Zeitschrift: L'Enseignement Mathématique
Herausgeber: Commission Internationale de l'Enseignement Mathématique
Band: 40 (1994)
Heft: 1-2: L'ENSEIGNEMENT MATHÉMATIQUE

Artikel: AN EXPOSITION OF POINCARÉ'S POLYHEDRON THEOREM
Autor: Epstein, David B.A. / Petronio, Carlo
DOI: https://doi.org/10.5169/seals-61108

L'Enseignement Mathématique, t. 40 (1994), p. 113-170

AN EXPOSITION OF POINCARÉ'S POLYHEDRON THEOREM

by David B.A. EPSTEIN and Carlo PETRONIO [1])

1. INTRODUCTION

Poincaré's Theorem is an important, widely used and well-known result. There are a number of expositions in the literature (see [Mas71, Sei75, Mor78, Apa86, Mas88]). However, as far as we know, there is no source which contains a completely satisfying proof which applies to all dimensions and all constant curvature geometries. There is a tendency for unnecessary hypotheses to be included, which are sometimes implied by the other hypotheses and sometimes unnecessarily restrict the range of validity of the theorem.

A feature of this paper is the emphasis on the algorithmic aspects of Poincaré's Theorem. This point of view was first stressed by [Ril83]. Riley's work is restricted to dimensions two and three, where various points become easier to analyze. We want procedures which will tell us whether a given finite set of finite-sided convex polyhedra and face-pairings do or do not give rise to an orbifold, or, equivalently to a tessellation of $\mathbf{X}^n$. Such procedures have been exploited to remarkable effect — readers are referred to the paper by Bob Riley just cited, and to other contributions by him.

Riley's computer programs start with a list of group generators, given numerically, and attempt to find a fundamental domain for the group. The procedure goes through a check on a putative fundamental domain, along the lines explained in this paper. An essential further feature of his programs, and of similar programs by others, is that it incorporates another procedure for improving the guess on the shape of the fundamental domain, if the original guess fails. This second feature is not addressed in this paper.

This paper elaborates notes of lectures given by Epstein in 1992 at Warwick University. We will assume that the reader is familiar with the elementary definitions of euclidean, spherical and hyperbolic spaces and their geodesic

[1]) This research was supported by SERC and NSF.

subspaces. We denote these spaces by $\mathbf{E}^n$, $\mathbf{S}^n$ and $\mathbf{H}^n$. When we want to denote one of these three spaces, without specifying which, we will call it $\mathbf{X}^n$. Note that $\mathbf{E}^1$ and $\mathbf{H}^1$ are isometric to each other and locally isometric to $\mathbf{S}^1$.

DEFINITION 1.1 ($\mathbf{X}$-subspace). An $\mathbf{X}$-subspace of $\mathbf{X}^n$ is a copy of $\mathbf{X}^i$, for some i with $0 \leqslant i \leqslant n$, embedded geodesically. In the case when $\mathbf{X}$ is spherical, every 0-dimensional $\mathbf{X}$-subspace is a copy of $\mathbf{S}^0$ in $\mathbf{S}^n$, embedded as a pair of antipodal points.

When reading this paper, it will be helpful for the reader to understand the concept of a manifold or orbifold modelled on one of these spaces. The reader is referred to [Thu, Thu80] for these definitions. [BP92] is also a useful source of background matter.

Suppose G is a discrete group of isometries of $\mathbf{X}^n$. Then one can fix a point $p \in \mathbf{X}^n$ such that no element of G fixes p, and define the *Dirichlet domain of G with centre p*. This is the set of $x \in \mathbf{X}^n$ such that, for all $g \in G$, $d(x,p) \leqslant d(x,gp)$. The Dirichlet domain is a convex polyhedron (see Definition 2.1). It has a finite number of faces in many important cases (for example if the quotient of $\mathbf{X}^n$ by G is compact), but in general may have an infinite number of faces, even if G is cyclic as in Example 1.3. It is a fundamental domain for the action of G on $\mathbf{X}^n$.

EXAMPLE 1.2. Consider the free abelian group on two generators acting on the plane by translation by (m, n), where m and n are integers. The Dirichlet domain centred on any point is a unit square. If the free abelian group acts by translation but does not have a pair of generators acting by orthogonal translations, then the Dirichlet domain has six sides.

EXAMPLE 1.3 (infinitely many faces). Take the isometry of $\mathbf{E}^3$ which is translation along the z-axis followed by an irrational rotation about the z-axis. Let G be the cyclic infinite group generated by this isometry. The Dirichlet domain centred on any point not on the z-axis has an infinite number of faces.

Each element $g \in G$ gives rise to a hyperplane A_g of $\mathbf{X}^n$, consisting of points which are equidistant from p and $g(p)$. Let F_g be the intersection of the Dirichlet domain with A_g. If F_g has dimension $n-1$, it is a face of the Dirichlet domain, and each face of dimension $n-1$ is equal to F_g for some g. If F_g is a face, so is $F_{g^{-1}}$. The element g sends $F_{g^{-1}}$ to F_g and is called a *face-pairing* of the Dirichlet domain.

Now suppose one is presented with a convex polyhedron and, for each face, an isometry pairing it with another face. Poincaré's Theorem is concerned with the question "Can this be the fundamental domain and face-pairings for the action of a discrete group of isometries?" It turns out that this question can be answered with conditions which are surprisingly simple to check, and the answer is the content of Poincaré's Theorem. If the polyhedron has a finite number of faces, the conditions for Poincaré's Theorem can be checked algorithmically.

Here is another point of view on Poincaré's Theorem. A manifold or orbifold modelled on $\mathbf{X}^n$ can be cut along $(n-1)$-dimensional geodesic subspaces to obtain a single convex polyhedron (the fundamental domain), as in the case of the Dirichlet domain. However, it may sometimes be convenient not to use a single polyhedron. We could for example take an arbitrary fundamental domain which is not necessarily convex (and does not necessarily have geodesic faces). This fundamental domain can be approximated by a union of convex pieces. The more convoluted the fundamental domain, the greater the number of convex pieces we might need for a reasonable approximation. As has been pointed out in [Bow93], in dimensions greater than four a geometrically finite discrete hyperbolic group may have a fundamental domain which needs to be built up from a finite number of finite-sided convex pieces — one such does not suffice.

In [Bow93] one finds an example of a four-dimensional hyperbolic manifold which has a fundamental domain with a finite number of faces, all geodesic, but such that no Dirichlet region has finitely many faces. Probably Bowditch's example has the property that every convex fundamental domain, even if not a Dirichlet domain, has to have an infinite number of faces. Since it is nice to use convex building blocks — for example, they can easily be specified using a finite set of real numbers — we would probably want to decompose the fundamental domain in such a case into a finite number of convex polyhedra.

Now suppose we are given a set of convex polyhedra with face-pairings. The role of Poincaré's Theorem is to determine whether this situation can arise from a manifold or orbifold by cutting along geodesic codimension-one subspaces. In each case there is also the problem of determining the associated (fundamental) group from the combinatorial data presented.

Poincaré's Theorem can be used to construct many interesting examples of groups acting on hyperbolic or euclidean space or on the sphere, and many interesting manifolds and orbifolds modelled on one of these spaces. Readers are referred to [Thu, Thu80] for such examples.

There are several reasons why it is better to use several convex building blocks than only one. Firstly, as we have already pointed out, this is necessary if we are to deal with all geometrically finite groups. Secondly many of the most interesting examples are constructed using more than one piece, for example the two ideal regular hyperbolic tetrahedra used to give a complete hyperbolic structure to the figure-eight complement (see [Thu, Thu80]). Thirdly the hypotheses come up naturally in the proof; if one starts with a single convex piece, the natural inductive proof inexorably leads one to consider glueing together several convex pieces in lower dimensions. Fourthly, it may be convenient to use a non-convex fundamental domain, rather than a convex fundamental domain. The non-convex fundamental domains that arise in practice can be cut into a finite number of convex pieces, making our hypotheses applicable.

One way in which our treatment differs from all previous treatments, is that we do not assume we start with an embedded fundamental domain. Instead the fundamental domain is expressed as the union of convex cells, each of which can be separately embedded, without knowing to begin with that their union can be embedded. For example, suppose we are given three planar wedges of angle $5\pi/6$, $6\pi/7$ and $7\pi/8$ with face-pairings glueing them together. The union of these pieces cannot form a fundamental domain, because their union after glueing cannot be embedded. The point here is whether this non-embeddability or embeddability needs to be checked beforehand. Our proof shows that the usual checks for Poincaré's Theorem, in the case where there is only one convex piece, in any case imply the embeddability of the potential fundamental domain, so no special separate check is necessary. In this case the extra necessary checking is easy, but in a more complicated situation, the algorithm presented here could lead to significant saving of time and complication.

2. Convex polyhedra

Let $\mathbf{X}^n$ be hyperbolic, euclidean or spherical n-dimensional space. A hyperplane (that is, a codimension-one $\mathbf{X}$-subspace) divides $\mathbf{X}^n$ into two components; we will call the closure of either of them a *half-space* in $\mathbf{X}^n$. Any $\mathbf{X}$-subspace is the intersection of hyperplanes, and *vice versa*.

Definition 2.1 (convex polyhedron). A connected subset P of $\mathbf{X}^n$ is called a *convex polyhedron* if it is the intersection of a family $\mathscr{H}$ of half-spaces with the property that each point of P has a neighbourhood meeting at most a finite number of boundaries of elements of $\mathscr{H}$. A convex polyhedron in $\mathbf{X}^n$ is said to be *thick* in $\mathbf{X}^n$ if it has non-empty interior.

REMARK 2.2 (antipodal points). In $\mathbf{H}^n$ and $\mathbf{E}^n$ any two points are joined by a unique geodesic segment, so the same property holds in any intersection of half-spaces. In particular intersections of half-spaces are connected. In $\mathbf{S}^n$, we have to make do with a slightly weaker form of this, in which any two points x and y, such that $d(x, y) < \pi$, are joined by a unique shortest geodesic segment, in any intersection of half-spaces. In $\mathbf{S}^n$ a pair of antipodal points can be obtained as the intersection of $n + 1$ half-spaces. Furthermore one can easily check that if an intersection P of half-spaces in $\mathbf{S}^n$ does not enjoy the property that any two points of P are joined by a geodesic arc within P, then P must be a pair of antipodal points. A single point in $\mathbf{S}^n$ is of course an intersection of half-spaces. So the only intersection of a locally finite family of half-spaces which is not a convex polyhedron is a pair of antipodal points in the sphere.

LEMMA 2.3 (interior). *An intersection P of half-spaces in $\mathbf{X}^n$ either has non-empty interior in $\mathbf{X}^n$ or is contained in a hyperplane. Moreover, if the interior of P is not empty, it is dense in P.*

Proof of 2.3. We may suppose that $P \neq \varnothing$. Let $\mathscr{S}$ be the set of non-empty $\mathbf{X}$-subspaces S of $\mathbf{X}^n$ such that $P \cap S$ has non-empty S-interior (V, say) and such that V is dense in $P \cap S$. Clearly $\mathscr{S}$ has a 0-dimensional member, so it is not empty. Let S be a maximal element of $\mathscr{S}$.

We claim that $P \subset S$. Otherwise, let $x \in P \backslash S$ and let S' be a minimal $\mathbf{X}$-subspace containing both x and S. Let $V \subset P \cap S$ be the S-interior of $P \cap S$. By definition V is not empty.

In the spherical case the antipodal point to x is not in $V \subset S$, since $x \notin S$. So for any point in V, there exists a unique shortest geodesic path joining it to x.

The whole "cone" based on V with vertex x is contained in $P \cap S'$ and this easily implies that x and $P \cap S$ are in the closure of the S'-interior of $P \cap S'$. This argument can be repeated for all $x \in (P \cap S') \backslash S$. Hence $S' \in \mathscr{S}$, which gives a contradiction.

Our claim is proved and the conclusion follows. $\square$

We define the *dimension* of an intersection P of half-spaces in $\mathbf{X}^n$ (in particular of a convex polyhedron) as the smallest integer i such that P is contained in an i-dimensional $\mathbf{X}$-subspace of $\mathbf{X}^n$. Lemma 2.3 shows that P is then thick in this subspace and the subspace is uniquely determined. A non-empty intersection of a convex polyhedron in $\mathbf{X}^n$ with an $\mathbf{X}$-subspace S of $\mathbf{X}^n$ is either a convex polyhedron in S or possibly a pair of antipodal points in the spherical case.

Let P be a convex polyhedron in $\mathbf{X}^n$. We define the *relative boundary* ∂P of P to be the topological boundary of P in S where S is the unique $\mathbf{X}$-subspace of $\mathbf{X}^n$ in which P is thick. The *relative interior* of P, denoted $\mathrm{RelInt}(P)$, is defined to be $P \backslash \partial P$. Both "interior" and "boundary" of P coincide with the topological interior and boundary respectively if and only if P is thick.

Let P be a convex polyhedron. A subset Q of ∂P is said to be a *codimension-one face* of P if P is thick in $\mathbf{X}^n$, $Q = P \cap S$ for some hyperplane S of $\mathbf{X}^n$, and Q is thick in S. (An exception has to be made when P is a semicircle and ∂P is a pair of antipodal points. In that case, we insist that Q is equal to one of the boundary points.) If $i \geqslant 2$, the *codimension-i* faces of P are defined inductively as codimension-one faces of codimension-$(i-1)$ faces of P. If P is thick in $\mathbf{X}^n$, a codimension-i face of P is a convex polyhedron of dimension $n - i$. Each codimension-i face of a convex polyhedron is contained in a face of codimension $i - 1$.

LEMMA 2.4 (boundary). *Let P be a thick convex polyhedron in $\mathbf{X}^n$ which is the intersection of a locally finite family $\mathscr{H}$ of half-spaces. Then*

$$\partial P = \bigcup_{H \in \mathscr{H}} P \cap \partial H .$$

Proof of 2.4. Let $x \in \partial P$ and let U be an open neighbourhood of x. Let $\{H_1, \ldots, H_k\}$ be the set of elements of $\mathscr{H}$ whose boundary meets U. If U is small then k is finite, and we may assume that $x \in \partial H_i$ for $1 \leqslant i \leqslant k$. We must have $k \geqslant 1$, for, if $k = 0$, x would be in the interior of P in $\mathbf{X}^n$.

Conversely, if $x \in P \cap \partial H$ for some $H \in \mathscr{H}$, then x is in the topological boundary of P in $\mathbf{X}^n$. $\square$

PROPOSITION 2.5 (essential faces). *Let P and $\mathscr{H}$ be as in Lemma 2.4. Set*

$$\mathscr{M} = \left\{ H_0 \in \mathscr{H} : P \neq \bigcap_{H \in \mathscr{H} \backslash \{H_0\}} H \right\} .$$

Then:

(a) *P is the intersection of the elements of $\mathscr{M}$;*

(b) *the elements of* $\mathscr{M}$ *are characterized as the elements* H_0 *of* $\mathscr{H}$ *such that* $P \cap \partial H_0$ *is thick in* ∂H_0*;*

(c) *the set* $\mathscr{M}$ *of half-spaces depends only on* P *and not on* $\mathscr{H}$.

Note that neither Proposition 2.5 nor Lemma 2.4 need be true when the family of half-spaces is not locally finite. For example, the closed unit ball in $\mathbf{R}^n$ is the intersection of a countable family of half-spaces, none of whose boundaries meets the closed unit ball.

Proof of 2.5. Any element of $\mathscr{H} \setminus \mathscr{M}$ can be omitted from $\mathscr{H}$ without affecting P. It follows that any finite number of elements of $\mathscr{H} \setminus \mathscr{M}$ can be omitted without affecting P. Let P' be the intersection of the elements of $\mathscr{M}$. Then $P \subset P'$. If P' is not connected, then P' must consist of two antipodal points and P must be a single point. But this contradicts the definition of $\mathscr{M}$, and so P' is connected. By the local finiteness property, every point of P has a neighbourhood U such that $P \cap U = P' \cap U$. This shows that P is an open subset of P'. Since P' is connected and P is a non-empty closed subset of $\mathbf{X}^n$, $P = P'$.

Assume that $H_0 \in \mathscr{M}$. Let P_0 be the intersection of the elements of $\mathscr{H} \setminus \{H_0\}$ and choose $x \in P_0 \setminus P$. Consider an open set U internal to P, and let C be the cone over U with vertex x. As shown in Figure 1, $C \cap \partial H_0$ is contained in P and has non-empty interior in ∂H_0, which implies that $P \cap \partial H_0$ is thick in ∂H_0.

Conversely, if x is in the ∂H_0-interior of $P \cap \partial H_0$, the only half-space containing P and having x on its boundary is H_0. Therefore, if H_0 is omitted,

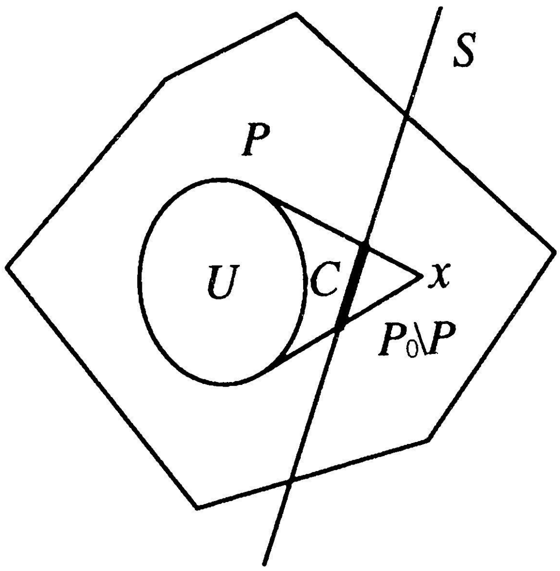

FIGURE 1.
Thick intersections.
If a half-space is essential for the definition of a polyhedron then its intersection with the polyhedron is thick. In the diagram the boundary ∂H_0 of H_0 is denoted by S.

x becomes an interior point of the intersection of half-spaces. So $H_0 \in \mathscr{M}$. The same argument proves that the elements of $\mathscr{M}$ can be characterized independently of $\mathscr{H}$ as the half-spaces H containing P and such that $P \cap \partial H$ is thick in ∂H. $\square$

The elements of the set $\mathscr{M}$ described in Proposition 2.5 are called the *essential half-spaces* of P. According to Proposition 2.5, the essential half-spaces are exactly those whose boundaries contain codimension-one faces of P. Lemma 2.4 implies the following result.

COROLLARY 2.6 (union of faces). *The boundary of a thick convex polyhedron in* $\mathbf{X}^n$ *is the union of its codimension-one faces.*

LEMMA 2.7 (codimension-two faces). *If* P *is a convex polyhedron in* $\mathbf{X}^n$ *and* C *is a codimension-two face of* P *there exist exactly two codimension-one faces of* P *containing* C.

Proof of 2.7. Without loss of generality we can assume P is thick in $\mathbf{X}^n$. Let S be the codimension-two subspace containing C. We may suppose that P is defined by its essential half-spaces. It follows from our definition of a face that there exist at least two essential half-spaces H_1 and H_2 whose boundary contains S. So C is contained in the codimension-one faces $P \cap \partial H_1$ and $P \cap \partial H_2$. Conversely if a codimension-one face $P \cap \partial H$ contains C then ∂H contains S. But it is easily checked (see Figure 2) that there cannot be three essential half-spaces whose boundaries have a codimension-two subspace in common. $\square$

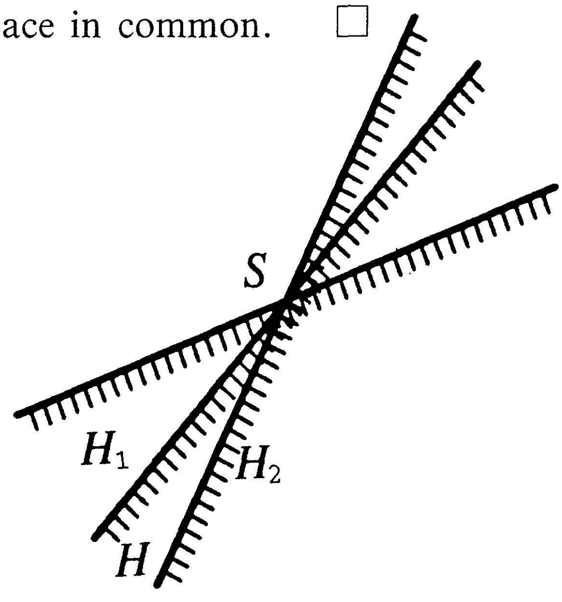

FIGURE 2.
Inessential half-spaces.
If three hyperplanes meet in a codimension-two subspace one of the corresponding half-spaces is not essential.

