Thurston's Eight Model Geometries

Nachiketa Adhikari

May 2016

Abstract

Just like the distance in Euclidean space, we can assign to a manifold a notion of distance. Two manifolds with notions of distances can be topologically homoemorphic, but geometrically very different. The question then arises: are there any “standard” or “building block” manifolds such that any manifold is either a quotient of these manifolds or somehow “made up” of such quotients? If so, how do we define such standard manifolds, and how do we discover how many of them there are?

In this project, we ask, and partially answer, these questions in three dimensions. We go over the material necessary for understanding the proof of the existence and sufficiency of Thurston’s eight three-dimensional geometries, study a part of the proof, and look at some examples of manifolds modeled on these geometries. On the way we try to throw light on some of the basic ideas of differential and Riemannian geometry.

Done as part of the course *Low Dimensional Geometry and Topology*, under the supervision of Dr Vijay Ravikumar. The main sources used are [7] and [6].

Contents

1 Preliminaries **2**
- 1.1 Basics . . . 2
- 1.2 Foliations . . . 3
- 1.3 Bundles . . . 4

2 Model Geometries **8**
- 2.1 What is a model geometry? . . . 8
- 2.2 Holonomy and the developing map . . . 9
- 2.3 Compact point stabilizers . . . 11

3 Thurston's Theorem **14**
- 3.1 In two dimensions . . . 14
- 3.2 In three dimensions . . . 14
  - 3.2.1 Group discussion . . . 15
  - 3.2.2 The theorem . . . 16

4 The Eight Geometries **22**
- 4.1 $S^2 \times \mathbb{R}$ . . . 22
- 4.2 $\mathbb{H}^2 \times \mathbb{R}$ . . . 23
- 4.3 $\widetilde{SL_2}(\mathbb{R})$ . . . 23
- 4.4 *Nilgeometry* . . . 24
- 4.5 *Solvegeometry* . . . 25

5 Conclusion **26**

Bibliography **27**

Chapter 1

Preliminaries

In this chapter we provide some basic definitions and prove some foundational results which will be needed later on.

1.1 Basics

Definition 1.1. If G is a group acting on a manifold X via diffeomorphisms, then a manifold M is called a **(G, X)-manifold** if

1. there is an open cover $\{U_\alpha\}$ of M and a family $\{\phi_\alpha : U_\alpha \to V_\alpha\}$ of diffeomorphisms onto open sets $V_\alpha \subset X$, and

2. if $U_\alpha \cap U_\beta \neq \phi$, then there exists a $g \in G$ such that $g \cdot x = \phi_\alpha \circ \phi_\beta^{-1}(x)$ for all $x \in V_\alpha \cap V_\beta$, which is to say, each transition map is given by the restriction of an element of G.

Example 1.2. A **Euclidean manifold** is a (G, X)-manifold with $X = \mathbb{R}^n$ and G a subgroup of $\mathsf{Isom}(\mathbb{R}^n)$, the group of isometries of $\mathbb{R}^n$. Similar definitions hold for **elliptic** ($X = S^n$) and **hyperbolic** ($X = \mathbb{H}^n$) manifolds.

Definition 1.3. A **Riemannian metric** on a manifold M is a family of positive definite inner products

$$g_p : T_pM \times T_pM \to \mathbb{R}, \qquad p \in M$$

such that, for all smooth vector fields X, Y on M, the function $p \mapsto g_p(X_p, Y_p)$ is smooth. Through integration, this inner product allows one to define lengths of curves on a manifold and hence endows it with a distance function. The manifold thus obtained is a metric space and is called a **Riemannian manifold**.

Example 1.4. The usual inner product on $\mathbb{R}^n$ is a Riemannian metric which gives rise to a distance function known to us as the Euclidean metric.

The following result will be useful later.

Proposition 1.5. *If $f : M \to N$ is an immersion and N has a Riemannian metric, then that metric can be pulled back to endow M with a Riemannian metric.*

Proof. Suppose μ is the metric on N. Then we define a symmetric bilinear form $f^*\mu$ on M by

$$(f^*\mu)(v, w) = \mu(f_*(v), f_*(w)).$$

This is, in general, not positive definite. But if f is an immersion, f_* has trivial kernel, so it is positive-definite, and $f^*\mu$ is indeed a metric on M. □

We now give an informal idea of the concept of *curvature*. First, we define the curvature of a circle as the inverse of its radius, and that of a line as zero. Then, given a curve on a surface and a point on the curve, there is a unique line or circle on the surface which most closely approximates the curve near that point, namely the osculating circle. The curvature of the curve at that point is defined to be the curvature of that line or circle.
Finally, at any point on a surface, there is a vector normal to the surface (we assume our surface is embedded in some ambient space with an inner product), and various planes containing that vector. The intersection of one such plane with the surface forms a curve, whose curvature is called a *normal curvature*. The maximum and minimum normal curvatures at a point are called the *principal curvatures*, denoted κ_1, κ_2, at that point. We are now ready for

Definition 1.6. The **Gaussian curvature** of a surface at a point is $\kappa_1\kappa_2$, the product of the principal curvatures at that point.

It is clear from the definition that the Gaussian curvature is an invariant of the metric. In fact, if the metric is multiplied by k, the curvature gets multiplied by $1/k^2$.

Definition 1.7. The **sectional curvature** of a Riemannian manifold, defined with respect to the tangent plane at a point, is the Gaussian curvature of the surface which is tangent to that plane. This definition depends on the tangent plane chosen.
We say that a Riemannian manifold has **constant sectional curvature** if the curvature is independent not only of the choice of tangent plane, but also of the choice of the point we choose to compute the curvature at. In this case we have

Theorem 1.8. *The only complete, simply connected Riemannian n-manifolds with constant sectional curvature are $\mathbb{R}^n, S^n$ and $\mathbb{H}^n$.*

A proof of this theorem can be found in [2].

1.2 Foliations

A foliation is a way of splitting up a manifold into lower-dimensional “slices”. We make this precise in

Definition 1.9. A p-dimensional **foliation** $\mathcal{F}$ of an n-dimensional manifold M is a covering $\{U_i\}$ of M along with diffeomorphisms $\phi_i : U_i \to \mathbb{R}^n$ such that the transition functions ϕ_{ij} take the form

$$\phi_{ij}(x, y) = (\phi^1_{ij}(x), \phi^2_{ij}(x, y)),$$

where x denotes the first $n - p$ coordinates and y the last p coordinates, which is to say that ϕ^1_{ij} is a map from $\mathbb{R}^{n-p}$ to itself and ϕ^2_{ij} is a map from $\mathbb{R}^n$ to $\mathbb{R}^p$.

In each chart U_i, the "p-planes" corresponding to the p-planes $\{(x, y) : x = a\}$ (for various fixed constants a) are p-dimensional immersed submanifolds of M. These piece together globally (that is, over all charts) to form maximal connected injectively immersed submanifolds which are called the **leaves** of $\mathcal{F}$.

Another way of thinking of a foliation is as a collection of pairwise disjoint p-dimensional connected immersed submanifolds that cover M, so that each $x \in M$ has a neighborhood U homeomorphic to $\mathbb{R}^n$ to which each leaf is either disjoint, or intersects in subspaces which map to the p-planes of $\mathbb{R}^n$ described above.

Example 1.10. A trivial example of a foliation is of $\mathbb{R}^n$ by the affine subspaces $\mathbb{R}^p$. Thus the lines (planes) $x = a$, where a varies over $\mathbb{R}$, foliate $\mathbb{R}^2$ ($\mathbb{R}^3$).

