UNIVERSITÉ DE PROVENCE

Année 2004

N^o attribué par la bibliothèque:

THÈSE

pour obtenir le grade de

DOCTEUR DE L'UNIVERSITÉ DE PROVENCE

Discipline: *Mathématiques*

École doctorale Mathématiques et Informatique de Marseille - E.D. n^o 184

présentée et soutenue publiquement par

Roland Karl Walter ROEDER

le 21 Mai, 2004

Titre:

Le Théorème d'Andreev sur polyèdres hyperboliques

JURY:

M. John Hamal HUBBARD, *Université de Provence et Cornell University*, Directeur de thèse
M. Adrien DOUADY, Université Paris-Sud 11, President
M. Jean-Pierre OTAL, *Université de Lille 1*, Rapporteur
M. Michel BOILEAU, *Université Paul Sabatier*, Rapporteur
M. Hamish SHORT, *Université de Provence*
M. Jérôme LOS, *Université de Provence*

Remerciements ~ Thanks

I would like to thank first of all John Hamal Hubbard, my thesis director for getting me interested and involved in the beautiful subject of hyperbolic geometry. Quite some time ago, he suggested that I write a computer program that impliments Andreev's Theorem using the ideas from Andreev's proof. I was given a version of this proof written by J. H. Hubbard and Bill Dunbar as a starting point. From the ideas in that manuscript, I successfully wrote a program following Andreev's proof, which *unsuccessfully* computed hyperbolic polyhedra. From here, I found the error in Andreev's proof and was able to correct the program so that it worked and I was able eventually to write these thesis consisting of a correct proof of Andreev's Theorem.

However, all of the work that I have done has come from ideas that I learned from that manuscript, and this thesis followes exactly the same outline of proof. Furthermore, sections 2, 3, and 4 consist of only relatively minor modifications and improvements of work that I know Hubbard and Dunbar worked very hard to achieve.

I thank both John Hubbard and Bill Dunbar for their work on that preliminary manuscript, and for their subsequent help and interest in my corrections of the proof. I also thank John Hubbard for spending hours upon hours reading this thesis and making suggestions for improvement, both mathematical and written.

I thank the rapoteurs Jean-Pierre Otal, Michel Boileau, and Mark Baker for taking the time to read this thesis and for their interest in it. I also thank the remaining members of my jury, Adrien Douady, Hamish Short, and Jérôme Los for their time and interest in the thesis.

I also thank the other members of the mathematical community in Marseille who have taken the time to help me and to share their ideas with me. In particular, I thank Pascal Hubert for inviting me to give a seminar about my research at LUNINY.

I thank the Department of Defense of the United States for supporting the first part of my graduate studies with a National Defense Science and Engineering Fellowship. I also thank the National Science Foundation of the United States for their support of my studies this year in Marseille via an Integrative Graduate Research and Training (IGERT) fellowship.

On the personal side, I thank my parents Stephen and Phoebe Roeder for their love and support from the beginning and my sister Adrienne Roeder for her support. I would also like to thank my girlfriend Johanna Kaufman for her love and support.

I also thank the Hubbard family for helping to smooth out the difficulties of living abroad, for their many invitations to dinner, and for inviting me on a delightful skiing trip to the Alps.

À la mer, la vent et les vagues,
à tous ces bonnes choses.

Résumé

E. M. Andreev a publié en 1970 une classification des polyèdres hyperboliques compacts de dimension trois dont les angles dièdres sont non-obtus [3]. Etant donné une description combinatoire d’un polyèdre C, le Théorème d’Andreev dit que les angles dièdres possibles sont exactement décrits par cinq classes d’inégalités linéaires. Le Théorème d’Andreev démontre également que le polyèdre résultant est alors unique à isométrie hyperbolique près.

D’une part, le Théorème de Andreev est évidemment un énoncé intéressant de la géométrie de l’espace hyperbolique en dimension 3; d’autre part c’est un outil essentiel dans la preuve du Théorème d’Hyperbolization de Thurston pour les variétés Haken de dimension 3. Il est d’ailleurs remarquable à quel point la démonstration d’Andreev rappelle (en plus simple) la démonstration de Thurston.

La démonstration d’Andreev contient une erreur importante. Nous corrigeons ici cette erreur et nous fournissons aussi une nouvelle preuve lisible des autres parties de la preuve, car le papier d’Andreev a la réputation d’être “illisible”. Nous fournissons aussi une classification des tétraèdres hyperboliques; c’est un cas particulier qui n’est pas couvert par le Théorème d’Andreev, et effectivement le résultat est assez différent, car dans ce cas l’ensemble des angles dièdres possibles n’est pas convexe, même dans le cas où ces angles sont non-obtus.

Discipline: *Géométrie*

MOTS CLES: géométry hyperbolique, polyèdre hyperbolique, groups Kleiniens.

Andreev’s Classification of Hyperbolic Polyhedra

Abstract

In 1970, E. M. Andreev published a classification of all three dimensional compact hyperbolic polyhedra having non-obtuse dihedral angles [3]. Given a combinatorial description of a polyhedron, C, Andreev’s Theorem provides five classes of linear inequalities, depending on C, for the dihedral angles, which are necessary and sufficient conditions for the existence of a hyperbolic polyhedron realizing C with the assigned dihedral angles. Andreev’s Theorem also shows that the resulting polyhedron is unique, up to hyperbolic isometry.

Andreev’s Theorem is both an interesting statement about the geometry of hyperbolic 3 dimensional space, as well as a fundamental tool used in the proof for Thurston’s Hyperbolization Theorem for 3 dimensional Haken manifolds. It is also remarkable to what level the proof of Andreev’s Theorem resembles (in a simpler way) the proof of Thurston.

