CHAPTER 4

Hyperbolic structures and triangulations

In chapter 3, we learned that hyperbolic structures lead to developing maps and holonomy, and that the developing map is a covering map if and only if the hyperbolic stucture is complete.

In this chapter, we wish to compute explicit complete hyperbolic structures on 3-manifolds, again with our primary examples being knot complements. One of the most straightforward ways to find a hyperbolic structure is to first triangulate the manifold, or subdivide it into tetrahedra, and then to put a hyperbolic structure on each tetrahedron, ensuring the tetrahedra glue to give a $(\mathrm{PSL}(2,\mathbb{C}),\mathbb{H}^3)$-structure whose developing map is complete. This method of computing hyperbolic structures has been studied by many, and in particular was implemented on the computer by J. Weeks as part of his 1985 PhD thesis [**Weeks, 1985**]. Here we will describe the conditions required to obtain a complete hyperbolic structure via triangulations, and as usual, work through examples.

4.1. Geometric triangulations

In chapter 3, we defined topological and geometric polygonal decompositions of 2-manifolds. We can extend these notions to 3-manifolds by considering decompositions into ideal polyhedra. In chapter 1, we obtained topological ideal polyhedral decompositions for knot complements. For many applications, including those later in this chapter, it simplifies matters greatly to consider decompositions into ideal tetrahedra.

Definition 4.1. Let M be a 3-manifold. A *topological ideal triangulation* of M is a combinatorial way of gluing truncated tetrahedra (ideal tetrahedra) so that the result is homeomorphic to M. Truncated parts will correspond to the boundary of M. As before, a gluing should take faces to faces, edges to edges, etc.

Example 4.2. The figure-8 knot has a topological ideal triangulation consisting of two ideal tetrahedra, as we saw in exercise 1.7 in chapter 1.

For a given knot complement, it is relatively easy to find topological ideal triangulations. For example, starting with any polyhedral decomposition,

March 29, 2018

choose an ideal vertex v and cone to that vertex: i.e. add edges between v and all other ideal vertices, between any two edges meeting v add an ideal triangle (adding an additional edge opposite v if necessary), and between three triangles meeting v add an ideal tetrahedron. Split off the resulting tetrahedra. This reduces the collection of polyhedra to a collection with at least one fewer ideal vertex. Hence after repeating a finite number of times, we are left with a collection of topological tetrahedra.

4.1.1. An extended example: the 6_1 knot. We work out an example for the 6_1 knot carefully. We will see how to decompose the complement into five tetrahedra. (In fact, the complement of the 6_1 knot can be decomposed into four tetrahedra, but we won't bother simplifying further here.)

We start with a polyhedral decomposition of the 6_1 knot. We use the decomposition obtained using the methods of chapter 1. The result is shown in figure 4.1, with the knot on the left, the top polyhedron in the center, and the bottom polyhedron on the right. Recall all polyhedra are viewed from the outside; that is the ball of the polyhedron is behind the projection plane in each figure. In this example, oriented edges are labeled 1 through 6.

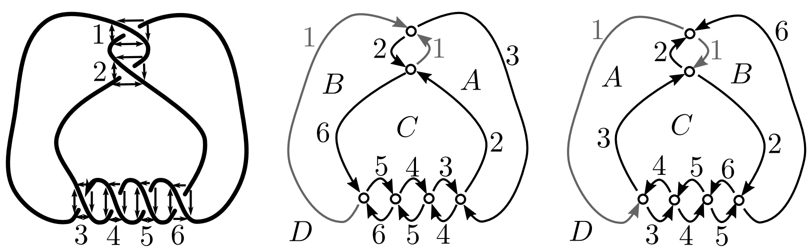

FIGURE 4.1. Left to right: The 6_1 knot, the top polyhedron, the bottom polyhedron

Collapse all bigons, identifying edges 1 and 2, and 3 through 6. New edges and orientations are shown in figure 4.2.

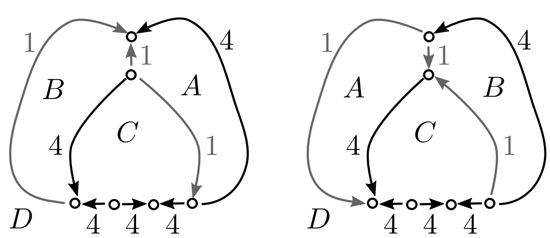

FIGURE 4.2. Polyhedra for 6_1 knot with bigons collapsed

We cone the top polyhedron to the vertex in the center. This subdivides faces C and D into triangles, shown in figure 4.3 in both top and bottom polyhedra.

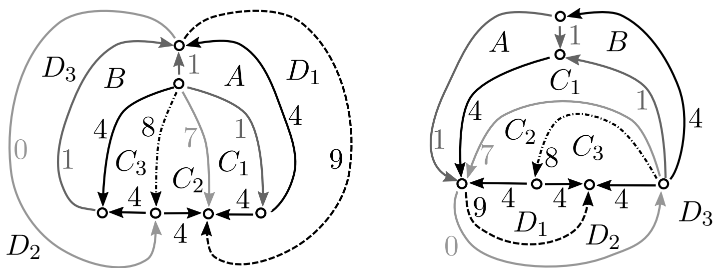

FIGURE 4.3. A subdivision of faces C and D in the top polyhedron (left) leads to a subdivision of the bottom (right)

Continuing the subdivision in the top polyhedron, two edges meeting in the center vertex bound an ideal triangle; three triangles bound a tetrahedron. Thus edges labeled 1, 7, 9 bound an ideal triangle E_1; edges labeled 1, 8, 0 bound an ideal triangle E_2. Triangles A, C_1, D_1, and E_1 bound an ideal tetrahedron, as do triangles B, C_3, D_3, and E_2. When we split off these tetrahedra a single tetrahedron remains. All tetrahedra making up the top polyhedron are shown in figure 4.4.

