is satisfied automatically. Because the non-hinge indices lie on concave-down parabolas, the concavity condition is also satisfied. (If 1 is not a hinge index, we need to ensure that the first parabolic segment lies close enough to ℓ to satisfy the concavity condition at w_1, but this can always be done because $m < w_1 - w_0$.) Thus it remains to check the hinge condition.

Because the line ℓ has a positive slope, $w_0 < w_1 < w_i$ for all $i > 1$. In particular, if 1 is a hinge index, the hinge condition at $i = 1$ can be rearranged to read $w_0 > w_2 - w_1$. But $w_0 > m$ by hypothesis, and we can make the difference $w_2 - w_1$ arbitrarily close to m. For any other hinge index i, bringing the adjacent parabolas closer to ℓ will bring the difference $|w_{i+1} - w_{i-1}|$ arbitrarily close to $2m$. But

$$2m \;=\; m + m \;<\; w_0 + (w_1 - w_0) \;=\; w_1 \;<\; w_i$$

for all $i > 1$. Thus the hinge condition is satisfied, and the rational tangle has an angle structure by Lemma 6.6. □

Remark 6.13. In the special case when E has no nodes (and thus K is a two-bridge link), an argument just like the one in Lemma 6.12 guarantees the existence of angle structures. In this scenario, $w_1 = w_{c-1} = \frac{\pi}{2}$. If we place all hinge indices on a horizontal line at height $\frac{\pi}{4}$ and connect them by pieces of parabolas lying above the line, the hinge and concavity conditions are satisfied immediately. This gives a quick proof that all two-bridge links whose word Ω has both letters R and L (i.e. all two-bridge links with at least two twist regions) are hyperbolic.

For interior branches, there is a statement of similar generality to Lemma 6.12. However, we will only need a simpler statement.

Lemma 6.14. *Let $B \subset E$ be an interior branch of length c, with edges labeled by pleating angles $w_0, \ldots, w_c$. Chose a value of $w_0 = w_c \in (0, \frac{\pi}{2})$. If $c \geq 2$, also choose a value of $w_1 = w_{c-1} \in (w_0, \frac{\pi}{2})$. Given any such choice, there exist pleating angles $w_2, \ldots, w_{c-2}$ that give the product region of B an angle structure.*

Proof. If $c = 1$, the product region contains no tetrahedra by Remark 5.7, and thus there is nothing to prove. If $c = 2$ or $c = 3$, both the hinge and concavity conditions are satisfied because w_1 and w_{c-1} are equal, and both are greater than w_0 and w_c. Thus, we can assume that $c \geq 4$.

If $c \geq 4$, we construct a horizontal line ℓ from $(1, w_1)$ to $(c-1, w_{c-1})$. As in the last lemma, we place all hinge indices on this line and connect them by pieces of parabolas that lie above the line. Thus the hinge and concavity conditions, as needed, are immediately satisfied for all $i \in [2, c-2]$. For $i = 1$ or $i = c-1$, we can satisfy the hinge and/or concavity conditions by ensuring that the first and last parabolas lie sufficiently close to ℓ. □

Now, assume that the expanded tree E is *balanced*; i.e., E represents a balanced diagram. Recall from Definition 4.17 that this is a condition on neighboring terminal branches that meet at a node of E. For alternating nodes, the definition says that two short branches (i.e., two terminal branches of length 2) cannot be adjacent. For non-alternating nodes, each triple of neighboring edges at v must satisfy the condition

$$\frac{1}{c_{i-1}} + \frac{1}{c_i} + \frac{1}{c_{i+1}} \; < \; 1,$$

where for terminal branches $c_i = \ell(B_i)$ is the length of the branch, or $c_i = \infty$ otherwise.

