6. Constructing hyperbolic manifolds

Recall that a hyperbolic manifold is a (G, X)-manifold, where X is $\mathbb{H}^n$ and $G = \mathrm{Isom}(\mathbb{H}^n)$. Since G is a group of Riemannian isometries, G acts rigidly. Hence, the pseudogroup generated by G is composed of diffeomorphisms between open subsets of $\mathbb{H}^n$ such that the restriction to any component is the restriction of an isometry of $\mathbb{H}^n$. This is so important that we include it as a proposition.

Proposition 6.1. *Let M be a hyperbolic manifold. Then M is a topological manifold with a cover by open sets U_i, together with open maps $\phi_i \colon U_i \to \mathbb{H}^n$ (known as charts) which are homeomorphisms onto their images, such that if $U_i \cap U_j \neq \emptyset$, then for each component X of $U_i \cap U_j$,*

$$\phi_j \circ \phi_i^{-1} \colon \phi_i(X) \to \phi_j(X)$$

is the restriction of a hyperbolic isometry.

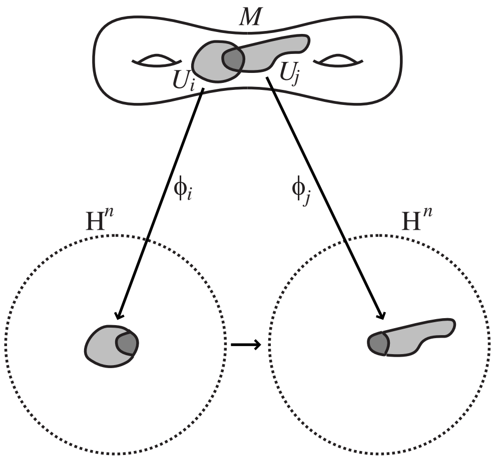

Figure 10.

Since G is a group of Riemannian isometries, a hyperbolic manifold inherits a Riemannian metric. The following characterises hyperbolic manifolds among all Riemannian manifolds.

Proposition 6.2. *Hyperbolic n-manifolds are precisely those Riemannian manifolds, each point of which has a neighbourhood isometric to an open subset of $\mathbb{H}^n$.*

Proof. In one direction this is straightforward. A hyperbolic manifold M has a chart into $\mathbb{H}^n$ which, by the way that the Riemannian metric on M is constructed, is an isometry to open subset of $\mathbb{H}^n$.

In the other direction, suppose that M is a Riemannian manifold, such that each point of M has a neighbourhood isometric to an open subset of $\mathbb{H}^n$. We take each such isometry to be a chart for M. We need to check that the transition maps lie in the pseudogroup generated by $\mathrm{Isom}(\mathbb{H}^n)$. Suppose that $\phi_i \colon U_i \to \mathbb{H}^n$ and $\phi_j \colon U_j \to \mathbb{H}^n$ are charts that overlap. Let X be some component of $\phi_j(U_i \cap U_j)$. Then $\phi_i \circ \phi_j^{-1}|_X$ is a isometry between open subsets of $\mathbb{H}^n$. Let x be some point of X. Then, by Theorem 2.1, there is some isometry h of $\mathbb{H}^n$ such that $h(x) = \phi_i \circ \phi_j^{-1}(x)$ and $(Dh)_x = (D\phi_i \circ D\phi_j^{-1})_x$. By Theorem 1.5 of the Introduction to Riemannian Manifolds, $\phi_i \circ \phi_j^{-1}|_X$ is the restriction of h. $\square$

Hyperbolic manifolds are obtained by gluing bits of hyperbolic space together via hyperbolic isometries. It is best to use 'nice' bits of hyperbolic space, for example, polyhedra.

Definition. A *k-dimensional hyperplane* in $\mathbb{H}^n$ is the image of $D^k \subset D^n$ after an isometry $D^n \to D^n$. A *half-space* is the closure in D^n of one component of the complement of a codimension one hyperplane.

Definition. A *polyhedron* in $\mathbb{H}^n$ is a compact subset of $\mathbb{H}^n$ that is the intersection of a finite collection of half-spaces. The *dimension* of a polyhedron is the smallest dimension of a hyperplane containing it. We will usually restrict attention to *non-degenerate* polyhedra in $\mathbb{H}^n$ which are those with dimension n. A *face* of a polyhedron P is the intersection $P \cap \pi$, where π is a codimension one hyperplane in $\mathbb{H}^n$ such that P is disjoint from one component of $\mathbb{H}^n - \pi$. Note that a face is a degenerate polyhedron. A *facet* is a face with codimension one. The *vertices* of P are the dimension zero faces. An *ideal polyhedron* is the intersection of a finite number of half-spaces in $\mathbb{H}^n$ whose closure in $\mathbb{H}^n \cap S^{n-1}_\infty$ intersects S^{n-1}_∞ in a finite number of points, and which has no vertices in $\mathbb{H}^n$.

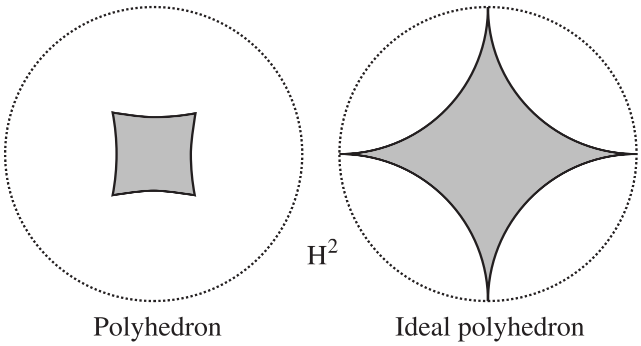

Figure 11.

Suppose now that M is obtained by gluing a collection of n-dimensional (possibly ideal) hyperbolic polyhedra $P_1, \ldots, P_m$ by identifying their facets in pairs via isometries between the facets. Let P be the disjoint union $P_1 \cup \ldots \cup P_m$, and let $q\colon P \to M$ be the quotient map. Note that $q|_{P-\partial P}$ is a homeomorphism and so $q(P - \partial P)$ inherits a hyperbolic structure. *Sometimes* M will be a hyperbolic manifold. But *sometimes* M may fail even to be a manifold. The following is a criterion which ensures that the hyperbolic structure on $q(P - \partial P)$ extends over all of M.

Theorem 6.3. *Suppose that each point $x \in M$ has a neighbourhood U_x and an open mapping $\phi_x\colon U_x \to B_{\epsilon(x)}(0) \subset D^n$ which is a homeomorphism onto its image, which sends x to 0 and which restricts to an isometry on each component of $U_x \cap q(P - \partial P)$. Then M inherits a hyperbolic structure.*

Proof. By reducing $\epsilon(x)$ if necessary, we can ensure that the closure of each component of $U_x \cap q(P - \partial P)$ contains x. Then ϕ_x will be the charts for M. These maps determine, for each $x_i \in q^{-1}(x)$, an isometry $h_{x_i}\colon B_{\epsilon(x)}(x_i) \to B_{\epsilon(x)}(0) \subset D^n$ such that $h_{x_i}|_{P-\partial P} = \phi_x \circ q$. We must check that if X is a component of $U_x \cap U_y$, then $\phi_y \phi_x^{-1}\colon \phi_x(X) \to \phi_y(X)$ is the restriction of a hyperbolic isometry. By assumption, this is true for each component of $\phi_x(X \cap q(P - \partial P))$. We must ensure that these isometries agree over all of $\phi_x(X)$. Any two points of $\phi_x(X \cap q(P - \partial P))$ are joined by a path in $\phi_x(X)$ which avoids (the image under $\phi_x \circ q$ of) the faces with dimension less than $n - 1$. Hence, we need only check that if z lies in $\phi_x(q(\partial P) \cap X)$ but not in a face of dimension less than $n - 1$, then then $\phi_y \phi_x^{-1}$ is an isometry in a neighbourhood of z. Let z_1 and z_2 be $q^{-1}\phi_x^{-1}(z)$. The component of $q^{-1}(U_x)$ (respectively, $q^{-1}(U_y)$) containing z_i contains a single point x_i of $q^{-1}(x)$

(respectively, y_i of $q^{-1}(y)$). Let F_i be the facet containing z_i and let $k\colon F_1 \to F_2$ be the identification isometry between the facets. Note that x_i and y_i also lie in F_i. There are two possible ways of extending k to an isometry of $\mathbb{H}^n$. Pick the one so that the following commutes:

$$\begin{tikzcd}
\mathbb{H}^n \arrow[rr, "k"] \arrow[dr, "h_{x_1}"'] & & \mathbb{H}^n \arrow[dl, "h_{x_2}"] \\
& D^n &
\end{tikzcd}$$

Hence, running down the left-hand side of the following diagram is the same as running down the right-hand side (where the maps are defined):

$$\begin{tikzcd}
& D^n \arrow[dl, "h_{x_1}^{-1}"'] \arrow[dr, "h_{x_2}^{-1}"] & \\
\mathbb{H}^n \arrow[rr, "k"] \arrow[dr, "h_{y_1}"'] & & \mathbb{H}^n \arrow[dl, "h_{y_2}"] \\
& D^n &
\end{tikzcd}$$

This ensures that $\phi_y \phi_x^{-1}$ is a well-defined isometry in a neighbourhood of z.

