5. Hyperbolic geometry.

Here we give a brief outline of some of the main features of hyperbolic geometry. Again, this will serve mainly as a source of examples and motivation, and we will not give detailed proofs. One-dimensional hyperbolic space is just the real line, so we begin in dimension 2. The main ideas generalise to higher dimensions.

5.1. The hyperbolic plane.

We describe the "Poincaré model" for the hyperbolic plane. For this it is convenient to use complex coordinates. Let $D = \{z \in \mathbf{C} \mid |z| < 1\}$. Suppose $\alpha : I \longrightarrow D$ is a smooth path. We write $\alpha'(t) \in \mathbf{C}$ for the complex derivative at t. Thus, $|\alpha'(t)|$ is the "speed" at time t. The euclidean length of α is thus given by the formula $l_E(\alpha) = \int_I |\alpha'(t)| dt$. This is equal to the "rectifiable" length as defined in Section 3.

We now modify this by the introduction a scaling factor, $\lambda : D \longrightarrow (0, \infty)$. The appropriate formula is: $\lambda(z) = 2/(1 - |z|^2)$. The hyperbolic length of α is thus given by $l_H(\alpha) = \int_I \lambda(\alpha(t))|\alpha'(t)|dt$.

Note that as z approaches ∂D in the euclidean sense, then $\lambda(z) \to \infty$. Thus close to ∂D, things big in hyperbolic space may look very small to us in euclidean space. Indeed, since $\int_0^\infty \frac{2}{1-x^2} dx = \infty$, one needs to travel an infinite hyperbolic distance to approach ∂D. For this reason, ∂D is often referred to as the *ideal* boundary — we never actually get there.

Given $x, y \in D$, write $\rho(x, y) = \inf\{l_H(\alpha)\}$ as α varies over all smooth paths from x to y. In fact, the minimum is attained — there is always a smooth geodesic from x to y. The remark about the ideal boundary in the previous paragraph boils down to saying that this metric is complete. Moreover, if we want to get between two points x and y as quickly as possible, it would seem a good idea to move a little towards the centre of the disc, in the euclidean sense. Thus we would expect hyperbolic geodesics approach the middle of the disc relative to their euclidean counterparts.

For a more precise analysis, we need the notion of a *Möbius transformation*. This is a map $f : \mathbf{C} \cup \{\infty\} \longrightarrow \mathbf{C} \cup \{\infty\}$ of the form $f(z) = \frac{az+b}{cz+d}$ for constants $a, b, c, d \in \mathbf{C}$ with $ad - bc \neq 0$. We set $f(\infty) = a/c$ and $f(-c/d) = \infty$ (though we don't really have to worry about ∞ here). It is usual to normalise so that $ad - bc = 1$. A Möbius transformation is bijective, and one sees easily that the set of such transformations forms a group under composition. Since it is complex analytic, any Möbius tranformation is conformal, i.e. it preserves angles.

Here are some useful observations about Möbius transformations, which we can leave as exercises:

Exercises

(1) A Möbius transformation sends euclidean circles to euclidean circles, where we allow a straight line union ∞ to be a "circle". (Warning: it need not preserve centres of circles.)

(2) If $d = \bar{a}$ and $c = \bar{b}$ (the complex conjugates) and $|a|^2 - |b|^2 > 0$ then $f(D) = D$. (We shall normalise so that $|a|^2 - |b|^2 = 1$.) In fact, any Möbius tranformation preserving D must have this form.

(3) Such an f (as in (2)) is an isometry of (D, ρ). For this, one needs to check that if α is a smooth path, then $l_H(f \circ \alpha) = l_H(\alpha)$. This follows from the formula, $\lambda(f(z))|f'(z)| = \lambda(z)$, which can be verified by direct calculation.

(4) If $z, w \in D$, then there is some such f sending z to w. (Without loss of generality, $w = 0$.)

