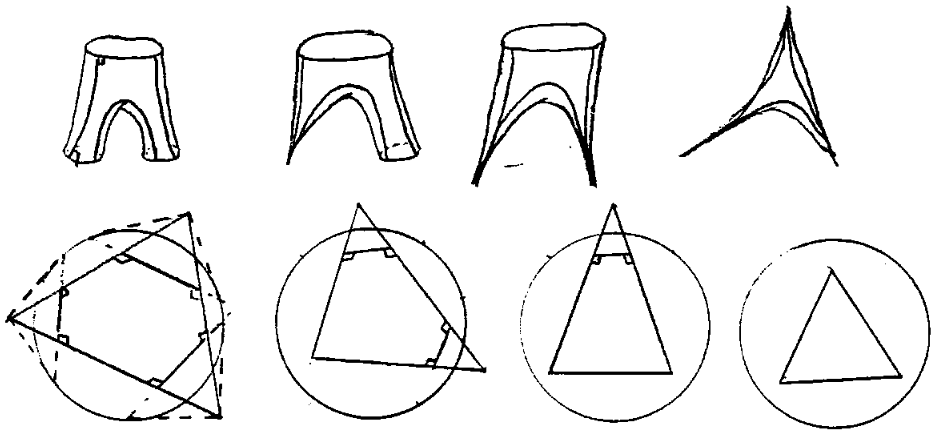

The orbifold $D^2_{(\,;m_1,\dots,m_l)}$ also can be decomposed into "generalized triangles,"

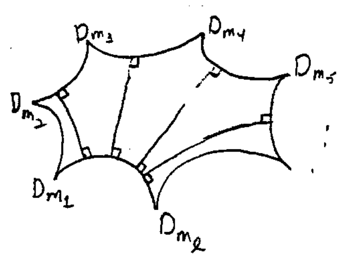

for instance in the pattern above. One immediately sees that the orbifold has hyperbolic structures (provided $\chi < 0$) parametrized by the lengths of the cuts; that is, $(\mathbb{R}_+)^{l-3}$. Special care must be taken when, say, $m_1 = m_2 = 2$. Then one of 13.25
the cuts must be omitted, and an 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:

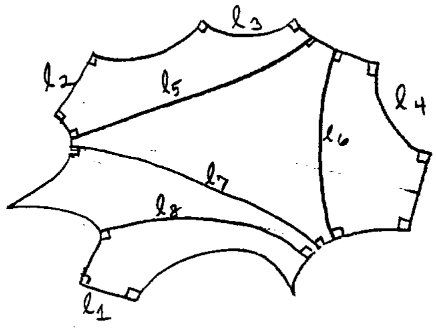

The annulus with more than one corner reflector on one boundary component should be dissected, as below, into $D_{(\,;n_1,\dots,n_k)}$ and an annulus with two order two corner reflectors. $D^2_{(n;m_1,\dots,m_l)}$ is analogous.

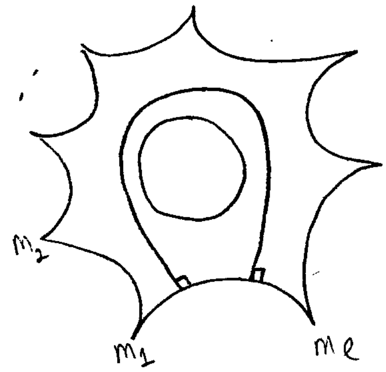

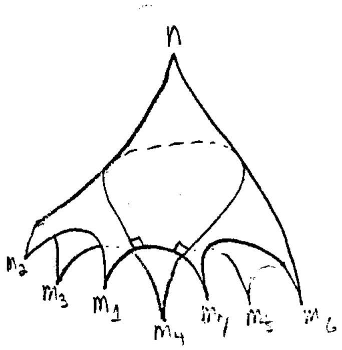

13.26

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:

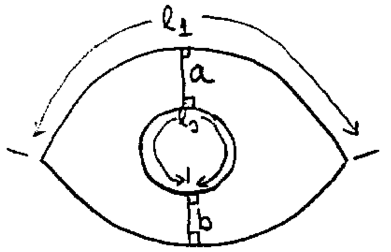

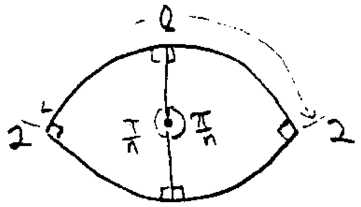

(The two all right pentagons agree on a and b, so they are congruent; thus they are determined by their edges of length $l_1/2$ and $l_2/2$). Similarly, $D^2_{(n;\,2,2)}$ is determined by one edge length, provided $n > 2$. $D^2_{(2;\,2,2)}$ is not hyperbolic. However, it has a degenerate hyperbolic structure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