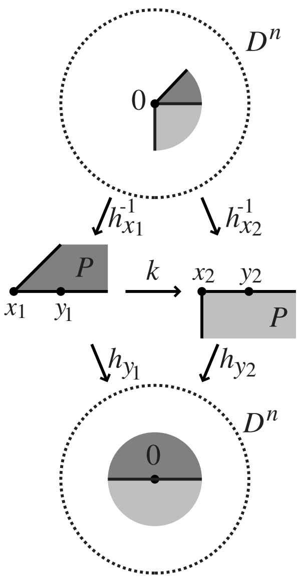

Figure 12.

There is one further thing to check: that M actually is a topological manifold. We must check that M is Hausdorff (which is clear) and that M has a countable basis of open sets. This is straightforward: we can refine $\{U_x : x \in M\}$ to a countable cover. Then, each U_x is homeomorphic to an open subset of $\mathbb{R}^n$, which has a countable basis of open sets. □

A similar theorem (with the same proof) also works for spherical and Euclidean structures.

This theorem reduces the problem of finding a hyperbolic structure to a problem one dimension lower. Define the *link* of a point x in a polyhedron P to be

$$\{v \in T_x\mathbb{H}^n : ||v|| = 1 \text{ and } \exp_x(\lambda v) \in P \text{ for some } \lambda > 0\}.$$

This is a polyhedron in the unit sphere in $T_x\mathbb{H}^n$. Isometries between facets of hyperbolic polyhedra induce isometries between facets of the links of identified points. The existence of a hyperbolic structure on the quotient space M is equivalent to each point of M having a link that is isometric to the $(n-1)$-sphere.

Here is a sample application of this theorem. Recall that every closed orientable 2-manifold is homeomorphic to one (and only one) of the surfaces F_0, F_1, ..., where $F_0 = S^2$ and F_k is obtained from F_{k-1} by removing the interior of an embedded 2-disc and then attaching a handle as follows:

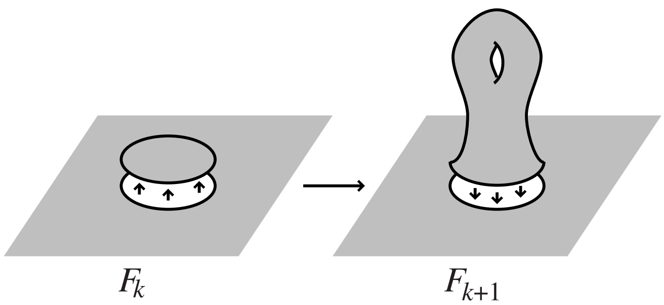

Figure 13.

Theorem 6.4. *F_k admits a hyperbolic structure for all $k \geq 2$.*

We shall see later that neither the 2-sphere nor the torus admits a hyperbolic structure.

Proof. For $k \geq 1$, F_k is obtained from a polygon with $4k$ sides by identifying their sides as follows:

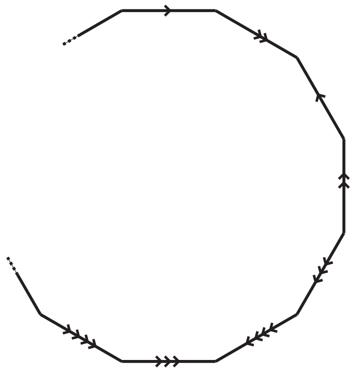

Figure 14.

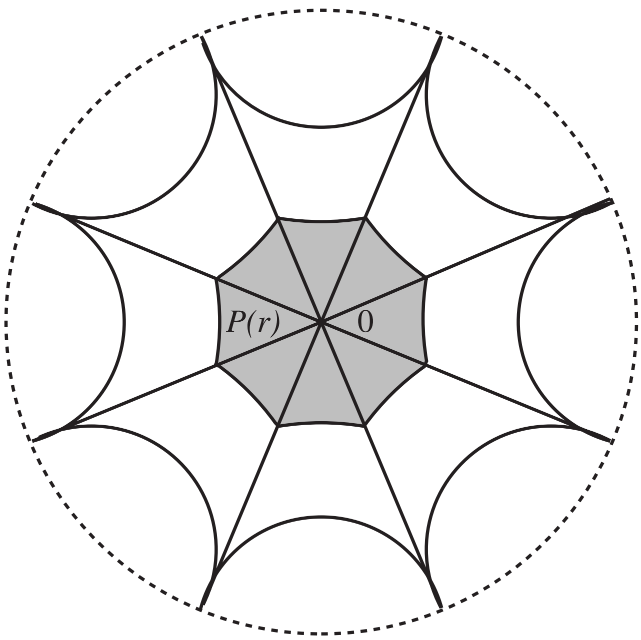

Figure 15.

Realise this as a polyhedron in $\mathbb{H}^2$, as follows. Draw $4k$ geodesics emanating from the origin 0 in D^2, with the angle between adjacent geodesics being $2\pi/4k$. Place a vertex on each geodesic, each a hyperbolic distance r from 0. Let $P(r)$ be the polyhedron having these $4k$ points as vertices. Let $\beta(r)$ be the interior angle at each vertex.

Claim. For some r, $\beta(r) = 2\pi/4k$.

Proof of theorem from claim. Glue the sides of $P(r)$ together to form F_k. Let

$q\colon P(r) \to F_k$ be the quotient map. We check that the requirements of Theorem 6.3 are satisfied, and hence that this determines a hyperbolic structure on F_k. If x lies in $q(P(r) - \partial P(r))$ then it automatically has a neighbourhood U_x as required.

If x lies in $q(\partial P(r))$ but is not in the image of a vertex of $P(r)$, then $q^{-1}(x)$ is two points. For each point y of $q^{-1}(x)$, $B_{\epsilon(x)}(y) \cap P(r)$ is isometric to half an $\epsilon(x)$-ball in D^n, providing $\epsilon(x)$ is sufficiently small. We may map these two half balls to a whole $\epsilon(x)$-ball in D^n as required.

There is a single point x lying in the image under q of the vertices. Since $\beta(r) = 2\pi/4k$, we may map $\epsilon(x)$-neighbourhoods of the $4k$ vertices homeomorphically to an $\epsilon(x)$-ball in D^n, in such a way that the homeomorphism restricts to an isometry on $q(P(r) - \partial P(r))$.

Proof of claim. Note that β is a continuous function of r. We examine the behaviour of $\beta(r)$ as $r \to 0$ and $r \to \infty$.

Perform a Euclidean scale h based at 0 which takes each vertex to a point on S^1_∞. This map preserves angles, but is *not* a hyperbolic isometry. As $r \to 0$, the sides of $h(P(r))$ approximate Euclidean straight lines. Hence, $\beta(r)$ tends to the interior angle of a regular Euclidean $4k$-gon, which is $\pi(1 - 1/2k)$.

As $r \to \infty$, the sides of $h(P(r))$ approximate hyperbolic geodesics. Hence, $\beta(r)$ tends to the angle between two geodesics emanating from the same point in S^1_∞. Since both geodesics are at right angles to S^1_∞, the angle between them is zero. Hence, $\beta(r) \to 0$ as $r \to \infty$.

Since $k > 1$, we have $0 < 2\pi/4k < \pi(1 - 1/2k)$. So, there is a value of r for which $\beta(r) = 2\pi/4k$. $\square$

7. The figure-eight knot complement

The following is a diagram of the figure-eight knot L. We will construct a hyperbolic structure on $S^3 - L$.

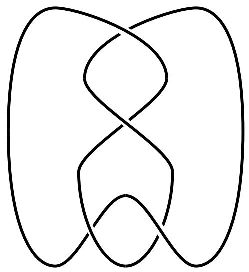

Figure 16.

Consider the following diagram of two regular tetrahedra in $\mathbb{R}^3$. There is a unique way to glue the faces in pairs (via Euclidean isometries) so that the edges (and their orientations) match. Let M be the space formed by this gluing.

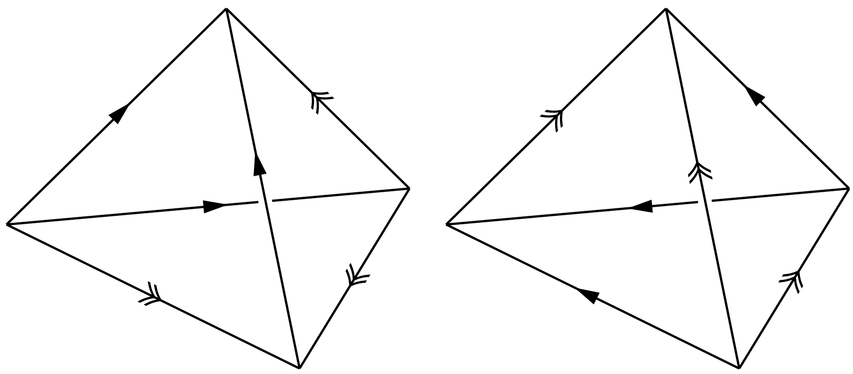

Figure 17.

