IISER Pune

Mathematics Project Report

Hyperbolic Knot Theory

By
Sriram Raghunath
Reg. No: 20151036

Under the guidance of
Prof. Tejas Kalelkar

May 4, 2019

1 Introduction

A knot is a piecewise linear embedding $K\colon S^1 \to S^3$. The image of this embedding (also denoted as K) has a tubular neighbourhood around it in S^3, denoted as $N(K)$, which is homeomorphic to the interior of the solid torus. The complement of the knot is the compact 3-manifold $S^3 \setminus N(K)$, which has a torus boundary. We will often use the term knot complement to denote the interior of $S^3 \setminus N(K)$; what we mean should be clear from the context.

Two knots K_1 and K_2 are said to be equivalent if there exists an orientation preserving homeomorphism $h\colon S^3 \to S^3$ such that $h(K_1) = K_2$. Gordon and Luecke proved that two knots are equivalent if and only if there exists an orientation preserving homeomorphism between the knot complements. So, invariants that can distinguish between the knot complements will give rise to knot invariants.

Many knots have a complement in the 3-sphere that can be given a hyperbolic structure; such knots are called hyperbolic knots. Studying this hyperbolic structure can give us a lot of information about these knots. One approach to get a hyperbolic structure on a knot complement is by triangulating the knot complement. We can decompose the knot complement into ideal tetrahedra, giving a hyperbolic metric on the individual tetrahedra and if the gluing of these tetrahedra satisfies certain conditions, then the knot complement will have a hyperbolic structure.

In this project, we study how we can decompose alternating knot complements into ideal polyhedra which can be algorithmically obtained from the knot diagram, based on an approach by Menasco. We then try to understand the classification of hyperbolic isometries with a view towards understanding the hyperbolic structure on knot complements. We also study the ideas of developing map and holonomy and how they can be used to determine completeness of hyperbolic structures. None of the material in this project is our original work; we have drawn our material from the books “Hyperbolic Geometry and Knot Theory” by Jessica Purcell, “The geometry and topology of three manifolds” by William P. Thurston, and “Geometry of surfaces” by John Stillwell to understand the topics in this report (the full citations are given in the references).

2 Polyhedral decomposition of the knot complement

In this section, we try to understand an approach to the polyhedral decomposition of the figure 8 knot complement and how this approach can be extended to decompose any alternating link complement into polyhedra. The decomposition for the figure 8 knot was first given by Thurston and the Menasco generalized the approach to all link complements. Our presentation will be similar to the one in the e-book ‘Hyperbolic Geometry and Knot Theory’ by Jessica Purcell.

Here, we use the term polyhedron to refer to a 3-ball with vertices and edges marked on its boundary, such that the faces that these vertices and edges demarcate simply connected regions with disjoint interiors. This definition will allow polyhedra to contain monogons and bigons, that is, faces bordered by only one and two edges respectively. An ideal polyhedron is just a polyhedron without any of its vertices.

2.1 An overview of the method

We will decompose the figure 8 knot complement into two ideal tetrahedra. The knot is embedded in S^3, which we visualize as $\mathbb{R}^3 \cup \{\infty\}$. Consider the knot in three space placed on the XY plane, such that wherever there are overcrossings, the knot goes above the XY plane and wherever there are undercrossings, the knot goes below the XY plane. Insert one ideal edge (an edge with its vertices removed) in the space for every crossing of the knot, running from the overstrand to the understrand. Imagine two balloons (3-balls) expanding from the point at ∞, one from above the XY plane and one from below it. When the balloons meet the XY plane, they will give rise to faces and edges on the surface of the balloon (3-ball). The faces will correspond to the regions cut out by the graph of the knot diagram, and the edges we inserted will border these faces. These two balls with the correct edges identified in pairs will give back the knot complement. We explain this in more detail below.

2.2 The expansion of the balloon into a crossing

We show in the pictures below how the balloon expands into different regions of the knot. When the balloon expands into an isolated strand, it just wraps over the strand and the strand leaves an indentation (or a hole) on the surface of the balloon, since the knot itself does not belong to the knot complement.

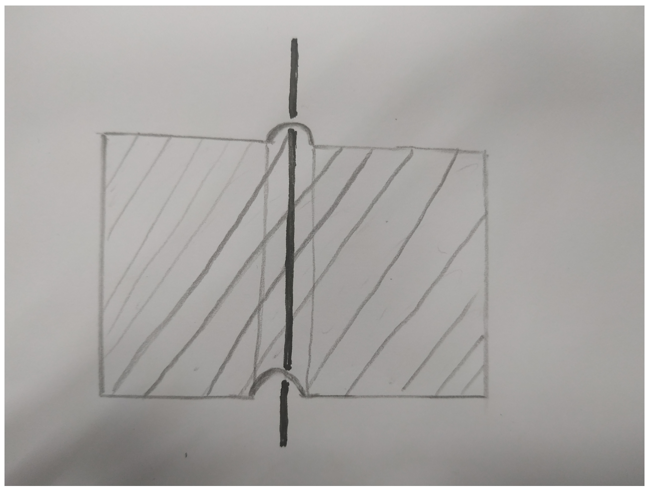

Figure 1: This is how the balloon expands into a single strand

When the balloon expands into a crossing, it wraps around the strands of the knot as indicated in the diagram below. We had inserted an ideal edge in the knot complement which runs from the overstrand to the understrand. We have now split this edge into the two orange edges on the surface of the balloon (3-ball); these edges will be identified later. We see that the crossing gives rise to four faces, the faces S and T are adjacent via an orange edge, and the faces U and V are adjacent via the other orange edge.

This is just the picture for the balloon on the top. For the balloon on the bottom, we have to view the knot from below. Then what were previously overstrands will now appear as understrands and vice versa. Previously the edge had split from the overstrand to the understrand, now it will split in the opposite way so that now faces S and V are adjacent to each other and faces T and U are adjacent to each other. Note that we now have four copies of each edge, two copies on the top balloon and two on the bottom balloon. If we do this procedure at every crossing and shrink the arcs of the knot to ideal vertices, we will obtain the two polyhedra which will give us the knot complement when all the faces and edges are identified correctly.

2.3 The figure 8 knot complement

We will now see how to implement this procedure for the figure 8 knot. We first explain how to obtain the top polyhedron. First we need to draw the knot diagram and sketch the faces and edges into the knot diagram (shown in figure 3 (a)). The knot diagram gives rise to six faces on the surface of the 3-ball, labeled A, B, C, D, E and F. The face D lies above the knot and contains the point at ∞. As explained before, we have split each edge running from the overstrand to the understrand into two edges which will be identified

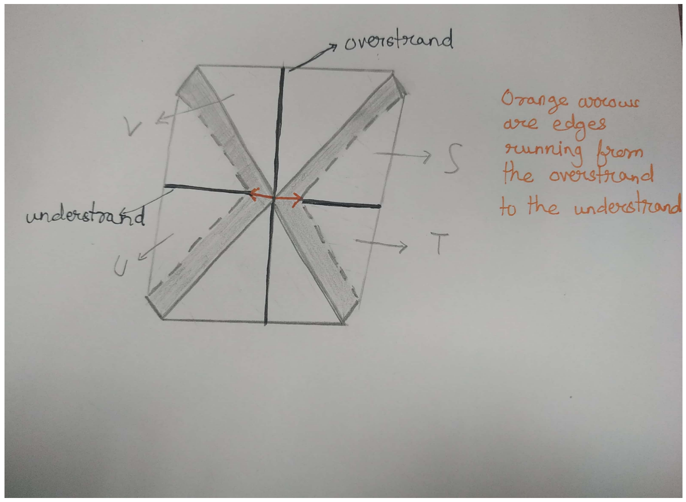

Figure 2: This is how the balloon expands into a crossing

later.

Now, the arcs of the knot (from one crossing to the next) do not belong to the knot complement, so we can shrink each arc to an ideal vertex. We can visualize this as the surface stretching to fill in the gap left by the arcs of the knot, till each arc reduces to a point and the edges have stretched to meet at this point. We shrink each arc to the point of crossing and extend each edge on the surface of the 3-ball to meet at these ideal vertices. At the end, we will obtain the top polyhedron, which is shown in figure 4 (a).

Now we construct the bottom polyhedron. First we draw the knot diagram as seen from below, where the overstrands of the top knot diagram become understrands and vice versa. The edges run in the same direction as before, which means that in this diagram, they run from the understrand to the overstrand. As we explained earlier, the edges now split in the other way and this can be seen from figure 3(b). The face D now lies below the knot, and contains the point at ∞. Again we shrink the arcs of the knot to ideal vertices by extending the edges. This will give us the bottom polyhedron, shown in figure 4(b).

We see that the top and bottom polyhedron are essentially just the graph of the knot diagram drawn on the surface of the 3-ball. The vertices of the polyhedron corresponding to the crossings of the knot diagram and the edges of the polyhedron correspond to the arcs of the knot diagram. Notice that in both the bottom and top polyhedron, a given face is bordered by the same edges, and this can be seen by understanding how the balloon expands into a crossing (figure 2). The edge inserted at every crossing borders all the faces which share that vertex (crossing), only the adjacency of faces changes in the top and bottom polyhedron. For example, the faces A and D are adjacent to each other via an orange edge and the faces E and B are adjacent to each other via another orange edge in the top polyhedron. However, in the bottom polyhedron, the faces A and E are adjacent via an orange edge and the faces B and D are adjacent via another orange edge.

We also see that the faces in the top and bottom polyhedron are identified to each other by rotation. For example, face A on the top polyhedron is identified with the face A on the bottom via clockwise rotation, and face C on the top is identified with face C on the bottom via anticlockwise rotation. Adjacent faces are identified by opposite rotation; this is what Thurston calls "gear" rotation. The knot diagram then admits a checkerboard colouring; the white faces on the top are rotated anticlockwise and the shaded faces on the top are rotated clockwise before identification with the corresponding face on the bottom polyhedron. We explain this rotation using the example of face A and the orange edge. When we go from the top knot

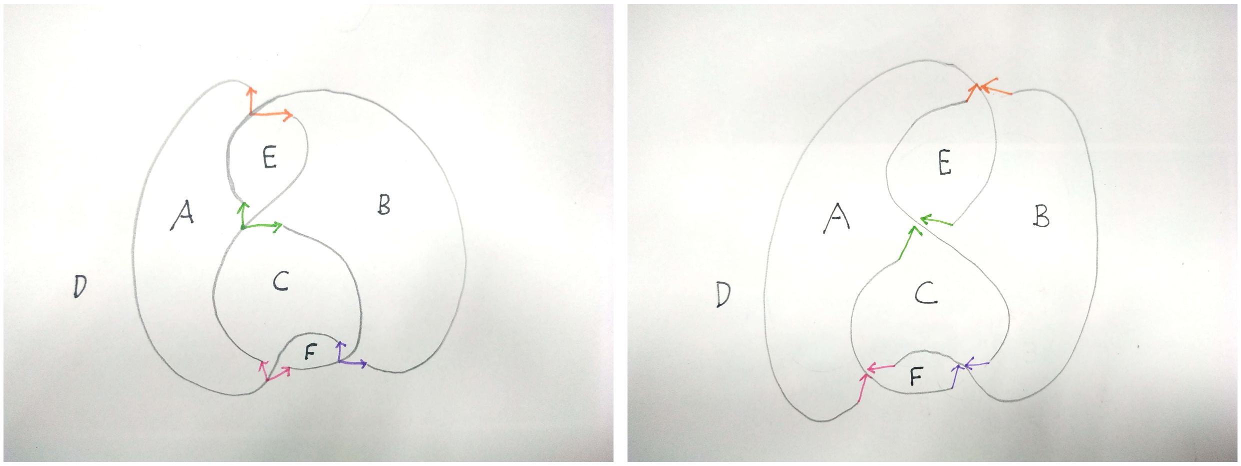

(a) Knot viewed from the top

(b) Knot viewed from the bottom

Figure 3: Knot diagrams with faces and edges sketched for the top and bottom polyhedra