Let $n \geqslant 2$. A *dihedral region with corner* S is defined to be the intersection of two half-spaces, whose boundaries intersect in a subspace S of codimension two. The *dihedral angle* of the dihedral region is defined to be the angle between the boundaries. This is measured by taking a two-dimensional subspace orthogonal to S and seeing what angle is marked out on it by the boundaries. If we think of one half-space as first and the other as second, and if we orient the orthogonal plane, then the dihedral angle θ is signed and $0 < |\theta| < \pi$. The definition can be extended to the case where the boundaries of the half-spaces coincide. If the half-spaces themselves coincide, the angle is defined (ambiguously) to be $\pm\pi$, and if the half-spaces have the same boundary, but are otherwise disjoint, the angle is defined to be zero.

DEFINITION 2.8 (convex cell). A *convex cell* is a slight generalization of a convex polyhedron in $\mathbf{X}^n$; it is a convex polyhedron whose proper faces may have been subdivided. Formally, a convex cell is a convex polyhedron P in $\mathbf{X}^n$, together with a locally finite collection of convex polyhedra $\{P_i\}_{i \in I}$ satisfying the following conditions:

(a) The relative interiors of P and of the P_i, $(i \in I)$, form a disjoint covering of P.

(b) For each $i \in I$, P_i together with $\{P_j \mid j \in I,\ P_j \subset \partial P_i\}$ is a convex cell. (This definition is not circular since the dimension of P_i is smaller than that of P.)

The P_i are called the *faces* of the convex cell. By abuse of notation, we will often denote the convex cell by P, without mentioning the P_i. The most obvious example of a convex cell is a convex polyhedron, together with all its proper faces. A convex cell is said to be thick in $\mathbf{X}^n$ if the underlying polyhedron is thick in $\mathbf{X}^n$.

We now present some lemmas which will be useful in the sequel.

LEMMA 2.9 (positive distance 1). *Two disjoint affine subspaces of* $\mathbf{E}^n$ *have positive distance from each other.*

Proof of 2.9. Consider the orthogonal projection to an orthogonal complement of one of the subspaces, and note that distances are not increased. It follows that we can assume that one of the subspaces is a point, in which case the conclusion is obvious. □

LEMMA 2.10 (positive distance 2). *Let* S, T *be affine subspaces of* $\mathbf{E}^n$ *and let* $S \cap T = V \neq \varnothing$. *We assume that* $S \neq V$. *Let* $\varepsilon > 0$ *and define* $\hat{S}_\varepsilon = \{s \in S : d(s, V) \geqslant \varepsilon\}$. *Then* $d(\hat{S}_\varepsilon, T) > 0$.

Proof of 2.10. Assume first that the intersection V is a point. We may take $V = \{0\}$ with respect to the usual coordinates of $\mathbf{R}^n \cong \mathbf{E}^n$. As s varies in $S \setminus \{0\}$ and t varies in T, the distance between $s/\| s \|$ and t is bounded away from zero by compactness of the unit sphere in S. This proves the result when V is a point.

Now consider the general case. Let π be the projection on some orthogonal complement of V. Then

$$d(T, \hat{S}_\varepsilon) = d(\pi T, \pi \hat{S}_\varepsilon) > 0$$

as we see from the case where V is a point. $\square$

PROPOSITION 2.11 (positive distance 3). *Let A and B be disjoint convex cells in the sphere or in euclidean space, each having only a finite number of faces. Then they are a positive distance apart.*

Proof of 2.11. This fact is obvious in the sphere, by compactness.

We prove the assertion by induction on the sum of the dimensions of A and B, which we denote by m. The case $m = 0$ is obvious, so we assume that $m > 0$ and that the assertion is true for all integers less than m. Assume by contradiction that there exist sequences $\{a_i\} \subset A$ and $\{b_i\} \subset B$ such that $d(a_i, b_i) \to 0$.

First of all we can assume that there is a $\delta > 0$, such that, for all i, $d(a_i \partial A) \geqslant \delta$; otherwise, using the fact that there are only finitely many faces, we can find a subsequence (which we denote by $\{a_i\}$ as well) and a proper face F of A such that $d(a_i, F) \to 0$; if we choose $\tilde{a}_i \in F$ such that $d(\tilde{a}_i, a_i) \to 0$, we have $d(\tilde{a}_i, b_i) \to 0$. The induction hypothesis applies to the faces F and B, proving that they meet, and this is a contradiction. Similarly, we can assume that the distance between the b_i's and ∂B is bounded away from 0; we can assume the same bound δ works for both.

Now, let S and T be the minimal subspaces containing A and B respectively. We claim that $S \not\subset T$ and $T \not\subset S$. Suppose for example that $S \subset T$, and choose i so that $d(a_i, b_i) < \delta$. Then $a_i \in T \cap B_\delta(b_i) \subset B$, which is false. So we assume that $S \neq T$. Lemma 2.9 implies that $V = S \cap T \neq \varnothing$, and Lemma 2.10 implies that $d(a_i, V) \to 0$. Then we can find $\{v_i\} \subset V$ such that $d(v_i, a_i) \to 0$, and hence $d(v_i, b_i) \to 0$. Since A is thick in S, as soon as $d(v_i, a_i) \leqslant \delta$ we have $v_i \in A$, and similarly if $d(v_i, b_i) \leqslant \delta$ we have $v_i \in B$. This is a contradiction. $\square$

LEMMA 2.12 (constant multiple). *Let P be a convex polyhedron in $\mathbf{X}^n$ with only a finite number of faces. Let E be a face of P and* *let S be the subspace of $\mathbf{X}^n$ containing E in which E is thick. Then there exists a constant $k > 0$ such that, for $x \in P$, $d(x, S) \geqslant k \cdot d(x, E)$. (We make an exception of the case where S is a pair of antipodal points — the result may then be false.)*

Proof of 2.12. We will obtain a contradiction by assuming that there exists a sequence (x_i) in $P \backslash E$ for which $d(x_i, S)/d(x_i, E) \to 0$. In the spherical case $d(x_i, E) \leqslant \pi$, so $d(x_i, S) \to 0$. Using compactness, it follows that $d(x_i, E)$ also converges to zero. Therefore we may restrict our attention to an approximately euclidean local picture. So we assume from now on that we are in the hyperbolic or euclidean case.

Given $x \in P \backslash E$, let y be the nearest point in S and let z be the nearest point in E. Let E_0 be the face of E containing z in its relative interior. The geodesic xz is orthogonal to E_0 and xy is orthogonal to S. Moreover the segment yz meets E only at z.

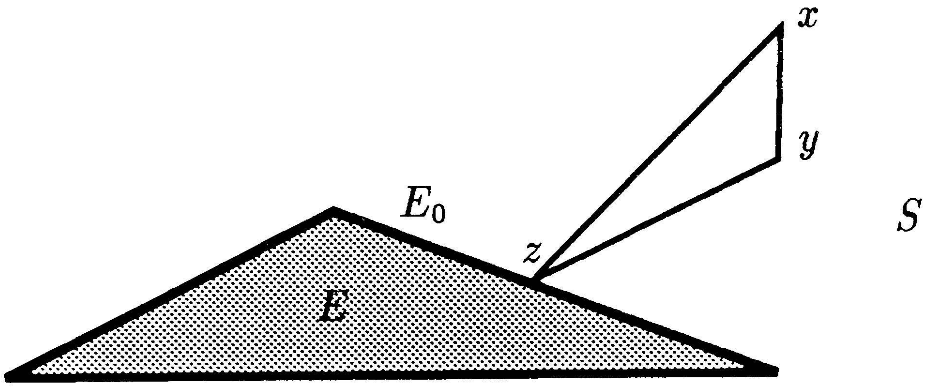

FIGURE 3.
Distance to face and subspace.
This picture illustrates the proof of Lemma 2.12. E is a face which is thick in the subspace S. The point z is the nearest point in E to x, and the smallest face containing z is E_0. The point y is the nearest point in S to x.

In our proof by contradiction, we obtain a sequence $x_i \in P \backslash E$ and corresponding sequences y_i and z_i, defined as above, such that $d(x_i, y_i)/d(x_i, z_i)$ converges to zero. This means that the angle between the segment $x_i z_i$ and S converges to zero. Since there are only a finite number of faces, we may assume that z_i lies in the relative interior of the same E_0 for each i. For each i, without changing the angle $\angle\, x_i z_i y_i$, we may now, without loss of generality, move x_i nearer to z_i along the ray $x_i z_i$, keeping z_i fixed and moving y_i correspondingly; z_i remains the nearest point of E. This moves y_i along the ray $y_i z_i$. The ratio $d(x_i, y_i)/d(x_i, z_i)$ is unaltered by the movement in the euclidean case, and is decreased in the hyperbolic case. We may therefore assume that, for each essential half-space H of P such that ∂H does not contain E_0 and for each i, $d(x_i, y_i) \leqslant d(x_i, z_i) < d(z_i, \partial H)/2$. Therefore $d(z_i, y_i) < d(z_i, \partial H)$.

It follows that for each essential half-space H of P, such that $y_i \notin H$, ∂H must contain E_0. Other half-spaces now become irrelevant, and we may assume that the boundary of each essential half-space of P contains E_0. The local geometry is therefore unchanged as z_i moves in $\mathrm{RelInt}(E_0)$, and we may assume without loss of generality that $z_i = z \in \mathrm{RelInt}(E_0)$ is independent of i.

Since only the angles are important, we may now move x_i along the ray zx_i, and assume that $d(x_i, z) = r$ is independent of i. Then z remains the nearest point of E to x_i and $d(x_i, y_i)$ tends to zero. We see that x_i converges (after taking a subsequence) to a point $x \in S$. It follows that $x \in P \cap S$ and therefore $x \in E$. But this contradicts the fact that z is the nearest point of E to x_i. This contradiction proves the result. □

LEMMA 2.13 (positive distance 4). *Let P, Q be convex polyhedra with a finite number of faces in $\mathbf{X}^n$ and let $E = P \cap Q$ be a non-empty face of both P and Q. We assume that $E \neq P$. Let $\delta > 0$ and denote by $N_\delta(E)$ the δ-neighbourhood of E. Then, provided δ is small enough so that P is not contained in $N_\delta(E)$, we have $d(P \setminus N_\delta(E), Q) > 0$.*

Proof of 2.13. The result is clearly true in the spherical case, so we assume that $\mathbf{X}^n$ is hyperbolic or euclidean space. Let S be the subspace containing E in which E is thick. By Lemma 2.12 there exists $\delta' > 0$ such that $P \setminus N_\delta(E) \subset P \setminus N_{\delta'}(S)$. Let P' be the intersection of the essential half-spaces of P whose boundary contains E and define Q' similarly. It is easy to check that $P' \cap Q' = S$ by considering a neighbourhood of a point of $\mathrm{RelInt}(E)$. Now $P \subset P'$ and $Q \subset Q'$, so it is sufficient to check that

$$d(P' \setminus N_{\delta'}(S), Q') > 0\,.$$

But now, everything is invariant under isometries which preserve S and act trivially in the direction normal to S, so we can work in an orthogonal complement to S.

Hence we only have to check that the result holds when P is replaced by P', Q is replaced by Q' and $P' \cap Q' = E = S$ is a point. We argue by contradiction. Let $x_i \in P' \setminus N_{\delta'}(E)$ and $y_i \in Q'$, be sequences such that $d(x_i, y_i)$ converges to zero. We may clearly assume that $d(y_i, E) > \delta'/2$ for each i. The rays Ex_i and Ey_i (extended indefinitely) therefore converge to the same ray, which must lie in both P' and Q'. This contradicts the fact that $P' \cap Q' = E$. □

DEFINITION 2.14 (link). Let $(P, \{P_i\}_{i \in I})$ be a thick convex cell in $\mathbf{X}^n$, and let E be one of the faces of P — that is, E is equal to one of the P_j, $(j \in I)$. Let $J \subset I$ be the set of indices j such that P_j contains E. Let $p \in \mathrm{RelInt}(E)$. Let S_p be a sphere in $\mathbf{X}^n$ with centre p, whose radius is chosen small enough so that it only meets faces of P which contain E. By a change of scale, S_p can be identified with $\mathbf{S}^{n-1}$. The *link* of p in P is defined to be a convex cell in $\mathbf{S}^{n-1}$, given by $S_p \cap P$, with the face structure given by $S_p \cap P_j$. There is one exceptional situation we need to discuss, when E is one-dimensional. In that case, $S_p \cap E$ consists of two points, and this gives rise to two zero-dimensional faces in the link, not one. Note that if E is a point, then, for each $j \in J$, $P_j \cap S_p$ is a convex polyhedron in S_p — the exceptional case of two antipodal points cannot arise since E is in the relative boundary of P_j.

Notice that it does not matter where we choose $p \in \mathrm{RelInt}(E)$, as there is an isometry between the links given by two different choices. This means that up to isometry the link depends only on E and not on p.

3. CONDITIONS FOR POINCARÉ'S THEOREM

We describe in this section various conditions which come up when we are given a set of convex cells and instructions for glueing them together: our basic objective (see Remark 3.6) is to make orbifolds or manifolds from these building blocks. Alternatively, we can express our basic objective as constructing a tessellation of hyperbolic or euclidean space or the sphere.

Let $n \geqslant 2$. Let $\mathscr{P}$ be a countable or finite set of thick convex cells in $\mathbf{X}^n$.

REMARK 3.1.

(a) In fact we are only interested in the members of $\mathscr{P}$ up to isometry, and all our considerations must take this into account. This means that any $P \in \mathscr{P}$ may be replaced by $\psi(P)$, where $\psi \in \mathrm{Isom}(\mathbf{X}^n)$, and this must not affect any of our considerations in an essential way.

(b) Strictly speaking, the set $\mathscr{P}$ is an indexed set — that is, we allow repetition. One could avoid this, using Remark 3.1(a), by moving each repeated convex cell a little to a different place, but that seems artificial.

We denote by $\mathscr{F}(\mathscr{P})$ the set of all pairs (F, P) as P varies over $\mathscr{P}$ and F varies over the codimension-one faces of P. Notice that two faces of different convex cells could be geometrically coincident, but nonetheless they must be viewed as distinct according to Remark 3.1(a).

CONDITION 3.2 (Pairing). Suppose we are given maps $R : \mathcal{F}(\mathcal{P}) \to \mathcal{F}(\mathcal{P})$ and $A : \mathcal{F}(\mathcal{P}) \to \mathrm{Isom}(\mathbf{X}^n)$ with the following properties:

(a) $R : \mathcal{F}(\mathcal{P}) \to \mathcal{F}(\mathcal{P})$ is an involution, that is $R \circ R$ is the identity.

(b) Let $(F, P) \in \mathcal{F}(\mathcal{P})$ and let $R(F, P) = (F', P')$. Then $A(F, P) \in \mathrm{Isom}(\mathbf{X}^n)$ maps F onto F' and maps the interior of P to the other side of F' from the interior of P'.

(c) $A(F, P)$ gives an isomorphism between the face structure of F and the face structure of F'.

(d) For each $(F, P) \in \mathcal{F}(\mathcal{P})$, $A(R(F, P)) = A(F, P)^{-1}$.

In that case, we say that (R, A) is a *face-pairing* for $\mathcal{P}$, and say the condition Pairing$(\mathcal{P}, R, A)$ is satisfied. (R, A) is also known as *glueing data.*

REMARK 3.3 (order two). In case $R(F, P) = (F, P)$ Condition 3.2(d) implies that $A(F, P)$ is a mapping of order two. Note that in this special situation $A(F, P)$ is not necessarily the reflection in the face F, though that is a common application of this theory.

EXAMPLE 3.4 (triangle example). Consider an equilateral triangle P in $\mathbf{E}^2$, and let $\mathcal{P} = \{P\}$. In this case a face-pairing is an isometry sending an edge to itself or another edge. For each pair of edges there are four such isometries of $\mathbf{E}^2$, but two of the four are excluded by Condition 3.2(b). This enables one to easily list all possible sets of face-pairings. (In fact there are twenty distinct sets of face-pairings.)

CONDITION 3.5 (connected). Connected$(\mathcal{P}, R)$ is the condition that, given any two convex cells P and P' in $\mathcal{P}$, there exists a finite sequence of elements $\{(F_i, F_i', P_i)\}_{i=1,\dots,k}$ with $P_i \in \mathcal{P}$ and F_i and F_i' codimension-one faces of P_i, such that $P_1 = P$, $P_k = P'$ and $R(F_i', P_i) = (F_{i+1}, P_{i+1})$ for $i \geqslant 1$. This condition means that any two elements of $\mathcal{P}$ are joined by a sequence of face-pairings.

REMARK 3.6 (basic objective). If Pairing$(\mathcal{P}, R, A)$, we can glue up $\mathcal{P}$ and obtain an identification space $Q = Q(\mathcal{P}, R, A)$. If we remove the $(n-2)$-skeleton, we obtain a manifold M modelled on $\mathbf{X}^n$ which falls into pieces if we remove the $(n-1)$-skeleton; each piece is the interior of some $P \in \mathcal{P}$. The universal cover of M is also divided into cells, each of which is isometric to (the interior of) some element $P \in \mathcal{P}$. If Connected$(\mathcal{P}, R)$, then M is connected, and its universal cover is mapped into $\mathbf{X}^n$ by the developing map. Different cells in the universal cover will in general correspond to the same P, because M is not simply connected. The developing map is uniquely defined, once the map is fixed on one component of the inverse image of one element of $\mathscr{P}$. Roughly speaking, our basic objective is to find conditions such that the closures in $\mathbf{X}^n$ of the images of the cells of the universal cover tessellate $\mathbf{X}^n$.

More precisely, we start with countably many copies of the elements of $\mathscr{P}$ and lay them out in $\mathbf{X}^n$ one by one. Each new copy has to be glued to a free face of what is already laid out, using the appropriate (conjugate of the) face-pairing. If at any stage overlapping of interiors occurs, or if the boundaries intersect, but not in a common face, or if a face of the new copy coincides with some existing free face, but not according to one of the given face-pairings, then the process fails. The process succeeds if we end with a locally finite tessellation of the whole of $\mathbf{X}^n$. The process might continue indefinitely without failure at any finite stage, for example covering a proper subspace of $\mathbf{X}^n$, and it will have failed if at the end it does not give a locally finite tessellation of all of $\mathbf{X}^n$.

We now describe some more conditions which arise in considering Poincaré's Theorem. Suppose Pairing$(\mathscr{P}, R, A)$. Let $(F_1, P_1) \in \mathscr{F}(\mathscr{P})$ and let C_1 be a codimension-one face of F_1. Let F_1' be the other codimension-one face of P_1 containing C_1 (see Lemma 2.7). Let $R(F_1', P_1) = (F_2, P_2)$ and let $g_1 = A(F_1', P_1)$. Note that $C_2 = g_1(C_1)$ is a codimension-one face of F_2, so it is a codimension-two face of P_2, and hence there exists only one other codimension-one face of P_2 containing C_2. We call this face F_2'. Set $g_2 = A(F_2', P_2)$ and continue in the same way, obtaining a sequence $\{\sigma_i = (P_i, C_i, F_i, F_i', g_i)\}_{i=1,2,\ldots}$. We have $g_i = A(F_i', P_i)$ and $g_{i-1} \circ \ldots \circ g_1(C_1) = C_i$. The sequence is determined once one has chosen P_1, F_1 and C_1.

CONDITION 3.7 (FirstCyclic). FirstCyclic$(\mathscr{P}, R, A)$ is the condition that, for each $(F_1, P_1) \in \mathscr{F}(\mathscr{P})$ and for each codimension-one face C_1 of F_1, there is some $r \geqslant 1$ such that $\sigma_{r+1} = \sigma_1$. The minimal $r \geqslant 1$ with this property is called the *first cycle length* of (C_1, F_1, P_1).

REMARK 3.8.

(a) The condition $\sigma_{r+1} = \sigma_1$ is obviously equivalent to the conditions $P_{r+1} = P_1$, $C_{r+1} = C_1$ and $F_{r+1} = F_1$.

