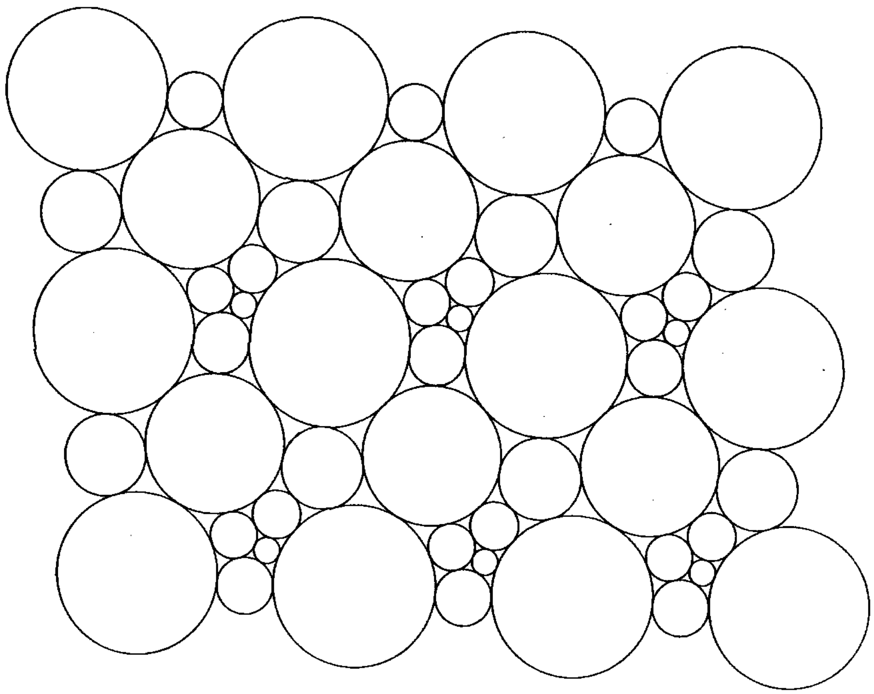

13.48.b

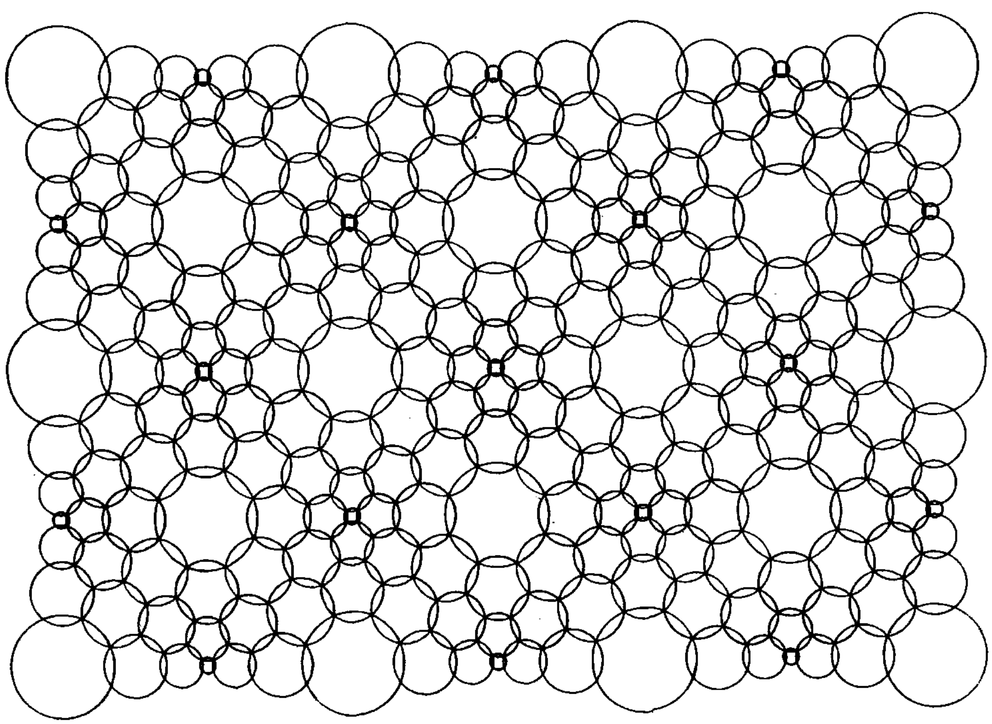

13.49

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.)