This gives M a cell complex. However, M is not a manifold. The vertices of the tetrahedra are all identified to a single vertex v. A small neighbourhood of v is a cone on a torus.

However, every other point has a neighbourhood homeomorphic to an open ball in $\mathbb{R}^3$. So $M - v$ is a 3-manifold.

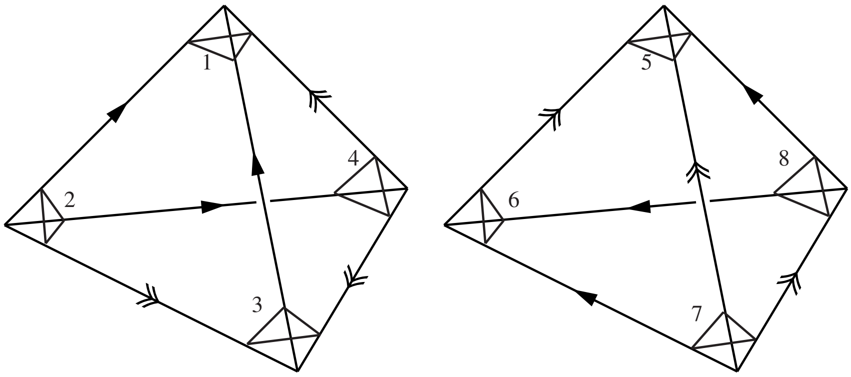

The boundary of an ε neighbourhood of v

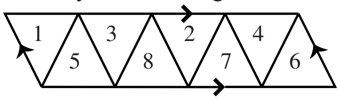

Figure 18.

Theorem 7.1. *$M - v$ is homeomorphic to $S^3 - L$, where L is the figure-eight knot.*

Proof. First consider the following cell complex K^1 embedded in S^3. The 1-cells are labelled 1 to 6 and are assigned an orientation.

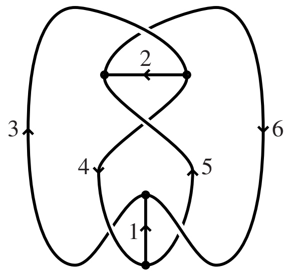

Figure 19.

Now attach to these four 2-cells, to form a cell complex K^2.

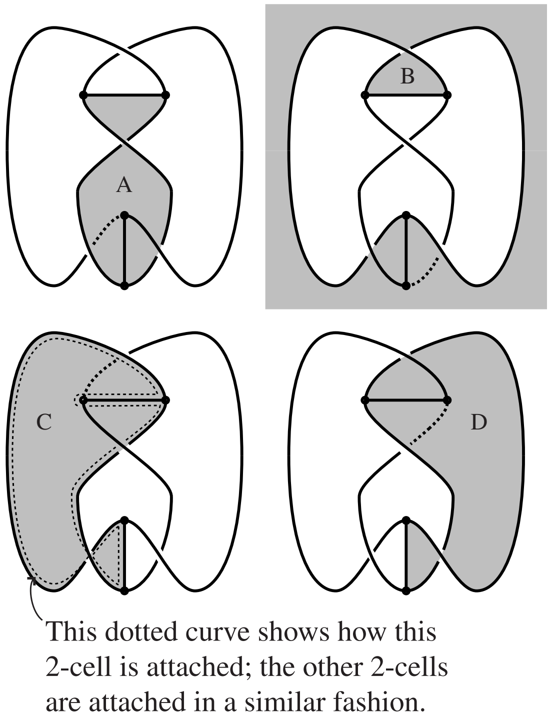

Figure 20.

Claim 1. $S^3 - K^2$ is homeomorphic to two 3-balls.

Before we prove this claim, the following will be useful.

Claim 2. There is a homeomorphism h between $S^3 - K^1$ and the complement of the following cell complex K_1^1.

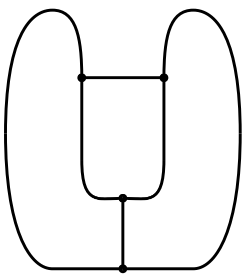

Figure 21.

The homeomorphism is supported in a neighbourhood of 1-cells 1 and 2. We focus on cell 1 (the other case being similar).

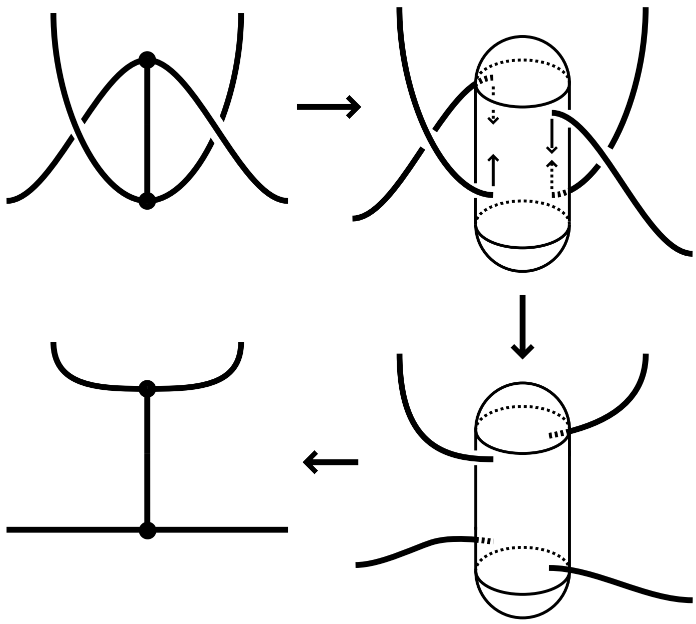

Figure 22.

In the first step above, we thicken cell 1 to a small closed 3-ball. In the second step, we slide the endpoints of the 1-cells attached to this ball. In the final step, we shrink the ball again to a single cell. This proves Claim 2.

The homeomorphism h takes the 2-cells of K^2 to the following 2-cells, which each lies in the plane of the diagram.

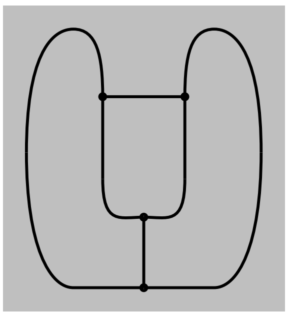

Figure 23.

Hence, $S^3 - K^2$ is homeomorphic to the complement of the above complex, which is two open 3-balls. This proves Claim 1.

Viewing these open 3-balls as the interiors of two 3-cells, we see that K^2

extends to a cell complex K^3 for S^3. The boundaries of the two 3-cells are attached onto K^2 according to the following diagram. Each 3-cell is a ball B glued onto K^2 via a map $\partial B \to K^2$. The following specifies this map for each 3-cell.

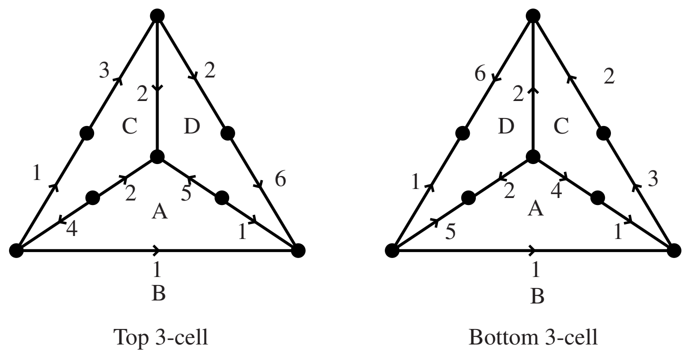

Figure 24.

The 0-cells and the 1-cells 3, 4, 5 and 6 combine to form the figure-eight knot L. If we collapse these cells to a point v, and then remove v, the result is the same as simply removing L from S^3. We therefore need to show that the cell complex we obtain by collapsing L to a point is M. It has a single 0-cell v, two 1-cells (coming from 1-cells 1 and 2), four 2-cells (coming from A, B, C and D), and two 3-cells. The attaching maps of the 2-cells and the 3-cells can be deduced from Figure 24, and are readily seen to give the required cell complex for M. Hence, $M - v$ is homeomorphic to $S^3 - L$. $\square$

Definition. An *ideal n-simplex* is the ideal polyhedron determined by $n+1$ points on S^{n-1}_∞. An ideal 3-simplex is also known as an *ideal tetrahedron*.

Remark. If an ideal polyhedron is determined by some points V on S^{n-1}_∞, then we will often abuse terminology by calling V its *vertices*.

Definition. An ideal n-simplex is *regular* if, for any permutation of its vertices, there is a hyperbolic isometry which realises this permutation.

Construct a regular ideal tetrahedron as follows. Let V be the points in S^2_∞ which are the vertices of a regular Euclidean tetrahedron centred at the origin of D^3. Let Δ be the hyperbolic ideal tetrahedron determined by V. Then Δ is regular because any permutation of the points of V is realised by an orthogonal map of $\mathbb{R}^3$ which is hyperbolic isometry.

