Part 2

Tools, techniques, and families of examples

CHAPTER 7

Twist knots and augmented links

In this chapter, we study a class of hyperbolic knots that have some of the simplest geometry, namely twist knots. This class includes the figure-8 knot, the 5_2 knot, and the 6_1 knot that we have encountered so far. We also generalize to give examples of hyperbolic knots and links whose geometry is relatively explicit. This will equip us with many examples.

7.1. Twist knots and Dehn fillings

We first define twist knots and show that they have a geometric limit that is a link complement.

DEFINITION 7.1. A *twist region* of a diagram of a knot is a maximal portion of the knot diagram where two strands twist around each other, as in figure 7.1.

More precisely, recall that we may consider a diagram of a knot as a 4–valent graph with over–under crossing information at each vertex. A twist region is a string of bigon regions in the diagram graph, arranged end–to–end at their vertices, that is maximal in the sense that there are no additional bigon regions meeting the vertices on either end. A single crossing adjacent to no bigons is also a twist region. We will further restrict so that all twist regions are alternating; if not there is an obvious simplification of the diagram removing crossings from the twist region.

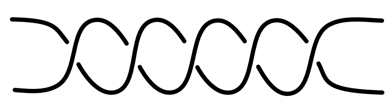

FIGURE 7.1. A twist region of a diagram

The condition that twist regions be maximal ensures that there is only one way to put together exactly two twist regions in a diagram.

DEFINITION 7.2. The *twist knot* $J(2, n)$ is the knot with a diagram consisting of exactly two twist regions, one of which contains two crossings, and the other contains $n \in \mathbb{Z}$ crossings. The direction of crossing depends on the sign of n.

Twist knots $J(2, 2)$, $J(2, 3)$, $J(2, 4)$, and $J(2, 5)$ are shown in figure 7.2.

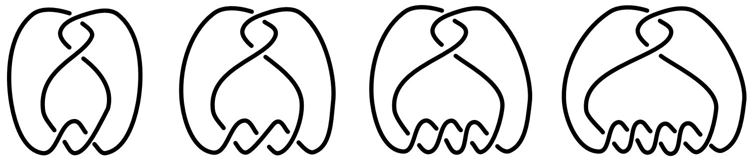

FIGURE 7.2. Twist knots $J(2,2)$ (the figure–8 knot), $J(2,3)$ (the 5_2 knot), $J(2,4)$ (the 6_1 or Stevedore knot), and $J(2,5)$

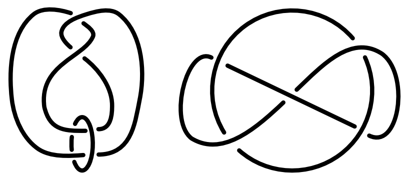

FIGURE 7.3. Two diagrams of the Whitehead link

DEFINITION 7.3. The *Whitehead link* is the link shown in figure 7.3. Note the two links shown are isotopic.

We will show in proposition 7.6 that the complement of the Whitehead link is hyperbolic.

PROPOSITION 7.4. *The complement of the twist knot $J(2,n)$ is obtained by Dehn filling the hyperbolic manifold isometric to the complement of the Whitehead link.*

PROOF. The proof uses topological properties of the sphere S^3 and the solid torus. Recall first that the sphere S^3 is the union of two solid tori whose cores are linked exactly once, but each core alone is unknotted.

The diagram of the Whitehead link on the left of figure 7.3 has a component at the bottom that is unknotted and does not cross itself. The complement of this component in S^3 is a solid torus. Note then that the other component is a knot in a solid torus, as shown on the left of figure 7.4.

Now we apply a homeomorphism to the solid torus, which we view as $S^1 \times D^2$. There is a homeomorphism given by slicing along a disk $\{x\} \times D^2$ of the solid torus, rotating one full time, then gluing back together. This homeomorphism is shown in the center of figure 7.4.

The homeomorphism replaces the original link in the solid torus by a link with two additional crossings. By applying the homeomorphism repeatedly, we see that the complement of the Whitehead link is homeomorphic to the complement of the link with any even number of crossings encircled by the unknotted component. In particular, it is homeomorphic to the complement of the link $J(2,2k) \cup U$, where U is a single unknotted component. By the Mostow–Prasad rigidity theorem (theorem 6.1), these link complements have isometric hyperbolic structures.

To obtain the knot $J(2,2k)$, attach a solid torus to $S^3 - (J(2,2k) \cup U)$, filling in U in a trivial way to give $S^3 - J(2,2k)$. Thus $J(2,2k)$ is obtained

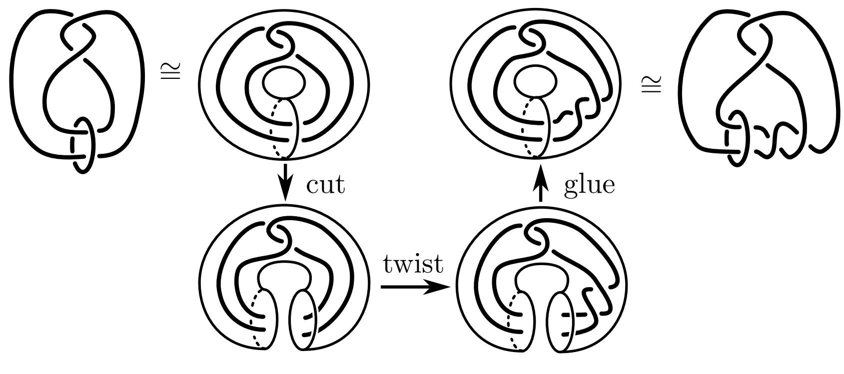

FIGURE 7.4. The Whitehead link complement is homeomorphic to a knot in a solid torus, which we cut, twist, and reglue. The result is homeomorphic to the complement of $J(2,2) \cup U$

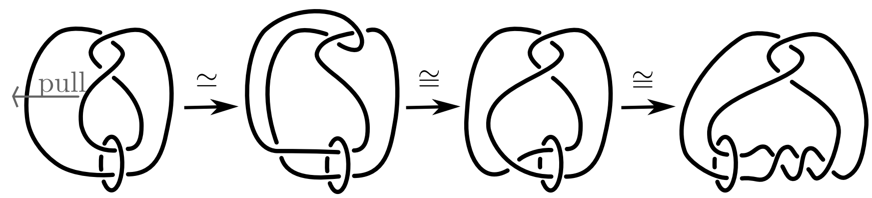

FIGURE 7.5. A sequence of homeomorphisms of the Whitehead link complement

from a manifold isometric to the complement of the Whitehead link by Dehn filling.