(5) If $p, q, r \in \partial D$ are distinct, and $p', q', r' \in \partial D$ are distinct, and the orientation of p, q, r is the same as that of p', q', r', then there is some such f with $f(p) = p'$, $f(q) = q'$ and $f(r) = r'$. Here is one way to see this. First show there is a (unique) Möbius transformation taking any three distinct points of $\mathbf{C} \cup \{\infty\}$ to any other three distinct points. (Since they form a group, we could take one set to be $\{0, 1, \infty\}$.) By (1), if all six points lie in ∂D then the Möbius transformation must preserve ∂D, since three points determine a euclidean circle. The condition about orientation is needed so that f sends the interior of D to the interior, rather than the exterior.

Putting together (3) and (4), we see that (D, ρ) is *homogeneous* that is, there is an isometry taking any point to any other point. It thus looks the same everywhere. In fact, any rotation about the origin is clearly an isometry (and a Möbius tranformation). Thus (D, ρ) is also *isotropic* — it looks the same in all directions. It thus shares these properties with the euclidean plane.

Now by symmetry it is easily seen that any euculidean diameter of D, (for example, the interval $(-1, 1)$) is a bi-infinite geodesic with respect to the metric ρ. Indeed it is the unique geodesic between any pair of points on it. Under isometries of the above type it is mapped onto arcs of euclidean circles othogonal to ∂D. Since any pair of points of D lie on such a circle, we see that all geodesics must be of this type, and so we conclude (Figure 5a):

Proposition 5.1 : *Bi-infinite geodesics in the Poincaré disc are arcs of euclidean circles othogonal to the ∂D (including diameters of the disc).* $\diamondsuit$

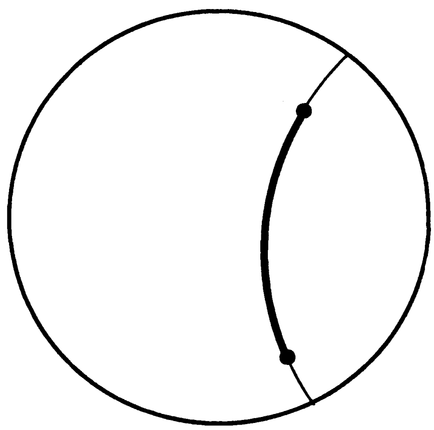

Figure 5a.

We remark that, in fact, all orientation preserving isometries of (D, ρ) are Möbius transformations of the above type. This is not very hard to deduce given our description of geodesics, but we shall not formally be needing it.

The isometry type of space we have just constructed is generally referred to as the *hyperbolic plane*, and denoted $\mathbf{H}^2$. It has many descriptions. The one we have given is called the *Poincaré model*

Hyperbolic geometry has its roots in attempts to understand the "parallel postulate" as formulated by Euclid about 2300 years ago. Its discovery, due to Bolyai, Lobachevsky and Gauß in the 1830s is one of the great landmarks in mathematics. Varous explicit models of hyperbolic geometry were subsequently discovered by Beltrami, Poincaré, Klein etc.

5.2. Some properties of the hyperbolic plane.

(1) In general angles in hyperbolic geometry are "smaller" than in the corresponding situation in Euclidean geometry. For example, if T is a triangle with angles p, q, r, then $p+q+r < \pi$. In fact, one can show that the area of T is $\pi - (p+q+r)$. One can allow for one or more of the vertices to lie in the ideal boundary, ∂D, in which case the corresponding angle is deemed to be 0. An *ideal triangle* is one where all three vertices are ideal. Its area is π.

(2) Triangles are "thin". One way of expressing this is to say that there is some fixed constant $k > 0$, so that if T is any triangle there is some point, $x \in \mathbf{H}^2$, whose distance from all three sides is at most k. (One can, in fact, take x in the interior of T.) To verify this, one can first deal with case of an ideal triangle. By exercise (5) above, we see that there is a hyperbolic isometry carrying any ideal triangle to any other. In particular, we can suppose that its vertices are three equally spaced points in ∂D. A calculation now shows that the centre, 0, of this triangle is a distance $\frac{1}{2}\log 3$ from each of the three sides. In fact, the largest disc in the interior of an ideal triangle is the one of radius $\frac{1}{2}\log 3$ about the centre. It touches all three sides. Now, for a general triangle, take a disc in the interior of maximal radius, r, say. This must touch all three sides, otherwise we could make it bigger. The centre of the disc is thus a distance r from all three sides. Now by pushing the three vertices of the triangle towards the ideal boundary, we can place our triangle inside an ideal triangle. It therefore follows that $r \leq \frac{1}{2}\log 3$, and so the same constant, $k = \frac{1}{2}\log 3$, works for all triangles (Figure 5b).