Remark 6.15. An easy but useful observation is that for positive integers c_1, c_2,

$$\frac{1}{c_1} + \frac{1}{c_2} < 1 \quad \text{implies} \quad \frac{1}{c_1} + \frac{1}{c_2} \leq \frac{5}{6}\,.$$

We are now ready to assign pleating angles to the edges of a balanced tree E. We pick a value of $\varepsilon > 0$, indeterminate for now, and proceed as follows.

(1) Every terminal edge $e \subset E$ is labeled $\frac{\pi}{2}$.

(2) Every edge $e \subset E$ connecting two nodes is labeled $\frac{\pi}{12} - 3\varepsilon$.

(3) Let $B \subset E$ be an interior branch of length c. We label $w_0 = w_1 = \frac{\pi}{18} + \varepsilon$. If $c \geq 2$, we label $w_1 = w_{c-1} = \frac{\pi}{9} + \varepsilon$. If $c \geq 4$, by Lemma 6.14 there exist pleating angles $w_2, \ldots, w_{c-2}$ that give the product region of B an angle structure.

(4) Let $v \in E$ be an alternating node adjacent to edges $e_0^1, \ldots e_0^k$. Some of these edges may have already been labeled because they connect to another node or are part of an interior branch; all remaining edges adjacent to v are part of terminal branches. If B_i is a short terminal branch, we give e_0^i the pleating angle $w_0^i = \frac{\pi}{4} + 2\varepsilon$. If B_i is a terminal branch of length $c_i \geq 3$, we give e_0^i the pleating angle $w_0^i = \frac{\pi}{4} + \varepsilon$ and give the next edge e_1^i the pleating angle $w_1^i = \frac{3\pi}{8} + \varepsilon$. We can then compute that

$$m_i \;=\; \frac{\frac{\pi}{2} - w_1^i}{c_i - 2} \;=\; \frac{\frac{\pi}{8} - \varepsilon}{c_i - 2} \;\leq\; \frac{\pi}{8} - \varepsilon \;<\; \min(w_0^i,\, w_1^i - w_0^i)\,,$$

satisfying the condition of Lemma 6.12. Thus there are choices for the remaining pleating angles of the branch that give the rational tangle an angle structure.

(5) Let $v \in E$ be a non-alternating node adjacent to edges $e_0^1, \ldots e_0^k$. For each terminal branch B_i of length c_i, we give e_0^i the pleating angle $w_0^i = \frac{\pi}{2c_i} + \varepsilon$. If $c_i \geq 3$, we give the next edge e_1^i the pleating angle $w_1^i = \frac{\pi}{c_i} + \varepsilon$. We can then compute that

$$m_i \;=\; \frac{\frac{\pi}{2} - w_1^i}{c_i - 2} \;=\; \frac{\frac{\pi}{2} - \frac{\pi}{c_i} - \varepsilon}{c_i - 2} \;<\; \frac{\pi}{2} \cdot \frac{1 - \frac{2}{c_i}}{c_i - 2} \;=\; \frac{\pi}{2c_i}\,,$$

satisfying the condition of Lemma 6.12. Thus there are choices for the remaining pleating angles of the branch that give the rational tangle an angle structure.

Theorem 6.16. *Let E be a balanced expanded tree representing an arborescent link K, and let the edges of E be labeled with pleating angles as above. Then there is a value of $\varepsilon > 0$ such that these edge labels give the polyhedral decomposition of $S^3 \setminus K$ an angle structure.*

Proof. By Theorem 6.11, the pleating angles assigned to edges of E define dihedral angles for all the polyhedra and prisms. Furthermore, we have already checked that the pleating angles of the branches of E give angle structures for the corresponding product regions and rational tangles. Thus it remains to show that the dihedral angles of each prism satisfy the definition of an angled polyhedron. We will consider two cases, depending on the node $v \in E$ associated to the prism P.

Case 1: v is an alternating node (which may or may not be exceptional). Let $e_0^1, \dots e_0^k$ be the edges adjacent to v. By Theorem 6.11, the above choices of pleating angles give the following dihedral angles on the prism.

