Editor's Preface

These notes are based primarily on a course given by Thurston during the 1977-1978 academic year at Princeton University. The course he had given the previous year had begun with a description of certain topological properties of orientable 3-dimensional manifolds, particularly those that are irreducible (every embedded 2-sphere bounds a ball) and contain a non-trivial embedded, orientable ("incompressible") surface. He described homogeneous geometric structures that arise in dimension three, particularly hyperbolic structures, and conjectured that such 3-manifolds have a canonical decomposition into pieces that have geometric structures. During the first semester he gave an outline of how one might attempt an inductive proof. At the beginning of the second semester, he announced that he had solved the conjecture, except when the manifold fibered over the circle. In particular, he had shown (with this exception) that any compact, irreducible, atoroidal 3-manifold containing an incompressible surface has a complete hyperbolic structure. By the end of the second semester he had solved the fibering case.

Thurston's courses during the next two years dealt with many of the ideas, examples, and techniques surrounding this version of the Geometrization Conjecture (typically called the Geometrization of Haken manifolds). Chapters 1-9 of these notes cover the material that Thurston presented during 1977-1978. Chapters 11 and 13 (there are no chapters 10 and 12) contain some material from the Fall semester of 1978-79.

Thurston started the first year by writing up notes himself, but after a while he enlisted Bill Floyd and me to help. The process was fast and furious. Thurston would lecture on Tuesdays and Thursdays. The notes for that week were prepared over the weekend and handed out to the class the following Tuesday. There was no attempt (or time) to prepare a polished manuscript. I don't think that any of us had any idea how widely distributed the notes ultimately would be. It's just as well since that would likely have led to a more self-conscious product. Instead, what was written was a reasonable facsimile of the way Thurston was presenting his amazing mathematics to us, in real time.

The underlying theme for the course was to show the intimate relationship between low-dimensional topology and geometric structures. The first three chapters of the notes develop the elementary structure of hyperbolic and elliptic geometry and give the general definition of a geometric ("(G, X)") structure. But more importantly they present several concrete examples of knot and link complements as constructed topologically by gluing together ideal polyhedra. This would later allow Thurston to explicitly describe finite volume hyperbolic structures on them. This approach gave a visceral connection between the geometry and the topology of these spaces.

Chapter 4 has ended up being one of the more influential sections of the notes. It is basically all about one example, the figure-eight knot complement. In Chapter 1, Thurston had shown how it could be decomposed into two ideal tetrahedra glued along their faces. He explains how hyperbolic ideal tetrahedra can be parametrized by a single complex parameter. A choice of two complex parameters determines the shapes of the two tetrahedra which can be glued isometrically along their faces. A single polynomial equation must be satisfied for the hyperbolic structure to extend smoothly over the edges; a second equation must be satisfied for the structure to be complete.

It is at this stage that Thurston had a really remarkable idea. By Mostow-Prasad any complete, finite volume hyperbolic structure is unique up to isometry so one can't expect to construct a continuous family of complete structures. However, by satisfying only smoothness conditions along the edges, one has a complex family of *incomplete* structures. For an infinite number of parameter values the incomplete structures can be completed to a smooth hyperbolic structure on different closed manifolds obtained topologically from the figure-eight knot complement by attaching a solid torus along its boundary ("hyperbolic Dehn filling"). Magically, from a single example an infinite number of closed hyperbolic manifolds are created. The 3-dimensional hyperbolic universe suddenly became much larger.

The chapter ends with a purely topological result which is striking in its own right. A (2-sided) incompressible surface in a closed orientable 3-manifold is an embedded closed surface of genus at least one whose fundamental group injects into the fundamental group of the 3-manifold. It is proved that of all the Dehn fillings on the figure-eight knot (parametrized by pairs of relatively prime integers) only six contain an incompressible surface. Before this the only 3-manifolds with infinite fundamental group having no incompressible surface were special Seifert fibered spaces. Furthermore, an infinite number (in fact all) of the hyperbolic manifolds obtained by Dehn filling have no incompressible surfaces. Although it is not explicitly mentioned, this provides an infinite number of examples of hyperbolic 3-manifolds whose existence does not follow from the Haken geometrization theorem.