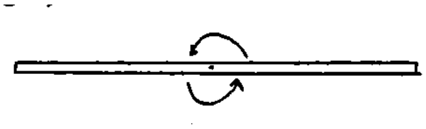

This is consistent with the way in which it arises in considering hyperbolic structures, in the dissection of $D^2_{(2;m_1,\dots,m_l)}$. One can cut such an orbifold along the perpendicular arc from the elliptic point to an edge, to obtain $D^2_{(\,;\,2,2,m_1,\dots,m_l)}$. In the case of an annulus with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

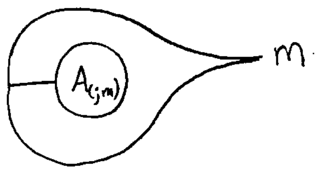

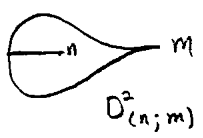

note first that it is symmetric, since it can be dissected into an isosceles "triangle." Now, from a second dissection, we see hyperbolic structures are paremetrized by the length of the boundary component without the reflector.

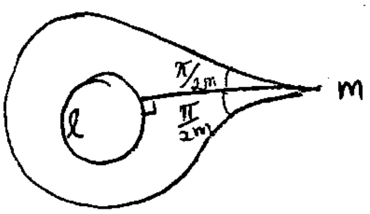

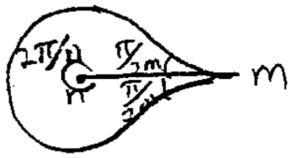

By the same argument, $D^2_{(n;m)}$ has a unique hyperbolic structure.

All these pieces can easily be reassembled to give a hyperbolic structure on O. $\square$

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. $\square$

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

$$\begin{tikzcd} \tilde{U} \times F \arrow[r] \arrow[d] & p^{-1}(U) \arrow[d] \\ \tilde{U} \arrow[r] & U \end{tikzcd}$$

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.

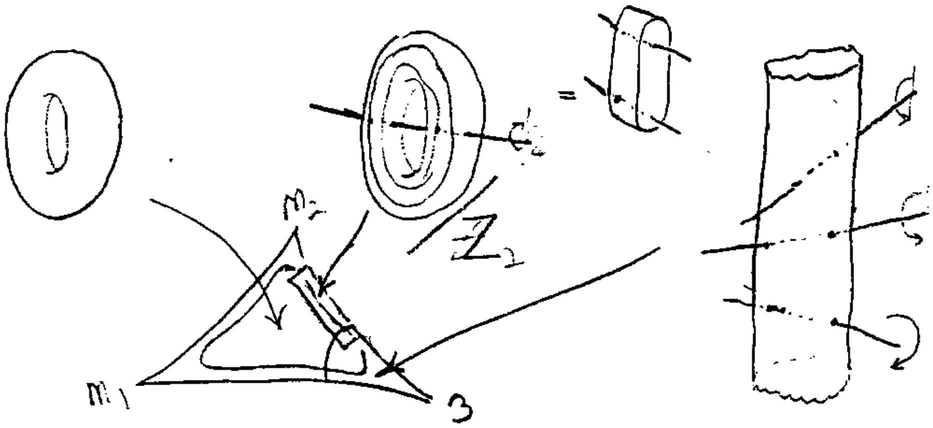

Over the non-singular part of O, we have a solid torus. Over an edge, we have $mI \times I$, with fibers folded into mI; nearby figures go once around these mI's. Above a corner reflector of order n, the fiber is folded into mI. The fibers above the nearby edges weave up and down n times, and nearby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