- Exceptional edges. We adopt the convention that, if v is an exceptional node, the exceptional edge is e_0^1.
  - If e_0^1 is part of an interior branch of length c_1, $\delta_1 = \frac{\pi}{18} + \varepsilon$. If $c_1 = 2$, $\alpha_1 = \frac{5\pi}{9} + \varepsilon$; otherwise, $\alpha_1 = \frac{11\pi}{18} + \varepsilon$.
  - If e_0^1 is part of a terminal branch, $\delta_1 = \frac{\pi}{4} + \varepsilon$ and $\alpha_1 = \frac{7\pi}{8} + \varepsilon$. (This occurs when K is a Montesinos link. Because we have added crossings to a rational tangle that began with at least two crossings, $c_1 \geq 3$.)
- Non-exceptional edges:
  - If e_0^i is part of an interior branch of length c_i, $\delta_i = \frac{\pi}{18} + \varepsilon$. If $c_i = 2$, $\alpha_i = \frac{4\pi}{9}$; otherwise, $\alpha_i = \frac{\pi}{2}$.
  - If e_0^i is part of a terminal branch of length $c_i = 2$, $\delta_i = \alpha_i = \frac{\pi}{4} + 2\varepsilon$. If $c_i \geq 3$, $\delta_i = \frac{\pi}{4} + \varepsilon$ and $\alpha_i = \frac{\pi}{4} + \varepsilon$.
  - If e_0^i connects to another node, $\delta_i = 0$ and $\alpha_i = \frac{5\pi}{12} + 3\varepsilon$.

To prove that P is an angled polyhedron, we check the conditions of Theorem 5.14.

(1) $(\alpha_i - \delta_i) + (\alpha_j - \delta_j) < \pi$ for all $i \neq j$. We may restrict our attention to the case when v is exceptional, because otherwise $\alpha_i - \delta_i < \frac{\pi}{2}$ for all i.

If the exceptional edge e_0^1 is part of an interior branch, the value of $\alpha_j - \delta_j$ is largest for $j = 1$, where $\alpha_1 - \delta_1 \leq \frac{5\pi}{9}$. Meanwhile, for $i \neq 1$, $\alpha_i - \delta_i < \frac{4\pi}{9}$ in all cases (if e_0^i connects to another node, we need to make ε sufficiently small).

If the exceptional edge e_0^1 is part of a terminal branch, the value of $\alpha_j - \delta_j$ is also largest for $j = 1$, where $\alpha_1 - \delta_1 = \frac{5\pi}{8}$. In this case, K is a Montesinos link, and all edges at v belong to terminal branches. Thus, for all $i \neq 1$, $\alpha_i - \delta_i < \frac{\pi}{8}$, and the condition is satisfied.

(2) $\sum_{i=1}^{k} \alpha_i > \pi$. Observe that $\alpha_i \geq \frac{3\pi}{8}$ for all i, except when e_0^i is part of a short terminal branch. Thus, if none of the terminal branches at v is short, we are done because $k \geq 3$. On the other hand, if e_0^i is part of a short branch, we note that E is balanced and thus two short branches cannot be adjacent. In this case, $\alpha_i = \frac{\pi}{4} + 2\varepsilon$ while $\alpha_{i\pm1} \geq \frac{3\pi}{8}$, and the sum adds up to more than π.

(3) $(\alpha_i - \delta_i) + (\alpha_{i+1} - \delta_{i+1}) > 0$. Observe that $\alpha_i \geq \frac{\pi}{4} + 2\varepsilon$ for all i, with equality if and only if e_0^i is part of a short terminal branch. Similarly, $\delta_j \leq \frac{\pi}{4} + 2\varepsilon$, with equality if and only if e_0^j is part of a short terminal branch. When $j = i \pm 1$, at least one of the two inequalities must be strict, because E is balanced. Thus $\alpha_i > \delta_{i+1}$ and $\alpha_{i+1} > \delta_i$.