So far our proof only works for $J(2,n)$ with n even. Now we consider the case of the knot $J(2,2k+1)$, with odd second component. We may isotope the Whitehead link, starting with the diagram on the left of figure 7.3, to reverse the two crossings at the top, and insert a crossing encircled by the unknotted component at the bottom. This is shown in figure 7.5, left.

Following that figure, we may then reflect the diagram in the plane of projection, reversing all the crossings. This is a homeomorphism of the knot complement, hence an isometry. Now just as in the even case, we may insert any even number of crossings into the two strands encircled by the unknotted component. To obtain $J(2,2k+1)$, simply Dehn fill the unknotted component in the obvious way. □

COROLLARY 7.5. *The complement of the Whitehead link is a geometric limit of* $S^3 - J(2,n)$.

PROOF. Because they are obtained by Dehn filling the complement of the Whitehead link, all but finitely many link complements $S^3 - J(2,n)$ lie in any given neighborhood of infinity in the Dehn surgery space for a cusp of the complement of the Whitehead link. Theorem 6.24 implies that the Whitehead link is therefore a geometric limit of these manifolds. □

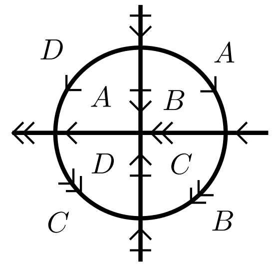

Figure 7.6. Shown is the boundary of an ideal octahedron (one vertex at infinity). Pairing faces as shown gives the complement of the Whitehead link

In order to study the geometry of twist knots, we study the geometry of the geometric limit, the Whitehead link complement.

Proposition 7.6. *The complete hyperbolic structure on the complement of the Whitehead link is obtained by gluing faces of a regular ideal octahedron, with the face pairings as shown in figure 7.6.*

A *regular ideal octahedron* is the ideal octahedron in $\mathbb{H}^3$ with all dihedral angles equal to $\pi/2$.

Proof. The fact that the Whitehead link complement is obtained by face pairings of an ideal octahedron can be readily seen by applying the methods of chapter 1 to the diagram of the Whitehead link on the right of figure 7.3. After collapsing bigons, we obtain two ideal polyhedra with four triangular faces and one quadrilateral face. Glue the quadrilaterals to obtain an ideal octahedron. The form is shown in figure 7.6. We leave the details for exercise 7.2.

In a regular ideal octahedron, all dihedral angles are $\pi/2$, so horospheres intersect a neighborhood of each ideal vertex in a square. We need to check that the face pairings give a hyperbolic structure in this case. Note first that every point in the interior of an octahedron and in the interior of a face of the octahedron has a neighborhood isometric to a ball in $\mathbb{H}^3$. We need to show that each point on an edge also has such a neighborhood, and then lemma 3.6 will imply that the gluing is a manifold with a (possibly incomplete) hyperbolic structure.

Note first that each of the edges (there are three) is glued four times. Thus the total angle around each edge will be $4\pi/2 = 2\pi$. This is not quite enough to show that each point on an edge has a neighborhood isometric to a ball in $\mathbb{H}^3$, because composing the gluings around an edge may introduce nontrivial translation or scale. To show that this does not happen, consider each end of an ideal edge within a cusp. Any horosphere intersects a neighborhood of an ideal vertex of the regular ideal octahedron in a Euclidean square. Under the developing map, squares can only patch together in squares to give a tiling of the universal cover of each cusp by Euclidean squares. There are four squares meeting around a vertex in the cusp corresponding to one of our ideal edges. Note that the squares cannot be scaled or sheared. It follows that edges glue up without shearing singularities, and the structure is hyperbolic.

To show that the structure is complete, we use theorem 4.10: the structure is complete if and only if for each cusp, the induced structure on the boundary is Euclidean. But as already noted, each cusp is tiled by Euclidean squares corresponding to intersections of a horosphere with an ideal vertex of the regular ideal octahedron. Under the developing map, squares can only patch together to give a Euclidean structure: there will be no rotation or scale. Thus the hyperbolic structure must be complete. □

In chapter 9, we will obtain a formula to calculate the volume of a regular hyperbolic ideal octahedron. For now, we state that the volume is a constant $v_{\text{oct}} = 3.66\ldots$.

Corollary 7.7. *The volume of any hyperbolic twist knot is universally bounded*

$$\text{vol}(J(2,n)) \leq v_{\text{oct}},$$

and as $n \to \infty$*,* $\text{vol}(J(2,n)) \to v_{\text{oct}}$.

Proof. The Dehn filling bound follows immediately from Jørgensen's theorem, theorem 6.25. The convergence follows from theorem 6.24. □

We have not yet discussed which twist knots are hyperbolic. We have seen that the figure-8 knot is hyperbolic, and similar methods can be used to show each of the knots in figure 7.2 are hyperbolic. More generally, we will see in chapter 11 (or by other methods in chapter 10) that all twist knots $J(2,n)$ with $n \geq 2$ or $n \leq -3$ are hyperbolic. When $n = 1$ or -2, the standard diagram of $J(2,n)$ can be easily reduced to a diagram with only a single twist region, which is not hyperbolic, and when $n = -1$ its diagram can be easily reduced to that of the unknot, which is also not hyperbolic. All other twist knots are hyperbolic.

7.2. Double twist knots and the Borromean rings

The results of the previous section generalize immediately to knots and links with exactly two twist regions, but with any number of crossings in either twist region.

Definition 7.8. The *double twist knot or link* $J(k,\ell)$ is the knot or link with a diagram consisting of exactly two twist regions, one of which contains k crossings, and the other contains ℓ crossings, for $k,\ell \in \mathbb{Z}$. See figure 7.7. Note that $J(k,\ell)$ is a knot if and only if at least one of k,ℓ is even; otherwise it is a link with two components.

Just as for twist knots, double twist knots are obtained by Dehn filling a simple link complement.