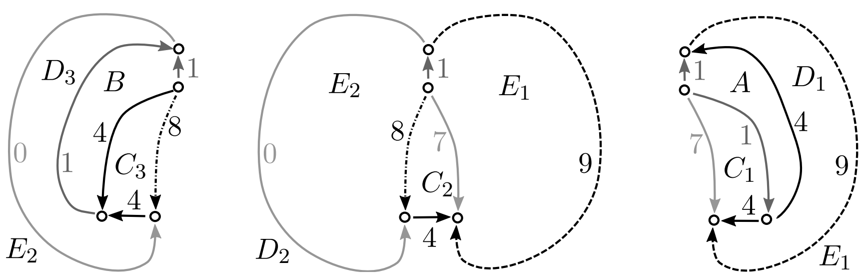

FIGURE 4.4. The top polyhedron splits into the three tetrahedra shown

Now we split the bottom polyhedron into tetrahedra. However, first, observe in figure 4.3 that edges labeled 7 and 0 in the bottom polyhedron run between the same two ideal vertices. Thus these two edges should be flattened and identified in the bottom polyhedron. While we could do that now in one step, we believe it is more geometrically clear how to flatten and identify if we first cut off ideal tetrahedra from the bottom polyhedron.

So first, note there will be an ideal triangle E_3 with edges labeled 4, 7, and 1, and this cuts off an ideal tetrahedron with sides A, B, C_1, E_3. Similarly there is an ideal triangle E_4 with edges 7, 9, and 4, cutting off an ideal tetrahedron with sides C_2, C_3, E_4, and D_1. These two tetrahedra, as well as the remnant of the bottom polyhedron, are shown in figure 4.5.

Notice that the object on the right of figure 4.5 is not a tetrahedron: edges labeled 7 and 0 in that polyhedron form a bigon, which collapses to a single edge which we label 7. When we do the collapse, the faces E_4 and D_2

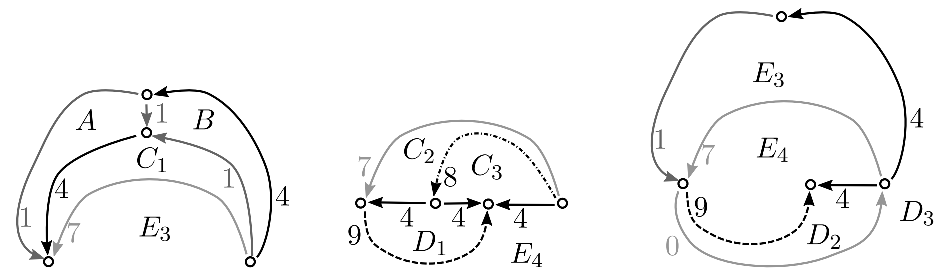

FIGURE 4.5. Splitting off two tetrahedron in the bottom polyhedron

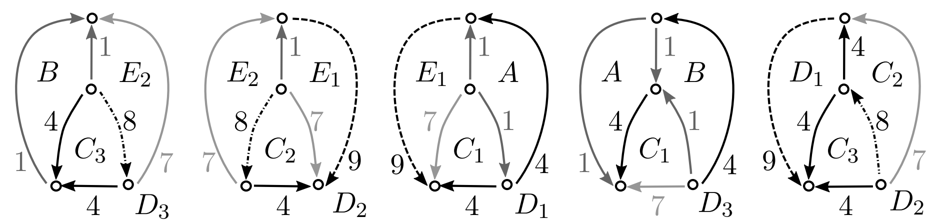

FIGURE 4.6. Five tetrahedra which glue to give the complement of the 6_1 knot

collapse to a single triangle, which we will label D_2. The faces E_3 and D_3 also collapse to a single triangle, which we will label D_3.

When we have finished, we have five tetrahedra that glue to give the complement of the 6_1 knot. All five tetrahedra with their edges and faces labeled are shown in figure 4.6.

4.1.2. Geometric ideal triangulations.

DEFINITION 4.3. A *geometric ideal triangulation* of M is a topological ideal triangulation such that each tetrahedron has a (positively oriented) hyperbolic structure, and the result of gluing is a smooth manifold with a complete metric.

As of the writing of this book, it is still an open question as to whether every hyperbolic 3-manifold actually admits a geometric ideal triangulation. It is known that every cusped hyperbolic 3-manifold can be decomposed into convex ideal polyhedra [**Epstein and Penner, 1988**]: we will go through this in chapter 14. However, subdividing this decomposition into tetrahedra may create degenerate tetrahedra — actual topological tetrahedra (as opposed to the object on the right of figure 4.5), but tetrahedra that are flat in the hyperbolic structure on M. There are known examples of generalized spaces with singularities that do not admit geometric triangulations [**Choi, 2004**].

4.2. Edge gluing equations

In chapter 3, we saw that a gluing of hyperbolic polygons has a hyperbolic structure if and only if the angle sum around each finite vertex is 2π (lemma 3.7). There are similar conditions for a gluing of hyperbolic tetrahedra. We now need to consider gluing around an edge.