Example 1.11. The cylinder $S^1 \times \mathbb{R}$, is foliated by the circles transverse to the lines $\{x\} \times \mathbb{R}$. This is a 1-dimensional foliation and each circle is a leaf. The lines $\{x\} \times \mathbb{R}$ are themselves the leaves of another 1-dimensional foliation.

Example 1.12. The above example is in fact an instance of a general fact: a foliation on X descends to any (G, X)-manifold as long as the elements of G take leaves to leaves. So the horizontal lines in the plane descend to a foliation of the torus by the "horizontal" circles (Figure 1.2), the vertical lines descend to "vertical" circles, and lines with irrational slopes are dense in the torus. All three are foliations.

Example 1.13. Any nowhere vanishing vector field has an associated 1-dimensional foliation, the leaves being the flow lines.

Example 1.14. This is a non-example. Concentric circles centered at the origin do not form a foliation of the plane, for the simple reason that, though each is a connected immersed submanifold, the origin does not lie on any such circle, so they don't cover the plane. If we deem the origin to be a leaf by itself, then the foliation wouldn't have a consistent dimension.

1.3 Bundles

Fiber bundles are generalizations of the concepts of vector and tangent bundles that we're familiar with. They look locally like product spaces, up to adjustments by group actions. More formally, we have

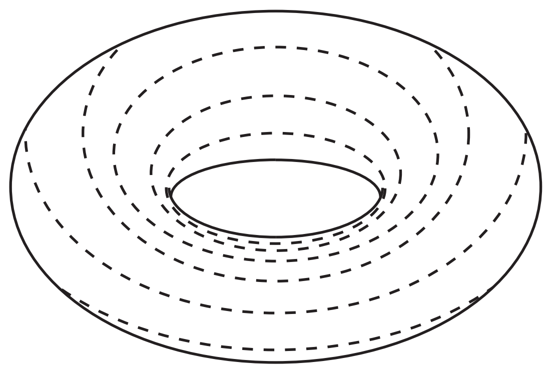

Figure 1.1: The torus foliated by "horizontal" circles

Definition 1.15. A **fiber bundle** (E, B, G, X, π) (also called a *fiber bundle* with structure group G and fiber X) consists of the following:

- A *total space* E
- A *base space* B
- A map $\pi : E \to B$, called the *bundle projection*
- A topological group G with a left action on the space X
- A *local trivialization*, that is, a covering of B by open sets U_i with homeomorphisms $\phi_i : \pi^{-1}(U_i) \to U_i \times X$ such that the following diagram commutes:

$$\begin{array}{ccc} \pi^{-1}(U_i) & \xrightarrow{\phi_i} & U_i \times X \\ & \searrow^{\pi} \quad \swarrow_{(u,x)\mapsto u} & \\ & U_i & \end{array}$$

  Moreover, we require that the transition maps

$$\psi_{ij} = \phi_i \circ \phi_j^{-1} : (U_i \cap U_j) \times X \to (U_i \cap U_j) \times X$$

  are of the form

$$\psi_{ij}(u, x) = (u, g_{ij}(u) \cdot x),$$

where the maps $g_{ij} : U_i \cap U_j \to G$ are continuous, and subject to the *cocycle condition*: $\psi_{ij}\psi_{jk} = \psi_{ik}$ whenever this composition is defined.

A few examples will serve to illuminate and clarify this definition:

Example 1.16. The Möbius strip is a $(\mathbb{Z}_2, \mathbb{R})$-bundle over S^1, which is to say that $B = S^1, G = \mathbb{Z}_2, X = \mathbb{R}, E$ is the strip itself, and π is the usual projection from the strip to the circle. To see this, fix a point of the base space and fix an orientation of the fiber above it. Moving once along the circle and returning to our fixed point, we see that the orientation of the fiber is reversed. So we can cover the circle with two open sets, each homeomorphic to an interval, such that the fiber above each is homeomorphic to the product of an interval and $\mathbb{R}$, and one of the intersections of these open sets has all points having the same orientations of their fibers, while the other intersection sees the fibers over a given point having one orientation in the first open set, and another in the second open set, thus requiring an orientation-reversing action. In other words, $\mathbb{Z}_2$ acts on the fibers.

Example 1.17. As a contrast with the previous example, the cylinder is a fiber bundle with base and fiber equal to those of the Möbius strip, but the group being the trivial group.

Example 1.18. A generalization of the previous example is the **product bundle** or **trivial bundle**, where E is of the form $B \times X$, and G is just the trivial group.

Example 1.19. As mentioned earlier, vector bundles are special cases of fiber bundles. A (G, X)-bundle is a vector bundle if X is a vector space and $G = GL(X)$, its group of linear automorphisms. A futher special case of this is obtained when $X = \mathbb{R}^n$, the fiber above a point being identified with the tangent space at that point and G being identified with the derivatives of diffeomorphisms between open sets of B. This is of course the tangent bundle.

Example 1.20. On a Riemannian manifold M, the subset of the tangent bundle TM consisting of tangent vectors of unit length itself forms a bundle called the **unit tangent bundle** over M, denoted UTM.

Example 1.21. Fiber bundles form an important class of foliations. Given a fiber bundle (E, BG, X, π), where B is n-dimensional and Xm-dimensional, E is an $(m + n)$-dimensional manifold with a foliation whose leaves are the fibers.

Example 1.22. A rather handy class of bundles is obtained via the following construction: suppose M is a smooth manifold and $\phi : M \to M$ a diffeomorphism. Then the **mapping torus** M_ϕ is obtained from the cylinder $M \times [0, 1]$ by identifying the two ends $M \times \{0\}$ and $M \times \{1\}$ via ϕ. Then M_ϕ is an M-bundle over the circle with the action of the group generated by ϕ. The cylinder and Möbius strip are examples of this, with $M = \mathbb{R}$ and ϕ the identity and a reflection respectively, as are the torus and the Klein bottle, where $M = S^1$ and ϕ the identity and the antipodal map respectively.

We now define a very special kind of fiber bundle:

Definition 1.23. A **principal fiber bundle** is a fiber bundle where the fiber is the structure group itself, that is, $X = G$, and the action is the usual left-multiplication.
In this case, however, we also get a right action of G on E, locally given, for an open set $U \subset B$, by

$$(x, h) \mapsto (x, h \cdot g),\ x \in U; g, h \in G.$$

This local action can be glued together to give a global action of G on E since the right action commutes with the left action of G on G. This right action is thus defined naturally on a principal bundle. Note that this is not the same as the left-action: that is an action on the fibers (so, for example, nearby points on different fibers need not map to nearby points under left translation), while the right action is an action on the whole space.

Since in what follows we will have occasion to consider only such situations where spaces are smooth manifolds, maps are smooth and G is a Lie group acting via diffeomorphisms, we now specialize to such a situation.

Definition 1.24. Let (E, B, G, π) be a principal fiber bundle. For $p \in E$, let V_p be the subspace of T_pE consisting of vectors tangent to the fiber through p. A **connection** Γ is an assignment of a subspace H_p of T_pE to each $p \in E$ such that

1. H_p varies smoothly with p
2. $T_pE = H_p \bigoplus V_p$
3. For any $g \in G$, $(R_g)_*(H_p) = H_{p \cdot g}$ (where R_g denotes the right action of G).