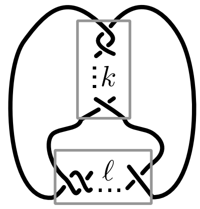

FIGURE 7.7. A double twist knot or link has two twist regions, one with k crossings and one with ℓ crossings

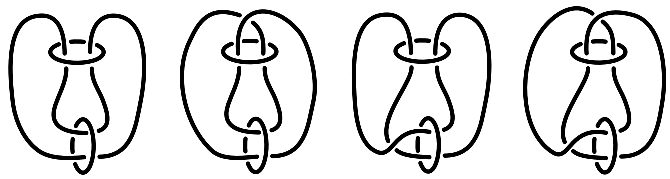

FIGURE 7.8. Complements of $J(k,\ell)$ are obtained by Dehn filling one of these four links. The link on the left is known as the Borromean rings

PROPOSITION 7.9. *The complement of the link $J(k,\ell)$ is obtained by Dehn filling the complement of one of the four links shown in figure 7.8, depending on the parity of k and ℓ.*

PROOF. The proof is nearly identical to that of proposition 7.4, except now it is done in two steps, since there are two unknotted components. Apply a homeomorphism of a solid torus as in figure 7.5 two times. The details are left to the reader. □

The link on the left of figure 7.8 is equivalent to a link more famously known as the *Borromean rings*; its more common diagram is shown in figure 7.18. We will call the other links of figure 7.8 the *Borromean twisted sisters*, and say the links are in the *Borromean family*. In fact, the middle two links are equivalent.

PROPOSITION 7.10. *The complements of the Borromean rings and the Borromean twisted sisters all admit complete hyperbolic structures obtained by gluing two regular ideal octahedra.*

PROOF. Because the Borromean rings has a diagram that is alternating, its complement can be split into ideal polyhedra using the methods of chapter 1. However, we present a new way to decompose links of the Borromean family that we will generalize below.

View the diagrams of figure 7.8 in three dimensions. The two link components in each diagram that will be Dehn filled to produce $J(k,\ell)$ should be viewed as lying perpendicular to the plane of the paper, which is the plane of projection $S^2 \subset S^3$. The other link component(s) should

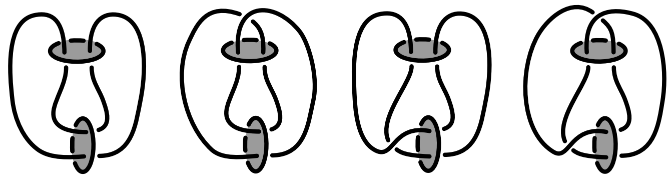

FIGURE 7.9. Shaded 2–punctured disks

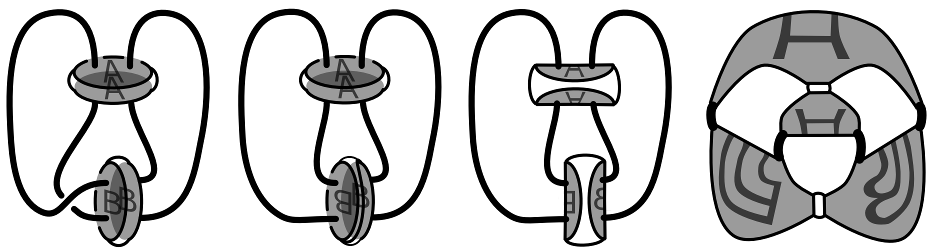

FIGURE 7.10. Left: slice 2–punctured disks up the middle (obtain parallel 2–punctured disks, shown here pulled apart). Middle left: Untwist single crossings. Middle right: Cut along plane of projection. Right: collapse remnants of the link to ideal vertices

be viewed as lying in the plane of projection except at crossings; when the component crosses itself it dips briefly above or below the plane of projection, then returns to the plane.

The components lying perpendicular to the plane of projection are unknotted, and each bounds a 2–punctured disk, shown as shaded in figure 7.9.

As the first step of the decomposition, slice each of these disks up the middle, replacing a single 2–punctured disk with two parallel copies of the 2–punctured disk. This move is shown on the left of figure 7.10.

Now if a 2–punctured disk is adjacent to a crossing in the plane of projection, the next step is to rotate that 2–punctured disk 180° to unwind the crossing, as in the middle left of figure 7.10. Note this rotation pulls the diagram along with it on one side, but the rotation is only performed on the 2–punctured disk adjacent to the crossing, not on the parallel 2–punctured disk. After this step, all crossings in the plane of projection have been removed.

Next, cut along the plane of projection, splitting the complement into two identical pieces as in the middle right of figure 7.10.

Finally, for each piece, collapse remnants of the link to ideal vertices, as on the right of figure 7.10. We claim the result in that figure is topologically an octahedron. To see this, note it has two ideal vertices colored white, coming from crossing circles, and four ideal vertices colored black, coming from the component of the link on the plane of projection. There are four shaded faces that all have three edges, hence all shaded faces are triangles.

There are four white faces, including the one running through the point at infinity in the plane of projection, and each of these white faces also has three edges, so each is a triangle. Thus the result is an ideal octahedron. Recall that there is actually another octahedron coming from our decomposition: the other octahedron comes from the region below the plane of projection after slicing along that plane. So the link complements in the Borromean family all decompose into two ideal octahedra.

Note that the face pairings of the two ideal octahedra that give back the original link complement will be different for the different links; they can be found by tracing backwards through the decomposition process above. To undo the step of cutting along the plane of projection, we glue matching white faces of the opposite octahedra together in pairs. To undo the step of slicing along 2–punctured disks, we glue remaining shaded triangles in pairs; however there are two options depending on whether or not we untwisted a crossing. If both parallel 2–punctured disks were adjacent to no crossings, then corresponding shaded triangles on the same ideal octahedron are glued across an ideal vertex (one of the white vertices of figure 7.10). If there was an adjacent crossing, then a shaded triangle on one octahedron is glued to the opposite shaded triangle on the other octahedron across the (white) ideal vertex.

Finally, to see that these link complements all admit a complete hyperbolic structure, we give each of the two octahedra the geometry of a hyperbolic regular ideal octahedron, then argue as in the proof of proposition 7.6. We check: each edge of the decomposition comes from the intersection of a 2–punctured disk with the plane of projection, and each edge class in the manifold is obtained by gluing four such edges. Thus the total angle around each edge will be $4(\pi/2) = 2\pi$. Again horospheres meet ideal vertices in Euclidean squares, and so the developing image cannot scale, shear, or rotate these squares. Thus edges glue without shearing singularities, and cusps are Euclidean. Hence the result is a complete hyperbolic structure. □

COROLLARY 7.11. *The volume of a double twist knot satisfies*

$$\mathrm{vol}(J(k,\ell)) < 2\,v_{\mathrm{oct}},$$

where $v_{\mathrm{oct}} = 3.66\ldots$ *is the volume of a regular ideal octahedron.*

7.3. Augmenting and highly twisted knots

The above procedure can be generalized.