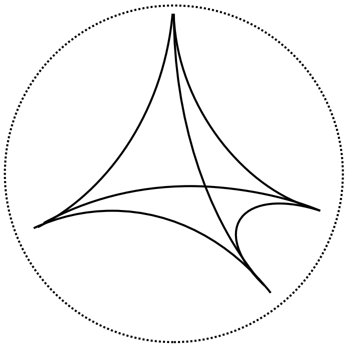

Figure 25.

Now glue two copies of Δ via isometries as specified by Figure 17. We will check that the conditions of Theorem 6.3 are satisfied and hence that this gives a hyperbolic structure on $S^3 - L$. As in Theorem 6.4, the only thing we have to check is that, for each point x lying in a 1-cell, the interior angles around x add up to 2π.

Note. The angle between two intersecting codimension one hyperplanes H_1 and H_2 in $\mathbb{H}^n$ is the same for all points of $H_1 \cap H_2$. This is because we may perform an isometry of D^n after which H_1 and H_2 both pass through 0. Then H_1 and H_2 become Euclidean hyperplanes, for which the assertion is clear.

Lemma 7.2. *Let F_1, F_2 and F_3 be three facets of an ideal tetrahedron in $\mathbb{H}^3$. Let β_{12}, β_{23} and β_{31} be the interior angles between these facets. Then $\beta_{12}+\beta_{23}+\beta_{31} = \pi$.*

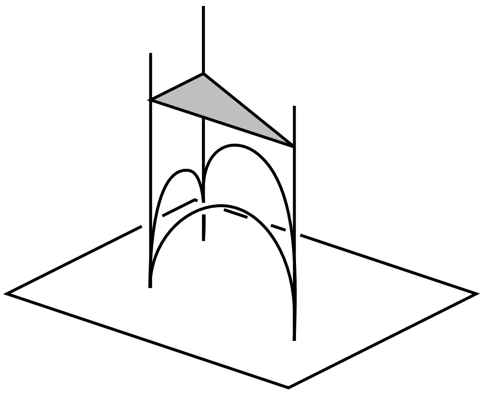

Figure 26.

Proof. Work with U^3. We may assume that the hyperplanes containing F_1, F_2 and F_3 are vertical Euclidean planes. Hence, β_{12}, β_{23} and β_{31} are the interior angles of a Euclidean triangle. □

Corollary 7.3. *If the ideal tetrahedron is regular, then* $\beta_{12} = \beta_{23} = \beta_{31} = \pi/3$.

For any point x of $M - v$ lying in a 1-cell, the 2-cells of M run past x six times. The six interior angles around x sum to 2π. Hence, we may construct the chart around x required by Theorem 6.3. Hence, $S^3 - L$ inherits a hyperbolic structure.

Remark. We have imposed a hyperbolic structure on the topological manifold $S^3 - L$. We have not shown that this is compatible with the smooth structure that $S^3 - L$ inherits from S^3 (although this is in fact true).

8. Gluing ideal tetrahedra

The construction of the hyperbolic structure on the figure-eight knot in the last section seems rather difficult to generalise. However, it is possible to construct many hyperbolic manifolds in this way, using ideal triangulations.

Definition. An *ideal triangulation* of a 3-manifold M is a way of constructing $M - \partial M$ from a collection of (topological) ideal tetrahedra by gluing their faces in pairs.

The following theorem (which we include without proof) demonstrates that there is no topological restriction to the existence of an ideal triangulation.

Theorem 8.1. *Any compact 3-manifold with non-empty boundary has an ideal triangulation.*

The aim of this section is to investigate when an ideal triangulation of a manifold can be used to impose a hyperbolic structure.

If the faces of an ideal tetrahedron in $\mathbb{H}^3$ are labelled 1 to 4, and β_{ij} is the angle between the faces i and j, then Lemma 7.2 gives the following equalities:

$$\beta_{23} + \beta_{34} + \beta_{42} = \pi$$
$$\beta_{34} + \beta_{41} + \beta_{13} = \pi$$
$$\beta_{41} + \beta_{12} + \beta_{24} = \pi$$
$$\beta_{12} + \beta_{23} + \beta_{31} = \pi$$

Adding the first two equations and subtracting the second two gives that $\beta_{12} = \beta_{34}$. Similarly, $\beta_{13} = \beta_{24}$ and $\beta_{14} = \beta_{23}$. Thus opposite edges have the same interior angle. An ideal tetrahedron in $\mathbb{H}^3$ therefore determines three interior angles α, β and γ adding up to π. All three angles appear at each vertex, and they cycle round the vertex the same way. Hence, if one considers such tetrahedra up to orientation-preserving isometry, the triple (α, β, γ) is well-defined up to cyclic permutation.

Lemma 8.2. *Ideal tetrahedra in* $\mathbb{H}^3$ *(up to orientation preserving isometry) are in one-one correspondence with triples of positive numbers adding to* π *(up to cyclic permutation).*

Proof. We have already seen how an ideal tetrahedron determines a triple. Any such triple can be realized by some ideal tetrahedron: find a Euclidean triangle in ∂U^3 with these interior angles and then consider the ideal tetrahedron with vertices being ∞ and the three corners of the triangle. Also, if two ideal triangles have the same interior angles, we can find an orientation-preserving isometry taking one to the other. Perform isometries taking one vertex of each tetrahedron to ∞. Then perform parabolic and loxodromic isometries that match up the remaining three vertices of each tetrahedron. □

There is an alternative way of describing these ideal tetrahedra. Given a triple (α, β, γ), consider a Euclidean triangle having these interior angles, and having vertices at 0, 1 and some point z in $\mathbb{C}$, where $\mathrm{Im}(z) > 0$. There is some ambiguity here since α, β or γ may be placed at the origin. Therefore, the following complex numbers all represent the same triple:

$$z, \ \frac{1}{1-z}, \ 1 - \frac{1}{z}.$$

Suppose now that M has an ideal triangulation, and that each ideal tetrahedron has been assigned interior angles α, β and γ. Each ideal tetrahedron then inherits a hyperbolic structure. The topological identification of facets of the tetrahedra is, according to the following lemma, realized by a unique isometry.

Lemma 8.3. *Let Δ and Δ' be ideal triangles in $\mathbb{H}^2$ with edges (e_1, e_2, e_3) and (e'_1, e'_2, e'_3). Then there is a unique isometry of $\mathbb{H}^2$ taking Δ to Δ' and the edges e_i to edges e'_i.*

Proof. Send the vertex joining e_1 and e_2 to ∞ in the upper half space model. Also send the vertex joining e'_1 and e'_2 to ∞. Then there is a unique hyperbolic isometry taking the remaining pairs of vertices to each other. $\square$

The hyperbolic structures on the tetrahedra then patch together to form a hyperbolic structure on the complement of the edges of M. When does this extend to a hyperbolic structure on M? The answer can be given in precise algebraic terms. At the i^{th} edge of M, let $w_{i1}, \dots, w_{ik}$ be the complex parameters of the tetrahedra around that edge. The choice of whether w_{ij} equals z, $1/(1-z)$ or $1 - 1/z$ is made so that the interior angle of the Euclidean triangle at the origin is the same as the interior angle at the edge i.

Theorem 8.4. *Let M be as above, with the extra condition that ∂M is a collection of tori. Then $M - \partial M$ inherits a hyperbolic structure if and only if for each edge i*

$$\prod_{j=1}^{k} w_{ij} = 1.$$

Proof. $M - \partial M$ is obtained from a collection of disjoint ideal tetrahedra in $\mathbb{H}^3$ by identifying pairs of faces. Let $e_1, \dots, e_k$ be the edges of these tetrahedra that are all identified to form the i^{th} edge e of M. We let $e_{k+1} = e_1$. The isometries between faces yield isometries $e_i \to e_{i+1}$. Hence we obtain an isometry

$$e_1 \to e_2 \to \dots \to e_1.$$

Case 1. $e_1 \to e_1$ is a non-zero translation along e_1.

Then each point on e has an infinite number of inverse-images in $e_1 \cup \dots \cup e_k$. So the quotient space is not a manifold.

Case 2. $e_1 \to e_1$ is a reflection.

This reflection has a fixed point. A neighbourhood of this point in the quotient space is a cone on $\mathbb{R}P^2$. Hence, the quotient space is not a manifold.

Case 3. $e_1 \to e_1$ is the identity.

In this case, the quotient space is indeed a manifold. We claim that this case holds if and only if

$$|\prod_{j=1}^{k} w_{ij}| = 1.$$

Place the tetrahedron containing e_1 in U^3 so that its vertices are ∞, 0, 1 and w_{i1}, and so that e_1 runs from 0 to ∞. Place the tetrahedron containing e_2 beside it, so that their faces are glued via the correct isometry. The vertices of this tetrahedra are ∞, 0, w_{i1} and $w_{i1}w_{i2}$. Continue this procedure for all k tetrahedra. In this way, we have ensured that all the gluing maps $e_i \to e_{i+1}$ are the identity for $1 \leq i < k$. The final gluing map $e_k \to e_1$ sends the points ∞, 0, $\prod_{j=1}^{k} w_{ij}$ to ∞, 0 and 1. It is therefore a loxodromy with invariant geodesic e_1. Thus $e_1 \to e_1$ is the identity if and only if $e_k \to e_1$ is the identity if and only if $|\prod_{j=1}^{k} w_{ij}| = 1$. This proves the claim.