We correct a fundamental error in Andreev's proof of existence and also provide a readable new proof of the other parts of the proof of Andreev's Theorem, because Andreev's paper has the reputation of being "unreadable". We also provide a classification of hyperbolic tetrahedra which is a special case that is not covered by Andreev's Theorem, and effectively a different manner of result, because in this case the set of possible dihedral angles is non-convex, even in the case where these angles are non-obtuse.

Disciplin: *Mathématiques*

KEY WORDS: hyperbolic geometry, hyperbolic polyhedra, Klienian groups.

Laboratoire d'Analyse, Topologie, Probabilités
UMR 6632
Centre de Mathématiques et Informatique
Université de Provence,
39 rue F. Joliot-Curie,
13453 Marseille cedex 13

Table des matières

Contents

1 Statement of Andreev's Theorem

Andreev's Theorem provides a complete characterization of compact hyperbolic polyhedra having non-obtuse dihedral angles. This classification is essential for proving Thurston's Hyperbolization theorem for Haken 3-manifolds and is also a particularly beautiful and interesting result in its own right. Complete and detailed proofs of Thurston's Hyperbolization for Haken 3-manifolds are available written in English by Jean-Pierre Otal [7] and in French by Michel Boileau [4].

In this paper, we prove Andreev's Theorem based on the main ideas from his original proof [3]. However, there is an error in Andreev's proof of existence. We explain this error in Section 6 and provide a correction. Although the other parts of the proof are proven in much the same way as Andreev proved them, we have re-proven them and re-written them to verify them as well as to make the overall proof of Andreev's Theorem clearer. Andreev's original proof has a reputation for being difficult to follow. We also include a classification of hyperbolic tetrahedra in the last section because their classification is not provided by Andreev's Theorem.

The reader may also wish to consider the three other similar results of Rivin and Hodgeson [8, 5], Thurston [9], and Marden and Rodin [6]. In [8], the authors prove a more general statement than Andreev's Theorem and in [5] Hodgeson deduces Andreev's Theorem as a consequence of their previous work. The proof in [8] is similar to the one presented here, except that the conditions classifying the polyhedra are written in terms of measurements in the De Sitter space, the space dual to the hyperboloid model of hyperbolic space. Although a beautiful result, the main drawback of this proof is that the last sections of the paper, which are necessary for their proof that such polyhedra exist, are particularly hard to follow.

The works of Marden and Rodin [6] and Thurston [9] consider configurations of circles with assigned overlap angles on the sphere and on surfaces of genus g with $g > 0$. Such a configuration of overlapping circles in the sphere corresponds directly to a configuration of hyperbolic planes in the conformal ball model of hyperbolic space. Thus, there is a direct connection between circle patters and hyperbolic polyhedra. The proof of Thurston [9] provides a classification of configurations of circles on surfaces of genus $g > 0$. The proof of Marden and Rodin [6] is an adaptation for the sphere of Thurston's circle packing approach and results in a theorem similar to Andreev's Theorem, but which is phrased entirely in terms of configurations of circles. Although Thurston's proof allows for arbitrary non-obtuse overlap angles between adjacent circles, the proof of Marden and Rodin limits certain angle sums, so that the patterns of overlapping circles derived cannot correspond to compact, hyperbolic polyhedra, the items classified by Andreev's Theorem.

We begin by defining hyperbolic 3-space and hyperbolic polyhedra. Then we will discuss the combinatorial properties of hyperbolic polyhedra and state Andreev's theorem.

The hyperboloid model of hyperbolic space

There are many models of hyperbolic n-dimensional space $\mathbb{H}^n$, each of which is isometrically isomorphic to the others. We define the hyperboloid model of hyperbolic n-dimensional space as the component of the subset of $\mathbb{R}^{n+1}$ given by the equation:

$$-x_0^2 + (x_1^2 + x_2^2 + \cdots + x_n^2) = -1$$

having $x_0 > 0$, with the Riemannian metric induced by the indefinite metric

$$-dx_0^2 + dx_1^2 + dx_2^2 + \cdots + dx_n^2.$$

We will denote the hyperboloid model of $\mathbb{H}^n$ by $\mathbb{H}^n_H$.

Calculations in $\mathbb{R}^n$ with the indefinite metric $-x_0^2 + (x_1^2 + x_2^2 + \cdots + x_n^2)$ will often be necessary, so we denote this space by $\mathbb{E}^{n,1}$. The majority of calculations in the hyperboloid model are actually done in $\mathbb{E}^{n,1}$.

Hyperbolic k-dimensional subspaces of $\mathbb{H}^n_H$ are defined to be the intersections $V \cap \mathbb{H}^n_H$, where V is a $(k+1)$-dimensional vector subspace of $\mathbb{E}^{n,1}$ that intersects $\mathbb{H}^n_H$. Notice that a k-dimensional subspace of $\mathbb{H}^n_H$ is canonically isomorphic to $\mathbb{H}^k_H$. Hyperbolic one and two dimensional subspaces will be referred to as hyperbolic lines and planes.

Hyperbolic space $\mathbb{H}^n_H$ can clearly be compactified by adding the set of rays to the lightcone C, which themselves clearly form a topological space $\partial\mathbb{H}^n_H$ homeomorphic to the sphere $\mathbb{S}^{n-1}$. We will refer to points in $\partial\mathbb{H}^n_H$ as *points at infinity* and refer to the compactification as $\overline{\mathbb{H}^n_H}$.