(4) $|\delta_{i-1} - \delta_{i+1}| < \alpha_i < \pi - |\delta_{i-1} - \delta_{i+1}|$. The first inequality follows because, as we have just argued, $\alpha_i > \delta_j$ for $j = i \pm 1$. Now, consider the second inequality. If K is a Montesinos link, every edge at v is part of a terminal branch; thus $\alpha_i \leq \frac{7\pi}{8} + \varepsilon$ while $|\delta_{i-1} - \delta_{i+1}| \leq \varepsilon$, for all i. If K is not a Montesinos link, $\alpha_i \leq \frac{2\pi}{3}$ and $\delta_j \leq \frac{\pi}{4} + \varepsilon$, for all i and j, giving the inequality.

(5) Condition (5) does not apply, because v is alternating.

Case 2: v is not an alternating node. By Theorem 4.22, this means that v is not exceptional. Let $e_0^1, \dots e_0^k$ be the edges adjacent to v; if e_0^i is part of a branch B_i, let e_1^i be the next edge of the branch. Then the above choices of pleating angles give the following dihedral angles on the prism:

- If e_0^i is part of an interior branch of length c_i, $\delta_i = \frac{\pi}{18} + \varepsilon$. If $c_i = 2$, $\alpha_i = \frac{4\pi}{9}$; otherwise, $\alpha_i = \frac{\pi}{2}$.
- If e_0^i is part of a terminal branch of length c_i, $\delta_i = \frac{\pi}{2c_i} + \varepsilon$. If $c_i = 2$, $\alpha_i = \frac{\pi}{4} + \varepsilon$; otherwise, $\alpha_i = \frac{\pi}{2} - \frac{\pi}{2c_i}$.
- If e_0^i connects to another node, $\delta_i = 0$ and $\alpha_i = \frac{5\pi}{12} + 3\varepsilon$.

To prove that P is an angled polyhedron, we check the conditions of Theorem 5.14.

(1) $(\alpha_i - \delta_i) + (\alpha_j - \delta_j) < \pi$ for all $i \neq j$. Since $\alpha_i - \delta_i < \frac{\pi}{2}$ for all i, this is immediate.

(2) $\sum_{i=1}^{k} \alpha_i > \pi$. We use the fact that $k \geq 3$. If e_0^i is part of a terminal branch for $i = 1, 2, 3$, then

$$\sum_{i=1}^{3} \alpha_i \;\geq\; \frac{3\pi}{2} \;-\; \frac{\pi}{2} \sum_{i=1}^{3} \frac{1}{c_i} \;>\; \pi,$$

because E is balanced. If two of these edges, say, e_0^1 and e_0^2, are part of terminal branches,

$$\alpha_1 + \alpha_2 \;\geq\; \pi - \frac{\pi}{2c_1} - \frac{\pi}{2c_2} \;\geq\; \frac{7\pi}{12},$$

because E is balanced and $\frac{1}{c_1} + \frac{1}{c_2} \leq \frac{5}{6}$. But then α_3 contributes at least $\frac{5\pi}{12} + 3\varepsilon$, bringing the total above π. Finally, if at most one of the edges e_0^i is part of a terminal branch, $\alpha_i > \frac{\pi}{4}$, while the other two edges contribute at least $\frac{5\pi}{12}$ each.

(3) $(\alpha_i - \delta_i) + (\alpha_{i+1} - \delta_{i+1}) > 0$. Condition (5) implies that $\alpha_i > \delta_{i+1}$ and $\alpha_{i+1} > \delta_i$.

(4) $|\delta_{i-1} - \delta_{i+1}| < \alpha_i < \pi - |\delta_{i-1} - \delta_{i+1}|$. Again, condition (5) is strictly stronger.

(5) $\delta_{i-1} + \delta_{i+1} < \alpha_i < \pi - (\delta_{i-1} + \delta_{i+1})$. The second inequality is never a problem: $\alpha_i \leq \frac{\pi}{2}$ for all i, while the sum $\delta_{i-1} + \delta_{i+1}$ can only add up to $\frac{\pi}{2}$ when both indices correspond to short terminal branches. But we can rule out this situation because E is balanced.