Suppose that this condition holds. We can then apply Theorem 6.3. A chart into $\mathbb{H}^3$ exists at each point of each edge of $M - \partial M$ if and only if the angles around each edge sum to 2π. Thus, in summary, $M - \partial M$ inherits a hyperbolic structure if and only if

$$|\prod_{j=1}^{k} w_{ij}| = 1 \text{ for each } i \text{ and} \tag{1}$$

$$\sum_{j=1}^{k} \arg(w_{ij}) = 2\pi \text{ for each } i. \tag{2}$$

We need to show that this is equivalent to

$$\prod_{j=1}^{k} w_{ij} = 1 \text{ for each } i. \tag{3}$$

Clearly, (1) and (2) imply (3). Also, (3) implies (1). Also, (3) implies that

$$\sum_{j} \arg(w_{ij}) = 2\pi N(i),$$

for positive integers $N(i)$. We need to show that $N(i) = 1$ for each i. Summing the above inequalities over all i, and noting that each of the six interior angles in each tetrahedron appears exactly once gives that

$$2\pi T(M) = 2\pi \sum_{i=1}^{E(M)} N(i),$$

where $T(M)$ is the number of tetrahedra of M and $E(M)$ is the number of edges of M. Thus, $T(M) \geq E(M)$, with equality if and only if $N(i) = 1$ for each i. The ideal triangulation of M induces a triangulation of ∂M with $V(\partial M)$ vertices, $E(\partial M)$ edges and $F(\partial M)$ faces. Since ∂M is a collection of tori

$$\begin{aligned} 0 = \chi(\partial M) &= V(\partial M) - E(\partial M) + F(\partial M) \\ &= V(\partial M) - 3F(\partial M)/2 + F(\partial M) \\ &= V(\partial M) - F(\partial M)/2 \\ &= 2E(M) - 2T(M). \end{aligned}$$

Hence, $E(M) = T(M)$. So, $N(i) = 1$ for each i. Therefore, (3) implies (2). □

We now apply Theorem 8.4 to the case where M is the exterior of the figure-eight knot. We assign complex numbers z_1 and z_2 to the two tetrahedra. There are two edges of M, giving the following equations:

$$\begin{aligned} 1 &= z_2 z_1 \left(1 - \frac{1}{z_2}\right) z_1 z_2 \left(1 - \frac{1}{z_1}\right) \\ 1 &= \left(\frac{1}{1 - z_2}\right)\left(1 - \frac{1}{z_1}\right)\left(\frac{1}{1 - z_2}\right)\left(\frac{1}{1 - z_1}\right)\left(1 - \frac{1}{z_2}\right)\left(\frac{1}{1 - z_1}\right). \end{aligned}$$

These equations are equivalent, since their product is $1 = 1$. To see this, note that

$$z\left(\frac{1}{1 - z}\right)\left(1 - \frac{1}{z}\right) = -1.$$

These equations can be written more neatly as

$$z_1(z_1 - 1)z_2(z_2 - 1) = 1.$$

The hyperbolic structure imposed on M in §7 was the case where $z_1 = z_2 = e^{i\pi/3}$. However, there is a one-complex-dimensional parametrisation of hyperbolic structures that arise by perturbing z_1 and z_2 from this value in such a way that the equation $z_1(z_1 - 1)z_2(z_2 - 1) = 1$ remains satisfied.

9. Completeness

We will prove the following theorem over the next few sections.

Theorem 9.1. *If M is a simply-connected complete hyperbolic n-manifold, then M is isometric to $\mathbb{H}^n$.*

Note that the any cover $\tilde{M}$ of a Riemannian manifold M inherits a Riemannian metric from M. Any covering transformations of $\tilde{M}$ are isometries. Note also that $\tilde{M}$ is complete if M is complete.

Corollary 9.2. *The universal cover of a complete hyperbolic n-manifold is isometric to $\mathbb{H}^n$.*

Theorem 9.1 is a special case of the following theorem, which we will prove over the next two sections.

Theorem 9.3. *Let G be a group of isometries of a simply-connected Riemannian manifold X. Let M be a complete simply-connected (G, X)-manifold. Then M is (G, X)-isomorphic (and hence isometric) to X.*

The hypothesis that M is complete comes in via the following well known result from Riemannian geometry.

Theorem 9.4. [Hopf-Rinow] *Let M be a connected Riemannian manifold. Then the following are equivalent:*

1. *M is complete as a metric space;*
2. *any geodesic $I \to M$ can be extended to a geodesic $\mathbb{R} \to M$;*
3. *for any $m \in M$, $\exp_m$ is defined on all of $T_m M$;*
4. *for some $m \in M$, $\exp_m$ is defined on all of $T_m M$;*
5. *any closed bounded subset of M is compact.*

The proof requires machinery from differential geometry that we have not developed. However, all we need is that (1) $\Rightarrow$ (2) when M is hyperbolic. This we now prove.

Proof. Suppose that M is a complete hyperbolic manifold. For each non-zero vector $v \in T_x M$, let $\alpha \colon I \to M$ be the geodesic with $\alpha(0) = x$ and $\alpha'(0) = v$, and where $I \subset \mathbb{R}$ is the maximal domain of definition of α. It is a consequence of Proposition 1.3 in the Introduction to Riemannian Manifolds that I is an open neighbourhood of 0. We will show that the completeness of M implies that I is closed and hence the whole of $\mathbb{R}$. Let t_i be a sequence of points in I, converging to some point $t_\infty \in \mathbb{R}$. This is a Cauchy sequence in $\mathbb{R}$. The fact that $||\alpha'(t)||$ is constant implies that $\alpha(t_i)$ is Cauchy in M. Hence, it converges to a point $y \in M$. Pick a chart $\phi: U \to D^n$ around y, where $\phi(y) = 0 \in D^n$. Then $\phi \circ \alpha$ is a Euclidean straight line approaching 0. Hence, it can be smoothly extended. □

The following is a useful way of checking completeness.

Proposition 9.5. *Let M be a metric space. Suppose that there is some family of compact subsets S_t of M (for $t \in \mathbb{R}_{>0}$) which cover M, such that S_{t+a} contains all points within distance a of S_t. Then M is complete.*

Proof. Any Cauchy sequence in M must be contained in S_t for some sufficiently large t. Hence, it converges, since S_t is compact. □

Corollary 9.6. *$\mathbb{H}^n$ is complete.*

Proof. Let $S_t = B_t(0) \subset D^n$. Then $B_t(0)$ is a closed Euclidean ball. (Here, we are implicitly applying Corollary 3.7 to geodesics through 0.) So, S_t is compact. □

Example. In the case of the hyperbolic structure on the figure-eight knot complement defined in §7, let $S_t \cap \Delta$ for each of the two ideal tetrahedra Δ be $B_t(0) \cap \Delta$, where 0 is the origin in D^3. These S_t satisfy the condition of Proposition 9.5. Hence, the hyperbolic structure is complete. Applying Corollary 9.2 to this case, we obtain the following purely topological corollary.

Corollary 9.7. *The universal cover of the figure-eight knot complement is homeomorphic to $\mathbb{R}^3$.*

According to Thurston's theorem, the complements of 'most' knots in S^3 admit a complete finite volume hyperbolic structure. Therefore, their universal covers are all homeomorphic to $\mathbb{R}^3$.

Example. Here is an example of an incomplete hyperbolic structure. Let

$$B = \{(x_1, x_2) \in U^2 : 1 \leq x_1 \leq 2\}.$$

Glue the two sides of B via the isometry $z \mapsto 2z$. The resulting space M is homeomorphic to $S^1 \times \mathbb{R}$, and inherits a hyperbolic structure. This is incomplete: here is a Cauchy sequence which does not converge. Let $z_i = (1, 2^i) \in U^2$. This is identified with the point $(2, 2^{i+1})$. So,

$$d(z_i, z_{i+1}) \leq d_{hyp}((2, 2^{i+1}), (1, 2^{i+1})) < 1/2^{i+1}.$$

So, this sequence is Cauchy in M. However, it does not converge to a point in M, since the x_2 co-ordinates tend to ∞.

Theorem 9.3 is proved by defining a local (G, X)-isomorphism (and hence a local isometry) from M to X. The existence of this local isometry together with the following proposition will prove Theorem 9.3.