The idea is this: assign to each point of E a vector space which is a complement of the tangent space to the fiber such that this assignment is smooth and invariant under the action of G. Thus Γ gives a horizontal subbundle of TE, the tangent bundle of E.
The projection $\pi : E \to B$ induces a linear map $\pi_* : T_pE \to T_{\pi(p)}B$ for each $p \in E$ which restricts to an isomorphism on H_p. Then it is easy to show that, corresponding to each vector field X on E, there exists a unique vector field X^* on E that corresponds to X under this isomorphism. It follows that X^* is horizontal and hence invariant under the right G-action.

Definition 1.25. A **connection form** ω for a connection Γ assigns, to each $p \in E$, the map $\omega_p : T_pE \to \mathbb{R}$ which takes a tangent vector at p to its inner product with a tangent vector X_p in V_p. The choice of such a vector X_p can be made smoothly over E, in such a way that ω is a smooth 1-form and $H_p = \ker \omega_p$ for all p.

Definition 1.26. The **curvature** of a connection ω is defined to be $\mathrm{d}\omega$.

Chapter 2

Model Geometries

We now turn to the definition of a model geometry, and provide some background as to why it is defined the way it is.

2.1 What is a model geometry?

Some manifolds, such as the torus and cylinder, have metrics inherited from the Euclidean space they are quotients of. Some, such as the sphere and projective plane, have no metric compatible with the Euclidean metric, but the projective plane does have a metric inherited from the sphere. This gives rise to the question: can all Riemannian manifolds (and their metrics) be obtained as quotients of some reference manifold(s)? If so, what are the properties these reference manifolds must possess? As a step towards answering these questions, we make

Definition 2.1. A **model geometry** (G, X) is a manifold X together with a Lie group G of diffeomorphisms of X such that

1. X is connected and simply connected;
2. G acts transitively on X with compact point stabilizers;
3. G is not contained in any larger group of diffeomorphisms of X with compact point stabilizers; and
4. there exists at least one compact (G, X)-manifold.

The standard examples of model geometries are $X = \mathbb{R}^n, S^n$ or $\mathbb{H}^n$ with its usual metric, and G its group of isometries.

Since there is no loss of generality in working with connected components individually, we may assume that X itself is connected. Any (G, X)-manifold is a $(G, \tilde{X})$-manifold as well, where $\tilde{X}$ is any covering space of X, so it makes sense to consider the general case where X is its own universal cover, that is,

simply connected. This explains the first requirement.
Similarly, since a (G, X)-manifold satisfying the first two requirements would be a (G', X)-manifold for $G \subset G'$, we are better off studying the most general case: the one where G is the largest such group. This explains the third requirement.
The fourth requirement is just to ensure we get some "interesting" manifolds based on a given geometry. After all, our aim is to classify the manifolds we're already familiar with.
Justifying the second requirement is the purpose of the rest of this chapter. The next section is devoted to developing some ideas we will need later.

2.2 Holonomy and the developing map

Given a group G acting via analytic diffeomorphisms on a space X, and a $(G, X)-$manifold M, we construct a map called the *developing map*, which, under certain nice conditions, gives us a subgroup Γ of G for which M is the quotient space X/Γ.

For the rest of the discussion, we work with a fixed basepoint m_0 of M, and a $(G, X)-$chart (U_0, ϕ_0) whose domain contains m_0. Let $\pi : \tilde{M} \to M$ be the universal cover of M. Then we can think of $\tilde{M}$ as the space of all homotopy classes of paths starting at m_0 and ending at some point of M.

Now suppose $[\sigma] \in \tilde{M}$, that is, σ is a path from m_0 to the point $\pi([\sigma])$ in M. Suppose the sets $U_0, U_1, \ldots, U_k$ cover this path, with $\pi([\sigma]) \in U_k$. It can be arranged that $U_i \cap U_j \neq \phi$ if and only if j and i are consecutive. On $U_0 \cap U_1$, the transition map $\phi_0 \circ \phi_1^{-1}$ is given by some element $g_1 \in G$. Since G acts via real analytic diffeomorphisms, this g_1 is unique. Similarly, on $U_1 \cap U_2$, the transition map $\phi_1 \circ \phi_2^{-1}$ is given by some element $g_2 \in G$ and so on till we get elements $g_1, g_2, \cdots g_k \in G$. So we get a well-defined smooth map ϕ on $U = U_0 \cup \cdots \cup U_k$, called the *analytic continuation* of ϕ_0 along σ:

$$\phi : U_0 \cup \cdots \cup U_k \to X$$
$$m \mapsto \begin{cases} \phi_0(m), & m \in U_0 \\ g_1 \cdot \phi_1(m), & m \in U_1 \\ \vdots & \\ g_1 \cdot g_2 \cdots g_k \cdot \phi_k(m), & m \in U_k \end{cases}$$

Since σ is a path in M and ϕ defines a smooth map from M to X, the image of σ in X is a path in X. But we have something stronger, namely that, the neighborhood U of σ maps to a neighborhood V of $\phi(\sigma)$ in X. If we have a path homotopy between σ and some path σ' which lies entirely inside U, then ϕ will take σ' and the homotopy to V and therefore $\phi(\sigma)$ and $\phi(\sigma')$ will be homotopic in X as well. In particular, they will also have the same endpoint. This argument can be extended to all paths homotopic to σ, so the map $g_1 \cdot g_2 \cdots g_k \cdot \phi_k$ does not

depend on the choice of a path in $[\sigma]$, and we can define $\phi_0^{[\sigma]} = g_1 \cdot g_2 \cdots g_k \cdot \phi_k$ unambiguously.
We now show that the map $\phi_0^{[\sigma]}$ does not depend on the choices $U_1, \ldots U_k$ made on the way. Suppose $(U_1', \phi_1'), \ldots, (U_k', \phi_k')$ is another choice of charts and the resulting map is ϕ'. Then both ϕ and ϕ' agree on U_0. Now suppose they agree up to some point m of σ. The open sets U_i and U_j' which contain m have nonempty intersection, and $(U_i, \phi|_{U_i})$ and $(U_j', \phi'|_{U_j'})$ are charts, so that $\phi' \circ \phi^{-1}$ is given by some $g \in G$. But ϕ and ϕ' agree on some neighborhood of the part of σ before m, so $\phi' \circ \phi^{-1}$ must coincide with the identity of G in some neighborhood of $\phi(m) = \phi'(m)$. Since G acts by real analytic diffeomorphisms, g is the identity and we have $\phi = \phi'$ in some neighborhood of m. Using connectedness of σ, it can easily be shown that they agree on the whole path, and it follows that $\phi_0^{[\sigma]}$ does not depend on the choice of charts.

The upshot of the preceding discussion is that, corresponding to each homotopy class of paths $[\sigma]$ in M, we get a unique point $x_{[\sigma]}$ of X (that is, the image of the endpoint of σ under $\phi_0^{[\sigma]}$), which varies smoothly with $[\sigma]$. The definition of the **developing map** $D : M \to X$ then follows: it is the map that takes $[\sigma]$ to $x_{[\sigma]}$ of X. It is clear that this is a smooth map, and also that, in a neighborhood of $[\sigma]$, we have a commutative diagram

$$\begin{array}{ccc} \tilde{M} & \xrightarrow{\pi} & M \\ & {\scriptstyle D}\searrow \qquad \swarrow{\scriptstyle \phi_0^{[\sigma]}} & \\ & X & \end{array}$$

Given a smooth map $f : N \to M$ of manifolds, it is easy to see that a $(G, X)-$structure on M can be pulled back to one on N by pulling back the charts on M via f. So, via $\pi : \tilde{M} \to M$, $\tilde{M}$ acquires a $(G, X)-$ structure, and D is a local $(G, X)-$ diffeomorphism.