Similarly, the first inequality is satisfied when e_0^i is not part of a terminal branch. In this situation, $\alpha_i \geq \frac{5\pi}{12} + 3\varepsilon$, while $\delta_{i-1} + \delta_{i+1} \leq \frac{5\pi}{12} + 2\varepsilon$.

If e_0^i is part of a terminal branch, $\alpha_i \geq \frac{\pi}{2} - \frac{\pi}{2c_i}$. If e_0^{i-1} and e_0^{i+1} also belong to terminal branches,

$$\alpha_i - \delta_{i-1} - \delta_{i+1} \;\geq\; \frac{\pi}{2} \;-\; \frac{\pi}{2} \sum_{j=i-1}^{i+1} \frac{1}{c_j} \;-\; 2\varepsilon \;>\; 0 \quad \text{for small } \varepsilon,$$

because E is balanced. If e_0^{i-1} is part of a terminal branch but e_0^{i+1} is not,

$$\alpha_i - \delta_{i-1} - \delta_{i+1} \;\geq\; \frac{\pi}{2} \;-\; \frac{\pi}{2}\left(\frac{1}{c_{i-1}} + \frac{1}{c_i}\right) \;-\; 2\varepsilon - \frac{\pi}{18} \;>\; 0,$$

because $\frac{1}{c_{i-1}} + \frac{1}{c_{i+1}} \leq \frac{5}{6}$ and $\delta_{i+1} \leq \frac{\pi}{18} + \varepsilon$. Finally, if neither e_0^{i-1} nor e_0^{i+1} is part of a terminal branch, observe that $\alpha_i > \frac{\pi}{4}$ and $\delta_{i-1} + \delta_{i+1} \leq \frac{\pi}{9} + 2\varepsilon$.

This completes the proof of Theorem 6.16, and thus also of Theorem 1.18. □

References

[1] Colin C. Adams, *Augmented alternating link complements are hyperbolic*, Low-dimensional topology and Kleinian groups (Coventry/Durham, 1984), London Math. Soc. Lecture Note Ser., vol. 112, Cambridge Univ. Press, Cambridge, 1986, pp. 115–130.

[2] Kenneth Baker, *Knots on once-punctured torus fibers*, Ph.D. thesis, University of Texas at Austin, 2004.

[3] Joan S. Birman, *Braids, links, and mapping class groups*, Princeton University Press, Princeton, N.J., 1974, Annals of Mathematics Studies, No. 82.

[4] Ken Chan, *Constructing hyperbolic* 3*-manifolds*, Undergraduate thesis with Craig Hodgson, University of Melbourne, 2002.

[5] J. H. Conway, *An enumeration of knots and links, and some of their algebraic properties*, Computational Problems in Abstract Algebra (Proc. Conf., Oxford, 1967), Pergamon, Oxford, 1970, pp. 329–358.

[6] Richard Crowell, *Genus of alternating link types*, Ann. of Math. (2) **69** (1959), 258–275.

[7] Mario Eudave-Muñoz, *On hyperbolic knots with Seifert fibered Dehn surgeries*, Topology Appl. **121** (2002), no. 1-2, 119–141.

[8] Mario Eudave-Muñoz and John Luecke, *Knots with bounded cusp volume yet large tunnel number*, J. Knot Theory Ramifications **8** (1999), no. 4, 437–446.

[9] David Futer and Jessica S. Purcell, *Links with no exceptional surgeries*, 2004, arXiv:math.GT/0412307.

[10] David Gabai, *Genera of the arborescent links*, Mem. Amer. Math. Soc. **59** (1986), no. 339, i–viii and 1–98.

[11] François Guéritaud, *On canonical triangulations of the mapping tori over the punctured torus*, 2004, arXiv:math.GT/0406242.