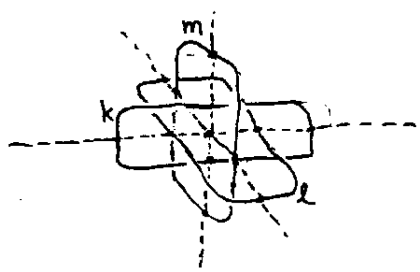

Thus, an orbifold O having the rings as elliptic axes of orders k, l and m is an eight-fold covering space of another orbifold, which has the combinatorial type of a cube.

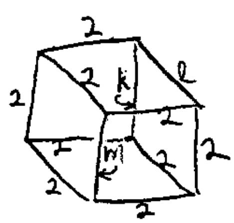

13.50

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

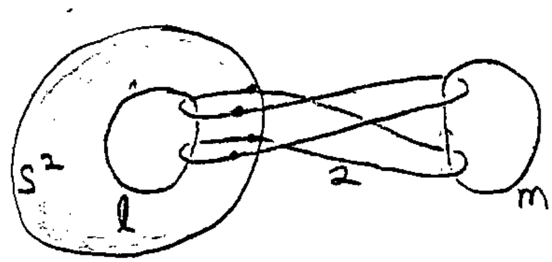

two halves, each fibering over two-orbifolds with boundary, but in incompatible ways (unless l or m is 2).

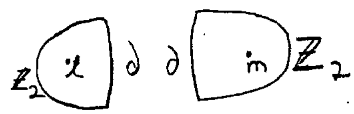

Base spaces of the fibrations

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 $\Theta(e) = 0$.

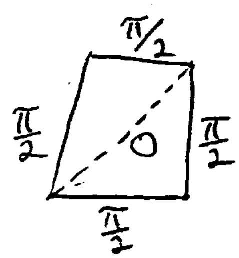

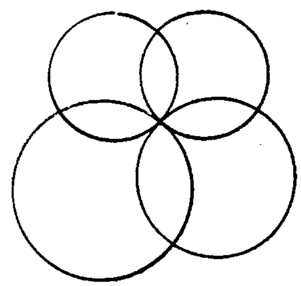

There is an extraneous tangency of circles here—in fact, all extraneous tangencies come from this situation. Henceforth, we assume τ has no quadrangles. The idea is to solve for the *radii* of the circles C_{v_1}. Given an arbitrary set of radii, we shall construct a Riemannian metric on S with cone 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

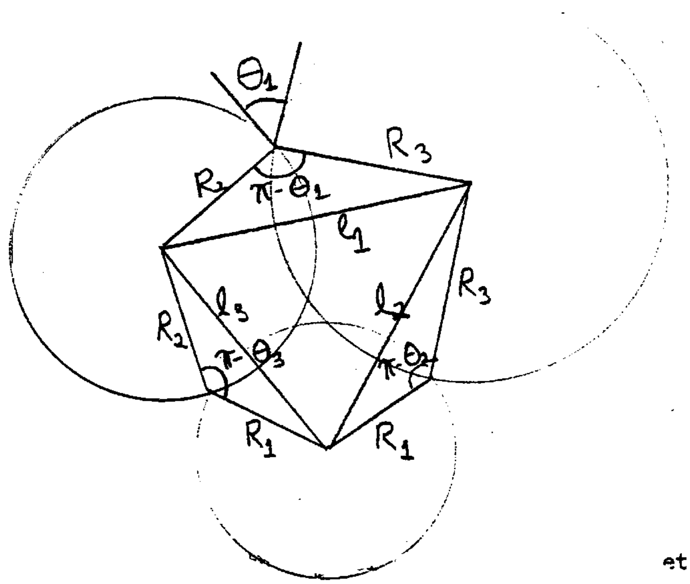

Proof of lemma. The length l_k of a side of the hypothetical triangle of centers of the circles is determined as the side opposite the obtuse angle $\pi - \theta_k$ in a triangle whose other sides are R_i and R_j. Thus, $\sup(R_i, R_j) < l_k \leq R_i + R_j$. The three numbers l_1, l_2 and l_3 obtained 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 Kda 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.

