William P. Thurston

The Geometry and Topology of Three-Manifolds

Electronic version 1.1 - March 2002
http://www.msri.org/publications/books/gt3m/

This is an electronic edition of the 1980 notes distributed by Princeton University. The text was typed in TEX by Sheila Newbery, who also scanned the figures. Typos have been corrected (and probably others introduced), but otherwise no attempt has been made to update the contents. Genevieve Walsh compiled the index.

Numbers on the right margin correspond to the original edition's page numbers.

Thurston's *Three-Dimensional Geometry and Topology*, Vol. 1 (Princeton University Press, 1997) is a considerable expansion of the first few chapters of these notes. Later chapters have not yet appeared in book form.

Please send corrections to Silvio Levy at levy@msri.org.

CHAPTER 3

Geometric structures on manifolds

A manifold is a topological space which is locally modelled on $\mathbb{R}^n$. The notion of what it means to be locally modelled on $\mathbb{R}^n$ can be made definite in many different ways, yielding many different sorts of manifolds. In general, to define a kind of manifold, we need to define a set $\mathcal{G}$ of gluing maps which are to be permitted for piecing the manifold together out of chunks of $\mathbb{R}^n$. Such a manifold is called a $\mathcal{G}$-manifold. $\mathcal{G}$ should satisfy some obvious properties which make it a *pseudogroup* of local homeomorphisms between open sets of $\mathbb{R}^n$:

(i) The restriction of an element $g \in \mathcal{G}$ to any open set in its domain is also in $\mathcal{G}$.
(ii) The composition $g_1 \circ g_2$ of two elements of $\mathcal{G}$, when defined, is in $\mathcal{G}$.
(iii) The inverse of an element of $\mathcal{G}$ is in $\mathcal{G}$.
(iv) The property of being in $\mathcal{G}$ is local, so that if $U = \bigcup_\alpha U_\alpha$ and if g is a local homeomorphism $g : U \to V$ whose restriction to each U_α is in $\mathcal{G}$, then $g \in \mathcal{G}$.

It is convenient also to permit $\mathcal{G}$ to be a pseudogroup acting on *any* manifold, although, as long as $\mathcal{G}$ is transitive, this doesn't give any new types of manifolds. See Haefliger, in Springer Lecture Notes #197, for a discussion.

A group G acting on a manifold X determines a pseudogroup which consists of restrictions of elements of G to open sets in X. A (G, X)-manifold means a manifold 3.2
glued together using this pseudogroup of restrictions of elements of G.

Examples. If $\mathcal{G}$ is the pseudogroup of local C^r diffeomorphisms of $\mathbb{R}^n$, then a $\mathcal{G}$-manifold is a C^r-manifold, or more loosely, a differentiable manifold (provided $r \geq 1$).

If $\mathcal{G}$ is the pseudogroup of local piecewise-linear homeomorphisms, then a $\mathcal{G}$-manifold is a PL-manifold. If G is the group of affine transformations of $\mathbb{R}^n$, then a $(G, \mathbb{R}^n)$-manifold is called an *affine* manifold. For instance, given a constant $\lambda > 1$ consider an annulus of radii 1 and $\lambda + \epsilon$. Identify neighborhoods of the two boundary components by the map $x \to \lambda x$. The resulting manifold, topologically, is T^2.

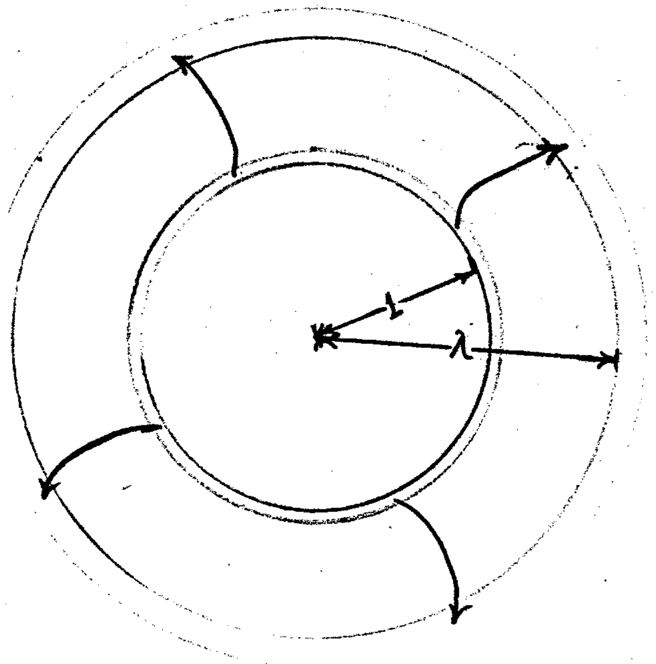

3.3

Here is another method, due to John Smillie, for constructing affine structures on T^2 from any quadrilateral Q in the plane. Identify the opposite edges of Q by the orientation-preserving similarities which carry one to the other. Since similarities preserve angles, the sum of the angles about the vertex in the resulting complex is 2π, so it has an affine structure. We shall see later how such structures on T^2 are intimately connected with questions concerning Dehn surgery in three-manifolds.

The literature about affine manifolds is interesting. Milnor showed that the only closed two-dimensional affine manifolds are tori and Klein bottles. The main unsolved question about affine manifolds is whether in general an affine manifold has Euler characteristic zero.

If G is the group of isometries of Euclidean space E^n, then a (G, E^n)-manifold is called a *Euclidean* manifold, or often a *flat* manifold. Bieberbach proved that a Euclidean manifold is finitely covered by a torus. Furthermore, a Euclidean structure automatically gives an affine structure, and Bieberbach proved that closed Euclidean manifolds with the same fundamental group are equivalent as affine manifolds. If G is the group $O(n+1)$ acting on elliptic space $\mathbb{P}^n$ (or on S^n), then we obtain *elliptic* manifolds.

CONJECTURE. Every three-manifold with finite fundamental group has an elliptic structure.