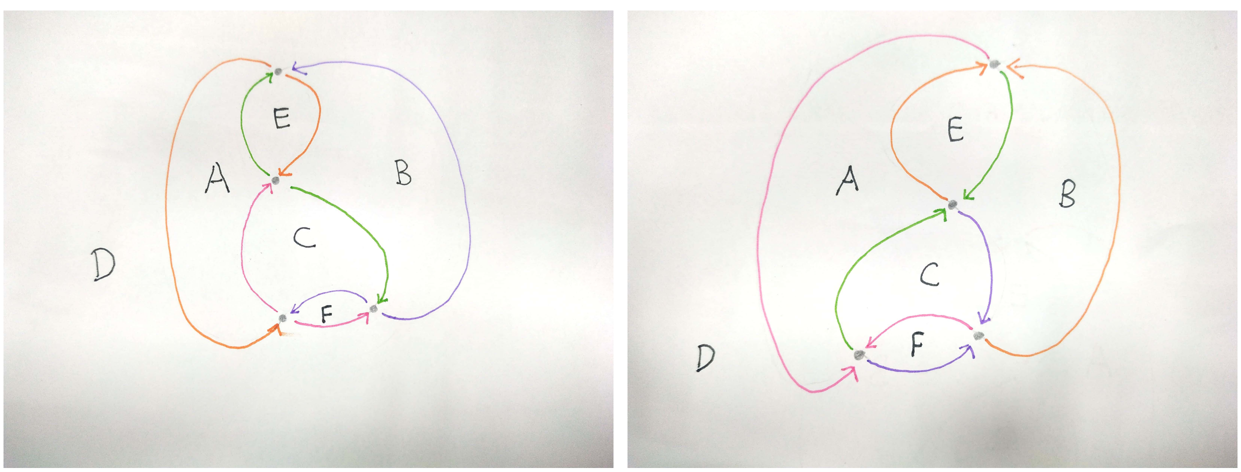

(a) The top polyhedron

(b) The bottom polyhedron

Figure 4: The top and bottom polyhedra

diagram (figure 3(a)) to the top polyhedron (figure 3(b)), we extend the orange edge bordering face A to the next vertex of face A (in the anticlockwise direction). Whereas, when we go from the bottom knot diagram (figure 3(a)) to the bottom polyhedron (figure 3(b)), we extend the orange edge bordering face A to the previous vertex of face A (in the anticlockwise direction). All other edges bordering face A undergo similar rotation, leading to the clockwise rotation of face A. Similarly, every face gets rotated, and adjacent faces are rotated in the opposite direction.

The actual reason behind gear rotation is that at each vertex, pairs of edges of the same colour are rotated clockwise from the top polyhedron to the bottom polyhedron, because of the opposite splitting of edges in the top polyhedron and the bottom polyhedron. This means that if the orange edge is rotated clockwise for face A, it gets rotated anticlockwise for face E. This can be seen in figure 6.

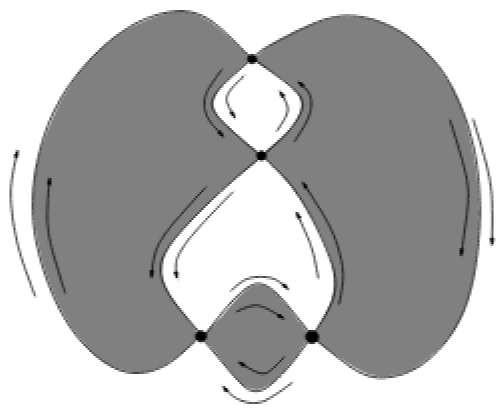

Figure 5: Checkerboard colouring and gear rotation for the figure 8 knot diagram - Image from pg.13, Purcell, J. (2019). Hyperbolic Geometry and Knot Theory. [1]

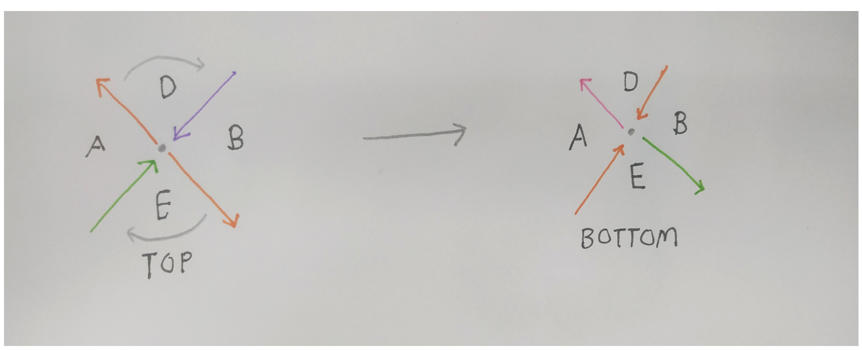

Figure 6: Reason behind gear rotation

Now, if we identify all the faces in the top polyhedron to the bottom polyhedron after gear rotation, and then identify the edges of the same colour in the resulting space, we will recover the figure 8 knot complement. However, both the top and bottom polyhedron have bigons (faces with only two sides). So we contract these bigons on the surface of both the polyhedra and identify the edges that border these bigons. So, we will contract the bigons E and F and identify the orange and green edge with each other (with opposite orientation) and the pink and purple edge with each other (with opposite orientation). The resulting tetrahedra are shown in figure 7, where we have replaced all green edges by orange edges and pink edges by purple edges. If we identify the faces of these tetrahedra with rotation and identify the edges of the same colour in the resulting complex, we will recover the figure 8 knot complement. Note that we need not always obtain tetrahedra as a result of polyhedral decomposition; different knots may give rise to different polyhedra.

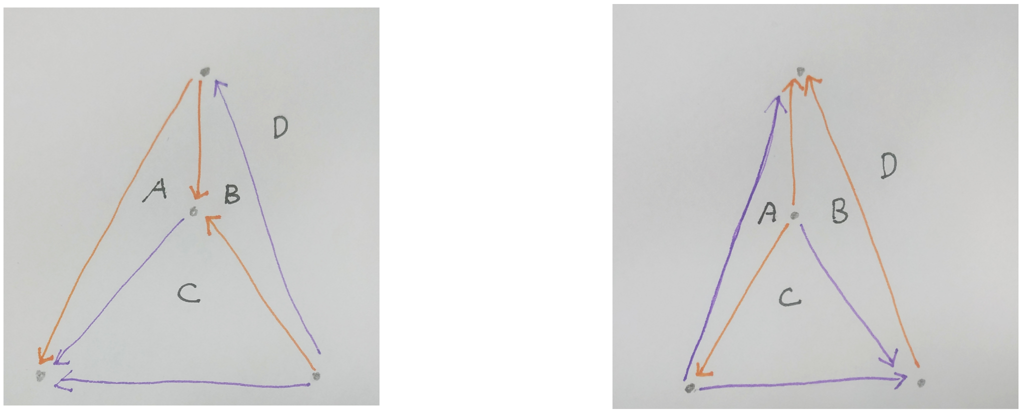

(a) The top tetrahedron

(b) The bottom tetrahedron

Figure 7: The top and bottom tetrahedra in the polyhedral decomposition of the figure 8 knot

2.4 Generalization of the algorithm to all alternating knots

All the features we noted for the polyhedral decomposition of the figure 8 knot work the same way for any alternating knot. We can construct a decomposition of the knot complement by using the alternating knot diagram. The projection graph of this diagram will be a 4-valent graph, so as in the case of the figure 8 knot, we will have two edges that are coming into every vertex, and two edges that go out of it. In the case of the figure eight knot, each edge corresponds to an arc of the knot starting at an underpass and ending at an overpass. This is only possible in an alternating knot diagram as the arcs of the knot will alternate between passing under and passing over from one crossing to the next. Thus, the polyhedra for the decomposition are again just obtained by labelling the surface of the 3-ball with the graph of the knot diagram. The faces in the top and bottom polyhedron will again be identified after gear rotation and the knot diagram will admit a checkerboard colouring; the same explanation works in this case.

For the case of non-alternating knots, some modification in this procedure is required. In the future, we plan to understand the modification required in this case.

3 Classification of isometries of the hyperbolic plane

To understand hyperbolic structures on 3-manifolds, it is essential to understand the isometries of hyperbolic space. In this section, we study the classification of isometries of $\mathbb{H}^2$. The presentation of material will be similar to that in the book 'Geometry of Surfaces' by John Stillwell. We have also referred to the book 'Fuchsian groups' by Svetlana Katok to understand the classification of isometries of $\mathbb{H}^2$ as hyperbolic, elliptic and parabolic isometries.

3.1 The upper half plane model and the disk model

The most frequently used model of the hyperbolic plane is the upper half plane, given by $\mathbb{H}^2 = \{z \in \mathbb{C}\colon z = x + iy, y > 0\}$. The metric on $\mathbb{H}^2$ is given by

$$ds^2 = \frac{dx^2 + dy^2}{y^2}$$

or in terms of the complex coordinate z

$$ds = \frac{|dz|}{Im(z)}$$

We describe some properties of the hyperbolic plane obvious from this metric. Fix a point in the hyperbolic plane, say i. Hyperbolic distance between this fixed point and any other point z becomes larger and tends to infinity as z becomes closer to the real axis or goes away to ∞ in the plane. The real axis along with the point at ∞ form the boundary of the hyperbolic plane, $\partial\mathbb{H}^2 = \{x + iy\colon y = 0\} \cup \{\infty\}$. Along any horizontal line $y = \alpha$, the hyperbolic length is just the Euclidean length divided by α. Some isometries of $\mathbb{H}^2$ which are obvious from this metric are

1. $t_\alpha(z) = z + \alpha$, where α is a real number, this corresponds to a Euclidean translation.
2. $d_\rho(z) = \rho z$, where ρ is a positive real number, this corresponds to a Euclidean dilation
3. $\bar{r}_{OY}(z) = -\bar{z}$, which is reflection in the Y axis.