(b) Instead of starting with P_1, F_1 and C_1, we could instead start with P_i, F_i and C_i, or with P_i, F_i' and C_i. Instead of getting the r-tuple

$(\sigma_1, \ldots, \sigma_r)$, we would get a cyclic permutation of it, or a cyclic permutation of $(\sigma_r', \ldots, \sigma_1')$, where $\sigma_i' = (P_i, C_i, F_i', F_i, g_{i-1}^{-1})$ for $1 \leqslant i \leqslant r$ and the indices are interpreted mod r.

(c) FirstCyclic($\mathscr{P}$, R, A) clearly has to be satisfied if our basic objective is to be achieved (see Remark 3.6). Note however that complications arise if we do not insist on local finiteness in the definition of a tessellation, when formulating our basic objective. For example, in $\mathbf{E}^2$, we could glue together a countable number of wedges, such that the sum of the wedge angles is 2π. Such a construction would not give the whole of $\mathbf{E}^2$, but would leave a single ray uncovered: is this a tessellation? The meaning of the word "tessellation" does not suffer from such ambiguities when one insists on local finiteness of the face structure.

CONDITION 3.9 (finite). Finite($\mathscr{P}$) is the condition that $\mathscr{P}$ is finite and that each element of $\mathscr{P}$ has only a finite number of faces. This is one of the usual conditions imposed for Poincaré's Theorem, but it is clearly not essential for our basic objective (see Remark 3.6). However, this condition is essential if one wishes to check all the conditions by a finite mechanical procedure.

Clearly, if Finite($\mathscr{P}$) then FirstCyclic($\mathscr{P}$, R, A).

CONDITION 3.10 (SecondCyclic). SecondCyclic($\mathscr{P}$, R, A) is the condition that for each $(F_1, P_1) \in \mathscr{F}(\mathscr{P})$ and for each codimension-one face C_1 of F_1, there exists $r \geqslant 1$ such that $\sigma_{r+1} = \sigma_1$ and the restriction of $g_r \circ \ldots \circ g_1$ to C_1 is the identity. The minimal $r \geqslant 1$ with this property is called the *second cycle length* of (C_1, F_1, P_1). Even if FirstCyclic($\mathscr{P}$, R, A), the second cycle length may be infinite (see Example 3.32 or Example 3.17).

The reader is referred to Examples 3.15, 3.16, 3.17 and 3.18, which may provide a better understanding of the cycle conditions.

REMARK 3.11.

(a) According to Remark 3.1(a), we need to check that our condition is not changed by the replacement of one of the convex cells $P \in \mathscr{P}$ by $\psi(P)$ for some $\psi \in \mathrm{Isom}(\mathbf{X}^n)$. In fact, suppose $(F', P') \in \mathscr{F}(\mathscr{P})$ and $R(F', P') = (F'', P'')$. Then $A(F', P')$ must be replaced by:

- $\psi \circ A(F', P') \circ \psi^{-1}$ if $P' = P$ and $P'' = P$;
- $A(F', P') \circ \psi^{-1}$ if $P' = P$ and $P'' \neq P$;
- $\psi \circ A(F', P')$ if $P' \neq P$ and $P'' = P$;
- $A(F', P')$ if $P' \neq P$ and $P'' \neq P$.

It follows that the mapping $g = g_r \circ \ldots \circ g_1$ to which SecondCyclic$(\mathscr{P}, R, A)$ refers is either unchanged or is replaced by $\psi \circ g \circ \psi^{-1}$ under the replacement of P by $\psi(P)$, so the condition is well-defined.

(b) Just as in the case of first cycles, the second cycle length will be the same for each (F_i, P_i) and (F'_i, P_i) which occurs in the cycle. The mapping g of Remark 3.11(a) has to be replaced by g^{-1} when (F_1, P_1) is replaced by (F'_1, P_1). As we have seen in Remark 3.11(a), g is only defined in an intrinsic way up to conjugation, because each of the convex cells is only defined up to isometry. If we start with (F_i, P_i) instead of (F_1, P_1), then g once again changes by a conjugation.

(c) FirstCyclic$(\mathscr{P}, R, A)$ and SecondCyclic$(\mathscr{P}, R, A)$ clearly have to be satisfied if our basic objective is to be achieved (see Remark 3.6).

LEMMA 3.12 (cycles and rotations). *We use the notation introduced above and assume that* $r \geqslant 1$ *is the second cycle length of* (C_1, F_1, P_1). *Let* θ_i *be the the dihedral angle of* P_i *along* C_i *for* $i = 1, \ldots, r$. *Then the isometry* $g = g_r \circ \ldots \circ g_1$ *of* $\mathbf{X}^n$ *is a rotation through an angle* $\Sigma\theta_i$ *around the codimension-two subspace of* $\mathbf{X}^n$ *containing* C_1.

Proof of 3.12. We denote by S the codimension-two subspace containing C_1. Note that g is necessarily the identity on S, since C_1 has non-empty S-interior and $g \mid C_1$ is the identity by hypothesis. We only need to prove that g preserves the orientation.

Consider in $\mathbf{X}^n$ the convex cells $P_1, g_1^{-1}(P_2), \ldots, (g_1^{-1} \circ \ldots \circ g_r^{-1})(P_{r+1})$; they have a common codimension-two face

$$C_1 = g_1^{-1}(C_2) = \ldots = (g_1^{-1} \circ \ \ldots \circ g_r^{-1})(C_{r+1}) .$$

Moreover, according to Condition 3.2(b), $(g_1^{-1} \circ \ldots \circ g_{i-1}^{-1})(P_i)$ and $(g_1^{-1} \circ \ldots \circ g_i^{-1})(P_{i+1})$ lie on opposite sides of the common codimension-one face $(g_1^{-1} \circ \ldots \circ g_{i-1}^{-1})(F'_i) = (g_1^{-1} \circ \ldots \circ g_i^{-1})(F_{i+1})$. Fix an orientation for two-dimensional subspaces normal to S. If we assume that the angle from F_1 to F'_1 is positive, then the angle from $F'_1 = g_1^{-1}(F_2)$ to $g_1^{-1}(F'_2)$ is also positive. By induction the angle from $(g_1^{-1} \circ \ldots \circ g_r^{-1})(F_{r+1})$ to $(g_1^{-1} \circ \ldots \circ g_r^{-1})(F'_{r+1})$ is positive. But $F_{r+1} = F_1$ and $F'_{r+1} = F'_1$, and hence $g_1^{-1} \circ \ldots \circ g_r^{-1}$ preserves the orientation, as required. □

CONDITION 3.13 (ThirdCyclic). ThirdCyclic$(\mathscr{P}, R, A)$ is the condition that for all $(F_1, P_1) \in \mathscr{F}(\mathscr{P})$ and for all codimension-one faces C_1 of F_1, if $r \geqslant 1$ is the second cycle length, the mapping g described in Lemma 3.12 is a rotation through an angle of the form $2\pi/m$ for some non-zero $m \in \mathbf{Z}$.

REMARK 3.14. It follows from Remark 3.11(a) that this condition and the absolute value of m are both independent of $i (1 \leqslant i \leqslant r)$ and of whether one starts with (F_i, P_i) or (F_i', P_i). The condition is necessary if our basic objective (see Remark 3.6) is to be achieved. However, we have to proceed carefully, as the following example shows. We take a wedge in $\mathbf{E}^2$, with angle $2\pi/3$. If the face-pairings are reflections, then the sum of angles which occurs in Condition 3.13 is $4\pi/3$, which is not of the required form. Note that the images of the wedge do tile $\mathbf{E}^2$. However, this tessellation is not consistent with the face-pairings (see Figure 4).

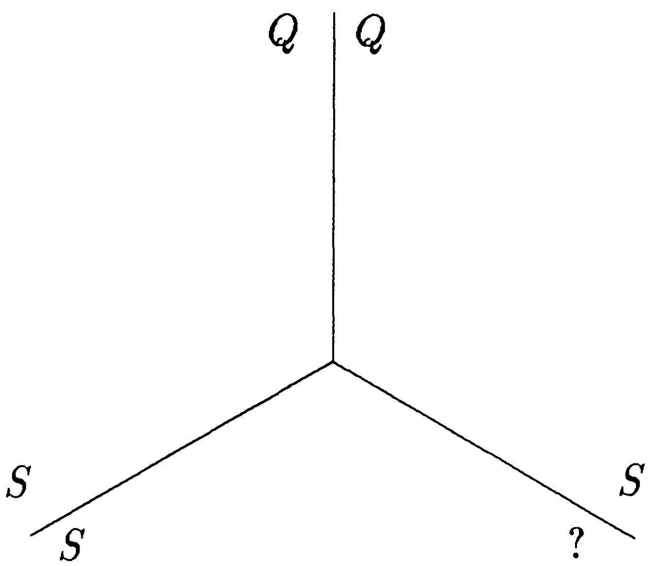

FIGURE 4.
Reflection in the sides of a wedge.
The different images seem to tessellate.
But if we take the face-pairings into account we find an inconsistency.

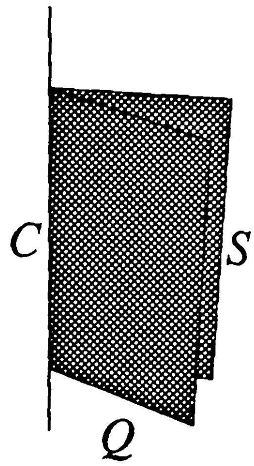

FIGURE 5.
Dihedral region.
This shows a dihedral region in $\mathbf{E}^3$, which is the only member of $\mathscr{P}$ in Examples 3.15, 3.16, 3.17 and 3.18.

EXAMPLE 3.15 (cyclic example 1). In $\mathbf{E}^3$ let P be the dihedral region with angle φ shown in Figure 5, and let $\mathscr{P} = \{P\}$. Let the codimension-one faces of P be Q and S, intersecting in the codimension-two face C. We set $R(Q, P) = (Q, P)$ and $R(S, P) = (S, P)$ and we define $A(Q, P)$ (respectively $A(S, P)$) to be the reflection in the plane containing Q (respectively S). Pairing($\mathscr{P}, R, A$) follows. Moreover, as illustrated in Figure 6, both first and second cycle lengths are equal to two. Then (by Lemma 3.12), ThirdCyclic($\mathscr{P}, R, A$) is equivalent to the condition $\varphi = \pi/m$ for some non-zero $m \in \mathbf{Z}$.

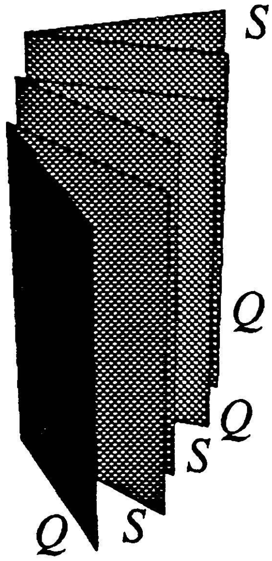

FIGURE 6.
Reflection face-pairings.
This illustrates Example 3.15. The first two cyclic conditions hold with $r = 2$.

EXAMPLE 3.16 (cyclic example 2). Let $\mathscr{P}$ be as in Example 3.15, set $R(Q, P) = (S, P)$ and define $A(Q, P)$ as the rotation through an angle φ around C; Pairing($\mathscr{P}, R, A$) is of course satisfied and FirstCyclic($\mathscr{P}, R, A$), SecondCyclic($\mathscr{P}, R, A$) both hold with $r = 1$ (see Figure 7). Hence, using

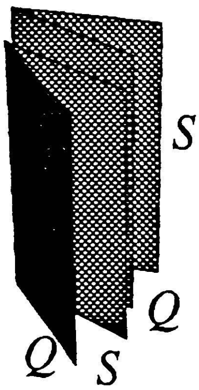

FIGURE 7.
Rotation face-pairing.
This illustrates Example 3.16. The first two cyclic conditions hold with $r = 1$.

Lemma 3.12, we see that ThirdCyclic($\mathcal{P}$, R, A) is equivalent to the condition that $\varphi = 2\pi/m$ for some non-zero $m \in \mathbf{Z}$.

EXAMPLE 3.17 (cyclic example 3). Let $\mathcal{P}$ and R be as in Example 3.16 and define $A(Q, P)$ as the composition of the rotation through an angle φ around C with a non-zero translation parallel to C. Then Pairing($\mathcal{P}$, R, A) is satisfied, FirstCyclic($\mathcal{P}$, R, A) is satisfied with $r = 1$ but SecondCyclic($\mathcal{P}$, R, A) is not satisfied.

EXAMPLE 3.18 (cyclic example 4). Let $\mathcal{P}$ and R be as in Example 3.16 and define $A(Q, P)$ as the composition of the rotation through an angle φ around C with the reflection in a plane orthogonal to C; Pairing($\mathcal{P}$, R, A) is satisfied. As shown in Figure 8, FirstCyclic($\mathcal{P}$, R, A) is satisfied with $r = 1$ (and hence for all $r \geqslant 1$), while SecondCyclic($\mathcal{P}$, R, A) is satisfied with $r = 2$. As in Example 3.15, ThirdCyclic($\mathcal{P}$, R, A) is equivalent to the condition that $\varphi = \pi/m$ for some non-zero $m \in \mathbf{Z}$.

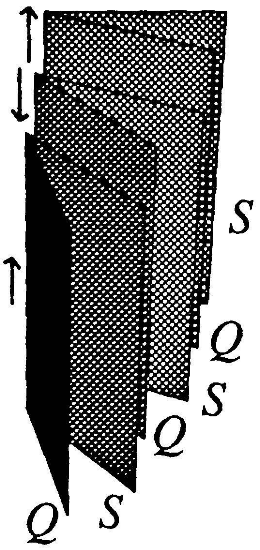

FIGURE 8.
Rotation plus flip face-pairing.
This illustrates Example 3.18.
The first cyclic condition holds with $r = 1$ and the second one with $r = 2$.

CONDITION 3.19 (Cyclic). Cyclic($\mathcal{P}$, R, A) is the conjunction of FirstCyclic($\mathcal{P}$, R, A), SecondCyclic($\mathcal{P}$, R, A) and ThirdCyclic($\mathcal{P}$, R, A).

We now introduce two more conditions, each of which involves the metric structure of the elements of $\mathcal{P}$.

CONDITION 3.20 (FirstMetric). FirstMetric($\mathcal{P}$) is the condition that there should exist a number $\varepsilon > 0$ such that for all elements P of $\mathcal{P}$ and for all faces E_1, E_2 of P, if $E_1 \cap E_2 = \varnothing$ then $d(E_1, E_2) \geqslant \varepsilon$ (where d denotes the usual distance between subsets of a metric space).

EXAMPLE 3.21 (not FirstMetric). FirstMetric($\mathscr{P}$) is not necessary for our basic objective to be achieved (see Remark 3.6). For example, take any tessellation of the euclidean plane by triangles. We can insert small triangles around the vertices, making the size of the inserted triangles tend to zero as one goes to infinity, as in Figure 9.

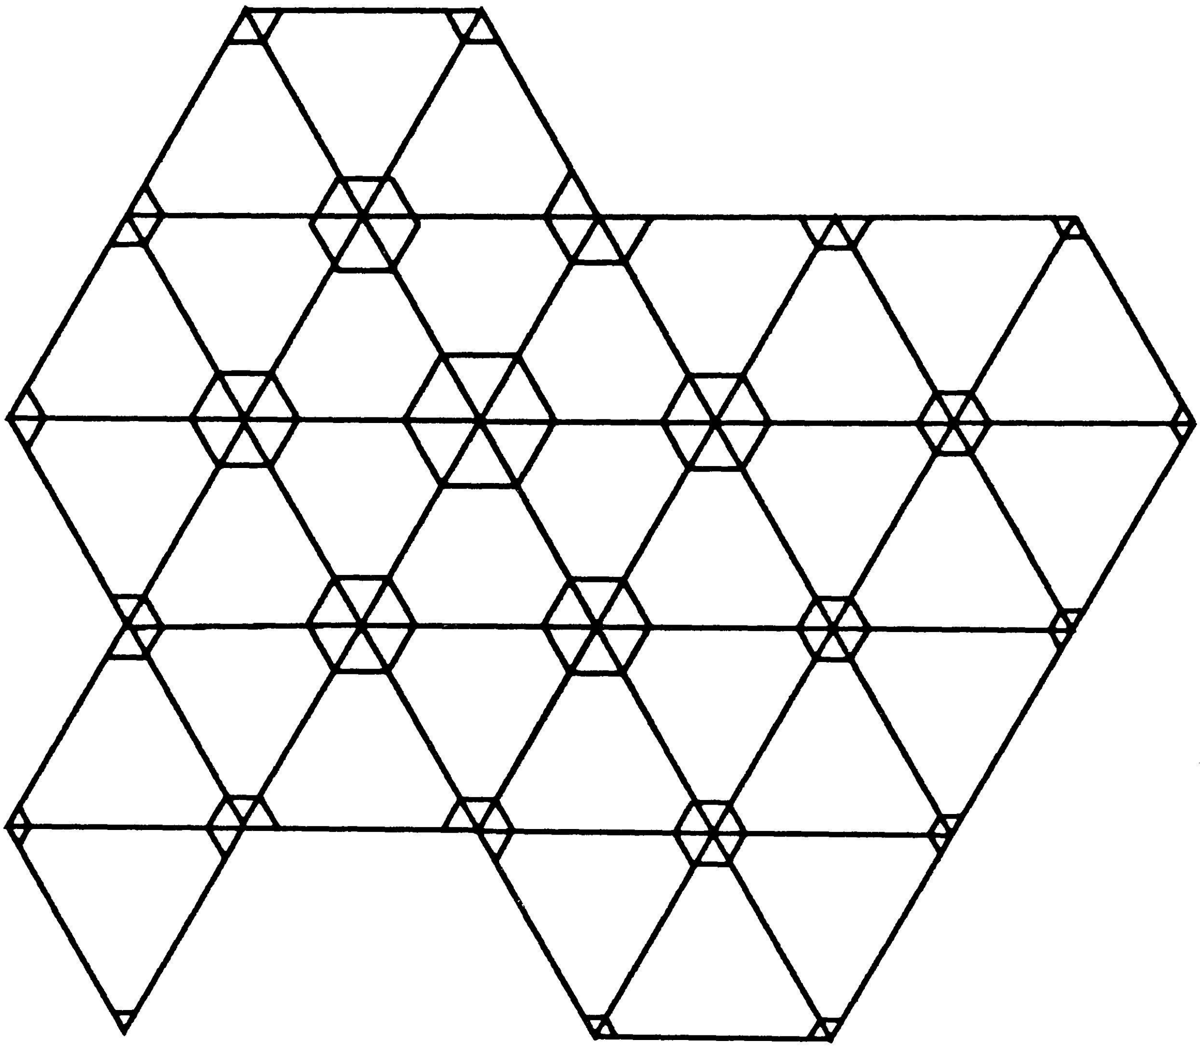

FIGURE 9.
Tessellation of $\mathbf{E}^2$. This illustrates Example 3.21.

CONDITION 3.22 (SecondMetric). SecondMetric($\mathscr{P}$) is the condition that, given any $\delta > 0$, there should exist $\mu(\delta) > 0$ with the following property. Suppose $P \in \mathscr{P}$ and E and F are faces of P such that $E \cap F \neq \varnothing$ and $E \not\subset F$. If x is a point of E at distance at least δ from ∂E, then $d(x, F) \geqslant \mu(\delta)$.

CONDITION 3.23 (Metric). Metric($\mathscr{P}$) is the conjunction of FirstMetric($\mathscr{P}$) and SecondMetric($\mathscr{P}$).

Condition 3.22 (SecondMetric) may appear to be strictly stronger than Condition 3.20 (FirstMetric), but it is not. For example, in $\mathbf{H}^3$, take P to be the intersection of four half-spaces, whose boundaries meet at a point z at infinity. We arrange for the intersection of P with a horosphere centred at p to be a square. Then SecondMetric($\mathscr{P}$) is satisfied but not FirstMetric($\mathscr{P}$).