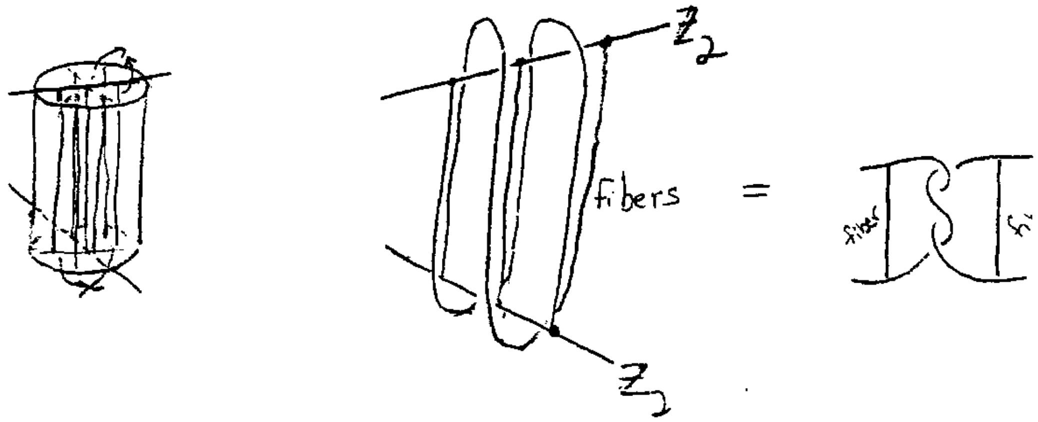

13.32

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

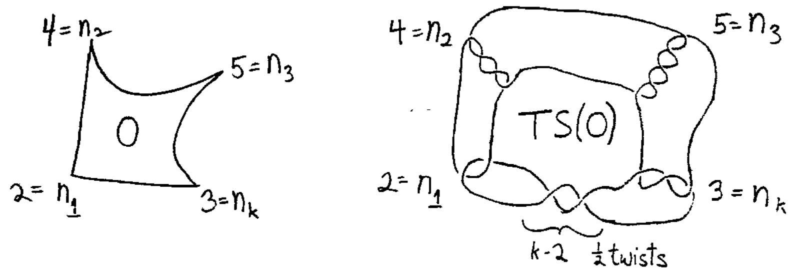

When O is a Riemannian orbifold, this gives $T_1(O)$ a canonical flow, the geodesic flow. For the Euclidean orbifolds with X_O a polygon, this flow is physically realized (up to friction and spin) by the motion of a billiard ball. The flow is tangent to the singular locus. Thus, the phase space for the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

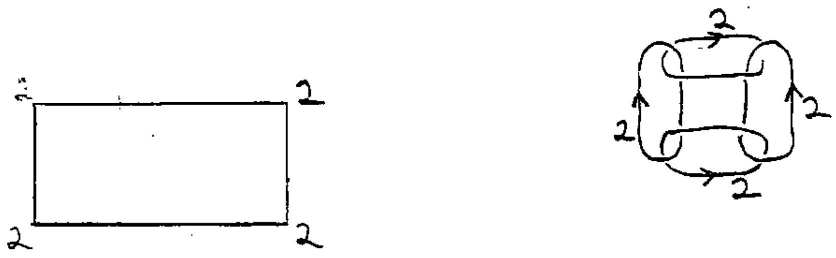

There are two invariant annuli, with boundary the singular locus, corresponding to trajectories orthogonal to a side. The other trajectories group into invariant tori.

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

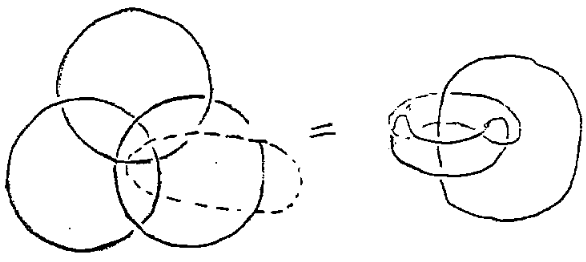

You can obtain many other examples via symmetries and covering spaces. For instance, the Borromean rings above have a three-fold axis of symmetry, with quotient orbifold:

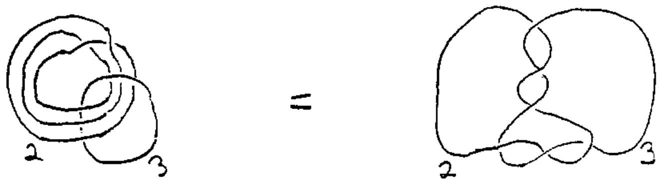

We can pass to a two-fold cover, unwrapping around the $\mathbb{Z}_3$ elliptic axis, to obtain the figure-eight knot as a $\mathbb{Z}_3$ elliptic axis.

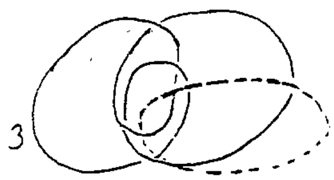

This is a Euclidean orbifold, whose fundamental group is generated by order 3 rotations in main diagonals of two adjacent cubes (regarded as fundamental domains for example 13.1.5).

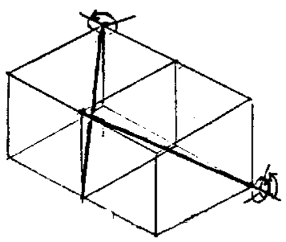

When O is elliptic, then all geodesics are closed, and the geodesic flow comes from a circle action. It follows that $T_1(O)$ is a fibration in a second way, by projecting to the quotient space by the geodesic flow! For instance, the singular locus of $T_1(D^2_{(2,3,5)})$ is a torus knot of type $(3,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