DEFINITION 7.12. For any twist region of any knot diagram, a new link is obtained by adding a single unknotted link component to the diagram, encircling the two strands of the link component. The link is said to be *augmented*. The added link component is called a *crossing circle*. We will refer to the original link components as *knot strands*.

When a crossing circle is added to each twist region of the diagram, the link is said to be *fully augmented*.

The complement of an augmented link is homeomorphic to the complement of the link with any even number of crossings added to or removed from the twist region, by the same argument illustrated in figure 7.4. Thus the complement of a fully augmented link is homeomorphic to the complement of the fully augmented link with one or zero crossings adjacent to each crossing circle. When there is one crossing adjacent to a crossing circle, we say the crossing circle is adjacent to a *half-twist*.

The four links of figure 7.8 are all examples of fully augmented links. The decomposition of proposition 7.10 goes through more generally for all fully augmented links. This decomposition appears in the appendix to [**Lackenby, 2004**] by Agol and D. Thurston. These links have very beautiful geometric properties, explored further in [**Futer and Purcell, 2007**], [**Purcell, 2007**], [**Purcell, 2008**], and in the survey article [**Purcell, 2011**]. Some of these results are included below, modeled off the exposition in [**Purcell, 2011**].

THEOREM 7.13. *A fully augmented link decomposes into two identical ideal polyhedra with the following properties.*

1. *Faces of the polyhedra can be checkerboard colored. White faces correspond to regions of the plane of projection. Shaded faces are all triangles, and come from 2–punctured disks bounded by crossing circles, which we call* crossing disks.
2. *Ideal vertices are all 4–valent.*
3. *Gluing the polyhedra identifies exactly four edges to a single edge class in the link complement.*

PROOF. The decomposition is obtained very similarly to that in the proof of proposition 7.10. First, each crossing circle bounds a 2–punctured disk, which we shade. Slice along these 2–punctured disks, splitting each into two parallel 2–punctured disks. Next, apply a 180° rotation to those 2–punctured disks adjacent to a crossing, unwinding the crossing. Then slice along the plane of projection, splitting the complement into two identical pieces. Finally, shrink remnants of the link to ideal vertices. We check that each item of the theorem holds.

First, note faces are already checkerboard colored, with shaded faces coming from 2–punctured disks and white faces coming from the plane of projection. Note that edges of the decomposition come from intersections of white and shaded faces. There are exactly three edges bordering each shaded face, so each shaded face is a triangle.

Ideal vertices of the polyhedra come from remnants of the link. For those ideal vertices coming from a component of the link in the plane of projection, the ideal vertex will be adjacent to two edges coming from the 2–punctured disk on one of its ends, and two edges coming from the 2–punctured disk on its other end. Thus it is 4–valent. An ideal vertex coming from a crossing circle is also adjacent to four edges: two from each point where the link component meets the plane of projection.

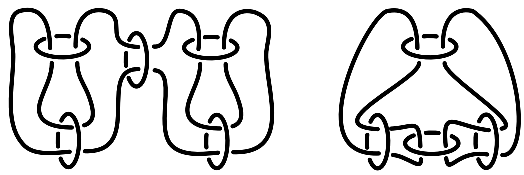

FIGURE 7.11. Left: a fully augmented link with a diagram that is not prime. Right: a fully augmented link that is not reduced. Removing one of the parallel crossing circles will give a reduced link

Finally, note that each edge class contains four edges: two in each polyhedron lying on the parallel copies of the 2–punctured disk. □

Just as with the family of Borromean rings, we can show that many of these links are hyperbolic, but not all. For example, if a fully augmented link has only one crossing circle, then the polyhedral decomposition of theorem 7.13 will have white bigon regions, and collapsing these will collapse the entire polyhedron to a triangle (exercise). Similarly, if there are parallel crossing circles then there will be white bigon faces. We wish to rule these out.

DEFINITION 7.14. A fully augmented link is called *reduced* if the following hold.

(1) Its diagram is connected.
(2) Its diagram is *prime*, i.e. any closed curve meeting the diagram twice bounds a region on one side with no crossings.
(3) None of its crossing circles are parallel. That is, there are no closed curves in the diagram meeting exactly two crossing circles and exactly two white faces. See figure 7.11.

Reduced fully augmented links come from adding crossing circles to links with reduced diagrams as in the sense of the following definition.

DEFINITION 7.15. A diagram is *twist-reduced* if whenever a simple closed curve γ meets the diagram exactly twice in two crossings, running from one side of each crossing to the opposite side, then the curve γ bounds a portion of the diagram containing a twist region.

We will encounter twist-reduced diagrams again, for example in definition 11.10. Meanwhile, the following gives a way of building large numbers of reduced fully augmented links.

LEMMA 7.16. *Let K be a link with a connected, prime, twist-reduced diagram. Then the fully augmented link obtained from K by adding crossing circles to each twist region gives a reduced fully augmented link.*

PROOF. Adding crossing circles to twist regions of a diagram does not change whether it is prime or connected. If the resulting fully augmented

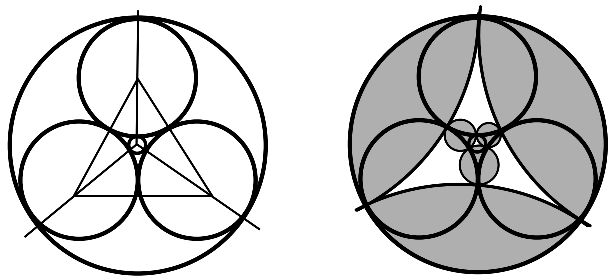

FIGURE 7.12. Left: A circle packing and its intersection graph. Right: Gray circles meeting white circles of a circle packing

link is not reduced, there must be two parallel crossing circles. Thus in the original K, there are two distinct twist regions in the diagram with the property that when crossing circles are added around them, the crossing circles are parallel. Then an isotopy of one of the crossing circles to the other traces out two arcs on the plane of projection disjoint from K. Straightening these, and drawing arcs across K over the crossing circle defines a closed curve in the diagram whose boundary meets the diagram exactly twice, once in each of the two distinct twist regions. Isotope slightly to give a curve in K contradicting the definition of a twist-reduced diagram. □

LEMMA 7.17. *The polyhedra in the decomposition of theorem 7.13 admit a hyperbolic structure in which all dihedral angles are $\pi/2$ provided the fully augmented link is reduced and contains at least two crossing circles.*

The proof of the lemma uses circle packings.