Now suppose $[\sigma] \in \pi_1(M)$. Then ϕ_0 and $\phi_0^{[\sigma]}$ have nonempty intersection (the basepoint m_0 must lie in both), hence $\phi_0^{[\sigma]} = g_{[\sigma]} \cdot \phi_0$ for some $g_{[\sigma]} \in G$. This $g_{[\sigma]}$ is defined to be the **holonomy** of $[\sigma]$. The map $[\sigma] \mapsto g_{[\sigma]}$ is a homomorphism from $[\sigma] \in \pi_1(M)$ into G, whose image is called the **holonomy group** Γ of M.

Definition 2.2. M is called a **complete** $(G, X)-$**manifold** if the deveoping map D is a covering map.

If X is simply connected, we obtain that D is a diffeomorphism because a simply connected space is homeomorphic to any cover. We can then assume that X itself is the universal cover of M, and obtain the following result (which is useful but we will not have much occasion to use)

Proposition 2.3. *If G is a group of analytic diffeomorphisms of a simply connected space X, then any complete $(G, X)-$manifold may be reconstructed from its holonomy group Γ as X/Γ.*

Proof. We use the standard result from algebraic topology: any space is the quotient of its universal cover by the group of deck transformations. If we can show that Γ coincides with the group of deck transformations A, we would be done. To show this, it is enough to show that Γ preserves fibers (so $\Gamma \subset A$), and acts transitively on some fiber (so $A \subset \Gamma$).
The latter is straightforward: take the fiber of x_0 itself. Any point $\tilde{x_0}$ in the fiber is the image of x_0 under $\phi_0^{[\sigma]}$ for some σ, and since $\phi_0^{[\sigma]} = g_{[\sigma]} \cdot \phi_0$, we have that $\tilde{x_0} = g_\sigma \cdot \phi_0(x_0)$, so the action of Γ on this fiber is transitive.
To answer the former, we ask ourselves: what is the fiber of a point $x \in M$? It is the set of points in X which are endpoints of lifts of paths from x_0 to x. Suppose τ is a path from x_0 to x whose lift in X is $\tilde{\tau}$. Suppose σ is a loop at x_0 whose lift is $\tilde{\sigma}$. Note that $\tilde{\sigma}$ need not be a loop in X. The endpoints of $\tilde{\sigma\tau}$ (the lift of $\sigma\tau$) and $\tilde{\tau}$ are two points in the fiber of x. If we apply $g_{[\sigma]}$ to $\tilde{\tau}$, we get a path $\tilde{\tau}'$ starting at the endpoint of $\tilde{\sigma}$ such that $\tilde{\sigma}\tilde{\tau}'$ is a lift of $\sigma\tau$. Since any path has a unique lift, the endpoints of $\tilde{\sigma\tau}$ and $\tilde{\sigma}\tilde{\tau}'$ must coincide, that is, $g_{[\sigma]}$ acts on the endpoint of $\tilde{\tau}$ to give the endpoint of $\tilde{\sigma\tau}$, which is in the same fiber. Since τ and σ were arbitrary, it follows that Γ preserves fibers, and the proof is complete. □

2.3 Compact point stabilizers

We will use the ideas developed in the previous section to understand why we require point stabilizers to be compact in Definition 2.1. One of the reasons is so that X has a G-invariant Riemannian metric:

Lemma 2.4. *Let G act transitively on a manifold X. Then X admits a G-invariant Riemannian metric if and only if, for some $x \in X$, the image of G_x in $GL(T_xX)$ has compact closure. (G_x denotes the stabilizer of x in G.)*

Proof. Since a Riemannian metric is a non-degenerate symmetric bilinear form, then for G to preserve it, G_x must map to a subgroup of $O(T_xX)$ (the orthogonal matrices; the map here is $g \mapsto (L_g)_*$, where L_g denotes left-multiplication by g). Since $O(T_xX)$ is compact, the image of G_x under this map has compact closure.

Conversely, suppose G_x has compact closure H_x under this map. Any compact group can be given a finite measure, called the *Haar measure*, that is invariant under left or right translations of the group. Let $<,>$ be any positive definite symmetric bilinear form on T_xX, and define

$$(u, v) = \int_{H_x} < gu, gv > dg, \qquad u, v \in T_xX$$

where dg is the Haar measure on H_x. Then (u, v) is an inner product on T_xX invariant under the action of G_x. Since G acts transitively, this inner product can be propagated to T_yX for any $y \in X$, and we thus get a G-invariant Riemannian metric on X. □

Further explanation of the requirement of compact point stabilizers will need some exposition.

Definition 2.5. A **geodesic** is a curve $\gamma : [0,1] \to M$ such that, for x, y close enough in $[0,1]$, the shortest path joining $\gamma(x)$ and $\gamma(y)$ in M coincides with $\gamma([x,y])$.

From the result on existence and uniqueness of the solution of a second-order ordinary differential equation, we know that, given $v \in T_pM$, there is a unique geodesic γ_v such that $\gamma_v(0) = p$ and $\gamma'_v(0) = v$. Here γ_v is a map from (a,b) to M, where a, b depend on M, v and p, and could even be $\pm\infty$.

Definition 2.6. The **exponential map** of γ_v is defined by

$$\mathsf{exp}_p(v) = \gamma_v(1).$$

This map need not be defined for a given pair p, v.

Since exp varies smoothly based on p and v, it can be shown that exp is well-defined on some open set containing 0 in T_pM. It can then be shown that $(\mathsf{exp}_p)_*|_0$ is the identity on a small neighborhood of 0, hence that, for any $p \in M$, there is a neighborhood V of 0 in T_pM and a neighborhood U of p in M such that $\mathsf{exp}_p : V \to U$ is a diffeomorphism. It follows that, given p, there exists an $\epsilon > 0$ such that if $d(p,q) < \epsilon$, then there is a unique geodesic joining p and q inside the ball $B(p,\epsilon)$. (We say, in this case, that $B(p,\epsilon)$ is *ball-like and convex.* Detailed proofs of the foregoing statements may be found in [3]. We are now ready for

Proposition 2.7. *Let G be a Lie group acting transitively on X such that G_x is compact for some (hence all) $x \in X$. Then every closed (that is, compact) (G,X)-manifold M is complete.*

Proof. By Lemma 2.4, X has a G-invariant Riemannian metric. Let M be a closed (G,X)-manifold. Using local charts and partitions of unity, as well as Proposition 1.5, this metric can be pulled back to give M a Riemannian metric which is invariant under any (G,X)-map. This metric can further be pulled back to give $\tilde{M}$, the universal cover of M, a Riemannian metric. By the preceding paragraph, we can get, around each $p \in M$, an ϵ_p-neighborhood that is ball-like and convex. By compactness of M, we can find an ϵ that works for all $p \in M$. Further, since G acts transitively on X, we can assume all ϵ-neighborhoods in X are ball-like and convex.

Take $y \in \tilde{M}$. The developing map D maps $B(y,\epsilon)$ isometrically onto $B(D(y),\epsilon)$: it is enough to show that it is an isometry and that it is injective. It is an isometry because the metric on $\tilde{M}$ is defined as the pullback of the metric on X under D; it is injective because, if $D(y) = D(y')$, then the geodesic from y to y' maps to a geodesic through $D(y)$ and $D(y')$ in X. Since these are the same point , and since any two points have a unique geodesic in X, we must have $y = y'$.