Proposition 9.8. *Let $h\colon M \to N$ be a local isometry between Riemannian manifolds, where M is complete. Then h is a Riemannian covering map.*

Proof. Let x be any point in N, and let y_i be the points of $h^{-1}(x)$. By Proposition 1.3 of the Introduction to Riemannian Manifolds, there is some r such that $\exp_x$ maps $B_r(0)$ diffeomorphically onto $B_r(x)$. Let $U_i = \exp_{y_i}(B_r(0))$. This is well-defined since $\exp_{y_i}$ is defined on all of $T_{y_i}M$ by the Hopf-Rinow theorem. Recall from Proposition 1.4 of the Introduction to Riemannian Manifolds that the following diagram commutes:

$$\begin{array}{ccc} T_{y_i}M & \xrightarrow{(Dh)_{y_i}} & T_xN \\ \Big\downarrow \exp_{y_i} & & \Big\downarrow \exp_x \\ M & \xrightarrow{h} & N \end{array}$$

Claim. $h|_{U_i}$ is a diffeomorphism onto its image.

The top and right arrows in the above diagram are diffeomorphisms when restricted to $B_r(0)$. Hence the bottom arrow must be a diffeomorphism when restricted to U_i.

Claim. $U_i \cap U_j = \emptyset$ if $i \neq j$.

If U_i and U_j overlap at a point m, there are vectors v_i and v_j in $T_{y_i}M$ and $T_{y_j}M$ with length at most r such that $m = \exp_{y_i}(v_i) = \exp_{y_j}(v_j)$. Using the above commutative diagram and the fact that $\exp_x$ is injective on $B_r(0)$, we deduce that

$(Dh)_{y_i}(v_i) = (Dh)_{y_j}(v_j)$. Hence, for all $t > 0$, $h \circ \exp_{y_i}(tv_i) = h \circ \exp_{y_j}(tv_j)$. By considering t near 1 and using the fact that h is injective near m, the geodesics $\exp_{y_i}(tv_i)$ and $\exp_{y_j}(tv_j)$ have the same derivative at $t = 1$. Hence they agree for all t. In particular, $y_i = \exp_{y_i}(0) = \exp_{y_j}(0) = y_j$. This proves the claim.

Claim. $\bigcup_i U_i = h^{-1}(B_r(x))$

Clearly, $\bigcup_i U_i \subset h^{-1}(B_r(x))$. To establish the opposite inclusion, consider a point $m \in M$ that is sent to a point in $B_r(x)$. Some geodesic of length $l < r$ runs from $h(m)$ to x, with velocity vector v at $h(m)$. There is a geodesic emanating from m with derivative $(Dh)_m^{-1}(v)$. Since M is complete, the Hopf-Rinow theorem gives that this geodesic is still defined after time l. By that stage it has reached some y_i. Hence, m lies in U_i.

This sequence of claims establishes that h is a covering map. □

10. The developing map

The goal of this section is prove Theorem 9.3. In fact, we will prove the following stronger result.

Theorem 10.1. *Let G be a group acting rigidly on a manifold X. Let M be a simply-connected (G, X)-manifold. Then any connected chart $\phi: U_0 \to X$ can be extended to a local (G, X)-isomorphism $D: M \to X$, known as a developing map.*

Note that in the above, we did not assume that X was a Riemannian manifold, complete or incomplete.

Proof. Start with a chart $\phi_0: U_0 \to X$, where U_0 is connected. Pick a basepoint $x_0 \in U_0$. We wish to define $D(x)$ for all points $x \in M$. There is a path $\alpha: [0, 1] \to M$ from x_0 to x. We will now define a corresponding path $\beta: [0, 1] \to X$, such that $\beta = \phi_0 \circ \alpha$ wherever the maps are defined. In particular, $\beta(0) = \phi_0(x_0)$. We will define $D(x)$ as $\beta(1)$.

For each point $t \in [0, 1]$, pick a chart around $\alpha(t)$. The inverse image of these charts forms an open cover of $[0, 1]$. By replacing each set with its connected components, we obtain a new open cover of $[0, 1]$. Since $[0, 1]$ is compact, there is a finite subcover $\{(t_0^-, t_0^+) \dots, (t_m^-, t_m^+)\}$ of $[0, 1]$. The interval (t_i^-, t_i^+) is open, half-open or closed in $\mathbb{R}$ as appropriate. Each (t_i^-, t_i^+) is a connected component of $\alpha^{-1}(U_i)$ for some chart $\phi_i\colon U_i \to X$. We may assume that (t_0^-, t_0^+) is the component of $\alpha^{-1}(U_0)$ containing 0. We may also assume that, if $i < j$, then $t_i^- \leq t_j^-$ and $t_i^+ \leq t_j^+$. This implies that $t_i^+ > t_{i+1}^-$ for all i.

The interval $(t_i^-, t_i^+) \cap (t_{i+1}^-, t_{i+1}^+)$ maps to a path in M, lying entirely in a connected component V_i of $U_i \cap U_{i+1}$. This has an associated transition map $\phi_{i+1} \circ \phi_i^{-1}$, which is the restriction of an element $h_i\colon X \to X$ of G. For $t \in [0,1]$, we pick (t_i^-, t_i^+) containing t and define $\beta(t)$ as

$$\beta(t) = h_0^{-1} \circ \ldots \circ h_{i-1}^{-1} \circ \phi_i \circ \alpha(t).$$

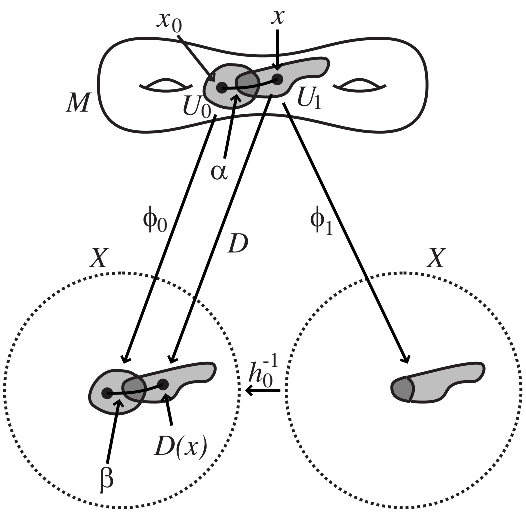

Figure 28.

We now show that this is independent of the choice of (t_i^-, t_i^+) containing t. If $t \in (t_k^-, t_k^+)$, with $k > i$, then $t \in (t_j^-, t_j^+)$ for all $i \leq j \leq k$. Then $\alpha(t)$ lies in $V_i, V_{i+1}, \ldots, V_{k-1}$. On these sets, the following maps are equal: $\phi_i = h_i^{-1} \circ \phi_{i+1}$, …, $\phi_{k-1} = h_{k-1}^{-1} \circ \phi_k$. Hence,

$$h_0^{-1} \circ \ldots \circ h_{i-1}^{-1} \circ \phi_i \circ \alpha(t) = h_0^{-1} \circ \ldots \circ h_{k-1}^{-1} \circ \phi_k \circ \alpha(t).$$

We now show that β is independent of the choice of open cover of $[0,1]$. Suppose that $\{(\hat{t}_0^-, \hat{t}_0^+), \ldots, (\hat{t}_{\hat{m}}^-, \hat{t}_{\hat{m}}^+)\}$ is another cover of $[0,1]$, with each $(\hat{t}_i^-, \hat{t}_i^+)$ being a connected component of $\alpha^{-1}(\hat{U}_i)$ for some chart $\hat{\phi}_i\colon \hat{U}_i \to X$. Let $\hat{h}_i$, $\hat{V}_i$ and $\hat{\beta}$ be as above. We claim that the maps

$$\phi_0$$

$$h_0^{-1} \circ \phi_1$$

$$h_0^{-1} \circ h_1^{-1} \circ \phi_2$$

$$\dots$$

$$h_0^{-1} \circ \dots \circ h_{m-1}^{-1} \circ \phi_m$$

and

$$\phi_0$$

$$\hat{h}_0^{-1} \circ \hat{\phi}_1$$

$$\hat{h}_0^{-1} \circ \hat{h}_1^{-1} \circ \hat{\phi}_2$$

$$\dots$$

$$\hat{h}_0^{-1} \circ \dots \circ \hat{h}_{m-1}^{-1} \circ \hat{\phi}_m$$

are equal on sets on which they are defined. If not, there is an infimal value of t (t_{inf}, say) such that $t \in (t_i^-, t_i^+) \cap (\hat{t}_k^-, \hat{t}_k^+)$ and

$$h_0^{-1} \circ \dots h_{i-1}^{-1} \circ \phi_i \neq \hat{h}_0^{-1} \circ \dots \hat{h}_{k-1}^{-1} \circ \hat{\phi}_k.$$

Then there is a $t < t_{\mathrm{inf}}$ such that $t \in (t_{i-1}^-, t_{i-1}^+) \cap (t_i^-, t_i^+) \cap (\hat{t}_k^-, \hat{t}_k^+)$ or $(t_i^-, t_i^+) \cap (\hat{t}_{k-1}^-, \hat{t}_{k-1}^+) \cap (\hat{t}_k^-, \hat{t}_k^+)$. Suppose the former. Then,

$$h_0^{-1} \circ \dots h_{i-2}^{-1} \circ \phi_{i-1} = \hat{h}_0^{-1} \circ \dots \hat{h}_{k-1}^{-1} \circ \hat{\phi}_k.$$

But

$$h_0^{-1} \circ \dots h_{i-2}^{-1} \circ \phi_{i-1} = h_0^{-1} \circ \dots h_{i-1}^{-1} \circ \phi_i.$$

The crucial fact we are using here is that if two elements of G agree on an open set, then they are equal, since G acts rigidly.