DEFINITION 7.18. A *circle packing* is a connected collection of circles with disjoint interiors. The *intersection graph* of a circle packing is the graph with a vertex at the center of each circle, and an edge between vertices whenever the corresponding circles are tangent.

Figure 7.12 shows an example of a circle packing and most of its intersection graph on the left — the vertex of the intersection graph in the unbounded region has been omitted.

THEOREM 7.19 (Circle packing theorem). *Let G be a finite planar graph that is simple, meaning G has no loops and no multiple edges between a pair of vertices. Then G is (isotopic to) the intersection graph of a circle packing on S^2. If G is a triangulation of S^2, then the circle packing is unique up to Möbius transformation.*

Theorem 7.19 is also known as the Koebe–Andreev–Thurston theorem. It was first proved by Koebe [**Koebe, 1936**]. We will use it here without giving its proof, as the proof is somewhat unrelated to the topic at hand.

PROOF OF LEMMA 7.17. Consider a polyhedron P coming from theorem 7.13 for a reduced fully augmented link with at least two crossing circles.

Edges and vertices of the polyhedron form a graph Γ on S^2. Form a new graph G on S^2 by taking a vertex for each white face of Γ, and an edge between vertices of G if two white faces are adjacent across an ideal vertex of Γ.

If we superimpose G on P, then notice that each region of G will contain exactly one shaded triangular face of P. Thus G is a triangulation of S^2. We show that G has no loops and no multiple edges.

Suppose first that G has a loop. Then the edge of G forming the loop can be superimposed on P to run from a white face, through an ideal vertex of P, then back to the same white face. White faces correspond to regions of the diagram, and ideal vertices correspond to remnants of the link. Thus there is a closed curve γ on the link diagram that runs from a region back to itself crossing over a single component of the link diagram. Because the link diagram consists of closed curves, this is possible only if the curve γ runs along a crossing circle from one white region back to the same white region. Pushing off the crossing circle slightly, this contradicts the fact that the diagram is prime.

Now suppose that the graph G has a multi-edge. Then there is a pair of white faces W_1 and W_2 of P and a pair of ideal vertices v_1 and v_2 such that v_1 and v_2 are both adjacent to W_1 and W_2. Form a loop in P running from W_1 through v_1 to W_2, then back through v_2 to W_1. This loop corresponds to a loop γ in the diagram meeting the regions on the plane of projection corresponding to W_1 and W_2 and meeting two distinct link components between those regions. If the link components came from components on the plane of projection, then this contradicts the fact that the diagram is prime. If the link components came from crossing circles, then it contradicts the fact that the diagram is reduced. If one link component lies in the plane of projection and the other is a crossing circle, then we may slide slightly off the crossing circle to obtain a loop meeting exactly three components in the plane of projection. This is impossible for a closed curve and closed link components.

It follows that G is a finite, simple, planar graph that is a triangulation of S^2. The circle packing theorem, theorem 7.19, implies that there is a unique circle packing of S^2 with G as its intersection graph. View S^2 as the boundary at infinity of $\mathbb{H}^3$. The circle packing of G is then a circle packing on $\partial\mathbb{H}^3$. Each Euclidean circle on $\partial\mathbb{H}^3$ is the boundary of a plane in $\mathbb{H}^3$. Color these planes *white*.

Because the intersection graph of the circle packing is a triangulation, regions complementary to the circle packing meet exactly three circles from the packing. There is a unique Euclidean circle running through the three points of tangency of the circle packing. Again this defines a geodesic plane in $\mathbb{H}^3$. This plane will intersect the white planes at right angles. Color this plane *gray*. See figure 7.12.

For each white plane, remove from $\mathbb{H}^3$ the region bounded by that plane that is disjoint from the other white planes. Similarly for each gray plane.

The result is a right–angled hyperbolic ideal polyhedron that is isomorphic to P, proving the lemma. □

THEOREM 7.20. *The complement of a reduced fully augmented link with at least two crossing circles admits a complete hyperbolic structure, which is obtained by putting a right–angled structure on each of the polyhedra of theorem 7.13.*

PROOF. By lemma 7.17, there exists a right–angled ideal hyperbolic polyhedron with the combinatorics of one of the polyhedron of theorem 7.13. We give each of the polyhedra of theorem 7.13 the hyperbolic structure of this right–angled hyperbolic polyhedron, and glue by corresponding face–pairing isometries to obtain the fully augmented link.

To show this admits a complete hyperbolic structure, we need to show the angle around each edge is 2π, that there is no shearing around edges, and that the cusps are all Euclidean. Because each edge class contains four edges, and each edge has dihedral angle $\pi/2$, the angle sum around each edge is 2π.

Now consider cusps. Any horosphere meets an ideal vertex of the right–angled polyhedron in a rectangle. The developing image of a cusp is obtained by gluing these rectangles according to the gluing isometries on the faces. Note that a white face is glued by a reflection to the identical white face on the opposite polyhedron, so gluing across white sides of a rectangle does not scale or rotate. But then the gluing across shaded faces cannot scale or rotate either. Hence the developing image of each cusp is a tiling of the plane by Euclidean rectangles. Thus around each vertex there cannot be shearing, and the structure on the cusp must be Euclidean. So this gives the complete hyperbolic structure on the fully augmented link. □

COROLLARY 7.21. *In a reduced fully augmented link with at least two crossing circles, each shaded 2–punctured disk bounded by a crossing circle is a totally geodesic surface embedded in the link complement. The* white surface, *obtained by gluing together regions corresponding to regions on the plane of projection (white faces), is also a totally geodesic surface embedded in the hyperbolic link complement. Moreover, these shaded 2-punctured disks and white surfaces meet at right angles whenever they intersect.*

PROOF. In the polyhedral decomposition, these surfaces become white and shaded faces, which are straightened to portions of geodesic planes to obtain the hyperbolic structure. Thus we know that these surfaces are *pleated*, i.e. they decompose into ideal polygons, each of which is totally geodesic. In general pleated surfaces are bent along the edges bounding each polygon, so they are not necessarily totally geodesic. However, in this case, white faces meet shaded faces at angle $\pi/2$, thus in the gluing, white faces glue to white faces with angle π, i.e. no bending, and similarly for shaded faces. If follows that these surfaces are totally geodesic. □

7.4. Cusps of fully augmented links

For many applications in later chapters, it will be useful to know more explicit information about the geometry of fully augmented links, particularly the geometry of their *cusps*. Recall from theorem 4.10 that each cusp admits a Euclidean structure. In this section, we will determine properties of that Euclidean structure for fully augmented links. The exposition is similar to that in [**Futer and Purcell, 2007**].

