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 13

Orbifolds

As we have had occasion to see, it is often more effective to study the quotient manifold of a group acting freely and properly discontinuously on a space rather than to limit one's image to the group action alone. It is time now to enlarge our vocabulary, so that we can work with the quotient spaces of groups acting properly discontinuously but not necessarily freely. In the first place, such quotient spaces will yield a technical device useful for showing the existence of hyperbolic structures on many three-manifolds. In the second place, they are often simpler than three-manifolds tend to be, and hence they often give easy, graphic examples of phenomena involving three-manifolds. Finally, they are beautiful and interesting in their own right.

13.1. Some examples of quotient spaces.

We begin our discussion with a few examples of quotient spaces of groups acting properly discontinuously on manifolds in order to get a taste of their geometric flavor.

EXAMPLE 13.1.1 (A single mirror). Consider the action of $\mathbb{Z}_2$ on $\mathbb{R}^3$ by reflection in the $y-z$ plane. The quotient space is the half-space $x \geq 0$. Physically, one may imagine a mirror placed on the $y-z$ wall of the half-space $x \geq 0$. The scene as viewed by a person in this half-space is like all of $\mathbb{R}^3$, with scenery invariant by the $\mathbb{Z}_2$ symmetry.

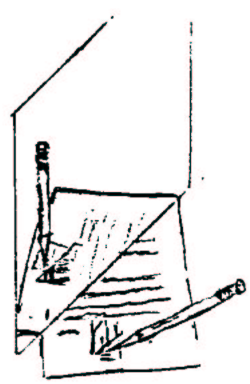

13.2

Example 13.1.2 (A barber shop). Consider the group G generated by reflections in the planes $x = 0$ and $x = 1$ in $\mathbb{R}^3$. G is the infinite dihedral group $D_\infty = \mathbb{Z}_2 * \mathbb{Z}_2$. The quotient space is the slab $0 \leq x \leq 1$. Physically, this is related to two mirrors on parallel walls, as commonly seen in a barber shop.

Example 13.1.3 (A billiard table). Let G be the group of isometries of the Euclidean plane generated by reflection in the four sides of a rectangle R. G is isomorphic to $D_\infty \times D_\infty$, and the quotient space is R. A physical model is a billiard table. A collection of balls on a billiard table gives rise to an infinite collection of balls on $\mathbb{R}^2$, invariant by G. (Each side of the billiard table should be one ball diameter larger than the corresponding side of R so that the *centers* of the balls can take any position in R. A ball may intersect its images in $\mathbb{R}^2$.)

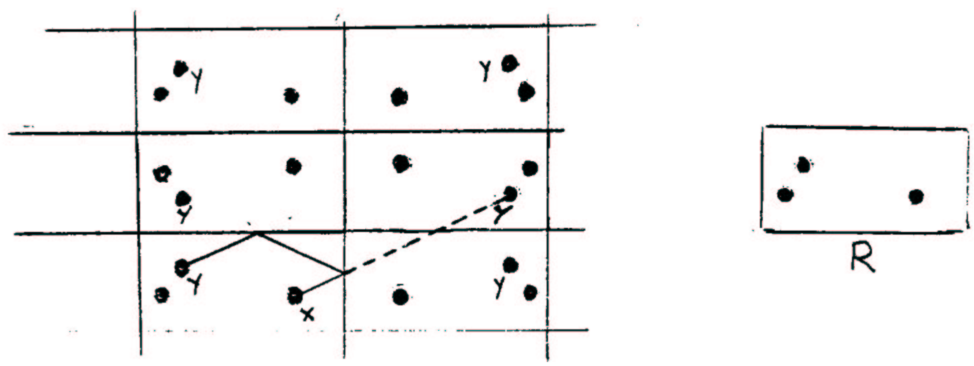

Ignoring spin, in order to make ball x hit ball y it suffices to aim it at any of the images of y by G. (Unless some ball is in the way.)

Example 13.1.4 (A rectangular pillow). Let H be the subgroup of index 2 which preserves orientation in the group G of the preceding example. A fundamental do- 13.3
main for H consists of two adjacent rectangles. The quotient space is obtained by identifying the edges of the two rectangles by reflection in the common edge.

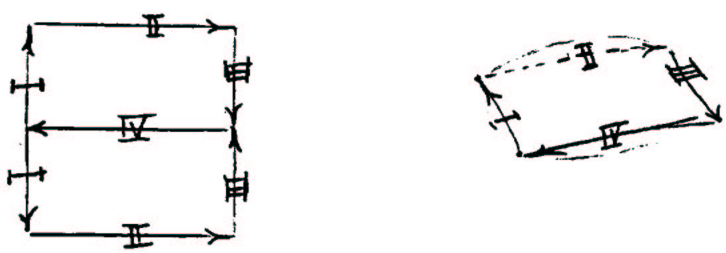

Topologically, this quotient space is a sphere, with four distinguished points or singular points, which come from points in $\mathbb{R}^2$ with non-trivial isotropy ($\mathbb{Z}_2$). The sphere inherits a Riemannian metric of 0 curvature in the complement of these 4 points, and it has curvature $K_{p_i} = \pi$ concentrated at each of the four points p_i. In other words, a neighborhood of each point p_i is a cone, with cone angle $\pi = 2\pi - K_{p_i}$.

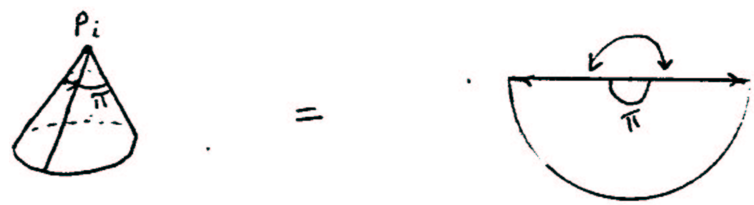

Exercise. On any tetrahedron in $\mathbb{R}^3$ all of whose four sides are congruent, every geodesic is simple. This may be tested with a cardboard model and string or with strips of paper. Explain.

Example 13.1.5 (An orientation-preserving crystallographic group). Here is one more three-dimensional example to illustrate the geometry of quotient spaces. Con- 13.4
sider the 3 families of lines in $\mathbb{R}^3$ of the form $(t, n, m+\frac{1}{2})$, $(m+\frac{1}{2}, t, n)$ and $(n, m+\frac{1}{2}, t)$ where n and m are integers and t is a real parameter. They intersect a cube in the unit lattice as depicted.

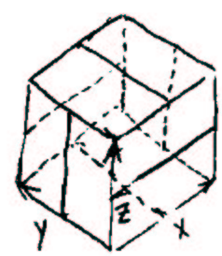

Let G be the group generated by 180° rotations about these lines. It is not hard to see that a fundamental domain is a unit cube. We may construct the quotient space by making all identifications coming from non-trivial elements of G acting on the faces of the cube. This means that each face must be folded shut, like a book. In doing this, we will keep track of the images of the axes, which form the singular locus.

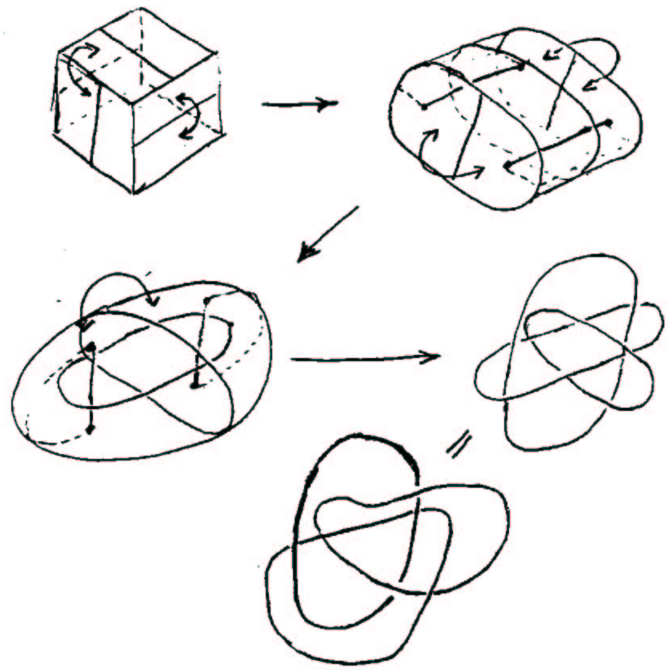

13.5

As you can see by studying the picture, the quotient space is S^3 with singular locus consisting of three circles in the form of Borromean rings. S^3 inherits a Euclidean structure (or metric of zero curvature) in the complement of these rings, with a cone-type singularity with cone angle π along the rings.

In these examples, it was not hard to construct the quotient space from the group action. In order to go in the opposite direction, we need to know not only the quotient space, but also the singular locus and appropriate data concerning the local behavior of the group action above the singular locus.

13.2. Basic definitions.

An *orbifold** O is a space locally modelled on $\mathbb{R}^n$ modulo finite group actions. Here is the formal definition: O consists of a Hausdorff space X_O, with some additional structure. X_O is to have a covering by a collection of open sets $\{U_i\}$ closed under finite intersections. To each U_i is associated a finite group Γ_i, an action of Γ_i on an open subset $\tilde{U}_i$ of $\mathbb{R}^n$ and a homeomorphism $\varphi_i : U_i \approx \tilde{U}_i/\Gamma_i$. Whenever $U_i \subset U_j$,

*This terminology should not be blamed on me. It was obtained by a democratic process in my course of 1976-77. An orbifold is something with many folds; unfortunately, the word "manifold" already has a different definition. I tried "foldamani," which was quickly displaced by the suggestion of "manifolded." After two months of patiently saying "no, not a manifold, a manifol*dead*," we held a vote, and "orbifold" won.

there is to be an injective homomorphism

$$f_{ij} : \Gamma_i \hookrightarrow \Gamma_j$$

and an embedding

$$\tilde{\varphi}_{ij} : \tilde{U}_i \hookrightarrow \tilde{U}_j$$

equivariant with respect to f_{ij} (i.e., for $\gamma \in \Gamma_i$, $\tilde{\varphi}_{ij}(\gamma x) = f_{ij}(\gamma)\tilde{\varphi}_{ij}(x)$) such that the 13.6
diagram below commutes.†

$$\begin{tikzcd}
\tilde{U}_i \arrow[r, "\tilde{\varphi}_{ij}"] \arrow[d] & \tilde{U}_j \arrow[d] \\
\tilde{U}_i/\Gamma_i \arrow[r, "\varphi_{ij} = \tilde{\varphi}_{ij}/\Gamma_i"] & \tilde{U}_j/\Gamma_i \arrow[d, "f_{ij}"] \\
 & \tilde{U}_j/\Gamma_j \\
U_i \arrow[uu, "\varphi_i"] \arrow[r, phantom, "\subset"] & U_j \arrow[u, "\varphi_j"']
\end{tikzcd}$$

We regard $\tilde{\varphi}_{ij}$ as being defined only up to composition with elements of Γ_j, and f_{ij} as being defined up to conjugation by elements of Γ_j. It is not generally true that $\tilde{\varphi}_{ik} = \tilde{\varphi}_{jk} \circ \tilde{\varphi}_{ij}$ when $U_i \subset U_j \subset U_k$, but there should exist an element $\gamma \in \Gamma_k$ such that $\gamma\tilde{\varphi}_{ik} = \tilde{\varphi}_{jk} \circ \tilde{\varphi}_{ij}$ and $\gamma \cdot f_{ik}(g) \cdot \gamma^{-1} = f_{jk} \circ f_{ij}(g)$.

Of course, the covering $\{U_i\}$ is not an intrinsic part of the structure of an orbifold: two coverings give rise to the same orbifold structure if they can be combined consistently to give a larger cover still satisfying the definitions.

A $\mathcal{G}$-orbifold, where $\mathcal{G}$ is a pseudogroup, means that all maps and group actions respect $\mathcal{G}$. (See chapter 3).

EXAMPLE 13.2.1. A closed manifold is an orbifold, where each group Γ_i is the trivial group, so that $\tilde{U} = U$.

EXAMPLE 13.2.2. A manifold M with boundary can be given an orbifold structure mM in which its boundary becomes a "mirror." Any point on the boundary has a neighborhood modelled on $\mathbb{R}^n/\mathbb{Z}_2$, where $\mathbb{Z}_2$ acts by reflection in a hyperplane. 13.7

†The commutative diagrams in Chapter 13 were made using Paul Taylor's diagrams.sty package (available at ftp://ftp.dcs.qmw.ac.uk/pub/tex/contrib/pt/diagrams/). —SL

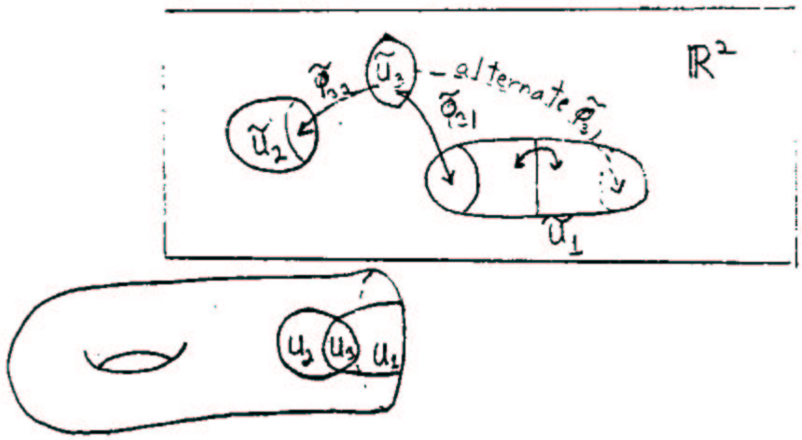

Proposition 13.2.1. *If M is a manifold and Γ is a group acting properly discontinuously on M, then M/Γ has the structure of an orbifold.*

Proof. For any point $x \in M/\Gamma$, choose $\tilde{x} \in M$ projecting to x. Let I_x be the isotropy group of $\tilde{x}$ (I_x depends of course on the particular choice $\tilde{x}$.) There is a neighborhood $\tilde{U}_x$ of $\tilde{x}$ invariant by I_x and disjoint from its translates by elements of Γ not in I_x. The projection of $U_x = \tilde{U}_x/I_x$ is a homeomorphism. To obtain a suitable cover of M/Γ, augment some cover $\{U_x\}$ by adjoining finite intersections. Whenever $U_{x_1} \cap \ldots \cap U_{x_k} \neq \emptyset$, this means some set of translates $\gamma_1 \tilde{U}_{x_1} \cap \ldots \cap \gamma_k \tilde{U}_{k_k}$ has a corresponding non-empty intersection. This intersection may be taken to be

$$\overbrace{U_{x_1} \cap \cdots \cap U_{x_k}}^{\sim},$$

with associated group $\gamma_1 I_{x_1} \gamma_1^{-1} \cap \cdots \cap \gamma_k I_{x_k} \gamma_k^{-1}$ acting on it. □

The orbifold mM arises in this way, for instance: it is obtained as the quotient space of the $\mathbb{Z}_2$ action on the double dM of M which interchanges the two halves. 13.8

Henceforth, we shall use the terminology M/Γ to mean M/Γ as an orbifold.

Note that each point x in an orbifold O is associated with a group Γ_x, well-defined up to isomorphism: in a local coordinate system $U = \tilde{U}/\Gamma$, Γ_x is the isotropy group of any point in $\tilde{U}$ corresponding to x. (Alternatively Γ_x may be defined as the smallest group corresponding to some coordinate system containing x.) The set $\Sigma_O = \{x | \Gamma_x \neq \{1\}\}$ is the *singular locus* of O. *We shall say that O is a manifold when $\Sigma_O = \emptyset$. Warning.* It happens much more commonly that the underlying space X_O is a topological manifold, especially in dimensions 2 and 3. Do not confuse properties of O with properties of X_O.

The singular locus is a closed set, since its intersection with any coordinate patch is closed. Also, it is nowhere dense. This is a consequence of the fact that a non-trivial homeomorphism of a manifold which fixes an open set cannot have finite order. (See Newman, 1931. In the differentiable case, this is an easy exercise.)

When M in the proposition is simply connected, then M plays the role of universal covering space and Γ plays the role of the fundamental group of the orbifold M/Γ, (even though the underlying space of M/Γ may well be simply connected, as in the examples of §13.1). To justify this, we first define the notion of a covering orbifold.

Definition 13.2.2. A *covering orbifold* of an orbifold O is an orbifold $\tilde{O}$, with a projection $p : X \to X_O$ between the underlying spaces, such that each point $x \in X_O$ has a neighborhood $U = \tilde{U}/\Gamma$ (where $\tilde{U}$ is an open subset of $\mathbb{R}^n$) for which each 13.9
component v_i of $p^{-1}(U)$ is isomorphic to $\tilde{U}/\Gamma_i$, where $\Gamma_i \subset \Gamma$ is some subgroup. The isomorphism must respect the projections.

Note that the underlying space $X_{\tilde{O}}$ is not generally a covering space of X_O.

As a basic example, when Γ is a group acting properly discontinuously on a manifold M, then M is a covering orbifold of M/Γ. In fact, for any subgroup $\Gamma' \subset \Gamma$, M/Γ' is a covering orbifold of M/Γ. Thus, the rectangular pillow (13.1.4) is a two-fold covering space of the billiard table (13.1.3).

Here is another explicit example to illustrate the notion of covering orbifold. Let S be the infinite strip $0 \le x \le 1$ in $\mathbb{R}^2$; consider the orbifold mS. Some covering spaces of S are depicted below.

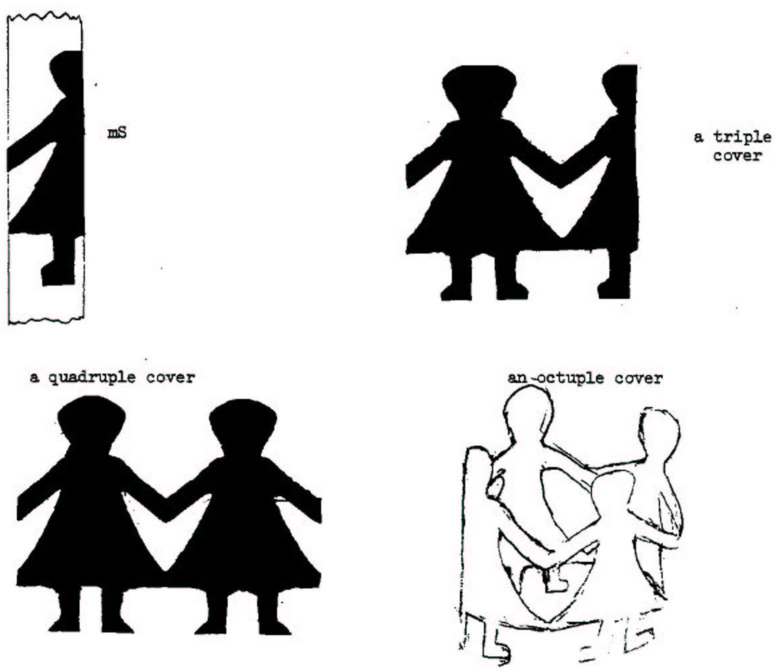

Definition 13.2.3. An orbifold is *good* if it has some covering orbifold which is a manifold. Otherwise it is *bad*.

The teardrop is an example of a bad orbifold. The underlying space for a teardrop is S^2. Σ_O consists of a single point, whose neighborhood is modelled on $\mathbb{R}^2/\mathbb{Z}_n$, where $\mathbb{Z}_n$ acts by rotations.

By comparing possible coverings of the upper half with possible coverings of the lower half, you may easily see that the teardrop has no non-trivial connected coverings.

Similarly, you may verify that an orbifold O with underlying space $X_O = S^2$ having only two singular points associated with groups $\mathbb{Z}_n$ and $\mathbb{Z}_n$ is bad, unless $n = m$. The orbifolds with three or more singular points on S^2, as we shall see, are always good. For instance, the orbifold below is S^2 modulo the orientation-preserving symmetries of a dodecahedron.

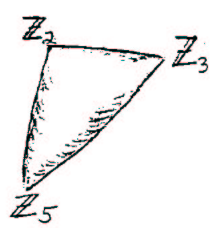

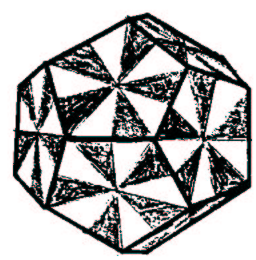

13.11

Proposition 13.2.4. *An orbifold O has a universal cover $\tilde{O}$. In other words, if $* \in X_O - \Sigma_O$ is a base point for O,*

$$\tilde{O} \xrightarrow{p} O$$

is a connected covering orbifold with base point $\tilde{*}$ which projects to $*$, such that for any other covering orbifold

$$\tilde{O}' \xrightarrow{p'} O$$

with base point $\tilde{*}'$, $p'(\tilde{*}') = *$, there is a lifting $q : \tilde{O} \to \tilde{O}'$ of p to a covering map of $\tilde{O}'$.

$$\begin{tikzcd}
\tilde{O} \arrow[dd, "p"'] \arrow[dr, "q"] & \\
& \tilde{O}' \arrow[dl, "p'"] \\
O &
\end{tikzcd}$$

The universal covering orbifold $\tilde{O}$, in some contexts, is often called the universal branched cover. There is a simple way to prove 13.2.4 in the case Σ_O has codimension 2 or more. In that case, any covering space of O is determined by the induced covering space of $X_O - \Sigma_O$ as its metric completion. Whether a covering Y space of $X_O - \Sigma_O$ comes from a covering space of O is a local question, which is expressed algebraically by saying that $\pi_1(Y)$ maps to a group containing a certain obvious normal subgroup of $\pi_1(X - \Sigma_O)$.