This conjecture is a stronger version of the Poincaré conjecture; we shall see the logic shortly. All known three-manifolds with finite fundamental group certainly have elliptic structures. 3.4

As a final example (for the present), when G is the group of isometries of hyperbolic space H^n, then a (G, H^n)-manifold is a *hyperbolic* manifold. For instance, any surface of negative Euler characteristic has a hyperbolic structure. The surface of genus two is an illustrative example.

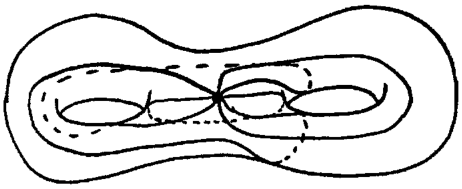

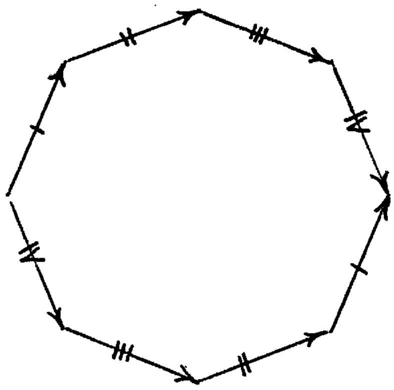

Topologically, this surface is obtained by identifying the sides of an octagon, in the pattern above, for instance. An example of a hyperbolic structure on the surface is obtained form any hyperbolic octagon whose opposite edges have equal lengths and whose angle sum is 2π, by identifying in the same pattern. There is a regular octagon with angles $\pi/4$, for instance. 3.5

3.6

3.1. A hyperbolic structure on the figure-eight knot complement.

Consider a regular tetrahedron in Euclidean space, inscribed in the unit sphere, so that its vertices are on the sphere. Now interpret this tetrahedron to lie in the projective model for hyperbolic space, so that it determines an ideal hyperbolic simplex: combinatorially, a simplex with its vertices deleted. The dihedral angles of the hyperbolic simplex are $60°$. This may be seen by extending its faces to the sphere at infinity, which they meet in four circles which meet each other in $60°$ angles.

By considering the Poincaré disk model, one sees immediately that the angle made by two planes is the same as the angle of their bounding circles on the sphere at infinity.

Take two copies of this ideal simplex, and glue the faces together, in the pattern described in Chapter 1, using Euclidean isometries, which are also (in this case) hyperbolic isometries, to identify faces. This gives a hyperbolic structure to the resulting manifold, since the angles add up to $360°$ around each edge. 3.7

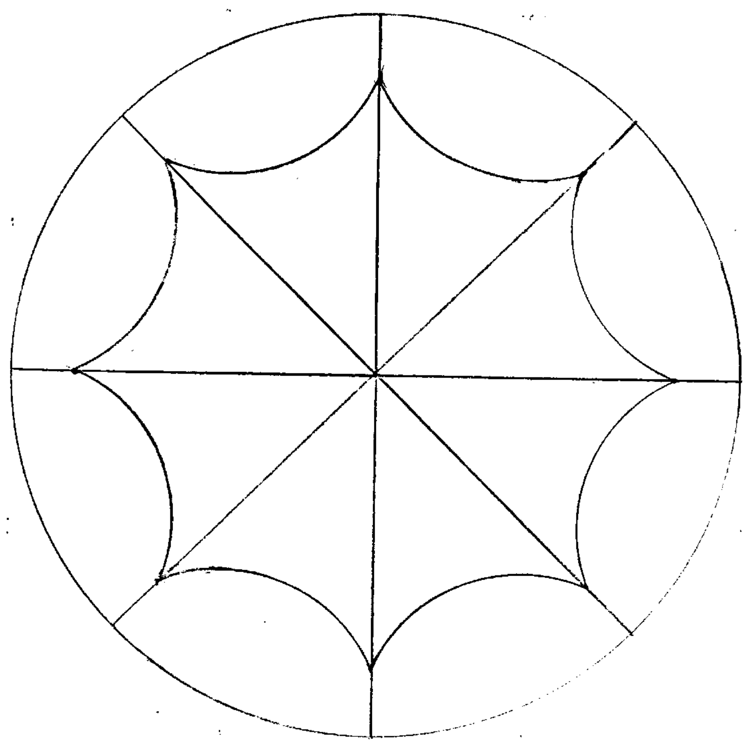

A regular octagon with angles $\pi/4$,
whose sides can be identified to give a surface of genus 2.

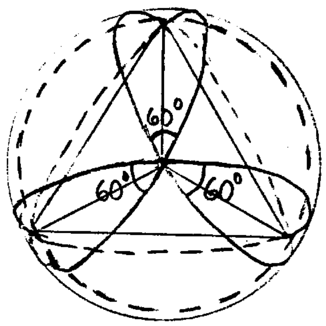

A tetrahedron inscribed in the unit sphere, top view.

According to Magnus, *Hyperbolic Tesselations*, this manifold was constructed by Gieseking in 1912 (but without any relation to knots). R. Riley showed that the figure-eight knot complement has a hyperbolic structure (which agrees with this one). This manifold also coincides with one of the hyperbolic manifolds obtained by an arithmetic construction, because the fundamental group of the complement of the figure-eight knot is isomorphic to a subgroup of index 12 in $\mathrm{PSL}_2(\mathbb{Z}[\omega])$, where ω is a primitive cube root of unity.

3.2. A hyperbolic manifold with geodesic boundary.

Here is another manifold which is obtained from two tetrahedra. First glue the two tetrahedra along one face; then glue the remaining faces according to this diagram:

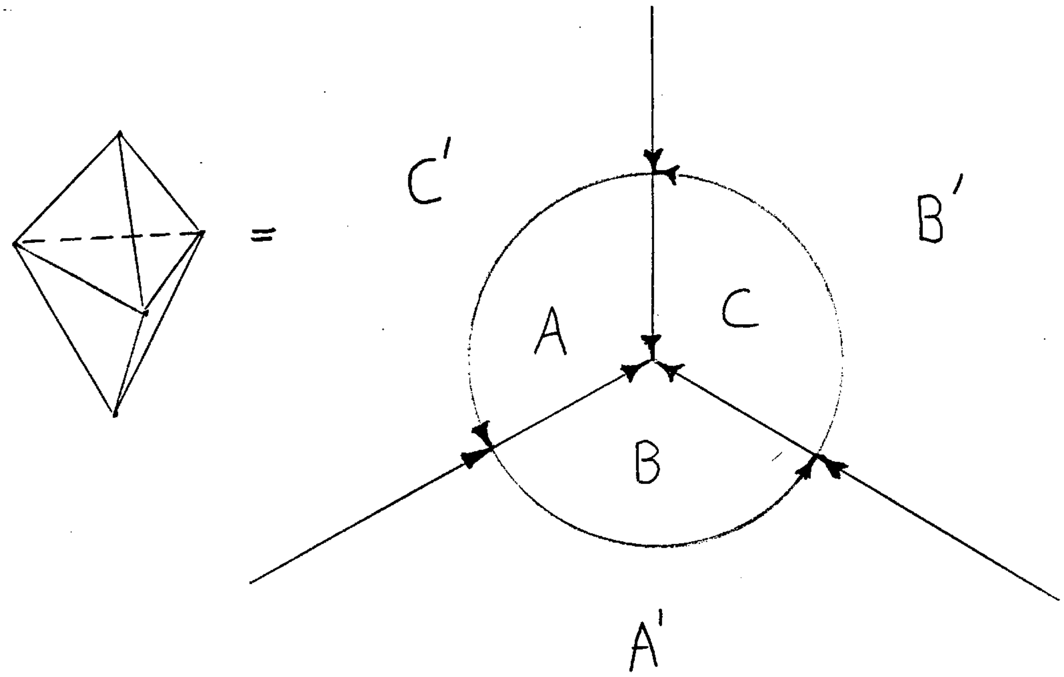

3.8

In the diagram, one vertex has been removed so that the polyhedron can be flattened out in the plane. The resulting complex has only one edge and one vertex. The manifold M obtained by removing a neighborhood of the vertex is oriented with boundary a surface of genus 2.

Consider now a one-parameter family of regular tetrahedra in the projective model for hyperbolic space centered at the origin in Euclidean space, beginning with the tetrahedron whose vertices are on the sphere at infinity, and expanding until the edges are all tangent to the sphere at infinity. The dihedral angles go from 60° to 0°, so somewhere in between, there is a tetrahedron with 30° dihedral angles. Truncate this simplex along each plane $v^{\perp}$, where v is a vertex (outside the unit ball), to obtain a *stunted simplex* with all angles 90° or 30°:

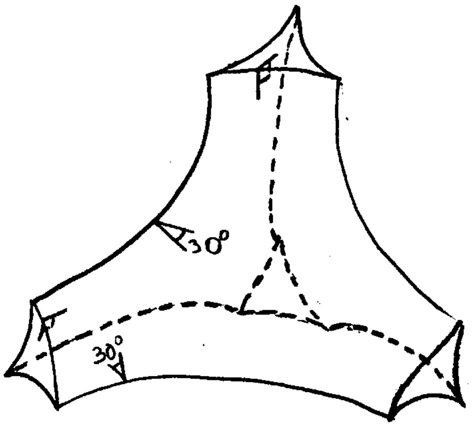

3.9

Two copies glued together give a hyperbolic structure for M, where the boundary of M (which comes from the triangular faces of the stunted simplices) is totally geodesic. A *closed hyperbolic three-manifold* can be obtained by doubling this example, i.e., taking two copies of M and gluing them together by the "identity" map on the boundary.

3.3. The Whitehead link complement.

The Whitehead link may be spanned by a two-complex which cuts the complement into an octahedron, with vertices deleted:

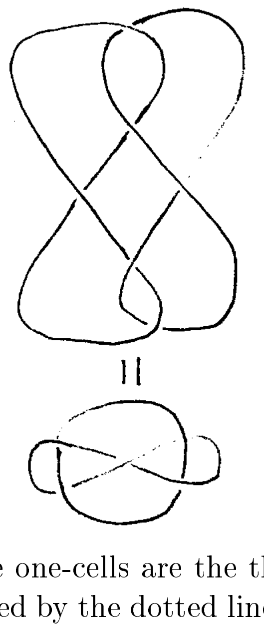

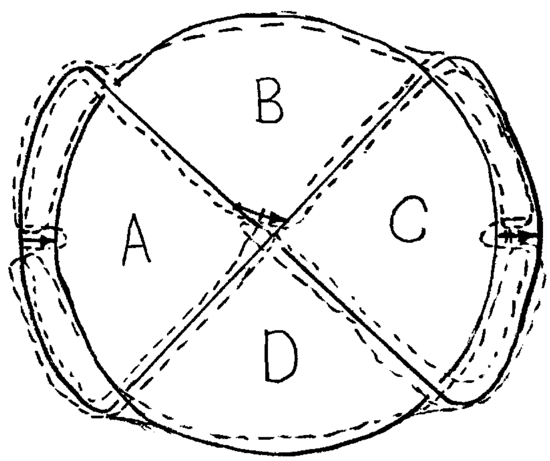

The one-cells are the three arrows, and the attaching maps for the two-cells are
indicated by the dotted lines. The three-cell is an octahedron (with vertices deleted), 3.10
and the faces are identified thus:

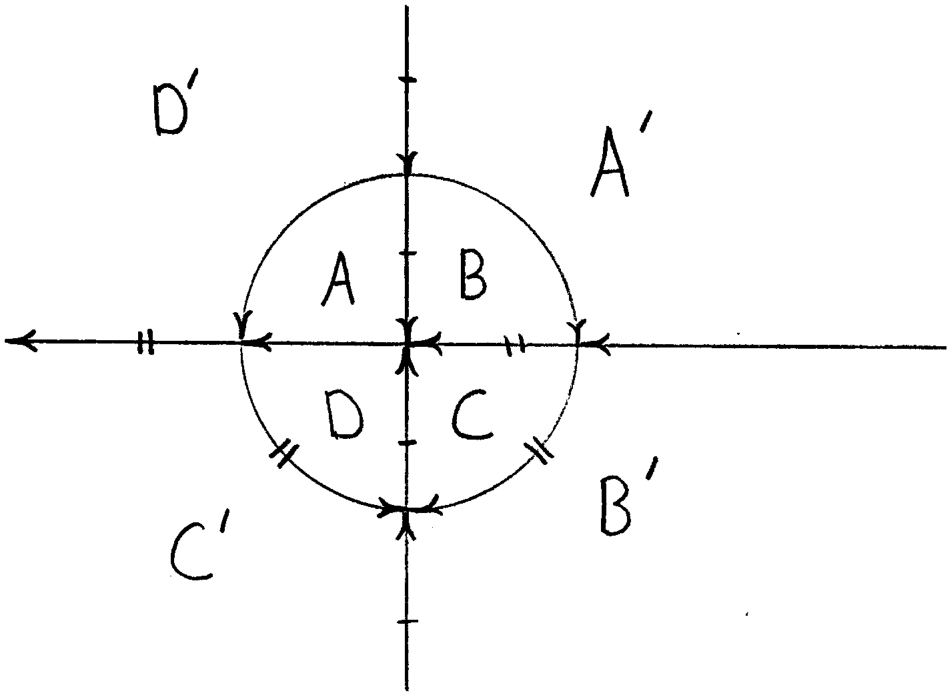

A hyperbolic structure may be obtained from a Euclidean regular octahedron inscribed in the unit sphere. Interpreted as lying in the projective model for hyperbolic space, this octahedron is an ideal octahedron with all dihedral angles 90°.

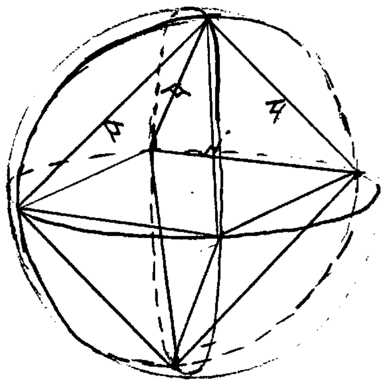

3.11

Gluing it in the indicated pattern, again using Euclidean isometries between the faces (which happen to be hyperbolic isometries as well) gives a hyperbolic structure for the complement of the Whitehead link.

3.4. The Borromean rings complement.

This is spanned by a two-complex which cuts the complement into two ideal octahedra:

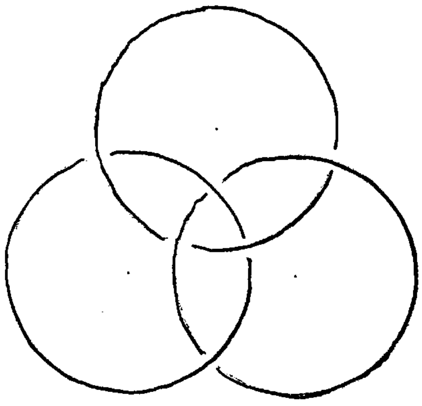

Borromean rings

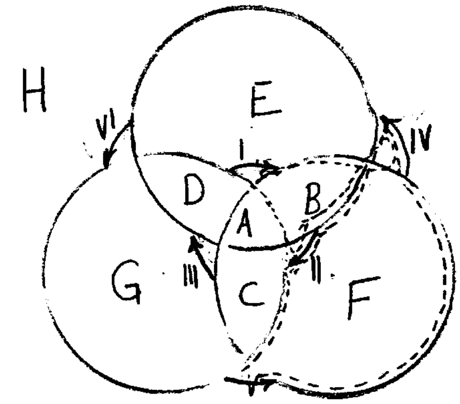

A spanning 2-complex

Here is the corresponding gluing pattern of two octahedra. Faces are glued to their corresponding faces with 120° rotations, alternating in directions like gears.

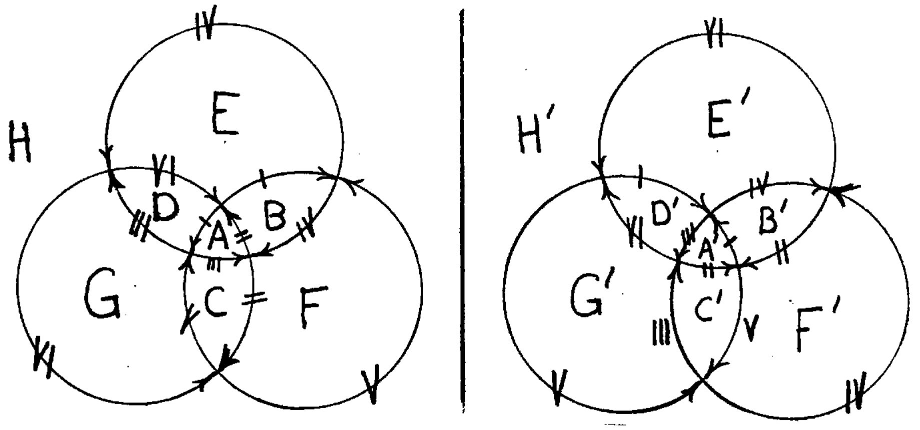

3.12

3.5. The developing map.

Let X be any real analytic manifold, and G a group of real analytic diffeomorphisms of X. Then an element of G is completely determined by its restriction to any open set of X.

Suppose that M is any (G, X)-manifold. Let $U_1, U_2, \dots$ be coordinate charts for M, with maps $\phi_i : U_i \to X$ and transition functions γ_{ij} satisfying

$$\gamma_{ij} \circ \phi_i = \phi_j.$$

In general the γ_{ij}'s are local G-diffeomorphisms of X defined on $\phi_i(U_i \cap U_j)$ so they are determined by locally constant maps, also denoted γ_{ij}, of $U_i \cap U_j$ into G.

Consider now an analytic continuation of ϕ_1 along a path α in M beginning in U_1. It is easy to see, inductively, that on a component of $\alpha \cap U_i$, the analytic continuation of ϕ_1 along α is of the form $\gamma \circ \phi_i$, where $\gamma \in G$. *Hence, ϕ_1 can be analytically continued along every path in M.* It follows immediately that there is a global analytic continuation of ϕ_1 defined on the universal cover of M. (Use the definition of the universal cover as a quotient space of the paths in M.) This map,

$$D : \tilde{M} \to X,$$

is called the *developing map.* D is a local (G, X)-homeomorphism (i.e., it is an immersion inducing the (G, X)-structure on $\tilde{M}$.) D is clearly unique up to composition 3.13
with elements of G.