Consider the universal cover of the complement of a hyperbolic fully augmented link. By corollary 7.21, the universal cover will contain the lift of embedded totally geodesic white surfaces, which will be a collection of disjoint totally geodesic planes that we color white in $\mathbb{H}^3$. It will also contain the lifts of embedded totally geodesic shaded 2-punctured disks bounded by crossing circles. These will also be totally geodesic planes in $\mathbb{H}^3$ and we call them *shaded*. The white planes and shaded planes meet at right angles in $\mathbb{H}^3$. They cut out all the translates of the two ideal polyhedra of theorem 7.13 under the developing map.

Apply an isometry so that the boundary $\widetilde{T}$ of a neighborhood of the point at infinity in $\mathbb{H}^3$ projects under the covering map to a cusp torus T of the fully augmented link. Because each link component meets both white and shaded surfaces, in the universal cover we will see vertical planes corresponding to white and shaded surfaces running into the point at infinity, meeting $\widetilde{T}$ in a rectangular lattice.

If we forget the fact that the edges of the lattice have lengths, but consider each rectangle on $\widetilde{T}$ as a topological object with two opposite shaded sides and two opposite white sides, then we obtain the following.

LEMMA 7.22. *Let T be a cusp torus of a fully augmented link, with universal cover $\widetilde{T}$ tiled by rectangles coming from white and shaded surfaces. Let s denote a step along a shaded surface between two white surfaces, and let w denote a step along a white surface between two shaded ones. Then a fundamental domain for T is given as follows.*

- *If T comes from a crossing circle without a half-twist, then it has meridian w and longitude $2s$.*
- *If T comes from a crossing circle with a half-twist, it has meridian $w \pm s$ (depending on the direction of the twist) and longitude $2s$.*
- *If T comes from a knot strand, i.e. a component that is not a crossing circle, then it has meridian $2s$ and longitude $nw + ks$, where n is the number of twist regions met by the strand, with multiplicity, and k is some integer.*

PROOF. From the construction of the polyhedral decomposition of $S^3 - L$, each crossing circle gives rise to an ideal vertex of each polyehdron. Thus a fundamental domain for a crossing circle consists of two rectangles, given by neighborhoods of the corresponding 4-valent ideal vertices.

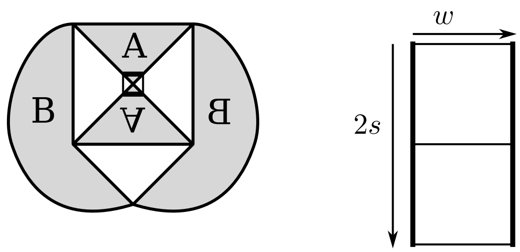

FIGURE 7.13. A fundamental region for a crossing circle.

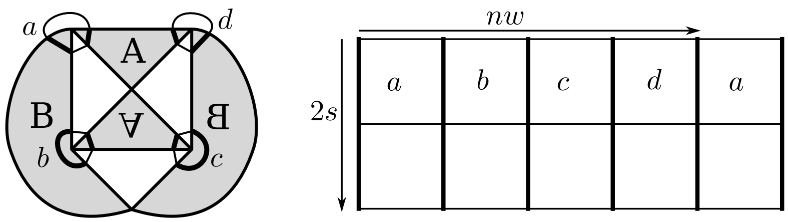

FIGURE 7.14. A fundamental region for a knot strand with no half twists.

In the case that there are no half-twists, the shaded faces adjacent to the ideal vertex are glued to each other. Thus an arc running along a white face has its endpoints glued into a meridian, and thus the meridian in this case is w. As for the longitude, a white face on one polyhedron is glued to a white face on the other. Thus a longitude steps along two shaded sides, one on one polyhedron and one on the other, before closing up. See figure 7.13.

For a knot strand K meeting no half-twists, there will be one ideal vertex of one polyhedron, hence one rectangular vertex neighborhood, for each portion of K between adjacent crossing circles. These rectangles are glued end to end along shaded faces coming from the crossing disks to complete a longitude. Thus there will be n such rectangles, and a longitude is given by n steps along white faces, or nw. There will be n identical rectangles glued end to end in the other polyhedron. These two blocks of n rectangles will be glued along their white faces to form a $2 \times n$ block, making up the fundamental domain of K. A meridian is given by two steps along shaded faces. See figure 7.14.

If there are half-twists, then the gluing changes along shaded faces at half-twists. A shaded triangle on one polyhedron will be glued to the opposite shaded triangle on the other polyehdron. This introduces shearing into the fundamental domain, as in figure 7.15. Since the shearing only occurs as shaded faces are glued, it does not affect the longitude of a crossing circle or the meridian of a knot strand: these are both $2s$. However, it will adjust a meridian of a crossing circle by adding $\pm s$, and it will adjust the longitude of a knot strand by adding $\pm s$ for each half-twist. Thus the longitude of a knot strand becomes $nw + ks$ for some integer k. $\square$

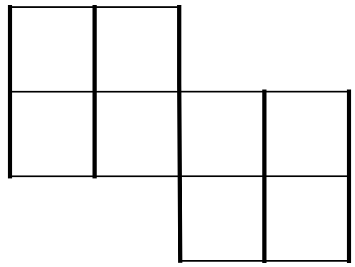

FIGURE 7.15. Adding a half twist shifts the gluing along the shaded faces, shearing the fundamental domain.

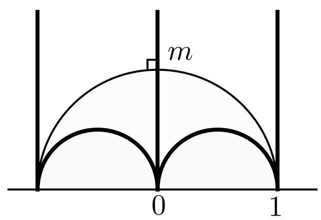

FIGURE 7.16. When an ideal triangle in $\mathbb{H}^2$ has vertices at 0, 1, and ∞, one of its midpoints will lie in $\mathbb{H}^2$ at height 1.

Lemma 7.22 is purely topological. We now wish to give geometric information on the rectangles forming the cusps. To do so, we need to find more explicit embedded cusp neighborhodds of the cusps of a fully augmented link. An embedded cusp neighborhood lifts to a disjoint collection of horoballs in the universal cover $\mathbb{H}^3$, one for each ideal vertex of each translate of the ideal polyhedra under the developing map. We will find an embedded cusp neighborhood by finding a collection of embedded horoballs about ideal vertices of the polyhedra forming a fully augmented link.