Now define $D(x)$ as $\beta(1)$. We need to show that this is independent of the choice of α. If $\hat{\alpha}$ is another path from x_0 to x, then there is a homotopy $H\colon [0,1] \times [0,1] \to M$ between α and $\hat{\alpha}$, keeping their endpoints fixed, since M is simply-connected. Pick a cover $\mathcal{C}$ of charts for M, one being U_0, the remainder being connected open sets disjoint from x_0. Subdivide $[0,1]$ into intervals,

$[0/N, 1/N], [1/N, 2/N]$, etc. Since $[0,1] \times [0,1]$ is compact, we may pick N large enough so that each square $[y/N, (y+1)/N] \times [z/N, (z+1)/N]$ lies within $H^{-1}(U)$ for some U of $\mathcal{C}$. So, $H([y/N, (y+1)/N] \times [z/N, (z+1)/N])$ lies within a compact subset K of U. Hence, there is a collection of paths $\alpha = \alpha_0, \alpha_1, \ldots, \alpha_k = \hat{\alpha}$, such that α_j and α_{j+1} differ only by a homotopy which alters the path within a compact subset K of U, keeping their endpoints fixed. By removing K from all the charts of $\mathcal{C}$ other than U, we may assume that that U is one of the charts U_i in the definition of β_j and β_{j+1}, and that the sets $U_0, \ldots, U_m, V_0, \ldots, V_{m-1}$ are the same for both α_j and α_{j+1}. So, β_j and β_{j+1} differ by a homotopy keeping their endpoints fixed. This does not alter $D(x)$. Thus, D is a well-defined map.

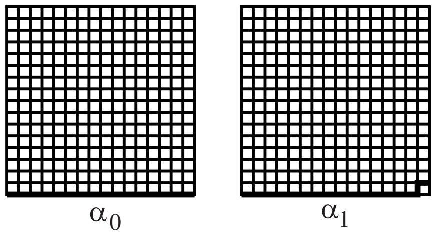

Figure 29.

We now show that D is a local (G, X)-isomorphism in a neighbourhood of each point $x \in M$. By definition, $D|_{U_0} = \phi_0$. So, we may assume that $x \neq x_0$. Pick a connected chart U_m around x not containing x_0, and let U be a connected open neighbourhood of x whose closure lies in U_m. Cover M by charts, one of which is U_m, the remainder being disjoint from U. Pick any path α from x_0 to x. Then, using the above cover,

$$D(x) = h_0^{-1} \circ \ldots \circ h_{m-1}^{-1} \circ \phi_m \circ \alpha(1) = h_0^{-1} \circ \ldots \circ h_{m-1}^{-1} \circ \phi_m(x),$$

where $h_i \colon X \to X$ are the relevant transition maps and $\phi_m \colon U_m \to \mathbb{H}^n$ is the chart. Define $D|_U$ using extensions of α within U. Then, $D|_U = h_0^{-1} \circ \ldots \circ h_{m-1}^{-1} \circ \phi_m$. Hence, $D|_U$ is a (G, X)-isomorphism onto its image. □

Example. Let $B = \{(x_1, x_2) \in U^2 : 1 \leq x_1 \leq 2\}$ and let M be the incomplete hyperbolic manifold obtained by gluing the two sides of B via the isometry $z \mapsto 2z$.

Let $p: \tilde{M} \to M$ be the universal cover. The sets

$$A_0 = \{(x_1, x_2) \in U^2 : 1 < x_1 < 2\}$$
$$A_1 = \{(x_1, x_2) \in U^2 : 1 \le x_1 < 3/2\}$$
$$\cup \{(x_1, x_2) \in U^2 : 3/2 < x_1 \le 2\}$$

form charts for M. These lift to charts for $\tilde{M}$. Let $\phi_0: U_0 \to A_0$ be the initial chart. Then, the associated developing map has image $\{(x_1, x_2) \in U^2 : 0 < x_1\}$.

The following lemma implies that, given the initial chart $\phi_0: U_0 \to X$, the developing map is unique.

Lemma 10.2. *Let G be a group acting rigidly on a space X. Let f_1 and f_2 be local (G, X)-isomorphisms between (G, X)-manifolds M and N, where M is connected. If f_1 and f_2 agree on some open set, then $f_1 = f_2$.*

Proof. Consider the set

$$V = \{x \in M : f_1 \text{ and } f_2 \text{ agree in some open neighbourhood of } x\}.$$

Clearly, V is open. We will now show that it is closed and so the whole of M. Consider a sequence of points $x_i \in V$ tending to x_∞. There are connected charts $\phi_M: U_M \to X$ and $\phi_N: U_N \to X$ around x_∞ and $f_1(x_\infty) = f_2(x_\infty)$. Then for $i = 1$ and 2, there are elements g_i of G such that $g_i = \phi_N \circ f_i \circ \phi_M^{-1}$, where this equality holds in a set U where the maps are defined. For i sufficiently large, x_i belongs to $\phi_M^{-1}(U)$. In some neighbourhood of x_i, f_1 and f_2 agree. So g_1 and g_2 agree on some open set. Therefore, $g_1 = g_2$, since G is rigid. So, $f_1 = f_2$ on $\phi_M^{-1}(U)$. Therefore x_∞ is in V. □

Let M be a (G, X)-manifold, where G is rigid. Let $p: \tilde{M} \to M$ be the universal cover. Give $\tilde{M}$ its inherited (G, X)-structure. Then each covering transformation $\tau: \tilde{M} \to \tilde{M}$ is a (G, X)-isomorphism. Pick a basepoint x_0 in $\tilde{M}$ and a connected chart $\phi: U_0 \to X$ around it. Let $D: \tilde{M} \to X$ be the associated developing map. Note that $\phi \circ \tau: \tau^{-1}(U_0) \to X$ is a chart around $\tau^{-1}(x_0)$. We may use this chart to define D in a neighbourhood of $\tau^{-1}(x_0)$. Then $D|_{\tau^{-1}(U_0)} = g_\tau^{-1} \circ \phi \circ \tau$, for some element g_τ of G. So, the following commutes on $\tau^{-1}(U_0)$:

$$\begin{array}{ccc} \tilde{M} & \xrightarrow{\tau} & \tilde{M} \\ \downarrow D & & \downarrow D \\ X & \xrightarrow{g_\tau} & X \end{array}$$

By Lemma 10.2, this diagram commutes on all of M. If we compose two covering transformations τ and σ, then (by pasting two commutative diagrams together) $g_{\sigma\tau} = g_\tau g_\sigma$. So, we have a group homomorphism

$$\eta: \pi_1(M) \to G$$

known as the *holonomy*. It depends on the choice of chart $\phi: U_0 \to X$ and of basepoint. Different choices lead to a holonomy which differs by conjugation by an element of G.

In the case where X is a simply-connected Riemannian manifold, G is a group of isometries and M is complete, then $D: \tilde{M} \to X$ is a (G, X)-isomorphism. This identifies $\tilde{M}$ with X. Hence, we have the following.

Proposition 10.3. *Let X be a simply-connected Riemannian manifold and let G be a group of isometries. Let M be a complete (G, X)-manifold. Then $\eta(\pi_1(M))$ is a group of covering transformations of X, and M is the quotient $X/\eta(\pi_1(M))$.*

Now consider the case where $X = \mathbb{H}^n$ and $G = \mathrm{Isom}(\mathbb{H}^n)$, and where M is complete. Then, η sends each non-trivial element of $\pi_1(M)$ to an isometry of $\mathbb{H}^n$ which fixes no point of $\mathbb{H}^n$.

Corollary 10.4. *The holonomy homomorphism η is injective and its image contains no elliptic isometries other than the identity.*

11. Topological properties of complete hyperbolic manifolds

Lemma 11.1. *$\mathbb{H}^n$ has infinite volume.*

Proof. The volume of U^n is

$$\int_{(x_1,\ldots,x_n)\in U^n} \frac{1}{x_n^n} dx_1 \ldots dx_n = \infty \qquad \square$$

Proposition 11.2. *A complete finite volume hyperbolic manifold M has infinite fundamental group.*

Proof. If not, the universal cover $p: \tilde{M} \to M$ would be finite-to-one. Then $\tilde{M}$ would have finite volume. But $\tilde{M}$ is isometric to $\mathbb{H}^n$. $\square$

Corollary 11.3. *S^n does not admit a hyperbolic structure for $n > 1$.*

Recall Thurston's geometrisation theorem:

Theorem. [Thurston] *Let M be a compact orientable irreducible atoroidal 3-manifold-with-boundary, such that ∂M is a non-empty collection of tori. Then either $M - \partial M$ has a complete finite volume hyperbolic structure, or M is homeomorphic to one of the following exceptional cases:*

1. $S^1 \times [0, 1] \times [0, 1]$
2. $S^1 \times S^1 \times [0, 1]$
3. *the space obtained by gluing the faces of a cube as follows: arrange the six faces into three opposing pairs; glue one pair, by translating one face onto the other; glue another pair, by translating one face onto the other and then rotating through π about the axis between the two faces.*

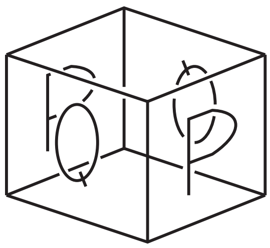

Figure 30.