However, these isometries have a different geometric meaning in the hyperbolic case than their apparent Euclidean geometric meaning and we will discuss this in more detail later. The hyperbolic plane also has isometries corresponding to rotation, but these are more easily described in terms of another model of the hyperbolic plane, the conformal disc model $\mathbb{D}^2$. The conformal disc $\mathbb{D}^2$ is just the unit disc in the complex plane. The metric on the conformal disc is given by

$$ds = \frac{|2d\omega|}{1 - |\omega^2|}$$

where ω is the complex coordinate in the unit disc. It is clear that the hyperbolic distance of a point from the centre of the disc becomes larger and larger and tends to infinity as the point approaches the boundary of the disc. The boundary of the hyperbolic plane in this model is given by $\partial\mathbb{D}^2 = \{\omega \in \mathbb{D}^2\colon |\omega| = 1\}$. In this model, it is clear that the boundary of the hyperbolic plane is homeomorphic to a circle.

The conformal disc model can be obtained from the upper half plane model by inversion in the circle centred at $-i$ with radius $\sqrt{2}$, followed by reflection in the X axis. As inversion and reflection preserve the magnitude of angles and invert their sign, their composition preserves the signed angle between two curves and hence $\mathbb{D}^2$ is called the conformal disc model. The transformation J from the upper half plane $\mathbb{H}^2$ to the conformal disc $\mathbb{D}^2$ is given by

$$\omega = J(z) = \frac{iz + 1}{z + i}$$

Isometries of $\mathbb{D}^2$ are then the J-conjugates of the isometries of $\mathbb{H}^2$. Some isometries that are natural in the conformal disc model of the hyperbolic plane are

1. $r_\theta(\omega) = e^{i\theta}\omega$, where θ is any real number. This corresponds to a hyperbolic rotation.
2. $\bar{r}(\omega) = \bar{\omega}$. This corresponds to an inversion about the unit circle in the upper half plane $\mathbb{H}^2$ as

$$J^{-1}\bar{r}J(z) = \frac{1}{\bar{z}}$$

As inversion in the unit circle is an isometry of $\mathbb{H}^2$, inversion in any circle with centre on the real axis is an isometry, since inversion in any circle is just inversion in the unit circle composed with the isometries t_α and d_ρ (which are Euclidean translations and dilations). We define $\mathbb{H}^2$-reflections to be inversions about any circle centred on the real axis along with Euclidean reflections about lines parallel to the the Y axis. Then the geodesics of the upper half plane will be the fixed point sets of the hyperbolic reflections. We know that the fixed point sets of Euclidean reflections about vertical lines are the same vertical lines and that of inversions about circles centred on the real axis will be the same circles. Thus, we infer that the geodesics of the upper half plane are vertical lines perpendicular to the real axis and semicircles centred on the real axis.

In the conformal disc model $\mathbb{D}^2$, the reflections are given by J-conjugates of $\mathbb{H}^2$ reflections. Thus, the $\mathbb{D}^2$-reflections are inversions in circles perpendicular to the boundary of the disc and reflections about diameters. And this means that the geodesics in $\mathbb{D}^2$ are circular arcs perpendicular to the boundary of the disc and diameters of the disc.

3.2 The three reflections theorem

In this section, we prove that the $\mathbb{H}^2$-reflections generate all the isometries of $\mathbb{H}^2$, in particular, we show that any $\mathbb{H}^2$-isometry is the product of at most three $\mathbb{H}^2$-reflections. This theorem is similar to the theorem that any Euclidean isometry can be realized as the product of at most three reflections, and the proof of the theorem for $\mathbb{H}^2$ also parallels the proof in the Euclidean case. The main upshot of this theorem is that we can easily prove that the isometries of $\mathbb{H}^2$ form a group, since any $\mathbb{H}^2$-reflection is invertible. The theorem will also help us to classify the isometries of $\mathbb{H}^2$.

To prove this theorem, we need to prove a few preliminary lemmas.

Lemma 3.1. *Any arbitrary segment PQ of a geodesic in $\mathbb{H}^2$ can be made to lie on the Y axis via isometries*

Proof. We first use an isometry of the form t_α to translate the point P on to the Y axis. We can then use a hyperbolic rotation about the point P to bring the point Q on to the Y axis; this isometry clearly fixes P. This hyperbolic rotation rotates points which are equidistant from the point P through a fixed angle; its fixed sets are circles whose hyperbolic center is P (hyperbolic circles are the same as Euclidean circles, but their centre differs from the Euclidean centre).

As these isometries map geodesics to geodesics and preserve angles, we infer that the image of PQ under the composition of these two isometries is a part of a geodesic perpendicular to the X axis such that the geodesic intersects the Y axis at the points P and Q. But this geodesic can only be the Y axis, hence the image of PQ lies on the Y axis. □

In Euclidean geometry, a line segment between two points is the shortest path between them. We can prove a similar statement for the hyperbolic plane. It is easy to see that given any two points in $\mathbb{H}^2$, there is a unique geodesic segment joining them. If the two points lie on a vertical line, then the vertical line segment is this geodesic segment. Otherwise, we can join the points P an Q by a straight line and construct its perpendicular bisector; the point of intersection of this bisector with the X axis is the centre of the geodesic semicircle and the arc joining P and Q along the geodesic semicircle is the required geodesic segment. Using this construction, we can easily observe that the geodesic segment joining P and Q must be unique.

Lemma 3.2. *The segment of the geodesic joining P and Q is the shortest path between them in $\mathbb{H}^2$*

Proof. By Lemma 3.1, we can use isometries to make PQ lie on the Y axis. Suppose C is any other path from P to Q. Each infinitesimal segment of the curve C corresponds to an infinitesimal segment of PQ (just project the segment to the Y axis) which has shorter length, since it has no displacement in the X direction. As a result, the length of C is given by

$$l(C) = \int_C \frac{\sqrt{dx^2 + dy^2}}{y} \geq \int_{PQ} \frac{dy}{y} = l(PQ)$$

Therefore, the geodesic segment PQ is the shortest path between the points P and Q. □

We now prove the triangle inequality for $\mathbb{H}^2$.

Theorem 3.3. *Suppose P, Q and R are any three points in $\mathbb{H}^2$. Then they satisfy the relation*

$$l(PQ) + l(QR) \geq l(PR)$$

and the equality holds if and only if all the points lie on the same geodesic with Q between P and R.

Proof. Using suitable isometries, we can make the segment PR lie on the Y axis. Now if the point Q also lies on the Y axis, the relation is clearly satisfied, with equality holding if Q lies in between P and R. Suppose now that Q does not lie on the Y axis. Suppose S is the projection of Q on the Y axis and suppose S lies in between P and R. Then we claim that

$$\mathbb{H}^2\text{-length of } PQ > \mathbb{H}^2\text{-length of } PS$$

$$\mathbb{H}^2\text{-length of } QR > \mathbb{H}^2\text{-length of } SR$$

If this claim is true, then we clearly have a strict inequality in the relation. To prove the claim, consider an infinitesimal segment of PQ and project it onto the Y axis; it will correspond to an infinitesimal segment of PS. Clearly, the length of the infinitesimal segment on PQ is greater than the length of the corresponding segment on PS, as there is no displacement in the X direction along PS, and the displacement in the Y direction is the same for both the segments. Thus, the length of PQ is the result of integrating the lengths of the infinitesimal segments along PQ, which will clearly be greater than the result of integrating the lengths of the infinitesimal segments along PS, which is the length of PS.

In the cases where S does not lie between P and R, Q must either lie above the segment PR or below it. In this case, the length of PQ itself will be greater than the length of PR by a similar argument as above. □

In Euclidean geometry, the set of points equidistant from two points P and P' is a line, and reflection in that line exchanges P and P'. This fact is crucial to the proof of the three reflections theorem. We now prove an analogous fact for $\mathbb{H}^2$.

Lemma 3.4. *The set of points equidistant from two points P and P' in $\mathbb{H}^2$ is a geodesic, and $\mathbb{H}^2$-reflection in that geodesic exchanges P and P'*

Proof. We will first use isometries to bring the points P and P' to positions where they are mirror images of each other in the Y-axis. We first rotate P' about P using a hyperbolic rotation about P until they both have the same Y coordinate. We then use a suitable isometry t_α to make P and P' equidistant from the Y axis.

We now claim that the Y axis is the set of all points equidistant from P and P'. Clearly, reflection in the Y axis exchanges P and P'. Suppose Q is equidistant from P and P' and Q does not lie on the Y axis; let Q' be the image of Q under reflection about the Y axis.

$$\begin{aligned}
\text{length of } P'Q' &= \text{length of } PQ \text{ (by reflection)} \\
&= \text{length of } P'Q \text{ (by assumption)} \\
&= \text{length of } P'R + \text{length of } RQ \text{ (as R lies on the geodesic segment between P and Q)} \\
&= \text{length of } P'R + \text{length of } RQ' \text{ (by reflection)}
\end{aligned}$$

This clearly contradicts the strict inequality which we get from the triangle inequality (Theorem 3.3), since R does not lie on the geodesic joining P' and Q'. Hence, the only points equidistant from P and P' are all the points on the X axis and reflection in the X axis exchanges P and P'. This proves the theorem. □

We now have all the statements necessary to prove the three reflections theorem.

Theorem 3.5 (Three reflections theorem). *Any $\mathbb{H}^2$-isometry is the product of at most three $\mathbb{H}^2$-reflections*

Proof. We will first show that any point in $\mathbb{H}^2$ is determined by its distance from any three non-collinear points (by collinear points, we mean points lying on the same geodesic). Suppose to the contrary that there exist two points P and P' such that their distances form three non-collinear points A, B and C are equal. Then Lemma 3.4 will imply that A, B and C are collinear, contrary to our initial assumption.

We now show that any isometry ϕ of $\mathbb{H}^2$ is determined by its effect on three non-collinear points. Suppose we know how ϕ acts on any three non-collinear points A, B and C in $\mathbb{H}^2$. Then, as ϕ is an isometry, it preserves lengths, and we can infer from the strict inequality in the triangle inequality that $\phi(A)$,$\phi(B)$ and $\phi(C)$ are also non-collinear. Now, the image of any other point P under ϕ is determined, as the distance of $\phi(P)$ from the non-collinear points $\phi(A)$,$\phi(B)$ and $\phi(C)$ is the same as the distance of P from A, B and C, which is given to us. But any point is uniquely determined by its distance from three non-collinear points, so we know how ϕ acts on P given its action on any other three non-collinear points.

We will now show that any isometry ϕ of $\mathbb{H}^2$ is the product of at most three reflections. Let A, B and C be any three non-collinear points in $\mathbb{H}^2$. There are four cases:

Case 1: Suppose ϕ fixes all three points. Then ϕ must be the identity isometry as for any point P, $\phi(P)$ must have the same distance from A, B and C as P does, which means that $\phi(P) = P$. The identity is the product of zero reflections.

Case 2: Suppose ϕ fixes A and B. Then A and B are equidistant from C and $\phi(C)$, so a reflection ρ in the geodesic passing through A and B will exchange C with $\phi(C)$, while leaving A and B invariant. Now $\rho \circ \phi$ fixes all points, and hence it is the identity, which means that $\phi = \rho^{-1} = \rho$, so ϕ is the product of one reflection (since inverse of a reflection is the same reflection). The other cases also follow the same idea.

Case 3: Suppose ϕ fixes only A. A reflection τ in the geodesic equidistant from B and $\phi(B)$ which passes through A exchanges B and $\phi(B)$. Now $\tau \circ \phi$ fixes both A and B, so from case 2, we know that $\tau \circ \phi = \rho$, where ρ is a reflection. Hence $\phi = \tau \circ \rho$, so ϕ is the product of two reflections.