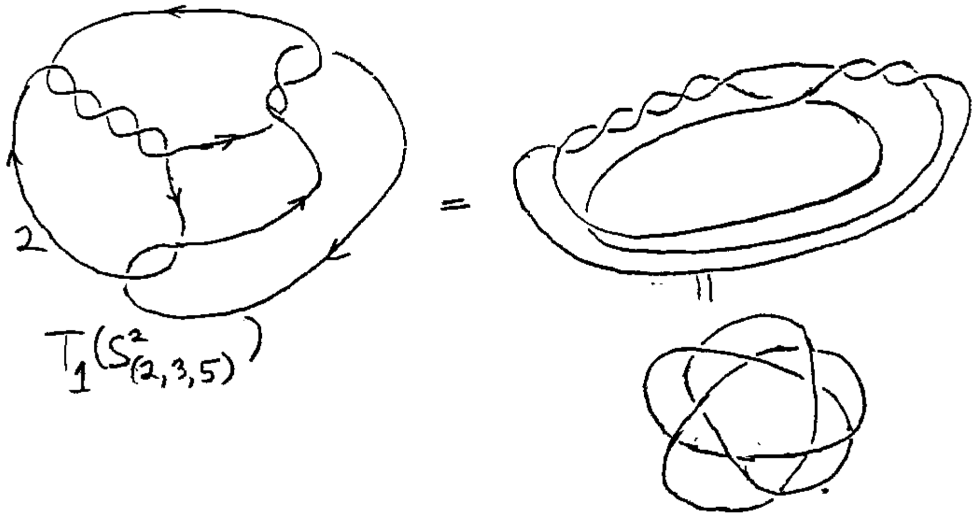

Therefore, it also fibers over $S^2_{(2,3,5)}$. In general, an oriented three-orbifold which fibers over a two-orbifold, with general fiber a circle, is determined by three kinds of information:

(a) The base orbifold.

(b) For each elliptic point or corner reflector of order n, an integer $0 \leq k < n$ 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

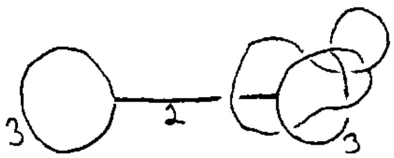

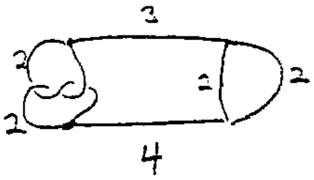

Frequently in three dimensions it is easy to see that certain orbifolds are good but hard to prove much more about them. For instance, the orbifolds with singular locus a knot or link in S^3 are always good: they always have finite abelian covers by manifolds.

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

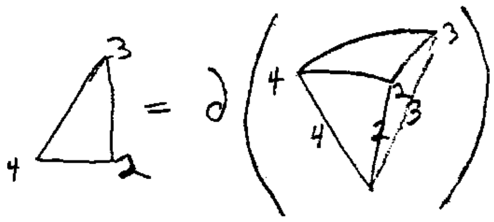

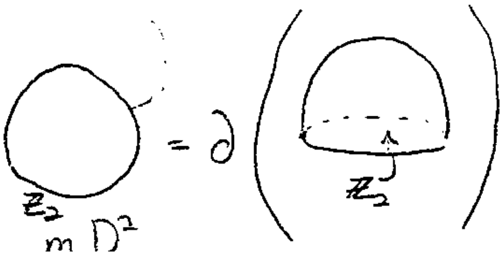

An *incompressible suborbifold* of a three-orbifold O, when X_O is oriented, is a two-suborbifold $O' \subset O$ with $\chi(O') \le 0$ such that every one-suborbifold $O'' \subset O'$ which bounds an elliptic suborbifold of $O - O'$ bounds an elliptic suborbifold of O'. O is *Haken* if it is irreducible 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 13.38
F, and meeting only edges of Σ^1_O incident to $\bar{F}$.

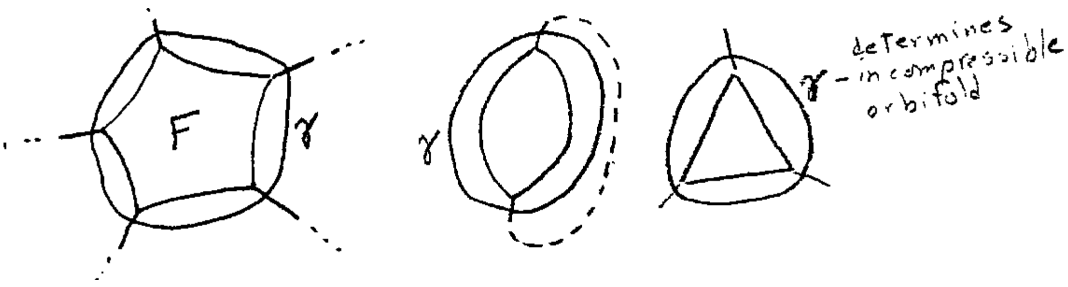

If γ meets no edges, then $\Sigma^1_O = \partial F$ (since O is irreducible) and O is the suspension of $D^2_{(;n,n)}$. The next case is that γ meets two edges of order n; then they must really be the same edge, and O is the suspension of an elliptic orbifold $D^2_{(;n,n_1,n_2)}$. If γ meets three edges, 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.

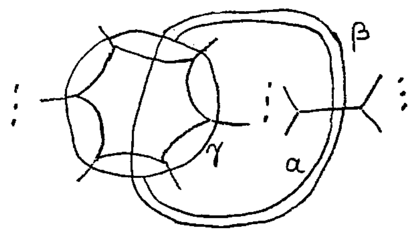

If α determines a "compression," then β determines a triangle orbifold.