(3) A (round) circle of radius r in $\mathbf{H}^2$ has length $2\pi \sinh r$. A round disc $B(r)$, of radius r has area $2\pi(\cosh r - 1)$. (This is an exercise in

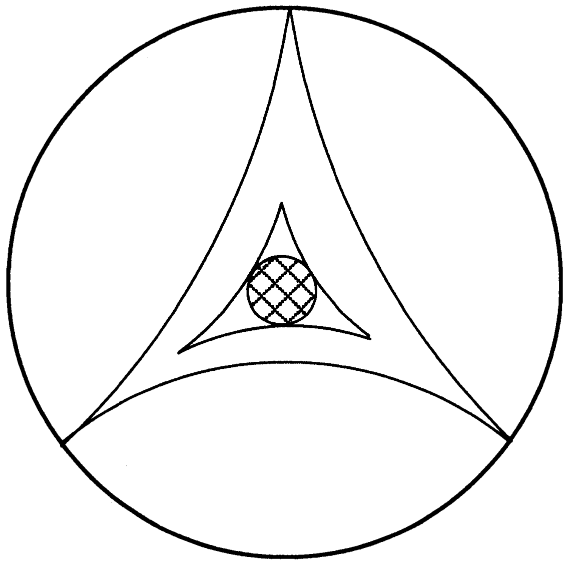

Figure 5b.

integration: note that a hyperbolic circle about the origin, 0, of the Poincaré model is also a euclidean circle.) We note in particular, that $\text{area}(B(r)) \leq \text{length}(\partial B(r))$. In fact, if B is any topological disc in $\mathbf{H}^2$, one can show with a little bit of work that $\text{area}(B) \leq \text{length}(\partial B)$. (It turns out that the round disc is the "worst case", but this is much harder to show.) This is in contrast to the euclidean plane, where such bounds are quadratic. Inequalities of this sort are called "isoperimetric inequalities". We will briefly mention this again in Section 6.

5.3. Tessellations of $\mathbf{H}^2$.

Suppose that $n \in \mathbf{N}$, $n \geq 3$. The regular euclidean n-gon has all angles equal to $(1 - \frac{2}{n})\pi$. If $0 < \theta < (1 - \frac{2}{n})\pi$, then one can construct a regular hyperbolic n-gon with all angles equal to θ. This can be seen simply by using a continuity argument: start with a very small regular n-gon centred at the origin of the Poincaré disc. Now push all the vertices out to the ideal boundary at a uniform rate. The angles must all tend to 0. Thus at some intermediate point, they will all equal θ. (In fact the angles decrease monotonically, so the polygon is unique.) Now if θ has the form $2\pi/m$ for some $m \in \mathbf{N}$ $n \geq 3$, we get a regular tessellation of the hyperbolic plane by repeatedly reflecting the polygon in its edges. (This requires some sort of formal argument, for example, applying "Poincaré's Theorem", but there are lots of pretty computer images to demonstrate that it works.) Note that the condition $\frac{2\pi}{m} < (1 - \frac{2}{n})\pi$ reduces to $\frac{1}{m} + \frac{1}{n} < \frac{1}{2}$. Thus we get:

Proposition 5.2 : *If $m, n \in \mathbf{N}$ with $\frac{1}{m} + \frac{1}{n} < \frac{1}{2}$, then there is a regular tessellation of the hyperbolic plane by regular n-gons so that m such n-gons meet at every vertex.* $\diamond$

We remark that in the euclidean situation, the corresponding condition is $\frac{1}{m} + \frac{1}{n} = \frac{1}{2}$. In this case, we just get the three familiar tilings where $(m, n) = (3, 6), (4, 4), (6, 3)$.

5.4. Surfaces.

This has consequences for the geometry and topology of surfaces. For simplicity, we consider here only closed orientable surfaces. These are classified by their "genus", which is a non-negative integer.