Let T be an ideal tetrahedron embedded in $\mathbb{H}^3$. Any ideal tetrahedron has six edges. If we select any one, say e, we may choose an isometry of $\mathbb{H}^3$ taking the endpoints of e to 0 and ∞, and sending a third vertex to $1 \in \mathbb{C} \subset \partial_\infty \mathbb{H}^3$. This choice uniquely determines the isometry. The fourth vertex of T will be mapped to some $z' \in \mathbb{C}$. We may assume that z' has positive imaginary part, for if not, apply an isometry of $\mathbb{H}^3$ rotating around the geodesic from 0 to ∞ and rescaling so that z' maps to 1. In this case, the image of 1 under this isometry will be a complex number with positive imaginary part.

Definition 4.4. For an ideal tetrahedron T embedded in $\mathbb{H}^3$, and edge e of that tetrahedron, define the number $z(e)$ in $\mathbb{C}$ to be the complex number with positive imaginary part obtained by applying the unique isometry of $\mathbb{H}^3$ that takes the vertices of e to 0 and ∞, takes another vertex to 1, and takes the final vertex of T to $z(e)$. This is called the *edge invariant* of e.

Remark 4.5. Note that it is possible to map an ideal tetrahedron to $\mathbb{H}^3$ so that three vertices map to 0, ∞, and 1, and the fourth maps to a point on the real line. In this case, the tetrahedron produced does not have a hyperbolic structure. If the fourth vertex is not 0 or 1, it is said to be *flat*. If the fourth vertex is 0 or 1, it is *degenerate*. Similarly, a fourth vertex mapped to infinity is a degenerate tetrahedron. An ideal triangulation of a hyperbolic 3-manifold with flat or degenerate tetrahedra is not a geometric ideal triangulation. When looking for geometric triangulations, we must rule out such tetrahedra. Similarly, for geometric triangulations, all edge invariants of all tetrahedra must have positive imaginary part. This ensures the tetrahedra are *positively oriented*. Finally, the procedure above always chooses an edge invariant with positive imaginary part. However, when we glue many tetrahedra together, at times it is impossible to simultaneously choose all edge invariants to have positive imaginary part; some may have negative negative imaginary part. Such a tetrahedron is a *negatively oriented* tetrahedron.

Edge invariants of an ideal tetrahedron determine each other, in the following way.

Lemma 4.6. *Let T be an ideal tetrahedron with edge e_1, mapped so that vertices of T lie at ∞, 0, 1, and $z(e_1)$. Then T has the following additional edge invariants.*

- *The edge e_1' opposite e_1, with vertices 1 and $z(e_1)$, has edge invariant $z(e_1') = z(e_1)$.*
- *The edge e_2 with vertices ∞ and 1 has edge invariant*

$$z(e_2) = \frac{1}{1 - z(e_1)}.$$

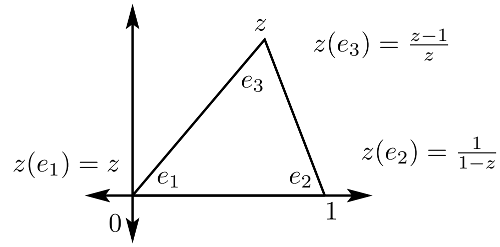

FIGURE 4.7. Edge invariants

- *The edge e_3 with vertices ∞ and $z(e_1)$ has edge invariant*
$$z(e_3) = \frac{z(e_1)-1}{z(e_1)}.$$

Thus we have the following relationships for these edge invariants.
$$z(e_1)z(e_2)z(e_3) = -1, \quad \textit{and} \quad 1 - z(e_1) + z(e_1)z(e_3) = 0$$

PROOF. The proof is obtained by considering isometries of $\mathbb{H}^3$ that move the different edges of T onto the geodesic from 0 to ∞. For ease of notation, we set $z = z(e_1)$.

For the first part, we label one more edge. Let e_3' be the edge of T opposite e_3. So e_3' has endpoints 0 and 1. Note there is a geodesic γ in $\mathbb{H}^3$ that meets the edges e_3 and e_3' orthogonally. An elliptic isometry rotating about γ by angle π maps 0 to 1 and 1 to 0, and maps ∞ to z and z to ∞, thus it preserves T. It takes the edge e_1' with endpoints 1 and z to an edge with endpoints 0 and ∞. Hence $z(e_1') = z$.

To determine $z(e_2)$, we apply a Möbius transformation fixing ∞, taking 1 to 0, and taking z to 1. This transformation is given by
$$w \mapsto \frac{w-1}{z-1}.$$
It sends 0 to $-1/(z-1)$. Thus $z(e_2) = 1/(1-z)$.

As for the edge e_3 running from z to ∞, to determine its edge invariant we apply a Möbius transformation fixing ∞, sending z to 0, and sending 0 to 1. This is given by
$$w \mapsto \frac{w-z}{-z}.$$
It sends 1 to $(1-z)/(-z)$. Thus $z(e_3) = (z-1)/z$. □

The three edge invariants of a tetrahedron are shown in figure 4.7.

Now consider a gluing of ideal tetrahedra. Fix an edge e of the gluing, and let T_1 be a tetrahedron which has edge e_1 glued to e. Put T_1 in $\mathbb{H}^3$ with the edge e_1 running from 0 to ∞, with a third vertex at 1, and the fourth vertex at $z(e_1)$, where $z(e_1)$ has positive imaginary part. The gluing identifies each face of T_1 with another face. Let F_1 denote the face of T_1

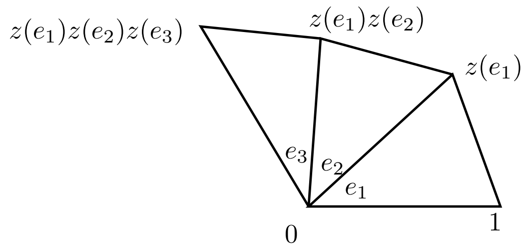

FIGURE 4.8. Vertices of attached triangles.

with vertices 0, $z(e_1)$, and ∞. This is glued to a face F_1' in some tetrahedron T_2, where the edge e_2 in T_2 glues to e.