Case 4: Suppose ϕ does not fix any of A, B or C. Then a reflection η in the geodesic equidistant from A and $\phi(A)$ will exchange A and $\phi(A)$, so $\eta \circ \phi$ will fix A. This means we are back in case 3, so $\eta \circ \phi = \tau \circ \rho$ which implies that $\phi = \eta \circ \tau \circ \rho$, so ϕ is the product of three reflections

We have exhausted all possible cases, and hence every isometry of $\mathbb{H}^2$ is the product of at most three reflections. ☐

Clearly, the composition of any two isometries is an isometry. Now, since isometries are generated by reflections, every isometry is invertible too (as reflections are self inverse). Hence, we have

Corollary 3.5.1. *The isometries of $\mathbb{H}^2$ form a group which is generated by the $\mathbb{H}^2$-reflections. We denote this group by Isom($\mathbb{H}^2$).*

3.3 Isometries are Mobius transformations

We have seen that every isometry of $\mathbb{H}^2$ is the product of at most three reflections. We can classify these isometries as orientation preserving (product of an even number of reflections) and orientation reversing (product of an odd number of reflections). These sets are well defined as any reflection reverses the sign of angles, hence the product of an odd number of reflections can never equal the product of an even number of reflections. We can see that the orientation preserving isometries form a subgroup of Isom($\mathbb{H}^2$), which we denote by $\mathrm{Isom}^+(\mathbb{H}^2)$. Thus the group of isometries is partitioned into two cosets, $\mathrm{Isom}^+(\mathbb{H}^2)$ (the orientation preserving isometries) and $\bar{r}_{OY}\mathrm{Isom}^+(\mathbb{H}^2)$ (the orientation reversing isometries), where $\bar{r}_{OY}$ is reflection in the Y axis.

As $\mathbb{H}^2$-reflections are just inversions in circles centred on the X axis and reflections in vertical lines, we know how to express them as complex functions. We will use this to find a canonical form for all isometries of $\mathbb{H}^2$.

Theorem 3.6. 1. *All orientation preserving isometries of $\mathbb{H}^2$ are of the form*

$$f(z) = \frac{az+b}{cz+d} \text{ where } a, b, c, d \in \mathbb{R}, \text{ such that } ad - bc = 1$$

This can be expressed as

$$f(z) = \begin{bmatrix} a & b \\ c & d \end{bmatrix} \begin{bmatrix} z \\ 1 \end{bmatrix} \text{ where } det \begin{bmatrix} a & b \\ c & d \end{bmatrix} = 1$$

Thus, orientation preserving isometries of $\mathbb{H}^2$ are given by Mobius transformations with real coefficients and they are isomorphic to $PSL(2, \mathbb{R})$.

2. *All orientation reversing isometries are of the form*

$$f(z) = \frac{a\bar{z}+b}{c\bar{z}+d} \text{ where } a, b, c, d \in \mathbb{R}, \text{ such that } ad - bc = -1$$

which can again be expressed as

$$f(z) = \overline{\begin{bmatrix} a & b \\ c & d \end{bmatrix} \begin{bmatrix} z \\ 1 \end{bmatrix}} \text{ where } det \begin{bmatrix} a & b \\ c & d \end{bmatrix} = -1$$

Proof. Let f be any isometry of $\mathbb{H}^2$. Suppose f is the product of only one reflection. Then f must either be reflection in a vertical line, in which case it can be composed with a suitable t_α to give $\bar{r}_{OY}$. This means that f is of the form $f(z) = -\bar{z} + \alpha$, which clearly satisfies the statement of the theorem. Suppose f is inversion in a circle centred at the point k on the real axis of radius r, we can compose it with suitable Euclidean dilations and translations to make it inversion in the unit circle centred at the origin ($z \mapsto 1/\bar{z}$). Hence, the formula for the inversion is

$$f(z) = \frac{k\bar{z} + r^2 - k^2}{\bar{z} - k}$$

The determinant of this transformation is $-r^2$, so if we multiply both numerator and denominator by $1/r$, the determinant will be -1.

Now, if f is the product of two $\mathbb{H}^2$-reflections, then the transformation will act as

$$z \mapsto \begin{bmatrix} a & b \\ c & d \end{bmatrix} \begin{bmatrix} \bar{z} \\ 1 \end{bmatrix} \mapsto \begin{bmatrix} a' & b' \\ c' & d' \end{bmatrix} \times \overline{\begin{bmatrix} a & b \\ c & d \end{bmatrix} \begin{bmatrix} \bar{z} \\ 1 \end{bmatrix}} = \begin{bmatrix} a' & b' \\ c' & d' \end{bmatrix} \begin{bmatrix} a & b \\ c & d \end{bmatrix} \begin{bmatrix} z \\ 1 \end{bmatrix}, \text{ where } a, b, c, d \in \mathbb{R}$$

This is clearly of the form in the theorem, as the product of two matrices with determinant -1 has determinant 1.

Now suppose f is the product of three reflections, then $f\bar{r}_{OY}$ will be an orientation preserving isometry, which we have shown to be of the form

$$f\bar{r}_{OY}(z) = \frac{az+b}{cz+d}$$

then as

$$r_{OY}(z) = -\bar{z}$$

$$f(z) = \frac{-a\bar{z} + b}{-c\bar{z} + d}$$

which is clearly of the form in the theorem, as the determinant of the transformation f is -1, since $ad-bc = 1$. As all isometries of $\mathbb{H}^2$ can be expressed as the product of at most three reflections, we have shown that all isometries are of the form stated in the theorem.

Now, any function f which is of one of the two kinds given in the theorem is an isometry. This is because the transformed metric ds' is given by

$$\begin{aligned} ds' &= \left| \frac{d(\frac{az+b}{cz+d})}{Im(\frac{az+b}{cz+d})} \right| \\ &= \frac{\frac{ad-bc}{|cz+d|^2}|dz|}{\frac{(ad-bc)(z-\bar{z})}{2i|cz+d|^2}} \text{ (after simplification)} \\ &= \frac{|dz|}{Im(z)} \\ &= ds \end{aligned}$$

A similar calculation will show that any function of the second kind is also an isometry. We will now show that any function of the first kind is an orientation preserving isometry, that is, it is the product of an even number of $\mathbb{H}^2$-reflections. This will also show that any function f of the second kind is an orientation reversing isometry, since $f\bar{r}_{OY}$ will again be of the first kind, which is an orientation preserving isometry. To prove the functions of the first kind are orientation preserving, we attempt to write them in terms of familiar isometries. Let us first assume that $c \neq 0$

$$\begin{aligned} f(z) &= \frac{az + b + ad/c - ad/c}{cz + d} \\ &= \frac{a}{c} + \frac{b - ad/c}{cz + d} \\ &= \frac{a}{c} + \frac{bc - ad}{c(cz + d)} \\ &= \frac{a}{c} - \frac{1}{c(cz + d)} \end{aligned}$$

If $c > 0$, then f is the composition of dilations like $z \mapsto cz$, translations like $z \mapsto z + d$, $z \mapsto z + a/c$, and the two $\mathbb{H}^2$- reflections $z \mapsto 1/\bar{z}$, $z \mapsto -\bar{z}$ (we are calling these functions translations and dilations only because they resemble Euclidean translations and dilations, they have a different interpretation in hyperbolic geometry). Since translations t_α and dilations d_ρ are both products of two $\mathbb{H}^2$- reflections, the function f itself is a product of an even number of reflections, and hence is an orientation preserving isometry. If $c < 0$, we start out by writing f as $f(z) = \frac{-az-b}{-cz-d}$, and carry out the same procedure. And if $c = 0$, $f(z) = az/d + b/d$, which is the product $t_{b/d}d_{a/d}$ (as $ad = 1$, $a/d > 0$).

We have shown that every orientation preserving isometry (or orientation reversing isometry) of $\mathbb{H}^2$ is of the form stated in the theorem, and that any function of the kinds given in the theorem are orientation preserving isometries (or orientation reversing isometries). We see that the orientation preserving isometries correspond to matrices of $SL(2, \mathbb{R})$, and since both signs of the matrix give the same isometry, we actually get that $\text{Isom}^+(\mathbb{H}^2)$ is isomorphic to $PSL(2, \mathbb{R})$. □

3.4 Classification of isometries - a geometric approach

In this section, we use the tools developed so far to give a classification of isometries of the hyperbolic plane which parallels the classification in the Euclidean case. We will use the three reflections theorem crucially in this classification. Before we prove the theorem, we will prove some useful lemmas about the fixed points of orientation reversing isometries which will aid us in our proof.

Lemma 3.7. *Every orientation reversing isometry has two fixed points on the boundary of the hyperbolic plane ($\partial\mathbb{H}^2$).*

Proof. We have proved in theorem 3.6 that every orientation reversing isometry has the form

$$f(z) = \frac{a\bar{z}+b}{c\bar{z}+d} \text{ where } a,b,c,d \in \mathbb{R}, \text{ such that } ad-bc=-1$$

So, to find out the fixed points of f on the boundary of $\mathbb{H}^2$, which is $\mathbb{R}\cup\{\infty\}$, we need to solve the equation

$$x = \frac{ax+b}{cx+d}$$

since for points on the boundary, we have $\bar{x} = x$. If $c = 0$, then we get two solutions ∞ and $b/(a-d)$. If $c \neq 0$, we will get a quadratic equation $cx^2 + (d-a)x - b = 0$ which has solutions given by

$$\begin{aligned} x &= \frac{-(d-a) \pm \sqrt{(d-a)^2+4bc}}{2a} \\ &= \frac{-(d-a) \pm \sqrt{(a+d)^2+4}}{2a} \text{ simplified using } ad-bc=-1 \end{aligned}$$

Since the discriminant is clearly positive, we have two distinct solutions to the equation and hence two fixed points on $\partial\mathbb{H}^2$. □

Theorem 3.8 (Classification of isometries)**.** *Any isometry of the hyperbolic plane has one of the following forms*

1. *A rotation r_θ*
2. *A limit rotation t_α*
3. *A translation d_ρ*
4. *A glide reflection (product of an $\mathbb{H}^2$-reflection and a translation)*

Proof. We first classify the orientation preserving isometries of $\mathbb{H}^2$. As any orientation preserving isometry is the product of two $\mathbb{H}^2$-reflections in two distinct lines (except the identity) we will need to understand how pairs of geodesic lines in $\mathbb{H}^2$ behave. For any pair of lines L and M in $\mathbb{H}^2$, only the following three cases are possible:

Case 1: The lines L and M intersect in $\mathbb{H}^2$.

Case 2: The lines L and M do not intersect in $\mathbb{H}^2$, but they intersect on the boundary of $\mathbb{H}^2$, in which case they are called asymptotic lines.

Case 3: The lines L and M neither intersect in $\mathbb{H}^2$ nor on the boundary of $\mathbb{H}^2$, in which case they are called ultra-parallel lines.

We now elaborate how the isometry looks like in each of these cases.