Throughout this paper we will be primarily interested in three dimensional hyperbolic space because will study 3-dimensional hyperbolic polyhedra.

One can check that the hyper-plane orthogonal to a vector $\mathbf{v} \in E^{3,1}$ intersects $\mathbb{H}^3_H$ if and only if $\langle \mathbf{v}, \mathbf{v} \rangle > 0$. Let $\mathbf{v} \in E^{3,1}$ be a vector with $\langle \mathbf{v}, \mathbf{v} \rangle > 0$, and define

$$P_{\mathbf{v}} = \{\mathbf{w} \in \mathbb{H}^3_H | \langle \mathbf{w}, \mathbf{v} \rangle = 0\}$$

to by the hyperbolic plane orthogonal to $\mathbf{v}$; and the corresponding closed half space:

$$H^+_{\mathbf{v}} = \{\mathbf{w} \in \mathbb{H}^3_H | \langle \mathbf{w}, \mathbf{v} \rangle \geq 0\}.$$

Notice that given two planes $P_{\mathbf{v}}$ and $P_{\mathbf{w}}$ in $\mathbb{H}^3_H$ with $\langle \mathbf{v}, \mathbf{v} \rangle = 1$ and $\langle \mathbf{w}, \mathbf{w} \rangle = 1$, they:

- intersect in a line if and only if $\langle \mathbf{v}, \mathbf{w} \rangle^2 < 1$, in which case their dihedral angle is $\arccos(-\langle \mathbf{v}, \mathbf{w} \rangle)$.
- intersect in a single point at infinity if and only if $\langle \mathbf{v}, \mathbf{w} \rangle^2 = 1$, in this case their dihedral angle is 0.

A *hyperbolic polyhedron* is an intersection

$$P = \bigcap_{i=0}^{n} H^{+}_{\mathbf{v_i}}$$

having non-empty interior.

Although the hyperboloid model of hyperbolic space is very natural, it is not easy to visualize, since the ambient space is four dimensional. At least three other models of hyperbolic space are in common use: the *Klein model*, the *conformal ball model*, and the *upper half-space model*. Each of these models has its own merits and its own weaknesses. A given theorem or calculation may be difficult in one model, while obvious in a different model. Within this paper, we will only use the hyperboloid model, the conformal ball model, and the upper half-space model.

The conformal ball model of hyperbolic space

Let B^n by the n-dimensional ball given by $x_1^2 + \cdots + x_n^2 < 1$ in the plane $x_0 = 0$ of $E^{n,1}$. Let $\pi : \mathbb{H}^n_H \to B^n$ be the linear projection from $\mathbb{H}^n_H$ to B^n through the point $(-1, 0, \cdots, 0)$. This is shown for the case $n = 2$ in the diagram below.

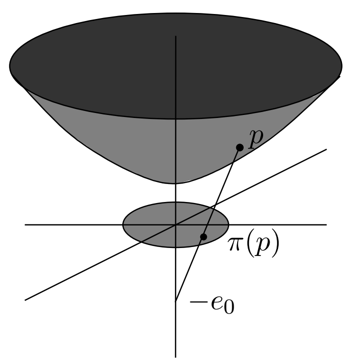

We define the conformal ball model to be the space B^n with the pull back metric induced by π^{-1} from the hyperboloid model. We denote the conformal ball model by $\mathbb{H}^n_C$. One can check that resulting metric is:

$$\frac{4(dx_1^2 + \cdots + dx_n^2)}{(1 - (x_1^2 + \cdots + x_n^2))^2}$$

and hence differs from the Euclidean metric on B by a conformal factor.

Within $\mathbb{H}^3_C$, hyperbolic planes correspond to Euclidean hemispheres and Euclidean planes in B^3 that meet ∂B^3 perpendicularly and hyperbolic lines correspond to Euclidean semi-circles and Euclidean lines that intersect ∂B^3 perpendicularly. Given an oriented hyperbolic plane P, the half-space defined by P consists of all of the points in B^3 on the side of W consistent with the orientation of W. Points at infinity in $\mathbb{H}^3_C$ correspond to points in the unit sphere ∂B^3.

See below for an image of a hyperbolic polyhedron depicted in the conformal ball model. The sphere at infinity is shown for reference.

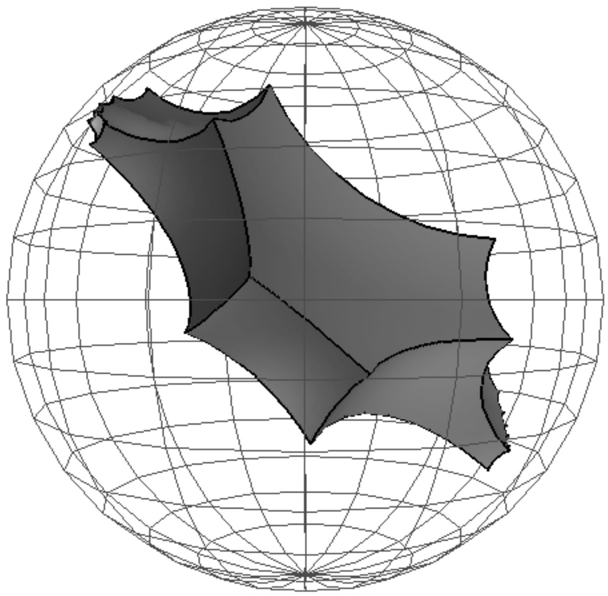

This hyperbolic polyhedron was displayed in the excellent computer program Geomview [2].