When O is a good orbifold, then it is covered by a simply connected manifold, M. It can be shown directly that M is the universal covering orbifold by proving that every covering orbifold is isomorphic to M/Γ', for some $\Gamma' \subset \Gamma$, where Γ is the group of deck transformations of M over O.

Proof of 13.2.4. One proof of the existence of a universal cover for a space X goes as follows.

Consider pointed, connected covering spaces

$$\tilde{X}_i \xrightarrow{p_i} X.$$

For any pair of such covering spaces, the component of the base point in the fiber product of the two is a covering space of both. 13.12

$$\begin{tikzcd}
 & \tilde{X}_3 \arrow[dl] \arrow[dr] & \\
\tilde{X}_1 \arrow[dr] & & \tilde{X}_2 \arrow[dl] \\
 & X &
\end{tikzcd}$$

(Recall that the fiber product of two maps $f_i : X_i \to X$ is the space $X_1 \times_X X_2 = \{(x_1, x_2) \in X_1 \times X_2 : f_1(x_1) = f_2(x_2)\}$.)

If X is locally simply connected, or more generally, if it has the property that every $x \in X$ has a neighborhood U such that every covering of X induces a trivial covering of U (that is, each component of $p^{-1}(U)$ is homeomorphic to U), then one can take the inverse limit over some set of pointed, connected covering spaces of X which represents all isomorphism classes to obtain a universal cover for X.

We can follow this same outline with orbifolds, but we need to refine the notion of fiber product. The difficulty is best illustrated by example. Two covering maps

$$S^1 = dI \xrightarrow{f_2} m_i \quad \text{and} \quad m_i \to m_i$$

are sketched below, along with the fiber product of the underlying maps of spaces.

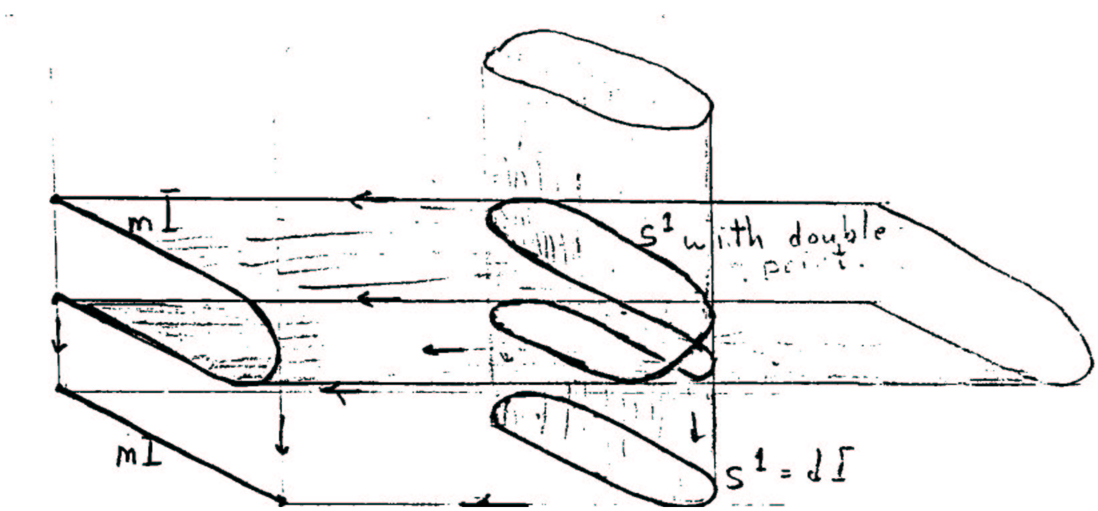

(This picture is sketched in $\mathbb{R}^3 = \mathbb{R}^2 \times_{\mathbb{R}_1} \mathbb{R}^2$.) The fiber product of spaces is a circle but with a double point. In the definition of fiber product of orbifolds, we must eliminate such double points, which always lie above Σ_O.

To do this, we work in local coordinates. Let $U \approx \tilde{U}/\Gamma$ be a coordinate system. 13.13
We may suppose that U is small enough so in every covering of O, $p^{-1}(U)$ consists

of components of the form $\tilde{U}/\Gamma'$, $\Gamma' \subset \Gamma$. Let

$$O_i \xrightarrow{p_i} O$$

be covering orbifolds ($i = 1, 2$), and consider components of $p_i^{-1}(U)$, which for notational convenience we identify with $\tilde{U}/\Gamma_1$ and $\tilde{U}/\Gamma_2$. Formally, we can write $\tilde{U}/\Gamma_1 = \{\Gamma_1 y \,|\, y \in \tilde{U}\}$. [It would be more consistent to use the notation $\Gamma_1 \backslash \tilde{U}$ instead of $\tilde{U}/\Gamma_1$]. For each pair of elements γ_1 and $\gamma_2 \in \Gamma$, we obtain a map

$$f_{\gamma_1,\gamma_2} : \tilde{U} \to \tilde{U}/\Gamma_1 \times \tilde{U}/\Gamma_2,$$

by the formula

$$f_{\gamma_1,\gamma_2} y = (\Gamma_1 \gamma_1 y, \Gamma_2 \gamma_2 y).$$

In fact, f_{γ_1,γ_2} factors through

$$\tilde{U}/\gamma_1^{-1}\Gamma_1\gamma_1 \cap \gamma_2^{-1}\Gamma_2\gamma_2.$$

Of course, f_{γ_1,γ_2} depends only on the cosets $\Gamma_2\gamma_1$ and $\Gamma_2\gamma_2$. Furthermore, for any $\gamma \in \Gamma$, the maps f_{γ_1,γ_2} and $f_{\gamma_1\gamma,\gamma_2\gamma}$ differ only by a group element acting on $\tilde{U}$; in particular, their images are identical so only the product $\gamma_1\gamma_2^{-1}$ really matters. Thus, the "real" invariant of f_{γ_1,γ_2} is the double coset

$$\Gamma_1\gamma_1\gamma_2^{-1}\Gamma_2 \in \Gamma_1\backslash\Gamma/\Gamma_2.$$

(Similarly, in the fiber product of coverings X_1 and X_2 of a space X, the components are parametrized by the double cosets $\pi_1 X_1 \backslash \pi_1 X / \pi_1 X_2$.) The fiber product of $\tilde{U}/\Gamma_1$ and $\tilde{U}/\Gamma_2$ over $\tilde{U}/\Gamma$, is defined now to be the disjoint union, over elements γ representing double cosets $\Gamma_1\backslash\Gamma/\Gamma_2$ of the orbifolds $\tilde{U}/\Gamma_1 \cap \gamma^{-1}\Gamma_2\gamma$. We have shown above how this canonically covers $\tilde{U}/\Gamma_1$ and $\tilde{U}/\Gamma_2$, via the map $f_{1,\gamma}$. This definition agrees with the usual definition of fiber product in the complement of Σ_O. These locally defined patches easily fit together to give a fiber product orbifold $O_1 \times_O O_2$. As in the case of spaces, a universal covering orbifold $\tilde{O}$ is obtained by taking the inverse limit over some suitable set representing all isomorphism classes of orbifolds. □

The universal cover $\tilde{O}$ of an orbifold O is automatically a regular cover: for any preimage of $\tilde{x}$ of the base point $*$ there is a deck transformation taking $\tilde{*}$ to $\tilde{x}$. 13.14

Definition 13.2.5. The *fundamental group* $\pi_1(O)$ of an orbifold O is the group of deck transformations of the universal cover $\tilde{O}$.

The fundamental groups of orbifolds can be computed in much the same ways as fundamental groups of manifolds. Later we shall interpret $\pi_1(O)$ in terms of loops on O.

Here are two more definitions which are completely parallel to definitions for manifolds.

Definition 13.2.6. *An orbifold with boundary* means a space locally modelled on $\mathbb{R}^n$ modulo finite groups and $\mathbb{R}^n_+$ modulo finite groups.

When X_O is a topological manifold, be careful not to confuse ∂X_O with ∂O or $X_{\partial O}$.

Definition 13.2.7. A *suborbifold* O_1 of an orbifold O_2 means a subspace $X_{O_1} \subset X_{O_2}$ locally modelled on $\mathbb{R}^d \subset \mathbb{R}^n$ modulo finite groups.

Thus, a triangle orbifold has seven distinct "closed" one-dimensional suborbifolds, up to isotopy: one S^1 and six mI's.

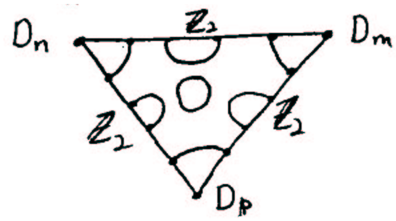

Note that each of the seven is the boundary of a suborbifold with boundary (defined in the obvious way) with universal cover D^2. 13.15

13.3. Two-dimensional orbifolds.

To avoid technicalities, we shall work with differentiable orbifolds from now on.

The nature of the singular locus of a differentiable orbifold may be understood as follows. Let $U = \tilde{U}/\Gamma$ be any local coordinate system. There is a Riemannian metric on $\tilde{U}$ invariant by Γ: such a metric may be obtained from any metric on $\tilde{U}$ by averaging under Γ. For any point $\tilde{x} \in \tilde{U}$ consider the exponential map, which gives a diffeomorphism from the ϵ ball in the tangent space at $\tilde{x}$ to a small neighborhood of $\tilde{x}$. Since the exponential map commutes with the action of the isotropy group of $\tilde{x}$, it gives rise to an isomorphism between a neighborhood of the image of $\tilde{x}$ in O, and a neighborhood of the origin in the orbifold $\mathbb{R}^n/\Gamma$, where Γ is a finite subgroup of the orthgonal group O_n.

Proposition 13.3.1. *The singular locus of a two-dimensional orbifold has these types of local models:*

(i) *The mirror: $\mathbb{R}^2/\mathbb{Z}_2$, where $\mathbb{Z}_2$ acts by reflection in the y-axis.*
(ii) *Elliptic points of order n: $\mathbb{R}^2/\mathbb{Z}_n$, with $\mathbb{Z}_n$ acting by rotations.*
(iii) *Corner reflectors of order n: $\mathbb{R}^2/D_n$, with D_n is the dihedral group of order $2n$, with presentation*

$$\langle a, b : a^2 = b^2 = (ab)^n = 1\rangle.$$

The generators a and b correspond to reflections in lines meeting at angle π/n. 13.16

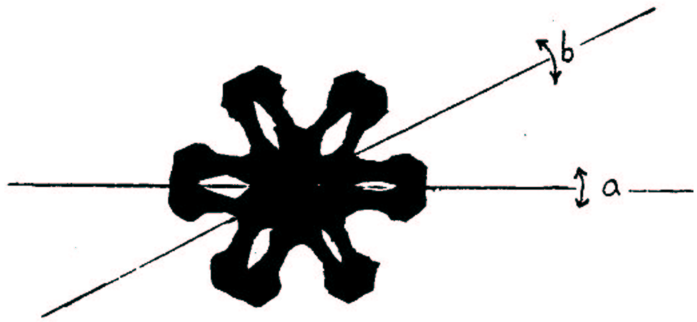

PROOF. These are the only three types of finite subgroups of O_2. □

It follows that the underlying space of a two-dimensional orbifold is always a topological surface, possibly with boundary. It is easy to enumerate all two-dimensional orbifolds, by enumerating surfaces, together with combinatorial information which determines the orbifold structure. From a topological point of view, however, it is not completely trivial to determine which of these orbifolds are good and which are bad.

We shall classify two-dimensional orbifolds from a geometric point of view. When G is a group of real analytic diffeomorphisms of a real analytic manifold X, then the elementary properties of (G, X)-orbifolds are similar to the case of manifolds (see §3.5). In particular a developing map

$$D : \tilde{O} \to X$$

can be defined for a (G, X)-orbifold O. Since we do not yet have a notion of paths in O, this requires a little explanation. Let $\{U_i\}$ be a covering of O by a collection of open sets, closed under intersections, modelled on $\tilde{U}_i/\Gamma_i$, with $\tilde{U}_i \subset X$, such that the 13.17
inclusion maps $U_i \subset U_j$ come from isometries $\tilde{\varphi}_{ij} : \tilde{U}_i \to \tilde{U}_j$. Choose a "base" chart $\tilde{U}_0$. When $U_0 \supset U_{i_1} \subset U_{i_2} \supset \cdots \subset U_{i_{2n}}$ is a chain of open sets (a simplicial path in the one-skeleton of the nerve of $\{U_i\}$), then for each choice of isometries of the form

$$\tilde{U}_0 \xleftarrow{\gamma_0 \tilde{\varphi}_{i_1,0}} \tilde{U}_{i_1} \xrightarrow{\gamma_2' \tilde{\varphi}_{i_1,i_2}} \tilde{U}_{i_2} \longleftarrow \cdots \longrightarrow \tilde{U}_{i_{2n}}$$

one obtains an isometry of $\tilde{U}_{i_{2n}}$ in X, obtained by composing the transition functions (which are globally defined on X). A covering space $\tilde{O}$ of O is defined by the covering $\{(\varphi, \varphi(\tilde{U}_i))\} \subset G \times X$, where φ is any isometry of $\tilde{U}_i$ obtained by the above construction.

These are glued together by the obvious "inclusion" maps, $(\varphi, \varphi\tilde{U}_i) \hookrightarrow (\psi, \psi\tilde{U}_j)$ whenever $\psi^{-1} \circ \varphi$ is of the form $\gamma_j \circ \tilde{\varphi}_{ij}$ for some $\gamma_j \in \Gamma_j$.

The reader desiring a picture may construct a "foliation" of the space $\{(x, y, g) \mid x \in X,\ y \in X_O,\ g$ is the germ of a G-map between neighborhoods of x and $y\}$. Any leaf of this foliation gives a developing map.

Proposition 13.3.2. *When G is an analytic group of diffeomorphisms of a manifold X, then every (G, X)-orbiifold is good. A developing map*

$$D : \tilde{O} \to X$$

and a holonomy homomorphism

$$H : \pi_1(O) \to G$$

are defined.

If G is a group of isometries acting transitively on X, then if O is closed or metrically complete, it is complete (i.e., D is a covering map). In particular, if X is simply connected, then $\tilde{O} = X$ and $\pi_1(O)$ is a discrete subgroup of G. 13.18

Proof. See §3.5. □

Here is an example. $\triangle_{2,3,6}$ has a Euclidean structure, as a 30°, 60°, 90° triangle. The developing map looks like this:

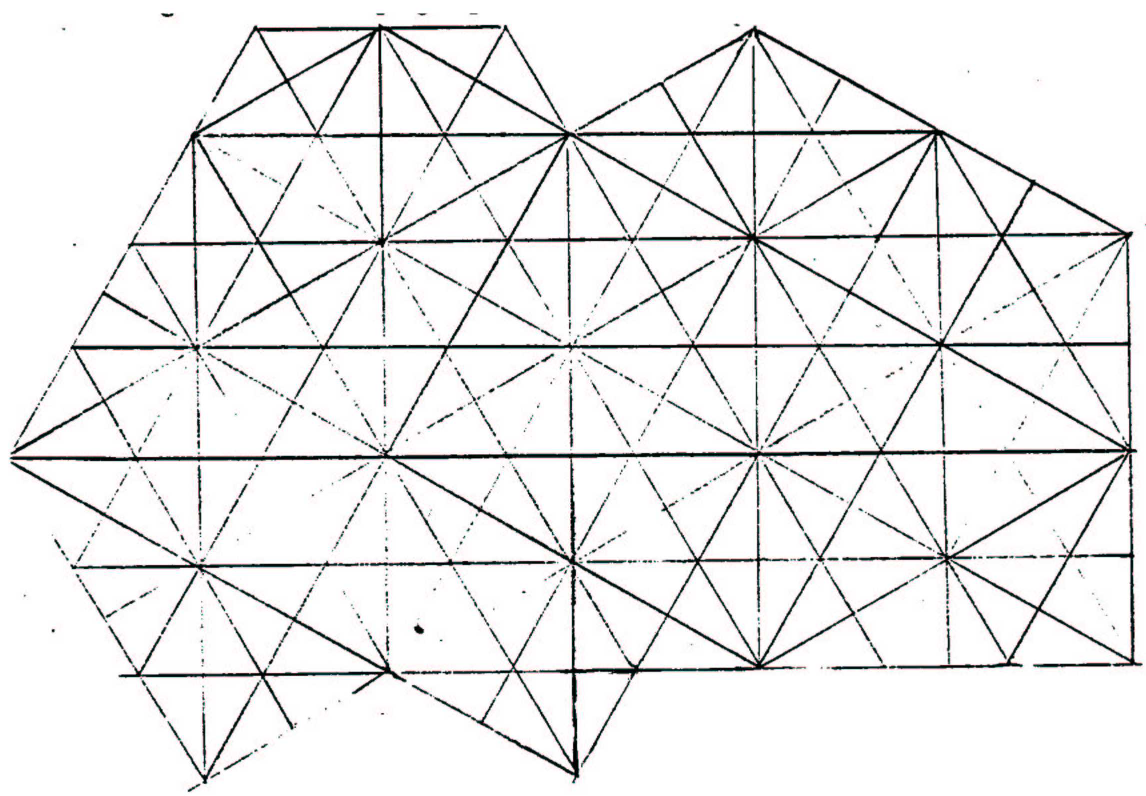

Here is a definition that will aid us in the geometric classification of two-dimensional orbifolds.

DEFINITION 13.3.3. When an orbifold O has a cell-division of X_O such that each open cell is in the same stratum of the singular locus (i.e., the group associated to the interior points of a cell is constant), then the *Euler number* $\chi(O)$ is defined by the formula 13.19

$$\chi(O) = \sum_{c_i} (-1)^{\dim(c_i)} \frac{1}{|\Gamma(c_i)|},$$

where c_i ranges over cells and $|\Gamma(c_i)|$ is the order of the group $\Gamma(c_i)$ associated to each cell. The Euler number is not always an integer.

The definition is concocted for the following reason. Define the *number of sheets* of a cover to be the number of preimages of a non-singular point.

PROPOSITION 13.3.4. *If $\tilde{O} \to O$ is a covering map with k sheets, then*

$$\chi(\tilde{O}) = k\chi(O).$$

PROOF. It is easily verified that the number of sheets of a cover can be computed by the ratio

$$\#\text{ sheets } = \sum_{\tilde{x} \ni p(\tilde{x})=x} (\,|\Gamma_x| \quad / |\Gamma_{\tilde{x}}|\,),$$

where x is *any* point. The formula [???] for the Euler number of a cover follows immediately. □

As an example, a triangle orbifold Δ_{n_1,n_2,n_3} has Euler number $\frac{1}{2}\big(\sum(1/n_i) - 1\big)$: here $+1$ comes from the 2-cell, three $-\frac{1}{2}$'s from the edges, and $1/(2n_i)$ from each vertex. 13.20

Thus, $\Delta_{2,3,5}$ has Euler number $+1/60$. Its universal cover is S^2, with deck transformations the group of symmetries of the dodecahedron. This group has order $120 = 2/(1/60)$. On the other hand, $\chi(\Delta_{2,3,6}) = 0$ and $\chi(\Delta_{2,3,7}) = -1/84$. These orbifolds cannot be covered by S^2.

The general formula for the Euler number of an orbifold O with k corner reflectors of orders $n_1, \ldots, n_k$ and l elliptic points of orders $m_1, \ldots, m_l$ is

13.3.4. $$\chi(O) = \chi(X_O) - \tfrac{1}{2}\sum(1 - 1/n_i) - \sum(1 - 1/m_i).$$

Note in particular that $\chi(O) \leq \chi(X_O)$, with equality if and only if O is the surface χ_O or if $O = m\chi_O$.

If O is equipped with a metric coming from invariant Riemannian metrics on the local models $\tilde{U}$, then one may easily derive the Gauss-Bonnet theorem,

13.3.5.
$$\int_O K\,dA = 2\pi\chi(O).$$

One way to prove this is by excising small neighborhoods of the singular locus, and applying the usual Gauss-Bonnet theorem for manifolds with boundary. For O to have an elliptic, parabolic or hyperbolic structure, $\chi(O)$ must be respectively positive, zero or negative. If O is elliptic or hyperbolic, then area $(O) = 2\pi|\chi(O)|$.

Theorem 13.3.6. *A closed two-dimensional orbifold has an elliptic, parabolic or hyperbolic structure if and only if it is good. An orbifold O has a hyperbolic structure if and only if $\chi(O) < 0$, and a parabolic structure if and only if $\chi(O) = 0$. An orbifold* 13.21
is elliptic or bad if and only if $\chi(O) > 0$.