Case 1: In this case, we can choose coordinates suitably to ensure that the two lines intersect at the origin of $\mathbb{D}^2$. Then the product of $\mathbb{H}^2$-reflections in these two lines will correspond to a rotation r_θ of $\mathbb{D}^2$. Such a rotation clearly has only one fixed point in $\mathbb{D}^2$ and no fixed point on the boundary of $\mathbb{D}^2$. It permutes the diameters of $\mathbb{D}^2$ and leaves any circle in $\mathbb{D}^2$ with the origin as its centre invariant.

Case 2: In this case, we can choose suitable coordinates to make the two lines intersect at ∞ in $\partial\mathbb{H}^2$. Then L and M are just vertical lines, and the product of reflections in L and M is a limit rotation t_α. A limit rotation fixes no points in $\mathbb{H}^2$ and fixes only the point at infinity on $\partial\mathbb{H}^2$. It permutes the vertical geodesics while keeping invariant the horocycles y =constant. We can understand why this isometry is termed limit rotation by looking at it in the $\mathbb{D}^2$ picture. The isometry then fixes one point on the boundary of the disk and it appears as if the other points are rotating about this 'limit point'.

Case 3: In this case, we can choose coordinates such that L is the Y axis in $\mathbb{H}^2$ and M is some semicircle centred on the X axis disjoint from L. we see that we can construct a geodesic N that is perpendicular to both L and M, so by transforming coordinates such that N is now the Y axis, we will get L and M to be concentric semicircles centred at the origin. The product of reflections in L and M will be of the form d_ρ, which is what we term as a hyperbolic translation. A translation does not fix any point in $\mathbb{H}^2$, but fixes two points on the boundary of $\mathbb{H}^2$ (0 and ∞ in this case) and leaves the Y axis invariant. It leaves all lines $y = mx$ passing through the origin invariant, and permutes geodesic semicircles centres at the origin.

We have thus classified all orientation preserving isometries. Now if f is an orientation reversing isometry, then it has two fixed points P and Q on the boundary of $\mathbb{H}^2$, by lemma 3.8. Let $\bar{r}_{PQ}$ be the reflection in the geodesic line PQ. Then $f\bar{r}_{PQ}$ is orientation preserving and it also fixes the points P and Q, since the reflection fixes P and Q. The only orientation preserving isometry of $\mathbb{H}^2$ which fixes two points is a translation d_ρ. Therefore, $f = d_\rho\bar{r}_{PQ}$, so f is the product of a hyperbolic translation and a reflection. Therefore, f has to be a glide reflection with the axis PQ. It fixes no points in $\mathbb{H}^2$, fixes the points P and Q on $\partial\mathbb{H}^2$, and leaves invariant the geodesic line PQ. $\square$

3.5 Classification of isometries of $\mathbb{H}^2$ - an algebraic approach

We have proved that the orientation preserving isometries of $\mathbb{H}^2$ are isomorphic to $PSL(2,\mathbb{R})$. In this section, we will use a more algebraic approach using matrices of $PSL(2,\mathbb{R})$ to classify all orientation preserving isometries as elliptic, parabolic and hyperbolic isometries. Any matrix of $PSL(2,\mathbb{R})$ is, up to a sign, of the form

$$A = \begin{bmatrix} a & b \\ c & d \end{bmatrix} \text{ where } det(A) = 1$$

We can use the trace of this matrix to classify the isometries. If the eigenvalues of A are λ and μ, we know that $\lambda\mu = 1$, that is $\mu = 1/\lambda$, as the product of the eigenvalues is the determinant and $det(A) = 1$. The trace of A is the sum of its eigenvalues, that is $Tr(A) = \lambda + 1/\lambda$. As we are working in $PSL(2,\mathbb{R})$, we are only concerned with the absolute value of the trace, and there are only three possibilities for it:

Case 1: $|Tr(A)| < 2$. These are called elliptic isometries.

Case 2: $|Tr(A)| = 2$. In this case, we call the isometry parabolic.

Case 3: $|Tr(A)| > 2$. Such isometries are called hyperbolic isometries.

We elaborate on the features of each of these isometries below and try to find their fixed points in $\mathbb{H}^2$ and on $\partial\mathbb{H}^2$. To find the fixed points, we need to solve the quadratic equation

$$\frac{az+b}{cz+d} = z$$

If $c = 0$, then $az + b = dz$ and we get two solutions for the equation, ∞ and $b/(a-d)$ (we are assuming that $a \neq d$, otherwise A will just be the identity matrix). In this case, as a and d are real numbers satisfying $ad = 1$, $|Tr(A)| = |a + d| > 1$, that is, the isometry is hyperbolic with two fixed points, 0 and ∞ on $\partial\mathbb{H}^2$. If $c \neq 0$, we will get a quadratic equation $cz^2 + (d-a)z - b = 0$ which has solutions given by

$$\begin{aligned} z &= \frac{-(d-a) \pm \sqrt{(d-a)^2 + 4bc}}{2a} \\ &= \frac{-(d-a) \pm \sqrt{(a+d)^2 - 4}}{2a} \text{ simplified using } ad - bc = 1 \end{aligned}$$

The discriminant of the equation is clearly $Tr(A)^2 - 4$, so the fixed points of this isometry will depend on the sign of the discriminant. The eigenvalues of the matrix A can be found by finding the roots of the characteristic polynomial of A, which is $x^2 - (a + d)x + 1$. This equation also has discriminant $Tr(A)^2 - 4$, and the eigenvalues of A and the canonical form of A depend on the sign of the discriminant.

Case 1: Elliptic isometries If $|Tr(A)| < 2$, then the discriminant of the equation is negative, and we get two fixed points on the complex plane which are complex conjugates of each other, and hence there is only one fixed point on the upper half plane $\mathbb{H}^2$. Also, since the discriminant is negative, we get that the eigenvalues are distinct and complex conjugates of each other. In this case, by using some linear algebra and the relations between the eigenvalues, we can determine that A must be conjugate in $SL(2, \mathbb{R})$ to a matrix of the form

$$A = \begin{bmatrix} cos\theta & sin\theta \\ -sin\theta & cos\theta \end{bmatrix} \text{ where } \theta \in \mathbb{R}$$

Thus, the invariant curves (in $\mathbb{R}^2$) of an elliptic isometry are ellipses, hence the terminology. We know that the only orientation preserving isometries of $\mathbb{H}^2$ which have one fixed point in $\mathbb{H}^2$ are the rotations r_θ, so elliptic isometries correspond to rotations in the previous classification.

Case 2: Parabolic isometries If $|Tr(A)| = 2$, then the discriminant is 0, so we get only one real solution to the equation. This means that there are no fixed points in $\mathbb{H}^2$, and one fixed point on the boundary of $\mathbb{H}^2$, which is $\mathbb{R} \cup \{\infty\}$. We also infer that A has only one eigenvalue, which is 1 (as $(a + d)/2 = 1$). Again, we can use basic linear algebra to deduce that A must be conjugate in $SL(2, \mathbb{R})$ to a matrix of the form

$$A = \begin{bmatrix} 1 & \lambda \\ 0 & 1 \end{bmatrix} \text{ where } \lambda \in \mathbb{R}$$

From the previous classification, we know that the only orientation preserving isometries of $\mathbb{H}^2$ which fix no point in $\mathbb{H}^2$ and only one point on the boundary of $\mathbb{H}^2$ are limit rotations. Thus parabolic isometries correspond to the limit rotations t_α of $\mathbb{H}^2$.

Case 3: Hyperbolic isometries Finally, if $|Tr(A)| > 2$, then the discriminant of the equation is positive and it will have two distinct real solutions. Again, this means that no points in $\mathbb{H}^2$ are fixed, but two points on the boundary $\partial\mathbb{H}^2$ are fixed. We can see that A will now have two real and distinct eigenvalues. By using some linear algebra, we deduce that A must be conjugate in $SL(2, \mathbb{R})$ to a matrix of the form

$$A = \begin{bmatrix} \lambda & 0 \\ 0 & 1/\lambda \end{bmatrix} \text{ where } \lambda \in \mathbb{R}, \lambda \neq 1$$

The invariant curves in the plane for such an isometry are hyperbolas, hence the nomenclature. We know that the only orientation preserving isometries of $\mathbb{H}^2$ which fix no points in $\mathbb{H}^2$, but two points on its boundary are the translations d_ρ. Hence, hyperbolic isometries correspond to translations from the previous classification.

We have thus described an algebraic method of classifying orientation preserving isometries and correlated it with the previous classification. We know that orientation reversing isometries are just given by composing orientation preserving isometries with the reflection $z \mapsto -\bar{z}$, so we can use the description of $\mathrm{Isom}^+(\mathbb{H}^2)$ to understand their features.

4 Geometric structures on manifolds

We wish to study knots by studying the geometric properties of the knot complement. To this end, we have described an algorithm for the polyhedral decomposition of alternating knots in the first section and classified the isometries of the hyperbolic plane in the second section. We will now try to understand the general theory behind how we can endow manifolds with a geometric structure and what conditions will ensure that the geometric structure on the manifold is complete (every Cauchy sequence on the manifold converges).

A topological polyhedral decomposition of a manifold M is the data of the polyhedra that make up M and the information of how to glue them together to obtain M. A topological polyhedral decomposition is said to be a geometric polyhedral decomposition if every polyhedron has a metric on it such that the gluing is by isometries and the resulting manifold M is complete. We want to understand when the knot complement possesses a geometric polyhedral decomposition, and the theory developed in this section will aid us in this goal. Most of the material in this section is drawn from the books by Purcell and Thurston.

4.1 Definition of geometric structure and examples

We first discuss a general way to describe geometric structures on manifolds.

Definition 4.1. *Consider a manifold X and let G be a group acting on it. A manifold M is said to have a (G, X) structure if for every point p of M, there is a chart (U, ϕ) around it that maps to X, where U is an open neighbourhood of p and ϕ is a homeomorphism from U onto its image $\phi(U)$. The charts must satisfy the following condition - whenever any two charts (U, ϕ) and (V, ψ) intersect, the transition map $\psi \circ \phi^{-1} \colon \phi(U \cap V) \to \psi(U \cap V)$ should be given by the restriction of some element of G on every connected component of $\phi(U \cap V)$.*

In what follows, we will mostly consider X and M to be real analytic manifolds and G to be a group of real analytic diffeomorphisms acting on X. We will also require that the charts from M to X be real analytic diffeomorphisms. We know that real analytic diffeomorphisms are uniquely determined by their restriction to any connected open set. Hence, the transition map on every connected component of $\phi(U \cap V)$ will be the restriction of some unique element of G. We will use this fact crucially later while defining the developing map.

We describe two different geometric structures on the torus to understand this definition better. For both examples, the manifold M under consideration is the torus and X is $\mathbb{R}^2$. For the first example, we let G be the group of isometries of Euclidean space, denoted by $\mathrm{Isom}(\mathbb{E}^2)$. We know that the universal cover of the torus is $\mathbb{R}^2$ and that it can be tiled by squares of side length 1. The square which has the origin as one of its vertices will be called the fundamental domain of the torus. We use this knowledge to give charts on the torus. Any point p on the torus lifts under the covering map to infinitely many points in $\mathbb{R}^2$, which form an integer lattice in $\mathbb{R}^2$. Consider an open disc around any lift of p in $\mathbb{R}^2$, where the radius of the disc is less than half. This disc will project to a neighbourhood U of p in the torus and we map this neighbourhood U to any such disc to give a chart around p. We can then clearly see (figure 8) that the transition maps are all given by integer translations which are clearly Euclidean isometries. We deduce that we can give an

(Isom($\mathbb{E}^2$), $\mathbb{E}^2$) structure on the torus obtained by gluing opposite sides of a square.This structure on the torus is called a Euclidean structure.

