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

Kapitel: 6. Completeness

DOI: https://doi.org/10.5169/seals-61108

Nutzungsbedingungen

Die ETH-Bibliothek ist die Anbieterin der digitalisierten Zeitschriften auf E-Periodica. Sie besitzt keine Urheberrechte an den Zeitschriften und ist nicht verantwortlich für deren Inhalte. Die Rechte liegen in der Regel bei den Herausgebern beziehungsweise den externen Rechteinhabern. Das Veröffentlichen von Bildern in Print- und Online-Publikationen sowie auf Social Media-Kanälen oder Webseiten ist nur mit vorheriger Genehmigung der Rechteinhaber erlaubt. Mehr erfahren

Conditions d'utilisation

L'ETH Library est le fournisseur des revues numérisées. Elle ne détient aucun droit d'auteur sur les revues et n'est pas responsable de leur contenu. En règle générale, les droits sont détenus par les éditeurs ou les détenteurs de droits externes. La reproduction d'images dans des publications imprimées ou en ligne ainsi que sur des canaux de médias sociaux ou des sites web n'est autorisée qu'avec l'accord préalable des détenteurs des droits. En savoir plus

Terms of use

The ETH Library is the provider of the digitised journals. It does not own any copyrights to the journals and is not responsible for their content. The rights usually lie with the publishers or the external rights holders. Publishing images in print and online publications, as well as on social media channels or websites, is only permitted with the prior consent of the rights holders. Find out more

Download PDF: 02.12.2025

ETH-Bibliothek Zürich, E-Periodica, https://www.e-periodica.ch

THEOREM 5.5 (Poincaré's Theorem Version 2). Suppose the hypotheses $Pairing(\mathcal{P}, R, A)$, $Connected(\mathcal{P}, R)$, $Cyclic(\mathcal{P}, R, A)$ and $Locally-Finite(\mathcal{P}, R, A)$ 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. \square

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 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 f at p. Otherwise we say that p is p is thin at p. If p is thin at p, it must have two faces p and p whose closures p and p meet in p only. In the Appendix, in this situation we refer to p as being p and p and p and p as being p and p and p and p and p as being p and p and p and p and p and p are the p and p and p and p are the p and p and p and p are the p and p and p are the p and p and p are the p are the p are the p and p are the p are the p and p are the p are the p and p are the p are the p are the p are the p and p are the p are the

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 $\operatorname{vol}(S^{n-2})/N$.

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 Q be the quotient space of the disjoint union of the 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 \mathbb{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 $\operatorname{Cyclic}(\mathcal{P},R,A)$ 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 $D_p: \mathbb{Z}_p \to \mathbb{H}^n$ 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 Q. If there is a consistent horosphere at p, then the image of G_p consists entirely of euclidean isometries of \mathbb{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 \mathscr{P} of hyperbolic convex cells satisfying Pairing(\mathscr{P} , R, A), Connected(\mathscr{P} , R), Cyclic(\mathscr{P} , R, A) and Finite(\mathscr{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 \overline{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 $P \in \mathcal{P}$. 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 $P \in \mathcal{P}$, the volume of the impression of P at p is either zero or is greater than $\operatorname{vol}(S^{n-2})/N$. 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 $A = \alpha \alpha' \alpha'' \dots$ 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 Q is hausdorff and that each point of \overline{Q} has a neighbourhood whose intersection with Q is complete. Under the conditions stated, G_p acts as a group of isometries of \mathbb{R}^{n-1} . From Poincaré's Theorem applied to \mathbb{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 [Bow 93] — 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 halfspace 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 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 $\overline{Q} \setminus Q$. 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 $P \in \mathcal{P}$, \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 $P \in \mathcal{P}$, $\Gamma_{P,p}$, 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.

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