REMARK 3.24 (Finite implies Metric). If Finite($\mathscr{P}$) holds then FirstMetric($\mathscr{P}$) is equivalent to the condition that any pair of disjoint faces of the elements of $\mathscr{P}$ have positive distance from each other. This is true in euclidean and spherical geometry but not necessarily true in hyperbolic geometry. For example, in the hyperbolic plane we take $\mathscr{P} = \{P\}$, where P is the region between two disjoint geodesics. If the geodesics meet at infinity then FirstMetric($\mathscr{P}$) is false. From Proposition 2.11 we see that Finite($\mathscr{P}$) implies FirstMetric($\mathscr{P}$), unless $\mathbf{X}^n = \mathbf{H}^n$. From Lemma 2.13 we see that Finite($\mathscr{P}$) implies SecondMetric($\mathscr{P}$) for all three geometries. Hence Finite($\mathscr{P}$) implies Metric($\mathscr{P}$) unless $\mathbf{X}^n = \mathbf{H}^n$.

SecondMetric($\mathscr{P}$) should be thought of as showing that the angle between faces is bounded below.

EXAMPLE 3.25. To make an example where FirstMetric($\mathscr{P}$) is satisfied, but not SecondMetric($\mathscr{P}$), we take a sequence of disjoint isoceles triangles T_i in $\mathbf{E}^2$, tending to infinity. T_i is chosen so that the apex angle tends to zero and the base of T_i always has length one, which means that the two equal sides have length tending to infinity. We can then complete this to a triangulation of $\mathbf{E}^2$ in which FirstMetric($\mathscr{P}$) is satisfied. SecondMetric($\mathscr{P}$) clearly fails.

Given a set S in $\mathbf{H}^n$ we denote by $\bar{S}$ the closure of S as a subset of $\bar{\mathbf{H}}^n$, and we refer to the points of $\bar{S} \cap \partial\mathbf{H}^n$ as the *points at infinity of S*.

LEMMA 3.26. *Let P be a convex cell in $\mathbf{H}^n$ with finitely many faces. Two disjoint faces of P can have at most one common point at infinity, and they are a positive distance apart if and only if they have no common point at infinity.*

Proof of 3.26. Let A and B be the two disjoint faces. If they have two common points at infinity, the geodesic joining them lies in both faces, contradicting the hypothesis that they are disjoint in $\mathbf{H}^n$. If A and B have a common point at infinity, then they are clearly zero distance apart. If, conversely, they are zero distance apart, then there are sequences $\{a_i\}$ in A and $\{b_i\}$ in B, such that $d(a_i, b_i)$ converges to zero. We may assume that the two sequences converge to the same point p at infinity. Then $p \in \bar{A} \cap \bar{B}$ as required. $\square$

EXAMPLE 3.27. Consider the polygon $P \subset \mathbf{H}^2$ given in the upper half-plane model by $[1, 2] \times (0, \infty)$. The two faces of P have a common point at infinity, and they are zero distance apart. Multiplication by two induces a face-pairing which satisfies Pairing($\mathscr{P}, R, A$) and Cyclic($\mathscr{P}, R, A$). However the images of P under the powers of the pairing cover only the right half of the half-plane.

DEFINITION 3.28 (codimension-i graph). For $i \geqslant 1$ we define $\mathscr{F}^i(\mathscr{P})$ to be the set of all pairs (E, P) where $P \in \mathscr{P}$ and E is a codimension-i face of P. So $\mathscr{F}^1(\mathscr{P}) = \mathscr{F}(\mathscr{P})$. Given a face-pairing (R, A) we define a graph $\Gamma^i(\mathscr{P}, R, A)$ which has a vertex for each element (E, P) of $\mathscr{F}^i(\mathscr{P})$ and an edge $e(E, F, P)$ for each triple with $E \subset F \subset P$, E a codimension-i face of P and F a codimension-one face of P. The edge $e(E, F, P)$ joins (E, P) to (E', P') if $R(F, P) = (F', P')$ and $E' = A(F, P)(E)$; we regard $e(E, F, P)$ as being the same edge as $e(E', F', P')$. Each component of $\Gamma^1(\mathscr{P}, R, A)$ consists of one or two vertices and one edge. FirstCyclic($\mathscr{P}, R, A$) is equivalent to the condition that each component of $\Gamma^2(\mathscr{P}, R, A)$ is finite.

CONDITION 3.29 (LocallyFinite). We now describe a condition which is clearly necessary for our basic objective (see Remark 3.6). In many situations, this condition does not need to be explicitly verified, since it follows from various subsets of the other conditions. LocallyFinite($\mathscr{P}, R, A$) is the condition that each component of $\Gamma^i(\mathscr{P}, R, A)$ is a finite graph. Clearly Finite($\mathscr{P}$) implies LocallyFinite($\mathscr{P}, R, A$).

If $n = 2$, LocallyFinite($\mathscr{P}, R, A$) is equivalent to FirstCyclic($\mathscr{P}, R, A$).

EXAMPLE 3.30 (not LocallyFinite). Pairing($\mathscr{P}, R, A$), Connected($\mathscr{P}, R$) and Cyclic($\mathscr{P}, R, A$) do not imply LocallyFinite($\mathscr{P}, R, A$). An example may be constructed as follows. For each integer $n > 0$, take the two-sphere of radius $1/n$ in $\mathbf{R}^3$ lying above the plane $z = 0$ and tangent to it at 0. These spheres cut $\mathbf{R}^3$ into a countable number of pieces. We can approximate each piece by a finite union of convex polyhedra, so that everything fits together in the same qualitative fashion as the spheres we have described. (We first approximate the spherical surfaces, and then cut up the regions.) In particular the origin appears as a point in each of the approximations. The result is not locally finite at the origin, though the other hypotheses are satisfied. Note that, with the obvious path metric induced by gluing the pieces together, the resulting space is a complete metric space; so completeness does not help, in this type of situation, in deducing local finiteness

REMARK 3.31 (stronger local finiteness). There is an alternative version of the local finiteness condition, used for example in [Mas71]: recall from Remark 3.6 that $Q(\mathscr{P}, R, A)$ is the quotient space of $\bigsqcup_{P \in \mathscr{P}} P$, the disjoint union of the convex cells in $\mathscr{P}$. We might assume that the inverse image under the quotient map of any point in $Q(\mathscr{P}, R, A)$ is finite. This obviously implies LocallyFinite$(\mathscr{P}, R, A)$. It will turn out that LocallyFinite$(\mathscr{P}, R, A)$ together with Cyclic$(\mathscr{P}, R, A)$ implies this stronger condition (see Theorem 4.14).

EXAMPLE 3.32 (irrational). Here is an example when the weaker condition of local finiteness is true, but not the stronger condition. Of course, Cyclic$(\mathscr{P}, R, A)$ is not true in this case. We take two codimension-one spherical subspaces of $\mathbf{S}^3$. These meet along a common $\mathbf{S}^1$. Let P be one of the four complementary three-dimensional regions, and let $\mathscr{P} = \{P\}$. Then P has two faces, each of which is a hemisphere. Suppose we glue one of these hemispheres to the other, inducing an irrational rotation on the common circle boundary. Then we have LocallyFinite$(\mathscr{P}, R, A)$ and Finite$(\mathscr{P})$, but the strong version of local finiteness just stated is false.

Another similar example in $\mathbf{H}^4$ is given as follows. Take the intersection of two half-spaces, such that the boundaries of these half-spaces intersect in a hyperbolic plane. There are two codimension-one faces F_1 and F_2, each of which is half of a three-dimensional hyperbolic space, and one codimension-two subspace S, which is a hyperbolic plane. We take as a face-pairing a rotation keeping the codimension-two face S pointwise fixed and taking F_1 to F_2, followed by an isometry T of $\mathbf{H}^4$. T sends S to itself and is elliptic, rotating S through an irrational angle. If we take $\mathbf{H}^4$ to be embedded as one sheet of the hyperboloid $\langle v, v \rangle = -1$ in a five-dimensional vector space with indefinite inner product of type (4, 1), then T is the identity on $S^{\perp}$. Cyclic$(\mathscr{P}, R, A)$ is false, LocallyFinite$(\mathscr{P}, R, A)$ and Finite$(\mathscr{P})$ are true, but the quotient space Q is not hausdorff.

4. DEVELOPING MAPS

As in the previous section, let $\mathscr{P}$ be a set of thick convex cells in $\mathbf{X}^n$, and let (R, A) satisfy Pairing$(\mathscr{P}, R, A)$. We define a graph $\Gamma(\mathscr{P}, R)$ in the following way. The vertices of the graph are the elements of $\mathscr{P}$. We have an edge, which we call either $e(F, P)$ or $e(F', P')$, joining P and P' if and only if $R(F, P) = (F', P')$. So there is one edge for each face-pairing. Clearly, Connected$(\mathscr{P}, R)$ if and only if $\Gamma(\mathscr{P}, R)$ is connected.

Now let T be a maximal tree in $\Gamma(\mathscr{P}, R)$. Consider the equivalence relation $\sim$ on the disjoint union $\bigsqcup_{P \in \mathscr{P}} P$ of the elements of $\mathscr{P}$, generated by $x \sim A(F, P)(x)$ if $(F, P) \in \mathscr{F}(\mathscr{P})$, $e(F, P) \subset T$ and $x \in F$. We define the space $Y(\mathscr{P}, R, A, T)$ and the quotient map $\pi_Y : \bigsqcup_{P \in \mathscr{P}} P \to Y(\mathscr{P}, R, A, T)$ by identifying each equivalence class to a point. We have Connected$(\mathscr{P}, R)$ if and only if $Y(\mathscr{P}, R, A, T)$ is (arcwise) connected. Since in T no edge is a loop, all elements of $\mathscr{P}$ are naturally embedded in $Y(\mathscr{P}, R, A, T)$ — that is, the restriction to any component of the domain of the projection $\bigsqcup_{P \in \mathscr{P}} P \to Y(\mathscr{P}, R, A, T)$ is injective. It is straightforward to see that $Y(\mathscr{P}, R, A, T)$ is contractible if it is connected — a deformation retraction to a point can be constructed inductively, cell by cell, working along the edges of T.

For the rest of this section we will assume that Pairing$(\mathscr{P}, R, A)$, Connected$(\mathscr{P}, R)$ and Cyclic$(\mathscr{P}, R, A)$ are satisfied.

The following lemma is easy to prove.

LEMMA 4.1 (developing Y). *For any choice of* $P_0 \in \mathscr{P}$ *there exists a unique mapping* $D_Y : Y(\mathscr{P}, R, A, T) \to \mathbf{X}^n$, *which we call the developing map associated to* $(\mathscr{P}, R, A, T)$, *with the following properties:*

- $D_Y|_{P_0}$ *is the identity;*
- *for each* $P \in \mathscr{P}$, $D_Y|_P$ *is the restriction of an isometry of* $\mathbf{X}^n$ *(which we denote by* ψ_P*);*
- *if* $(F, P) \in \mathscr{F}(\mathscr{P})$ *and* $e(F, P) \subset T$ *joins* P *to* P', *then* $\psi_{P'} A(F, P) = \psi_P$.

A different choice of the initial convex cell P_0 *or a different choice of the way it is embedded in* $\mathbf{X}^n$ *leads to the mapping* $\psi \circ D_Y$ *for some* $\psi \in Isom(\mathbf{X}^n)$.

Changing the positions of the convex cells $P \in \mathscr{P}$ *(see Remark 3.1(a)), we may take each* ψ_P *to be the identity and then* $A(F, P)$ *is the identity for each edge* $e(F, P)$ *in* T.

From now on, we will assume that ψ_P is the identity for each $P \in \mathscr{P}$.

DEFINITION 4.2. We define an abstract group $G(\mathscr{P}, R, A, T)$ as the group generated by the set of symbols:

$$\{\alpha_{(F,P)} : (F, P) \in \mathscr{F}(\mathscr{P})\}$$

subject to the following relations:

- $\alpha_{(F,P)} = \mathrm{id}$ if $e(F, P) \subset T$.

- if $(F, P) \in \mathcal{F}(\mathcal{P})$, $e(F, P) \not\subset T$ and $R(F, P) = (F', P')$ then we have the relation $\alpha_{(F,P)}\alpha_{(F',P')} = \mathrm{id}$. In particular, if $R(F, P) = (F, P)$ then $\alpha_{(F,P)}$ has order two.
- for each $P_1 \in \mathcal{P}$ and for each codimension-two face C_1 of P_1, in the notation of Conditions 3.7 and 3.13, we have the relation

$$(\alpha_{(F_r', P_r)} \cdots \alpha_{(F_1', P_1)})^m = \mathrm{id} .$$

REMARK 4.3. According to Remark 3.14, given $P_1 \in \mathcal{P}$ and a codimension-two face C_1 of P_1, we obtain an equivalent relation starting from either of the codimension-one faces of P_1 containing C_1, or from any of the faces F_i or F_i'.

LEMMA 4.4 (holonomy). *We assume Pairing*$(\mathcal{P}, R, A)$, *Connected*$(\mathcal{P}, R)$ *and Cyclic*$(\mathcal{P}, R, A)$. *For any choice of developing map* D_Y *associated to* $(\mathcal{P}, R, A)$, *there exists a unique homomorphism* $h: G(\mathcal{P}, R, A, T) \to Isom(\mathbf{X}^n)$ *with the following property: if* $(F, P) \in \mathcal{F}(\mathcal{P})$ *then* $h(\alpha_{(F,P)}) = A(F, P)$. *A different choice of* D_Y *leads to the homomorphism* $g \mapsto \psi h(g)\psi^{-1}$ *for some* $\psi \in Isom(\mathbf{X}^n)$.

Proof of 4.4. Given D_Y, the position in $\mathbf{X}^n$ of each $P \in \mathcal{P}$ is determined. For each $(F, P) \in \mathcal{F}(\mathcal{P})$, the face-pairing $A(F, P)$ is then also determined. We define $h(\alpha_{(F,P)}) = A(F, P)$. According to Pairing$(\mathcal{P}, R, A)$ and to Lemma 4.1 the relations defining G starting from the generators $\alpha_{(F,P)}$ hold for the corresponding $A(F, P)$'s in Isom$(\mathbf{X}^n)$, and then h can be extended to a homomorphism of the whole of $G(\mathcal{P}, R, A, T)$. Uniqueness is obvious. The last assertion is readily deduced from Lemma 4.1. □

DEFINITION 4.5. We abbreviate as follows: $Y = Y(\mathcal{P}, R, A, T)$ and $G = G(\mathcal{P}, R, A, T)$. We give G the discrete topology, and consider the space $G \times Y$ with the product topology. We consider on $G \times Y$ the equivalence relation $\sim$ generated by: $(g\alpha_{(F,P)}, x) \sim (g, A(F, P)(x))$ whenever $g \in G, (F, P) \in \mathcal{F}(\mathcal{P})$ and $x \in F \subset P \hookrightarrow Y$. We will denote by $Z = Z(\mathcal{P}, R, A, T)$ the quotient space of $G \times Y$ by this equivalence relation, and by $\pi_Z: G \times Y \to Z$ the quotient map.

REMARK 4.6 (Y not subset Z). It is false in general that the restriction to $\{\mathrm{id}\} \times Y$ of the projection $G \times Y \to Z$ is injective — see Example 4.12.

G acts on Z in an obvious way, and G acts on $\mathbf{X}^n$ via the homomorphism h.

LEMMA 4.7 (developing Z). *For any choice of* P_0 *in* $\mathcal{P}$, *there exists a unique G-equivariant mapping* $D_Z : Z \to \mathbf{X}^n$ *such that the element of* Z *represented by* $(g, y) \in G \times Y$ *is mapped to* $h(g)(D_Y(y))$, *where* D_Y *and* h *are given respectively by Lemma 4.1 and Lemma 4.4. A different choice of the initial convex cell and its position in* $\mathbf{X}^n$ *leads to the mapping* $\psi \circ D_Z$ *for some* $\psi \in Isom(\mathbf{X}^n)$.

Proof of 4.7. We only have to check that if $(g, y) \sim (g', y')$ in $G \times Y$ then $h(g)(D_Y(y)) = h(g')(D_Y(y'))$, and this is readily deduced from the definition of $\sim$ and from the definitions of D_Y and h. $\square$

COROLLARY 4.8 (P embeds in Z). *For each* $P \in \mathcal{P}$ *and* $g \in G$ *the mapping*

$$P \ni x \mapsto \pi_Z(g, x) \in Z$$

is injective.

REMARK 4.9 (Z independent of T). The definition of Z given above depends on T. However, this dependence is not real. To see this we define $\mathcal{G}$ to be a groupoid (a small category in which every morphism has a two-sided inverse). We take $\mathcal{P}$ to be the set of objects of $\mathcal{G}$. We take $\alpha_{(F,P)}$ to be a morphism from P to P', where $R(F, P) = (F', P')$. In general, the morphisms are formed from compositions of these, subject to the same relations as those used in the definition of G above, except that we now take $T = \varnothing$. We give the set M of morphisms of $\mathcal{G}$ the discrete topology, and we take the obvious topology on $\bigsqcup_{P \in \mathcal{P}} P$. To define Z, we fix $P_0 \in \mathcal{P}$, and let $M(P_0)$ be the set of morphisms with range P_0. We then take all pairs $(g, x) \in M(P_0) \times \bigsqcup_{P \in \mathcal{P}} P$, where $x \in P$, $P \in \mathcal{P}$ and $g : P \to P_0$. We identify $(g\alpha_{(F,P)}, x)$ with $(g, A(F, P)x)$, provided $x \in F \subset P$ and $g : P' \to P_0$, where $R(F, P) = (F', P')$. We define Z to be the identification space just defined. If we change P_0, the resulting Z is unchanged. An isomorphism between the two versions of Z is given by choosing a word in the $\alpha_{(F,P)}$'s relating the choices. The isomorphism is therefore determined up to the action of an element of M.

The only reason for using the definition given previously, in terms of a group, rather than that given now, is that the concept of a group is more familiar than the concept of a groupoid. The construction of the group G from the groupoid $\mathcal{G}$ is the standard construction of a group from a connected groupoid. We are therefore justified in writing $Z(\mathcal{P}, R, A)$ instead of $Z(\mathcal{P}, R, A, T)$, if the occasion demands.

REMARK 4.10 (cell structure of Z). For each $g \in M(P_0)$ and each face E of P, we obtain the subset $g(E) = \pi_Z(\{g\} \times E)$ of Z. To see that $g(E)$ is an isomorphic copy of E, we apply the developing map D_Z. So $g(E)$ is a convex cell of the same dimension as E. Since the identifications respect the face structure (see Condition 3.2(c)), we see that Z is the disjoint union of the relative interiors of these convex cells of various dimensions. Of course, g and E are not determined by the cell; $g(E)$ is just one representation. The left action of G preserves the cell structure of Z. If x and y are interior points of the same top-dimensional cell of Z and if $gx = y$ for some $g \in G$, then $x = y$ and g is the identity element.

It is easy to see that Connected($\mathscr{P}$, R) is equivalent to Z being (arcwise) connected.

DEFINITION 4.11 (boundary and interior of Y). We write $Y = Y(\mathscr{P}, R, A, T)$. The *boundary of* Y, denoted ∂Y, is the union of the faces F such that $(F, P) \in \mathscr{F}(\mathscr{P})$ and $e(F, P) \not\subset T$. The *interior of* Y is the complement of the boundary.

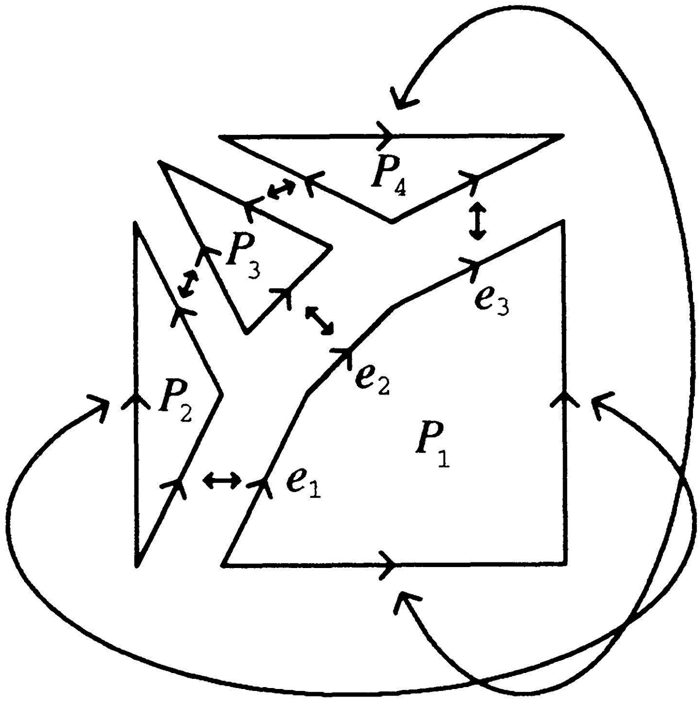

FIGURE 10.
Face pairings.
A set of polyhedra in Euclidean two-space, and a description of their face-pairings.