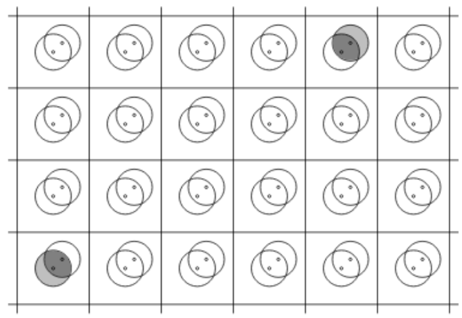

Figure 8: Euclidean structure on a torus - Image from pg. 31, Purcell, J. (2019). Hyperbolic Geometry and Knot Theory. [1]

We give another geometric structure on the torus, called an affine structure. For this, we can take the tiling of the plane by any quadrilateral of different side lengths. We take G to be the group of affine transformations on $\mathbb{R}^2$, that is transformations of the form $x \mapsto Ax + b$, where A is any invertible linear transformation. This will allow transition maps to not just include isometries like rotation and translation, but also scaling. Thus, the torus formed by identifying opposite sides of any quadrilateral by affine maps can be given a geometric structure which we call an affine structure. The affine structure on a torus is depicted in figure 9.

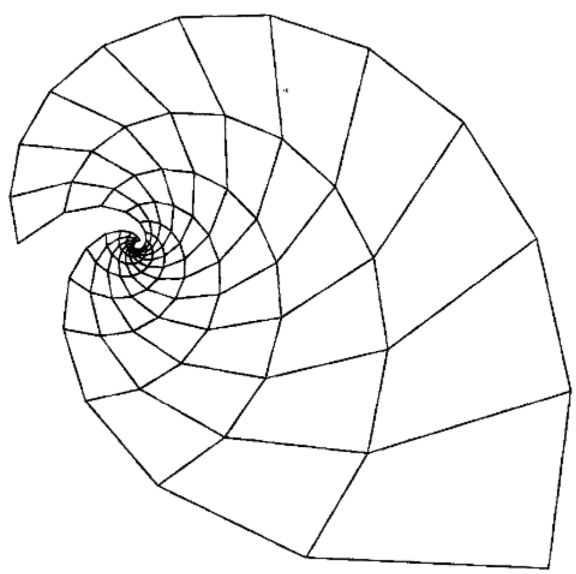

Figure 9: Developing an affine torus - Image from pg. 142, Thurston, Three dimensional geometry and topology (1997). [2]

As we can see, it is cumbersome to describe geometric structures on manifolds in terms of charts and much easier to describe them in terms of quotient spaces. We will mainly concern ourselves with describing geometric structures, especially hyperbolic structures (which are defined as (Isom($\mathbb{H}^n$), $\mathbb{H}^n$) structures), on the gluing of polyhedra. These hyperbolic structures will be very useful in describing and studying knot complements.

4.2 Hyperbolic structure on gluing of polygons

We will try to understand gluings of hyperbolic polygons - these are polygons in $\mathbb{H}^2$ with geodesic edges. We know that if polygonal gluings satisfy certain 'nice' properties - if the edges are identified in pairs and vertices go to vertices and edges go to edges - the quotient space will be a manifold (Massey describes this in more detail in his book 'Algebraic topology - an introduction'). In what follows, we assume that all gluings satisfy these nice properties and that the quotient space is a manifold. But when does a gluing of hyperbolic polygons admit a hyperbolic structure? The following lemma is an answer to this question. We will assume that the gluing is via isometries and in that case the quotient manifold M will inherit a metric from the hyperbolic polygons.

Lemma 4.1. *A gluing of hyperbolic polygons admits a hyperbolic structure if every point in the gluing has a neighbourhood isometric to a disk in $\mathbb{H}^2$. More generally, a gluing of hyperbolic polyhedra admits a hyperbolic structure if every point in the gluing has a neighbourhood isometric to a ball in $\mathbb{H}^n$.*

Proof. Suppose every point in the quotient manifold M has a neighbourhood U isometric to a ball in $\mathbb{H}^n$ under the isometry ϕ. Then, we can take (U, ϕ) to be the geometric chart neighbourhood around that point. Whenever two charts (U, ϕ) and (V, ψ) overlap, the transition map $\psi \circ \phi^{-1}$ will be a composition of isometries, hence it will be an isometry of $\mathbb{H}^n$. Thus the manifold M will get a hyperbolic structure. □

To apply this lemma, we need to understand what conditions are required for every point in a gluing of hyperbolic polygons to have a neighbourhood isometric to a disk in $\mathbb{H}^2$. Any point in the interior of one of the polygons clearly has a disk neighbourhood which does not meet the boundary of the polygon; this will give an isometry to a disk in $\mathbb{H}^2$. We claim that any point in the interior of an edge also has a neighbourhood isometric to a disk in $\mathbb{H}^2$. The edge is shared by exactly two polygons in the gluing, so we can apply isometries to each one of them so that that edge of the polygon is on the Y axis in $\mathbb{H}^2$ and by a suitable scaling (a hyperbolic translation d_ρ) we can make sure that the polygons meet up perfectly along the edge. Every point on the edge has two half disk neighbourhoods from each polygon and they will glue together perfectly to give a disk neighbourhood in $\mathbb{H}^2$ when placed in such a standard position. A problem arises only at the vertices; they need not always have such a disk neighbourhood. The following lemma tells us when such a disk neighbourhood will exist for a vertex.

Lemma 4.2. *A gluing of hyperbolic polygons has a neighbourhood isometric to a disk in $\mathbb{H}^2$ if and only if at every vertex of the gluing, the interior angles of the polygons sum to 2π.*

Proof. If every point of the gluing has a neighbourhood isometric to a disk in $\mathbb{H}^2$, then in particular, any vertex has such a disk neighbourhood. Since isometries preserve angles, we will get that the sum of interior angles around any vertex is 2π.

Suppose the gluing is such that the sum of interior angles around a vertex is 2π. We know from the discussion preceding the lemma that any point in the interior of one of the polygons or in the interior of one of the edges will always have a neighbourhood isometric to a disk. We can now again use isometries to make the vertex under consideration lie on the Y axis of $\mathbb{H}^2$ at a fixed point, say i. We will place the polygons one after the other in a cycle, gluing them at the respective edges till we reach the last polygon. The edge of the last polygon will then glue to the edge of the first one since the angle sum around the vertex is 2π. Thus the vertex will have a neighbourhood isometric to a disk in $\mathbb{H}^2$. □

4.3 The developing map and holonomy

We have seen the conditions required for a gluing of hyperbolic polyhedra to admit a hyperbolic structure. We will now try to understand the conditions on the polyhedra which will ensure that this hyperbolic structure is complete with respect to the quotient metric inherited from the polyhedra. For this, we will need to understand the concepts of developing map and holonomy.

Let M be a real analytic manifold and suppose it has a (G, X) structure, where X is also a real analytic manifold and G is a group of real analytic diffeomorphisms acting on X. Every point p of M has a chart (U, ϕ) which maps a neighbourhood U of p homeomorphically onto a neighbourhood $\phi(U)$ in X. We know that if two charts (U, ϕ) and (V, ψ) intersect, then the transition map $\psi \circ \phi^{-1}$ is given uniquely by the restriction of an element $g(y)$ of G in the connected component of y, where y is any point in the intersection. We can use this to extend the chart (U, ϕ) further in X by defining

$$\begin{aligned}\Phi(x) &= \phi(x) \text{ for } x \in U \\ &= g(y)\psi(x) \text{ for } x \in V\end{aligned}$$

Then we can see that the map is well defined on the intersection when $U \cap V$ is connected, but it may not be well defined if the intersection is not connected. The idea behind the developing map is to use the global extensions of the transition maps (which are given by elements of G) and to extend the information of the charts to define a map from all of M to X. However, there are difficulties with this approach. The intersection of two chart neighbourhoods need not be connected; also if the manifold M has any non-trivial loops (if M is not simply connected), we will face a problem when we develop charts along this loop, as when we come back to the starting point, there will be two different definitions of the map there. To solve these problems and define a map consistently, we will use the universal cover $\widetilde{M}$ of M.

To define the developing map, we use the construction of the universal cover as the space of homotopy classes of paths starting at a fixed base point in M. To define the map at any point of $\widetilde{M}$, we will develop the chart along the path corresponding to that point.

Let $\alpha \colon [0, 1] \to M$ be a path in M based at x_0 and suppose the path corresponds to the point $[\alpha]$ in $\widetilde{M}$. We will develop the chart information along this path. As the image of the path is compact, we can find finitely many chart neighbourhoods (U_i, ϕ_i) which cover the path and sub-intervals $[t_i, t_{i+1}]$ (where $0 = t_0 < t_1 < ... < t_n = 1$) of $[0, 1]$ such that $\alpha([t_i, t_{i+1}])$ lies completely in U_i for every i. The points $\alpha(t_i)$ belong to the intersection of the two charts U_{i-1} and U_i; let us denote these points as x_i. We know that the transition map $\phi_{i-1} \circ \phi_i^{-1}$ will be determined by a unique element in the neighbourhood of the point x_i; let this element be denoted by g_i. We develop the chart (U_0, ϕ_0) along the path as follows:

The chart (U_0, ϕ_0) gives us a map from $[0, t_1]$ to X, defined by $\phi_0(\alpha(t))$. We can extend this to a map Φ_1 from $[0, t_2]$ to X by defining

$$\begin{aligned}\Phi_1(t) &= \phi_0(\alpha(t)) && \text{for } t \in [0, t_1] \\ &= g_1\phi_1(\alpha(t)) && \text{for } t \in [t_1, t_2]\end{aligned}$$

Clearly, this map is well defined at the point t_1 as $g_1 phi_1(\alpha(t_1)) = \phi_0(\alpha(t_1))$, since g_1 corresponds to the transition map $\phi_0 \circ \phi_1^{-1}$ in the connected component of x_1.