DEFINITION 7.23. Let $T \subset \mathbb{H}^3$ be an ideal triangle. For each edge e of T, define the *midpoint* m of e to be the point such that the geodesic from m to the opposite ideal vertex is perpendicular to e. Note this point is unique. See figure 7.16.

For each edge e of the ideal polyhedral decomposition of a fully augmented link, define its *midpoint* to be the midpoint of that edge on one of the two ideal triangles adjacent to the edge. Note that since the two polyhedra are symmetric by a reflection in the white faces, both triangles adjacent to e have the same midpoint, so the midpoint of each edge is well-defined.

LEMMA 7.24. *Let L be a hyperbolic fully augmented link, with decomposition into ideal polyhedra P_1 and P_2. For each ideal vertex of P_i, there is a unique horoball meeting the midpoint of each edge through that ideal vertex. The collection of all such horoballs, intersected with P_i and P_j, glue to give an embedded cusp neighborhood of all the cusps of $S^3 - L$.*

PROOF. Place P_i in $\mathbb{H}^3$ so that the ideal vertex of interest lies at infinity, and so that one of the two shaded faces meeting the ideal vertex has its ideal vertices at 0, 1, and ∞ in $\mathbb{H}^3$. Note that the edges of that shaded face have midpoints at height 1, as in figure 7.16. Because the polyhedron is right-angled, the other shaded face will have ideal points at some points ci, $1+ci$, and ∞ for some $c \in \mathbb{R}$. Thus again the midpoints of these edges lie at height 1. Then the horoball of height 1 about infinity meets the midpoint of each edge through the ideal vertex.

The above discussion applies to any vertex of any polyhedron, and so this proves the first statement of the lemma. However, to show that these horoballs glue to give an embedded cusp neighborhood, we need to show that under the developing map, these horoballs have disjoint embedded interiors in $\mathbb{H}^3$.

Develop in a neighborhood of infinity. Since white faces are glued by reflection, the developing map takes P_i to a reflected copy of P_i, where the reflection is through the vertical plane determined by a white face meeting infinity. Note that the reflection isometry takes points at height 1 to points at height 1. Similarly, because the polyhedra are right angled, developing by gluing shaded faces produces shaded faces of the same width, and thus midpoints are height 1. Thus a horoball of height 1 through infinity will meet the midpoints of all edges through infinity under the developing image.

We claim that this horoball cannot meet any white faces besides those that have an ideal vertex at infinity. Consider a white face that does not meet infinity. It lies in a hemisphere with boundary a circle C on $\mathbb{C}$. Because the white surface is embedded in the fully augmented link complement, the lifts of this surface are disjointly embedded in $\mathbb{H}^3$. Thus the boundary circles of all lifts of white faces meet only at points of tangency corresponding to ideal vertices. Thus the circle C meets the boundaries of the vertical planes containing white faces meeting ∞ only in points of tangency. The vertical planes have boundary on $\mathbb{C}$ a collection of parallel vertical lines, and these lines must be exactly distance 1 apart. Then the diameter of C can be at most 1. It follows that the height of the hemisphere containing a white face that does not run through infinity must be at most 1/2; therefore the horoball at height 1/2 cannot meet it.

By an isometry, the previous argument applies to any ideal vertex. Thus we have proved that the horoballs through the midpoints of ideal vertices of P_i only meet white faces that run through the center of the horoball at infinity.

Suppose, by way of contradiction, that under the developing map one of these horoballs H centered at a point p in $\mathbb{C}$ has diameter strictly greater than 1, so that the collection of interiors of horoballs will not be embedded. Then p must lie on the boundary of one of the vertical white planes, else H intersects a vertical white plane in a compact region, giving a contradiction.

So p is an ideal vertex of a polyhedron meeting a white face on a vertical plane V, and some other white plane W. The boundary ∂W is a circle on

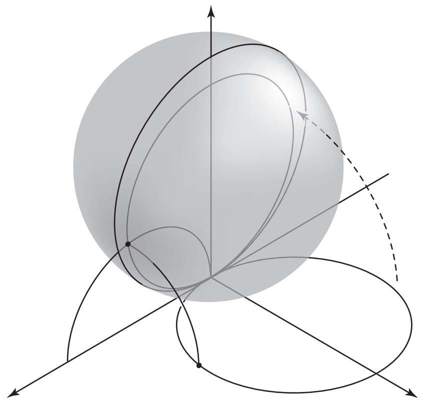

FIGURE 7.17. If H is centered at a point p on a vertical white plane, and has diameter greater than 1, then it must contain the midpoint of an edge through p. Figure from [**Futer and Purcell, 2007**].

$\mathbb{C}$ of diameter at most 1, as we have seen above. The vertex p also meets two shaded faces, and at least one of these, call it S, is not a vertical plane. Then $S \cap V$ is an ideal edge of a polyhedron, and it must have a midpoint. The midpoint is obtained by taking a perpendicular from a point on ∂W to the semicircle $S \cap V$ on the vertical plane V. The set of all points obtained by dropping a perpendicular from ∂W to V is a circle of diameter equal to the diameter of ∂W on the plane V; see figure 7.17. But H has diameter greater than 1, so this entire circle lies inside of H. This contradicts the fact that H does not contain any of the midpoints of edges through p.

Thus when we expand all horoballs to the midpoints of their adjacent edges, all those centered at points on $\mathbb{C}$ have diameter at most 1, while that at infinity has height exactly 1, so their interiors are embedded. Since the above discussion applies to any ideal vertex of any polyhedron, we conclude that under the developing map, interiors of all such horoballs are embedded, and thus the quotient under the covering map gives an embedded horoball neighborhood of each cusp of $S^3 - L$. □

COROLLARY 7.25. *Let L be a hyperbolic fully augmented link. There exists an embedded horoball neighborhood of the cusps of $S^3 - L$ such that, when measured in the induced Euclidean metric on the boundary of each cusp, the sides of the steps s and w (of lemma 7.22) have lengths $\ell(s) = 1$ and $\ell(w) \geq 1$.*

PROOF. If we place the ideal vertices of a shaded triangle at 0, 1, and ∞, then the midpoints of the edges from 0 to ∞ and from 1 to ∞ are of height 1. Thus the horoball neighborhood through these points is at height 1, and distance along the boundary of this horoball is just Euclidean distance.

Since the shaded triangle meets this plane in a line segment from 0 to 1, the length of the step s is $\ell(s) = 1$.