Suppose $x \in X$ and $D(y) \in B(x, \epsilon/2)$. Since $B(y, \epsilon)$ is isometric to a ball containing $B(x, \epsilon/2)$, it must contain a homeomorphic copy of $B(x, \epsilon/2)$. It follows that $B(y, \epsilon)$ is a disjoint union of such homeomorphic copies, that is, D evenly covers X. Therefore M is complete. $\square$

Chapter 3

Thurston's Theorem

In this chapter we state the main result and provide a partial proof. First, we deal with the two-dimensional version.

3.1 In two dimensions

Theorem 3.1. *There are exactly three two-dimensional model geometries:* $\mathbb{R}^2, S^2$ *and* $\mathbb{H}^2$*.*

Proof. Let (G, X) be a two-dimensional model geometry. Since we're working in two dimensions, there is exactly one tangent plane at each point and the sectional curvature is the Gaussian curvature. From the definition of a model geometry, we know that G acts transitively on X and that the metric on X is G-invariant. By transitivity, G takes the curvature at a point to the curvature at any other point, and since the curvature is an invariant of the metric, it must be the same at all points. By Theorem 1.8, X can only be one of $\mathbb{R}^2, S^2$ and $\mathbb{H}^2$ □

So in two dimensions, things are relatively straightforward. In three dimensions, however, we already see that, for example, $S^2 \times \mathbb{R}$ is a candidate for a model geometry which is not homeomorphic to either of $\mathbb{R}^3, S^3$ or $\mathbb{H}^3$. In fact, there are eight model geometries in three dimensions, and that is the content of the main theorem.

3.2 In three dimensions

Let (G, X) be a model three-dimensional geometry. We first embark on a discussion of the group G.

3.2.1 Group discussion

The whole point of having a group acting on our manifold X is that studying the group will give us information about X itself. Accordingly, we first take a look at the connected component of the identity of G. Call it G'.

Lemma 3.2. *The action of G' is transitive.*

Proof. Recall that the stabilizer of a point $x \in X$ is called G_x. We claim that it is enough to show that, for an arbitrary $x \in X$, $G'G_x = G$. For then, if $y \in X$ is any other point, and $g \cdot x = y$ for some $g \in G$ (the action of G is transitive), then g would be of the form $g'h$ for some $g' \in G'$ and $h \in G_x$. This would imply that $y = g \cdot x = g'h \cdot x = g' \cdot x$, that is, that there is an element in G' taking x to y. Since x and y were arbitrary, this would complete the proof.

So it remains to show that $G'G_x = G$. Since G acts transitively with stabilizer G_x, G/G_x is diffeomorphic to X. G', being the connected component of the identity, is an open normal subgroup of G, hence $G'G_x$ is itself an open subgroup of G. Its image under the open map $G \to G/G_x$ is an open set, as are the images of its cosets in G. We thus get a collection of disjoint open sets in G/G_x. However, this is diffeomorphic to X, which is connected. It follows that there is actually just one such open set. That is, $G'G_x = G$. ☐

Lemma 3.3. *G'_x is connected for any x.*

Proof. G'_x is the stabilizer of x in G'. Let $(G'_x)_0$ be the connected component of the identity in G'_x. Proving the lemma is equivalent to proving that $G'_x = (G'_x)_0$. By the previous lemma, the action of G' is transitive, so we have $G'/G'_x \simeq X$ for any $x \in X$. Instead of G'_x, if we quotient G' by $(G'_x)_0$, a subgroup of G'_x, the resulting map into X might not be injective, but it will still be surjective. We therefore have a projection $\pi : G'/(G'_x)_0 \to X$ with fiber $G'_x/(G'_x)_0$. Since $(G'_x)_0$ is the connected component of the identity, the quotient $G'_x/(G'_x)_0$ is discrete, and the map π is actually a covering space map. Since X is simply connected, the fiber over each point is exactly one point, so $G'_x = (G'_x)_0$. ☐

Lemma 3.4. *G'_x is closed in G for any x.*

Proof. Since G' is open, so is each of its cosets (they are all diffeomorphic to G'). Hence G', being the complement in G of the union of these cosets, is also closed (another way to see this is that G' is a connected component, hence clearly closed). G_x, being the inverse image of x under the map $G \to X(g \mapsto g \cdot x)$, is also closed. It follows that $G'_x = G_x \cap G'$ is also closed. ☐

The upshot of these lemmas is that G'_x is a connected closed subgroup of G. Since it is a closed subgroup, it is also a Lie subgroup. We have already seen that G must map into $O(3)$ if it is to preserve a Riemannian metric on X. Since the connected component of the identity in $O(3)$ is $SO(3)$, we see that G'_x is a connected Lie subgroup of $SO(3)$. We will use, without proof, a well-known fact: the only possibilities for G'_x are $SO(3), SO(2)$ and the trivial group. Accordingly, G_x is also a Lie group of the same dimension.

3.2.2 The theorem

As we saw in the previous section, there are three possibilities for what the point stabilizers can be. It turns out that these decide what X can be.

Theorem 3.5. *There are eight three-dimensional model geometries* (G, X)*:*

(a) If the point stabilizers are three-dimensional, then X *is* $\mathbb{R}^3, S^3$ *or* $\mathbb{H}^3$*.*

(b) If the point stabilizers are two-dimensional, then X *fibers over the two-dimensional model geometries. The connection orthogonal to the fibers has curvature zero (in which case* X *is* $S^2 \times \mathbb{R}$ *or* $\mathbb{H}^2 \times \mathbb{R}$*) or 1 (in which case* X *is nilgeometry (which fibers over* $\mathbb{R}^2$*) or* $\widetilde{SL_2}(\mathbb{R})$ *(which fibers over* $\mathbb{H}^2$*)).*

(c) If the point stabilizers are one-dimensional, X *is solvegeometry, which fibers over the line.*

Proof. Let us assume that X has already been endowed with a G-invariant Riemannian metric (a consequence of Lemma 2.4). We discuss the three possibilities in turn.

(a) The sectional curvature at any point of a manifold is defined in terms of the Gaussian curvatures of various 2-planes in the tangent space at the point, that is, the curvatures of the surfaces whose tangent spaces those planes are. To say that a manifold has constant sectional curvature is to say that, at any given point, the sectional curvatures of any two tangent planes are equal. Since X has a G-invariant Riemannian metric, and G acts transitively on X, a single point having constant sectional curvature would be equivalent to all points having constant sectional curvature.

    If $G'_x = SO(3)$, then any tangent plane at the point x can be taken to any other tangent plane at x. Why? Because we can think of each tangent plane as a great circle of $S^2 \subset T_xX \simeq \mathbb{R}^3$, and the action of $SO(3)$ on S^2 is by rotations, so any great circle can be taken to any other great circle (another way to see this is by rotating the normal to one plane to the normal to the other plane: this is always possible using an element of $SO(3)$). Since the metric, and hence the curvature, is G-invariant, it follows that x has constant sectional curvature, and hence the whole of X has constant sectional curvature.

    The only three-dimensional simply connected complete manifolds with constant sectional curvature are $\mathbb{R}^3$, $\mathbb{H}^3$ and S^3, hence these are the possible geometries for this case.

(b) If G' acts with stabilizer $SO(2)$, then at each $x \in X$, the tangent space T_xX (which is $\mathbb{R}^3$) contains a one-dimensional subspace which is fixed under the action of G'_x. Then we can fix a point $p \in X$ and a tangent vector V_p in this one-dimensional subspace of T_pX. Since the action of G' is transitive, taking $g_*(V_p)$ over all $g \in G$ gives a G'-invariant vector field V on X (this is well-defined since G_p fixes V_p). This is a nowhere-vanishing smooth vector field.