In this way, we can develop the chart along the path inductively by defining

$$\begin{aligned}\Phi_i(t) &= \Phi_{i-1}(\alpha(t)) && \text{for } t \in [0, t_i] \\ &= g_1 g_2 ... g_{i-1} g_i \phi_i(\alpha(t)) && \text{for } t \in [t_i, t_{i+1}]\end{aligned}$$

Suppose all maps Φ_k are well defined for k from 1 to i. Then the map Φ_{i+1} is also well defined, as at the point t_{i+1}

$$\begin{aligned}\Phi_{i+1}(t_{i+1}) &= g_1 g_2 ... g_{i-1} g_i g_{i+1} \phi_{i+1}(\alpha(t_{i+1})) \\ &= g_1 g_2 ... g_{i-1} g_i \phi_i(\alpha(t_{i+1})) && \text{as } g_{i+1} = \phi_i \circ \phi_{i+1}^{-1} \\ &= \Phi_i(t_{i+1}) && \text{by inductive hypothesis}\end{aligned}$$

Thus the map Φ_{i+1} is also well defined and hence by induction we have a well defined map $\Phi_{n-1} \colon [0,1] \to X$. Then we define the developing map $D \colon \widetilde{M} \to X$ as

$$D([\alpha]) = \Phi_{n-1}(1) = g_1 g_2 ... g_{n-1} \phi_{n-1}(\alpha(1))$$

This process of continuing a chart along a path is known as analytic continuation as we are crucially using the fact that elements of G are real analytic diffeomorphisms. We have made many arbitrary choices while defining the developing map, so we need to verify that the map is well defined and is independent of these choices.

Theorem 4.3. *1. For a fixed basepoint x_0, the definition of the developing map is independent of the path chosen in the path homotopy class $[\alpha]$, the charts (U_i, ϕ_i) chosen along the path and the points x_i chosen in the intersection of the charts.*

2. $D \colon \widetilde{M} \to X$ is a local diffeomorphism

3. If we choose a different basepoint and define another developing map D' in the same way as D, the resulting map D' will differ from D by composition with a unique element of G.

Proof. 1. Let us first show that the definition is independent of the charts chosen along the path. We first consider the simple case where another chart neighbourhood (V, ψ) is inserted between (U_{i-1}, ϕ_{i-1}) and (U_i, ϕ_i). In this case, let the transition maps $\phi_{i-1} \circ \psi^{-1}$ and $\psi \circ \phi_i^{-1}$ correspond to the group elements h and k respectively. Then as the point x_i belongs to all the three chart neighbourhoods, we have

$$g_i = \phi_{i-1} \circ \phi_i = \phi_{i-1} \circ \psi^{-1} \circ \psi \circ \phi_i^{-1} = hk$$

Thus, the new definition of the developing map is

$$\begin{aligned}D'([\alpha]) &= g_1 ... g_{i-1} hk g_{i+1} ... g_{n-1} \phi_{n-1}(\alpha(1)) \\ &= g_1 ... g_{i-1} g_i g_{i+1} ... g_{n-1} \phi_{n-1}(\alpha(1)) \\ &= D([\alpha])\end{aligned}$$

If we have two different sets of charts covering the path, we can refine both of them to a common set of charts (which is the union of both sets), and the previous calculation shows that the developing map for the refined set of charts is the same as the developing map for each individual set of charts.

We now show that the developing map is independent of the points chosen in the intersection of the chart neighbourhoods. We note that the portion of the path in the intersection of two charts lies in a single connected component, so the group elements determined by any other choice of points in the intersection will be the same as the group elements we have obtained before. We are using the fact that the group element is uniquely determined in each connected component. As a result, the developing map defined in this case will be the same as the map defined before.

Finally, we show that the definition is independent of choice of path in the path homotopy class $[\alpha]$. Suppose we take a path β which is the same homotopy class, then there will exist a path homotopy $F(s,t)$ taking α to β, where $F(0,t) = \alpha(t)$ and $F(1,t) = \beta(t)$. We can break this homotopy F into a series of smaller homotopies F_j such that the image of the point $(s, t_0$ remains in a single chart neighbourhood for the duration of every small homotopy. Hence, we can assume that we can go from α to β with just one such homotopy. Then the path traced by the point $\alpha(t_i)$ under the homotopy will be contained in the connected component of x_i, so the group elements determined by them will be the same. Hence, we will get the same definition for the developing map for both the choices of path.

2. This is clear as the covering projection $[\alpha] \mapsto \alpha(1)$ is a local diffeomorphism from $\widetilde{M}$, to M the chart map ϕ_{n-1} is a local diffeomorphism from a neighbourhood of M to X and all the g_i's are real analytic diffeomorphisms. The developing map is just the composition of all these maps, hence it is also a local diffeomorphism.

3. Suppose we choose a different base point x_1 instead of x_0. Let γ be a fixed path from x_0 to x_1. Then with the new choice of basepoint, the point $[\alpha]$ in $\widetilde{M}$ will correspond to the path $\gamma * \alpha$. We will now have additional charts covering the portion of the path given by γ, and they will contribute additional group elements h_i in the definition of the developing map. Thus the new developing map D' will be defined by

$$\begin{aligned} D'([\alpha]) &= h_1 ... h_k g_1 ... g_{n-1} \phi_{n-1}(\alpha(1)) \\ &= hD([\alpha]) && \text{where } h = h_1 ... h_k \end{aligned}$$

The product h of the elements h_i is unique; the proof of this is the same as the proof given in the first part of the theorem. Thus, the new developing map D' differs from D by composition with a unique element of G.

□

The definition of the developing map actually provides us with another chart around the endpoint $\alpha(1)$ of the path, which is given by

$$\phi_{[\alpha]}(x) = g_1 g_2 ... g_{n-1} \phi_{n-1}(x)$$

Suppose $[\alpha]$ is a loop, then the transition map from this chart around the basepoint $\alpha(0)$ of the loop to the original chart (U_0, ϕ_0) will be given by a unique element of G; this element is called the holonomy of the loop $[\alpha]$. Thus, heuristically, the holonomy measures the difference between the chart obtained at the end of analytic continuation along a loop and the initial chart. As the choice of charts does not affect the product $g_1 ... g_{n-1}$ (Theorem 4.3), we may as well choose (U_{n-1}, ϕ_{n-1}) to be the same as (U_0, ϕ_0). Then we can clearly see from the definition of $\phi_{[\alpha]}$ that the holonomy of $[\alpha]$ is just the product $g_1 ... g_{n-1}$.

Theorem 4.4. *Let $T_{[\alpha]}$ be the covering isomorphism of $\widetilde{M}$ corresponding to the loop $[\alpha]$ in $\pi_1(M)$. Then if $g_{[\alpha]}$ denotes the holonomy of the loop $[\alpha]$, then $g_{[\alpha]}$ is the unique element of G satisfying*

$$D \circ T_{[\alpha]} = g_{[\alpha]} \circ D$$

Thus, the holonomy map $\rho \colon \pi_1(M) \to G$ is a group homomorphism.

Proof. We know that the covering isomorphism $T_{[\alpha]}$ acts on $\widetilde{M}$ as

$$T_{[\alpha]}([\beta]) = [\alpha] * [\beta]$$

So, for an arbitrary point $[\beta]$ in $\widetilde{M}$, we get that

$$D \circ T_{[\alpha]}([\beta]) = D([\alpha] * [\beta]) = D([\alpha * \beta])$$

Now, $[\alpha]$ is a loop based at the beginning of the path $[\beta]$. So when we develop the chart along the path $\alpha * \beta$, we will get the product $g_1 ... g_n$ corresponding to the development along the path $[\beta]$, composed with the product $h_1...h_k$ which corresponds to the development along the loop $[\alpha]$. In the paragraph preceding the lemma, we have proved that the product $h_1...h_k$ is actually the holonomy $g_{[\alpha]}$ of the loop $[\alpha]$. So, we get,

$$\begin{aligned} D \circ T_{[\alpha]}([\beta]) &= D([\alpha * \beta]) \\ &= h_1...h_k g_1...g_n \phi_{n-1}(\beta(1)) \qquad \text{as the endpoint of } \alpha \text{ is the same as the endpoint of } \alpha * \beta \\ &= g_{[\alpha]} \circ D([\beta]) \end{aligned}$$

Suppose any other element g in G satisfies this property, we will get

$$g_{[\alpha]} \circ D = g \circ D$$

So, both g and $g_{[\alpha]}$ will agree on the image of D, and in particular on some open set in X (as D is a local diffeomorphism), hence they must be equal (as elements of G are uniquely determined by their restriction to any open set).

Now, we can use this relation to show that the holonomy $\rho\colon \pi_1(M) \to G$ is a group homomorphism. We know that $T_{[\alpha*\beta]} = T_{[\alpha]} \circ T_{[\beta]}$, so

$$\begin{aligned} D \circ T_{[\alpha*\beta]} &= D \circ T_{[\alpha]} \circ T_{[\beta]} \\ &= g_{[\alpha]} \circ D \circ T_{[\beta]} \\ &= g_{[\alpha]} \circ g_{[\beta]} \circ D \end{aligned}$$

We have proved that $\rho([\alpha] * [\beta])$ is the unique element of G which satisfies the above relation, and hence, we must have

$$\rho([\alpha] * [\beta]) = g_{[\alpha]} \circ g_{[\beta]}$$

This means that ρ is a group homomorphism. □

Suppose we define the developing map using another base point and obtain a new map $D'\colon \widetilde{M} \to X$. Then, we know that $D' = gD$, for some unique element g of G. Then the new holonomy map ρ' will just be the conjugation of ρ by g. This is clear from the previous statement and Theorem 4.4.

4.4 Completeness of gluing of hyperbolic polygons

We have proved in Theorem 4.2 that if in a gluing of hyperbolic polygons, around every vertex of the gluing the interior angles sum to 2π, then the quotient manifold will inherit a hyperbolic structure. If we glue ideal polygons, this condition does not impose any restrictions on the polygons as they do not have vertices. In this section, we will analyse the conditions required for the gluing of hyperbolic ideal polygons to give a complete hyperbolic structure on the resulting 2-manifold.

In the gluing of the hyperbolic ideal polygons, there will be equivalence classes of ideal vertices based on the gluing of the edges. Consider any ideal vertex of the quotient manifold (an equivalence class of ideal vertices). Suppose the ideal hyperbolic polygons P_0,P_1,...,P_{n-1} share the ideal vertex v and let them be arranged in an anti-clockwise direction around this vertex. Consider any horocycle h_0 in P_0 centred at the vertex v. It will meet the edge of the polygon P_1. Now, there exists a unique horocycle with a given centre passing through a given point, so we can extend the horocycle h_0 in P_1 to another horocycle h_1. We can continue extending the horocycle like this at every edge and we will get horocycles h_2, h_3, ..., h_{n-1}. We will eventually reach the initial polygon P_0, and now the extended horocycle h_n may not be the same as the initial horocycle h_0. But as they are both horocycles which share the same vertex in P_0, the distance between them will be constant; let us denote this constant by $d(v)$. We use the convention that $d(v) > 0$ if the horocycle h_n is closer to v than h_0 and $d(v) < 0$ if it is farther away. We need to show that the parameter $d(v)$ is well defined.

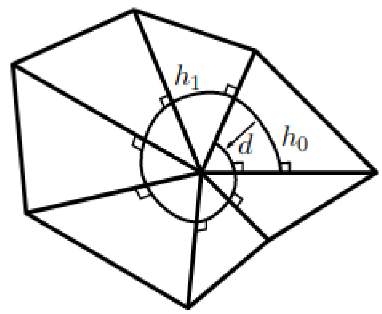

Figure 10: Extending horocycles to define d(v) - Image from pg. 38, Purcell, J. (2019). Hyperbolic Geometry and Knot Theory. [1]

Lemma 4.5. *The parameter $d(v)$ is well defined and is independent of the choice of the initial polygon or the initial horocycle*

Proof. Suppose we started in the same polygon but chose a different horocycle h_0'. Then h_0' must be at a constant distance x from h_0, so when we extend h_0' in P_1 to h_1', h_1' will also be at the same constant distance x from h_1 as P_0 and P_1 are identified isometrically along their common edge. So, each horocycle h_i' will be at a distance x from h_i, and hence h_n' will also be at a distance x from h_n. As a result, the distance between h_0 and h_n will be the same as the distance between h_0' and h_n', so the parameter $d(v)$ obtained will remain the same.