The upper half-space model of hyperbolic space

The model of hyperbolic space most commonly used and, perhaps the most intuitively immediate is the half-space model. However, this model is more difficult to relate to $\mathbb{H}^n_H$ than the previous models. Consider the differentiable mapping $i : D^n \to \mathbb{R}^n$ given by:

$$\mathbf{x} \mapsto 2\frac{\mathbf{x} + \mathbf{e}_n}{\|\mathbf{x} + \mathbf{e}_n\|^2} - \mathbf{e}_n$$

where $\mathbf{e}_n = (0, \cdots, 0, 1)$ and $\|.\|$ denotes the Euclidean norm in $\mathbb{R}^n$. One can easily check that i is a diffeomorphism from B^n to the set $H = \{\mathbf{x} \in \mathbb{R}^n | x_n > 0\}$. The half space model consists of the set H with the pullback metric induced by i^{-1} from $\mathbb{H}^3_C$. We will denote the upper half-space model by $\mathbb{H}^n_U$. One can check that metric on $\mathbb{H}^n_U$ is given by:

$$\frac{dx_1^2 + \cdots + dx_n^2}{x_n^2}.$$

Hence the metric from this model differs from the Euclidean metric by a conformal factor as well.

Within $\mathbb{H}^3_U$, hyperbolic planes are Euclidean hemispheres and Euclidean planes that intersect the plane $x_3 = 0$ perpendicularly. Hyperbolic lines consist of Euclidean semi-circles and Euclidean lines that intersect the plane $x_3 = 0$ perpendicularly. Given an oriented hyperbolic plane P in $\mathbb{H}^3_U$, the corresponding half-space is defined in the analogous way. Points at infinity in $\mathbb{H}^3_U$ correspond to points in the plane $x_3 = 0$ and one additional point which we will label ∞ corresponding to the one point compactification of this plane.

We will use the upper half-space model extensively because the positions of planes in $\mathbb{H}^3$ are uniquely defined by the positions of Euclidean lines and Euclidean circles in the plane $x_3 = 0$. The dihedral angle between a pair of intersecting planes corresponds exactly to the Euclidean angle between the pair of intersecting lines and circles in the plane $x_3 = 0$.

Throughout this paper we will often just refer to hyperbolic 3-space, $\mathbb{H}^3$, and only refer to a specific model of hyperbolic space when it is necessary.

Combinatorial properties of hyperbolic polyhedra and Andreev's Theorem

Some elementary combinatorial facts about hyperbolic polyhedra are essential before we can state Andreev's Theorem. Notice that a compact hyperbolic polyhedron P is topologically a 3-dimensional ball, and its boundary a 2-sphere $\mathbb{S}^2$. The face structure of P gives $\mathbb{S}^2$ the structure of a cell complex C whose faces correspond to the faces of P, and so forth.

Considering only hyperbolic polyhedra with non-obtuse dihedral angles simplifies the combinatorics of any such C:

Proposition 1.1 *(a) A finite vertex of a non-obtuse hyperbolic polyhedron P is the intersection of exactly 3 faces.*
(b) For such a P, we can compute the angles of the faces in terms of the dihedral angles; these angles are also $\leq \pi/2$.

Proof. Let v be a finite vertex where n faces of P meet. After an appropriate isometry, we can assume that v is the origin in the hyperbolic ball model, so that the faces at v are subsets of Euclidean planes through the origin. A small sphere centered at the origin will intersect P in a spherical n-gon Q whose angles are the dihedral angles between faces. Call these angles $\alpha_1, ..., \alpha_n$. The Gauss-Bonnet formula gives $\alpha_1 + \cdots + \alpha_n = \pi(n-2) + \text{Area}(Q)$. The restriction to $\alpha_i \leq \pi/2$ for all i gives $n\pi/2 \geq \pi(n-2) + \text{Area}(Q)$. Hence, $n\pi/2 < 2\pi$. We conclude that $n = 3$.

The edge lengths of Q are precisely the angles in the faces at the origin. Supposing that Q has angles $(\alpha_i, \alpha_j, \alpha_k)$ and edge lengths $(\beta_i, \beta_j, \beta_k)$ with the edge β_l opposite of angle α_l for each l, The Law of Cosines in spherical geometry gives that:

$$\cos(\beta_i) = \frac{\cos(\alpha_i) + \cos(\alpha_j)\cos(\alpha_k)}{\sin(\alpha_j)\sin(\alpha_k)}. \tag{1}$$

Hence, the face angles are calculable from the dihedral angles. They are non-obtuse, since the right hand side of the equation is positive for $\alpha_i, \alpha_j, \alpha_k$ non-obtuse.

(The reader should notice that this equation will be used heavily throughout this paper.) □

The fundamental axioms of incidence place the following, obvious, further restrictions on the complex C:

- Every edge of C belongs to exactly two faces.
- A non-empty intersection of two faces is either an edge or a vertex.
- Every face contains not fewer than three edges.

We will call any trivalent cell complex C on $\mathbb{S}^2$ that satisfies the three conditions above an *abstract polyhedron*. Notice that since C must be a trivalent cell complex on $\mathbb{S}^2$, its dual, C^*, has only triangular faces. The three other conditions above give that the dual complex C^* is a simplicial complex on $\mathbb{S}^2$. (Andreev refers to this dual complex as the *scheme of the polyhedron*.) The restriction to classifying hyperbolic polyhedra with non-obtuse dihedral angles provides a simplification that is necessary in the proof of Andreev's Theorem.

We call a simple closed curve Γ formed of k edges of C^* a *k-circuit* and if all of the endpoints of the edges of C intersected by Γ are distinct, we call such a circuit a *prismatic k-circuit*.