We are now going to investigate to what extent these conditions are necessary.

Theorem 11.4. *Any complete hyperbolic* 3*-manifold M is irreducible.*

Proof. Let S^2 be a smoothly embedded 2-sphere in M. Since S^2 is simply-connected, the universal cover $p\colon \mathbb{H}^3 \to M$ restricts to a homeomorphism on each component of $p^{-1}(S^2)$. Pick one component S^2_1 of $p^{-1}(S^2)$. Then, by the Schoenflies theorem, the closure of one component of $\mathbb{H}^3 - S^2_1$ is homeomorphic to a closed 3-ball B_1 which is therefore compact. Therefore, there can only be finitely many covering translates of S^2_1 in B_1. Pick one innermost in B_1. This bounds a closed ball B_2. The covering translates of B_2 are all disjoint. Therefore, p projects B_2 homeomorphically to a closed ball in M. The boundary of this ball is the original S^2. □

We now investigate atoroidality. In fact any complete finite volume hyperbolic manifold is atoroidal, but the proof of this requires some understanding of the ends of these manifolds. Instead, we focus on the closed case. Even here, we need some extra geometric concepts.

Definition. Let M be a Riemannian manifold. For each $x \in M$, define the *injectivity radius* at x to be

$$\mathrm{inj}(x) = \sup\{\epsilon : \exp_x \text{ is injective on } B_\epsilon(0) \subset T_xM\}.$$

Note that there is such an $\epsilon > 0$, by Proposition 1.3 of the Introduction to Riemannian Manifolds.

Proposition 11.5. *Let M be a complete hyperbolic manifold. Let $p\colon \mathbb{H}^n \to M$ be the universal cover. For each $x \in M$, pick a point $\tilde{x} \in p^{-1}(x)$. Let*

$$i(\tilde{x}) = \sup\{\epsilon : \gamma(B_\epsilon(\tilde{x})) \cap B_\epsilon(\tilde{x}) = \emptyset \textit{ for all } \gamma \in \eta(\pi_1(M)) - \mathrm{id}\}.$$

Then $i(\tilde{x}) = \mathrm{inj}(x)$.

Proof.

Claim. $i(\tilde{x}) \geq \mathrm{inj}(x)$.

Pick $\epsilon < \mathrm{inj}(x)$ arbitrarily close to $\mathrm{inj}(x)$. Suppose that $\gamma(B_\epsilon(\tilde{x})) \cap B_\epsilon(\tilde{x}) \neq \emptyset$, for some $\gamma \in \eta(\pi_1(M)) - \mathrm{id}$. Let z be a point in their intersection, which we may assume to be on the geodesic joining $\tilde{x}$ to $\gamma(\tilde{x})$. Then $z = \exp_{\tilde{x}}(v_1) = \exp_{\gamma(\tilde{x})}(v_2)$ for vectors v_1 and v_2 with length less than ϵ. Note that if $(T\gamma)_{\tilde{x}}(v_1) = v_2$, then γ would preserve the geodesic between $\tilde{x}$ and $\gamma(\tilde{x})$ and would reverse its orientation. This would imply that γ fixed some point in $\mathbb{H}^n$, which is impossible. Hence, $(Tp)_{\tilde{x}}(v_1) \neq (Tp)_{\gamma(\tilde{x})}(v_2)$. These are distinct vectors in T_xM with length less than ϵ which map to the same point $p(z)$ in M. This contradicts the definition of $\mathrm{inj}(x)$.

Claim. $\mathrm{inj}(x) \geq i(\tilde{x})$.

Pick $\epsilon < i(\tilde{x})$ arbitrarily close to $i(\tilde{x})$. Then $p|_{B_\epsilon(\tilde{x})}$ is an isometry onto its image. Hence, we may take its inverse ϕ to be a chart for M around x. Since $\exp_{\phi(x)}$ is injective, $\exp_x$ is injective on $B_\epsilon(0) \subset T_xM$. □

Corollary 11.6. *For a complete hyperbolic manifold M,* inj *is a continuous function on M.*

Proof. Let x_1 and x_2 be points in M. Let $\tilde{x}_1$ and $\tilde{x}_2$ be points in $p^{-1}(x_1)$ and $p^{-1}(x_2)$ such that $d(\tilde{x}_1, \tilde{x}_2) = d(x_1, x_2)$. Then, a ball of radius ϵ around $\tilde{x}_1$ contains a ball of radius $\epsilon - d(\tilde{x}_1, \tilde{x}_2)$ around $\tilde{x}_2$. So, by Proposition 11.5,

$$\mathrm{inj}(x_2) \geq \mathrm{inj}(x_1) - d(x_1, x_2).$$

The same is true with the rôles of x_1 and x_2 reversed. Hence,

$$|\mathrm{inj}(x_1) - \mathrm{inj}(x_2)| \leq d(x_1, x_2).$$

So, inj is continuous. $\square$

Corollary 11.7. *If M is a closed hyperbolic manifold, there is a positive lower bound on* $\mathrm{inj}(x)$ *for all $x \in M$.*

Proposition 11.8. *Let M be a closed hyperbolic manifold. Then the image of each non-trivial element of $\pi_1(M)$ under η is a loxodromic isometry.*

Proof. Let γ be non-trivial element of $\pi_1(M)$. By Corollary 10.4, $\eta(\gamma)$ is non-trivial and non-elliptic. It is therefore parabolic or loxodromic. If it is parabolic, then by Proposition 4.5, it is conjugate to the isometry as in Example 2 of §2 with $\lambda = 1$. Recall that this is a Euclidean isometry $h\colon U^n \to U^n$ which fixes the n^{th} co-ordinate. For $(x_1, \dots, x_n) \in U^n$,

$$d_{hyp}(h(x_1, \dots, x_n), (x_1, \dots, x_n)) < d_{Eucl}(h(x_1, \dots, x_n), (x_1, \dots, x_n))/x_n \to 0$$

as $x_n \to \infty$. So, $\mathrm{inj}(p(x_1, \dots, x_n)) \to 0$ as $x_n \to \infty$. This contradicts Corollary 11.7. $\square$

Lemma 11.9. *Let f and g be commuting functions from a set to itself. Then g maps the fixed point set of f to itself.*

Proof. If $f(x) = x$, then $g(x) = gf(x) = fg(x)$. $\square$

Theorem 11.10. *Let M be a closed hyperbolic n-manifold. Then no subgroup of $\pi_1(M)$ is isomorphic to $\mathbb{Z} \oplus \mathbb{Z}$.*

Proof. Let γ_1 and γ_2 be commuting elements of $\pi_1(M)$. Then, $\eta(\gamma_1)$ and $\eta(\gamma_2)$ are loxodromic, by Proposition 11.8. Since they commute, $\eta(\gamma_1)$ maps the fixed point set of $\eta(\gamma_2)$ to itself. It therefore preserves the geodesic α left invariant by $\eta(\gamma_2)$. It cannot reverse the orientation of α, since then it would have a fixed point on α.

Thus, $\eta(\gamma_1)$ has the same fixed point set as $\eta(\gamma_2)$. They both translate points on α some fixed hyperbolic distance along α. There is a uniform lower bound on this translation length for all non-identity elements of the group generated by $\eta(\gamma_1)$ and $\eta(\gamma_2)$. Therefore, some power of $\eta(\gamma_1)$ equals some power of $\eta(\gamma_2)$. So, γ_1 and γ_2 do not generate a subgroup isomorphic to $\mathbb{Z} \oplus \mathbb{Z}$. $\square$

Remark. The above theorem is false if 'closed' is replaced with 'complete and finite volume'. For example, it can be shown that for any knot K in S^3 other than the unknot, the map $i_*: \pi_1(\partial N(K)) \to \pi_1(S^3 - K)$ is injective, where i is the inclusion map of the boundary of the tubular neighbourhood $N(K)$ into $S^3 - K$. However, for 'most' knots K, $S^3 - K$ admits a complete finite volume hyperbolic structure. In these cases, $\eta(i_*\pi_1(\partial N(K)))$ is group of commuting parabolic isometries.

Corollary 11.11. *The n-manifold $S^1 \times \ldots \times S^1$ $(n > 1)$ does not admit a hyperbolic structure.*

Corollary 11.12. *A closed hyperbolic 3-manifold is atoroidal.*