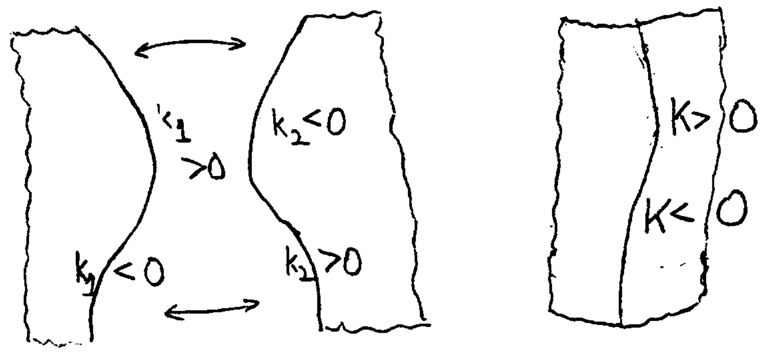

(The effect of gathering is more subtle—it is obtained by putting two lines infinitely close together, one with positive curvature and one with balancing negative curvature. Another instance of this is the boundary of a lens.)

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.*

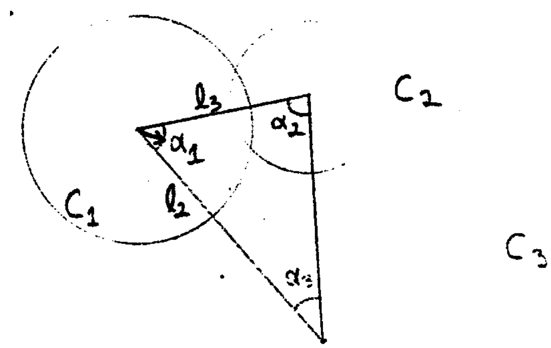

13.56

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.

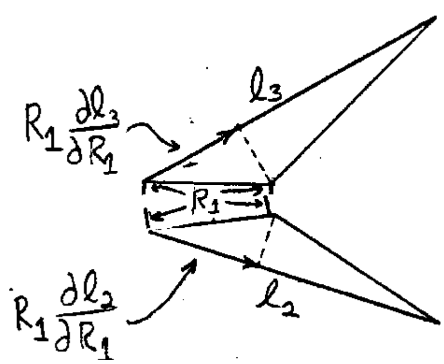

If v_1 denotes the center of C_1, then $\partial v_1/\partial R_1$ is determined as the vector whose orthogonal projections sides 2 and 3 are $\partial l_2/\partial R_1$ and $\partial l_3/\partial R_1$. Thus,