Theorem 1.2 Andreev's Theorem

Let C be an abstract polyhedron with more than 4 faces and suppose that non-obtuse angles α_i are given corresponding to each edge e_i of C. There is a unique compact hyperbolic polyhedron P, up to isometries of $\mathbb{H}^3$, whose faces realize C with dihedral angle α_i at each edge e_i if and only if:

1. *For each edge e_i, $0 < \alpha_i \leq \pi/2$.*
2. *Whenever 3 distinct edges e_i, e_j, e_k meet at a vertex, $\alpha_i + \alpha_j + \alpha_k > \pi$.*
3. *Whenever Γ is a prismatic 3-circuit intersecting edges e_i, e_j, e_k, $\alpha_i + \alpha_j + \alpha_k < \pi$.*
4. *Whenever Γ is a prismatic 4-circuit intersecting edges e_i, e_j, e_k, e_l, then $\alpha_i + \alpha_j + \alpha_k + \alpha_l < 2\pi$.*
5. *Whenever there is a four sided face bounded by edges e_1, e_2, e_3, e_4, enumerated successively, with edges $e_{12}, e_{23}, e_{34}, e_{41}$ entering the four vertices (edge e_{ij} connects to the ends of e_i and e_j), then:*

$$\alpha_1 + \alpha_3 + \alpha_{12} + \alpha_{23} + \alpha_{34} + \alpha_{41} < 3\pi$$

$$\alpha_2 + \alpha_4 + \alpha_{12} + \alpha_{23} + \alpha_{34} + \alpha_{41} < 3\pi$$

For a given C let E be the number of edges of C. The subset of $(0, \pi/2]^E$ satisfying these linear inequalities will be called the *Andreev Polytope*, A_C. Since A_C is determined by linear inequalities, it is convex.

It is worth noting that conditions (4) and (5) are only relevant when each of the angles listed in the inequality is $\pi/2$, since all angles are restricted to being non-obtuse.

Proposition 1.3 *If C is not the triangular prism, condition (5) of Andreev's Theorem is a consequence of conditions (3) and (4).*

Proof: Given a quadrilateral face, if the four edges leading from it form a prismatic 4-circuit, Γ_1, as depicted on the left hand side of the figure below, clearly condition (5) is a result of condition (4). Otherwise, at least one pair of the edges leading from it meet at a vertex. If only one pair meets at a point, we have the diagram below in the middle. In this case the curve Γ_2 is a prismatic 3-circuit, so that $\alpha_{3,4} + \alpha_{4,1} + \beta < \pi$, so that condition (5) is satisfied because both $\alpha_{3,4}$ and $\alpha_{4,1}$ cannot be $\pi/2$.

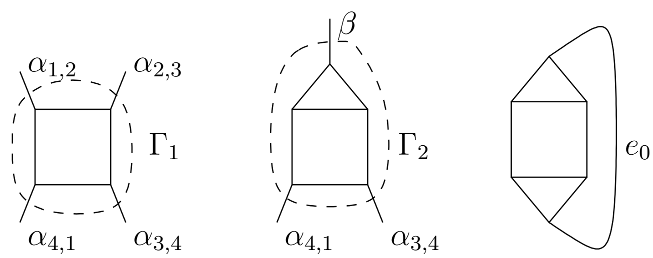

Otherwise, if two pairs of the edges leaving the quadrilateral face meet at vertices, we have the diagram on the right hand side. The only way to complete this diagram is with the edge labeled e_0, resulting in the triangular prism. □

Hence, we need only check condition (5) for the triangular prism, which corresponds to the only five faced C. The only C having fewer than five faces is the tetrahedron, which is dealt with as a special case in a special section at the end.

Given some C, it may be a difficult problem to determine whether $A_C = \emptyset$ and correspondingly, whether there are any hyperbolic polyhedra realizing C with non-obtuse dihedral angles. In fact, for the abstract polyhedron in the following figure, conditions (2) and (3) give respectively that $\alpha_1 + \cdots + \alpha_{12} > 3\pi$ and $\alpha_1 + \cdots + \alpha_{12} < 3\pi$. So, for this C, we have $A_C = \emptyset$. However, for more complicated C, it can be significantly harder to determine whether $A_C = \emptyset$.

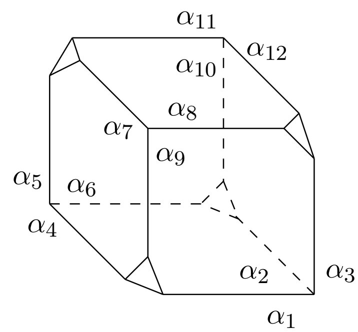

Luckily, there are special cases:

Corollary 1.4 *If there are no prismatic 3-circuits in C, there exists a unique hyperbolic polyhedron realizing C with dihedral angles $2\pi/5$.*

Proof. Since there are no prismatic 3-circuits in C condition (3) of the theorem is vacuous and clearly $\alpha_i = 2\pi/5$ satisfy conditions (1), (2), (4), and (5). □

The following two corollaries are essential for Thurston's hyperbolization Theorem:

Corollary 1.5 *If there are no prismatic 3-circuits and no prismatic 4-circuits in C, there exists a unique hyperbolic polyhedron realizing C with dihedral angles $\pi/2$.*

Proof. Conditions (1) and (2) are clearly satisfied setting all of the dihedral angles to $\pi/2$. Since there are no prismatic 3 or 4-circuits, conditions (3) and (4) are irrelevant. Condition (5) is also irrelevant since it needs only be checked for the triangular prism, which has a prismatic 3-circuit. □

Corollary 1.6 *If C is any (trivalent) abstract polyhedron we can construct a abstract polyhedron C' by subdividing the faces of C that can be realized by a polyhedron with all right angles.*

Proof.