EXAMPLE 4.12 (fundamental domain not embedded). Let $\mathscr{P}$ be the set of polyhedra in $\mathbf{E}^2$ shown in Figure 10, and let the face-pairing (R, A) be defined by the arrows in the picture, in such a way that the orientation of the edges is preserved. All the conditions described in Section 3 hold for $(\mathscr{P}, R, A)$. It is evident from Figures 10 and 11 the developing map $D_Y\colon Y \to \mathbf{E}^2$ is not injective.

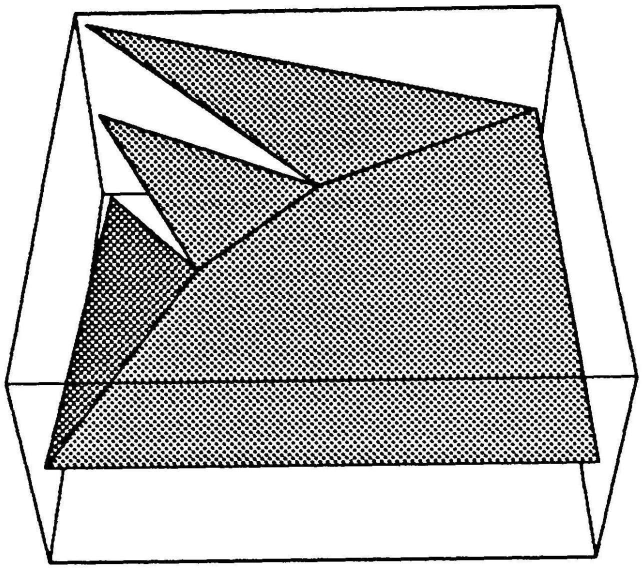

FIGURE 11.
The space Y.
We illustrate the space Y arising from Figure 10.

THEOREM 4.13 (modelled on $\mathbf{X}^n$). *Let $n \geqslant 2$, let $\mathscr{P}$ be a set of thick convex cells in $\mathbf{X}^n$ and let (R, A) be a face-pairing such that:*

(a) *Pairing$(\mathscr{P}, R, A)$;*

(b) *Connected$(\mathscr{P}, R)$;*

(c) *Cyclic$(\mathscr{P}, R, A)$;*

(d) *LocallyFinite$(\mathscr{P}, R, A)$ (recall from Condition 3.29 that this condition is automatically true if $n = 2$).*

Let T be a maximal tree in $\Gamma(\mathscr{P}, R)$, set $Y = Y(\mathscr{P}, R, A, T)$, $G = G(\mathscr{P}, R, A, T)$ and $Z = Z(\mathscr{P}, R, A, T)$, and let $D_Y\colon Y \to \mathbf{X}^n$, $h\colon G \to Isom(\mathbf{X}^n)$, $D_Z\colon Z \to \mathbf{X}^n$ be the developing maps as in Lemma 4.1, Lemma 4.4 and Lemma 4.7. Then Z is endowed with an $\mathbf{X}^n$-structure with respect to which $D_Z\colon Z \to \mathbf{X}^n$ is a local isometry. Also the convex cell structure of Z (see Remark 4.10) is locally finite. Furthermore the action of the group G on Z is proper discontinuous. Let p be a point in the interior of a top-dimensional cell P of Z. Then the stabilizer of p is trivial, and the orbit of p contains no other point of P.

This result will be proved by induction on n, assuming the following result in dimensions less than n. In Section 5, we will complete the induction by showing how Theorem 4.13 in dimension n implies Theorem 4.14 in dimension n.

THEOREM 4.14 (Poincaré's Theorem Version 1). *Let* $n \geqslant 2$, *let* $\mathscr{P}$ *be a set of thick convex cells in* $\mathbf{X}^n$ *and let* (R, A) *be a face-pairing such that:*

(a) *Pairing*$(\mathscr{P}, R, A)$;

(b) *Connected*$(\mathscr{P}, R)$;

(c) *Cyclic*$(\mathscr{P}, R, A)$;

(d) *Metric*$(\mathscr{P})$.

Let T *be a maximal tree in* $\Gamma(\mathscr{P}, R)$, *set*

$$Y = Y(\mathscr{P}, R, A, T), \quad G = G(\mathscr{P}, R, A, T), \quad Z = Z(\mathscr{P}, R, A, T)$$

and let $D_Y : Y \to \mathbf{X}^n$, $h : G \to Isom(\mathbf{X}^n)$, $D_Z : Z \to \mathbf{X}^n$ *be mappings as in 4.1, 4.4 and 4.7. Then the following conclusions hold:*

(e) *LocallyFinite*$(\mathscr{P}, R, A)$ *is true in its strong form (see Remark 3.31);*

(f) Z *is endowed with an* $\mathbf{X}^n$*-structure with respect to which* $D_Z : Z \to \mathbf{X}^n$ *is a (bijective) isometry;*

(g) $h : G \to Isom(\mathbf{X}^n)$ *is injective and its image is a discrete subgroup of* $Isom(\mathbf{X}^n)$;

(h) $D_Y : Y \to \mathbf{X}^n$ *is injective on the interior of* Y *(see Definition 4.11), so that* $D_Y(Y)$ *or its interior can be considered as a fundamental domain for the action of* $h(G)$ *on* $\mathbf{X}^n$, *depending on the precise definition of that concept;*

(j) *the convex cell structure of* Z *(see Remark 4.10) is locally finite.*

The hypotheses (and hence the conclusions) hold in particular if we add to conditions *Pairing*$(\mathscr{P}, R, A)$, *Connected*$(\mathscr{P}, R)$ *and* *Cyclic*$(\mathscr{P}, R, A)$ *either of the following additional conditions:*

(k) $\mathbf{X}^n = \mathbf{E}^n$ *or* $\mathbf{S}^n$ *and* *Finite*$(\mathscr{P})$;

(l) $\mathbf{X}^n = \mathbf{H}^n$, *Finite*$(\mathscr{P})$ *and* *FirstMetric*$(\mathscr{P})$;

Proof of 4.13. We will assume that Theorem 4.14 has been proved in dimensions less than n.

For $n = 2$, Theorem 4.13 is a consequence of Condition 3.19. To see this, note that each point in $\pi_Z^{-1}(z)$ lies in some C_j $(1 \leqslant j \leqslant r)$ of one particular cycle, where the notation comes from Condition 3.10. Let $m > 0$ be as in Condition 3.13. *A priori,* we do not know that there are m distinct copies in Z of each of the r dihedral regions at the various $C_j \subset P_j$, though we do know that there are no more than these mr regions around $z \in Z$, because of the way Z is constructed. The existence of D_Z shows that there are also no fewer than mr regions. This completes the proof for $n = 2$.

We now prove Theorem 4.13 for $n > 2$, assuming Theorem 4.14 in dimensions less than n. Let $\pi_Z(g, x) = z$, and let $x \in \mathrm{RelInt}(E)$, where E is a face of $P \in \mathscr{P}$ of codimension i. Since we are assuming LocallyFinite$(\mathscr{P}, R, A)$, we have a finite graph Γ_E which is the component of $\Gamma^i(\mathscr{P}, R, A)$ containing (E, P) as a vertex (see Definition 3.28). Each vertex of Γ_E is a pair of the form (E', P'), where E' is a codimension-i face of $P' \in \mathscr{P}$. The link (see Definition 2.14) of E' in P' is a convex cell in $\mathbf{S}^{n-1}$, which is well-defined up to isometry.

Let $\mathscr{P}_E$ be the finite collection of links arising from the finite set of vertices of Γ_E. These are convex polyhedra in $\mathbf{S}^{n-1}$, defined up to isometry. For each vertex (E', P') of Γ_E, we choose a point $u' \in \mathrm{RelInt}(E')$. We make no attempt to choose these points consistently — indeed, in general consistency of choice is not possible. The position of a link in $\mathbf{S}^{n-1}$ is determined by fixing an isomorphism between $\mathbf{R}^n$ and the tangent space at u'. The given face-pairing (R, A) induces a face-pairing (R_E, A_E) on $\mathscr{P}_E$ as follows. Suppose E' is a face of F', $R(F', P') = (F'', P'')$, and $A(F', P')$ is the corresponding face-pairing. Let $E'' = A(F', P')(E')$. Let $u' \in \mathrm{RelInt}(E')$ and $u'' \in \mathrm{RelInt}(E'')$ be the points we have chosen. We define the face-pairing $A_E(E', F', P')$ by applying $A(F', P')$ to the tangent space at u', and then parallel translating from $A(F', P')(u')$ to u''. This definition of the face-pairing is clearly independent (in the appropriate sense) of the choice of the points u' and u''. It is easy to check the truth of Pairing$(\mathscr{P}_E, R_E, A_E)$.

Connected$(\mathscr{P}_E, R_E)$ follows from the connectedness of Γ_E. Cyclic$(\mathscr{P}_E, R_E, A_E)$ follows immediately from Cyclic$(\mathscr{P}, R, A)$. Metric$(\mathscr{P}_E)$ follows from Remark 3.24, applied to $\mathscr{P}_E$. We apply Theorem 4.14 in dimension $n - 1$ to deduce that the developing map $Z(\mathscr{P}_E, R_E, A_E) \to \mathbf{S}^{n-1}$ is an isometry. The induction also tells us that the cell structure of $Z(\mathscr{P}_E, R_E, A_E)$ is finite.

We choose $z \in \mathrm{RelInt}(E)$, and identify E with $\mathrm{id}(E)$, in the notation of Remark 4.10. Each cell of $Z(\mathscr{P}_E, R_E, A_E)$ corresponds to a triple of the form (h, E', P') where h is a member of the finite groupoid $\mathscr{G}(\mathscr{P}_E, R_E, A_E)$. The face-pairings identify z with the point $\pi_Z(z')$, where $z' \in \mathrm{RelInt}(E')$ depends on (h, E', P'). Since the setup is finite, we can choose $\delta > 0$ simultaneously for all (h, E', P') so that the only faces of P' met by the δ-neighbourhood centred at z' are those that contain E'.

There is a map of $Z(\mathscr{P}_E, R_E, A_E)$ into the δ-neighbourhood of z in $Z(\mathscr{P}, R, A)$, since each of the groupoid relations relevant in the definition of the first space will also apply to the second. Any identification of a point of $\mathrm{RelInt}(E)$ with another point, when $Z(\mathscr{P}, R, A)$ is formed, can only be formed as a result of face-pairings which are also in (R_E, A_E). Thus the strong form of local finiteness (see 3.31) is satisfied by $(\mathscr{P}, R, A)$.

The composition of $D_Z: Z(\mathscr{P}, R, A) \to \mathbf{X}^n$ with the obvious map from $Z(\mathscr{P}_E, R_E, A_E)$ to $Z(\mathscr{P}, R, A)$ can be identified with the developing map $Z(\mathscr{P}_E, R_E, A_E) \to \mathbf{S}^{n-1}$ by a change of scale in the range. By induction, this developing map is an isometry. Therefore the obvious map of $Z(\mathscr{P}_E, R_E, A_E)$ to $Z(\mathscr{P}, R, A)$ is injective and the image of $Z(\mathscr{P}_E, R_E, A_E)$ is mapped injectively by D_Z. It follows easily that a neighbourhood of z in $Z(\mathscr{P}, R, A)$ is the cone on $\mathbf{S}^{n-1}$, which is mapped isometrically to $\mathbf{X}^n$ by D_Z. □

The main part of the induction step for Theorem 4.14 will be proved in Section 5. At this point, we prove only a small part of this result.

LEMMA 4.15 (locally finite). *LocallyFinite*$(\mathscr{P}, R, A)$ *follows from the hypotheses of Theorem 4.14 and the inductive hypothesis that Theorem 4.14 is true in dimensions less than* n.

Proof of 4.14. In the proof of Theorem 4.13 we used Locally-Finite$(\mathscr{P}, R, A)$ in order to show that the link of z is embedded in $Z(\mathscr{P}, R, A)$ and that the local picture is as we expect. Here we are trying to prove LocallyFinite$(\mathscr{P}, R, A)$, so the argument needs to be modified. Note that Metric$(\mathscr{P})$, which we are now assuming, implies SecondMetric$(\mathscr{P})$, which in turn implies Metric$(\mathscr{P}_E)$.

The version of Theorem 4.14 for $\mathbf{S}^{n-1}$ is already known inductively, and so we know that $Z(\mathscr{P}_E, R_E, A_E) = \mathbf{S}^{n-1}$. We deduce that the tessellation of $Z(\mathscr{P}_E, R_E, A_E)$ is finite. This means that we have proved the strong form of LocallyFinite$(\mathscr{P}, R, A)$ (see 3.31). □

5. DEFINING A METRIC

If Pairing$(\mathscr{P}, R, A)$ and Connected$(\mathscr{P}, R)$, we obtain the connected quotient space $Q = Q(\mathscr{P}, R, A)$ defined in Remark 3.6. We can define a "metric" on Q in the obvious way: Given two points z_1 and z_2 in Q, we join them with a special kind of path in Q. The path is divided into a finite number of subpaths, and each subpath is the image of a rectifiable path in some $P \in \mathscr{P}$. The distance between z_1 and z_2 is defined as the infimum over all such paths of the sum of the lengths of the subpaths. We get the same infimum if we restrict to subpaths starting and ending in the interior of a codimension-one face; furthermore we may insist that each subpath is a geodesic. (Of course, an exception may have to be made for z_1 and z_2 themselves.) The proof of this is left to the reader — it uses the fact that if two points $\bigsqcup_{P \in \mathscr{P}} P$ are identified, there is a finite sequence of face-pairings connecting them. The axioms for a metric space are easy to verify, except for the condition that $d(z_1, z_2) = 0$ implies that $z_1 = z_2$. Unfortunately, this condition is not always true even if Cyclic($\mathscr{P}, R, A$), as the following example shows.

EXAMPLE 5.1 (only a pseudometric). This example is a variant of Example 3.30. The example will arise from a decomposition of a certain open subset U of $\mathbf{R}^3$ into regions. We define $U = \{(x, y, z) \mid -z < x < z\}$ (which implies in particular that $z > 0$). The boundary of U is the union of two half-planes of slope ± 1, each containing the y-axis $x = z = 0$.

We now explain how to cut U into smaller regions. First we use a countable family of planes, each containing the y-axis, with slopes $1 + 1/m$ and $-1 - 1/m$, where m can be any positive integer. We also use the set of spheres in $\mathbf{R}^3$, lying above and tangent to the plane $z = 0$ at 0, with radii equal either to n or to $1/n$, for some positive integer n. This cuts upper half-space into an infinite number of pieces, parametrized by m and n. A single piece is bounded by (parts of) two half-planes, each with boundary the y-axis, and parts of two spheres, each tangent to the plane $z = 0$ at 0. The piece is closed, and contains 0.

As in the case of Example 3.30, the pieces described are not convex. However, the spherical surfaces can be approximated by finite unions of planar polygons, and then each region can be broken up into a finite union of convex polyhedra. So we have a qualitative description of a family $\mathscr{P}$ of convex polyhedra in $\mathbf{E}^3$, together with face-pairings. We have Pairing($\mathscr{P}, R, A$), Connected($\mathscr{P}, R$) and Cyclic($\mathscr{P}, R, A$). However, the point 0 gives rise to two distinct points in $Q(\mathscr{P}, R, A)$, and these points are zero distance apart. In fact $Q(\mathscr{P}, R, A)$ is not even hausdorff. Also $Z(\mathscr{P}, R, A) = Q(\mathscr{P}, R, A)$ in this particular case.

A very similar example could have been described in dimension two, but then it would not have been possible to satisfy Cyclic($\mathscr{P}, R, A$).

LEMMA 5.2 (metrizable). *Suppose Pairing($\mathscr{P}, R, A$), Connected($\mathscr{P}, R$), Cyclic($\mathscr{P}, R, A$) and LocallyFinite($\mathscr{P}, R, A$). Then $Q = Q(\mathscr{P}, R, A)$ is a metric space, with the metric defined as at the beginning of this section. Also $Z = Z(\mathscr{P}, R, A)$ has a metric defined in a similar way, and Z with this metric is locally isometric to $\mathbf{X}^n$. The topologies defined by these metrics are the appropriate quotient topologies.*

Proof of 5.2. We have already seen in Theorem 4.13 that Z is modelled on $\mathbf{X}^n$ under the given hypotheses and that the polyhedral cell structure of Z is locally finite. It follows immediately that the metric structure on Z given by piecewise rectifiable paths induces the correct topology on Z. We have also seen in 4.13 that G acts properly discontinuously on Z. It follows that $Q = Z/G$ is hausdorff with the quotient topology. Also every point in Q has a neighbourhood which is homeomorphic to the quotient of a disk in Z by a finite group of isometries. The radius function is invariant under the finite group, and therefore gives a map which does not increase distances from a neighbourhood of a point in Q to $[0, \delta]$. (This is proved by seeing that the radius function does not increase distances on the intersection of any $P \in \mathscr{P}$ with the inverse image of our neighbourhood.) From this it is easy to see that the metric on Q is indeed a metric, and that it induces the right topology. $\square$

Lemma 4.15 implies the following result.

COROLLARY 5.3. *The conclusions of Lemma 5.2 hold if we have Pairing*$(\mathscr{P}, R, A)$, *Connected*$(\mathscr{P}, R)$, *Cyclic*$(\mathscr{P}, R, A)$ *and Metric*$(\mathscr{P})$.

We are now able to prove Theorem 4.14. We are allowed to assume the truth of Theorem 4.13 in all dimensions up to and including dimension n.

Proof of 4.14. We will prove that there is an $\varepsilon > 0$ such that any point of Z has a neighbourhood in Z which is isometric to a ball in $\mathbf{X}^n$ of radius ε. We denote by Z^i the i-skeleton of Z, namely the union of the polyhedral cells of Z of dimension at most i. We prove, by induction on i for $0 \leqslant i \leqslant n$, that there is an $\varepsilon_i > 0$, such that each point $z \in Z$ satisfying $d(z, Z_i) < \varepsilon_i$ has a neighbourhood in Z which is isometric to an ε_i-ball in $\mathbf{X}^n$.

Suppose $0 \leqslant j \leqslant n$ and that the induction statement is known for $i < j$. We take $\varepsilon_i < \varepsilon_{i-1}/2$. Then if $d(z, Z^{i-1}) < \varepsilon_{i-1}/2$ the desired result is true for z. So we suppose that $d(z, Z^{i-1}) \geqslant \varepsilon_{i-1}/2$. We have already seen in Theorem 4.13 that z has a small neighbourhood in Z which is isometric to a ball in $\mathbf{X}^n$ with centre z. It is clear from the cone structure on the neighbourhood in z that we can take the ball to have radius r, where r is the distance from z to the union of the faces not containing z.

To proceed, recall that Remark 3.24 together with the hypothesis 4.14(k) gives us the condition Metric$(\mathscr{P})$ in the euclidean or spherical case. Also Remark 3.24 together with the hypothesis 4.14(l) imply Metric$(\mathscr{P})$ in the hyperbolic case. Looking through the statement of Theorem 4.14, we see that we may therefore assume Metric($\mathscr{P}$). It is clear that Metric($\mathscr{P}$) gives a lower bound for r in terms of ε_{i-1}.

Having found ε as promised, it is standard that the developing map $D_Z : Z \to \mathbf{X}^n$ is an isometry. For completeness, we give the proof. We first note that the image of D_Z is an open subset of $\mathbf{X}^n$, since D_Z is a local isometry (by Theorem 4.13). Using ε it is clear that the image is also closed, and is therefore the whole of $\mathbf{X}^n$. The inverse image in Z of the open ε-ball B centred at any point of $\mathbf{X}^n$ is a disjoint union of open subsets of Z, each mapped isometrically onto B. It follows that D_Z is a covering map. Since $\mathbf{X}^n$ is simply connected and Z is connected, D_Z is a homeomorphism and therefore an isometry. □

LEMMA 5.4 (completeness of Q and Z). *Under the same hypotheses as in Lemma 5.2, Q is complete if and only if Z is complete.*

Proof of 5.4. Suppose Q is complete. To deduce that Z is complete, consider a Cauchy sequence (x_n) in Z. Then $(\pi_{ZQ}(x_n))$ is a Cauchy sequence in Q, and therefore has a limit p. We take a small neighbourhood N of p, in particular a neighbourhood meeting only a finite number of polyhedral cells. The inverse image of N under π_{ZQ} is a union of components, each of which is isometric to a round ball in $\mathbf{X}^n$. The stabilizer in G of any such component is a finite group. The quotient of the component by this finite group gives N, and the inverse image of p in the component is a single point. By making N smaller, we may assume that there is an $\varepsilon > 0$ such that any two of these components are at least ε apart. From this we see that (x_n) must eventually stay in one of these components. It follows that (x_n) converges to a point in Z.