Now, we could put T_2 in $\mathbb{H}^3$ with vertices at 0, ∞, 1, and $z(e_2)$, but since we're gluing to T_1, we want the face F_1' to have vertices 0, ∞, and $z(e_1)$ rather than vertices 0, ∞, and 1. Thus to do the gluing, we apply an isometry of $\mathbb{H}^3$ fixing 0 and ∞, mapping 1 to $z(e_1)$. This takes the fourth vertex of T_2 to $z(e_1)z(e_2)$.

Continue attaching tetrahedra counterclockwise around e. The next tetrahedron attached will have vertices 0, ∞, $z(e_1)z(e_2)$, and $z(e_1)z(e_2)z(e_3) \in \mathbb{C}$. See figure 4.8. Eventually one of the tetrahedra will be glued to T_1 again. The fourth vertex of the final tetrahedron will be at $z(e_1)z(e_2)\cdots z(e_n)$.

THEOREM 4.7 (Edge gluing equations). *Let M^3 admit a topological ideal triangulation such that each ideal tetrahedron has a hyperbolic structure. The hyperbolic structures on the ideal tetrahedra induce a hyperbolic structure on the gluing, M, if and only if for each edge e,*

$$\prod z(e_i) = 1 \quad \text{and} \quad \sum \arg(z(e_i)) = 2\pi,$$

where the product and sum are over all edges that glue to e.

PROOF. The hyperbolic structure on the tetrahedra induces a hyperbolic structure on M if and only if every point in M has a neighborhood isometric to a ball in $\mathbb{H}^3$, by lemma 3.6. Consider a point on an edge. If it has a neighborhood isometric to a ball in $\mathbb{H}^3$ then the sum of the dihedral angles around the edge must be 2π. See figure 4.9. This sum of dihedral angles is $\sum \arg(z(e_i))$. Moreover there must be no nontrivial translation as we move around the edge. Since the last face of the last triangle glues to the triangle with vertices 0, 1, and ∞, this condition requires that $\prod z(e_i) = 1$.

Conversely, if we have $\prod z(e_i) = 1$ and $\sum \arg(z(e_i)) = 2\pi$, then any point on the edge under the gluing has a ball neighborhood isometric to a ball in $\mathbb{H}^3$. $\square$

The equations $\prod z(e_i) = 1$ (and restrictions $\sum \arg(z(e_i)) = 2\pi$) are called the *edge gluing equations*. We have one for each edge. However, since the three edge invariants of a tetrahedron are all determined by a single edge

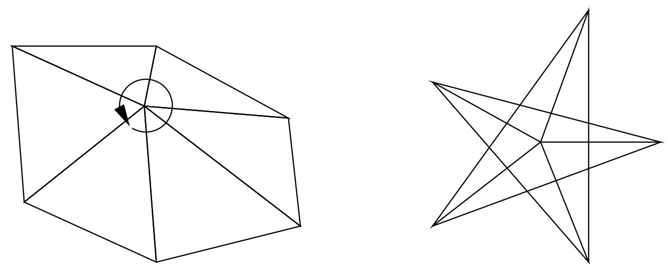

FIGURE 4.9. Left: Angle sum must be 2π. Right: An example of why this condition is important.

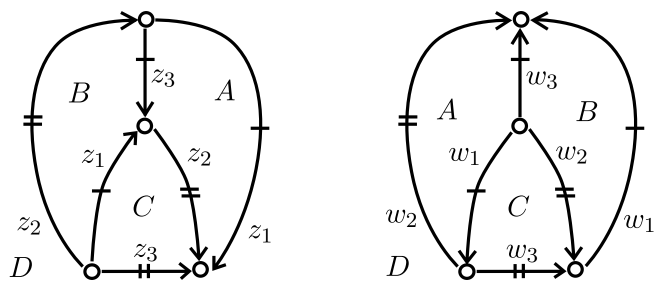

FIGURE 4.10. The ideal tetrahedra of the figure-8 knot complement.

invariant, lemma 4.6, one ideal tetrahedron contributes at most one unknown to the gluing equations.

EXAMPLE 4.8 (Edge gluing equations for the figure-8 knot). The figure-8 knot decomposes into two ideal tetrahedra. Choose the two tetrahedra to be *regular*. That is, all dihedral angles are $\pi/3$. We claim that this gives a hyperbolic structure on the figure-8 knot complement.

We wish to find all such structures.

Thurston worked through this example in detail in his notes; we recall his work here [**Thurston, 1979**, pages 50–52].

Figure 4.10 shows the two tetrahedra in the decomposition of the figure-8 knot complement, which we obtained in chapter 1. These tetrahedra come from the two ideal polyhedra that glue to give the figure-8 knot complement that we discussed in detail in chapter 1; see figure 1.8 and figure 1.10. The tetrahedra differ from those in chapter 1 in the following ways. First, we have collapsed the bigons. This gives two remaining edge classes, which we label with one tick mark and with two tick marks. Second, in figure 1.8, we viewed the top ideal polyhedron from the inside; that is, the ball of the polyhedron lay above the plane of projection. To be more consistent in viewing both top and bottom polyhedron, we have rotated our perspective such that now both tetrahedra are viewed from the outside.