Although G acts transitively on X in the cases of primary interest, this condition is not necessary for the definition of D. For example, if G is the trivial group and X is closed then closed (G, X)-manifolds are precisely the finite-sheeted covers of X, and D is the covering projection.

From this uniqueness property of D, we have in particular that for any covering transformation T_α of $\tilde{M}$ over M, there is some (unique) element $g_\alpha \in G$ such that

$$D \circ T_\alpha = g_\alpha \circ D.$$

Since $D \circ T_\alpha \circ T_\beta = g_\alpha \circ D \circ T_\beta = g_\alpha \circ g_\beta \circ D$ it follows that the correspondence

$$H : \alpha \mapsto g_\alpha$$

is a homomorphism, called the *holonomy* of M.

In general, the holonomy of M need not determine the (G, X)-structure on M, but there is an important special case in which it does.

DEFINITION. M is a *complete* (G, X)-manifold if $D : \tilde{M} \to X$ is a covering map. (In particular, if X is simply-connected, this means D is a homeomorphism.)

If X is similarly connected, then any complete (G, X)-manifold M may easily be reconstructed from the image $\Gamma = H(\pi_1(M))$ of the holonomy, as the quotient space X/Γ.

Here is a useful sufficient condition for completeness.

PROPOSITION 3.6. *Let G be a group of analytic diffeomorphisms acting transitively on a manifold X, such that for any $x \in X$, the isotropy group G_x of x is compact. Then every closed (G, X)-manifold M is complete.*

PROOF. Let Q be any positive definite on the tangent space $T_x(X)$ of X at some point x. Average the set of transforms $g(Q), g \in G_x$, using Haar measure, to obtain a quadratic form on $T_x(X)$ which is invariant under G_x. Define a Riemannian metric 3.14
$(ds^2)_y = g(Q)$ on X, where $g \in G$ is any element taking x to y. This definition is independent of the choice of g, and the resulting Riemannian metric is invariant under G.

Therefore, this metric pieces together to give a Riemannian metric on any (G, X)-manifold, which is invariant under any (G, X)- map.

If M is any closed (G, X)-manifold, then there is some $\epsilon > 0$ such that the ϵ-ball in the Riemannian metric on M is always convex and contractible. If x is any point in X, then $D^{-1}(B_{\epsilon/2}(x))$ must be a union of homeomorphic copies of $B_{\epsilon/2}(x)$ in $\tilde{M}$. D evenly covers X, so it is a covering projection, and M is complete. $\square$

For example, any closed elliptic three-manifold has universal cover S^3, so any simply-connected elliptic manifold is S^3. Every closed hyperbolic manifold or Euclidean manifold has universal cover hyperbolic three-space or Euclidean space. Such manifolds are consequently determined by their holonomy. 3.15

Even for G and X as in proposition 3.6, the question of whether or not a non- 3.16
compact (G, X)-manifold M is complete can be much more subtle. For example, consider the thrice-punctured sphere, which is obtained by gluing together two triangles minus vertices in this pattern:

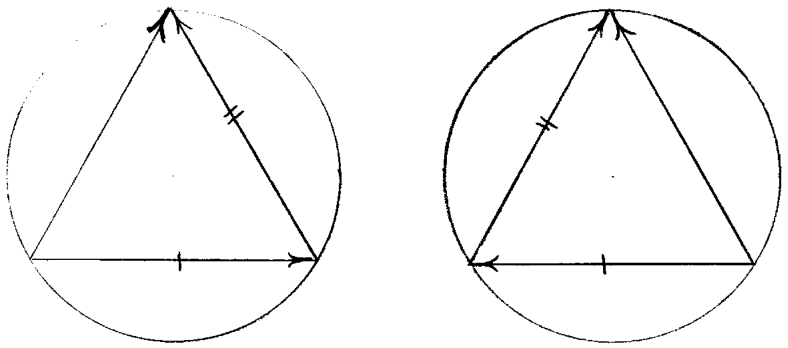

A hyperbolic structure can be obtained by gluing two ideal triangles (with all vertices on the circle at infinity) in this pattern. Each side of such a triangle is isometric to the real line, so a gluing map between two sides may be modified by an arbitrary translation; thus, we have a family of hyperbolic structures in the thrice-punctured sphere parametrized by $\mathbb{R}^3$. (These structures need not be, and are not, all distinct.) *Exactly one parameter value yields a complete hyperbolic structure*, as we shall see presently.

Meanwhile, we collect some useful conditions for completeness of a (G, X)-structure with (G, X) as in 3.6. For convenience, we fix some natural metrics on (G, X)- 3.17
structures.

Labelled 3.7.p to avoid multiple labels

Proposition 3.7. *With (G, X) as above, a (G, X)-manifold M is complete if and only if any of the following equivalent conditions is satisfied.*

(a) *M is complete as a metric space.*
(b) *There is some $\epsilon > 0$ such that each closed ϵ-ball in M is compact.*