Now suppose that Z is complete. To deduce that Q is complete, consider a Cauchy sequence (y_i) in Q. By moving each y_i a little, we may assume that it lies in the interior of a top-dimensional cell. By taking a subsequence, we may assume that $d(y_i, y_{i+1}) < 2^{-i}$ for each i. We may join y_i to y_{i+1} by a path in Q of length less than 2^{-i}, which avoids the $(n-2)$-skeleton of Q. This gives us a rectifiable path in Q from y_1, going through each of the points y_i. We now choose a point $z_1 \in Z$ in the inverse image of y_1. Since the path avoids the $(n-2)$-skeleton, there is a unique lift to Z of the path, starting at z_1. Since Z is complete, the path converges to a limit, which we call z_0. Since the projection map π_{ZQ} is continuous, it follows that (y_i) converges to the limit $\pi_{ZQ}(z_0)$. □

THEOREM 5.5 (Poincaré's Theorem Version 2). *Suppose the hypotheses Pairing*($\mathscr{P}, R, A$), *Connected*($\mathscr{P}, R$), *Cyclic*($\mathscr{P}, R, A$) *and Locally-Finite*($\mathscr{P}, R, A$) *are satisfied. If* Z *is complete, then it is isometric to* $\mathbf{X}^n$.

Proof of 5.5. Since Z is complete, all geodesics can be extended indefinitely. It follows that the developing map $D_Z : Z \to \mathbf{X}^n$ is a covering map. Since Z is connected, the developing map is an isometry. □

6. COMPLETENESS

In this section we discuss questions of completeness in more detail, in relation to the case of a finite number of finite-sided hyperbolic polyhedra. We have already seen in Theorem 4.14 that completeness follows from Finite($\mathscr{P}$) in the euclidean and spherical cases, so no special discussion is necessary in those cases. We also discuss the question of verifying the hypotheses of Poincaré's Theorem algorithmically, giving attention mainly to completeness in the hyperbolic case. We give a detailed account of other aspects of an algorithmic approach in Section 7. Such an algorithm only makes sense if a single real number is regarded as a single datum, as opposed to the Turing machine model where a real number is known only as a bitstring, and can therefore never be specified precisely. (In practice, Poincaré's Theorem is often used in connection with a group of matrices over an algebraic number field. In this case, the conventional Turing machine model can be used.) We need a mathematical model which allows addition, multiplication and division of two real numbers with perfect accuracy and in unit time. Such a model is discussed in [BSS89].

THEOREM 6.1. *There is an algorithm (in the sense of* [BSS89]) *which has a finite set* $\mathscr{P}$ *of convex polyhedra, each with a finite number of faces, and a set of face-pairings as its input, and as its output the answer to the question "Does this data define a tesselation of* $\mathbf{X}^n$*?" More precisely, "Does this data allow us to define* Z *and is* Z *isometric to* $\mathbf{X}^n$*?"*

The proof of the theorem just stated is discussed in more detail in Section 7; here we cover the main points only.

The various aspects of an algorithmic approach are fairly straightforward, with the exception of an algorithmic check that Z is complete. In order to check our conditions algorithmically, we are of course restricted to a finite set of data, and, as we have already said, a single real number is regarded as a single datum. We assume that we are given a finite set $\mathscr{P}$ of convex polyhedra in $\mathbf{X}^n$, each with a finite number of faces. We are also given a finite number of face-pairings. We can check Connected$(\mathscr{P}, R)$ and Cyclic$(\mathscr{P}, R, A)$, and then Z can be constructed. By 4.14(k) we know that Z is complete except in the hyperbolic case, where further checking is necessary. From now on we assume we are in the hyperbolic case. We will find necessary and sufficient conditions for completeness, which are algorithmically checkable.

If P is a convex polyhedron in $\mathbf{H}^n$, let $\bar{P}$ be the closure of P in the closure of hyperbolic space. If p is an ideal boundary point of P, let H_p be a horosphere centred at p, chosen so that the corresponding horoball is disjoint from each face of P which does not contain p in its closure. In the upper half-space model, with p the point at infinity, H_p is a horizontal plane. Each face whose closure contains p lies in a vertical half-plane, and every other face is contained in a hemisphere which is orthogonal to the boundary plane $\mathbf{R}_0^{n-1}$ of the upper half-space. We assume that none of the codimension-one faces of P lies in a hemisphere which meets H_p. We may regard $P \cap H_p$ as an $(n-1)$-dimensional euclidean convex polyhedron, in view of the fact that H_p is isometric to $\mathbf{R}^{n-1}$.

We define the *impression,* denoted $I(A)$, of an $(n-1)$-dimensional euclidean convex polyhedron A as the subset of S^{n-2} consisting of all directions with the property that a point moving along a line in that direction stays at a bounded distance from A. The distance between two directions is the angle between them. This definition is due to Brian Bowditch (see Appendix). Note that a euclidean similarity between euclidean convex polyhedra gives rise to an isometry of the associated impressions. The impression of a convex polyhedron either consists of two antipodal points, or is a connected convex polyhedron in S^{n-2}. The impression of a compact convex polyhedron is empty.

Returning to the case of a pair (P, p), where P is a hyperbolic convex polyhedron and p is a point in the ideal boundary of P, we see that we can identify the impression of $H_p \cap P$ with the set of tangent directions v at p to S^{n-1} for which there is a curve in $S^{n-1} \cap \bar{P}$ starting at p with non-zero derivative in the direction of v. We talk of the *impression of P at p.* If the impression has non-empty interior, we say that P is *fat* at p. Otherwise we say that P is *thin* at p. If P is thin at p, it must have two faces F_1 and F_2 whose closures $\bar{F}_1$ and $\bar{F}_2$ meet in p only. In the Appendix, in this situation we refer to (P, p) as being *non-pyramidal.*

Consider for example the region P in the upper half-space model of $\mathbf{H}^n$ lying between two parallel vertical codimension-one subspaces and let p be the point at infinity. Then the impression of P at p is an equatorial S^{n-3} in S^{n-2}, and P is thin at p.

As P varies over $\mathscr{P}$ and p varies over the ideal boundary of P, there are only a finite number of isometry types of impression of P at p. This is because the impression does not change as p varies in $X \setminus Y$, where X is a connected component of the ideal boundary of a face F, and Y is the set of ideal boundary points of the proper faces of F. In particular, there is an integer $N > 0$, such that, for each ideal point p of any $P \in \mathscr{P}$, the volume of the impression of P at p is either zero or is greater than $\mathrm{vol}(S^{n-2})/N$.

Suppose $\mathscr{P}$ is a finite collection of hyperbolic convex polyhedra, each with a finite number of faces. Suppose we are given a set of face-pairings which satisfy Pairing($\mathscr{P}, R, A$), Connected($\mathscr{P}, R$) and Cyclic($\mathscr{P}, R, A$). Let Z and Q be as in Definition 4.5 and Remark 3.6. Let $\bar{Q}$ be the quotient space of the disjoint union of the $\bar{P}$'s by (the extension to the ideal points of) the given face-pairings.

Given a pair (P, p), we develop Z into upper half-space, with p being sent to the point at infinity. The developing map D_Z is determined up to composition with a euclidean similarity of $\mathbf{R}^{n-1}$, acting as a hyperbolic isometry keeping the point at infinity fixed. We will restrict our attention to the development of pairs (P', p') such that p' is sent to the point at infinity. More precisely, having defined the developing map on a certain collection of n-cells of Z, we look only at those codimension-one faces of these n-cells which are mapped to vertical faces extending upwards to infinity, and extend the developing map across these faces.

Another way of thinking about the situation is to define a graph Γ_∞ as follows. The vertices of Γ_∞ are pairs of the form (P, p) where $P \in \mathscr{P}$ and p is an ideal point of P. For each face-pairing $A(F, P)$, such that p is an ideal point of F, Γ_∞ contains an edge from (P, p) to (P', p'), where $R(F, P) = (F', P')$ and $p' = A(F, P)(p)$. An edge from (P, p) to (P', p') arising from $A(F, P)$ is identified with the edge from (P', p') to (P, p) arising from $A(F', P')$. In general the number of vertices of Γ_∞ is uncountable. However, we are only interested in the components of this graph and each component has at most a countable number of vertices. We denote by $\Gamma_{P,p}$ the component of Γ_∞ containing the vertex (P, p).

In Example 3.32 we give an example where $\Gamma_{P,p}$ is countable, but not all the current hypotheses are satisfied. The appendix to this paper by Brian Bowditch shows that in fact $\Gamma_{P,p}$ is always finite under the current hypotheses, but the body of the paper will not assume this result (Theorem 10.1 (Bowditch)). Example 3.32 shows that $\Gamma_{P,p}$ can be infinite if Cyclic($\mathscr{P}, R, A$) is not satisfied. A famous example where $\Gamma_{P,p}$ has eight vertices (due to Gieseking, Riley and Thurston) is formed from two regular ideal hyperbolic tetrahedra by appropriate face-pairings to give a complete hyperbolic structure on the complement of a figure-eight knot. All eight vertices of the two tetrahedra are identified to a single ideal boundary point. In this case the restricted developing image (see Definition 6.2) entails four different versions of (T_1, p) and four different versions of (T_2, p), where T_1 and T_2 are the two ideal tetrahedra and p varies over the four ideal vertices.

Let p be an ideal boundary point of an n-cell P of Z and let $\Gamma_{P,p}$ be the associated graph. We define the subspace Z_p to be the smallest subspace of Z which is a union of cells, one of which is equal to P, and such that any vertical codimension-one face F of an n-cell of Z_p is also the face of another n-cell of Z_p on the other side of F. (Note that any vertical face of an n-cell in Z_p must extend upwards to infinity by convexity.) The face-pairings that come up are all associated to the edges of $\Gamma_{P,p}$.

DEFINITION 6.2 (restricted developing map). Let $D_p : Z_p \to \mathbf{H}^n$ be the restriction of D_Z. We call D_p the *restricted developing map* associated to p.

Each n-dimensional cell of Z_p is mapped to a convex polyhedron in upper half-space with at least one vertical codimension-one face which extends upwards to infinity. (To be completely precise, there is also the case where p is in the ideal boundary of P, but not in the ideal boundary of any proper face of P. In that case, $\Gamma_{P,p}$ consists of a single vertex, Z_p consists of one cell only; the impression of this cell at p is the whole of S^{n-2}, and there are no vertical codimension-one faces.)

We are not assuming, in the body of the paper, that $\Gamma_{P,p}$ is finite. In these circumstances, it is not to begin with clear, even in the case that Z is complete, that we can choose a single horosphere which is disjoint from all non-vertical faces in the restricted developing image. However, if we confine our attention to the image of only a finite number of cells of Z in the restricted developing image, we can take the horosphere high enough to achieve the desired disjointness property for the finite number of cells.

Since the horosphere centred at p is not unique, its intersection with P gives a euclidean convex polyhedron which is only determined up to similarity. This enables us to define a similarity $(n-1)$-manifold S_p associated to Z_p.

(Although we are not assuming that there is a horosphere disjoint from all the non-vertical faces, the similarity structure may be constructed locally.) Let G_p be the group generated by the face-pairings arising from vertical faces of Z_p, modulo the relations coming from codimension-two vertical faces. The image of G_p in the isometry group of the upper half-space model of $\mathbf{H}^n$ consists of isometries which fix the point at infinity. Its image is a group of similarity transformations which preserve the cell structure of S_p.

We say that Z_p has a *consistent horosphere* if we can choose a horizontal horosphere which lies above all non-vertical faces in the developing image of Z_p, and which is mapped to itself by each face-pairing corresponding to a vertical codimension-one face in the developing image of Z_p. This is equivalent to saying that there are well-defined horospheres in the quotient Q, such that the intersection of a horosphere with a cell P' of Q has exactly one component for each pair (P', p') such that p' and p are identified in $\bar{Q}$. If there is a consistent horosphere at p, then the image of G_p consists entirely of euclidean isometries of $\mathbf{R}^{n-1}$ and S_p can be identified with this consistent horosphere.

Let $\Gamma_{P,p}$ be the graph defined earlier in this section. This graph results from taking a vertex for each pair (P', p') corresponding to a cell of Z_p and an edge for each face-pairing corresponding to a vertical codimension-one face. (In general there will be many cells of Z_p, possibly an infinite number, corresponding to a single pair (P', p').)

THEOREM 6.3 (checking completeness). *Suppose we have a set $\mathcal{P}$ of hyperbolic convex cells satisfying Pairing$(\mathcal{P}, R, A)$, Connected$(\mathcal{P}, R)$, Cyclic$(\mathcal{P}, R, A)$ and Finite$(\mathcal{P})$. Then the following conditions are equivalent.*

(a) *Z is complete.*

(b) *For each $P \in \mathcal{P}$ and each boundary point p of P, Z_p has a consistent horosphere.*

(c) *For each $P \in \mathcal{P}$ and each boundary point p of P, one of the following two mutually exclusive situations prevails:*

   (i) *$\Gamma_{P,p}$ is finite, has some fat vertex and the group G_p is finite.*

   (ii) *For each pair (P', p'), such that p' and p have the same image in $\bar{Q}$, P' is thin at p'. The group G_p does not contain any hyperbolic or loxodromic elements.*

(d) $\bar{Q}$ *is hausdorff and each point of* $\bar{Q}$ *has a neighbourhood whose intersection with* Q *is complete.*

It is possible to check Condition 6.3(c) algorithmically.

Proof of 6.3. Equivalence of (b) and (c) is easy and we assume it (a proof of this fact is actually implicit in the argument we give below).

First suppose that Z is complete. Equivalently, the developing map $D_Z : Z \to \mathbf{X}^n$ is an isometry.

There are only a finite number of faces of the various $P \in \mathscr{P}$. It follows that there are only a finite number of peaks among the ideal points of P. This implies that the set of thin vertices of $\Gamma_{P,p}$ is finite.

Recall the definition of the integer $N > 0$: for each ideal point p of any $P \in \mathscr{P}$, the volume of the impression of P at p is either zero or is greater than $\mathrm{vol}(S^{n-2})/N$. Suppose that Z is complete and that Z_p has a cell corresponding to a pair (P', p') where P' is fat at p'.

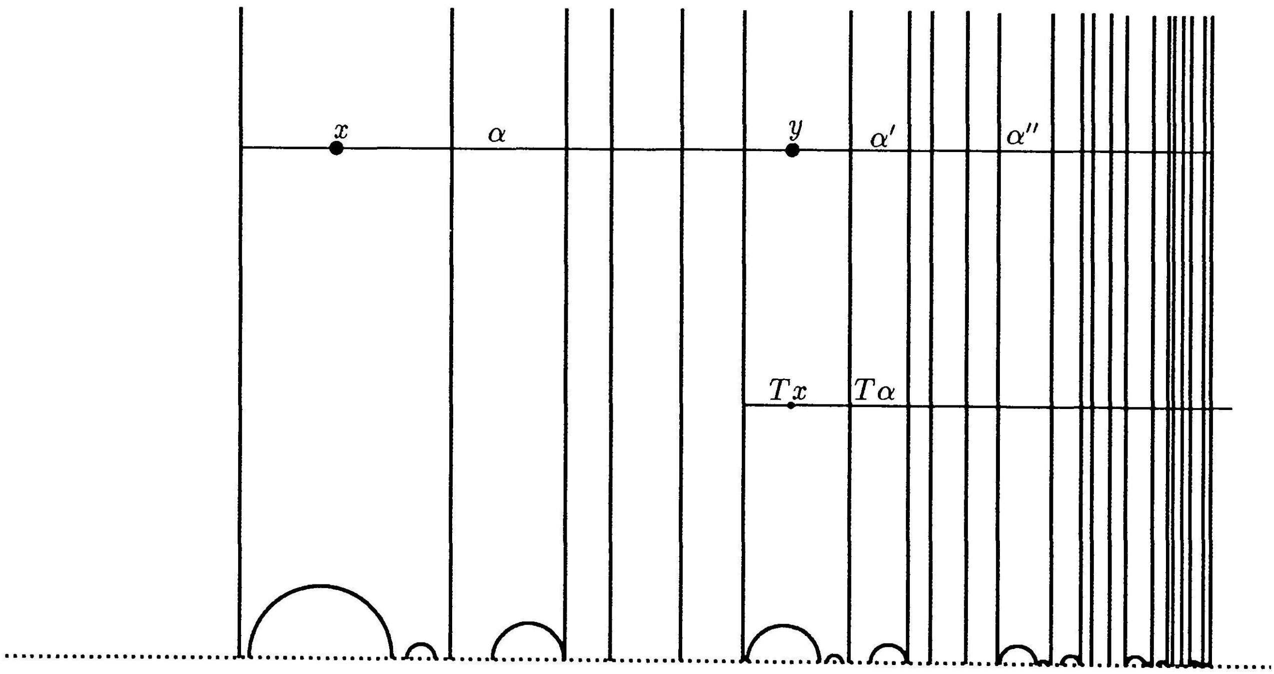

FIGURE 12.
Failing to construct a consistent horosphere.
This illustrates part of the proof of Theorem 6.3.

There can be at most N such cells in Z_p, for otherwise the images of two different n-cells of Z have developing images whose interiors intersect. But this would contradict the completeness of Z. It follows that Z_p must be finite, and so G_p must be finite. We deduce that there is a consistent horosphere.

Now suppose every cell of Z_p is a pair (P', p') such that P' is thin at p'. Then $\Gamma_{P,p}$ is a finite graph. In order to check that there is a consistent horosphere, we need only check that we can construct a consistent horosphere along each simple circuit in $\Gamma_{P,p}$ (that is, a circuit in which no vertex is repeated). We construct a horosphere following some circuit, and we check whether it matches up when we return. The holonomy map corresponding to the circuit is then an isometry of the upper half-space model of $\mathbf{H}^n$ which fixes the point at infinity. We want to show that if Z is complete, then this holonomy must preserve setwise each horosphere centred at infinity. If not, we may assume (by reversing the direction of the circuit if necessary) that the holonomy is a euclidean similarity T with λ as change of scale, $0 < \lambda < 1$.

We take a horizontal path α in upper half-space, following the circuit in $\Gamma_{P,p}$. This path goes from a point x in the interior of some n-cell C of Z_p to a point y in the interior of TC, such that Tx is directly below y at a height λ times that of y. We continue α with the path α' formed as follows. We take the horizontal path $T\alpha$ and translate it (translation in the euclidean meaning) upwards until the ends match at y. The euclidean length of α' is the same as the euclidean length of $T\alpha$, but the hyperbolic length is λ times the hyperbolic length of $T\alpha$, which is also λ times the hyperbolic length of α. Continuing in this way, we get a path $A = \alpha\alpha'\alpha''\ldots$ whose length is finite. This is a Cauchy path in Z which must have a limit in Z, since Z is complete. Since the cell structure of Z is locally finite, this means that A passes through only a finite number of codimension-one faces of Z. But by the construction of A, this is not the case.

This proves by contradiction that we can construct a consistent horosphere for Z_p. (Note that we may assume that α lies above all the non-vertical codimension-one planes containing codimension-one faces in the finite set of cells that it passes through. Therefore the same is true for $T\alpha$. Since α' lies at a higher level than $T\alpha$, the same is true for α'. Inductively A lies above all such planes bounding non-vertical faces of cells that it meets.)

Now we assume Condition 6.3(b) and show that $\bar{Q}$ is hausdorff and that each point of $\bar{Q}$ has a neighbourhood whose intersection with Q is complete. Under the conditions stated, G_p acts as a group of isometries of $\mathbf{R}^{n-1}$. From Poincaré's Theorem applied to $\mathbf{R}^{n-1}$, we see that the portion of Z_p above the consistent horosphere tessellates the part of upper half-space above the corresponding horizontal plane. Moreover, G_p acts on this tessellation effectively, as a discrete group of parabolic or elliptic transformations.