For each tetrahedron, we label each edge with a complex number z_i or w_i, to denote the edge invariant associated with that edge. Note that

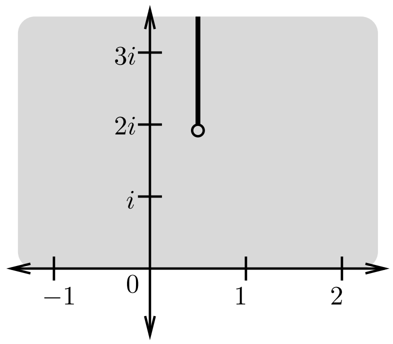

FIGURE 4.11. Solutions to edge gluing equations for the figure-8 knot complement are parametrized by the above region.

opposite edges in a tetrahedron have the same edge invariant. We also have relationships between z_1, z_2, and z_3 as in lemma 4.6, and similarly for w_1, w_2, and w_3.

There are two edge classes in the tetrahedra in figure 4.10, labeled with one or two tick marks on the edge. We obtain the edge gluing equations by taking the product of edge invariants for all edges identified with each edge class.

For the edge with one tick mark, we obtain the edge gluing equation

$$z_1^2\, z_3\, w_1^2\, w_3 = 1.$$

For the edge with two tick marks,

$$z_2^2\, z_3\, w_2^2\, w_3 = 1.$$

We set $z_1 = z$ and $w_1 = w$. From lemma 4.6, the first edge gluing equation gives

$$z^2 \left(\frac{z-1}{z}\right) w^2 \left(\frac{w-1}{w}\right) = 1,$$

or

$$z\,(z-1)\,w\,(w-1) = 1. \tag{4.1}$$

Solving for z in terms of w:

$$z = \frac{1 \pm \sqrt{1 + 4/(w(w-1))}}{2}.$$

We need the imaginary parts of z and w to be strictly greater than 0. For each value of w, there is at most one solution for z with positive imaginary part. The solution exists provided that the discriminant $1 + 4/(w(w-1))$ is not positive real. Thus solutions are parameterized by the region of $\mathbb{C}$ shown in figure 4.11 (see also exercise 4.6).

Notice that

$$z = w = \sqrt[3]{-1} = \frac{1}{2} + \frac{\sqrt{3}}{2}i$$

is one solution to the equations. We will see that this gives a complete hyperbolic structure on the complement of the figure-8 knot.

4.3. Completeness equations

Suppose now that M is a 3-manifold with torus boundary. In much of this section, we will assume that M admits a topological ideal triangulation, and moreover we have a solution to the edge gluing equations for this triangulation, thus M admits a hyperbolic structure. We need to consider cusps of the manifold to determine whether this is a complete structure or not.

Definition 4.9. Let M be a 3-manifold with torus boundary. Define a *cusp*, or *cusp neighborhood* of M to be a neighborhood of ∂M homeomorphic to the product of a torus and an interval, $T^2 \times I$. Define a *cusp torus* to be a torus component of ∂M, or the boundary of a cusp.

A hyperbolic structure on M induces an affine structure on the boundary of any cusp of M.

Theorem 4.10. *Let M be a 3-manifold with torus boundary and hyperbolic structure, i.e. with $(\mathrm{Isom}(\mathbb{H}^3), \mathbb{H}^3)$-structure. Then the structure on M is complete if and only if for each cusp of M, the induced structure on the boundary of the cusp is a Euclidean structure on the torus.*

Proof. Exercise 4.7. Hint: the proof is very similar to that of the analogous result in two dimensions, proposition 3.15. $\square$

Definition 4.11. Let M have a topological ideal triangulation. If we truncate the vertices of each ideal tetrahedron, we obtain a collection of triangles, each of which lies on the boundary of a cusp. Edges of each triangle inherit a gluing from the gluing of faces of the ideal tetrahedra. This gives a triangulation of each boundary torus, which we call a *cusp triangulation*.

An example for the figure-8 knot is shown in figure 4.12. The truncated ideal vertices give eight triangles, with labels a through h. These glue together on the boundary of the cusp to give a triangulation of the torus as shown. Note that the corner of each triangle is labeled with the edge invariant of the tetrahedron corresponding to the edge meeting that corner.

Figure 4.12 shows a fundamental region of the cusp triangulation. By tracing through gluings of cusp triangles, we may obtain the full developing image of the cusp torus. Theorem 4.10 states that the original manifold is complete if and only if the cusp tori are Euclidean, which will hold if and only if the holonomy maps for each element of $\pi_1(T)$ on each cusp torus T are pure Euclidean translations, without rotation or scale.

We can determine if holonomy maps are Euclidean translations directly from the cusp triangulation. Start with a triangle Δ whose vertices we may assume lie at 0, 1, and $z(e_1)$ in the complex plane $\mathbb{C}$. Let $\alpha \in \pi_1(T)$. Then the holonomy $\rho(\alpha)$ takes Δ to a new triangle, which appears in the