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. $\square$

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:

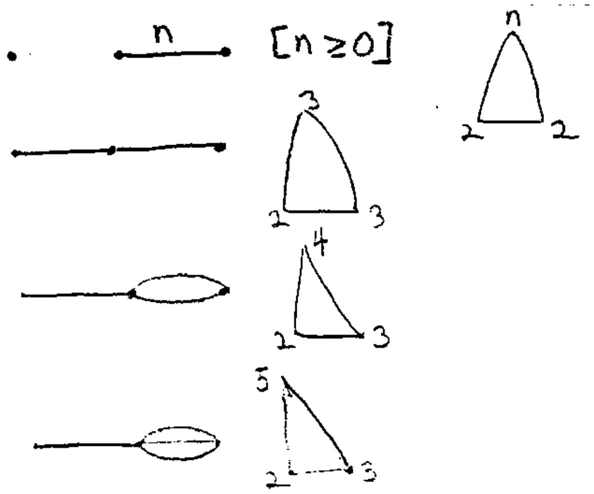

13.40

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:*

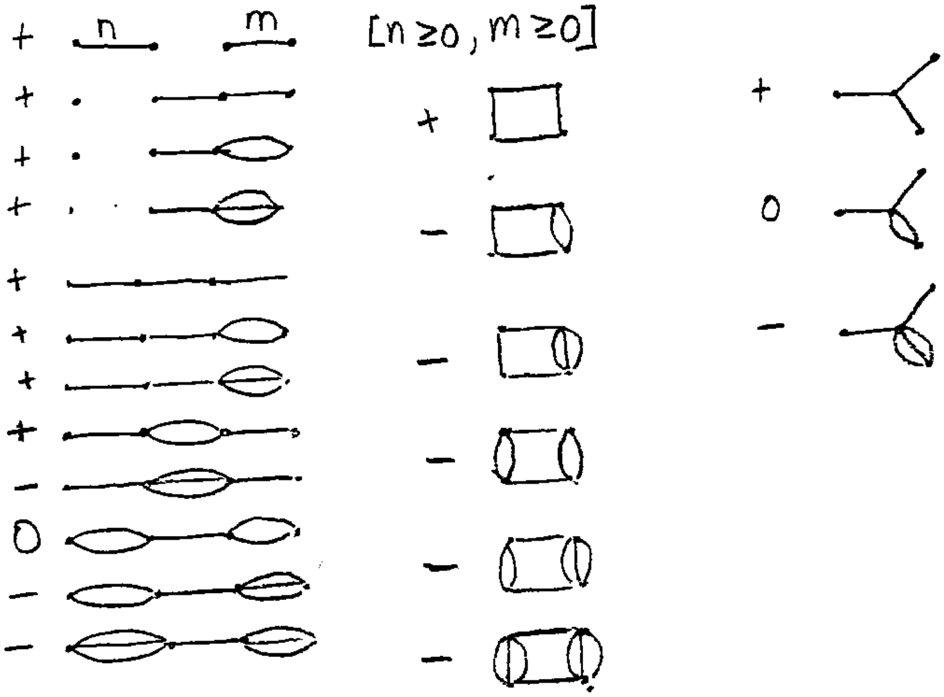

This statement may be slightly generalized to include non-compact orbifolds of the combinatorial type of a simplex with some vertices deleted.

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 struc-* 13.41
ture of finite volume. The three-dimensional examples are listed below:

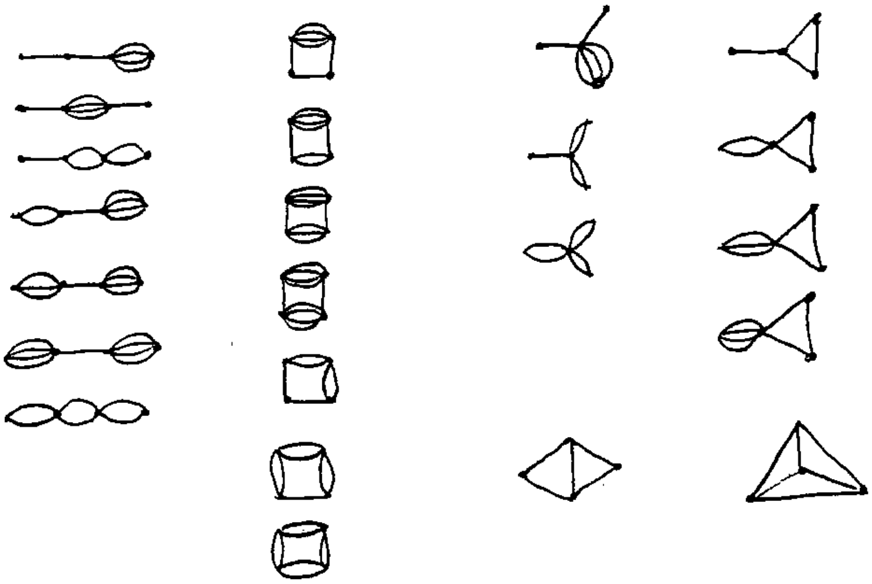

PROOF OF 13.5.3 AND 13.5.4. The method is to describe a simplex in terms of the quadratic form models. Thus, an n-simplex σ^n on S^n has $n+1$ hyperfaces. Each face is contained in the intersection of a codimension one subspace of E^{n+1} with S^n. Let $V_O, \dots, V_n$ be unit vectors orthogonal to 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 $X_1^2 + \dots + X_n^2 - X_{n+1}^2 = -1$ (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 $V_0, \dots, V_n$ 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 $V_O, \dots, V_n$ 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 $V_O, \dots, V_n$ 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

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

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

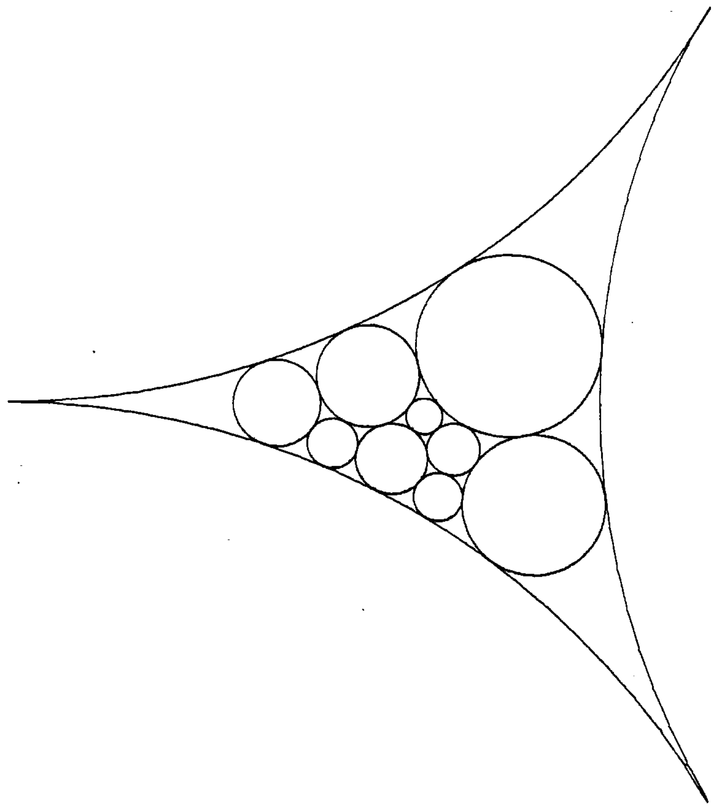

13.45

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.

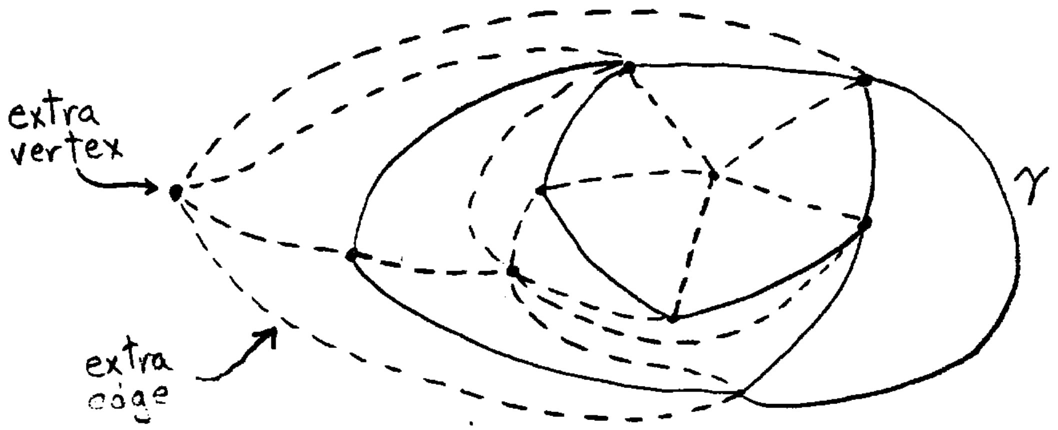

Let P be the polyhedron obtained by cutting off neighborhoods of the vertices of T, down to the middle of each edge of T.