All bad, elliptic and parabolic orbifolds are tabulated below, where

$$(n_1, \dots, n_k; m_1, \dots, m_l)$$

denotes an orbifold with elliptic points of orders $n_1, \dots, n_k$ (in ascending order) and corner reflectors of orders $m_1, \dots, m_l$ (in ascending order). Orbifolds not listed are hyperbolic.

- Bad orbifolds:
  - $X_O = S^2$: (n), (n_1, n_2) with $n_1 < n_2$.
  - $X_O = D^2$: $(\,;n)$, $(\,;n_1, n_2)$ with $n_1 < n_2$.
- Elliptic orbifolds:
  - $X_O = S^2$: $(\,)$, (n, n), $(2, 2, n)$, $(2, 3, 3)$, $(2, 3, 4)$, $(2, 3, 5)$.
  - $X_O = D^2$: $(\,;\,)$, $(\,;n, n)$, $(\,;2, 2, n)$, $(\,;2, 3, 3)$, $(\,;2, 3, 4)$, $(\,;2, 3, 5)$, $(n;\,)$, $(2;m)$, $(3;2)$.
  - $X_O = \mathbb{P}^2$: $(\,)$, (n).
- Parabolic orbifolds:
  - $X_O = S^2$: $(2, 3, 6)$, $(2, 4, 4)$, $(3, 3, 3)$, $(2, 2, 2, 2)$.
  - $X_O = D^2$: $(\,;2, 3, 6)$, $(\,;2, 4, 4)$, $(\,;3, 3, 3)$, $(\,;2, 2, 2, 2)$, $(2;2, 2)$, $(3;3)$, $(4;2)$, $(2;2;\,)$.
  - $X_O = \mathbb{P}^2$: $(2, 2)$.
  - $X_O = T^2$: $(\,)$
  - $X_O =$ Klein bottle: $(\,)$
  - $X_O =$ annulus: $(\,;\,)$
  - $X_O =$ Möbius band: $(\,;\,)$

13.21.a

The universal covering space of $D^2_{(;4,4,4)}$ and $S^2_{(4,4,4)} \cdot \pi_1(D^2_{(;4,4,4)})$ is generated by reflections in the faces of one of the triangles. The full group of symmetries of this tiling of H^2 is $\pi_1(D^2_{(;2,3,8)})$.

This picture was drawn with a computer by Peter Oppenheimer.

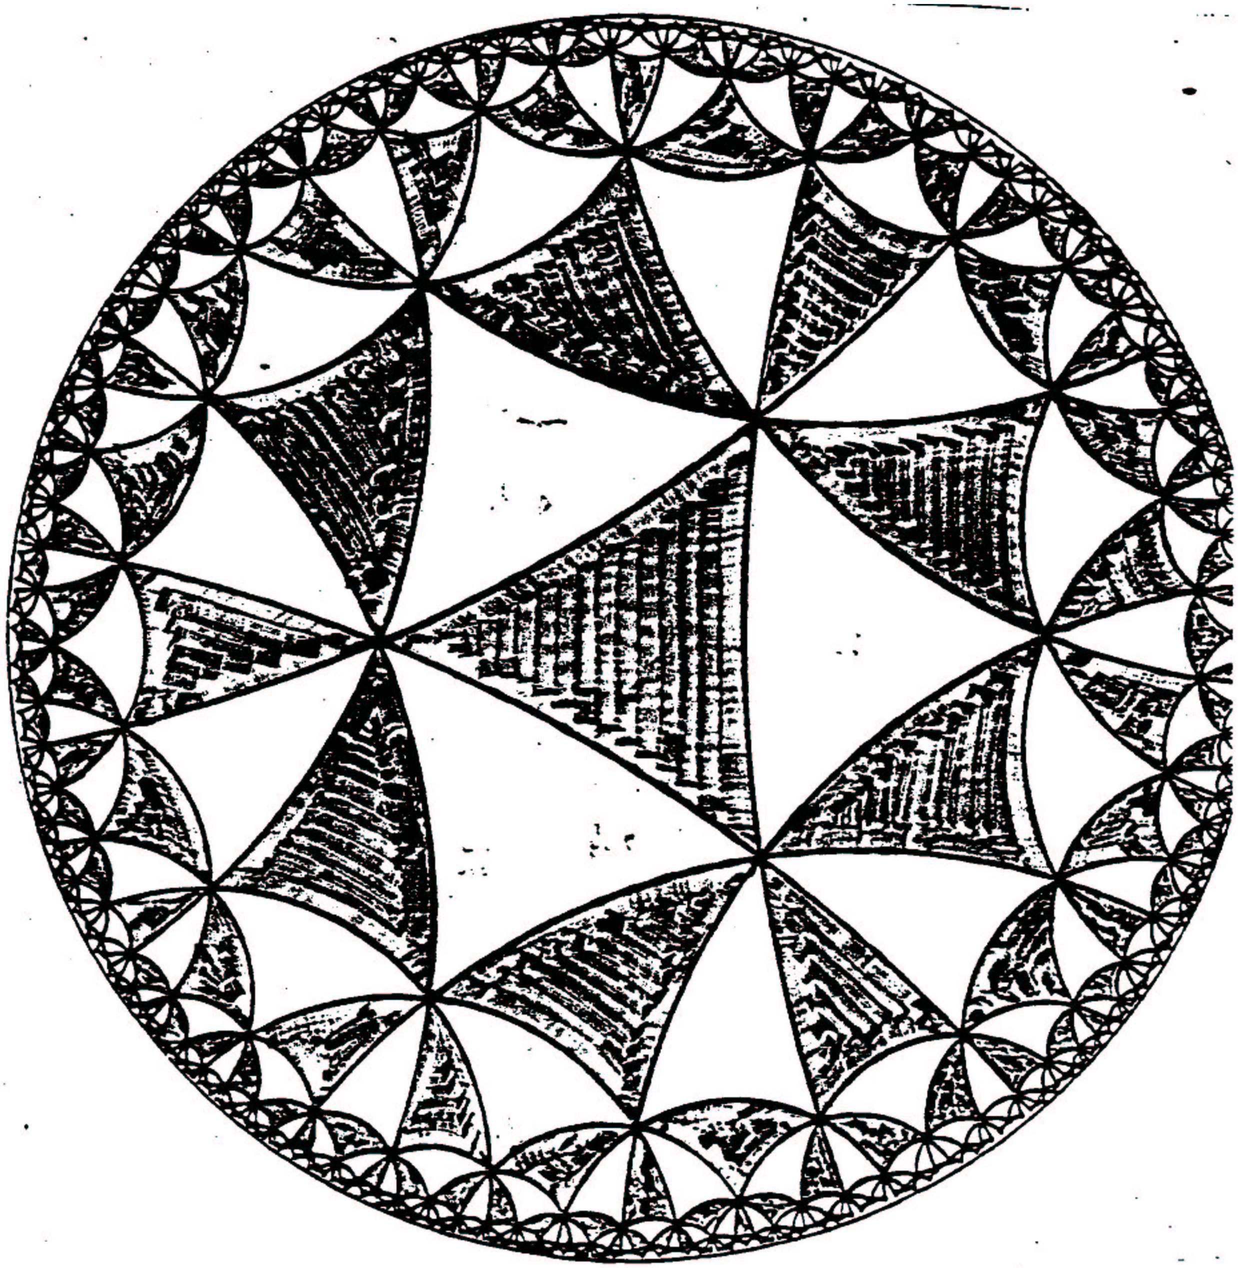

PROOF. It is routine to list all orbifolds with non-negative Euler number, as in 13.22
the table. We have already indicated an easy, direct argument to show the orbifolds listed as bad are bad; here is another. First, by passing to covers, we only need consider the case that the underlying space is S^2, and that if there are two elliptic

points their orders are relatively prime. These orbifolds have Riemannian metrics of curvature bounded above zero,

which implies (by elementary Riemannian geometry) that any surface covering them must be compact. But the Euler number is either $1 + 1/n$ or $1/n_1 + 1/n_2$, which is a rational number with numerator > 2.

Since no connected surface has an Euler number greater than 2, these orbifolds must be bad.

Question. What is the best pinching constant for Riemannian metrics on these orbifolds?

All the orbifolds listed as elliptic and parabolic may be readily identified as the quotient of S^2 or E^2 modulo a discrete group. The 17 parabolic orbifolds correspond to the 17 "wallpaper groups." The reader should unfold these orbifolds for himself, 13.23
to appreciate their beauty. Another pleasant exercise is to identify the orbifolds associated with some of Escher's prints.

Hyperbolic structures can be found, and classified, for orbifolds with negative Euler characteristics by decomposing them into primitive pieces, in a manner analogous to our analysis of Teichmüller space for a surface (§5.3). Given an orbifold O with $\chi(O) < 0$, we may repeatedly cut it along simple closed curves and then "mirror" these curves (to remain in the class of closed orbifolds) until we are left with pieces of the form below. (If the underlying surface is unoriented, then make the first cut so the result is oriented.)

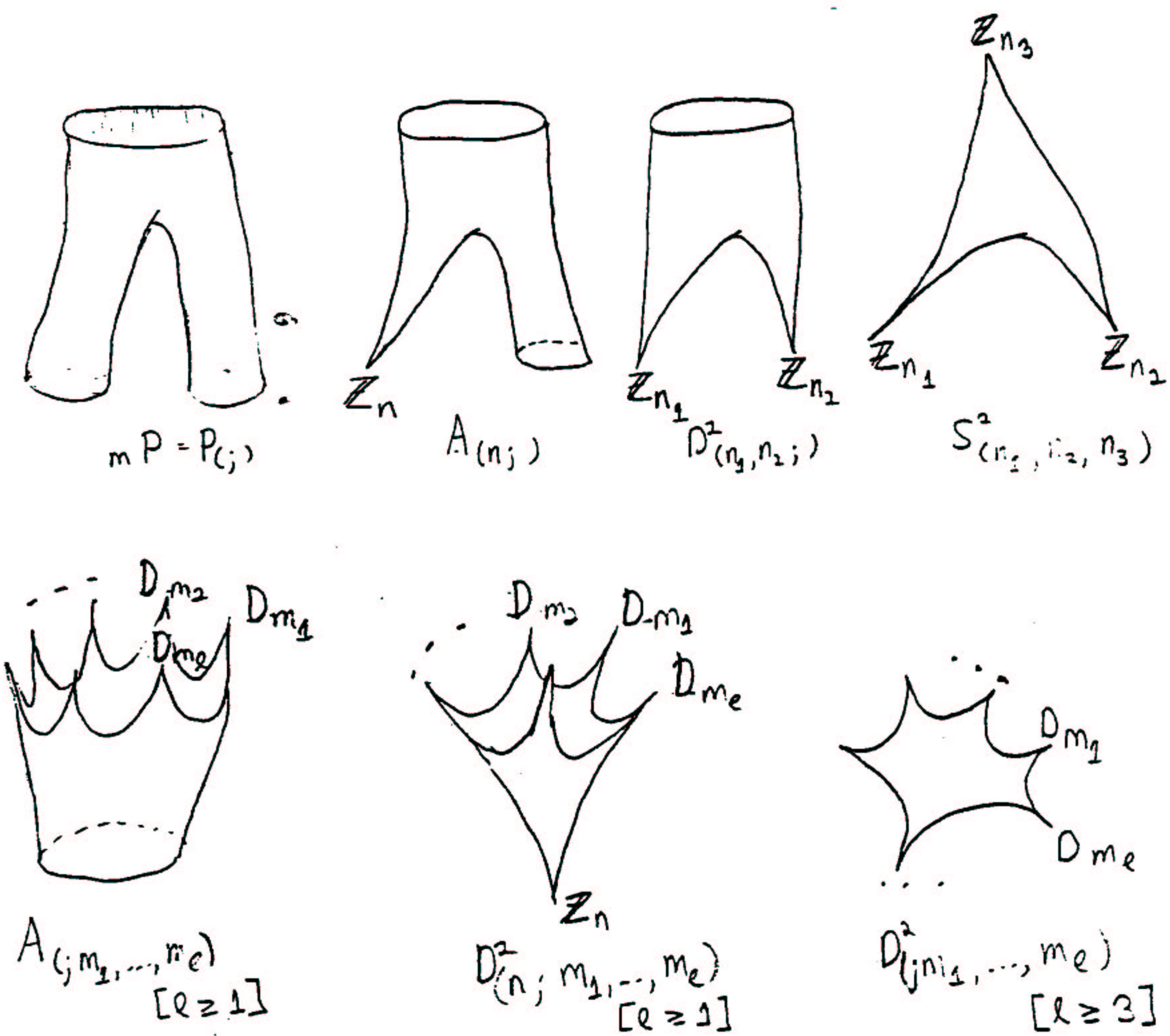

13.24

The orbifolds mP, $A_{(n\,;)}$ and $D_{(n_1,n_2\,;)}$ (except the degenerate case $A_{(2,2;)}$) and $S^2_{(n_1,n_2,n_3)}$ have hyperbolic structures paremetrized by the lengths of their boundary components. The proof is precisely analogous to the classification of shapes of pants in §5.3; one decomposes these orbifolds into two congruent “generalized triangles” (see §2.6).

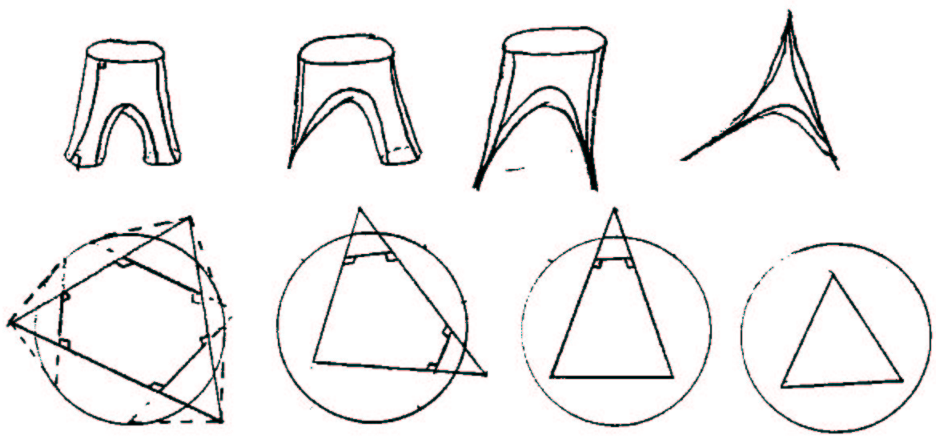

The orbifold $D^2_{(;m_1,\dots,m_l)}$ also can be decomposed into "generalized triangles,"

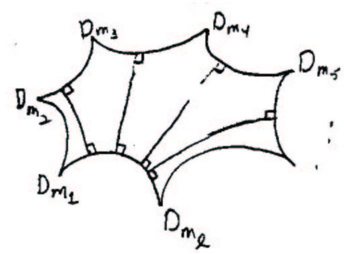

for instance in the pattern above. One immediately sees that the orbifold has hyperbolic structures (provided $\chi < 0$) parametrized by the lengths of the cuts; that is, $(\mathbb{R}_+)^{l-3}$. Special care must be taken when, say, $m_1 = m_2 = 2$. Then one of 13.25
the cuts must be omitted, and an edge length becomes a parameter. In general any disjoint set of edges with ends on order 2 corner reflectors can be taken as positive real parameters, with extra parameters coming from cuts not meeting these edges:

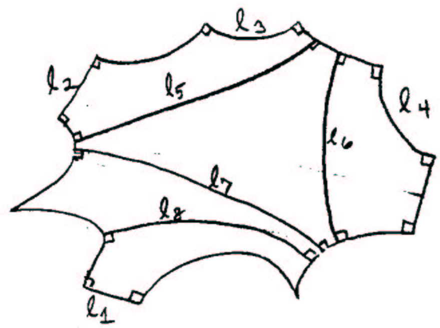

The annulus with more than one corner reflector on one boundary component should be dissected, as below, into $D_{(\,;n_1,\dots,n_k)}$ and an annulus with two order two corner reflectors. $D^2_{(n;\,m_1,\dots,m_l)}$ is analogous.

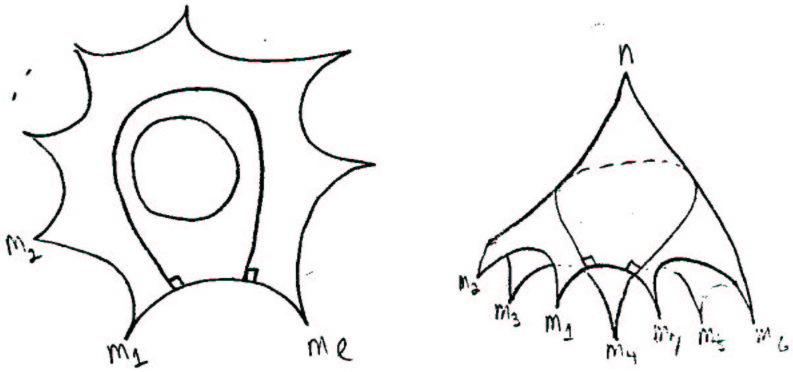

13.26

Hyperbolic structures on an annulus with two order two corner reflectors on one boundary component are parametrized by the length of the other boundary component, and the length of one of the edges:

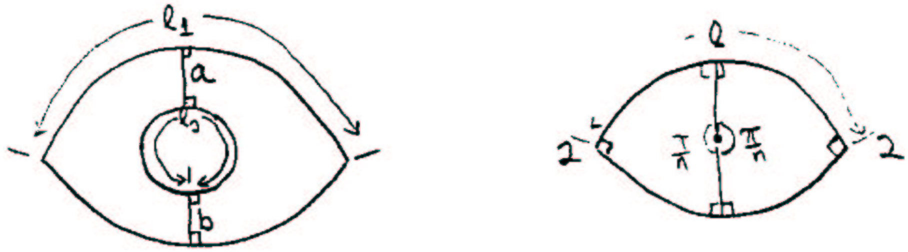

(The two all right pentagons agree on a and b, so they are congruent; thus they are determined by their edges of length $l_1/2$ and $l_2/2$). Similarly, $D^2_{(n;\,2,2)}$ is determined by one edge length, provided $n > 2$. $D^2_{(2;\,2,2)}$ is not hyperbolic. However, it has a degenerate hyperbolic structure as an infinitely thin rectangle, modulo a rotation of order 2—or, an interval.

This is consistent with the way in which it arises in considering hyperbolic structures, in the dissection of $D^2_{(2;m_1,\dots,m_l)}$. One can cut such an orbifold along the perpendicular arc from the elliptic point to an edge, to obtain $D^2_{(\,;\,2,2,m_1,\dots,m_l)}$. In the case of an annulus with only one corner reflector, 13.27

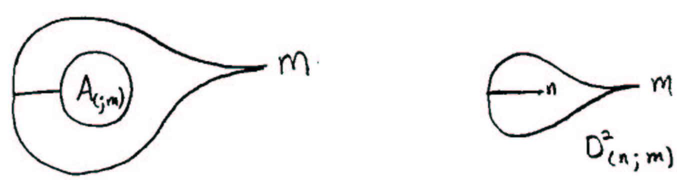

note first that it is symmetric, since it can be dissected into an isosceles "triangle." Now, from a second dissection, we see hyperbolic structures are paremetrized by the length of the boundary component without the reflector.

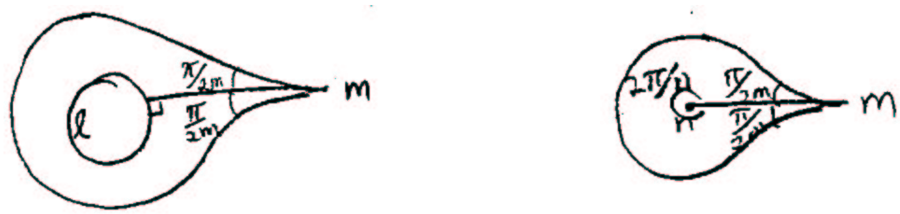

By the same argument, $D^2_{(n;m)}$ has a unique hyperbolic structure.

All these pieces can easily be reassembled to give a hyperbolic structure on O. □

From the proof of 13.3.6 we derive

Corollary 13.3.7. *The Teichmüller space $\mathcal{T}(O)$ of an orbifold O with $\chi(O) < 0$ is homeomorphic to Euclidean space of dimension $-3\chi(X_O) + 2k + l$, where k is the number of elliptic points and l is the number of corner reflectors.*

Proof. O can be dissected into primitive pieces, as above, by cutting along disjoint closed geodesics and arcs perpendicular to ∂X_O: i.e., one-dimensional hyperbolic 13.28
suborbifolds. The lengths of the arcs, and lengths and twist parameters for simple closed curves form a set of parameters showing that $\mathcal{T}(O)$ is homeomorphic to Euclidean space of some dimension. The formula for the dimension is verified directly for the primitive pieces, and so for disjoint unions of primitive pieces. When two circles are glued together, neither the formula nor the dimension of the Teichmüller space changes—two length parameters are replaced by one length parameter and one first parameter. When two arcs are glued together, one length parameter is lost, and the formula for the dimension decreases by one. □

13.4. Fibrations.

There is a very natural way to define the tangent space $T(O)$ of an orbifold O. When the universal cover $\tilde{O}$ is a manifold, then the covering transformations act on $T(\tilde{O})$ by their derivatives. $T(O)$ is then $T(\tilde{O})/\pi_1(O)$. In the general case, O is made up of pieces covered by manifolds, and the tangent space of O is pieced together from the tangent space of the pieces. Similarly, any natural fibration over manifolds gives rise to something over an orbifold.