The sphere (genus 0) clearly has "spherical geometry" — as the unit sphere in $\mathbf{R}^3$. It is simply connected. In other words, its fundamental group is trivial. Moreover, it is compact, and therefore quasi-isometric to a point. At this point, geometric group theorists start to lose interest in spherical geometry.

The torus, T, (genus 1) is a bit more interesting. We can think of the torus topologically as obtained by gluing together the opposite edges of the unit square, $[0, 1]^2$. Note that at the vertex we get a total angle of $4(\pi/2) = 2\pi$, so there is no singularity there. We get a metric on the torus which is locally euclidean (i.e. every point has a neighbourhood isometric to an open subset of the euclidean plane). Such a metric is often referred to simply as a "euclidean structure". The universal cover of the torus is then, in a natural way, identified with the euclidean plane, $\mathbf{R}^2$, with the fundamental group acting by translations. (Given the remarks on universal covers in Section 4, this in fact gives a proof that $\pi_1(T) \cong \mathbf{Z}^2$.) The square we used to construct T lifts to a regular square tiling of the plane. Note that the edges of the square project to loops representing generators, a, b of $\pi_1(T)$ (Figure 4b). Reading around the boundary of the square we see that $[a, b] = aba^{-1}b^{-1} = 1$. The 1-skeleton of the square tessellation of the plane can be identified with the Cayley graph of $\pi_1(T)$ with respect to these generators. In particular, we see that $\pi_1(T)$ is quasi-isometric to the euclidean plane.

We have made most of the above observations, in some form, before. We now move into new territory.

Let S be the closed surface of genus 2. We can construct S by taking a (regular) octagon and gluing together its edges. We do this according to the cyclic labelling $ABA^{-1}B^{-1}CDC^{-1}D^{-1}$, so that the first edge gets mapped to the third with opposite orientation etc. (Figure 5c).

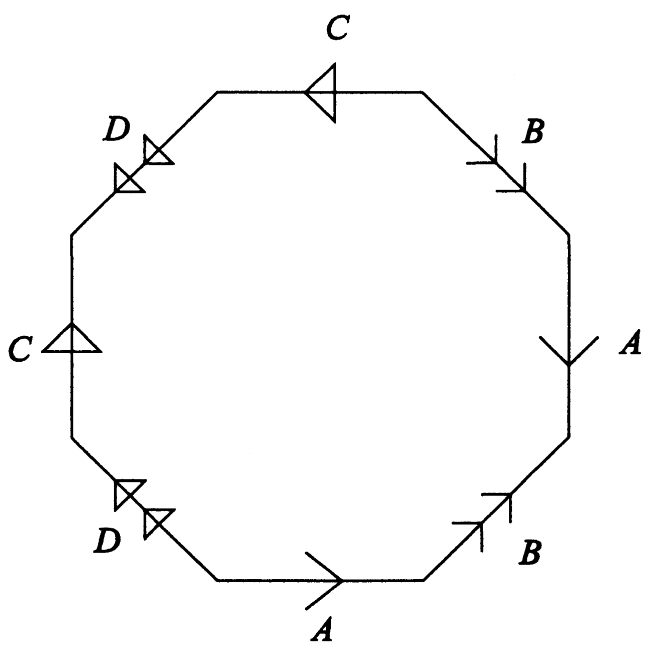

Figure 5c.

If we try the above constuction with a euclidean octagon we would end up with an angle of $8(3\pi/4) = 6\pi > 2\pi$ at the vertex, so our euclidean structure would be singular. It is therefore very natural to take instead the regular hyperbolic octagon all of whose angles are $\pi/4$. In this way we get a metric on S that is locally hyperbolic, generally termed a *hyperbolic structure* in S. The universal cover is $\mathbf{H}^2$ and our octagon lifts to a tessellation of the type $(m, n) = (8, 8)$ described above. The edges of the octagon project to loops a, b, c, d, and reading around the boundary, we see that $[a, b][c, d] = aba^{-1}b^{-1}cdc^{-1}d^{-1} = 1$ (Figure 5d).