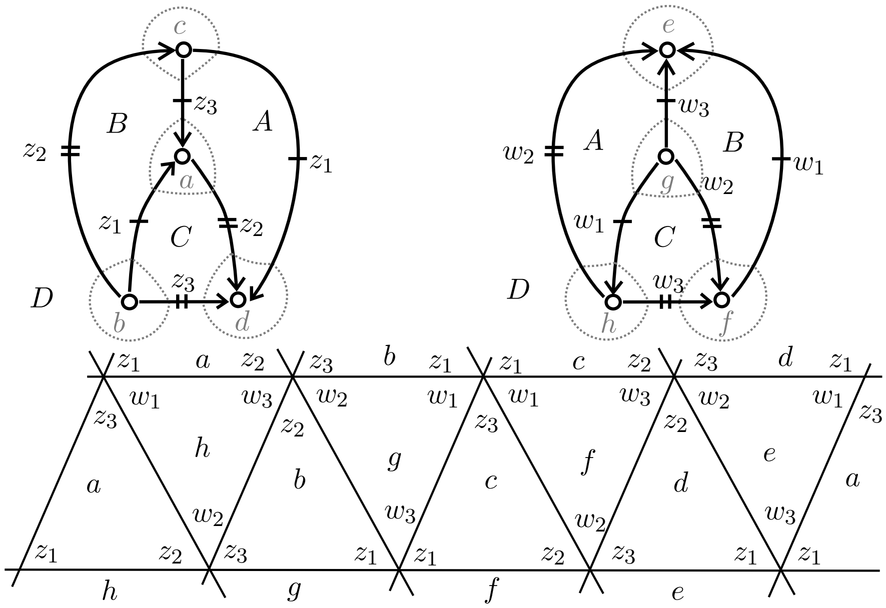

FIGURE 4.12. Finding the cusp triangulation of the figure-8 knot complement

developing image. The holonomy $\rho(\alpha)$ will be a Euclidean translation if and only if the triangle side from 0 to 1 of Δ is mapped to the side of a triangle of length 1 pointing in the same direction, without rotation (or scale). To determine whether this holds, we may follow the side of the triangle in the developing image, and obtain exactly its rotation and scale by considering the edge invariants that adjust its length and direction as it is adjusted in the cusp triangulation, as in figure 4.8. This can be described efficiently in the following way.

DEFINITION 4.12. Suppose M has a topological ideal triangulation, and let T be the boundary torus of a cusp of M. Let $[\alpha] \in \pi_1(T)$, so α is a loop on T in the homotopy class of $[\alpha]$. We associate a complex number $H(\alpha)$ to α as follows.

First, orient the loop α on T. The loop α can be homotoped to run through any triangle of the cusp triangulation of T monotonically, i.e. in such a way that it cuts off a single corner of each triangle it enters. Denote the edge invariants of the corners cut off by α by $z_1, z_2, \ldots, z_n$. Further associate to each corner a value $\epsilon_i = \pm 1$: if the i-th corner cut off by α lies to the left of α, set $\epsilon_i = +1$. If the corner lies to the right of α, set $\epsilon_i = -1$. Finally, set the value of $H(\alpha)$ to be

$$H([\alpha]) = \prod_{i=1}^{n} z_i^{\epsilon_i} \tag{4.2}$$

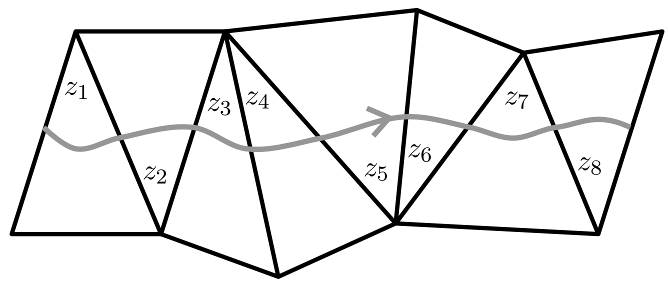

FIGURE 4.13. Example for determining $H([\alpha])$

EXAMPLE 4.13. An example cusp is shown in figure 4.13. For this example, the value of $H(\alpha)$ is given by

$$H(\alpha) = z_1\, z_2^{-1}\, z_3\, z_4\, z_5^{-1}\, z_6^{-1}\, z_7\, z_8^{-1}.$$

We will see that H is independent of homotopy class of α (exercise 4.11). For this reason, we sometimes denote the complex number by $H([\alpha])$, or evaluate it on a homotopy class rather than a curve.

EXAMPLE 4.14. For the figure-8 knot, there is a closed curve on the cusp torus running from the left side of the triangle labeled a on the left of figure 4.12 to the left side of the triangle labeled a on the right of that figure. Call this curve α. Then we can compute:

$$H(\alpha) = z_3\, w_2^{-1}\, z_2\, w_3^{-1}\, z_3\, w_2^{-1}\, z_2\, w_3^{-1} = \left(\frac{z_2\, z_3}{w_2\, w_3}\right)^2$$

Another closed curve runs from the base of the triangle labeled a on the left of figure 4.12 to the top of the triangle labeled h, also on the left of that figure. Call this curve β. Then we have:

$$H(\beta) = z_2^{-1}\, w_1 = \frac{w_1}{z_2}.$$

PROPOSITION 4.15 (Completeness equations). *Let T be the torus boundary of a cusp neighborhood of M, where M admits a topological ideal triangulation, and the ideal tetrahedra admit hyperbolic structures that satisfy the edge gluing equations (theorem 4.7). Let α and β generate $\pi_1(T)$. If $H(\alpha) = H(\beta) = 1$, then the ideal triangulation is a geometric ideal triangulation, i.e. the hyperbolic structure on M induced by the hyperbolic structure on the tetrahedra will be a complete structure.*

The equations $H(\alpha) = 1$ and $H(\beta) = 1$ are called the *completeness equations*.

PROOF OF PROPOSITION 4.15. By theorem 4.10, it suffices to show that the induced structure on T is Euclidean. To do so, it suffices to show that the holonomy elements $\rho(\alpha)$ and $\rho(\beta)$ are pure translations, with no rotation and scale. Thus we will show $\rho(\alpha)$ and $\rho(\beta)$ do not rotate or scale.