In Chapter 5 Thurston develops some of the general structure of hyperbolic 3-manifolds. In particular, there is the "thick-thin" decomposition which describes an explicit, simple structure for the subset where the injectivity radius is less than a certain value, the "Margulis constant". Using this Thurston proves a theorem of Jorgensen's about the structure of the set of hyperbolic 3-manifolds whose volumes are less than any given constant. Unlike the situation in higher dimensions where, in each dimension, there are only a finite number with volume below any constant, there can be an infinite number in dimension three. For example, all of the manifolds obtained by hyperbolic Dehn filling the figure eight knot complement have volumes strictly less than that of the figure-eight complement. Jorgensen's Theorem says that, for any given value $c > 0$, there exist a finite number of complete hyperbolic 3-manifolds of volume at most c from which all hyperbolic 3-manifolds with volume at most c can be obtained by Dehn filling. Also, in this chapter, Thurston shows that from any finite volume, non-compact hyperbolic 3-manifold an infinite number of closed hyperbolic manifolds can be obtained by Dehn filling. What seemed like special properties of the figure-eight complement in Chapter 4 are actually general phenomena.

Chapter 6 contains another surprising connection between topology and hyperbolic geometry in dimension three. Gromov defined a purely topological and abstract invariant, now called the "Gromov norm", by minimizing the sum of the sizes of the coefficients of any chain defining the fundamental cycle. Although it has useful properties, it appears hopeless to compute it explicitly. But Gromov proved that, for finite volume hyperbolic manifolds, the Gromov norm equals the hyperbolic volume times a dimensional constant. In this chapter Thurston gives a different proof, using a variant of Gromov's definition. Using properties of the Gromov norm under degree d mappings and Jorgensen's Theorem, it is shown that the set of volumes of hyperbolic 3-manifolds is countable and has the ordinal type of ω^ω. It further follows that the volume always goes strictly down under hyperbolic Dehn filling.

Chapter 7 is based on a lecture given by John Milnor in the course. He derives beautiful formulae for the volumes of many hyperbolic ideal polyhedra in terms of the Lobachevsky function and proves some identities for that function. Since many hyperbolic manifolds can be decomposed into such polyhedra, this provides a way to compute their volumes. This includes the figure-eight knot complement and the Whitehead link complement. Although the values of the functions are not known explicitly, they can be computed to high precision. Other examples include arithmetic hyperbolic manifolds. The volumes of these manifolds can be expressed in terms of zeta functions and this gives a relation between values of zeta functions and values of the Lobachevsky function. It is striking how, although this chapter is written in a different style from the others, it fits so naturally with the rest.

Chapters 8 and 9 begin the journey into the world of Kleinian groups, discrete subgroups of isometries of hyperbolic 3-space, where the focus is on those whose quotients have infinite volume. Kleinian groups play a central role in the proof of Haken geometrization. The proof is by an inductive process using the fact that Haken manifolds can be cut into simpler pieces that typically have infinite volume hyperbolic structures. There are large parameter spaces of structures on these pieces and the proof boils down to proving that they can be varied in a way that they can be glued back together to induce a hyperbolic structure on the original manifold. This is not discussed in the notes. Rather, as with the earlier material, the approach here is to begin with a few definitions and examples and go from there. Nonetheless, the content in these two chapters is so vast and so varied that it would be impossible to even begin to describe it. Instead I will pick out a number of objects that make their first appearance in these pages and have continued to be important players in this area.

Chapter 8 starts by discussing geometrically finite hyperbolic structures where the group action has a finite sided polyhedral fundamental domain. These typically have infinite volume. Thurston defines the *convex core* of the quotient by taking the convex hull of the limit set in hyperbolic space and taking the quotient under the group action. (We will assume that there is no torsion so that the quotient is a manifold.) It will have finite volume when the structure is geometrically finite and it contains all of the closed geodesics of the full quotient space. Unlike the situation in dimension two where the boundary of the convex core consists only of closed geodesics, the boundary in dimension three consists of surfaces that are totally geodesic but bent along a closed collection of simple geodesics called a geodesic lamination. The amount of bending determines a transverse measure and geodesic measured laminations first appear in the theory of Kleinian groups here. Perhaps even more surprising is the analysis of these analytic objects using the combinatorial device of a *train track* which has played an important role in surface topology and the study of the mapping class group.

Thurston specializes to the case where the Kleinian group is quasi-Fuchsian. These are isomorphic to a closed surface group of genus at least two and the convex core is homeomorphic to a surface times an interval. Its boundary consists of two surfaces, each bent along some geodesic lamination. He shows how to fill the quotient space with pleated surfaces, which are also totally geodesic, bent along a geodesic lamination but are no longer embedded. (In the notes they are called *uncrumpled surfaces*, a term that mercifully quickly fell out of favor.) Although such surfaces can be somewhat wild, they create a lot of internal structure that provide important control when limiting to a non-geometrically finite structure. Thurston shows that every geodesic lamination (where transverse measures may not exist and are ignored) arises as the bending lamination of some pleated surface (essentially unique). A similar structure exists for π_1-injective maps of a surface into any geometrically finite hyperbolic 3-manifold.