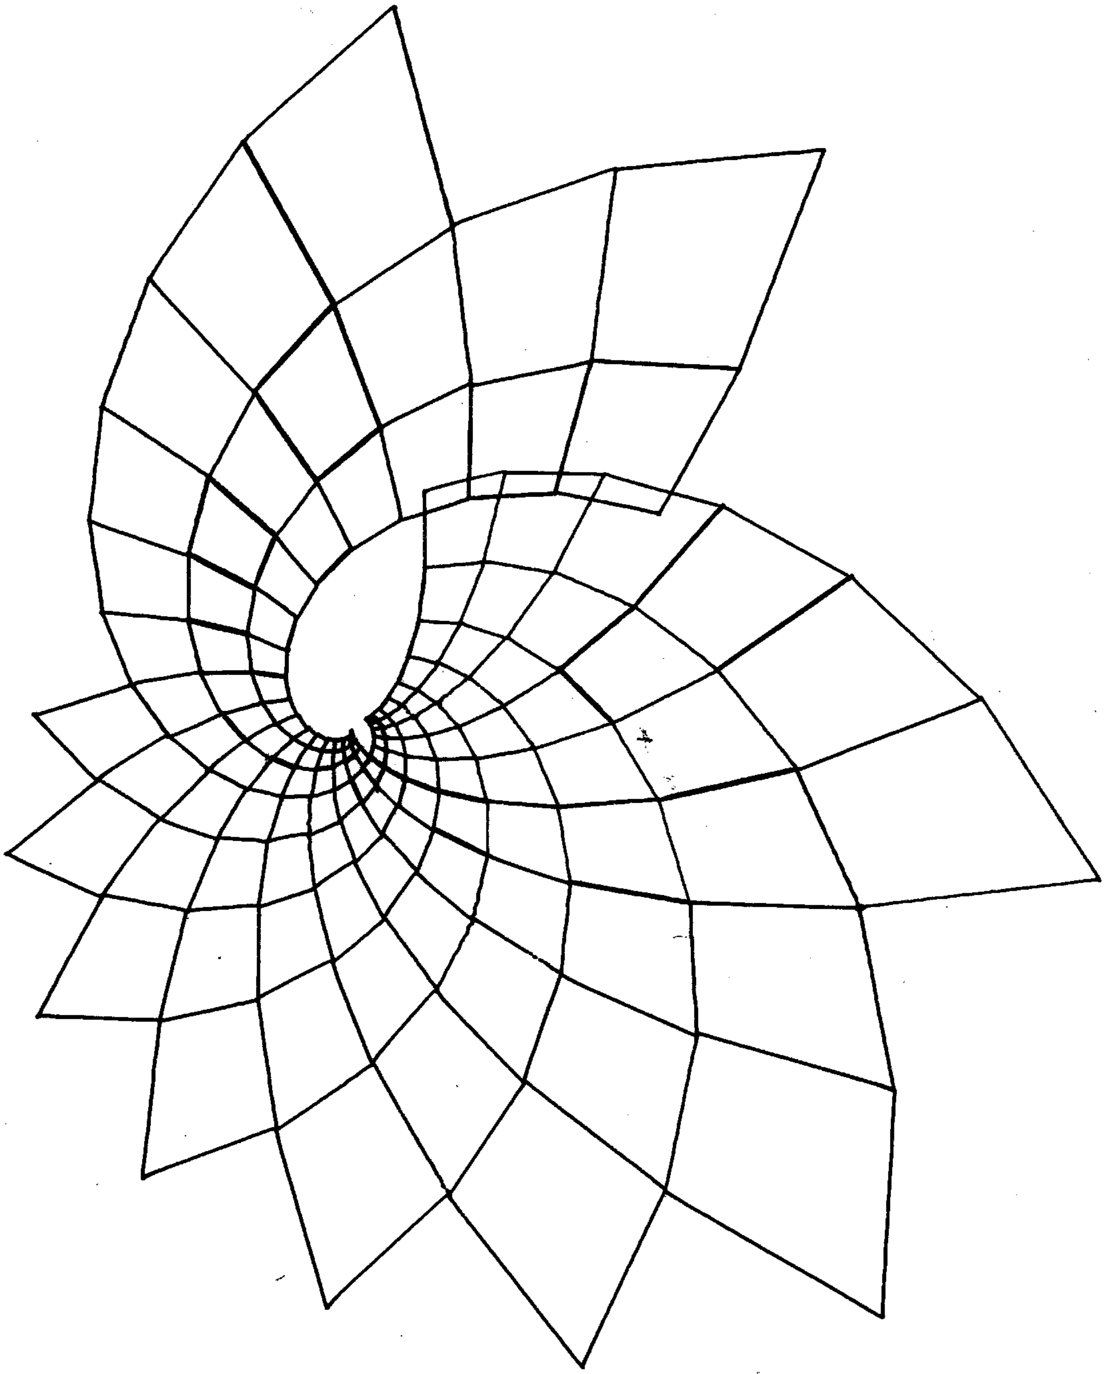

The developing map of an affine torus constructed from a quadrilateral (see p. 3.3). The torus is plainly not complete. *Exercise*: construct other affine toruses with the same holonomy as this one. (Hint: walk once or twice around this page.)

(c) *For every $k > 0$, all closed k-balls are compact.*
(d) *There is a family $\{S_t\}; t \in \mathbb{R}$, of compact sets which exhaust M, such that S_{t+a} contains a neighborhood of radius a about S_t.*

Proof. Suppose that M is metrically complete. Then $\tilde{M}$ is also metrically complete. We will show that the developing map $D : \tilde{M} \to X$ is a covering map by proving that any path α_t in X can be lifted to $\tilde{M}$. In fact, let $T \subset [0,1]$ be a maximal connected set for which there is a lifting. Since D is a local homeomorphism, T is open, and because $\tilde{M}$ is metrically complete, T is closed: hence, α can be lifted, so M is complete.

It is an elementary exercise to see that (b) $\iff$ (c) $\iff$ (d) $\implies$ (a). For any point $x_0 \in \tilde{X}$ there is some ϵ such that the ball $B_\epsilon(x)$ is compact; this ϵ works for all $x \in \tilde{X}$ since the group $\tilde{G}$ of (G, X)-diffeomorphisms of $\tilde{X}$ is transitive. Therefore X satisfies (a), (b), (c) and (d). Finally if M is a complete (G, X)-manifold, it is covered by $\tilde{X}$, so it satisfies (b). The proposition follows. $\square$

3.18

3.8. Horospheres.

To analyze what happens near the vertices of an ideal polyhedron when it is glued together, we need the notion of horospheres (or, in the hyperbolic plane, they are called horocycles.) A horosphere has the limiting shape of a sphere in hyperbolic space, as the radius goes to infinity. One property which can be used to determine the spheres centered at a point X is the fact that such a sphere is orthogonal to all lines through X. Similarly, if X is a point on the sphere at infinity, the *horospheres* “centered” at X are the surfaces orthogonal to all lines through X. In the Poincaré disk model, a hyperbolic sphere is a Euclidean sphere in the interior of the disk, and a horosphere is a Euclidean sphere tangent to the unit sphere. The point X of tangency is the center of the horosphere.