[12] Wolfgang Haken, *Theorie der Normalflächen*, Acta Math. **105** (1961), 245–375.

[13] Craig D. Hodgson and Steven P. Kerckhoff, *Universal bounds for hyperbolic Dehn surgery*, Ann. of Math. (2) **162** (2005), no. 1, arXiv:math.GT/0204345.

[14] Marc Lackenby, *Word hyperbolic Dehn surgery*, Invent. Math. **140** (2000), no. 2, 243–282.

[15] ———, *The canonical decomposition of once-punctured torus bundles*, Comment. Math. Helv. **78** (2003), no. 2, 363–384.

[16] ———, *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.

[17] William Menasco, *Closed incompressible surfaces in alternating knot and link complements*, Topology **23** (1984), no. 1, 37–44.

[18] John Milnor, *Hyperbolic geometry: the first 150 years*, Bull. Amer. Math. Soc. (N.S.) **6** (1982), no. 1, 9–24.

[19] José M. Montesinos, *Seifert manifolds that are ramified two-sheeted cyclic coverings*, Bol. Soc. Mat. Mexicana (2) **18** (1973), 1–32.

[20] G. D. Mostow, *Strong rigidity of locally symmetric spaces*, Princeton University Press, Princeton, N.J., 1973, Annals of Mathematics Studies, No. 78.

[21] Kunio Murasugi, *On the genus of the alternating knot. I, II*, J. Math. Soc. Japan **10** (1958), 94–105, 235–248.

[22] Walter D. Neumann and Don Zagier, *Volumes of hyperbolic three-manifolds*, Topology **24** (1985), no. 3, 307–332.

[23] Ulrich Oertel, *Closed incompressible surfaces in complements of star links*, Pacific J. Math. **111** (1984), no. 1, 209–230.

[24] Grisha Perelman, *The entropy formula for the Ricci flow and its geometric applications*, 2002, arXiv:math.DG/0211159.

[25] ———, *Ricci flow with surgery on three-manifolds*, 2003, arXiv:math.DG/0303109.

[26] Gopal Prasad, *Strong rigidity of* $\mathbf{Q}$*-rank* 1 *lattices*, Invent. Math. **21** (1973), 255–286.

[27] Józef H. Przytycki, *Incompressibility of surfaces after Dehn surgery*, Michigan Math. J. **30** (1983), no. 3, 289–308.

[28] Jessica S. Purcell, *Cusp shapes of hyperbolic link complements and Dehn filling*, Ph.D. thesis, Stanford University, 2004.

[29] Igor Rivin, *Euclidean structures on simplicial surfaces and hyperbolic volume*, Ann. of Math. (2) **139** (1994), no. 3, 553–580.

[30] ———, *A characterization of ideal polyhedra in hyperbolic* 3*-space*, Ann. of Math. (2) **143** (1996), no. 1, 51–70.

[31] Makoto Sakuma and Jeffrey Weeks, *Examples of canonical decompositions of hyperbolic link complements*, Japan. J. Math. (N.S.) **21** (1995), no. 2, 393–439.

[32] William P. Thurston, *The geometry and topology of three-manifolds*, Princeton Univ. Math. Dept. Notes, 1980, Available at http://www.msri.org/gt3m/.

[33] ———, *Three-dimensional manifolds, Kleinian groups and hyperbolic geometry*, Bull. Amer. Math. Soc. (N.S.) **6** (1982), no. 3, 357–381.

[34] ———, *Three-dimensional geometry and topology. Vol. 1*, Princeton Mathematical Series, vol. 35, Princeton University Press, Princeton, NJ, 1997, Edited by Silvio Levy.

[35] Ying-Qing Wu, *Dehn surgery on arborescent knots*, J. Differential Geom. **43** (1996), no. 1, 171–197.

[36] ———, *Dehn surgery on arborescent links*, Trans. Amer. Math. Soc. **351** (1999), no. 6, 2275–2294.