We know that any discrete group G_0 of euclidean transformations acting on a euclidean space E gives rise to a non-empty affine subspace W on which G_0 is either fixed or acts by translations. W is foliated by affine subspaces V which are minimal G_0-invariant affine subspaces of E. (See [Bow93].)

The next step is to form a *standard cusp region* — originally defined in [Bow93] — for G_p acting on the upper half-space model of $\mathbf{H}^n$ as we now explain. We fix a minimal G_p-invariant affine subspace V of $\mathbf{R}_0^{n-1}$, the boundary of upper half-space. (If G_p is finite, then V is a point.) Then a standard cusp region in our situation will be the set of points x in upper half-space whose euclidean distance from V is at least r, and r is chosen suitably large. In our case, we fix a representative n-cell in Z_p for each relevant pair (P', p'), and then ensure that our standard cusp region is small enough (r is large enough) so that it is disjoint from each non-vertical codimension-one plane containing a codimension-one face of the n-cell. Since there are only a finite number of such pairs (P', p'), this is easy to do. Any other cell which is in the developing image of Z_p is the image of one of our representatives under some element of G_p. Since the standard cusp region is G_p-invariant, the desired condition of disjointness from non-vertical faces holds for all cells of Z_p. The closure of the standard cusp region in closed hyperbolic space projects to a neighbourhood of the image of p in $\bar{Q}$. This neighbourhood is isomorphic to the quotient of the closure of the standard cusp region by G_p. It is easy to see that it has the desired completeness properties.

This proves the desired completeness property for all points of $\bar{Q} \setminus Q$. The completeness property for points of Q itself follows from the fact that Z is a manifold and Q is an orbifold covered by Z.

To see that $\bar{Q}$ is hausdorff, note that for each point of $\bar{Q} \setminus Q$, we have a sequence of (quotients of) standard cusp regions, whose intersection is a unique point of $\bar{Q}$.

Now suppose 6.3(d) is satisfied. Since $\bar{P}$ is compact for each $P \in \mathscr{P}$, $\bar{Q}$ is a compact hausdorff space. Therefore we have a finite covering of $\bar{Q}$ by sets whose intersection with Q is complete. It follows that Q is complete. Lemma 5.4 now shows that Z is complete.

This proves the equivalence of the conditions in Theorem 6.3. We still need to show that we can check for completeness algorithmically starting with the input data (R, A). Note first that we can count the number of peaks in $\Gamma_{P,p}$, and we know that we cannot have more than N fat vertices in the complete case. This gives an upper bound b_0 to the possible size of $\Gamma_{P,p}$.

As p varies within the set of ideal points associated to the interior of a face of some $P \in \mathscr{P}$, $\Gamma_{P,p}$, Z_p and D_p will be essentially unchanged. This means that we can reduce our search to a finite number of vertices (P, p) of Γ_∞. We focus attention on one of these cases. We start to explore $\Gamma_{P,p}$. If we find more than b_0 vertices, we know Z is not complete. Otherwise we find a generating set of circuits in $\Gamma_{P,p}$ and check for each of these that a consistent horosphere can be constructed. $\square$

7. ALGORITHMIC ASPECTS

We will now look more closely at the algorithmic aspects of Poincaré's Theorem. We wish to produce a mechanical procedure which takes as input a finite number of finite-sided convex polyhedra in $\mathbf{E}^n$ or $\mathbf{S}^n$ or $\mathbf{H}^n$, together with a finite set of face-pairings, and which outputs "Yes" or "No" to the question of whether these polyhedra and face-pairings give a tessellation of the appropriate space. In the case that the answer is "Yes", it also outputs a presentation for the group of symmetries of this tessellation with the given finite union of finite polyhedra as a fundamental domain.

What kind of mathematical model of a computing machine is necessary in order to carry out the procedure described in the preceding pages? It is not appropriate to use a Turing machine model. A Turing machine is not capable of taking as input a list of real numbers and coming out with the answer "Yes" or "No". We need to be able to handle real numbers not as sequences of bits but as entities. We need to be able to compare two real numbers for equality or inequality in a one-step operation, and likewise for addition and multiplication and division of real numbers.

Such a mathematical model has been described in [BSS89]. Their model is devoted to the study of polynomial and rational maps, and it is assumed that computation of a polynomial can be carried out in a single step. In most computations in hyperbolic or spherical geometry, trigonometric and hyperbolic trigonometric functions are likely to arise, and so it seems at first sight that a model of computation able to carry out only polynomial operations would not be relevant. However, in the case of Poincaré's Theorem it happens that the computation can be expressed in polynomial terms. Since the BSS scheme has been thought out and developed far enough to be a reasonable tool, we use it.

However, for more general computations in geometry, it seems that it would be more satisfactory to have a computational model with a library of functions, satisfying certain axioms. It might, for example, be assumed that any of the functions in the library could be computed with complete accuracy in a single step. We put this forward in the hope of encouraging someone to develop such an approach.

Let us go through the steps of the computation to see what kind of operations are necessary. We need to start by making a decision as to how to represent the input data. We first need to decide how to represent $\mathbf{E}^n$, $\mathbf{S}^n$ and $\mathbf{H}^n$. It is convenient in each case to embed the space in $\mathbf{R}^{n+1}$. In order to be able to change basis easily, we will describe the situation in a general manner.

Suppose we are given a positive definite symmetric $(n+1) \times (n+1)$ real matrix M_S defining a positive definite inner product on $\mathbf{R}^{n+1}$. We define $\mathbf{S}^n$ to be the set of vectors v of unit length with respect to this inner product. We will frequently represent a point in $\mathbf{S}^n$ by a non-zero vector which does not have unit length; conceptually this can be normalized, but computationally we will not normalize. The reason for avoiding normalization is that BSS machines are capable of polynomial operations, but not of taking square roots.

We take $\mathbf{E}^n$ to be an affine subspace of $\mathbf{R}^{n+1}$ which does not contain the origin. We can think of the subspace as specified by a non-zero linear map $\lambda : \mathbf{R}^{n+1} \to \mathbf{R}$ as follows: $\mathbf{E}^n = \{v : \lambda(v) = 1\}$. As in the case of $\mathbf{S}^n$, we assume that $\mathbf{R}^{n+1}$ has a positive definite inner product, given by a matrix M_E. We will often represent a point in $\mathbf{E}^n$ by an element $v \in \mathbf{R}^{n+1}$, such that $\lambda(v) > 0$, without supposing $\lambda(v) = 1$. Multiplying by a positive scalar, we find a vector in $\mathbf{E}^n$.

If we are given a real $(n+1) \times (n+1)$ symmetric matrix M_H with n positive eigenvalues and one negative eigenvalue, we obtain a non-degenerate indefinite inner product on $\mathbf{R}^{n+1}$ of type $(n, 1)$. We define $\mathbf{H}^n$ to be one sheet of the hyperboloid $\{v \in \mathbf{R}^{n+1} : \langle v, v \rangle = -1\}$. We specify such a sheet by fixing a linear map $\lambda : \mathbf{R}^{n+1} \to \mathbf{R}$, such that the sheet lies in the half-space $\lambda > 0$. A vector v such that $\langle v, v \rangle < 0$ and $\lambda(v) > 0$ represents a well-defined point of $\mathbf{H}^n$, obtained by multipling by a suitable positive scalar. However this scalar cannot be computed by our BSS machine, since the computation involves taking a square root.

If $\mathbf{X}^n$ is any of the three spaces, we specify a codimension-one $\mathbf{X}$-subspace by means of a single linear equation, and a general $\mathbf{X}$-subspace by means of a finite number of linear equations (with no constant term). The condition on the subspace in the hyperbolic case is that the coefficient vectors of the linear equations define a positive definite subspace with respect to M_H. In the euclidean case the condition is that the coefficient vector of λ, the linear map defining $\mathbf{E}^n$, is not linearly dependent on the set of coefficient vectors of the linear inequalities. In the spherical case there are no conditions. Such a subspace is therefore determined by a finite list of real numbers.

Face-pairings are represented by matrices. A half-space is determined by a linear inequality (with no constant term).

Each finite collection of linear equalities and inequalities (satisfying appropriate conditions to give a codimension-one subspace or a codimension-zero half-space) defines either the nullset or some i-dimensional convex polyhedron in $\mathbf{X}^n$. If there are exactly $n - i$ linearly independent equalities and if each of the half-spaces is essential, the defining collection of equalities and inequalities is minimal. If $i = n$, there is a unique minimal set of defining half-spaces (see Proposition 2.5). If $i < n$, the number of inequalities in a minimal collection is equal to the number of codimension-one faces, but neither the equalities nor the inequalities are uniquely determined. Any collection of equalities and inequalities defining the i-dimensional polyhedron can be transformed into a minimal collection by changing some of the inequalities to equalities and then omitting some of the equalities and inequalities. (For example the two conditions $\mu \geqslant 0$ and $\mu \leqslant 0$ are equivalent to the one condition $\mu = 0$.)

THEOREM 7.1 (BSS polyhedron computation). *There is a BSS program which carries out the following computation. We input* $\mathbf{E}^n$, $\mathbf{H}^n$ *or* $\mathbf{S}^n$, *represented by a real non-singular symmetric* $(n+1) \times (n+1)$ *matrix* M_E, M_H *or* M_S, *and a linear map* $\lambda : \mathbf{R}^{n+1} \to \mathbf{R}$. *We also input a finite set of linear equalities and inequalities defining codimension-one subspaces and codimension-zero half-spaces in* $\mathbf{E}^n$, $\mathbf{H}^n$ *or* $\mathbf{S}^n$ *respectively. The output from the program is the dimension* i *of the convex polyhedron defined by the intersection of these subspaces and half-spaces, the combinatorial structure of its faces, for each face a minimal subset of the defining equalities and inequalities (with some of the defining inequalites converted to equalities). The program also outputs for each face* F *an element* $x_F \in \mathbf{R}^{n+1}$ *representing a point in the relative interior of the face.*

Proof of 7.1. If $n = 0$ or $n = 1$, the result is clearly true. Inductively we assume the result is known for dimensions less than n.

If the collection of equalities and inequalities input includes one or more equality, then the result follows by induction on n. To see this, we transform by a matrix which changes one of the equalities to $x_{n+1} = 0$. This changes the matrix of the inner product and the coordinates of λ. In the hyperbolic case we next check that $(0, \ldots, 0, 1)$ is a positive vector (otherwise the plane $x_{n+1} = 0$ is not a plane in hyperbolic space). In the euclidean case, we check that λ does not have the form $cx_{n+1} = 0$ in the new coordinates. (If these checks fail, then the input data was inconsistent). We then apply the BSS program which has been inductively constructed. So we may asume that our input consists of inequalities only.

We now assume that we have j inequalities, and that we have constructed a program which gives the required output for any collection of $j-1$ inequalities satisfying the induction assumptions. Let $P \subset \mathbf{X}^n$ be the convex polyhedron defined by the first $j-1$ inequalities. Let $f \geqslant 0$ be the j-th inequality. We first construct a point x_f representing a point $y_f \in \mathbf{X}^n$, such that $f(x_f) = f(y_f) = 0$.

One of the following situations must hold, and we want to construct a BSS program to find which.

CONDITION 7.2 (situation for (P, f)).

(a) $P \subset \{f = 0\}$. In the other cases, we assume that P is not a subset of $\{f = 0\}$.

(b) $P \subset \{f > 0\}$.

(c) $P \subset \{f \geqslant 0\}$ and the codimension-one subspace $f = 0$ meets P.

(d) P meets $f > 0$ and $f < 0$.

(e) $P \subset \{f \leqslant 0\}$ and the codimension-one subspace $f = 0$ meets P.

(f) $P \subset \{f < 0\}$.

Assuming we know which case we are in, the inductive proof deals with all cases except Condition 7.2(d), when we need also to compute the new face structure and to find a representative for a point in the relative interior of each new face.

We proceed as follows, assuming that we are in case Condition 7.2(d). If P has no faces, then $P = \mathbf{X}^n$. The new polyhedron has two cells, namely $f \geqslant 0$ and $f = 0$. It is easy to find representatives for points in these two faces. (Solving linear equations can be done by row operations.)

If P does have faces, we first tackle the same problem for each proper face. Let F be a face and let S be the smallest $\mathbf{X}$-subspace containing F. If the plane $f = 0$ meets S, then either $f = 0$ contains S, which we can check by a linear independence computation (row operations), or $f = 0$ meets S in a codimension-one subspace of S. In both cases we can treat the problem by induction on n. If the plane $f = 0$ does not meet S, then take the point $x_F \in \operatorname{Int}(F)$ given by our induction, and evaluate $f(x_F)$. The value is either negative, in which case we must be in case 7.2(f) for the pair (F, f), or it is positive, in which case we must be in case 7.2(b) for (F, f).

The construction of the point in the relative interior of $\{f \geqslant 0\} \cap P$ in case 7.2(d) is as follows. By induction we find a point in the relative interior of $\{f = 0\} \cap P$. This means that all the equalities required for the definition of P are satisfied for this point, and all the inequalities required are satisfied as strict inequalities. We can therefore move the point a little so that the equalities and the strict inequalities continue to hold. In addition we move it in a direction away from $\{f = 0\}$ so as to increase f.

Now let us see how to recognize which case we have for (P, f). Using our minimal set of equalities and inequalities for P, a linear independence check (row operations) tells us whether or not we are in case 7.2(a).

We are in case 7.2(b) for (P, f), if firstly for each proper face F of P, (F, f) is in case 7.2(b), which shows that $f = 0$ does not meet ∂P, and secondly we check that that $x_f \notin P$. Case 7.2(f) is treated in the same way.

We are in case 7.2(c) if the following three conditions are satisfied: firstly for each proper face F of P we have case 7.2(a) or case 7.2(b) or case 7.2(c) for (F, f), secondly we are in case 7.2(a) for some face F, and thirdly $f(x_P) > 0$ for the point x_P already constructed in the relative interior of P. Similarly for the case 7.2(e), except that the signs are changed.

To see if we are in case 7.2(d), we change coordinates so that $f = 0$ becomes the plane $x_{n+1} = 0$, and then look at the intersection P' of this plane with P. We take a point in the relative interior of P' and check whether it is in the relative interior of P. $\square$

To complete the discussion of the algorithmic approach, suppose we are given a finite set $\mathscr{P}$ of finite-sided polyhedra and maps $R : \mathscr{F}(\mathscr{P}) \to \mathscr{F}(\mathscr{P})$ and $A : \mathscr{F}(\mathscr{P}) \to \operatorname{Isom}(\mathbf{X}^n)$, where $\mathscr{F}(\mathscr{P})$ is the set of codimension-one faces of the polyhedra in $\mathscr{P}$. There is obviously no problem in checking Pairing$(\mathscr{P}, R, A)$, Finite$(\mathscr{P})$ and Connected$(\mathscr{P}, R)$. To check Cyclic$(\mathscr{P}, R, A)$, we need to be more explicit about the form in which the face-pairings are given. We will assume each face-pairing is given by an $(n + 1) \times (n + 1)$ matrix of real numbers which preserves the appropriate structure. Then we can check Cyclic$(\mathscr{P}, R, A)$ by multiplying such matrices together. The fact that a certain product is the identity on a codimension-two face can be checked by a linear independence calculation, applied to the coefficient vectors of the planes defining the $\mathbf{X}$-subspace spanned by the face. The fact that the angle of rotation has the form $2\pi/m$ can be checked by seeing whether the m-th power of a certain group element is the identity. We can see approximately which values of m to use by means of floating point arithmetic.

Finally we indicate circumstances under which it seems that a Turing machine could do all the relevant checks. Suppose we are given a finite set of matrices each of which is an isometry of $\mathbf{X}^n$, and such that each entry is an algebraic number. We can hold the algebraic numbers in the computer by holding the coefficients of its irreducible polynomial, together with a floating point approximation to the number. Suppose we are also given a finite set of finite-sided polyhedra, given approximately using floating point numbers, together with face-pairings each of which is equal to one of our given matrices. We can then check the condition Cyclic$(\mathscr{P}, R, A)$ precisely, using integer arithmetic, by checking on a certain product of face-pairings. (We can use floating point arithmetic to see which words in the face-pairings need to have checks performed.)

8. Special cases

One case of Poincaré's Theorem which is often used is the case where there is a single element of $\mathscr{P}$ and all face-pairings are reflections. In that case completeness is a consequence of Lemma 5.4, provided the other axioms are satisfied. This enables a number of important examples to be constructed.

As a minor point, we note that it enables us to construct infinitely generated fuchsian groups with an arbitrary subset of the positive integers being the set of exponents of maximal cyclic subgroups. These and other applications of Poincaré's Theorem are well-known.

Poincaré's Theorem works in an especially simple way in dimension two. In this dimension, a face-pairing is called an *edge-pairing*. The following result is essentially due to de Rham [dR71].

THEOREM 8.1 (dimension two). *Suppose we have a finite set $\mathscr{P}$ of finite-sided polygons in $\mathbf{H}^2$ and an edge-pairing (R, A) of the boundary edges satisfying Pairing$(\mathscr{P}, R, A)$, Connected$(\mathscr{P}, R)$ and Cyclic$(\mathscr{P}, R, A)$. Then the quotient Q of $\bigsqcup_{P \in \mathscr{P}} P$ by the edge-pairing is a two-dimensional hyperbolic orbifold which is obtained from a complete orbifold with geodesic boundary by removing the compact boundary components. The hyperbolic structure on Q is induced in an obvious way from the hyperbolic structure on the hyperbolic polygons used to define it. The group G generated by the edge-pairings in the manner described in Definition 4.2 is discrete. If all the polygons are compact, then Q is a compact orbifold without boundary. (But it may have mirrors.)*

REMARK 8.2. The main feature of this result is that for the two-dimensional case it describes the quotient Q even when this is not complete. (For the conditions under which Q is complete the reader is referred to Lemma 5.4 and Theorem 6.3.)

REMARK 8.3. As an orbifold, the boundary of Q is empty, since each edge is paired to some other. If an edge is paired to itself by a reflection, Q has a corresponding mirror. This is a (possibly non-compact) boundary component of the underlying manifold, but not an orbifold boundary component. The completion of Q may well have orbifold boundary components which are not in Q itself: each of these is a circle.

Proof of 8.1. We first look at each ideal point p which is the end of two distinct boundary components of some $P \in \mathscr{P}$ (in this situation we will say p is a peak of P). We remove from P a small horodisk neighbourhood centred at p. If we glue the remaining pieces together using the edge-pairings, then the boundary horocycles do not necessarily match up and we obtain a hyperbolic orbifold T whose boundary is a union of topological circles and arcs, each of which is a finite union of horocyclic and geodesic arcs.

We now need to glue back the pieces we have cut out. We first glue together the horodisk pieces corresponding to a single boundary component of T, obtaining an orbifold B_i, and understand what B_i looks like. This can be done by using the horocyclic foliation of B_i.

Each piece constituting B_i is triangular, where two of the sides are geodesic rays which are asymptotic and one side is a horocyclic arc. When we glue these pieces together, several things can happen.

(a) The pieces glue in a cyclic fashion and the associated holonomy is parabolic. Then B_i is complete. Gluing B_i into place gives rise to a cusp in Q.

(b) The pieces glue in a cyclic fashion and the associated holonomy is hyperbolic. Then the developing image into the upper half-plane is as shown in Figure 13. In this case B_i is a cylinder, which at one end is incomplete, with the completion adding a compact geodesic, and at the other end is bounded by alternate horocyclic and geodesic arcs. (Actually, a suitable choice of the pieces at the beginning allows to reduce to the case of one geodesic arc and one horocyclic arc.) Gluing B_i into place gives an incomplete end for Q. The completion adds to Q a compact geodesic boundary component.

(c) The pieces do not glue together in a cyclic fashion. This means that in both directions one eventually reaches a geodesic ray which is glued either to itself (by a reflection) or to a geodesic ray whose point at infinity is not a peak.