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

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

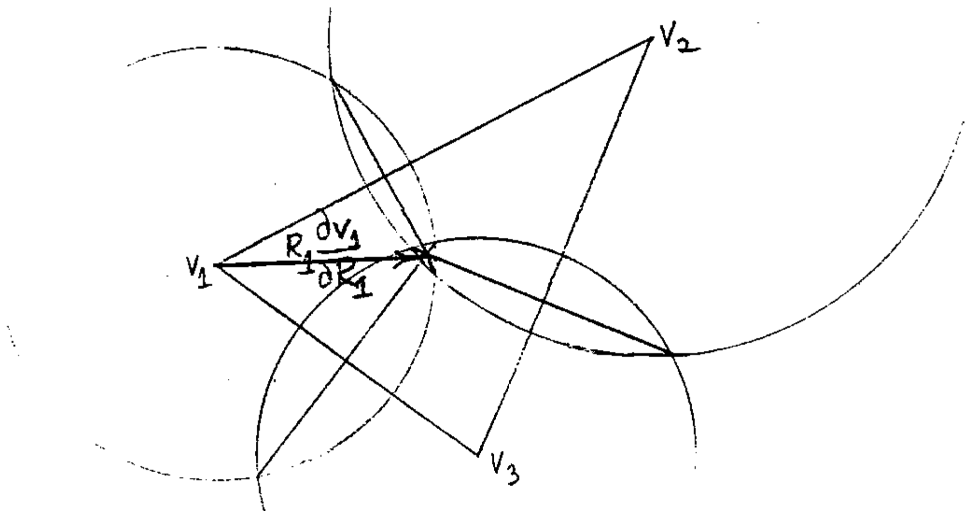

13.57

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

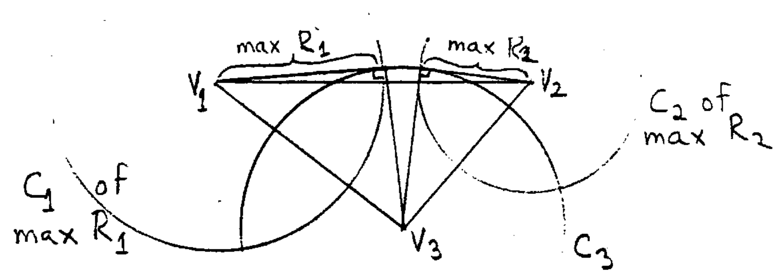

C_3 meets $\overline{v_1 v_2} \implies C_1$ and C_2 don't meet.

It follows that $\partial v_1/\partial R_1$ points to the interior of $\Delta v_1 v_2 v_3$.

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 $\beta_i + \gamma_i - \pi$.

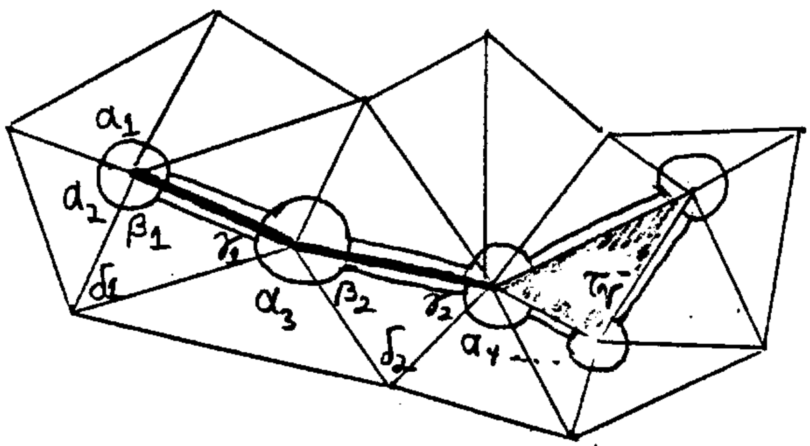

In view of 13.7.3, an angle such as α_1 increases in the R' metric. The change in β_1 and γ_1 is unpredictable. However, their sum must increase: first, let R_1 and R_2 decrease; $\pi - \delta_1 - (\beta_1 + \beta_2)$, which is the area of the triangle in the hyperbolic case, decreases or remains constant 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.

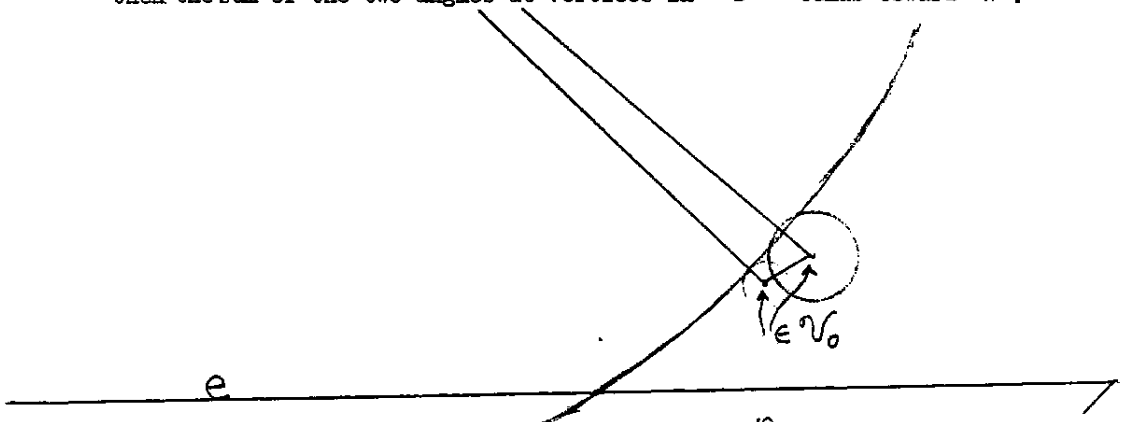

When a triangle of τ has only one vertex in $\mathcal{V}^0$, then the angle at that vertex tends toward the value $\pi - \Theta(e)$, where e is the opposite edge. Thus, the total curvature of $\partial S_{\mathcal{V}^0}$ tends toward the value $\sum_{e \in L(\tau_{\mathcal{V}^0})} (\pi - \Theta(e))$,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

