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 $D_Z \colon Z(\mathscr{P}, R, A) \to \mathbf{X}^n$ 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. $\quad\square$

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). $\quad\square$

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. $\square$

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 $\pi_{ZQ}(z_0)$. $\square$

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