Let ϕ_t be the flow of V at time t. We first prove

Lemma 3.6. *ϕ_t commutes with the action of G'.*

Proof. ϕ_t is the flow of V at time t. It suffices to show that, for any $g \in G$, $\psi_t = g \circ \phi_t \circ g^{-1}$ is also a flow of V. Then by the uniqueness of flow we would have $\phi_t = g \circ \phi_t \circ g^{-1}$ and hence the lemma. We have $\psi_t(p) = g \circ \phi_t \circ g^{-1}(g \cdot p) = g \cdot \phi_t(p)$. So $\psi_0(g \cdot p) = g \cdot \phi_0(p) = g \cdot p$. Moreover, $\psi_0'(g \cdot p) = (g \circ \phi_t \circ g^{-1})'|_0(g \cdot p) = (g \circ \phi_t)'|_0(p) = g_*(\phi_0(p)) = g_*(V_p) = V_{g \cdot p}$. To conclude, we have shown that $\psi_0(g \cdot p) = g \cdot p$ and $\psi_0'(g \cdot p) = V_{g \cdot p}$, which shows that ψ_t is a flow of V at time t, concluding our proof of the lemma. □

Since V is nowhere-vanishing and its integral curves are pairwise-disjoint, they form a one-dimensional foliation of X. Call this foliation $\mathcal{F}$. Suppose the points x and y lie on the same leaf F of $\mathcal{F}$, and suppose $g \in G'_x$. Then, since $y = \phi_t(x)$ for some t, we have $g \cdot y = g \cdot \phi_t(x) = \phi_t(g \cdot x) = y$, so that $g \in G'_y$. It follows that $G'_x = G'_y$.

Lemma 3.7. *If $h \in G'$ takes x to y in F, then h commutes with every element of $G'_x = G'_y$.*

Proof. Let H_F be the group of all elements of G' that keep F invariant (in fact, if any h takes some point of a leaf to another point on the same leaf, it must be in H_F). For $g \in G'_x$ and $h \in H_F$, the map $g \mapsto hgh^{-1}$ is an automorphism of $SO(2) \simeq S^1$. Since there are only two automorphisms of S^1, we get a continuous map from H_F to $\mathbb{Z}/2\mathbb{Z}$. If we can show that H_F is connected, then we would be done, because then hgh^{-1} would be the identity for every $g \in G'_x$. But G'_x is connected and since $H_F/G'_x \simeq F$, which is connected, we have that H_F itself is connected, and the proof is complete. □

Fix a time t and a point x on a leaf F. Let g_t be an element of G' that takes $\phi_t(x)$ back to x. Then $g_t \circ \phi_t$ fixes x. Moreover, since g_t and ϕ_t are both diffeomorphisms, $(g_t \circ \phi_t)_*$ is an automorphism of T_xX. Since $(\phi_t)_*$ and $(g_t)_*$ both leave V invariant, $(g_t \circ \phi_t)_*$ is the identity on the one-dimensional subspace of T_xX that is fixed by G'_x. Moreover, by the preceding results, it commutes with the elements of G'_x. Elementary matrix calculations show that such a map must be of the form

$$\begin{pmatrix} 1 & 0 & 0 \\ 0 & e & f \\ 0 & -f & e \end{pmatrix}, \quad e, f \in \mathbb{R}.$$

Such a matrix can be written as

$$\begin{pmatrix} 1 & 0 & 0 \\ 0 & e & f \\ 0 & -f & e \end{pmatrix} = \begin{pmatrix} 1 & 0 & 0 \\ 0 & \frac{e}{e^2+f^2} & \frac{f}{e^2+f^2} \\ 0 & \frac{-f}{e^2+f^2} & \frac{e}{e^2+f^2} \end{pmatrix} \begin{pmatrix} 1 & 0 & 0 \\ 0 & e^2+f^2 & 0 \\ 0 & 0 & e^2+f^2 \end{pmatrix},$$

that is, a composition of a rotation with a scaling.

Lemma 3.8. *$g_t \circ \phi_t$ can only be a rotation.*

Proof. We need to show that $g_t \circ \phi_t$ has no scaling component. Suppose it does. By assumption, we have a compact manifold M with the geometry of X. Passing to its orientation cover we can assume M is orientable and hence has a canonical volume form ω obtained from the Riemannian metric it inherits from X.

Therefore, $((g_t \circ \phi_t)^*\omega)_p = \lambda(p)\omega_{g_t\circ\phi_t(p)}$, with λ representing the scaling at p. Now $((g_t \circ \phi_t)^*\omega)_p = (\phi_t^*\omega)_p$ and $\omega_{g_t\circ\phi_t(p)} = \omega_{\phi_t(p)}$ since G' keeps the volume form invariant. So we have $(\phi_t^*\omega)_p = \lambda(p)\omega_{\phi_t(p)}$ for all $p \in M$. Suppose q is another point. Then, since the action of G' is transitive, there is some element $g \in G'$ such that $g \cdot p = q$. Then we have $\lambda(q)\omega_{\phi_t(p)} = g^*(\lambda(q)\omega_{\phi_t(q)}) = g^* \circ \phi_t^*(\omega_q) = \phi_t^* \circ g^*(\omega_q) = \phi_t^*(\omega_p) = \lambda(p)\omega_{\phi_t(p)}$, yielding $\lambda(p) = \lambda(q)$. Thus λ is constant over M and we have $(\phi_t^*\omega)_p = \lambda\omega_{\phi_t(p)}$ for some fixed positive real number λ.

In the current scenario, the field V preserves the volume of the compact manifold, hence we have

$$\mathrm{Vol}(M) = \int_M \omega = \int_M \phi_t^*\omega = \lambda \int_M \omega,$$

which forces λ to be 1 and hence the lemma. □

Thus $(g_t \circ \phi_t)_*$ is an isometry on T_xX. Since g_t is an isometry, ϕ_t is also an isometry. The same is true for any x and t, and therefore we have the

Proposition 3.9. *The flow of V is by isometries.*

We now show that the quotient of X by this foliation is itself a two-dimensional model geometry.

Lemma 3.10. *Any two leaves F_1 and F_2 have disjoint neighborhoods.*

Proof. Fix $x \in F_1$ and $y \in F_2$. Let U_1 and U_2 be disjoint neighborhoods of x and y. Let $G_1 = \{h \in G : h \cdot x \in U_1\}$, and analogously G_2 for y and U_2. Then, since the action of G' commutes with the flow, we have that $G_1 \cdot F_1$ and $G_2 \cdot F_2$ are disjoint neighborhoods of F_1 and F_2 respectively. □

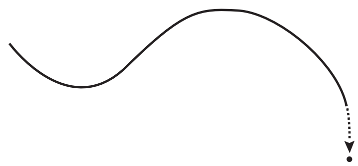

Figure 3.1: An integral curve approaching a point.

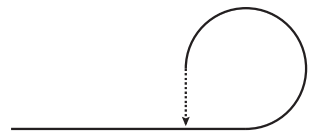

Figure 3.2: An integral curve accumulating on itself.