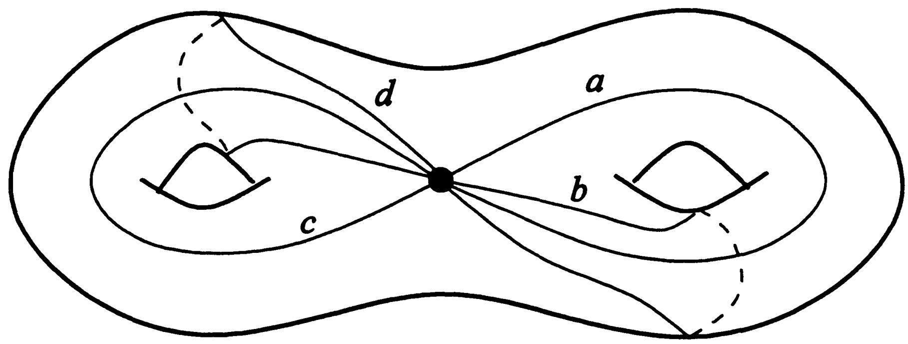

Figure 5d.

In fact, it turns out that $\langle a, b, c, d \mid [a, b][c, d] = 1 \rangle$ is a presentation for $\pi_1(S)$. Its Cayley graph can be identified with the 1-skeleton of our $(8, 8)$ tessellation of $\mathbf{H}^2$. We see that $\pi_1(S) \sim \mathbf{H}^2$. Indeed, for this, it is enough, by Theorem 3.6, to note that S is the quotient of a p.d.c. isometric action on $\mathbf{H}^2$.

More generally, if S is a closed surface of genus, $g \geq 2$, we get a similar story taking a regular $4g$-gon with cone angles $\pi/2g$. We get

$$\pi_1(S) \cong \langle a_1, b_1, a_2, b_2, \ldots, a_g, b_g \mid [a_1, b_1][a_2, b_2] \cdots [a_g, b_g] = 1 \rangle.$$

The Cayley graph is the 1-skeleton of a $(4g, 4g)$ tesselation of $\mathbf{H}^2$.

In summary we see:

Proposition 5.3 : *If S is a closed surface of genus at least 2, then $\pi_1(S) \sim \mathbf{H}^2$.* $\diamondsuit$

Remark: We just showed that such a surface admits some hyperbolic structure. However, there are lots of variations on this construction, and in fact, there are many different hyperbolic structures one could put on such a surface. Indeed there is a whole "Teichmüller space" of them, and Teichmüller theory is a vast subject in itself.

We also note:

Proposition 5.4 : *If S and S' are closed orientable surfaces of genus at least 2, then $\pi_1(S) \approx \pi_1(S')$.*

Proof : We can do this by a similar argument to Theorem 4.2. Imagine embedding the graph K_n in $\mathbf{R}^3$, and thickenning it up to a 3-dimensional object (called a "handlebody") whose boundary is a surface of genus $n+1$. (See Figure 5e, where $n = 5$.) Now we do essentially the same construction, to see that a surface of genus $p \geq 2$ and a surface of genus $q \geq 2$ are both covered by a surface of genus $pq - p - q + 2$. $\diamondsuit$

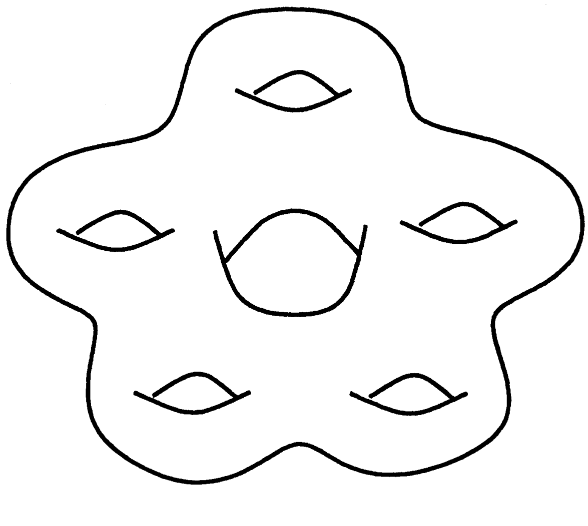

Figure 5e.

(Note that the above was a topological not a geometric construction. The covers need not respect any given hyperbolic structures.)