Concentric horocycles and orthogonal lines.

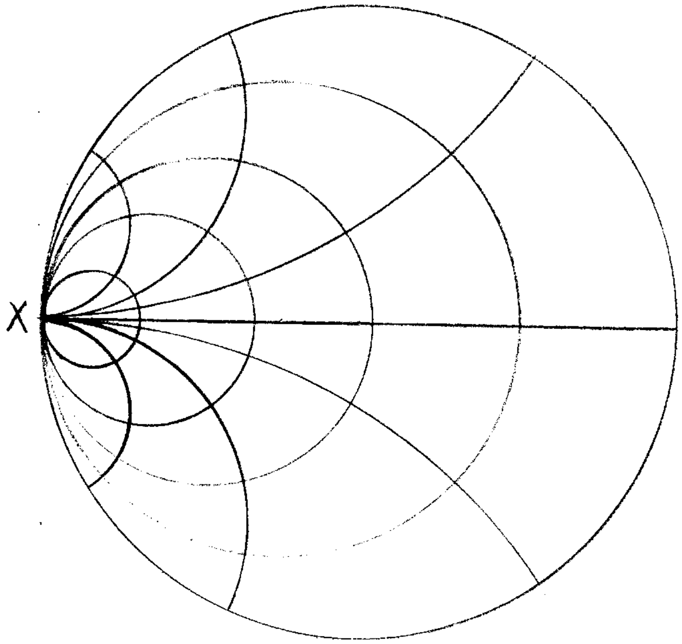

3.19

Translation along a line through X permutes the horospheres centered at X. Thus, *all horospheres are congruent.* The convex region bounded by a horosphere is a *horoball.* For another view of a horosphere, consider the upper half-space model. In this case, hyperbolic lines through the point at infinity are Euclidean lines orthogonal to the plane bounding upper half-space. A horosphere about this point is a horizontal Euclidean plane. From this picture one easily sees that a horosphere in H^n is isometric to Euclidean space E^{n-1}. One also sees that the group of hyperbolic isometries fixing the point at infinity in the upper half-space model acts as the group of similarities of the bounding Euclidean plane. One can see this action internally as follows. Let X be any point at infinity in hyperbolic space, and h any horosphere centered at X. An isometry g of hyperbolic space fixing X takes h to a concentric horosphere h'. Project h' back to h along the family of parallel lines through X. The composition of these two maps is a similarity of h.

Consider two directed lines l_1 and l_2 emanating from the point at infinity in the upper half-space model. Recall that the hyperbolic metric is $ds^2 = (1/x_n^2)\, dx^2$. This means that the hyperbolic distance between l_1 and l_2 along a horosphere is inversely proportional to the Euclidean distance above the bounding plane. The hyperbolic distance between points X_1 and X_2 on l_1 at heights of h_1 and h_2 is $|\log(h_2) - \log(h_1)|$. It follows that for any two concentric horospheres h_1 and h_2 which are a distance d apart, and any pair of lines l_1 and l_2 orthogonal to h_1 and h_2, the ratio of the distance 3.20

FIGURE 1. Horocycles and lines in the upper half-plane

between l_1 and l_2 measured along h_1 to their distance measured along h_2 is $\exp(d)$.

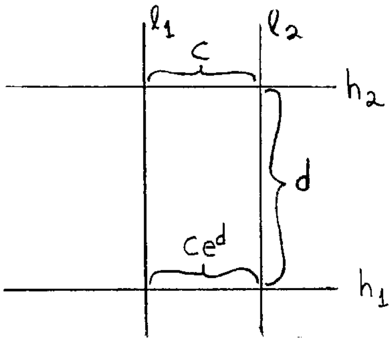

3.9. Hyperbolic surfaces obtained from ideal triangles.

Consider an oriented surface S obtained by gluing ideal triangles with all vertices at infinity, in some pattern. *Exercise: all such triangles are congruent.* (Hint: you can derive this from the fact that a finite triangle is determined by its angles—see 2.6.8. Let the vertices pass to infinity, one at a time.)

Let K be the complex obtained by including the ideal vertices. Associated with each ideal vertex v of K, there is an invariant $d(v)$, defined as follows. Let h be a horocycle in one of the ideal triangles, centered about a vertex which is glued to v and "near" this vertex. Extend h as a horocycle in S counter clockwise about v. It meets each successive ideal triangle as a horocycle orthogonal to two of the sides, until finally it re-enters the original triangle as a horocycle h' concentric with h, at a distance $\pm d(v)$ from h. The sign is chosen to be positive if and only if the horoball bounded by h' in the ideal triangle contains that bounded by h. 3.21

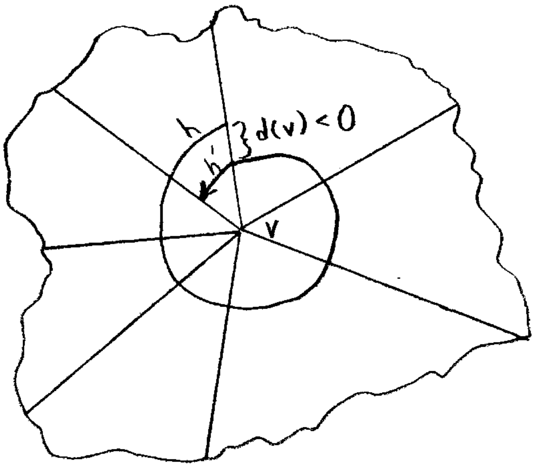

The surface S is complete if and only if all invariants $d(v)$ are 0. Suppose, for instance, that some invariant $d(v) < 0$. Continuing h further round v; the length of each successive circuit around v is reduced by a constant factor < 1, so the total length of h after an infinite number of circuits is bounded. A sequence of points evenly spaced along h is a non-convergent Cauchy sequence.

If all invariants $d(v) = 0$, on the other hand, one can remove horoball neighborhoods of each vertex in K to obtain a compact subsurface S_0. Let S_t be the surface obtained by removing smaller horoball neighborhoods bounded by horocycles a distance of t from the original ones. The surfaces S_t satisfy the hypotheses of 3.7(d) 1—hence S is complete.