Definition 13.4.1. A *fibration*, E, with generic fiber F, over an orbifold O is an orbifold with a projection

$$p : X_E \to X_O$$

between the underlying spaces, such that each point $x \in O$ has a neighborhood $U = \tilde{U}/\Gamma$ (with $\tilde{U} \subset \mathbb{R}^n$) such that for some action of Γ on F, $p^{-1}(U) = \tilde{U} \times F/\Gamma$ (where Γ acts by the diagonal action). The product structure should of course be consistent with p: the diagram below must commute. 13.29

$$\begin{array}{ccc} \tilde{U} \times F & \rightarrow & p^{-1}(U) \\ \downarrow & & \downarrow \\ \tilde{U} & \longrightarrow & U \end{array}$$

With this definition, natural fibrations over manifolds give rise to natural fibrations over orbifolds.

The *tangent sphere bundle* $TS(M)$ is the fibration over M with fiber the sphere of rays through O in $T(M)$. When M is Riemannian, this is identified with the unit tangent bundle $T_1(M)$.

Proposition 13.4.2. *Let O be a two-orbifold. If O is elliptic, then $T_1(O)$ is an elliptic three-orbifold. If O is Euclidean, then $T_1(O)$ is Euclidean. If O is bad, then $TS(O)$ admits an elliptic structure.*

Proof. The unit tangent bundle $T_1(S^2)$ can be identified with the grup SO_3 by picking a "base" tangent vector V_O and parametrizing an element $g \in SO_3$ by the image vector $Dg(V_O)$. SO_3 is homeomorphic to $\mathbb{P}^3$, and its universal covering group os S^3. This correspondence can be seen by regarding S^3 as the multiplicative group of unit quaternions, which acts as isometries on the subspace of purely imaginary quaternions (spanned by i, j and k) by conjugation. The only elements acting trivially are ± 1. The action of SO_3 on $T_1(S^2) = SO_3$ corresponds to left translation so that for an orientable $O = S^2/\Gamma$, $T_1(O) = T_1(S^2/\Gamma) = \Gamma \backslash SO_3 = \tilde{\Gamma} \backslash S^3$ is clearly elliptic. Here $\tilde{\Gamma}$ is the preimage of Γ in S^3. (Whatever Γ stands for, $\tilde{\Gamma}$ is generally called "the binary Γ"—e.g., the binary dodecahedral group, etc.)

When O is not oriented, then we use the model $T_1(S^2) = O_3/\mathbb{Z}_2$, where $\mathbb{Z}_2$ is generated by the reflection r through the geodesic determined by V_O. Again, the 13.29
action of O_3 on $T_1(S^2)$ comes from left multiplication on $O_3/\mathbb{Z}_2$. An element gr, with $g \in SO_3$, thus takes $g'V_O$ to $grg'rV_O$. But $rg'r = sg's$, where $s \in SO_3$ is $180°$ rotation of the geodesic through V_O, so the corresponding transformations of S^3,

$\tilde{g} \mapsto (\tilde{g}\tilde{s})\, \tilde{g}'\, (\tilde{s})$, are compositions of left and right multiplication, hence isometries.

For the case of a Euclidean orbifold O, note that T_1E^2 has a natural product structure as $E^2 \times S^1$. From this, a natural Euclidean structure is obtained on T_1E^2, hence on $T_1(O)$.

The bad orbifolds are covered by orbifolds $S^2_{(n)}$ or $S^2_{(n_1,n_2)}$. Then $TS(H)$, where H is either hemisphere, is a solid torus, so the entire unit tangent space is a lens space—hence it is elliptic. $TS(D^2_{(\,;n)})$, or $TSD^2_{(\,;n_1,n_2)}$, is obtained as the quotient by a $\mathbb{Z}_2$ action on these lens spaces. $\square$

As an example, $T_1(S^2_{(2,3,5)})$ is the Poincaré dodecahedral space. This follows immediately from one definition of the Poincaré dodecahedral space as S^3 modulo the binary dodecahedral group. In general, observe that $TS(O^2)$ is always a manifold if O^2 is oriented; otherwise it has elliptic axes of order 2, lying above mirrors and consisting of vectors tangent to the mirrors. In more classical terminology, the Poincaré dodecahedral space is a Seifert fiber space over S^2 with three singular fibers, of type $(2,1)$, $(3,1)$ and $(5,1)$.

When O has the combinatorial type of a polygon, it turns out that $X_{TS(O)}$ is S^3, with singular locus a certain knot or two-component link. There is an a priori reason to suspect that $X_{TS(O)}$ be S^3, since $\pi_1 O$ is generated by reflections. These reflections 13.31
have fixed points when they act on $TS(\tilde{O})$, so $\pi_1(X_{TS(O)})$ is the surjective image of $\pi_1 TS(\tilde{O})$. The image is trivial, since a reflection folds the fibers above its axis in half. Every easily producible simply connected closed three-manifold seems to be S^3. We can draw the picture of $TS(O)$ by piecing.

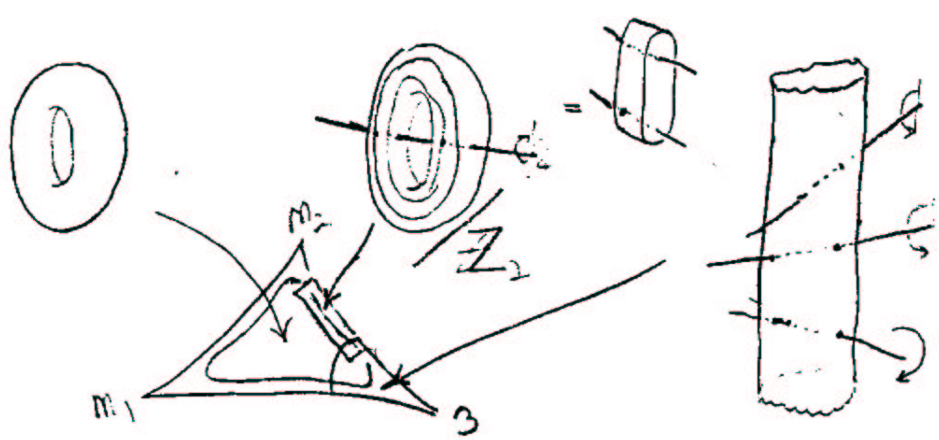

Over the non-singular part of O, we have a solid torus. Over an edge, we have $mI \times I$, with fibers folded into mI; nearby figures go once around these mI's. Above a corner reflector of order n, the fiber is folded into mI. The fibers above the nearby edges weave up and down n times, and nearby circles wind around $2n$ times.

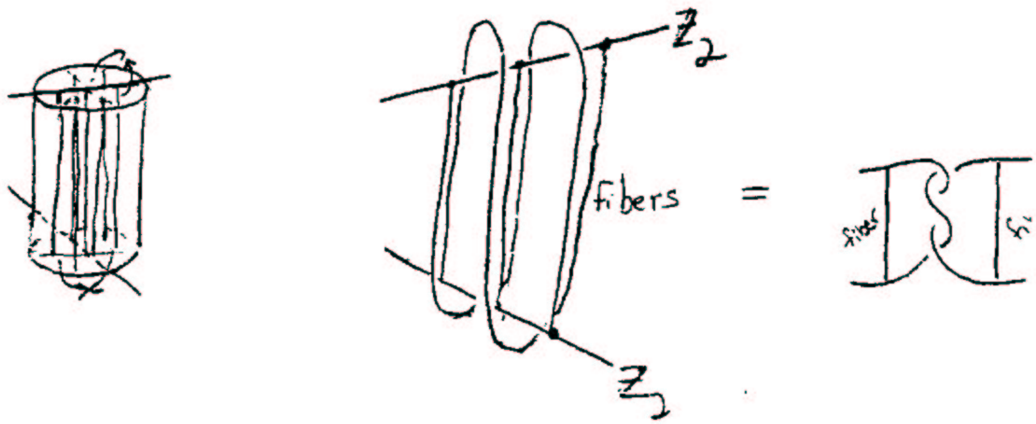

13.32

When the pieces are assembled, we obtain this knot or link:

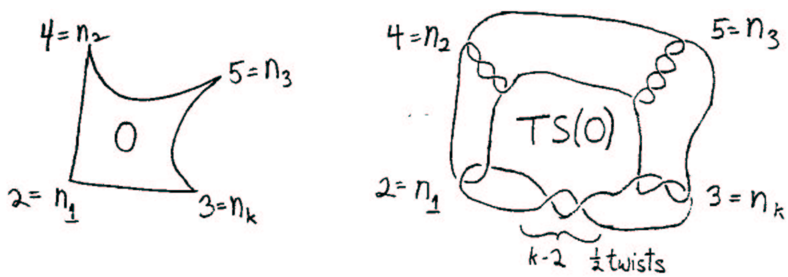

When O is a Riemannian orbifold, this gives $T_1(O)$ a canonical flow, the geodesic flow. For the Euclidean orbifolds with X_O a polygon, this flow is physically realized (up to friction and spin) by the motion of a billiard ball. The flow is tangent to the singular locus. Thus, the phase space for the familiar rectangular billiard table is S^3:

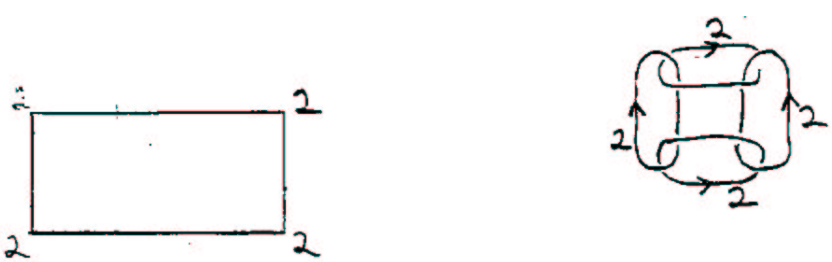

There are two invariant annuli, with boundary the singular locus, corresponding to trajectories orthogonal to a side. The other trajectories group into invariant tori.

Note the two-fold symmetry in the tangent space of a billiard table, which in the picture is 180° rotation about the axis perpendicular to the paper. The quotient orbifold is the same as example 13.1.5. 13.33

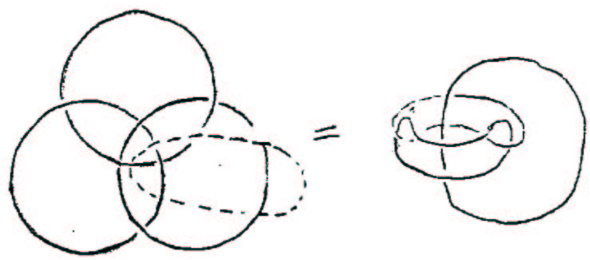

You can obtain many other examples via symmetries and covering spaces. For instance, the Borromean rings above have a three-fold axis of symmetry, with quotient orbifold:

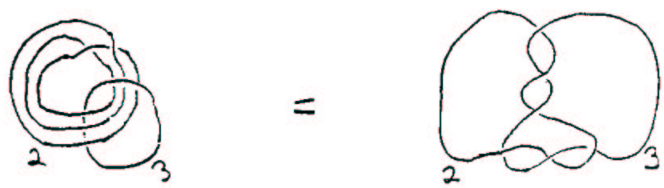

We can pass to a two-fold cover, unwrapping around the $\mathbb{Z}_3$ elliptic axis, to obtain the figure-eight knot as a $\mathbb{Z}_3$ elliptic axis.

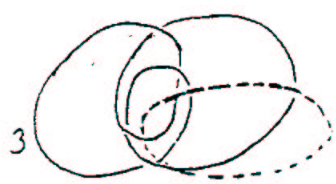

This is a Euclidean orbifold, whose fundamental group is generated by order 3 rotations in main diagonals of two adjacent cubes (regarded as fundamental domains for example 13.1.5).

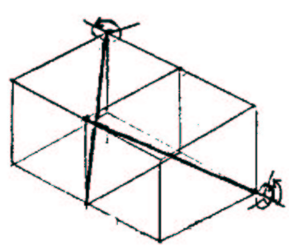

When O is elliptic, then all geodesics are closed, and the geodesic flow comes from a circle action. It follows that $T_1(O)$ is a fibration in a second way, by projecting to the quotient space by the geodesic flow! For instance, the singular locus of $T_1(D^2_{(2,3,5)})$ is a torus knot of type $(3,5)$: 13.34

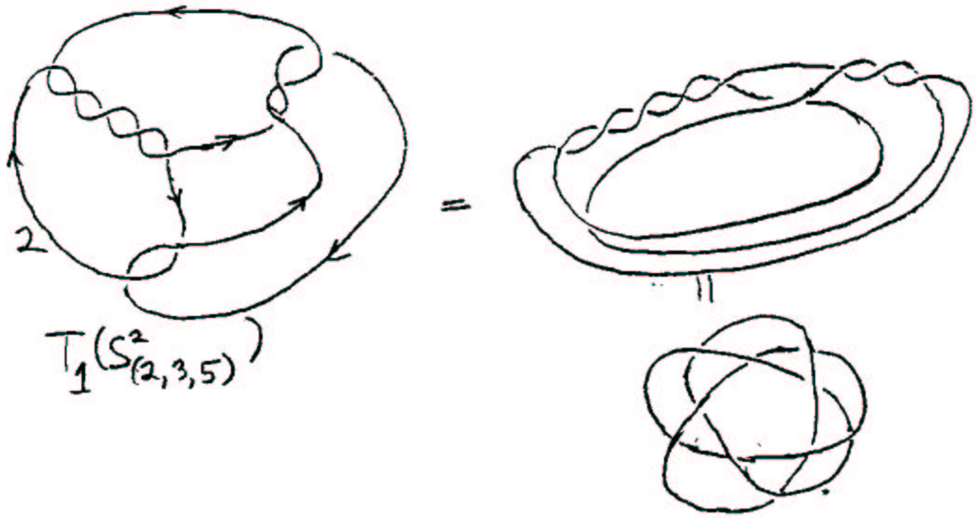

Therefore, it also fibers over $S^2_{(2,3,5)}$. In general, an oriented three-orbifold which fibers over a two-orbifold, with general fiber a circle, is determined by three kinds of information:

(a) The base orbifold.
(b) For each elliptic point or corner reflector of order n, an integer $0 \le k < n$ which specifies the local structure. Above an elliptic point, the $\mathbb{Z}_n$ action on $\tilde{U} \times S^1$ is generated by a $1/n$ rotation of the disk U and a k/n rotation of the fiber S^1. Above a corner reflector, the D_n action on $\tilde{U} \times S^1$ (with S^1 taken as the unit circle in $\mathbb{R}^2$) is generated by reflections of $\tilde{U}$ in lines making an angle of π/n and reflections of S^1 in lines making an angle of $k\pi/n$.
(c) A rational-valued Euler number for the fibration. This is defined as the obstruction to a rational section—i.e., a multiple-valued section, with rational weights for the sheets summing to one. (This is necessary, since there is not usually even a local section near an elliptic point or corner reflector).

The Euler number for $TS(O)$ equals $\chi(O)$. It can be shown that a fibration of non- 13.35
zero Euler number over an elliptic or bad orbifold is elliptic, and a fibration of zero Euler number over a Euclidean orbifold is Euclidean.

13.5. Tetrahedral orbifolds.

The next project is to classify orbifolds whose underlying space is a three-manifold with boundary, and whose singular locus is the boundary. In particular, the case when X_O is the three-disk is interesting—the fundamental group of such an orbifold (if it is good) is called a *reflection group*. It turns out that the case when O has the combinatorial type of a tetrahedron is quite different from the general case. Geometrically, the case of a tetrahedron is subtle, although there is a simple way to classify such orbifolds with the aid of linear algebra.

The explanation for this distinction seems to come from the fact that orbifolds of the type of a simplex are non-Haken. First, we define this terminology.

A closed three-orbifold is *irreducible* if it has no bad two-suborbifolds and if every two-suborbifold with an elliptic structure bounds a three-suborbifold with an elliptic structure. Here, an elliptic orbifold with boundary is meant to have totally geodesic boundary—in other words, it must be D^3/Γ, for some $\Gamma \subset O_3$. (For a non-oriented three-manifold,this definition entails being irreducible *and* $\mathbb{P}^2$-*irreducible*, in the usual terminology.) Observe that any three-dimensional orbifold with a bad suborbifold must itself be bad—it is conjectured that this is a necessary and sufficient condition for badness. 13.36

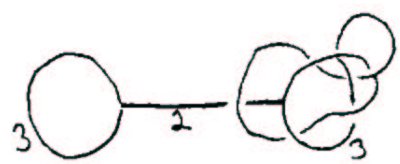

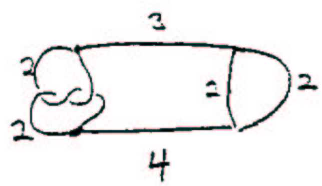

Frequently in three dimensions it is easy to see that certain orbifolds are good but hard to prove much more about them. For instance, the orbifolds with singular locus a knot or link in S^3 are always good: they always have finite abelian covers by manifolds.

Each elliptic two-orbifold is the boundary of exactly one elliptic three-orbifold, which may be visualized as the cone on it.

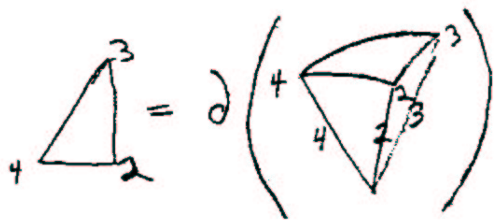

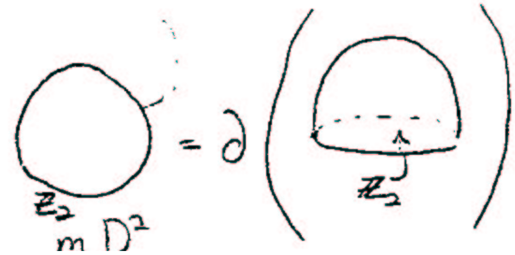

An *incompressible suborbifold* of a three-orbifold O, when X_O is oriented, is a two-suborbifold $O' \subset O$ with $\chi(O') \leq 0$ such that every one-suborbifold $O'' \subset O'$ which bounds an elliptic suborbifold of $O - O'$ bounds an elliptic suborbifold of O'. O is *Haken* if it is irreducible and contains an incompressible suborbifold.

Proposition 13.5.1. *Suppose $X_O = D^3$, $\Sigma_O = \partial D^3$. Then O is irreducible if and only if:*

(a) *The one-dimensional singular locus Σ^1_O cannot be disconnected by the removal of zero, one, or two edges, and* 13.37

(b) *if the removal of γ_1, γ_2 and γ_3 disconnects Σ^1_O, then either they are incident to a commen vertex or the orders n_1, n_2 and n_3 satisfy*

$$1/n_1 + 1/n_2 + 1/n_3 \leq 1.$$

Proof. For any bad or elliptic suborbifold $O' \subset O$, $X_{O'}$ must be a disk meeting Σ^1_O in $1, 2$ or 3 points. $X_{O'}$ separates X_O into two three-disks; one of these gives an elliptic three-orbifold with boundary O' if and only if it contains no one-dimensional parts of Σ_O other than the edges meeting $\partial X_{O'}$. For any set E of edges disconnecting Σ^1_O there is a simple closed curve on ∂X_O meeting only edges in E, meeting such an edge at most once, and separating $\Sigma^1_O - E$. Such a curve is the boundary of a disk in X_O, which determines a suborbifold. Any closed elliptic orbifold S^n/Γ of dimension $n \geq 2$ can be *suspended* to give an elliptic orbifold S^{n+1}/Γ of dimension $n+1$, via the canonical inclusion $O_{n+1} \subset O_{n+2}$. $\square$

Proposition 13.5.2. *An orbifold O with $X_O = D^3$ and $\Sigma_O = \partial D^3$ is Haken if and only if it is irreducible, it is not the suspension of an elliptic two-orbifold and it does not have the type of a tetrahedron.*

Proof. First, suppose that O satisfies the conditions. Let F be any *face* of O, that is a component of Σ_O minus its one dimensional part. The closure $\bar{F}$ is a disk or sphere, for otherwise O would not be irreducible. If F is the entire sphere, then O is the suspension of $D^2_{(;)}$. Otherwise, consider a curve γ going around just outside F, and meeting only edges of Σ^1_O incident to $\bar{F}$. 13.38

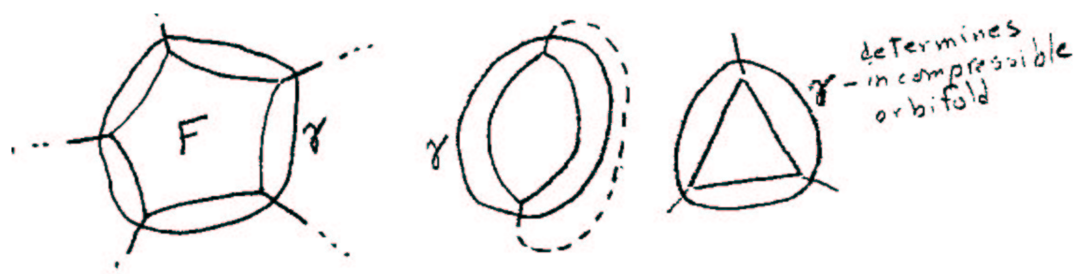

If γ meets no edges, then $\Sigma^1_O = \partial F$ (since O is irreducible) and O is the suspension of $D^2_{(;n,n)}$. The next case is that γ meets two edges of order n; then they must really be the same edge, and O is the suspension of an elliptic orbifold $D^2_{(;n,n_1,n_2)}$. If γ meets three edges, then γ determines a "triangle" suborbifold $D^2_{(;n_1,n_2,n_3)}$ of O. O' cannot be elliptic, for then the three edges would meet at a point and O would have the type of a tetrahedron. Since $D^2_{(;n_1,n_2,n_3)}$ has no non-trivial one-suborbifolds, it is automatically incompressible, so O is Haken. If γ meets four or more edges, then the two-suborbifold it determines is either incompressible or compressible. But if it is compressible, then an automatically incompressible triangle suborbifold of O can be constructed.