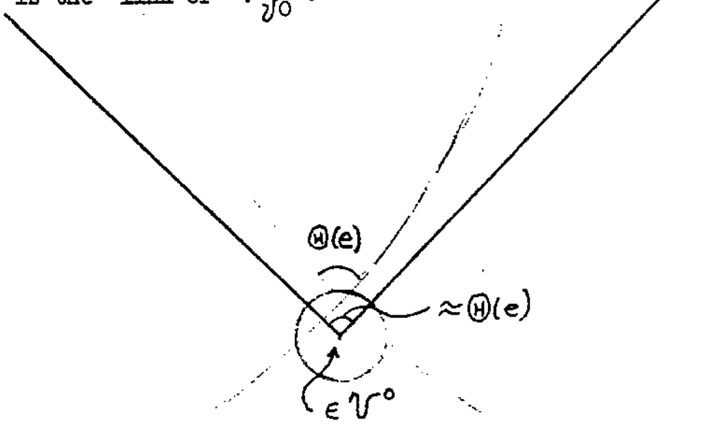

When a triangle of τ has only one vertex in $\mathcal{V}^O$, then the angle at that vertex tends toward the value $\pi - \Theta(e)$, where e is the opposite edge. Thus, the total curvature of $\partial S_{\mathcal{V}^O}$ tends toward the value

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

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:

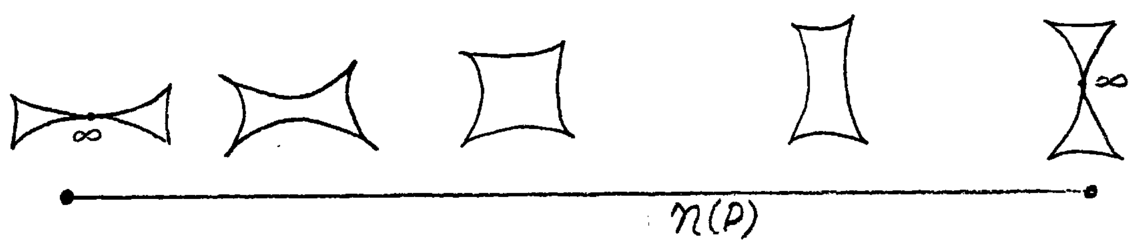

If there are two adjacent order 2 corner reflectors, the qualitative picture must be modified appropriately. For instance,

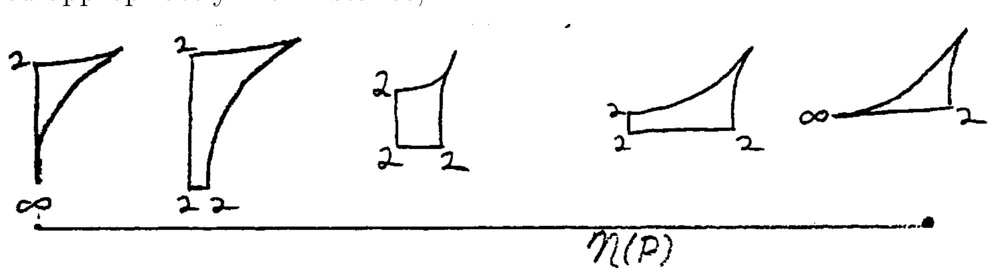

When P is a pentagon, $\mathcal{T}(P)$ is $\mathbb{R}^2$. There are five possible nodes, and the cell-structure is diagrammed below: 13.64