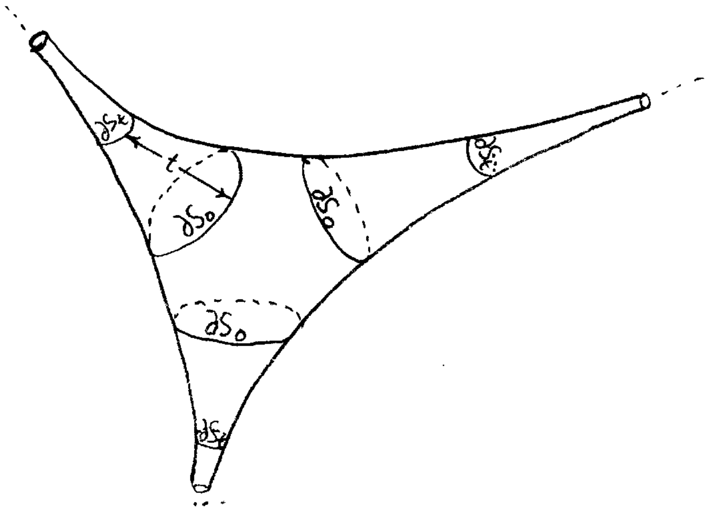

3.22

For any hyperbolic manifold M, let $\bar{M}$ be the metric completion of M. In general, $\bar{M}$ need not be a manifold. However, if S is a surface obtained by gluing ideal hyperbolic triangles, then $\bar{S}$ is a hyperbolic surface with geodesic boundary. There is one boundary component of length $|d(v)|$ for each vertex v of K such that $d(v) \neq 0$. $\bar{S}$ is obtained by adjoining one limit point for each horocycle which "spirals toward" a vertex v in K. The most convincing way to understand $\bar{S}$ is by studying the picture:

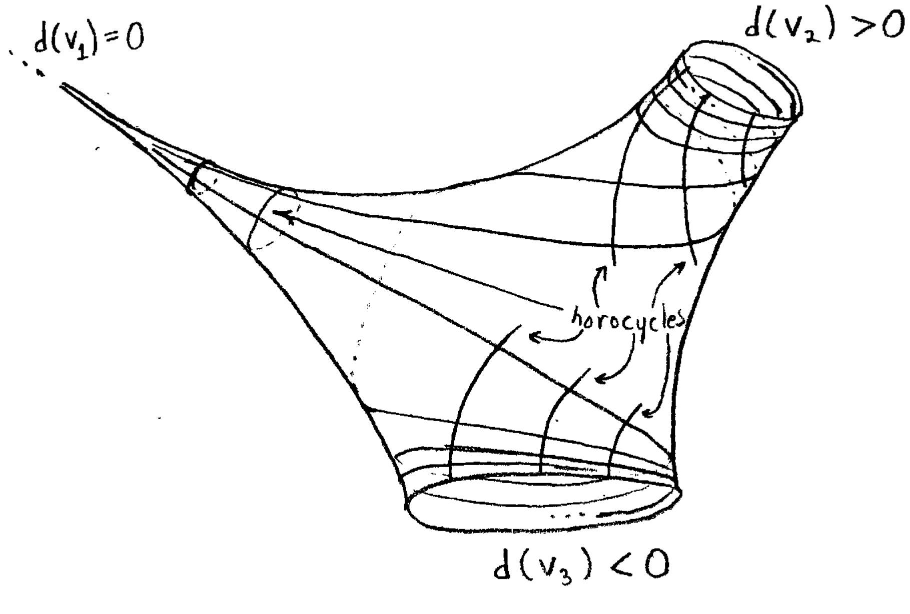

3.10. Hyperbolic manifolds obtained by gluing ideal polyhedra.

Consider now the more general case of a hyperbolic manifold M obtained by gluing together the faces of polyhedra in H^n with some vertices at infinity. Let K 3.23
be the complex obtained by including the ideal vertices. The link of an ideal vertex v is (by definition) the set $L(v)$ of all rays through that vertex. From 3.7 it follows that the link of each vertex has a canonical (similarities of E^{n-1}, E^{n-1}) structure, or similarity structure for short. An extension of the analysis in 3.9 easily shows that *M is complete if and only if the similarity structure on each link of an ideal vertex is actually a Euclidean structure*, or equivalently, if and only if the holonomy of these similarity structures consists of isometries. We shall be concerned mainly with dimension $n = 3$. It is easy to see from the Gauss-Bonnet theorem that any similarity two-manifold has Euler characteristic zero. (Its tangent bundle has a flat orthogonal connection). Hence, if M is oriented, each link $L(v)$ of an ideal vertex is topologically a torus. If $L(v)$ is not Euclidean, then for some $\alpha \in \pi_1 L(v)$, the holonomy $H(\alpha)$ is a contraction, so it has a unique fixed point x_0. Any other element $\beta \in \pi_1(L(v))$ must also fix x_0, since β commutes with α. Translating x_0 to 0, we see that the similarity two-manifold $L(v)$ must be a $(\mathbb{C}^*, \mathbb{C} - 0)$-manifold where $\mathbb{C}^*$ is the multiplicative group of complex numbers. (Compare p. 3.15.) Such a structure

is automatically complete (by 3.6), and it is also modelled on

$$(\tilde{\mathbb{C}}^*, \widetilde{\mathbb{C} - 0}),$$

or, by taking logs, on $(\mathbb{C}, \mathbb{C})$. Here the first $\mathbb{C}$ is an additive group and the second $\mathbb{C}$ is a space. Conversely, by taking exp, any $(\mathbb{C}, \mathbb{C})$ structure gives a similarity structure. $(\mathbb{C}, \mathbb{C})$ structures on closed oriented manifolds are easy to describe, being determined by their holonomy, which is generated by an arbitrary pair (z_1, z_2) of 3.24
complex numbers which are linearly independent over $\mathbb{R}$.

We shall return later to study the spaces $\bar{M}$ in the three-dimensional case. They are sometimes *closed* hyperbolic manifolds obtained topologically by replacing neighborhoods of the vertices by solid toruses.