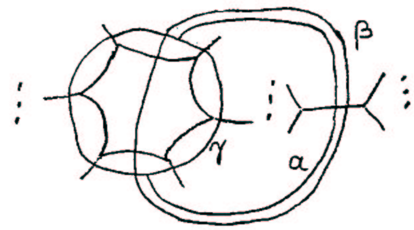

If α determines a "compression," then β determines a triangle orbifold.

The converse assertion, that suspensions of elliptic orbifolds and tetrahedral orbifolds are not Haken, is fairly simple to demonstrate. In general, for a curve γ on 13.39
∂X_O to determine an incompressible suborbifold, it can never enter the same face twice, and it can enter two faces which touch only along their common edge. Such a curve is evidently impossible in the cases being considered. □

There is a system of notation, called the *Coxeter diagram*, which is efficient for describing n-orbifolds of the type of a simplex. The Coxeter diagram is a graph, whose vertices are in correspondence with the $(n-1)$-faces of the simplex. Each pair of $(n-1)$-faces meet on an $(n-2)$-face which is a corner reflector of some order k. The corresponding vertices of the Coxeter graph are joined by $k-2$ edges, or alternatively, a single edge labelled with the integer $k-2$. The notation is efficient because the most commonly occurring corner reflector has order 2, and it is not mentioned. Sometimes this notation is extended to describe more complicated orbifolds with $X_O = D^n$ and $\Sigma_O \subset \partial D^n$, by using dotted lines to denote the faces which are not incident. However, for a complicated polyhedron—even the dodecahedron—this becomes quite unwieldy.

The condition for a graph with $n+1$ vertices to determine an orbifold (of the type of an n-simplex) is that each complete subgraph on n vertices is the Coxeter diagram for an elliptic $(n-1)$-orbifold.

Here are the Coxeter diagrams for the elliptic triangle orbifolds:

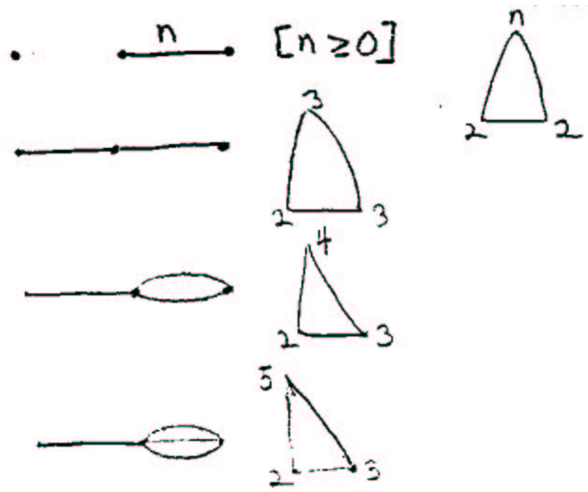

13.40

THEOREM 13.5.3. *Every n-orbifold of the type of a simplex has either an elliptic, Euclidean or hyperbolic structure. The types in the three-dimensional case are listed below:*

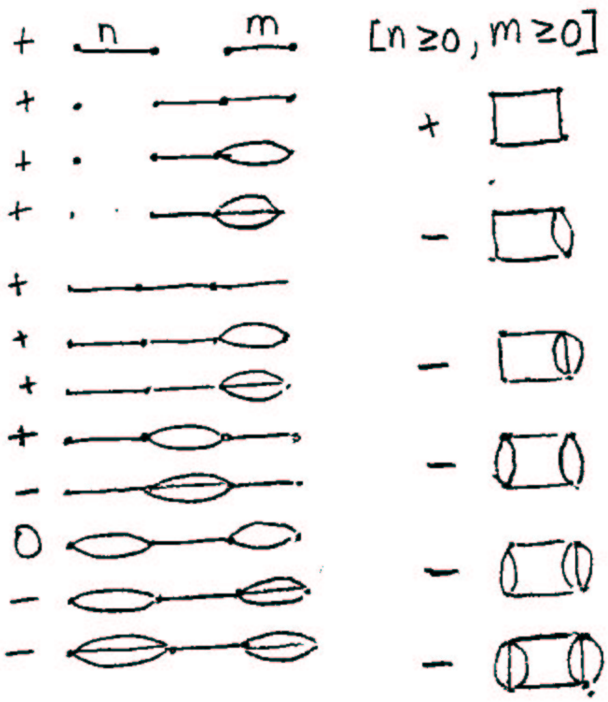

This statement may be slightly generalized to include non-compact orbifolds of the combinatorial type of a simplex with some vertices deleted.

Theorem 13.5.4. *Every n-orbifold which has the combinatorial type of a simplex with some deleted vertices, such that the "link" of each deleted vertex is a Euclidean orbifold, and whose Coxeter diagram is connected, admits a complete hyperbolic structure of finite volume. The three-dimensional examples are listed below:* 13.41

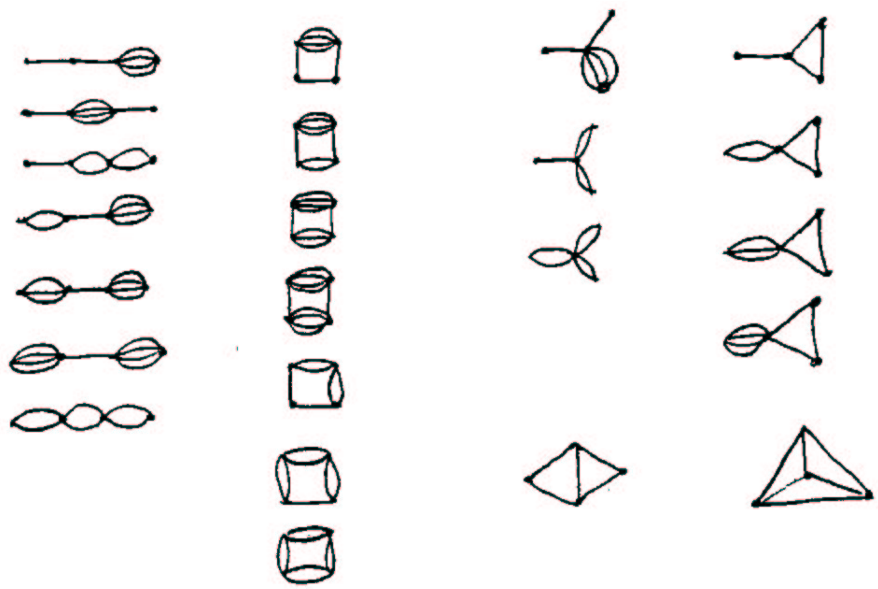

Proof of 13.5.3 and 13.5.4. The method is to describe a simplex in terms of the quadratic form models. Thus, an n-simplex σ^n on S^n has $n+1$ hyperfaces. Each face is contained in the intersection of a codimension one subspace of E^{n+1} with S^n. Let $V_O, \ldots, V_n$ be unit vectors orthogonal to these subspaces in the direction away from σ^n. Clearly, the V_i are linearly indpendent. Note that $V_i \cdot V_i = 1$, and when $i \neq j$, $V_i \cdot V_j = -\cos \alpha_{ij}$, where α_{ij} is the angle between face i and face j. Similarly, each face of an n-simplex in H^n contained in the intersection of a subspace of $E^{n,1}$ with the sphere of imaginary radius $X_1^2 + \cdots + X_n^2 - X_{n+1}^2 = -1$ (with respect to the standard inner product $X \cdot Y = \sum_{i=1}^n X_i \cdot Y_i - X_{n+1} \cdot Y_{n+1}$ on $E^{n,1}$). Outward vectors $V_0, \ldots, V_n$ orthogonal to these subspaces have real length, so they can be normalized to have length 1. Again, the V_i are linearly independent and $V_i \cdot V_j = -\cos \alpha_{ij}$ when $i \neq j$. For an n-simplex σ^n in Euclidean n-space, let $V_O, \ldots, V_n$ be outward unit 13.42
vectors in directions orthogonal to the faces on σ^n. Once again, $V_i \cdot V_j = -\cos \alpha_{ij}$.

Given a collection $\{\alpha_{ij}\}$ of angles, we now try to construct a simplex. For the matrix M of presumed inner products, with l's down the diagonal and $-\cos \alpha_{ij}$'s off the diagonal. If the quadratic form represented by M is positive definite or of type $(n, 1)$, then we can find an equivalence to E^{n+1} or $E^{n,1}$, which sends the basis vectors to vectors $V_O, \ldots, V_n$ having the specified inner product matrix. The intersection of the half-spaces $X \cdot V_i \leq O$ is a cone, which must be non-empty since the $\{V_i\}$ are linearly independent. In the positive definite case the cone intersects S^n in a simplex, whose dihedral angles β_{ij} satisfy $\cos\beta_{ij} = \cos\alpha_{ij}$, hence $\beta_{ij} = \alpha_{ij}$. In the hyperbolic case, the cone determines a simplex in $\mathbb{RP}^n$, but the simplex may not be contained in $H^n \subset \mathbb{RP}^n$. To determine the positions of the vertices, observe that each vertex v_i determines a one-dimensional subspace, whose orthogonal subspace is spanned by $V_O, \dots, \hat{V}_i, \dots, V_n$. The vertex v_i is on H^n, on the sphere at infinity, or outside infinity according to whether the quadratic form restricted to this subspace is positive definite, degenerate, or of type $(n-1, 1)$. Thus, the angles $\{\alpha_{ij}\}$ are the angles of an ordinary hyperbolic simplex if and only if M has type $(n, 1)$, and for each i the submatrix obtained by deleting the ith row and the corresponding column is positive definite. They are the angles of an ideal hyperbolic simplex (with vertices in H^n or S^{n-1}_∞) if and only if all such submatrices are either positive definite, or have rank $n-1$.

By similar considerations, the angles $\{\alpha_{ij}\}$ are the angles of a Euclidean n-simplex if and only if M is positive semidefinite of rank n. 13.43

When the angles $\{\alpha_{ij}\}$ are derived from the Coxeter diagram of an orbifold, then each submatrix of M obtained by deleting the i-th row and the i-th column corresponds to an elliptic orbifold of dimension $n-1$, hence it is positive definite. The full matrix must be either positive definite, of type $(n, 1)$ or positive semidefinite with rank n. It is routine to list the examples in any dimension. The sign of the determinant of M is a practical invariant of the type. We have thus proven theorem 13.5.

In the Euclidean case, it is not hard to see that the subspace of vectors of zero length with respect to M is spanned by $(a_0, \dots, a_n)$, where a_i is the $(n-1)$-dimensional area of the i-th face of σ.

To establish 13.5.4, first consider any submatrix M_i of rank $n-1$ which is obtained by deleting the i-th row and i-th column (so, the link of the i-th vertex is Euclidean). Change basis so that M_i becomes

$$\begin{bmatrix} 1 & & & & 0 \\ & \cdot & & & \\ & & \cdot & & \\ & & & \cdot & \\ & & & 1 & \\ 0 & & & & 0 \end{bmatrix}$$

using $(a_0, \dots, \hat{a}_i, \dots, a_n)$ as the last basis vector. When the basis vector V_i is put back, the quadratic form determined by M becomes

$$\left[\begin{array}{ccc|c} 1 & & 0 & \wr \\ & \ddots & & \wr \\ 0 & & 1 & \wr \\ & & 0 & -C \\ \hline \sim\sim\sim & \sim\sim\sim & -C & 1 \end{array}\right].$$

where $-C = -\sum_{j \ni j \neq i} a_i \cos \alpha_{ij}$ is negative since the Coxeter diagram was supposed to be connected. It follows that M has type $(n, 1)$, which implies that the orbifold is hyperbolic. □

13.44

13.6. Andreev's theorem and generalizations.

There is a remarkably clean statement, due to Andreev, describing hyperbolic reflection groups whose fundamental domains are not tetrahedra.

THEOREM 13.6.1 (Andreev, 1967). (a) *Let O be a Haken orbifold with*

$$X_O = D^3, \quad \Sigma_0 = \partial D^3.$$

Then O has a hyperbolic structure if and only if O has no incompressible Euclidean suborbifolds.

(b) *If O is a Haken orbifold with $X_O = D^3$—(finitely many points) and $\Sigma_O = \partial X_O$, and if a neighborhood of each deleted point is the product of a Euclidean orbifold with an open interval, (but O itself is not such a product) then O has a complete hyperbolic structure with finite volume if and only if each incompressible Euclidean suborbifold can be isotoped into one of the product neighborhoods.*

The proof of 13.6.1 will be given in §??.

COROLLARY 13.6.2. *Let γ be any graph in $\mathbb{R}^2$, such that each edge has distinct ends and no two vertices are joined by more than one edge. Then there is a packing of circles in $\mathbb{R}^2$ whose nerve is isotopic to γ. If γ is the one-skeleton of a triangulation of S^2, then this circle packing is unique up to Moebius transformation.*

A *packing* of circles means a collection of circles with disjoint interiors. The nerve of a packing is then a graph, whose vertices correspond to circles, and whose edges correspond to pairs of circles which intersect. This graph has a canonical embedding in the plane, by mapping the vertices to the centers of the circles and the edges to straight line segments which will pass through points of tangency of circles. 13.44a

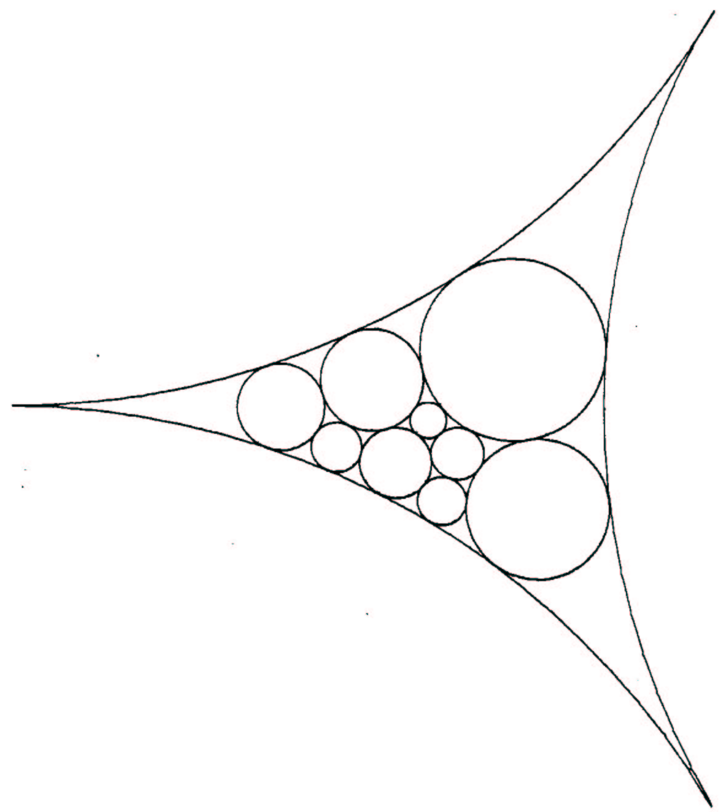

13.45

Proof of 13.6.2. We transfer the problem to S^2 by stereographic projection. Add an extra vertex in each non-triangular region of $S^2 - \gamma$, and edges connecting it to neighboring vertices, so that γ becomes the one-skeleton of a triangulation T of S^2.

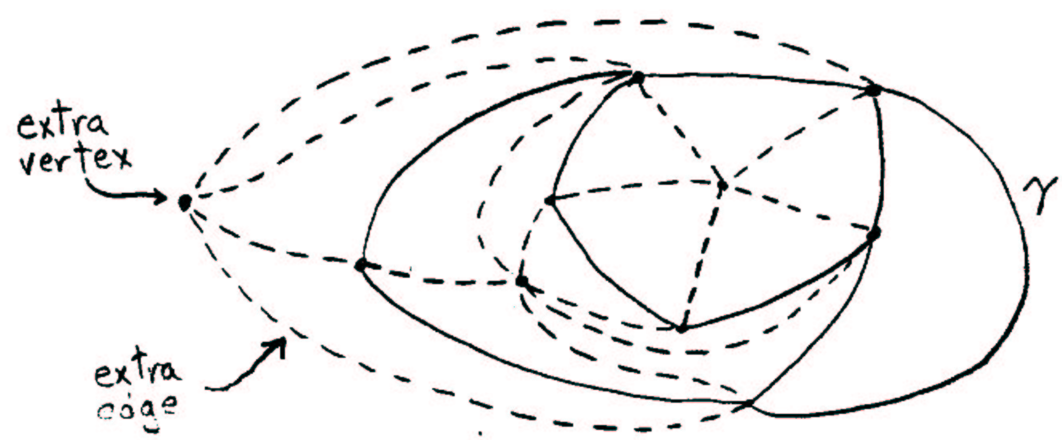

Let P be the polyhedron obtained by cutting off neighborhoods of the vertices of T, down to the middle of each edge of T.

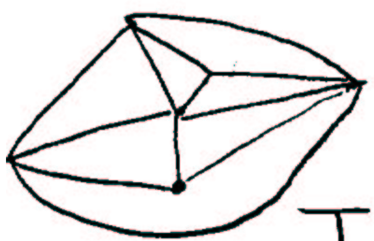

T

P

Let O be the orbifold with underlying space

$$X_O = D^3\text{-vertices of } P, \qquad \text{and} \qquad \Sigma^1_O = \text{ edges of } P,$$

each modelled on $\mathbb{R}^3/D_2$. For any incompressible Euclidean suborbifold O', ∂X_O must be a curve which circumnavigates a vertex. Thus, O satisfies the hypotheses of 13.6.1(b), and O has a hyperbolic structure. This means that P is realized as an ideal polyhedron in H^3, with all dihedral angles equal to 90°. The planes of the new faces of P (faces of P but not T) intersect S^2_∞ in circles. Two of the circles are tangent whenever the two faces meet at an ideal vertex of P. This is the packing required 13.46
by 13.6.2. The uniqueness statement is a consequence of Mostow's theorem, since the polyhedron P may be reconstructed from the packing of circles on S^2_∞. To make the reconstruction, observe that any three pairwise tangent circles have a unique common orthogonal circle. The set of planes determined by the packing of circles on S^2_∞, together with extra circles orthogonal to the triples of tangent circles coming from vertices of the triangular regions of $S^2 - \gamma$ cut out a polyhedron of finite volume combinatorially equivalent to P, which gives a hyperbolic structure for O. □

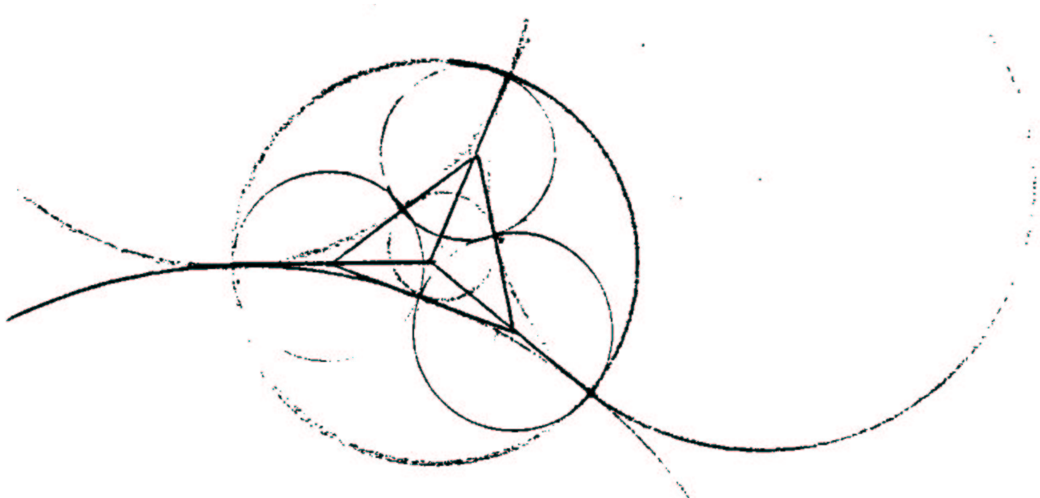

Remark. Andreev also gave a proof of uniqueness of a hyperbolic polyhedron with assigned concave angles, so the reference to Mostow's theorem is not essential.