Theorem 5.5 : *Suppose S and S' are closed orientable sufaces. If $\pi_1(S) \sim \pi_1(S')$ then $\pi_1(S) \approx \pi_1(S')$.*

If we believe that the euclidean plane is not quasi-isometric to the hyperbolic plane, then there are exactly three quasi-isometry classes — one each for the sphere, the torus, and all higher genus surfaces. The result then follows by Propostion 5.4. The fact that $\mathbf{R}^2 \not\sim \mathbf{H}^2$ will follow from our discussion of hyperbolicity in Section 6.

Alternatively you believe that a higher genus surface group is not (virtually) $\mathbf{Z}^2$, then the fact that $\mathbf{R}^2 \not\sim \mathbf{H}^2$ also follows from the q.i. invariance of virually abelian groups, cited but not proven, in Section 3. This is, however, much more work than is necessary for this result.

Fundamental groups of closed surfaces, other than the 2-sphere, are generally just referred to as *surface groups*. (Any non-orientable surface is double covered by an orientable surface, and so non-orientable surfaces can easily be brought into the above discussion.)

Fact: Any f.g. group quasi-isometric to a surface group is a virtual surface group.

The case of the torus was already discussed in Section 3. The hyperbolic case (genus at least 2) is a difficult result of Tukia, Gabai and Casson and Jungreis.

In fact any group quasi-isometric to a complete riemannian plane is a virtual surface group. This was shown by Mess (modulo the completion of the above theorem which came later).

5.5. 3-dimensional hyperbolic geometry.

Our construction of the Poincaré model makes sense in any dimension n, except that we don't have such convenient complex coordinates. In this case, we take the disc $D^n = \{\underline{x} \in \mathbf{R}^n \mid ||\underline{x}|| < 1\}$. We scale the metric by the same factor $\lambda(\underline{x}) = \frac{2}{1-||\underline{x}||^2}$. We get a complete geodesic metric, ρ, and the isometry class of (D, ρ) is referred to as "hyperbolic n-space", $\mathbf{H}^n$. It is homogeneous and isotropic. Its ideal boundary, ∂D, is an $(n-1)$-sphere. Once again, bi-infinite geodesics are arcs of euclidean circles (or diameters) orthogonal to ∂D. More generally any euclidean sphere of any dimension, meeting ∂D othogonally intersects D in a hyperbolic subspace isometric to $\mathbf{H}^m$ for some $m < n$ — there is an isometry of (D, ρ) that maps it to a euclidean subspace through the origin, thereby giving us a lower dimensional Poincaré model.

There has been an enormous amount of work on 3-dimensional hyperbolic geometry, going back over a hundred years. One can construct lots of examples of polyhedra and tessellations, though the situation becomes more complicated. One can also use these to construct examples of compact hyperbolic 3-manifolds. The "Seifert-Weber space" is a nice example made out of a dodecahedron. In the 1970s people, such as Riley, started to notice that "many" 3-manifolds admitted hyperbolic structures. In the late 1970s Thurston revolutionised the subject by making the conjecture that every compact 3-manifold can be cut into pieces in a natural way so that each piece has a geometric structure. There are eight geometries in dimension 3, but by far the richest source of examples is hyperbolic geometry. In 2003, Perelman claimed a proof of Thurston's conjecture, building on earlier work of Hamilton.

This is a vast subject, we won't have time to look into here. We'll just mention a few facts relevant to group theory.

Let M be a closed hyperbolic 3-manifold, and $\Gamma = \pi_1(M)$. Then Γ is finitely generated (in fact, finitely presented), and $\Gamma \sim \mathbf{H}^3$. Thus any two such groups are q.i. There are however examples of closed hyperbolic 3-manifolds which do not have any common finite cover. (In contrast to the 2-dimenional case, the hyperbolic structure on a closed 3-manifold is unique, and it follows that covers are forced to respect hyperbolic metrics.) The verification of this involves algebraic number theory, so we won't describe it here. The point is that a closed hyperbolic 3-manifold has associated to it "stable trace field", a finite extention of the rationals, which one can compute. If these are different, then the groups are incommensurable. (It is likely that, in some sense, one would expect a "random" pair of hyperbolic 3-manifolds to be incommensurable.) This is therefore again in contrast to the 2-dimensional case, where any two surface groups of genus at least 2 are commensurable.