Here, and many places later in this paper it will be easier to work in the dual complex C^*, since it is a simplicial complex. We will show how to add edges to C^* to eliminate every prismatic 3 or 4 circuit by increasing the number of edges in each such circuit. The first step is to do a barycentric subdivision on C^*. After this has been done, there can be no prismatic three circuits and no prismatic four circuits, other than those surrounding a four valent vertex of C^*. (The barycentric subdivision process creates many of these!) The figure below shows how preexisting prismatic 3 and 4-circuits are eliminated by the barycentric subdivision. The grey regions denote places where other simplicies could be.

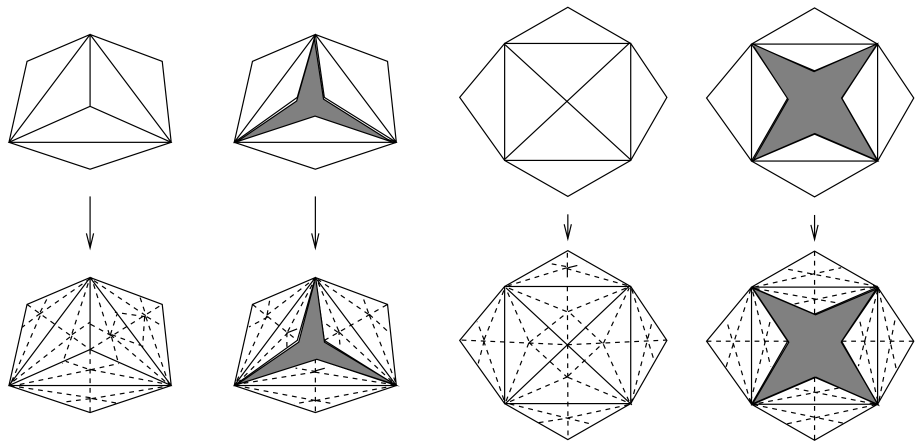

By the steps that we have already done, the types of remaining prismatic 3 or 4-circuits are prismatic 4-circuits surrounding a single 4-valent vertex that are created by the barycentric subdivision, as in the left hand side of the diagram below. As a result of the barycentric subdivision, none of the four triangles surrounding the outside of the prismatic 4-circuit can share a side. The right hand side of this diagram shows how to add more edges (the dashed ones) in a way that eliminates these prismatic 4-circuits without introducing any new ones.

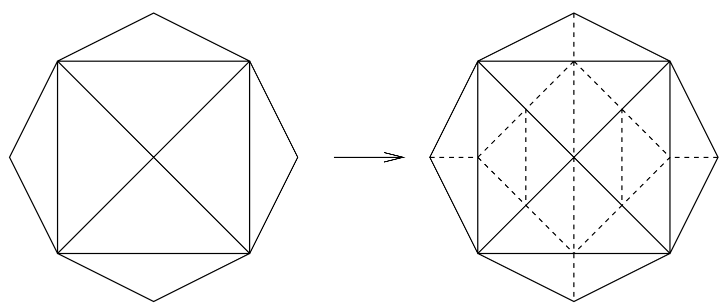

After doing this last step we have introduced new edges to C^* to eliminate every prismatic 3-circuit and every prismatic 4-circuit. The resulting complex, C'^*, can therefore be realized with all right angles by Corollary (1.5). $\square$

2 Setup of the Proof.

The proof of Andreev's Theorem uses a relatively common type of reasoning: you manufacture two manifolds of the same dimension: one, X, consisting of the geometric objects that you want to construct, and the other, Y, a subset of $\mathbb{R}^n$ consisting of various angles, lengths, etc. The space X should be viewed as unknown and the space Y as known.

You then consider the mapping $f : X \to Y$ which takes your geometric object, in X, and reads off its appropriate measurements, in Y. Of course, you need to show that the image is actually in Y, namely, that the constraints that you put on the coordinates of Y (typically something like the triangle inequality for the edges of a triangle) are indeed satisfied for each geometric object of X.

This map f will always be obviously continuous, and it is not too hard to show that it is proper and injective, hence a homeomorphism onto its image, which is a union of connected components of Y. This reduces the problem to showing that X is nonempty and that Y is connected, which are usually the hardest parts!

Given an abstract polyhedron C, define $\mathcal{P}_\mathcal{C}$ to be the set of compact hyperbolic polyhedra realizing C up to hyperbolic isometry. Let N, E, and V be the number of faces, edges, and vertices in C.

Proposition 2.1 *The space $\mathcal{P}_\mathcal{C}$ is a manifold of dimension $3N-6$ (perhaps empty).*

Proof.

Let $\mathcal{H}$ be the space of closed half spaces of $\mathbb{H}^3$; clearly $\mathcal{H}$ is a 3-dimensional manifold. The set of compact polyhedra with N faces, $\mathcal{O}_N$, is an open subset of the $3N$-dimensional manifold $\mathcal{H}^N$, and those whose combinatorial structure is a given abstract polyhedron forms some union of components of $\mathcal{O}_N$. Clearly the automorphisms of $\mathbb{H}^3$ act freely on the subset of $\mathcal{H}^N$ where all the subspaces are distinct, as soon as $N \geq 3$, hence they act freely on $\mathcal{O}_N$. So $\mathcal{O}_N/Aut(\mathbb{H}^3)$ is a manifold, and $\mathcal{P}_\mathcal{C}$ is an open subset of this manifold. This manifold has dimension $3N-6$, since $\mathcal{O}_N$ has dimension $3N$ and $Aut(\mathbb{H}^3)$ has dimension 6. $\Box$

In fact, we will restrict to the subset $\mathcal{P}_C^0$ of polyhedra with dihedral angles in $(0, \pi/2]$. Notice that $\mathcal{P}_C^0$ is not, a-priori, a manifold or even a manifold with boundary. All that we will need for the proof of Andreev's Theorem is that $\mathcal{P}_C$ is a manifold and that the subspace $\mathcal{P}_C^0$ is a metric space.