Corollary 13.6.3. *Let T be any triangulation of S^2. Then there is a convex polyhedron in $\mathbb{R}^3$, combinatorially equivalent to T whose one-skeleton is circumscribed about the unit sphere (i.e., each edge of T is tangent to the unit sphere). Furthermore, this polyhedron is unique up to a projective transformation of $\mathbb{R}^3 \subset \mathbb{P}^3$ which preserves the unit sphere.*

Proof of 13.6.3. Construct the ideal polyhedron P, as in the proof of 13.6.2. Embed H^3 in $\mathbb{P}^3$, as the projective model. The old faces of P (coming from faces of T) form a polyhedron in $\mathbb{P}^3$, combinatorially equivalent to T. Adjust by a projective transformation if necessary so that this polyhedron is in $\mathbb{R}^3$. (To do this, transform P so that the origin is in its interior.) □

Remarks. Note that the dual cell-division T^* to T is also a convex polyhedron in $\mathbb{R}^3$, with one-skeleton of T^* circumscribed about the unit sphere. The intersection $T \cap T^* = P$.

These three polyhedra may be projected to $\mathbb{R}^2 \subset \mathbb{P}^3$, by stereogrpahic projection, from the north pole of $S^2 \subset \mathbb{P}^3$. Stereographic projection is conformal on the tangent space of S^2, so the edges of T^* project to tangents to these circles. It follows that the vertices of T project to the centers of the circles. Thus, the image of the one-skeleton of T is the geometric embedding in $\mathbb{R}^2$ of the nerve γ of the circle packing.

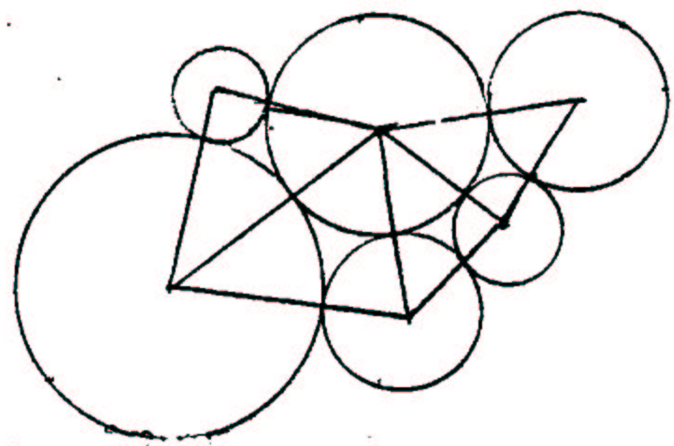

The existence of other geometric patterns of circles in $\mathbb{R}^2$ may also be deduced from Andreev's theorem. For instance, it gives necessary and sufficient condition for the existence of a family of circles meeting only orthogonally in a certain pattern, or 13.48
meeting at 60° angles.

One might also ask about the existence of packing circles on surfaces of constant curvature other than S^2. The answers are corollaries of the following theorems:

THEOREM 13.6.4. *Let O be an orbifold such that $X_O \approx T^2 \times [0,\infty)$, (with some vertices on $T^2 \times O$ having Euclidean links possibly deleted) and $\Sigma_O = \partial X_O$. Then O admits a complete hyperbolic structure of finite volume if and only if it is irreducible, and every incompressible complete, proper Euclidean suborbifold is homotopic to one of the ends.*

(Note that $mS^1 \times [0,\infty)$ is a complete Euclidean orbifold, so the hypothesis implies that every non-trivial simple closed curve on ∂X_O intersects Σ^1_O.)

THEOREM 13.6.5. *Let M^2 be a closed surface, with $\chi(M^2) < 0$. An orbifold O such that $X_O = M^2 \times [0,1]$ (with some vertices on $M^2 \times 0$ having Euclidean links possibly deleted), $\Sigma_O = \partial X_O$ and $\Sigma^1_O \subset M^2 \times O$. Then O has a hyperbolic structure if and only if it is irreducible, and every incompressible Euclidean suborbifold is homotopic to one of the ends.*

By considering $\pi_1 O$, O as in 13.6.4, as a Kleinian group in upper half space with $T^2 \times \infty$ at ∞, 13.6.4 may be translated into a statement about the existence of doubly periodic families of circles in the plane, or 13.48.a

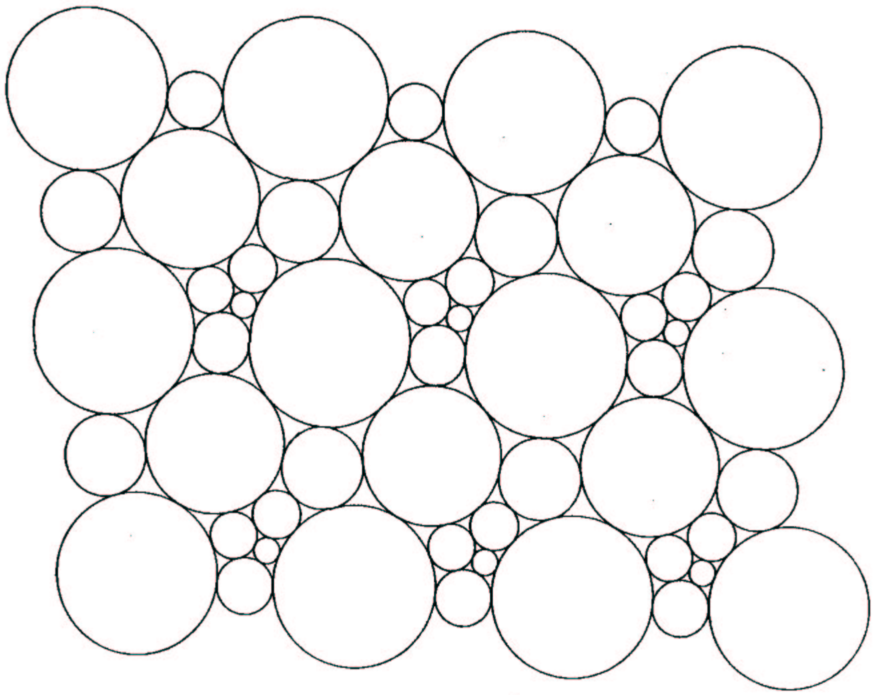

13.48.b

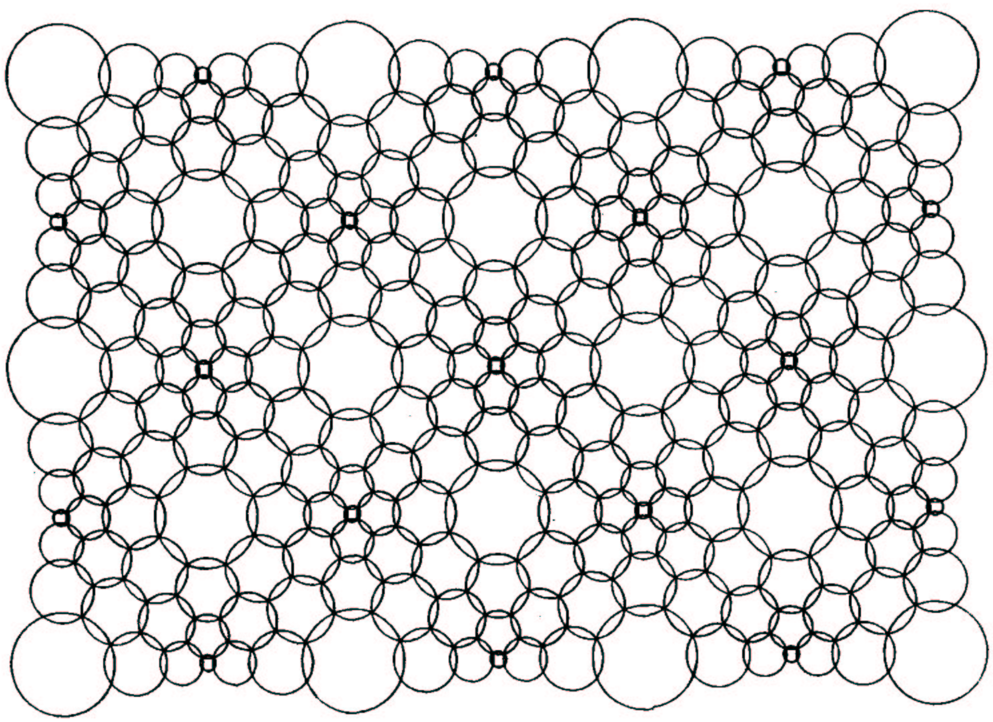

13.49

families of circles on flat toruses. Similarly, 13.6.5 is equivalent to a statement about families of circles in hyperbolic structures for M^2; in fact, since $M^2 \times 1$ has no one-dimensional singularities, it must be totally geodesic in any hyperbolic structure, so $\pi_1 M^2$ acts as a Fuchsian group. The face planes of $M^2 \times O$ give rise to a family of circles in the northern hemisphere of S^2_∞, invariant by this Fuchsian group, so each face corresponds to a circle in the hyperbolic structure for M^2.

Theorems 13.6.1, 13.6.4 and 13.6.5 will be proved in the next section, by studying patterns of circles on surfaces.

In example 13.1.5 we saw that the Borromean rings are the singular locus for a Euclidean orbifold, in which they are elliptic axes of order 2. With the aid of Andreev's theorem, we may find all hyperbolic orbifolds which have the Borromean rings as singular locus. The rings can be arranged so they are invariant by reflection in three orthogonal great spheres in S^3. (Compare p. 13.4.)

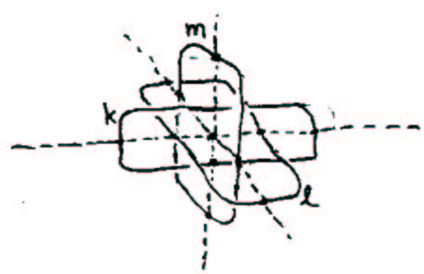

Thus, an orbifold O having the rings as elliptic axes of orders k, l and m is an eight-fold covering space of another orbifold, which has the combinatorial type of a cube.

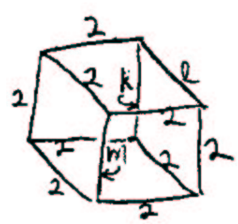

13.50

By Andreev's theorem, such an orbifold has a hyperbolic structure if and only if k, l and m are all greater than 2. If k is 2, for example, then there is a sphere in S^3 separating the elliptic axes of orders l and m and intersecting the elliptic axis of order 2 in four points. This forms an incompressible Euclidean suborbifold of O, which breaks O into

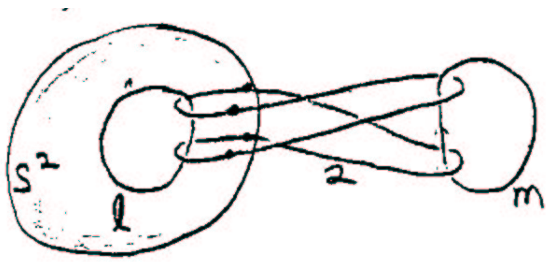

two halves, each fibering over two-orbifolds with boundary, but in incompatible ways (unless l or m is 2).

Base spaces of the fibrations

When $k = l = m = 4$, the fundamental domain, as in example 13.1.5, for $\pi_1 O$ acting on H^3 is a regular right-angled dodecahedron.

Any of the numbers k, l or m can be permitted to take the value ∞ in this discussion, to denote a parabolic cusp. When $l = m = \infty$, for instance, then O has a k-fold cover which is the complement of the untwisted $2k$-link chain D_{2k} of 6.8.7.

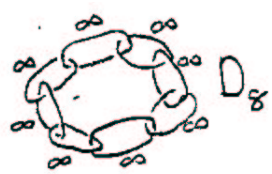

13.51

13.7. Constructing patterns of circles.

We will formulate a precise statement about patterns of circles on surfaces of non-positive Euler characteristic which gives theorems 13.6.4 and 13.6.5 as immediate consequences.

Theorem 13.7.1. *Let S be a closed surface with $\chi(S) \leq 0$. Let τ be a cell-division of S into cells which are images of immersions of triangles and quadrangles which lift to embeddings in $\tilde{S}$. Let $\Theta : \mathcal{E} \to [0, \pi/2]$ (where $\mathcal{E}$ denotes the set of edges of τ) be any function satisfying the conditions below:*

(i) *$\Theta(e) = \pi/2$ if e is an edge of a quadrilateral of τ.*

(ii) *If e_1, e_2, e_3 $[e_i \in \mathcal{E}]$ form a null-homotopic closed loop, and if $\sum_{i=1}^{3} \Theta(e_i) \geq \pi$, then these three edges form the boundary of a triangle of τ.*

(iii) *If* e_1, e_2, e_3, e_4 *form a null-homotopic closed loop and if* $\sum_{i=1}^{4} \Theta(e_i) = 2\pi (\Leftrightarrow \Theta(e_i) = \pi/2)$*, then the* e_i *form the boundary of a quadrilateral or of the union of two adjacent triangles.*

Then there is a metric of constant curvature on S*, uniquely determined up to a scalar multiple, a uniquely determined geometric cell-division of* S *isotopic to* τ *so that the edges are geodesics, and a unique family of circles, one circle* C_v *for each vertex* v *of* τ*, so that* C_{v_1} *and* C_{v_2} *intersect at a positive angle if and only if* v_1 *and* v_2 *lie on a common edge. The angles in which* C_{v_1} *and* C_{v_2} *meet are determined by the common edges: there is an intersection point of* C_{v_1} *and* C_{v_2} *in a two-cell* σ *if and only if* v_1 *and* v_2 *are vertices of* σ*. If* σ *is a quadrangle and* v_1 *and* v_2 *are diagonally opposite, then* C_{v_1} *is tangent to* C_{v_2}*; otherwise, they meet at an angle of* $\Theta(e)$*, where* 13.52
e *is the edge joining them in* σ*.*

Proof. First, observe that quadrangles can be eliminated by subdivision into two triangles by a new edge e with $\Theta(e) = 0$.

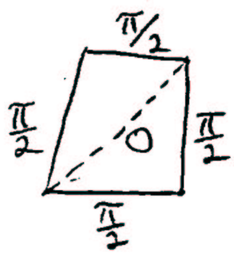

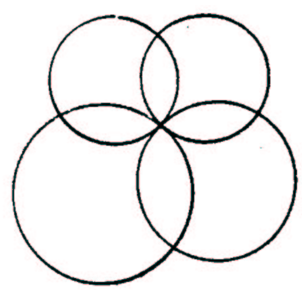

There is an extraneous tangency of circles here—in fact, all extraneous tangencies come from this situation. Henceforth, we assume τ has no quadrangles. The idea is to solve for the *radii* of the circles C_{v_1}. Given an arbitrary set of radii, we shall construct a Riemannian metric on S with cone type singularities at the vertices of τ, which has a family of circles of the given radii meeting at the given angles. We adjust the radii until S lies flat at each vertex. Thus, the proof is closely analogous to the idea that one can make a conformal change of any given Riemannian metric on a surface until it has constant curvature. Observe that a conformal map is one which takes infinitesimal circles to infinitesimal circles; the conformal factor is the ratio of the radii of the target and source circles.

Lemma 13.7.2. *For any three non-obtuse angles* θ_1, θ_2 *and* $\theta_3 \in [0, \pi/2]$ *and any three positive numbers* R_1, R_2*, and* R_3*, there is a configuration of* 3 *circles in both hyperbolic and Euclidean geometry, unique up to isometry, having radii* R_i *and meeting in angles* θ_i*.*

13.53

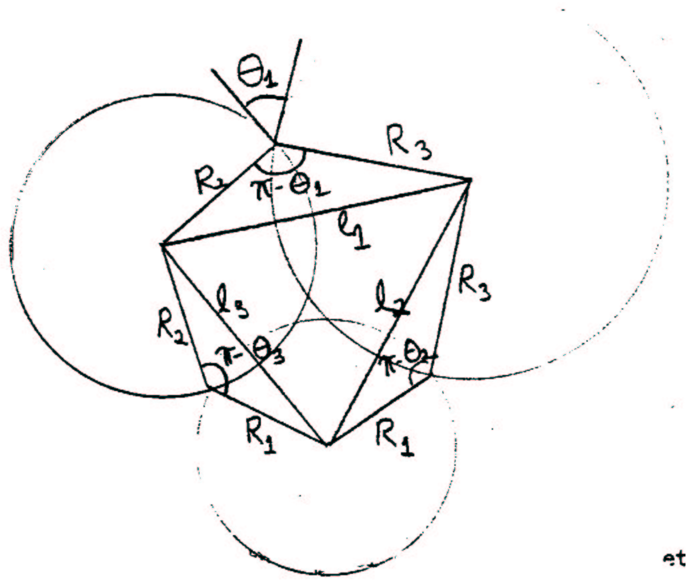

Proof of lemma. The length l_k of a side of the hypothetical triangle of centers of the circles is determined as the side opposite the obtuse angle $\pi - \theta_k$ in a triangle whose other sides are R_i and R_j. Thus, $\sup(R_i, R_j) < l_k \leq R_i + R_j$. The three numbers l_1, l_2 and l_3 obtained in this way clearly satisfy the triangle inequalities $l_k < l_i + l_j$. Hence, one can construct the appropriate triangle, which gives the desired circles. □

Proof of 13.7.1, continued. Let $\mathcal{V}$ denote the set of vertices of τ. For every element $R \in \mathbb{R}_+^{\mathcal{V}}$ (i.e., if we choose a radius for the circle about each vertex), there is a singular Riemannian metric, which is pieced together from the triangles of centers of circles with given radii and angles of intersetcion as in 13.7.2. The triangles are taken in H^2 or E^2 depending on whether $\chi(S) < 0$ or $\chi(S) = 0$. The edge lengths of cells of τ match whenever they are glued together, so we obtain a metric, with singularities only at the vertices, and constant curvature 0 or -1 everywhere else.

The notion of curvature can easily be extended to Riemannian surfaces with certain sorts of singularities. The curvature form $K da$ becomes a measure κ on such 13.54
a surface. Tailors are of necessity familiar with curvature as a measure. Thus, a seam has curvature $(k_1 - k_2) \cdot \mu$, where μ is one-dimensional Lebesgue measure and k_1 and k_2 are the geodesic curvatures of the two halves.

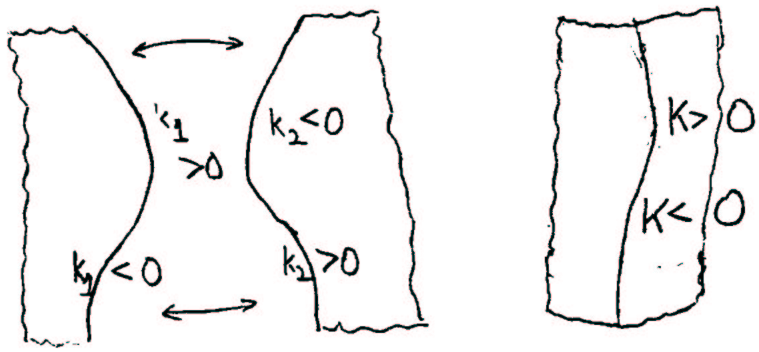

(The effect of gathering is more subtle—it is obtained by putting two lines infinitely close together, one with positive curvature and one with balancing negative curvature. Another instance of this is the boundary of a lens.)

More to the point for us is the curvature concentrated at the apex of a cone: it is $2\pi - \alpha$, where α is the cone angle (computed by splitting the cone to the apex and laying it flat). It is easy to see that this is the unique value consistent with the Gauss-Bonnet theorem.

Formally, we have a map

$$F : \mathbb{R}_+^{\mathcal{V}} \to \mathbb{R}^{\mathcal{V}}.$$

Given an element $R \in \mathbb{R}_+^{\mathcal{V}}$, we construct the singular Riemannian metric on S, as above; $F(R)$ describes the discrete part of the curvature measure κ_R on S, in other 13.55
words, $F(R)(v) = \kappa_R(v)$. Our problem is to show that O is in the image of F, for then we will have a non-singular metric with the desired pattern of circles built in.

When $\chi(S) = 0$, then the shapes of the Euclidean triangles do not change when we multiply R by a constant, so $F(R)$ also does not change. Thus we may as well normalize so that $\sum_{v \in \mathcal{V}} R(v) = 1$. Let $\Delta \subset R_+^{\mathcal{V}}$ be this locus—Δ is the interior of the standard $|\mathcal{V}| - 1$ simplex. Observe, by the Guass-Bonnet theorem, that

$$\sum_{v \in \mathcal{V}} \kappa_R(v) = 0.$$

Let $Z \subset \mathbb{R}^{\mathcal{V}}$ be the locus defined by this equation.