Chapter 9 begins to analyze the limiting behavior of sequences of geometrically finite structures. A sequence of representations of a finitely presented group into the group of isometries of hyperbolic space *converges algebraically* if the representations of each generator converge. If the sequence comes from a sequence of geometrically finite structures on a fixed underlying manifold, one can consider the corresponding sequence of convex cores. If the sequence of representations converges algebraically to a discrete representation, the limiting convex core may look very different geometrically from the approximates; the limit set of the limiting group can be much smaller than the Hausdorff limit of the approximating limit sets. Thurston defines a notion of a *geometric limit* of a sequence that includes limits of representations of sequences of elements in the group, not just single elements. If the representations converge algebraically, the geometric limit may contain the algebraic limit as a strict subgroup. When they are the same, the sequence is said to *converge strongly*. Thurston shows that the sequence of convex cores converge in a precise sense in this case and gives conditions that guarantee strong convergence.

In the case of a sequence of quasi-Fuchsian structures converging strongly, Thurston shows that the limiting convex core has interior diffeomorphic to a surface times the real line, with the full convex core having zero, one or two geometrically finite ends. The geometrically finite boundary components have the same structure as that of a quasi-Fuchsian group. When an end of the interior does not converge to a boundary component, there are pleated surfaces exiting that end whose bending laminations converge to a unique *ending lamination*. The arguments involved in establishing these properties form the foundation for the Geometrization of Haken manifolds. One sees versions of what came to be called the Covering Theorem and the Double Limit Theorem. And the existence of the ending lamination led Thurston to make his Ending Lamination Conjecture. (See the section on Thurston's work on hyperbolic 3-manifolds for a discussion of this history.)

The 1977-1978 academic year closed with Chapter 9. As Thurston writes at this point in the notes, (he was back to writing them himself since Bill Floyd and I had graduated) the following academic year did not simply continue from there but went in a seemingly different direction. It began with a lengthy investigation of orbifolds (Chapter 13). These objects, which were formally defined and named during 1976-1977, are introduced again, with a focus on low-dimensional examples. In dimension two, all but a few ("bad") orbifolds are seen to have spherical, Euclidean, or hyperbolic structures and the spaces of such structures are seen to have simple compactifications. In dimension three, Thurston first analyzes orbifolds whose underlying space is a 3-dimensional disk, with the singular locus on the boundary. When it is combinatorially a tetrahedron, the results are classical, the proof algebraic. In the remaining cases, the result is due to Andreev. Thurston interprets Andreev's result in terms of circle packings on S^2. He then proves a generalization to the case when the underlying space is a closed surface of non-positive Euler characteristic times an interval and the singular locus is contained in one component of the boundary. By a beautiful argument he proves this by proving the existence of certain circle packings on the closed surface. It is not explicitly mentioned, but the generalization of Andreev's theorem is a special case of the Haken geometrization theorem for certain orbifolds and was to serve as a kind of base case for the induction process.

Thurston wrote at the beginning of that academic year that he intended eventually to write about topics to be included in Chapters 10, 11 and 12, as well as an extended Chapter 9. Only Chapter 11 appeared. It is a short discussion of extending vector fields and other objects from the sphere at infinity to hyperbolic 3-space by a type of visual averaging. This was used in the latter part of the first semester to extend isotopies from the boundary to the interior of geometric structures during the process of analyzing the space of such structures on certain 3-manifolds. There were other topics covered that semester but they did not make it into the notes.

It is quite remarkable that, despite the fact that the impetus for these courses came from the Haken geometrization theorem, it was never really mentioned in the notes. This very much reflected Thurston's attitude about mathematics. He rarely was content with simply *proving* a result. Focusing on a particular theorem could be a distraction from the overall context in which it existed. Rather he would *surround* a theorem by providing a world full of examples and techniques in which it could not fail to be true. The point was to appreciate the world one had discovered.

As he wrote in the introduction:

I did not always know where I was going and the discussion tends to wander. The countryside is scenic, however, and it is fun to tramp around if you keep your eyes alert and don't get lost.

Perhaps even more than this, Thurston wanted to provide an accessible entry for a large variety of mathematicians so that they could share in his enjoyment. The point was to build a community who could inhabit that world. These notes are a good start.

Steve Kerckhoff
March 1, 2022