Lemma 3.11. *Any leaf F is an embedded image of S^1 or $\mathbb{R}$.*

Proof. Suppose not. Then, since F is actually an integral curve for V for all time, it must either approach some point p without actually passing through it (Figure 3.1), or accumulate somewhere on itself (Figure 3.2). In the former case, the point p is a part of some leaf $F' \neq F$, but then F' and F would have no disjoint neighborhoods, contrary to the previous lemma. In the latter case, suppose $q \in F$ is the accumulation point. Then we have $h_n \to e$ for some sequence $\{h_n\}$ in G'. Now consider some point $p' \in F$ "after" p (that, is, lying on the part between p and the part that accumulates on p). Since the group action effectively moves the integral curve "ahead", we must have $h_n \cdot p' \to p$ in X, contradicting the fact that $h_n \cdot p' \to e \cdot p' = p'$. □

From the above we also see that each leaf is closed. If we define $Y = X/\mathcal{F}$, it can be seen that Y itself is a manifold (Hausdorffness follows from Lemma 3.10; the other conditions are easy to check). Recall that we defined H_F to be the subgroup of G' that keeps F invariant. Since each leaf is closed, and the action of H_F is transitive, we see that H_F itself is closed, hence a Lie subgroup of G'. We see that the action of G' on Y is transitive, with stabilizer H_F. In other words, $Y \simeq G'/H_F$.

Being a quotient of X, Y is connected. Y is also simply connected. We see it thus: consider a loop centered at a point $H_F \cdot x \in Y$. It has multiple pre-images in X which are paths. Choose any one such path. The end-points of this path lie on the leaf F of x. Since X is simply connected and F is path-connected, this path is homotopic to a path between the two end-points, lying entirely inside F. Thus their images, the original loop and the identity map, are homotopic, and Y is simply connected.

Proposition 3.12. *Y inherits a Riemannian metric from X, and a transitive action of G' by isometries.*

Proof. Let $\pi : X \to Y$ be the quotient map. At each point $p \in X$, π_* is an isomorphism between the horizontal tangent space (that is, the plane orthogonal to the vector field) and the tangent space of the image of p in Y. Thus we can assign an inner product to $T_{\pi(p)}Y$ as $\langle \pi_*(u), \pi_*(v) \rangle = \langle u, v \rangle$. This is well-defined since $\langle u, v \rangle = \langle (\phi_t)_*(u), (\phi_t)_*(v) \rangle$ for all t, ϕ_t itself being an isometry. All properties required of this inner product to be a metric follow from the fact that it defines a metric on X.

The action of G' on Y is transitive because the action of G' on X was transitive to begin with. That G' preserves the metric follows from the fact that the metric itself is inherited from X. □

Thus we have shown that (G', Y) is itself a two-dimensional model geometry and hence

Proposition 3.13. *Y is either $\mathbb{R}^2$, $\mathbb{H}^2$ or S^2.*

We also note that X is a principal fiber bundle over Y, with fiber and structure group equal to either S^1 or $\mathbb{R}$: since the fibers of the bundle are the leaves of the foliation, clearly the fiber is either S^1 or $\mathbb{R}$. Moreover, since the transition maps on the intersection of two neighborhoods of leaves can only be given by translations. For example, if it is $\mathbb{R}$, given an open set U in Y, we can find an open set U' in X which is of the form $\phi_t(A)$ for some "local section" A of X with local trivialization map given by $\psi(\phi_t(a)) = \psi(a) + t$. If $\phi_t(a)$ too happens to be $\phi_s(b)$ for some open set V' corresponding to an open set V in Y, then the transition map from $\psi(U)$ to $\psi(V)$ would involve a translation by $s - t$.

It is clear that the plane field τ orthogonal to the foliation $\mathcal{F}$ is a connection. By a suitable rescaling, we can assume the curvature of this connection is -1, 0 or 1. Further, by choosing appropriate orientations for the base and the fiber, we can tackle the non-zero cases together, and we thus have two possibilities for the curvature: 0 and 1.

If the curvature is zero, we get a horizontal foliation transverse to the fibers. We thus get local homeomorphisms between the intersections of the leaves with open sets in X and open sets in Y. Considering the topology on a given leaf as its own topology and not the one inherited from X, we see that each leaf is a covering space of Y. Since Y is simply connected, each leaf is itself Y, and we thus get a global section of X.

It can be shown that a principal bundle is trivial if it admits a global section (for example, sending each element of the section to the identity element of the corresponding fiber in the trivial bundle gives a bundle isomorphism).

Thus we conclude that if the curvature is zero, the bundle is trivial. Passing to the universal cover, there are three possible cases: $\mathbb{H}^2 \times \mathbb{R}, S^2 \times \mathbb{R}$ and $\mathbb{R}^3$, two of which are new.

If the curvature is non-zero, we get what is called a *contact structure* (a detailed discussion of this and the next part of the proof is beyond the scope of our project; the reader is advised to refer to [7] to see the full proof). If Y has non-zero curvature, X can be taken as the universal cover of the unit tangent bundle of Y. So if $Y = S^2$, we get S^3 (which is not a new geometry), and if $Y = \mathbb{H}^2$, we get $\widetilde{SL_2}(\mathbb{R})$. If Y has zero curvature, we get *nilgeometry*.

(c) If G' acts with trivial stabilizer, then we have $G'/G'_x \simeq X$, so X is itself a Lie group. So the task becomes one of finding connected and simply connected three-dimensional Lie groups for which there is at least one subgroup H such that G/H is compact, and which are not any of the seven geometries obtained so far. It turns out that there is only one such Lie group, and the resulting geometry is called *solvegeometry*.

□

Chapter 4

The Eight Geometries

This chapter is dedicated to a discussion of some of the eight geometries. Since we are somewhat familiar with the Euclidean, hyperbolic and elliptic geometries, we will refrain from discussing $\mathbb{R}^3, S^3$ and $\mathbb{H}^3$, and instead concentrate on the others. We will need the following, a proof of which can be found in [4]:

Theorem 4.1. *Let Γ be a group acting on a manifold X. The quotient space X/Γ is a manifold if and only if Γ acts freely and properly discontinuously.*

4.1 $S^2 \times \mathbb{R}$

In this case $G = \mathsf{Isom}(S^2) \times \mathsf{Isom}(\mathbb{R}) = O(3) \times \left\{ \begin{pmatrix} a & b \\ 0 & 1 \end{pmatrix} : a \in \mathbb{Z}/2\mathbb{Z}, b \in \mathbb{R} \right\}$, the former being composed of rotations and reflections and the latter of translations and reflections. It turns out that there are only seven manifolds without boundary modeled on this geometry. We will identify all six non-trivial ones ($S^2 \times \mathbb{R}$ is the trivial one) of them as X/Γ for some $\Gamma \subset G$ in accordance with Theorem 4.1. First suppose Γ is generated by some (α, β), where $\alpha \in \mathsf{Isom}(S^2)$ and $\beta \in \mathsf{Isom}(\mathbb{R})$ (note that at least one of α and β must be fixed-point free, leaving us with very few options).

1. If α is the identity and β a translation, then $X/\Gamma \simeq S^2 \times S^1$. Moreover, we have a map $\pi : (S^2 \times \mathbb{R})/\Gamma \to S^1$ given by $\pi([(x, y)]) = [y]$, and a map $\phi : (S^2 \times \mathbb{R})/\Gamma \to S^1 \times S^2$ given by $\phi([(x, y)]) = ([y], [x]) = ([y], x)$, which makes $S^2 \times S^1$ a trivial bundle over S^1.

2. If α is the antipodal map and β a translation, then an equivalence class of $S^2 \times \mathbb{R}$ looks like $\{(x, y), (-x, y + a), (x, y + 2a), \dots\}$, so the resulting quotient is still $S^2 \times S^1$, but it is a non-trivial bundle over S^1, so we denote it $S^2 \tilde{\times} S^1$.