To find w, we note that there will be horoballs of diameter 1 centered at all the corners of the rectangle containing the step w. Two of these will be centered at 0 and 1, the other two at some ci and $1 + ci$ in $\mathbb{C}$, for some $c = \ell(w)$. Because the four horoballs are disjoint, we must have $\ell(w) \geq 1$. □

The above results lead to consequences on slope lengths of Dehn fillings.

THEOREM 7.26. *Let L be a hyperbolic fully augmented link, and let $C_1, \dots, C_k$ be crossing circles of L. Let s_j be a slope on $N(C_j)$ such that Dehn filling along s_j replaces the crossing circle C_j by a twist region with n_j crossings (with n_j even if and only if C_j is not adjacent to a half-twist). Then there is an embedded horoball neighborhood of all cusps of $S^3 - L$ such that on the boundary of each cusp, the length of s_j is at least $\ell(s_j) \geq \sqrt{n_j^2 + 1}$.*

PROOF. The slope of the Dehn filling that replaces a crossing circle with $2a_j$ crossings runs over one meridian and a_j longitudes.

If $n_j = 2a_j$ is even, then C_j meets no half-twist. Then lemma 7.22 implies that the slope s_j will have the form $w + 2a_j s$ or $w - 2a_j s$. Because w and s run in orthogonal directions, and each has length at least one by corollary 7.25, the length of $w \pm 2a_j s$ is at least $\sqrt{1 + (2a_j)^2} = \sqrt{1 + n_j^2}$, as claimed.

If $n_j = 2a_j + 1$ is odd, then C_j meets a half-twist, and lemma 7.22 implies that s_j has the form $w \pm s + 2a_j s$ or $w \pm s - 2a_j s$, with the signs the same: $s_j = w \pm (2a_j + 1)s$. Again because w and s are orthogonal, corollary 7.25 implies the length of s_j is at least $\sqrt{1 + (2a_j + 1)^2} = \sqrt{1 + n_j^2}$. □

7.5. Exercises

EXERCISE 7.1. In this exercise, you investigate the two diagrams of the Whitehead link shown in figure 7.3.

(1) Show by a sequence of diagrams that the two links in that figure are isotopic.
(2) Use SnapPy [**Culler et al., 2016**] to show that the two link complements are isometric. The check using SnapPy is not mathematically rigorous, but in this case the link has a special property: it is *arithmetic*. We will not define an arithmetic link here (we won't use the definition elsewhere), but a consequence of arithmeticity is that the program Snap [**Coulson et al., 2000**] can be used to give a mathematically rigorous certification that the two links shown are isometric.

EXERCISE 7.2. Use the methods of chapter 1 to prove that the complement of a Whitehead link can be decomposed into two ideal pyramids with a square base, which in turn can be glued to an ideal octahedron.

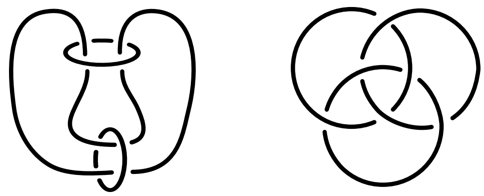

FIGURE 7.18. Two different diagrams of the Borromean rings are shown

EXERCISE 7.3. Shown on the left of figure 7.18 is the diagram of a link which we claim is the Borromean rings. Shown on the right is a more familiar diagram of the Borromean rings. There are several different ways to prove the complements of these hyperbolic manifolds are isometric.

(1) Show by a sequence of diagram moves that the links are isotopic. Why does this suffice to show the complements are isometric?
(2) Find a hyperbolic structure on each by hand, and show by hand that the manifolds are isometric. This will take some work, and sounds tedious. The exercise here is to think about why this will be tedious: list the steps involved.
(3) Use computational tools. Use SnapPy [**Culler et al., 2016**] to show they are isometric. As in exercise 7.1, these links are arithmetic, so you can check using Snap [**Coulson et al., 2000**] that the link complements are isometric, which gives a mathematically rigorous certification.

EXERCISE 7.4. The (p, q, r)–pretzel link. As p, q, r go to infinity, find geometric limits of pretzel links. Find a universal upper bound on their volumes.

EXERCISE 7.5. (Topology of the solid torus) A solid torus V is homeomorphic to $S^1 \times D^2$, where a specified homeomorphism $h\colon\ S^1 \times D^2 \to V$ is called a *framing*.

(a) A non-trivial simple closed curve in ∂V is called a *meridian* if it bounds a disk in V. Prove that if μ is a meridian, then for some framing $h\colon\ S^1 \times D^2 \to V$, $\mu = h(\{1\} \times \partial D^2)$.
(b) A non-trivial simple closed curve λ in ∂V is called a *longitude* if it represents a generator of $\pi_1(V) \cong \mathbb{Z}$. Prove that if λ is a longitude, then for some framing $h\colon\ S^1 \times D^2 \to V$, $\lambda = h(S^1 \times \{1\})$.
(c) Prove that there are infinitely many ambient isotopy classes of longitudes in a solid torus.

EXERCISE 7.6. For the Whitehead link, find slopes of Dehn filling giving the twist knot $J(2, n)$ for n even. Write them as $p\mu + q\lambda$, for relatively prime integers p and q, where μ is a meridian and λ is the longitude that bounds a disk in S^3. This is called the *standard longitude*.

Repeat for n odd, using the isometric link.

Exercise 7.7. Using the meridian and standard longitude as a basis for two boundary components of the exterior of the Borromean rings (i.e. take a longitude on each component that bounds a disk in S^3), find the slopes of the Dehn fillings of the Borromean rings that give $J(2k, 2\ell)$.

Repeat for $J(2k, 2\ell+1)$ and $J(2k+1, 2\ell+1)$.

Exercise 7.8. The simplest fully augmented link has a single crossing circle; it comes from augmenting a knot with only one twist region. Show that when we apply the decomposition of this chapter to the fully augmented link with only one crossing circle, the result is not a decomposition into two ideal polyhedra. What does the decomposition give?

Exercise 7.9. Prove a result analogous to theorem 7.26 for knot strand cusps. If K_i is a knot strand cusp of a hyperbolic fully augmented link, and s_i is a slope on K_i that represents a nontrivial filling (i.e. s_i is not a meridian), then the length of s_i is at least m_i, where m_i denotes the number of crossing disks that K_i intersects, counted with multiplicity.

Exercise 7.10. In chapter 10 we will consider a class of links called *two-bridge links* which have twist regions arranged in two rows, illustrated in figure 10.2. Show that the complement of the fully augmented link coming from a 2-bridge link can be obtained by gluing a collection of regular ideal octahedra. How many regular ideal octahedra?