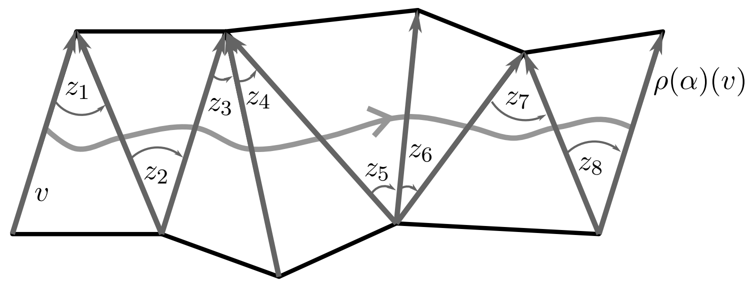

FIGURE 4.14. A path of vectors in the proof of proposition 4.15

To show this, let Δ be a triangle met by the curve α used in defining the complex number $H(\alpha)$, and suppose α meets a side e_1 of Δ. Let v be a vector with length equal to the length of e_1, pointing in the direction of e_1 such that the oriented curve α and the vector e_1 are oriented according to the right hand rule. This is true of the vector v shown on the far left of figure 4.14.

The holonomy $\rho(\alpha)$ is Euclidean if and only if the image of v under $\rho(\alpha)$ still has length v, and points in the same direction as v. We determine the effect of holonomy by considering what happens to v in each triangle of the cusp triangulation.

We may rotate v around a vertex of the triangle Δ meeting e_1, and scale, so that the result lines up with a second edge e_2 of the triangle, having the same length and direction as e_2. We know exactly how the rotation and scale is determined when the vertex of the triangle is labeled with edge invariant z_1: if we rotate in a counterclockwise direction, v is adjusted by multiplication by z_1, as in figure 4.8. If we rotate in a clockwise direction, v is adjusted by multiplication by $1/z_1$.

Now, our path α cuts off exactly one corner of each triangle it meets. This defines a path of edges of triangles, namely, starting with v, at each step we have a vector lying on the side of a triangle where α enters that triangle. In this triangle, rotate through the corner cut off by α to produce a new vector pointing in the direction of the side where α exits. An example path of such vectors is shown in figure 4.14. When α returns to the initial triangle Δ, the final vector of this path will be parallel to the image of Δ under $\rho(\alpha)$. Then $\rho(\alpha)$ will be a Euclidean transformation if and only if the final vector in the path has length and direction identical to that of v.

On the other hand, the final length and direction of the vector $\rho(\alpha)v$ is given by the product of edge invariants at the corners of each triangle in the path of edges, with edge invariant either multiplied or divided depending on whether the rotation is in the counterclockwise or clockwise direction, respectively. This is exactly the complex number $H(\alpha)$. Thus $\rho(\alpha)$ is Euclidean if and only if $H(\alpha) = 1$.

The same argument applies to $H(\beta)$ and $\rho(\beta)$. Since the holonomy group of the cusp is generated by $\rho(\alpha)$ and $\rho(\beta)$, the cusp will be Euclidean if and only if $H(\alpha) = H(\beta) = 1$. $\square$

EXAMPLE 4.16. Returning to the example of the figure-8 knot, in example 4.14, we found that completeness equations are given by

$$H(\alpha) = \left(\frac{z_2\, z_3}{w_2\, w_3}\right)^2 \quad \text{and} \quad H(\beta) = \frac{w_1}{z_2}.$$

Lemma 4.6 implies that these can be rewritten in terms of variables z and w alone, as

$$H(\alpha) = \left(\frac{1}{1-z} \cdot \frac{z-1}{z} \cdot \frac{1-w}{1} \cdot \frac{w}{w-1}\right)^2 = \left(\frac{w}{z}\right)^2$$

and

$$H(\beta) = w\,(1-z) \tag{4.3}$$

If the hyperbolic structure is complete, then by proposition 4.15, $H(\alpha) = H(\beta) = 1$, so $z = w$.

From equation (4.1), $(z(z-1))^2 = 1$. From equation (4.3), $z(z-1) = -1$. Hence the only possibility is $z = w = \frac{1}{2} + i\frac{\sqrt{3}}{2}$.

4.4. Computing hyperbolic structures

Given a triangulation of a 3-manifold M with torus boundary, we may determine a complete hyperbolic structure on M by solving the edge gluing and completeness equations. However, note this amounts to solving a complicated system of nonlinear equations. Consequently, it is difficult to use these to find exact hyperbolic structures on infinite families of manifolds.

However in practice, topological triangulations, edge gluing equations, and completeness equations can be found very efficiently by computer for specific, finite examples. Resulting nonlinear system of equations can then be solved numerically. The first software to find hyperbolic structures on knots and 3-manifolds was the program SnapPea, written by Weeks [**Weeks, 1985**] (see also [**Weeks, 2005**]). This program has allowed researchers to run experiments on large classes of hyperbolic 3-manifolds, making observations and testing conjectures, and has been influential in a great deal of results on hyperbolic structures on knots and 3-manifolds. The SnapPea kernel is now part of a program maintained by Culler, Dunfield, Goerner, and others, reincarnated as SnapPy, and available for free download [**Culler et al., 2016**]. This new program includes much additional functionality, and still remains an excellent tool for research in hyperbolic knot theory.