3. If α is the antipodal map and β the identity, then $(S^2 \times \mathbb{R})/\Gamma \simeq \mathbb{R}P^2 \times \mathbb{R}$, which is a trivial line bundle over the projective plane. This case is analogous to the first one.

4. If α is the antipodal map and β a reflection, then we get, analogously to the second case, a non-trivial line bundle over the projective plane, denoted $\mathbb{R}P^2 \tilde{\times} \mathbb{R}$.

5. Now suppose Γ is generated by two elements: (α_1, β_1) and (α_2, β_2). If α_1 is the antipodal map, α_2 the identity, β_1 the identity and β_2 a translation, then (α_1, β_1) and (α_2, β_2) commute, so the identifications in both components happen independently, and we get $\mathbb{R}P^2 \times S^1$, the trivial circle bundle over the projective plane.

6. If α_1 and α_2 are both the antipodal map and β_1 and β_2 distinct reflections, then one easily checks that the manifold obtained is $\mathbb{R}P^3 \# \mathbb{R}P^3$, the connected sum of two projective spaces.

4.2 $\mathbb{H}^2 \times \mathbb{R}$

Since there are infinitely many manifolds modeled on $\mathbb{H}^2$, there are infinitely many manifolds with the geometry of $\mathbb{H}^2 \times \mathbb{R}$ ($S_g \times S^1$ or $S_g \times \mathbb{R}$, for example, where S_g is the orientable surface of genus g).

In general, the mapping torus M_ϕ, where M is a hyperbolic surface and ϕ an isometry of M, is a manifold with this geometry.

$\mathbb{H}^2 \times \mathbb{R}$ is foliated by the lines $\{x\} \times \mathbb{R}$. The isometries of $\mathbb{R}$ leave these lines invariant. The isometries of $\mathbb{H}^2$ don't leave the lines invariant, but they map these lines to other such lines. It follows that the elements of $\mathsf{Isom}(\mathbb{H}^2 \times \mathbb{R}) = \mathsf{Isom}(\mathbb{H}^2) \times \mathsf{Isom}(\mathbb{R})$ keep the foliation invariant, and if we quotient $\mathbb{H}^2 \times \mathbb{R}$ by any subgroup of this group, the resulting space will still be foliated, by either lines or circles (depending on whether an $\mathbb{R}$-translation was involved). This leads us to

Definition 4.2. A space foliated by circles is called a **Seifert fibered space**.

(This is not a rigorous definition, but will serve our purposes. For the full definition, please see [6]).

Now suppose Γ is a discrete, fixed-point free group of isometries of $\mathbb{H}^2 \times \mathbb{R}$ giving us a manifold with this geometry (we do not consider the non-discrete case here). Then so is $\Gamma \cap \mathsf{Isom}(\mathbb{R})$, which therefore must be $\{1\}$ or $\mathbb{Z}$ ($\mathbb{Z}$ is the only nontrivial discrete closed subgroup of $\mathbb{R}$). If it is $\{1\}$, we get a line bundle. If it is $\mathbb{Z}$, we get a Seifert fibered space.

4.3 $\widetilde{SL_2(\mathbb{R})}$

As we saw in the proof of Theorem 3.5, this geometry is actually the universal cover of $U\mathbb{H}^2$, the unit tangent bundle of $\mathbb{H}^2$. One way to see this is that

$U\mathbb{H}^2 \simeq PSL_2(\mathbb{R})$, the orientation-preserving isometries of the hyperbolic plane. This is because $PSL_2(\mathbb{R})$ acts on $U\mathbb{H}^2$ transitively (the action on $\mathbb{H}^2$ is clearly transitive, while a rotation can take any unit tangent vector at a point to any other unit vector at that point), and since rotations are the only orientation-preserving isometries of $\mathbb{H}^2$ with fixed points, the corresponding actions on $U\mathbb{H}^2$ have no fixed points (a nontrivial rotation cannot fix all unit tangent vectors).

As is evident from the preceding discussion or from the proof of Theorem 3.5, the unit tangent bundle of any hyperbolic surface (a circle bundle, hence a Seifert fibered space)is an example of a compact manifold with this geometry.

4.4 *Nilgeometry*

This geometry is isometric to what is known as the *Heisenberg group*:

$$\left\{ \begin{pmatrix} 1 & x & y \\ 0 & 1 & z \\ 0 & 0 & 1 \end{pmatrix} : x, y, z \in \mathbb{R} \right\}$$

This is a nilpotent group (hence the name *nilgeometry*), and has the structure of a line bundle over the plane. An example of a manifold with this geometry is what is known as the *integral Heisenberg group*:

$$\left\{ \begin{pmatrix} 1 & x & y \\ 0 & 1 & z \\ 0 & 0 & 1 \end{pmatrix} : x, y, z \in \mathbb{R}/\mathbb{Z} \right\},$$

which is homeomrphic to the 3-torus, but not isometric to it. For further clarification and examples, we need a slight digression.

Definition 4.3. An **Anosov diffeomorphism** is an automorphism of the 2-torus that splits the tangent bundle into two subbundles, one expanding and one contracting with respect to some Riemannian metric. Equivalently, it is an automorphism given by an element of $SL_2(\mathbb{Z})$ whose eigenvalues are real and distinct.

(Again, this is not a fully rigorous definition. One may be found in [1].)

Definition 4.4. A **torus bundle** is a mapping torus M_ϕ with M the 2-torus and ϕ a diffeomorphism of the 2-torus.

A proof of the following may be found in [ref].

Proposition 4.5. *Geometrically, there are three types of torus bundles M_ϕ:*

(1) If ϕ has finite order, then M_ϕ has the geometry of $\mathbb{R}^3$.

(2) If ϕ is a power of a Dehn twist, then M_ϕ is modeled on nilgeometry.

(3) If ϕ is an Anosov diffeomorphism, then M_ϕ is modeled on solvegeometry.

4.5 *Solvegeometry*

This geometry fibers over the line with the plane as the fiber. It is the only geometry with discrete point stabilizers (the stabilizer is D_8, the dihedral group of order 8). An example of a manifold with this geometry can be obtained from part (3) of Proposition 4.5.

Chapter 5

Conclusion

We have thus studied the proof of Thurston's theorem and understood a little about the eight three-dimensional geometries. The main result involving these is

Thurston's geometrization conjecture Every oriented prime 3-manifold can be cut along tori, so that the interior of each of the resulting manifolds has a geometric structure with finite volume (that is, one of the eight geometries).

This was proved in 2003 by G. Perelman [5]. Corollaries of this result include the following:

Thurston's elliptization conjecture A closed 3-manifold with finite fundamental group is spherical.

Poincaré conjecture Every simply connected, closed 3-manifold is homeomorphic to the 3-sphere.

Bibliography

[1] D. V. Anosov. Geodesic flows on closed Riemannian manifolds of negative curvature. volume 90 of *Trudy Mat. Inst. Steklov.*, 1967.

[2] Manfredo do Carmo. *Riemannian Geometry.* Birkhäuser Basel, 1992.

[3] Shoshichi Kobayashi and Katsumi Nomizu. *Foundations of Differential Geometry.* Wiley, 1996.

[4] John M. Lee. *Introduction to Smooth Manifolds.* Wiley, Springer-Verlag New York.

[5] John W. Morgan. Recent progress on the Poincaré conjecture and the classification of 3-manifolds. *Bull. Amer. Math. Soc.*, 2005.

[6] Peter Scott. The Geometries of 3-Manifolds. *Bull. London Math. Soc.*, 1983.

[7] William P. Thurston. *Three-Dimensional Geometry and Topology, Volume 1.* 1997.