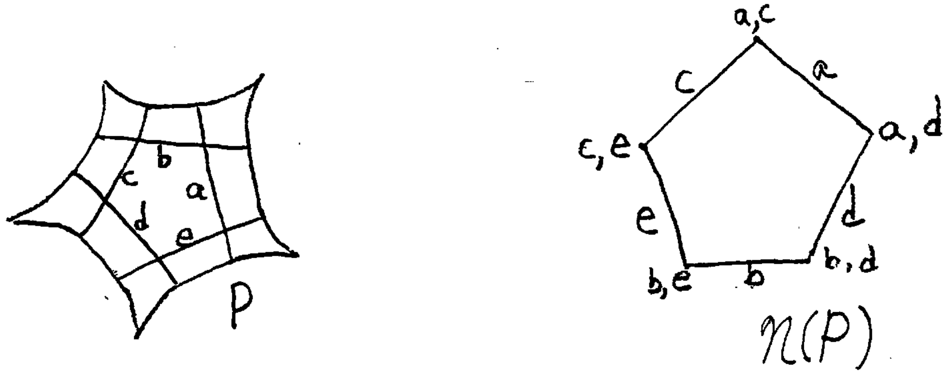

When there is only one node, the pentagon is pinched into a quadrilateral and a triangle, so there is still one degree of freedom.

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

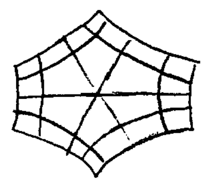

Each single node pinches the hexagon into a pentagon and a triangle, or into two quadrilaterals, so its associated 2-cell is a pentagon or a square. The cell division of ∂D^3 is diagrammed below:

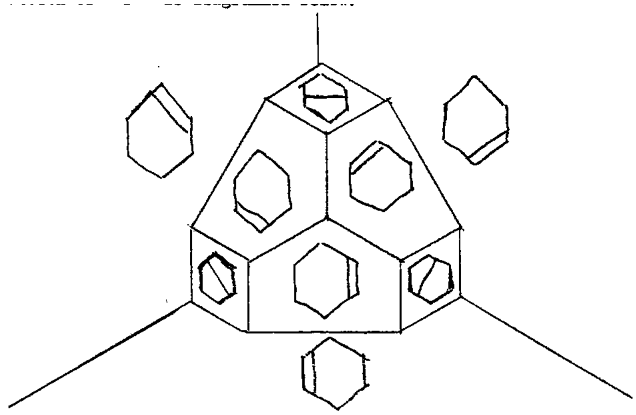

(The zero and one-dimensional cells are parametrized by the union of the nodes of 13.65
the incident 2-cells.)

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 $\mathcal{PL}(P)^{\ddagger}$ 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.

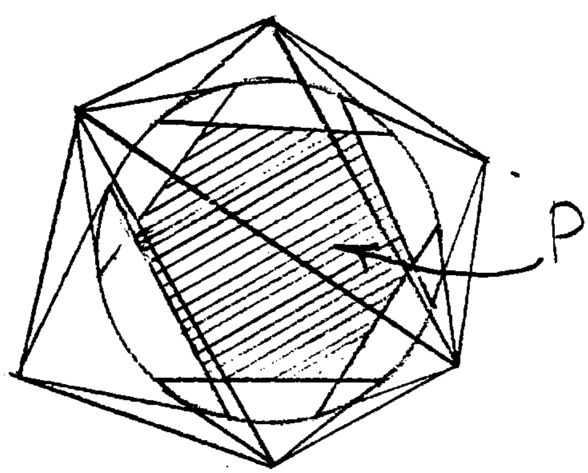

Take the model P now to be a regular polygon in $\mathbb{R}^2 \subset \mathbb{R}^3$. Let V be the vertex set. For any function $f : V \to \mathbb{R}$, let C_f be the convex hull of the set of points

[‡] 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. □

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

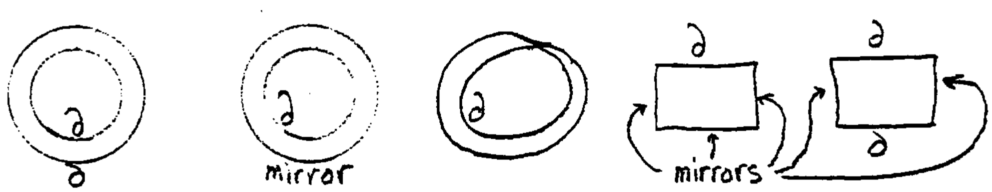

but for the orbifolds we are considering only the last two can occur.

(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.