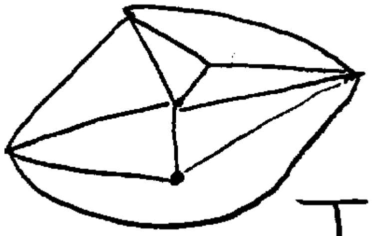

T

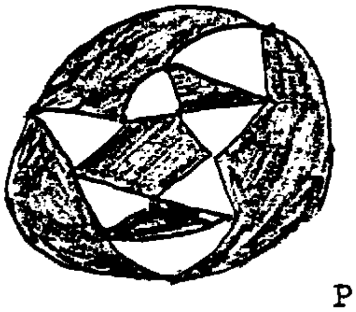

P

Let O be the orbifold with underlying space

$$X_O = D^3\text{-vertices of } P, \quad \text{and} \quad \Sigma_O^1 = \text{ edges of } P,$$

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. □

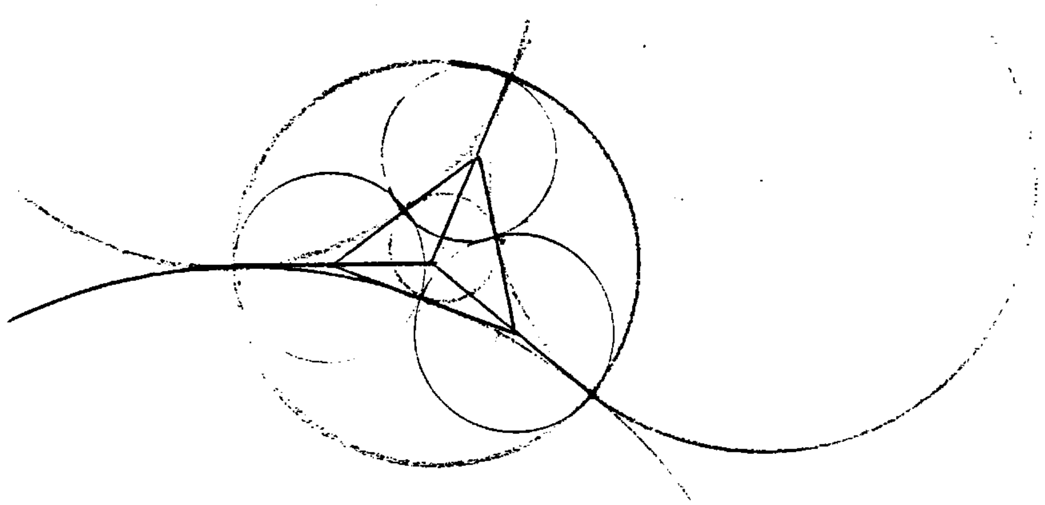

REMARK. Andreev also gave a proof of uniqueness of a hyperbolic polyhedron with assigned concave angles, so the reference to Mostow's theorem is not essential.

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.

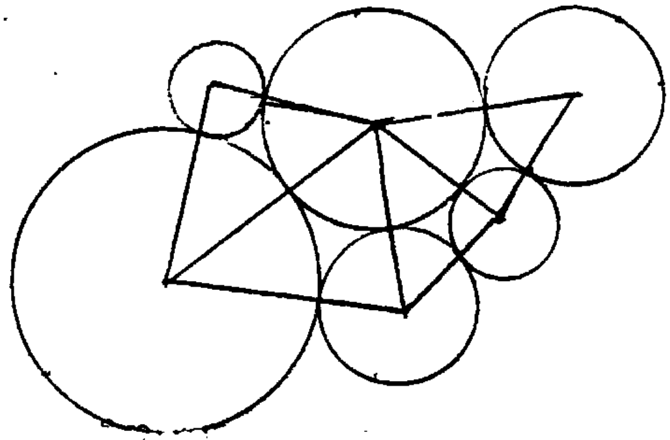

The existence of other geometric patterns of circles in $\mathbb{R}^2$ may also be deduced from Andreev's theorem. For instance, it gives necessary and sufficient condition for the existence of a family of circles meeting only orthogonally in a certain pattern, or 13.48
meeting at 60° angles.

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.*

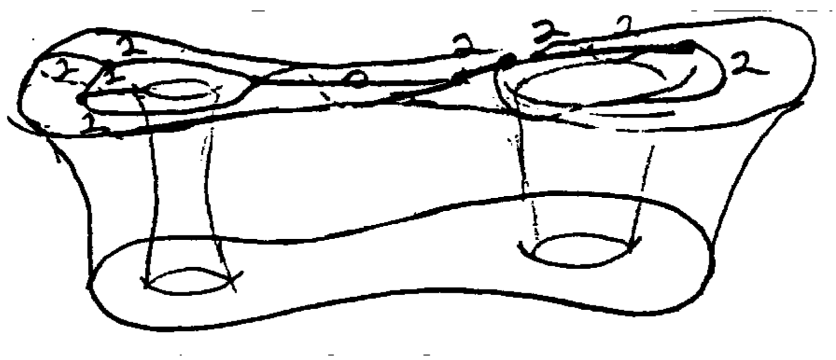

By considering $\pi_1 O$, O as in 13.6.4, as a Kleinian group in upper half space with $T^2 \times \infty$ at ∞, 13.6.4 may be translated into a statement about the existence of doubly periodic families of circles in the plane, or 13.48.a