If $\chi(S) < 0$, then changing R by a constant does make a difference in κ. In this case, let $\Delta \subset \mathbb{R}_+^{\mathcal{V}}$ denote the set of R such that the associated metric on S has total area $2\pi\,|\chi(S)|$. By the Gauss-Bonnet theorem, $\Delta = F^{-1}(Z)$ (with Z as above). As one can easily believe, Δ intersects each ray through O in a unique point, so Δ is a simplex in this case also. This fact is easily deduced from the following lemma, which will also prove the uniqueness part of 13.7.1:

Lemma 13.7.3. *Let C_1, C_2 and C_3 be circles of radii R_1, R_2 and R_3 in hyperbolic or Euclidean geometry, meeting pairwise in non-obtuse angles. If C_2 and C_3 are held* *constant but C_1 is varied in such a way that the angles of intersection are constant but R_1 decreases, then the center of C_1 moves toward the interior of the triangle of centers.*

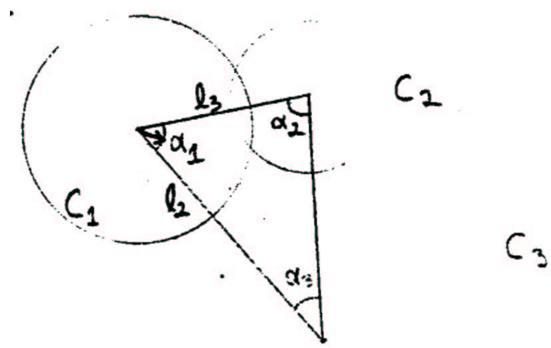

13.56

Thus we have

$$\frac{\partial \alpha_1}{\partial R_1} < 0 \quad , \quad \frac{\partial \alpha_2}{\partial R_1} > 0 \quad , \quad \frac{\partial \alpha_3}{\partial R_1} > 0,$$

where the α_i are the angles of the triangle of centers.

Proof of 13.7.3. Consider first the Euclidean case. Let l_1, l_2 and l_3 denote the lengths of the sides of the triangle of centers. The partial derivatives $\partial l_2/\partial R_1$ and $\partial l_3/\partial R_1$ can be computed geometrically.

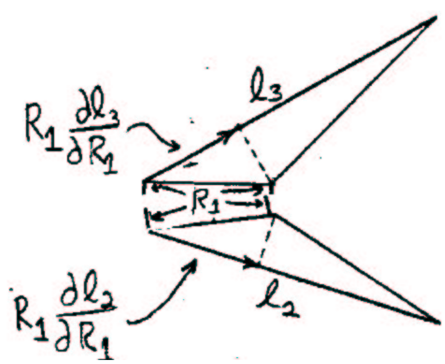

If v_1 denotes the center of C_1, then $\partial v_1/\partial R_1$ is determined as the vector whose orthogonal projections sides 2 and 3 are $\partial l_2/\partial R_1$ and $\partial l_3/\partial R_1$. Thus,

$$R_1 \, \frac{\partial v_1}{\partial R_1}$$

is the vector from v_1 to the intersection of the lines joining the pairs of intersection points of two circles.

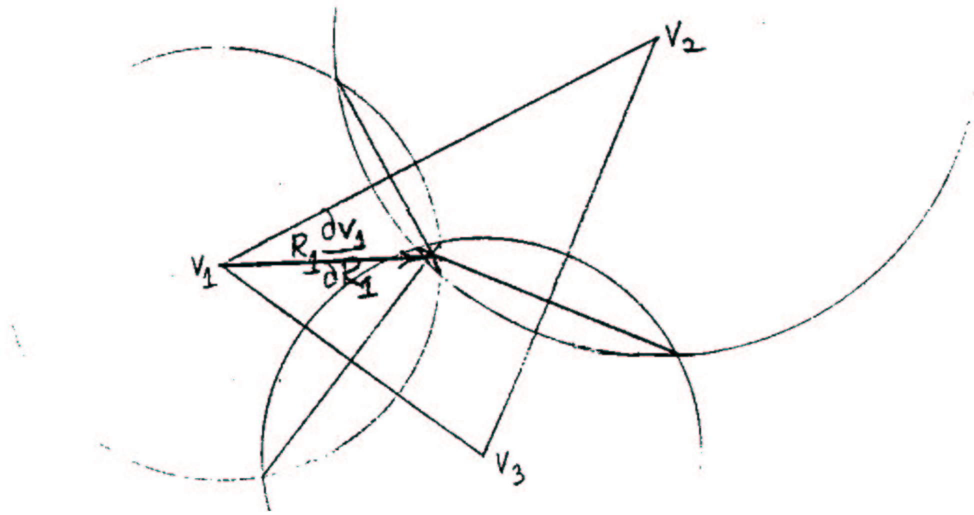

13.57

When all angles of intersection of circles are acute, no circle meets the opposite side of the triangle of centers:

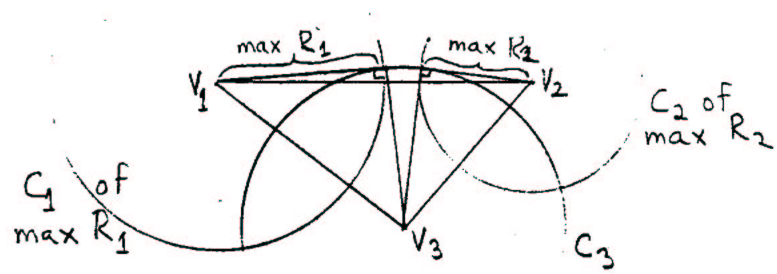

C_3 meets $\overline{v_1v_2} \implies C_1$ and C_2 don't meet.

It follows that $\partial v_1/\partial R_1$ points to the interior of $\Delta v_1v_2v_3$.

The hyperbolic proof is similar, except that some of it takes place in the tangent space to H^2 at v_1. □

Continuation of proof of 13.7.1. From lemma 13.7.3 it follows that when all three radii are increased, the new triangle of centers can be arranged to contain the old one. Thus, the area of S is monotone, for each ray in $\mathbb{R}_+^{\mathcal{V}}$. The area near 0 is near 0, and near ∞ is near $\pi \times$ (# triangles $+ 2\#$ quadrangles); thus the ray intersects $\Delta = F^{-1}(Z)$ in a unique point.

It is now easy to prove that F is an embedding of Δ in Z. In fact, consider any two distinct points R and $R' \in \Delta$. Let $\mathcal{V}^- \subset \mathcal{V}$ be the set of v where $R'(v) < R(v)$. Clearly $\mathcal{V}^-$ is a proper subset. Let $\tau_{\mathcal{V}^-}$ be the subcomplex of τ spanned by $\mathcal{V}^-$. ($\tau_{\mathcal{V}^-}$ consists of all cells whose vertices are contained in $\mathcal{V}^-$). Let $S_{\mathcal{V}^-}$ be a small neighborhood of $\tau_{\mathcal{V}^-}$. We compare the geodesic curvature of $\partial S_{\mathcal{V}^-}$ in the two metrics. To do this, we

may arrange $\partial S_{\mathcal{V}^-}$ to be orthogonal to each edge it meets. Each arc of intersection of $\partial S_{\mathcal{V}^-}$ with a triangle having one vertex in $\mathcal{V}^-$ contributes approximately α_i to the 13.58
total curvature, while each arc of intersection with a triangle having two vertices in $\mathcal{V}^-$ contributes approximately $\beta_i + \gamma_i - \pi$.

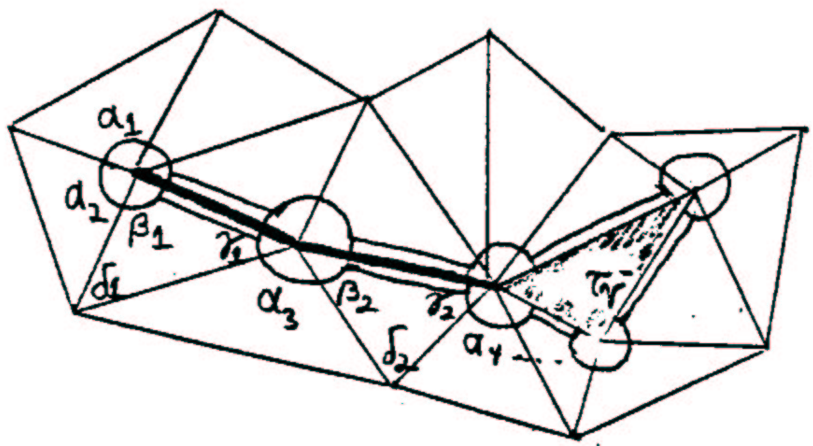

In view of 13.7.3, an angle such as α_1 increases in the R' metric. The change in β_1 and γ_1 is unpredictable. However, their sum must increase: first, let R_1 and R_2 decrease; $\pi - \delta_1 - (\beta_1 + \beta_2)$, which is the area of the triangle in the hyperbolic case, decreases or remains constant but δ_1 also decreases so $\beta_1 + \gamma_1$ must increase. Then let R_3 increase; by 13.7.3, β_1 and γ_1 both increase. Hence, the geodesic curvature of $\partial S_{\mathcal{V}^-}$ increases.

From the Gauss-Bonnet formula,

$$\sum_{v \in \mathcal{V}^-} \kappa(v) = \int_{\partial S_{\mathcal{V}^-}} d_g \, ds - \int_{S_{\mathcal{V}^-}} K \, dA + 2\pi\chi(S_{\mathcal{V}'})$$

it follows that the total curvature at vertices in $\mathcal{V}^-$ must decrease in the R' metric. (Note that the area of $S_{\mathcal{V}^-}$ decreases, so if $k = -1$, the second term on the right decreases.) In particular, $F(R) \neq F(R')$, which shows that F is an embedding of Δ. 13.59

The proof that O is in the image of F is based on the same principle as the proof of uniqueness. We can extract information about the limiting behavior of F as R approaches $\partial\Delta$ by studying the total curvature of the subsurface $S_{\mathcal{V}^O}$, where $\mathcal{V}^O$ consists of the vertices v such that $R(v)$ is tending toward O. When a triangle of τ has two vertices in $\mathcal{V}^O$ and the third not in $\mathcal{V}^O$, then the sum of the two angles at vertices in $\mathcal{V}^O$ tends toward π.

The proof that 0 is in the image of F is based on the same principle as the proof of uniqueness. We can extract information about the limiting behaviour of F as R approaches $\partial\Delta$ by studying the total curvature of the subsurface $S_{\mathcal{V}^0}$, where $\mathcal{V}^0$ consists of the vertices v such that $R(v)$ is tending toward 0. When a triangle of τ has two vertices in $\mathcal{V}^0$ and the third not in $\mathcal{V}^0$, then the sum of the two angles at vertices in $\mathcal{V}^0$ tends toward π.

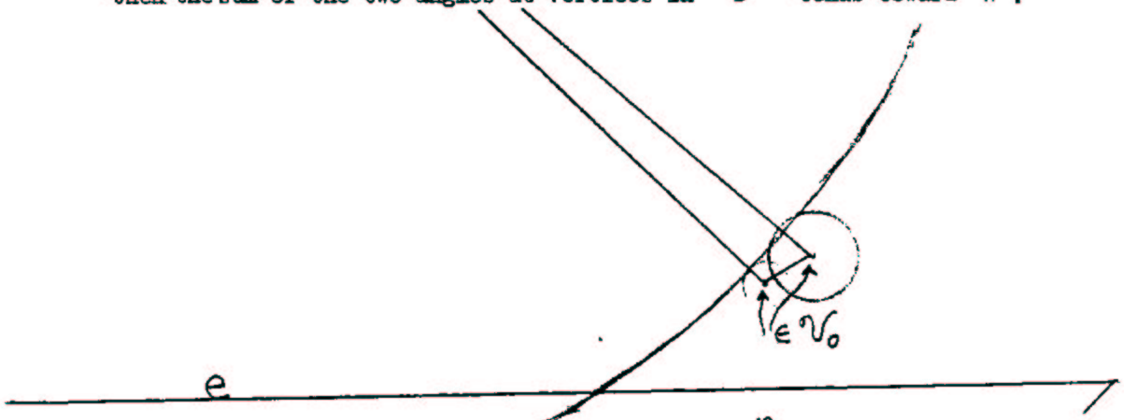

When a triangle of τ has only one vertex in $\mathcal{V}^0$, then the angle at that vertex tends toward the value $\pi - \Theta(e)$, where e is the opposite edge. Thus, the total curvature of $\partial S_{\mathcal{V}^0}$ tends toward the value $\sum_{e \in L(\tau_{\mathcal{V}^0})} (\pi - \Theta(e))$, where $L(\tau_{\mathcal{V}^0})$ is the "link of $\tau_{\mathcal{V}^0}$."

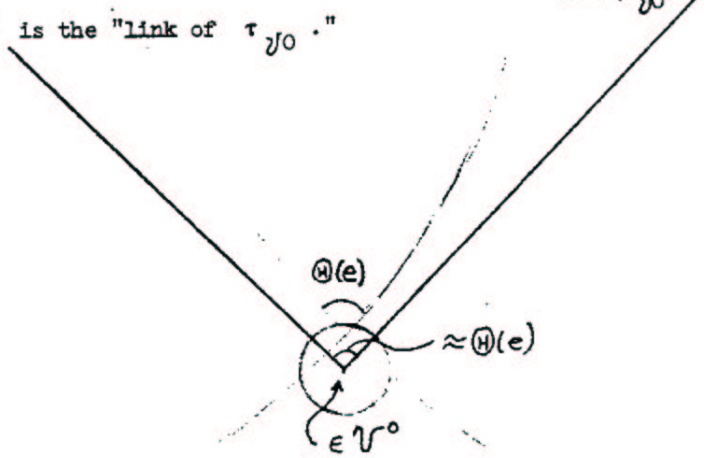

When a triangle of τ has only one vertex in $\mathcal{V}^O$, then the angle at that vertex tends toward the value $\pi - \Theta(e)$, where e is the opposite edge. Thus, the total curvature of $\partial S_{\mathcal{V}^O}$ tends toward the value

$$\sum_{e \in L(\tau_{\mathcal{V}^O})} \big(\pi - \Theta(e)\big),$$

where $L(\tau_{\mathcal{V}^O})$ is the "link of $\tau_{\mathcal{V}^O}$."

The Gauss-Bonnet formula gives 13.60

$$\operatorname{Lim} \sum_{v \in \mathcal{V}^O} \kappa(v) = -\sum_{e \in L(\tau_{\mathcal{V}^O})} \big(\pi - \Theta(e)\big) + 2\pi\chi(S_{\mathcal{V}^O}) < 0.$$

(Note that area $(S_{\mathcal{V}^O}) \to 0$.) To see that the right hand side is always negative, it suffices to consider the case that $\tau_{\mathcal{V}^O}$ is connected. Unless $\tau_{\mathcal{V}^O}$ has Euler characteristic one, both terms are non-positive, and the sum is negative. If $L(\tau_{\mathcal{V}^O})$ has length 5 or more, then

$$\sum_{e \in L(\tau_{\mathcal{V}^O})} \pi - \Theta(e) > e\pi,$$

so the sum is negative. The cases when $L(\tau_{\mathcal{V}^O})$ has length 3 or 4 are dealt with in hypotheses (ii) and (iii) of theorem 13.7.1.

When $\mathcal{V}'$ is any proper subset of $\mathcal{V}^O$ and $R \in \Delta$ is an arbitrary point, we also have an inequality

$$\sum_{v \in \mathcal{V}'} \kappa_R(v) > -\sum_{e \in L(\tau_{\mathcal{V}'})} \big(\pi - \Theta(e)\big) + 2\pi\chi(S_{\mathcal{V}'}).$$

This may be deduced quickly by comparing the R metric with a metric R' in which $R'(\mathcal{V}')$ is near 0. In other words, the image $F(\Delta)$ is contained in the interior of the polyhedron $P \subset Z$ defined by the above inequalities. Since $F(\Delta)$ is an open set whose boundary is ∂P, $F(\Delta) =$ interior (P). Since $O \in \operatorname{int}(P)$, this completes the proof of 13.7.1, and also that of 13.6.4, and 13.6.5. □

Remarks. This proof was based on a practical algorithm for actually constructing patterns of circles. The idea of the algorithm is to adjust, iteratively, the radii of the circles. A change of any single radius affects most strongly the curvature at that vertex, so this proces converges reasonably well. 13.61

The patterns of circles on surfaces of constant curvature, with singularities at the centers of the circles, have a three-dimensional interpretation. Because of the inclusions $\operatorname{isom}(H^2) \subset \operatorname{isom}(H^3)$ and $\operatorname{isom}(E^2) \subset \operatorname{isom}(H^3)$, there is associated with such a surface S a hyperbolic three-manifold M_S, homeomorphic to $S \times \mathbb{R}$, with cone type singularities along (the singularities of S) $\times \mathbb{R}$. Each circle on S determines a totally geodesic submanifold (a "plane") in M_S. These, together with the totally geodesic surface isotopic to S when S is hyperbolic, cut out a submanifold of M_S with finite volume—it is an orbifold as in 13.6.4 or 13.6.5 but with singularities along arcs or half-lines running from the top to the bottom.

Corollary 13.7.4. *Theorems* 13.6.4 *and* 13.6.5 *hold when S is a Euclidean or hyperbolic orbifold, instead of a surface.* (*The orbifold O is to have only singularities as in* 13.6.4 *or* 13.6.5*, plus* (singularities of S) $\times I$ *or* (singularities of S) $\times [0,\infty)$.)

Proof. Solve for pattern of circles on S in a metric of constant curvature on S—the underyling surface of S will have a Riemannian metric with cone type singularities of curvature $2\pi(1/n-1)$ at elliptic points of S, and angles at corner reflectors of S.

An alternative proof is to find a surface $\tilde{S}$ which is a finite covering space of the orbifold S, and find a hyperbolic structure for the corresponding covering space $\tilde{O}$ of O. The existence of a hyperbolic structure for O follows from the uniqueness of the hyperbolic structure on $\tilde{O}$ thence the invariance by deck transformations of $\tilde{O}$ over O. □ 13.62

13.8. A geometric compactification for the Teichmüller spaces of polygonal orbifolds

We will construct hyperbolic structures for a much greater variety of orbifolds by studying the quasi-isometric deformation spaces of orbifolds with boundary whose underlying space is the three-disk. In order to do this, we need a description of the limiting behavior of conformal structure on its boundary. We shall focus on the case when the boundary is a disjoint union of polygonal orbifolds. For this, the greatest clarity is attained by finding the right compactifications for these Teichmüller spaces.

When M is an orbifold, $M_{[\epsilon,\infty)}$ is defined to consist of points x in M such that the ball of radius $\epsilon/2$ about x has a finite fundamental group. Equivalently, no loop through x of length $< \epsilon$ has infinite order in $\pi_1(M)$. $M_{(0,\epsilon]}$ is defined similarly. It does *not*, in general, contain a neighborhood of the singular locus. With this definition, it follows (as in §5) that each component of $M_{(0,\epsilon]}$ is covered by a horoball or a uniform neighborhood of an axis, and its fundamental group contains $\mathbb{Z}$ or $\mathbb{Z}\oplus\mathbb{Z}$ with finite index.

In §5 we defined the geometric topology on sequences of hyperbolic three-manifolds of finite volume. For our present purpose, we want to modify this definition slightly. First, define a *hyperbolic structure with nodes* on a two-dimensional orbifold O to be a complete hyperbolic structure with finite volume on the complement of some one-dimensional suborbifold, whose components are the *nodes.* This includes the case when there are no nodes. A topology is defined on the set of hyperbolic structures with nodes, up to diffeomorphisms isotopic to the identity on a given surface, by saying that M_1 and M_2 have distance $\leq \epsilon$ if there is a diffeomorphism of O [isotopic to the identity] whose restriction to $M_{1[\epsilon',\infty)}$ is a (e^ϵ)-quasi-isometry to $M_{2[\epsilon',\infty)}$. Here, ϵ' is some fixed, small number. 13.63

Remark. The related topology on hyperbolic structures with nodes *up to diffeomorphism* on a given surface is always compact. (Compare Jørgensen's theorem, 5.12, and Mumford's theorem, 8.8.3.) This gives a beautiful compactification for the modular space $\mathcal{T}(M)/\operatorname{Diff}(M)$, which has been studied by Bers, Earle and Marden and Abikoff. What we shall do works because a polygonal orbifold has a finite modular group.

For any two-dimensional orbifold O with $\chi(O) < 0$, let $\mathcal{N}(O)$ be the space of all hyperbolic structures with nodes (up to isotopy) on O.

Theorem 13.8.1. *When P is an n-gonal orbifold, $\mathcal{N}(P)$ is homeomorphic to the* (*closed*) *disk, D^{n-3}, with interior $\mathcal{T}(P)$. It has a natural cell-structure with open cells parametrized by the set of nodes* (*up to isotopy*).

Here are the three simplest examples.