As a set, B_i is identified to a triangle in $\mathbf{H}^2$ with one vertex at infinity, two sides which are asymptotic geodesic half-lines and the third

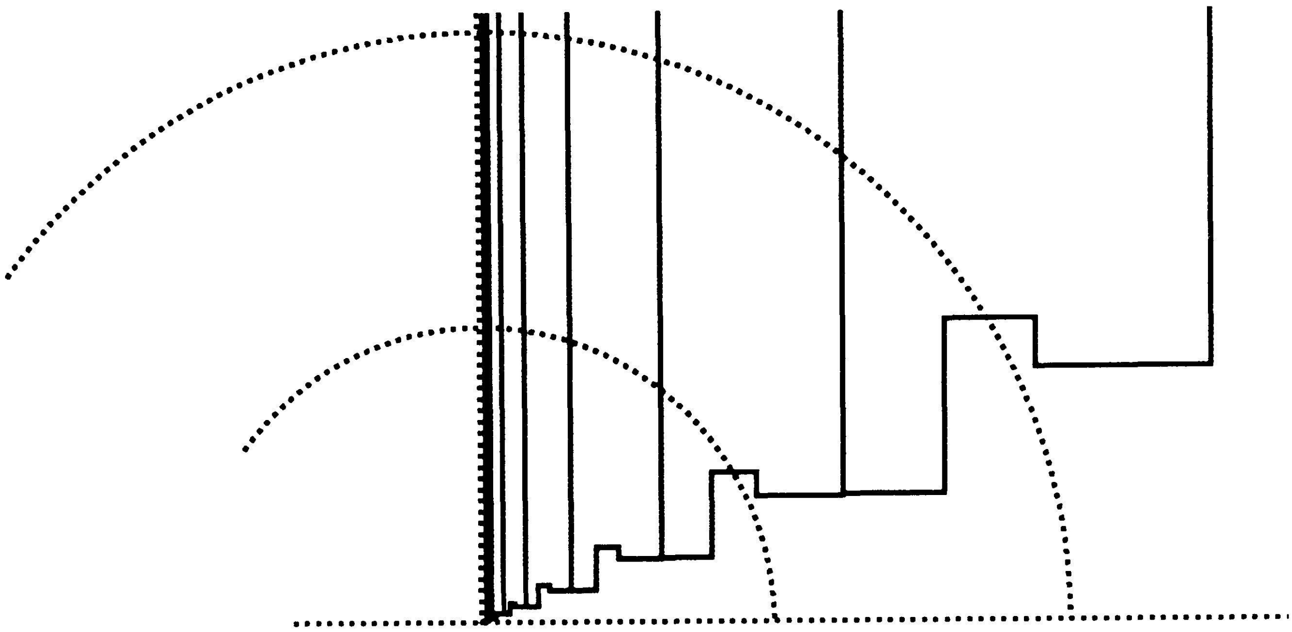

FIGURE 13.
Hyperbolic holonomy.
This picture shows the developing image associated with B_i in the situation described in (b) in the proof of 8.1. The dotted lines show an alternative fundamental domain which shows the structure of B_i as a cylinder more clearly.

side which consists of alternate horocyclic and geodesic arcs (a suitable choice of the pieces at the beginning actually allows to obtain only one horocyclic arc and no geodesic arcs). As a hyperbolic orbifold, B_i is either isomorphic to the triangle or obtained from the triangle by assuming that one or both the geodesic half-lines are mirrors. In particular, B_i is complete. The situation is shown in Figure 14.

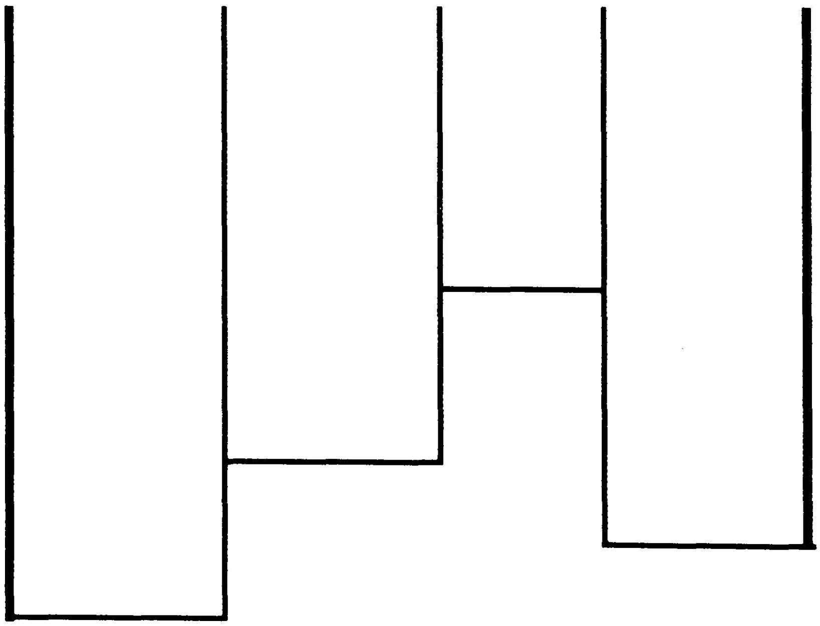

FIGURE 14.
Mirrors.
One or both the vertical thick edges may represent mirrors, as described in (c) in the proof of 8.1.

A consistent horocycle always exists in this case, so gluing B_i into place gives us a metric in Q which is complete near B_i. We can also describe what Q looks like near B_i. If we call wedge a region in $\mathbf{H}^2$ bounded by two geodesic half-lines with common origin, then a subset of Q which contains B_i is obtained as follows: for every half-infinite geodesic side s of B_i which is not a mirror, glue a suitably thin wedge to B_i by identifying s with one of the sides of the wedge.

This description of the possible situations implies the conclusion of the proof. □

9. Literature review

It seems to the authors that a minimal requirement for a satisfactory treatment of Poincaré's Theorem is that it should apply directly to the case of a finite-sided Dirichlet domain resulting from the action of a discontinuous group of isometries on one of the three constant curvature geometries $\mathbf{S}^n$, $\mathbf{E}^n$ and $\mathbf{H}^n$. Furthermore the hypotheses should be easy to verify, and extraneous hypotheses should not be included. We review the literature with these criteria in mind.

The first versions of Poincaré's Theorem were published in [Poi82], covering the two-dimensional version, and [Poi83], covering the three-dimensional version. These are reprinted in Volume Two of [Poi52]. It is clear that Poincaré understood very well what was going on. However, the papers are not easy to read. In particular, the reader of the three-dimensional case is referred to the treatment of the two-dimensional case for proofs; this is fully acceptable for a trail-blazing paper, but not satisfactory in the long term.

There are a number of reasonable published versions of Poincaré's Theorem in dimension two. Of these, we would single out the version by de Rham [dR71] as being particularly careful and easy to read. Most published versions of Poincaré's Theorem applying to all dimensions are unsatisfactory for one reason or another. The most satisfactory version is [Sei75], due to Seifert. The proofs are careful and rigorous, but rather long. Poincaré's Theorem is proved in all dimensions and for all three constant curvature geometries. The treatment is not constructive in several aspects, specially when it comes to completeness. There is some discussion of conditions which are equivalent to completeness in the hyperbolic case, which are closer to being constructive. However this discussion is limited to the finite volume case. Seifert's treatment also contains unnecessary restrictions, which, for example, would prevent his version being directly applicable to the Dirichlet domain applied to a rotation through π about a fixed point in dimension two. (One would first have to subdivide the boundary of the Dirichlet domain, since Seifert assumes that the map to the quotient space is injective on each face of the given polyhedron.)

The treatment in [Mas88] is difficult to understand. For example in H.9 on page 75, it is claimed that a metric is defined in a certain way, and this fact is said to be "easy to see", but it seems to us an essential and non-trivial point, which is not so easy to see, particularly when the group generated by the face-pairings is not discrete. Maskit's proof does not use induction on dimension, which seems to us essential for a simple and clear treatment. We refer in particular to the assertions that certain maps are homeomorphisms on page 77. The Proposition in IV.1.6 on page 79 of this book is incorrect — a counter-example is given in Example 9.1 — because there are no infinite cycles or infinite edges according to the definitions in the book. As in the case of Seifert's paper, the constructive aspect is ignored, and the question of completeness is handled in an entirely non-constructive way. Maskit's local finiteness condition is more demanding than ours, and Seifert's is more demanding than Maskit's.

In [Ril83], there is a statement of Poincaré's Theorem with no proof, and [Sei75] is cited. Unfortunately, Riley fails to take into account Seifert's restriction to the finite volume case. This leads him to a statement of Poincaré's Theorem, which implies that if two parallel vertical planes in upper half-space are matched by a hyperbolic isometry, then the infinite cyclic group thus generated is discrete.

Maskit's paper [Mas71] contains a nice discussion of completeness, though again it is not a constructive approach. He limits his discussion to hyperbolic space in dimensions two and three. We are not confident that the arguments in the paper are complete. For example, there seems to be an assumption that the quotient of a metric space, such that the inverse image of any point is finite, is again metric. This is false, as is shown by identifying x with $-x$ in $[-1, 1]$, provided $0 \leqslant x < 1$. A slight variation of this gives a counter-example in which the inverse image of a point is always equal to two points.

EXAMPLE 9.1 (incomplete example). Take a quadrilateral in the euclidean plane with no two sides parallel, and multiply with $(0, \infty)$. Embed this in the upper half-space model of $\mathbf{H}^3$, with the quadrilateral embedded in a horizontal horosphere, and the factor $(0, \infty)$ corresponding to vertical straight lines. This gives us a convex hyperbolic polyhedron P with four faces.

The quadrilateral gives rise to two commuting orientation-preserving euclidean similarities which identify opposite sides. These similarities can be regarded as hyperbolic isometries which are face-pairings for P. They do not generate a discrete group of isometries of $\mathbf{H}^3$. Maskit's paper [Mas71] and his book [Mas88] both contain statements implying that this group of isometries is discrete.

There is a discussion of Poincaré's Theorem in Beardon's paper [Bea83]. Beardon concentrates on $\mathbf{H}^2$, with a single compact convex polygon. Questions of completeness are not treated.

Morokuma's paper [Mor78] is another paper which is difficult to read. If the definitions in this paper are taken literally, then the statement of the main theorem implies that a closed ball of finite radius is equal to the whole of hyperbolic space. This is because a closed ball is the intersection of a collection of half-spaces, each containing the ball in its interior, and as a consequence a closed ball is a polyhedron with no faces. The paper contains a great deal of notation, which, to our way of thinking, obscures the ideas. On occasion the author seems to assume the main point of what needs to be proved. For example, on page 163 of his article, the statement "$\tau^{-1}p' \in F'_{k+1}$ namely $F'_{k+1} = F'$" would not be true if Morokuma's group Γ were not discrete. But at this point he is trying to prove discreteness.

Apanasov's paper [Apa86] is yet another paper which is difficult to read. Apanasov allows non-convex polyhedra. To see the consequences of Apanasov's definitions, consider the Poincaré disk model for $\mathbf{H}^2$. According to his definitions, the union of the closed first and third quadrants is a polyhedron with two one-dimensional faces, namely the x- and y-axes. There are no codimension-two faces. The intersection of two faces of a polyhedron does not need to be a face. It is not clear to us what is meant by Condition IV on page 474 of the English translation of Apanasov's paper. As a general comment on this paper, it seems as though much of what one should prove in Poincaré's Theorem are presented as hypotheses, rather than as conclusions.

An earlier paper by Aleksandrov, [Ale54], also makes many parts of Poincaré's Theorem into hypotheses rather than conclusions.

A proof of Poincaré's theorem in the special case of a single polyhedron with each face-pairing equal to the reflection in that face is given in [dlH91]; this proof has the same inductive structure as the proof given in our paper. The only condition to check is that the angles at codimension 2 faces have the form π/m for some integer m. This version of Poincaré's theorem is readily deduced from Theorem 5.5 using 5.4; in fact LocallyFinite is obvious in this case and the quotient space is complete as it is identified with the polyhedron itself.

10. Appendix

This appendix contains results due to Brian Bowditch, published here with his permission.

We recall that a finite-sided closed convex cell of $\mathbf{H}^{n+1}$ is said to be *pyramidal* at an ideal point p if any two faces whose closures contain p meet in $\mathbf{H}^{n+1}$. The intersection of such a convex cell with a horosphere centred at p is a euclidean finite-sided closed convex cell of dimension n (provided the horosphere only meets faces which have p as an ideal point). One way to see this is to use the upper half-space model with p equal to the point at infinity. Conversely, given a convex finite-sided n-dimensional euclidean cell, we can think of this cell as lying in a horosphere which is a horizontal subspace in the upper half-space model. This gives rise to an $(n+1)$-dimensional hyperbolic convex cell, by taking the intersection of vertical half-spaces determined by the half-spaces defining the euclidean convex cell. We use the names "pyramidal" and "non-pyramidal" for convex euclidean cells if the corresponding hyperbolic cells are pyramidal or non-pyramidal respectively. A euclidean convex cell is non-pyramidal if and only if it has disjoint faces. If a euclidean cell is pyramidal, then there is a face which is the intersection of all other faces, that is there is a unique minimal face. A pyramidal euclidean n-cell is the product of an i-dimensional cell with the cone on a spherical $(n-i-1)$-dimensional cell. (The cone point is placed at the centre of the $(n-i-1)$-dimensional sphere.)

Let M be a connected euclidean similarity n-dimensional manifold which is the union of a locally finite set of closed subsets $\{X_i\}$. Each X_i has an induced similarity structure which is isomorphic to that of a closed finite-sided euclidean convex polyhedron. There are only a finite number of distinct similarity classes of X_i. The intersection of any face of any X_i with any face of any X_j is a common face of each. This implies that M has the structure of a locally finite polyhedral cell complex. Let G be a group of similarities of M which preserve the cell structure. Suppose that the number of orbits of non-pyramidal polyhedra is finite.

Theorem 10.1 (Bowditch). *Under the above assumptions, the number of orbits of cells is finite. Moreover, the number of orbits is bounded in terms of the number of orbits of non-pyramidal cells and the geometry (up to similarity) of the* X_i.

Bowditch has suggested that if there is one or more pyramidal polyhedral cell, then one should be able to prove that G is a finite group. It would follow that G consists of euclidean isometries and that M contains only a finite number of cells. This conjecture remains open.

Proof of 10.1. Let X be the union of the non-pyramidal cells in M, and let Y be the union of cells which meet X. Note that $X \subset Y$.

Now suppose there is a top-dimensional cell which is not in Y and let σ be its unique minimal face. Then σ is similar to $\mathbf{R}^i$ for some i. If α is any cell meeting σ, then $\sigma \subset \alpha$ since σ is minimal. Clearly α is not in X. Therefore σ is the unique minimal face of α. We have seen above that α is the product of σ and the cone on a convex subset $\mathbf{S}^{n-i-1}$. It follows that the union of the cells meeting σ is the product of σ with the cone on $\mathbf{S}^{n-i-1}$. It follows that the cell structure of M is finite, G is a finite group and $X = \varnothing$. The other possibility is that $Y = M$.

Let $K \subset X$ be a finite union of cells such that $GK = X$. The cell structure of M is locally finite, with a bound for the number of cells in any small neighbourhood being given by the geometry of the X_i. The number of cells of M which meet K is bounded by the number of cells of K and the maximum possible number of cells meeting a fixed small neighbourhood of any fixed point of K. This gives an upper bound for the number of orbits of cells of M under the action of G in case $Y = M$. If $X = Y = \varnothing$, then the number of cells of M is bounded by the geometry of the X_i. $\square$

We apply Theorem 10.1 to find out a litte more about the spaces that arise in Poincaré's Theorem. Suppose the hypotheses Pairing($\mathscr{P}, R, A$), Connected($\mathscr{P}, R$), Finite($\mathscr{P}$) and Cyclic($\mathscr{P}, R, A$) are satisfied for a set of convex cells (see Definition 2.8) in $\mathbf{H}^n$. To each convex cell we adjoin the ideal points, so as to obtain a compact space. The face-pairings are defined on the closures of the faces. Let $\bar{Q}$ be the quotient of the disjoint union of the extended cells by the face-pairings, endowed with the quotient topology.

THEOREM 10.2. *$\bar{Q}$ is a compact hausdorff space.*

Proof of 10.2. Let X be the disjoint union of the closures of the convex cells. So X is compact and hausdorff. We first show that the inverse image of a point under the quotient map $X \to \bar{Q}$ is a finite set. This is clear from Theorem 4.13 for any point which is not an ideal point. For an ideal point p, we can construct a similarity manifold to which Theorem 10.1 applies, by developing a horosphere centred at p into $\mathbf{R}^{n-1}$. More details, which will help the interested reader with the construction of the similarity manifold, are given in the discussion of Definition 6.2.

A pyramidal cell in $\mathbf{R}^{n-1}$ corresponds to a convex cell in $\mathbf{H}^n$ together with an ideal point p in its boundary, such that any two faces with closures containing p meet inside $\mathbf{H}^n$. A non-pyramidal cell corresponds to a convex cell in $\mathbf{H}^n$ and an ideal point p contained in the closures of two non-intersecting faces of the convex cell. The hypothesis needed in order to apply Theorem 10.1, that there are only a finite number of orbits of non-pyramidal cells, comes from the fact that there are only a finite number of pairs of faces and therefore only a finite number of pairs of non-intersecting faces which meet at infinity.

It follows that the inverse image in X of any point of $\bar{Q}$ is finite. Moreover the number of points in the inverse image is bounded by a fixed integer N. Two points $x, y \in X$ are mapped to the same point of $\bar{Q}$ if and only if there is a sequence $(x_0, \ldots, x_n)$ such that $x = x_0$, $y = x_n$ and $x_{i+1} = A(F_i)(x_i)$, where $x_i \in F_i$ and $x_{i+1} \in R(F_i)$. (Here (R, A) is the glueing data.) We may take $n \leqslant N$. It follows easily from compactness and the finiteness of the situation that the map $X \to \bar{Q}$ is closed. Therefore $\bar{Q}$ is hausdorff. $\square$

REFERENCES

[Ale54] ALEKSANDROV, A.D. Filling space by polyhedra. *Vestnik. Leningrad Univ. Math. 2* (1954), 33-34.

[Apa86] APANASOV, B.N. Filling a space by polyhedra and deformation of incomplete hyperbolic structures. *Siberian Math. J. 27* (1986), 473-485. (English translation.)

[Bea83] BEARDON, A.F. *The Geometry of Discrete Groups.* Springer-Verlag, 1983.

[Bow93] BOWDITCH, B. Geometrical finiteness for hyperbolic groups. *J. Funct. Anal. 113* (1993), 245-317.

[BP92] BENEDETTI, R. and C. PETRONIO. *Lectures on hyperbolic geometry.* Springer-Verlag, 1992.

[BSS89] BLUM, L., M. SHUB and S. SMALE. On a theory of computation and complexity over the real numbers: NP-completeness, recursive functions and universal machines. *Bull. Amer. Math. Soc. 21* (1989), 1-46.

[dlH91] DE LA HARPE, P. An invitation to Coxeter groups. In: *Group theory from a geometrical viewpoint,* Ghys-Haefliger-Verjovsky editors, World Scientific Publishers, Singapore, 1991.

[dR71] DE RHAM, G. Sur les polygones générateurs de groupes fuchsiens. *Enseign. Math. 17* (1971), 49-61.

[Mas71] MASKIT, B. On Poincaré's theorem for fundamental polygons. *Adv. Math. 7* (1971), 219-230.

[Mas88] —— *Kleinian Groups.* Springer-Verlag, 1988.

[Mor78] MOROKUMA, T. A characterization of fundamental domains of discontinuous groups acting on real hyperbolic spaces. *J. Fac. Sci. Univ. Tokyo Section 1A Math. 25* (1978), 157-183.

[Poi82] POINCARÉ, H. Théorie des groupes fuchsiens. *Acta Math. 1* (1882), 1-62.
[Poi83] —— Mémoire sur les groupes Kleinéens. *Acta Math. 3* (1883), 49-92.
[Poi52] —— *Collected Works*. Gauthier-Villars, 1952.
[Ril83] RILEY, R. Applications of a computer implementation of Poincaré's theorem on fundamental polyhedra. *Math. Comp.* (1983), 607-632.
[Sei75] SEIFERT, H. Komplexe mit Seitenzuordnung. *Göttinger Nachrichten* (1975), 49-80.
[Thu] THURSTON, W.P. The geometry and topology of three-dimensional manifolds. (In preparation.)
[Thu80] —— *The Geometry and Topology of Three-Manifolds.* Princeton University Mathematics Department, 1980. (Thurston's original notes.)

(Reçu le 10 novembre 1993)

David B.A. Epstein
University of Warwick (Coventry, England)
and Geometry Center (Minneapolis, USA)

Carlo Petronio
Università di Pisa (Italy)