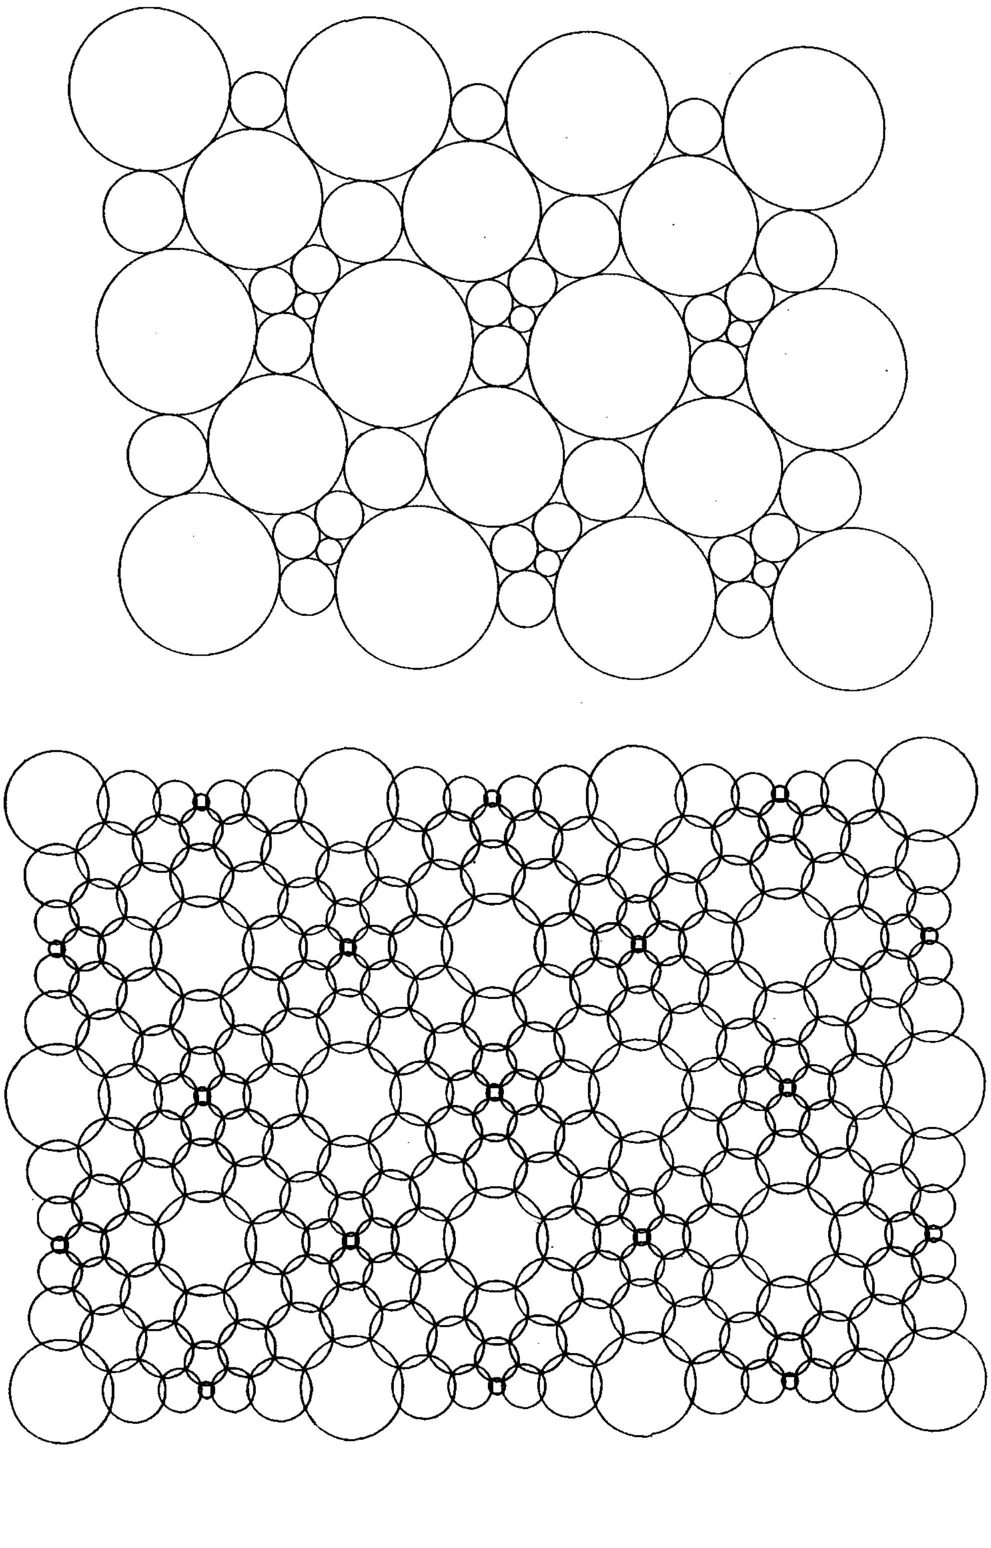

13.48.b