Using the fact that the edge graph of C is trivalent, one can check that E, the number of edges of C, is the same as the dimension of $\mathcal{P}_C$. Since exactly three edges enter each vertex and each edge enters exactly two vertices, $2V = 3E$. The Euler characteristic gives $N - E + V = N - E + 2/3E = 2$ implying $E = 3(N-2)$, the dimension of $\mathcal{P}_C$.

Given any $P \in \mathcal{P}_C$ let $\alpha(P) = (\alpha_1, \alpha_2, \alpha_3, ...)$ be the E-tuple consisting of the dihedral angles of P at each edge (according to some fixed numbering of the edges of C). This map α is obviously continuous with respect to the topology on $\mathcal{P}_C$, which it inherits from its manifold structure.

So, we have $\alpha : \mathcal{P}_C \to \mathbb{R}^E$ and we will use the general idea of proof presented in the beginning of this section. The goal is to show that α restricted to $\mathcal{P}_C^0$ is a homeomorphism onto A_C. This will prove Andreev's Theorem. Notice that A_C is a convex subset of $\mathbb{R}^E$, so A_C is connected. Of course we have to first check that $\alpha(\mathcal{P}_C^0) \subset A_C$.

3 The inequalities are satisfied.

This section begins the proof of Andreev's Theorem.

Proposition 3.1 *Given $P \in \mathcal{P}_C^0$, the dihedral angles $\alpha(P)$ satisfy conditions (1-5).*

We will need the following two lemmas about the basic properties of hyperbolic geometry.

Lemma 3.2 *Suppose that three planes $P_{\mathbf{v_1}}, P_{\mathbf{v_2}}, P_{\mathbf{v_3}}$ intersect pairwise in $\mathbb{H}^3$ with non-obtuse dihedral angles α, β, and γ. Then, $P_{\mathbf{v_1}}, P_{\mathbf{v_2}}, P_{\mathbf{v_3}}$ intersect at a vertex in $\overline{\mathbb{H}^3}$ if and only if $\alpha + \beta + \gamma \geq \pi$. The planes intersect in $\mathbb{H}^3$ if and only if the inequality is strict.*

Proof.

The planes intersect in a point of $\overline{\mathbb{H}^3}$ if and only if the subspace spanned by $\mathbf{v_1}, \mathbf{v_2}, \mathbf{v_3}$ is positive semi-definite, so that the orthogonal is a negative semi-definite line of $E^{1,3}$. If the inner product on this line is negative, the line defines a point of intersection with the hyperboloid model. Otherwise, the inner product on the line is zero, this line corresponds to a point in $\partial\mathbb{H}^3$, since the line is then in the cone to which the hyperboloid is asymptotic. The symmetric matrix defining the inner product is

$$\begin{bmatrix} 1 & \langle \mathbf{v_1}, \mathbf{v_2} \rangle & \langle \mathbf{v_1}, \mathbf{v_3} \rangle \\ \langle \mathbf{v_1}, \mathbf{v_2} \rangle & 1 & \langle \mathbf{v_2}, \mathbf{v_3} \rangle \\ \langle \mathbf{v_1}, \mathbf{v_3} \rangle & \langle \mathbf{v_2}, \mathbf{v_3} \rangle & 1 \end{bmatrix} = \begin{bmatrix} 1 & -\cos\alpha & -\cos\beta \\ -\cos\alpha & 1 & -\cos\gamma \\ -\cos\beta & -\cos\gamma & 1 \end{bmatrix}$$

where α, β, and γ are the dihedral angles between the pairs of faces $(P_{\mathbf{v_1}}, P_{\mathbf{v_2}})$, $(P_{\mathbf{v_1}}, P_{\mathbf{v_3}})$, and $(P_{\mathbf{v_2}}, P_{\mathbf{v_3}})$, respectively.

Since the principle minor is positive definite for $0 < \alpha \leq \pi/2$, it is enough to find out when the determinant

$$1 - 2\cos\alpha\cos\beta\cos\gamma - \cos^2\alpha - \cos^2\beta - \cos^2\gamma$$

is non-negative.

A bit of trigonometric trickery (we used complex exponentials) shows that the expression above can be rewritten

$$-4\cos\left(\frac{\alpha+\beta+\gamma}{2}\right)\cos\left(\frac{\alpha-\beta+\gamma}{2}\right)\cos\left(\frac{\alpha+\beta-\gamma}{2}\right)\cos\left(\frac{-\alpha+\beta+\gamma}{2}\right) \quad (2)$$

Let $\delta = \alpha + \beta + \gamma$. When $\delta < \pi$, (2) is strictly negative, when $\delta = \pi$, (2) is clearly zero, and when $\delta > \pi$ (2) is strictly positive. Hence the inner product on the space spanned by $\mathbf{v_1}, \mathbf{v_2}, \mathbf{v_3}$ is positive semidefinite if and only if $\delta \geq \pi$. It is positive definite if and only if $\delta > \pi$.