Now if we start with a different polygon P_i with a different choice of horocycle, then again it will be at a constant distance x from h_i. If we continue extending the horocycle and reach the polygon P_0, then the extended horocycle will be at a distance of x from h_n, which is itself at a distance of $d(v)$ from h_0. If we continue extending the horocycle till we come back to the polygon P_i again, the extended horocycle will be at a distance x from the horocycle h_{n+i}, which is itself at a distance $d(v)$ from h_i. So we will again get that the distance between the final horocycle and the initial one is $d(v)$. □

We will now use this parameter $d(v)$ to determine when the gluing of hyperbolic ideal polygons has a complete hyperbolic structure.

Theorem 4.6. *Suppose M is a 2-manifold obtained by gluing hyperbolic ideal polygons isometrically along their edges. Then M is complete with respect to the quotient metric if and only if $d(v) = 0$*

Proof. Suppose $d(v) > 0$ (if $d(v) < 0$ then the entire argument works the same way, we just have to extend the horocycles in the opposite direction). Then we can take the sequence of intersection points of the horocycles with the edges meeting v. We claim that this is a Cauchy sequence that doesn't converge. Suppose the ideal polygons are all placed in $\mathbb{H}^2$ with the ideal vertex v at ∞, then the horocycles will all be horizontal lines. After one full circuit of the ideal vertex, since $d(v) > 0$, the horocycles will continue extending in $\mathbb{H}^2$, but at a higher Y-coordinate, since they are at a constant distance d from the previous circuit. Then we can clearly see that the length of each full circuit of the horocycles around ∞ will be a constant factor smaller than the length of the previous horocycle, and the sum of lengths of all the circuits will converge to a finite value. This will mean that the length of the path along the horocycle between successive points of the sequence will tend to 0, and hence the distance between those points will also tend to 0. Thus the sequence of intersection points is a Cauchy sequence in M. However, this sequence does not converge, as for any neighbourhood in M, there will be infinitely many points of the sequence that lie outside the neighbourhood.

Suppose $d(v) = 0$ for all ideal vertices v in the gluing. Then for every ideal vertex, all horocycles will close up around it. Choose some set of horocycles that close up around every ideal vertex and remove the interior horoball for each horocycle. The resultant manifold will be a compact manifold with boundary which we denote by S. For each positive real number t, let S_t be the manifold obtained by removing the interior horoballs of horocycles at a distance t from the original set of horocycles (the distance is measured towards the vertex). Then we have a increasing collection of compact subsets S_t of M satisfying $\cup_{t\in\mathbb{R}} S_t = M$, where the sets S_t are such that $S_{t+\epsilon}$ contains an ϵ neighbourhood of S_t. Then any Cauchy sequence must be contained in some S_t for large enough t, so by the compactness of S_t, the Cauchy sequence must converge. Hence, M is complete. □

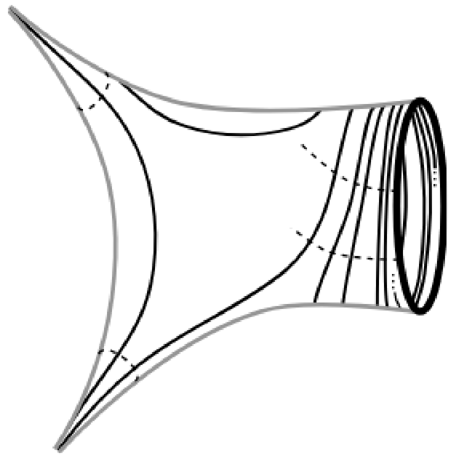

Figure 11: The geodesics for an incomplete gluing (solid lines) appear to rotate infinitely many times around the boundary and the horocycles (dotted lines) converge to the boundary) - Image from pg. 43, Purcell, J. (2019). Hyperbolic Geometry and Knot Theory. [1]

4.5 Criteria for completeness in terms of the developing map

Suppose M is a manifold with a (G, X) structure, where X is a Riemannian manifold and G is a group of isometries of M. Then, M will inherit a metric from X due to the geometric structure. This can be seen in the following way. As $D\colon \widetilde{M} \to X$ is a local diffeomorphism, we can use D to pull back the metric on X to a well defined metric on $\widetilde{M}$. Now as G is a subgroup of isometries of X, the metric on X is G invariant. This will show that the metric on $\widetilde{M}$ will be invariant under the action of $\pi_1(M)$. We know that M is the quotient of $\widetilde{M}$ under the action of $\pi_1(M)$. So M will inherit the metric from $\widetilde{M}$. We know that $\pi_1(M)$ acts on $\widetilde{M}$ by covering isomorphisms $T_{[\alpha]}$, where $[\alpha] \in \pi_1(M)$. Consider two tangent vectors v and w in the tangent space of some point in $\widetilde{M}$

$$\begin{aligned}
\langle T_{[\alpha]_*}v, T_{[\alpha]_*}w\rangle_{\widetilde{M}} &= \langle D_* \circ T_{[\alpha]_*}v, D_* \circ T_{[\alpha]_*}w\rangle_X \\
&= \langle g_{[\alpha]_*} \circ D_*v, g_{[\alpha]_*} \circ D_*w\rangle_X && \text{since } D \circ T_{[\alpha]} = g_{[\alpha]} \circ D \\
&= \langle D_*v, D_*w\rangle_X && \text{since G is a group of isometries of X} \\
&= \langle v, w\rangle_{\widetilde{M}} && \text{as the metric on } \widetilde{M} \text{ is the pullback of the metric on X}
\end{aligned}$$

Thus, the metric on $\widetilde{M}$ is invariant under the action of $\pi_1(M)$ and M inherits a Riemannian metric from $\widetilde{M}$. We wish to understand when this metric on M is complete.

Theorem 4.7. *Let M be a manifold with (G, X) structure, where X is a complete Riemannian manifold and G is a group of isometries acting on X. Then M is complete with the metric inherited from X if and only if the developing map $D\colon \widetilde{M} \to X$ is a covering map*

Proof. Let us assume that the developing map $D\colon \widetilde{M} \to X$ is a covering map. Consider any Cauchy sequence $\{x_n\}_{n=1}^{\infty}$ in M. After a sufficiently large value of n, all terms of the sequence will be contained in an evenly covered open neighbourhood in M. By definition of the metric on M, the covering projection from $\widetilde{M}$ to M is a local isometry, so the sequence $\{x_n\}$ in M will lift to a Cauchy sequence $\{\tilde{x}_n\}$ in $\widetilde{M}$. Again, by definition, $D\colon \widetilde{M} \to X$ is a local isometry, so the Cauchy sequence $\{\tilde{x}_n\}$ in $\widetilde{M}$ will project to a Cauchy sequence $D(\{\tilde{x}_n\})$ in X. Now as X is complete, the Cauchy sequence $D(\{\tilde{x}_n\})$ will converge to a point y in X. Now, as D is a covering map, y has an evenly covered neighbourhood containing all but finitely many terms of the sequence $D(\{\tilde{x}_n\})$. So we can lift y to a point $\tilde{y}$ which lies in the same neighbourhood where all but finitely many points of the sequence $\{\tilde{x}_n\}$ lie and we will see that the sequence $\{\tilde{x}_n\}$ converges to $\tilde{y}$. Now, if y is the projection of $\tilde{y}$ to M, then again any neighbourhood of y will contain all but finitely many terms of the sequence $\{x_n\}$, so the sequence $\{x_n\}$ must converge to y. Thus M must be complete.

Now suppose M is complete. We need to show that $D\colon \widetilde{M} \to X$ is a covering map. To prove this we will show that any path α_t in M lifts to a unique path $\widetilde{\alpha_t}$ in $\widetilde{M}$. We are using the fact that any local diffeomorphism which has unique path lifting property must be a covering map.

As M is complete, $\widetilde{M}$ must also be complete. Consider a Cauchy sequence $\{\tilde{x}_n\}$ in $\widetilde{M}$. As the covering projection is a local isometry, the sequence will project to a Cauchy sequence $\{x_n\}$ in M, and since M is complete, $\{x_n\}$ will converge to a point x in M. Now, consider an evenly covered open neighbourhood of x; it will contain all but finitely many points of the sequence $\{x_n\}$. Let $\tilde{x}$ be the lift of x in $\widetilde{M}$ which is in the same neighbourhood as the terms of the sequence $\{\tilde{x}_n\}$. Any neighbourhood of $\tilde{x}$ will contain all but finitely many points of the sequence $\{\tilde{x}_n\}$, so the sequence $\{\tilde{x}_n\}$ converges to $\tilde{x}$. Hence, $\widetilde{M}$ is complete.

Now let α_t be a path in X. As the developing map is a local diffeomorphism, we can lift some portion of the path α_t to some neighbourhood in $\widetilde{M}$. Let t_0 be the maximal time for which the path α_t lifts to $\widetilde{M}$ and suppose $t_0 \neq 1$. Now as $\widetilde{M}$ is complete, the points on the lifted path corresponding to $[0, t_0]$ will converge to a point $\tilde{x}$ and we can extend the lifting by defining $\widetilde{\alpha}(t_0) = \tilde{x}$. But as D is a local diffeomorphism, we can lift the path to extend to $[0, t_0 + \epsilon]$. This contradicts the maximality of t_0, so we must have $t_0 = 1$, that is, the path α_t lifts to all of $[0, 1]$. By construction and due to D being a local diffeomorphism, we see that this lifting must be unique. Hence, $D\colon \widetilde{M} \to X$ must be a covering map.

□

5 Conclusion

We have seen an algorithm to decompose the complement of alternating knots into ideal polyhedra using the knot diagram. We have also studied the basic ideas of plane hyperbolic geometry and the classification of isometries of $\mathbb{H}^2$. We have then defined geometric structures on manifolds and studied the conditions which ensure that the geometric structure on the manifold gives a complete metric on it. We plan to use these tools to decompose the complement of knots into geometric ideal tetrahedra and study Thurston's gluing consistency and gluing completeness equations, which will tell us when a gluing of hyperbolic ideal tetrahedra gives us a complete hyperbolic 3-manifold.

6 References

1. William P. Thurston, The geometry and topology of three-manifolds, Princeton Univ. Math. Dept. Notes, 1979. Available at: http://www.msri.org/communications/books/gt3m.

2. Purcell, J. (2019). Hyperbolic Geometry and Knot Theory. [ebook]
   Available at: http://users.monash.edu/ jpurcell/hypknottheory.html

3. Stillwell, J., 1992. Geometry of surfaces. 1st ed. New York: Springer-Verlag.

4. Katok, S. (1992). Fuchsian groups. Chicago: University of Chicago Press, pp.23-24.