If P is a quadrilateral, then $\mathcal{T}(P)$ is $\mathbb{R}$. There are two possible nodes. $\mathcal{N}(P)$ looks like this:

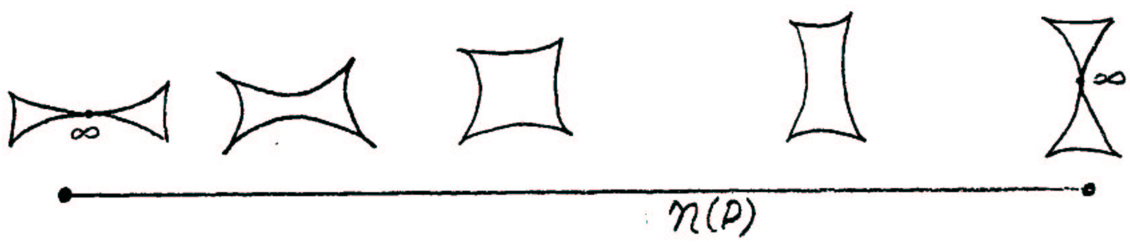

If there are two adjacent order 2 corner reflectors, the qualitative picture must be modified appropriately. For instance,

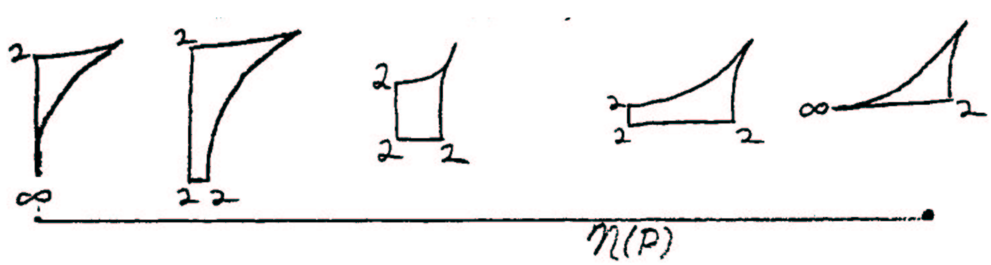

When P is a pentagon, $\mathcal{T}(P)$ is $\mathbb{R}^2$. There are five possible nodes, and the cell-structure is diagrammed below: 13.64

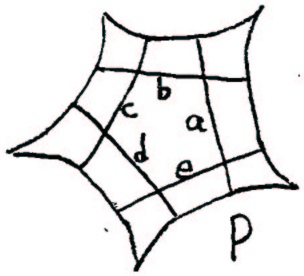

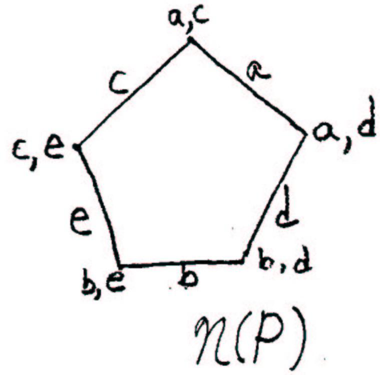

When there is only one node, the pentagon is pinched into a quadrilateral and a triangle, so there is still one degree of freedom.

When P is a hexagon, there are 9 possible nodes.

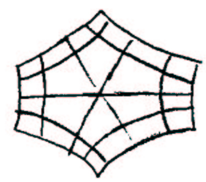

Each single node pinches the hexagon into a pentagon and a triangle, or into two quadrilaterals, so its associated 2-cell is a pentagon or a square. The cell division of ∂D^3 is diagrammed below:

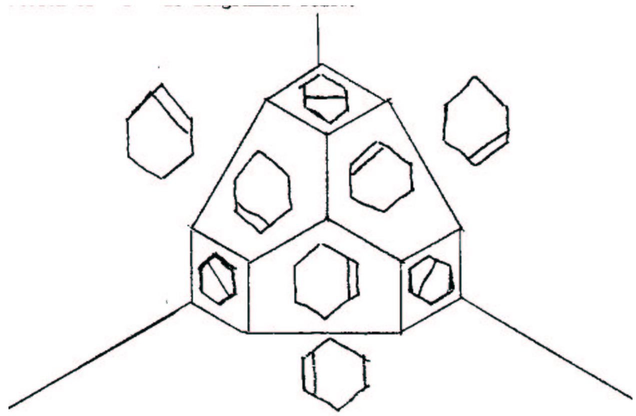

(The zero and one-dimensional cells are parametrized by the union of the nodes of 13.65
the incident 2-cells.)

Proof of 13.8.1. It is easy to see that $\mathcal{N}(P)$ is compact by familiar arguments, as in 5.12 and 8.8.3, for instance. In fact, choose ϵ sufficiently small so that $P_{(0,\epsilon]}$ is always a disjoint union of regular neighborhoods of short arcs. Given a sequence $\{P_i\}$, we can pass to a subsequence so that the core one-orbifolds of the components of $P_{i(0,\epsilon]}$ are constant. Extend this system of arcs to a maximal system of disjoint geodesic arcs $\{\alpha_1, \dots, \alpha_k\}$. The lengths of all such arcs remain bounded in $\{P_i\}$ (this follows from area considerations), so there is a subsequence so that all lengths converge—possibly to zero. But any set of $\{l(\alpha_i) | l(\alpha_i) \geq 0\}$ defines a hyperbolic structure with nodes, so our sequence converges in $\mathcal{N}(P)$.

Furthermore, we have described a covering of $\mathcal{N}(P)$ by neighborhoods diffeomorphic to quadrants, so it has the structure of a manifold with corners. Change of coordinates is obviously differentiable. Each stratum consists of hyperbolic structures with a prescribed set of nodes, so it is diffeomorphic to Euclidean space (this also follows directly from the nature of our local coordinate systems.)

Theorem 13.8.1 follows from this information. Here is a little overproof. An explicit homeomorphism to a disk can be constructed by observing that $\mathcal{PL}(P)$[‡] has a natural triangulation, which is dual to the cell structure of $\partial\mathcal{N}(P)$. This arises from the fact that any simple geodesic on P must be orthogonal to the mirrors, so a geodesic lamination on P is finite. The simplices in $\mathcal{PL}(P)$ are measures on a maximal family of geodesic one-orbifolds. 13.66

A projective structure for $\mathcal{PL}(P)$—that is, a piecewise projective[§] homeomorphism to a sphere—can be obtained as follows (compare Corollary 9.7.4). The set of geodesic laminations on P is in one-to-one correspondence with the set of cell divisions of P which have no added vertices. Geometrically, in fact, a geometric lamination extends in the projective (Klein) model to give a subdivision of the dual polygon.

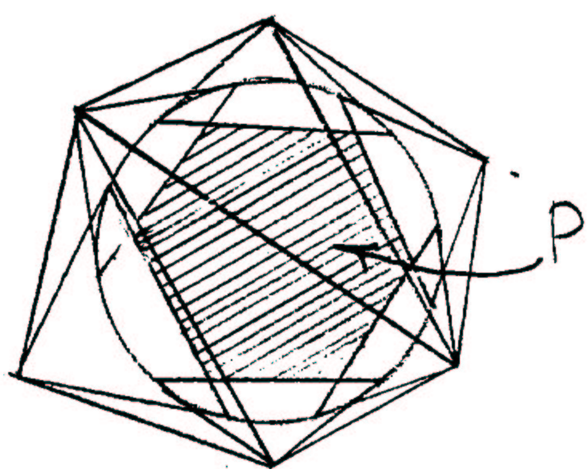

Take the model P now to be a regular polygon in $\mathbb{R}^2 \subset \mathbb{R}^3$. Let V be the vertex set. For any function $f : V \to \mathbb{R}$, let C_f be the convex hull of the set of points

[‡]For definition, and other information, see p. 8.58

[§]See remark 9.5.9.

obtained by moving each vertex v of P to a height $f(v)$ (positive or negative along the perpendicular to $\mathbb{R}^2$ through v). The "top" of C_f gives a subdivision of P. The nature of this subdivision is unchanged if a function which extends to an affine function from $\mathbb{R}^2$ to $\mathbb{R}$ is added to f. Thus, we have a map $\mathbb{R}^V/\mathbb{R}^3 \to \mathcal{GL}(P)$. To lift the map to measured laminations, take the directional derivative at O of the bending measure for the top of the convex hull, in the direction f. The global description of this map is that a function f is associated to the measure which assigns to each edge e of the bending locus the change in slope of the intersection of the faces adjacent to e with a plane perpendicular to e.

It is geometrically clear that we thus obtain a piecewise linear homeomorphism, 13.67

$$e : \mathcal{ML}(P) \approx \mathbb{R}^{V-3} - 0.$$

The set of measures which assigns a maximal value of 1 to an edge gives a realization of $\mathcal{PL}(P)$ as a convex polyhedral sphere Q in $\mathbb{R}^{V-3}$. The dual polyhedron Q^*—which is, by definition, the set of vectors $X \in \mathbb{R}^{V-3}$ such that $\sup_{y \in Q} X \cdot Y = 1$—is the boundary of a convex disk, combinatorially equal to $\mathcal{N}(P)$. This seems explicit enough for now. $\square$

13.9. A geometric compactification for the deformation spaces of certain Kleinian groups.

Let O be an orbifold with underlying space $X_O = D^3$, $\Sigma_O \subset \partial D^3$, and $\partial\Sigma_O$ a union of polygons.

We will use the terminology *Kleinian structure* on O to mean a diffeomorphism of O to a Kleinian manifold $B^3 - L_\Gamma/\Gamma$, where Γ is a Kleinian group.

In order to describe the ways in which Kleinian structures on O can degenerate, we will also define the notion of a *Kleinian structure with nodes* on O. The nodes are meant to represent the limiting behavior as some one-dimensional suborbifold S becomes shorter and shorter, finally becoming parabolic. We shall see that this happens only when S is isotopic in one or more ways to ∂O; the geometry depends on the set of suborbifolds on ∂O isotopic to S which are being pinched in the conformal geometry of ∂O. To take care of the various possibilities, nodes are to be of one of these three types:

(a) An incompressible one-suborbifold of ∂O.
(b) An incompressible two-dimensional suborbifold of O, with Euler characteristic zero and non-empty boundary. In general, it would be one of these five: 13.68

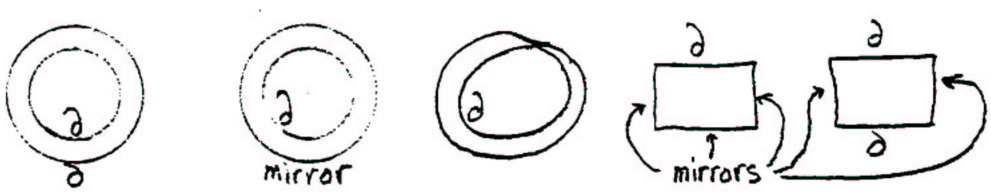

but for the orbifolds we are considering only the last two can occur.

(c) An orbifold T modelled on $P_{2k} \times \mathbb{R}$, $k > 2$ where P_{2k} is a polygon with $2k$ sides. The sides of P_{2k} are to alternate being on ∂O and in the interior of O. (Cases a and b could be subsumed under this case by thickening them and regarding them as the cases $k = 1$ and $k = 2$.)

A Kleinian structure with nodes is now defined to be a Kleinian structure in the complement of a union of nodes of the above types, neighborhoods of the nodes in being horoball neighborhoods of cusps in the Kleinian structures. Of course, if O minus the nodes is not connected, each component is the quotient of a separate Kleinian group (so our definition was not general enough for this case).

Let $\mathcal{N}(O)$ denote the set of all Kleinian structure with nodes on O, up to homeomorphisms isotopic to the identity. As for surfaces, we define a topology on $\mathcal{N}(O)$, by saying that two structures K_1 and K_2 have distance $\leq \epsilon$ if there is a homeomorphism between them which is an $e^\epsilon -$ quasi-isometry on $K_{1[\epsilon,\infty)}$ intersected with the convex hull of K_1.

Theorem 13.9.1. *Let O be as above with O irreducible and ∂O incompressible. If O has one non-elementary Kleinian structure, then $\mathcal{N}(O)$ is compact. The conformal structure on ∂O is continuous, and it gives a homeomorphism to a disk,*

$$\mathcal{N}(O) \approx \mathcal{N}(\partial O).$$

Note: The necessary and sufficiently conditions for existence of a Kleinian structure will be given in [???] or they can be deduced from Andreev's theorem 13.6.1. 13.69
We will use 13.6.1 to prove existence.

Proof. We will study the convex hulls of the Kleinian structures with nodes on O. (When the Kleinian structure is disconnected, this is the union of convex hulls of the pieces.)

Lemma 13.9.2. *There is a uniform upper bound for the volume of the convex hull, H, of a Kleinian structure with nodes on O.*

Proof of 13.9.2. The bending lamination for ∂O has a bounded number of components. Therefore, H is (geometrically) a polyhedron with a bounded number of faces, each with a bounded number of sides. Hence the area of the boundary of the polyhedron is bounded. Its volume is also bounded, in view of the isoperimetric inequality,

$$\text{volume }(S) \leq 1/2 \text{ area}(\partial S)$$

for a set $S \subset H^3$. (cf. §5.11). □

Theorem 13.9.1 can now be derived by an adaptation of the proof of Jørgensen's theorem (5.12) to the present situation. It can also be proved by a direct analysis of the shape of H. We will carry through this latter course to make this proof more concrete and self-contained.

The first observation is that H can degenerate only when some edges of H become very long. When a face of H has vertices at infinity, "length" is measured here as the distance between canonical neighborhoods of the vertices. In fact, if the edges of H remain bounded in length, the faces remain bounded in shape by (§13.8, for 13.70
instance; the components of ∂H can be treated as single faces for this analysis). If we view X_H as a convex polyhedron in H^3 then as long as a sequence $\{H_i\}$ has all faces remaining bounded in shape, there is a subsequence such that the polyhedra $\{X_{H_i}\}$ converge, in the sense that the maps of each face into H^3 converge. One possibility is that the limiting map of X_H has a two-dimensional image: this happens in the case of a sequence of quasi-Fuchsian groups converging to a Fuchsian group, and we do not regard the limit as degenerate. The significant point is that two silvered faces of H (faces of H not on ∂H) which are not incident (along an edge or at a cusp) cannot come close together unless their diameter goes to infinity, because any points of close approach are deep inside $H_{(0,\epsilon]}$.

We can obtain a good picture of the degeneration which occurs as an edge becomes very long by the following analysis. We will consider only edges which are not in the interior of ∂H. Since the area of each face of H is bounded, any edge e of H which is very long must be close and nearly parallel, for most of its length all but a bounded part, of its length, on both sides, to other edges of its adjacent faces.

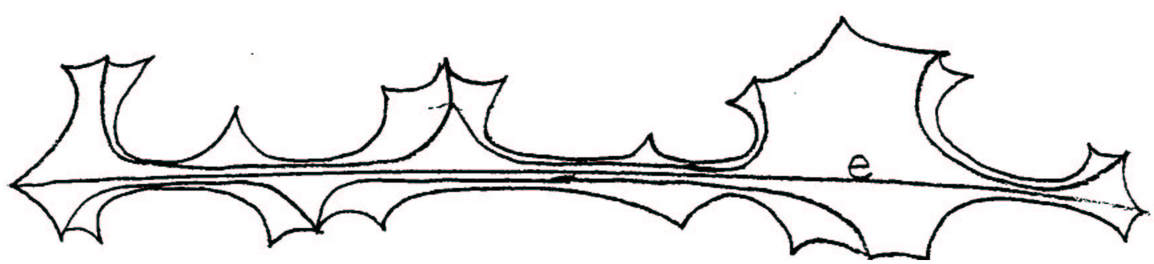

Similarly, these nearly parallel edges must be close and nearly parallel to still more edges on the far side from e. How long does this continue? Remember that H has an angle at each edge. In fact, if we ignore edges in the interior of ∂H, no angle exceeds 90°. Special note should be made here of the angles between ∂H and mirrors 13.71

of H: the condition for convexity of H is that ∂H, together with its reflected image, is convex, so these angles also are $\leq 90°$. (If they are strictly less, then that edge of ∂H is part of the bending locus, and consequently it must have ends on order 2 corner reflectors.) Since H is geometrically a convex polyhedron, the only way that it can be bent so much along such closely spaced lines is that it be very thin. In other words, along most of the length of e, the planes perpendicular to $e \subset X_H \subset H^3$ intersect XH in a small polygon, which represents a suborbifold. It has $2, 3$ or 4 intersections with edges of XH not interior to ∂H.

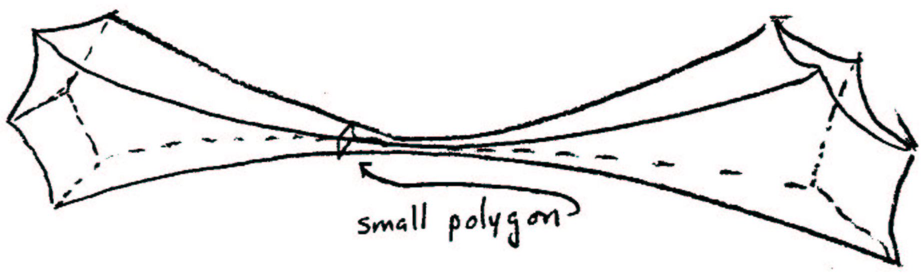

By area-angle considerations, this small suborbifold must have non-negative Euler characteristic. We investigate the cases separately.

(a) $\chi = 0, \quad \partial = \emptyset$

(i) 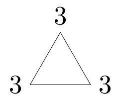

This is automatically incompressible, and since it is closed, it must be homotopic to a cusp. But this is supposed to be avoided by keeping our investigations away from the vertices of faces of P.

(ii) 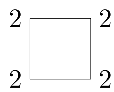

Either it is incompressible, and avoided as in (i), or compressible, so it is homotopic to some edge of H.

But since it is small, it must be very close to that edge. This contradicts the way it was chosen—or, in any case, it can account for only a small part 13.72
of the length of e.

(b) $\chi = 0, \quad \partial \neq \emptyset$:

(i) 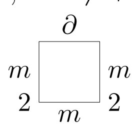

(ii) 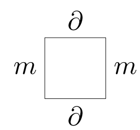

where m denotes a mirror.

These can occur either as small ∂-incompressible suborbifolds (representing incipient two-dimensional nodes) or as small ∂-compressible orbifolds, representing the boundary of a neighborhood of an incipient one-dimensional node.

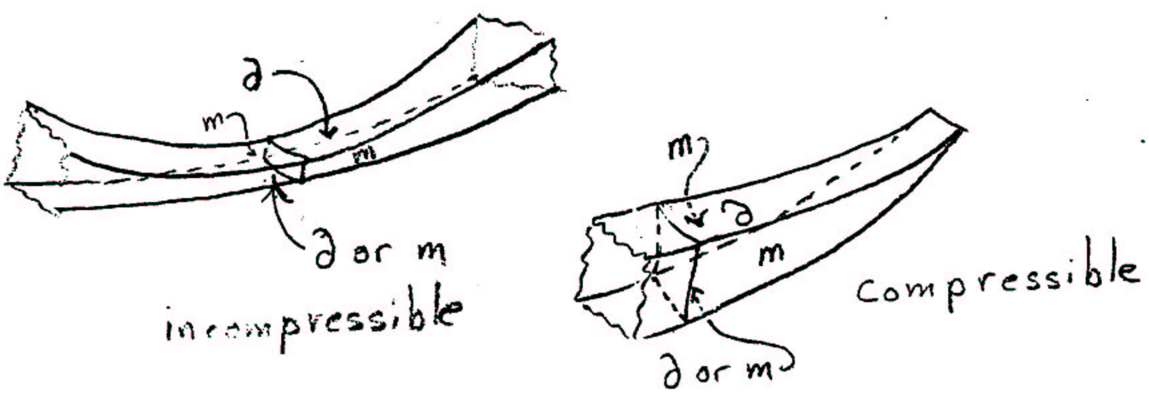

(c) $\chi > 0$. This can occur, since O is irreducible and ∂O incompressible.

We now can see that H is decomposed into reasonably wide convex pieces, joined together along long thin spikes whose cross-sections are two-dimensional orbifolds with boundary. There also may be some long thin spikes representing neighborhoods of short one-suborbifolds (arcs) of ∂O.

$H_{(0,\epsilon]}$ contains all the long spikes. It may also intersect certain regions between spikes, where two silvered faces of H come close together. If so, then $H_{(0,\epsilon]}$ contains the entire region, bounded by spikes (since each edge of the two nearby faces comes to a spike within a bounded distance, as we have seen). 13.73

The fundamental group of that part of H must be elementary: in other words, all faces represent reflections in planes perpendicular to or containing a single axis.

It should by now be clear that $\mathcal{N}(O)$ is compact. By [???], Kleinian structures with nodes of a certain type on O are parametrized, if they exist, by conformal structures with nodes of the appropriate type on ∂O. Given a Kleinian structure with nodes, K, and a nearby element K' in $\mathcal{N}(O)$, theer is a map with very small dilation from all but a small neighborhood of the nodes in ∂K to $\partial K'$, covering all but a long thin neck; this implies that $\partial K'$ is near ∂K in $\mathcal{N}(\partial O)$. Therefore, the map from $\mathcal{N}(O)$ to $\mathcal{N}(\partial O)$ is continuous. Since $\mathcal{N}(O)$ is compact, the image is all of $\mathcal{N}(\partial O)$. Since the map is one-to-one, it is a homeomorphism. $\square$

To be continued. . . .