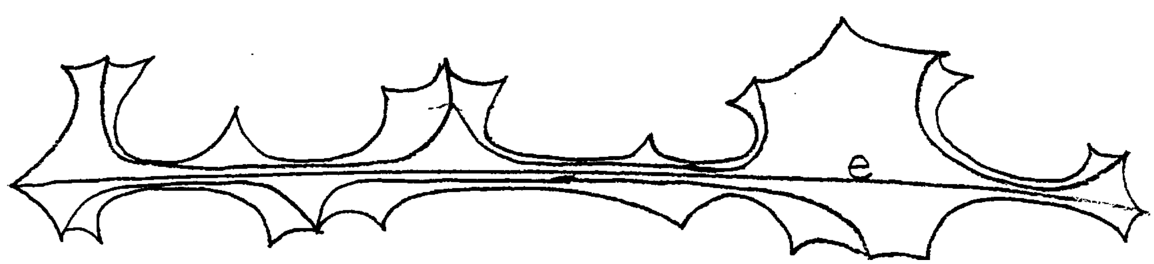

Similarly, these nearly parallel edges must be close and nearly parallel to still more edges on the far side from e. How long does this continue? Remember that H has an angle at each edge. In fact, if we ignore edges in the interior of ∂H, no angle exceeds 90°. Special note should be made here of the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

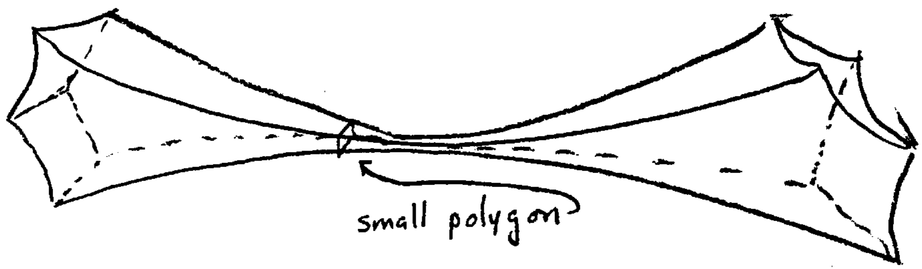

By area-angle considerations, this small suborbifold must have non-negative Euler characteristic. We investigate the cases separately.

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

(i) 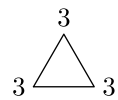

This is automatically incompressible, and since it is closed, it must be homotopic to a cusp. But this is supposed to be avoided by keeping our investigations away from the vertices of faces of P.

(ii) 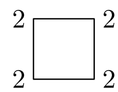

Either it is incompressible, and avoided as in (i), or compressible, so it is homotopic to some edge of H.

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

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

(i) 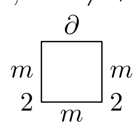

(ii) 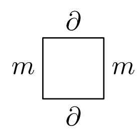

where m denotes a mirror.

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.

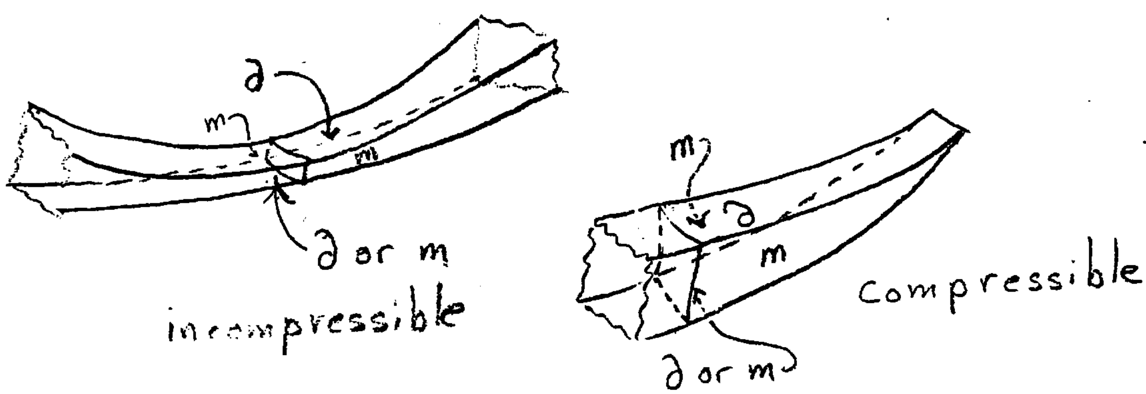

(c) $\chi > 0$. This can occur, since O is irreducible and ∂O incompressible.

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 . . . .