One issue in the past with finding a hyperbolic structure via SnapPea (SnapPy) is that it would only give a numerical approximation to a hyperbolic structure, and there was no guarantee that the manifold would be actually provably hyperbolic. This has been addressed in a few ways. The program Snap [**Coulson et al., 2000**] deduces exact solutions from the numerical approximations, which can be used to prove hyperbolicity. In another direction, Moser used analytic techniques to prove that a solution to edge gluing and completeness equations exists in a small neighborhood of an approximate solution [**Moser, 2009**]. In [**Hoffman et al., 2016**], interval arithmetic is used to prove hyperbolic structures exist when a structure is computed numerically. Thus using these tools, we can often prove that if SnapPy computes a hyperbolic structure on a knot complement, then the knot is indeed hyperbolic.

4.5. Exercises

Exercise 4.1. Write down the edge gluing equations (not completeness equations) for the 6_1 knot, using the ideal tetrahedra of example 4.1.1. Make appropriate substitutions such that your equations contain exactly one variable per tetrahedron.

Exercise 4.2. Notice that for both the figure-8 knot complement and for the 6_1 knot, we had exactly the same number of edges as tetrahedra in the ideal triangulation.

(a) Prove that this will always be true. That is, prove that if M is any 3-manifold with (possibly empty) boundary consisting of tori, then for any topological ideal triangulation of M, the number of edges of the triangulation will always equal the number of tetrahedra.

(b) Since we have one unknown per ideal tetrahedra, part (a) implies that the number of gluing equations will equal the number of unknowns. However, in fact the gluing equations are always redundant. Prove this fact.

Exercise 4.3. In chapter 1, we found a polyhedral decomposition of the 5_2 knot complement (without bigons). Split this into a topological ideal triangulation of the knot complement.

Exercise 4.4. Using the ideal tetrahedra of exercise 4.3, or otherwise, write down all edge invariants and all edge gluing equations, one variable per tetrahedron.

Exercise 4.5. Find a topological ideal triangulation of the 6_3 knot, edge invariants, and edge gluing equations.

Exercise 4.6. Check that figure 4.11 does indeed parametrize the space of hyperbolic structures on the figure-8 knot complement. What is the equation of the vertical ray shown in that picture?

Exercise 4.7. Prove theorem 4.10: the hyperbolic structure on M is complete if and only if for each cusp of M, the induced structure on the boundary of the cusp is a Euclidean structure on the torus.

Exercise 4.8. For the topological triangulation of the 5_2 knot of exercise 4.3:

(a) Find the triangulation of the cusp. Label a fundamental domain, and meridian and longitude.
(b) Write down completeness equations.

EXERCISE 4.9. Find the cusp triangulation for the complement of the 6_1 knot from example 4.1.1.

EXERCISE 4.10. Find completeness equations for the 6_1 or 6_3 knot.

EXERCISE 4.11. Suppose M admits an ideal triangulation that satisfies the edge gluing equations.

(a) In definition 4.12, we claimed that for any closed curve α in a torus boundary component of ∂M, we could homotope α in such a way that it cuts off a single corner of each triangle that it meets. Prove this.
(b) Show that $H([\alpha])$ is independent of the choice of α in the homotopy class of $[\alpha]$. In particular, if α is homotoped to run through different triangles, the value of $H([\alpha])$ is unchanged.

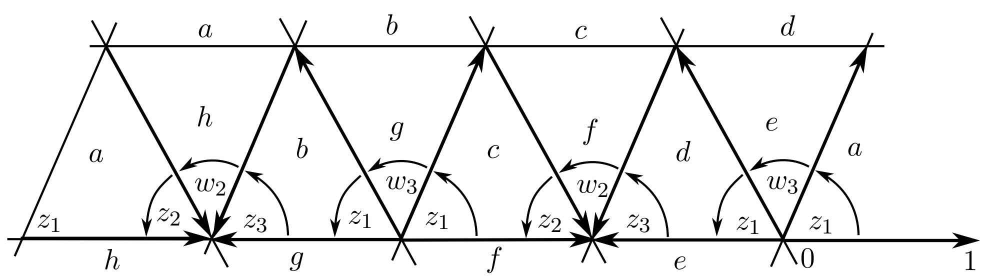

FIGURE 4.15. Path of vectors going from e_1, the oriented edge from 0 to 1, to $\rho(-\alpha)(e_1)$

EXERCISE 4.12. In Thurston's 1979 notes [**Thurston, 1979**], he computed completeness equations for the figure-8 knot using a method similar to our proof of proposition 4.15. Namely, he found a path of vectors from an edge on a triangle Δ to the same edge on $\rho(-\alpha)(\Delta)$ and $\rho(\beta)(\Delta)$. His path of vectors for $\rho(\beta)$ agrees with ours. His path of vectors for $\rho(-\alpha)$ is different from our path for $\rho(\alpha)$, and is shown in figure 4.15.

(a) Prove that the completeness equation obtained from Thurston's path of vectors is equivalent to our completeness equation.
(b) More generally, prove that if we replace our path of vectors used to construct the complex number $H([\alpha])$ by any other path of vectors obtained by rotating around vertices of the cusp triangulation, with same starting and ending vectors, then the equation we obtain from multiplying (and dividing) by edge invariants corresponding to the path of vectors gives a completeness equation that is equivalent to $H([\alpha]) = 1$.

Exercise 4.13. What breaks down when you try to find triangulations and edge gluing equations for non-hyperbolic knots and links, such as the trefoil or the $(2,4)$-torus link?

Exercise 4.14. Use the computer program SnapPy to determine which of the knots with seven or fewer crossings admit a hyperbolic structure [**Culler et al., 2016**]. For those that do admit a hyperbolic structure, use SnapPy to find the cusp triangulation of the knot. Obtain a screen shot of this information, which should include cusp triangles as well as a fundamental parallelogram for the cusp.