Then it is easy to see that the three planes $P_{\mathbf{v_1}}, P_{\mathbf{v_2}}, P_{\mathbf{v_3}} \subset \mathbb{H}^3$ intersect at a point in $\overline{\mathbb{H}^3}$ if and only if they intersect pairwise in $\mathbb{H}^3$ and the sum of the dihedral angles $\delta \geq \pi$. It is also clear that they intersect at a finite point if and only if the inequality is strict. $\square$

Lemma 3.3 *Let $P_1, P_2, P_3 \subset \mathbb{H}^3$ be planes carrying faces of a polyhedron P that has all dihedral angles $\leq \pi/2$.*
(a) If P_1, P_2, P_3 intersect at a point in $\mathbb{H}^3$, then the point $p = P_1 \cap P_2 \cap P_3$ is a vertex of P.
(b) If P_1, P_2, P_3 intersect at a point in $\partial\mathbb{H}^3$, then P is not compact, and the point of intersection is in the closure of P.

Proof. (a) Consider what we see in the plane P_1. Let H_i^+ be the half space bounded by P_i which contains the interior of P, and let $Q = P_1 \cap H_2^+ \cap H_3^+$. If $p \notin P$, then let U be the component of $Q - P$ that contains p in its closure. This is a non-convex polygon; let $p, p_1, ..., p_k$ be its vertices. The exterior angles of U at $p_1, ..., p_k$ are the angles of the face of P carried by P_1, hence $\leq \pi/2$ by part (b) of Proposition 1.1. See the following figure:

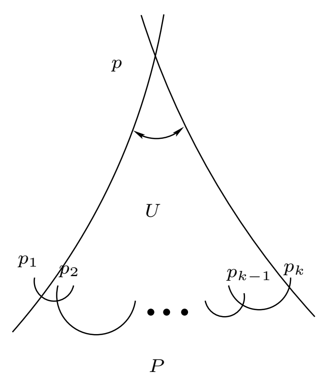

Suppose that $\alpha_1, ...\alpha_k$ are the angles of P at $p_1, ..., p_k$, and let α be the angle at p. Then the Gauss-Bonnet formula tells us that:

$$(\pi - \alpha) + \alpha_1 - ((\pi - \alpha_2) + \cdots + (\pi - \alpha_{k-1})) + \alpha_k - \text{Area}(U) = 2\pi,$$

which can be rearranged to read

$$(\alpha_1 + \alpha_k - \pi) - \alpha - \sum_{j=2}^{k-1}(\pi - \alpha_j) = \text{Area}(U).$$

This is clearly a contradiction. All of the terms on the left are non-positive, and $\text{Area}(U) > 0$.

If p is at infinity (i.e. $\alpha = 0$), this expression is still a contradiction, proving part (b). □

Proof of Proposition 3.1.

For condition (1), notice that if two adjacent faces intersect at dihedral angle 0, they intersect at a point at infinity. If this were the case, P would be non-compact. In addition, the dihedral angle between adjacent faces is $\leq \pi/2$ by hypothesis.

For condition (2), let x be a vertex of P. Since P is compact, $x \in \mathbb{H}^3$ and by Lemma 3.2 part (a), the sum of the dihedral angles between the three planes intersecting at x must be $> \pi$.

For condition (3), note first that by Lemma 3.2 if three faces forming a 3-circuit have dihedral angles summing to a number $\geq \pi$, then they meet in $\overline{\mathbb{H}}^3$. If they meet at a point in $\mathbb{H}^3$, by Lemma 3.3 part (a) this point is a vertex of p, so these three faces do not form a prismatic 3-circuit. Alternatively, if the three planes meet in $\partial\mathbb{H}^3$ by Lemma 3.3 part (b) the P is non-compact, contrary to assumption. Hence, any three faces forming a prismatic 3-circuit in P must have dihedral angles summing to $< \pi$.

For condition (4), let $H_{\mathbf{v_1}}, H_{\mathbf{v_2}}, H_{\mathbf{v_3}}, H_{\mathbf{v_4}}$ be a prismatic 4-circuit; obviously it satisfies condition (4) unless all of the dihedral angles are $\pi/2$, so we suppose that they are. We will assume the normalization $\langle \mathbf{v_i}, \mathbf{v_i} \rangle = 1$ for each i. The Gram matrix

$$Q = \begin{bmatrix} \langle \mathbf{v_1}, \mathbf{v_1} \rangle & \langle \mathbf{v_1}, \mathbf{v_2} \rangle & \langle \mathbf{v_1}, \mathbf{v_3} \rangle & \langle \mathbf{v_1}, \mathbf{v_4} \rangle \\ \langle \mathbf{v_2}, \mathbf{v_1} \rangle & \langle \mathbf{v_2}, \mathbf{v_2} \rangle & \langle \mathbf{v_2}, \mathbf{v_3} \rangle & \langle \mathbf{v_2}, \mathbf{v_4} \rangle \\ \langle \mathbf{v_3}, \mathbf{v_1} \rangle & \langle \mathbf{v_3}, \mathbf{v_2} \rangle & \langle \mathbf{v_3}, \mathbf{v_3} \rangle & \langle \mathbf{v_3}, \mathbf{v_4} \rangle \\ \langle \mathbf{v_4}, \mathbf{v_1} \rangle & \langle \mathbf{v_4}, \mathbf{v_2} \rangle & \langle \mathbf{v_4}, \mathbf{v_3} \rangle & \langle \mathbf{v_4}, \mathbf{v_4} \rangle \end{bmatrix} =$$

$$\begin{bmatrix} 1 & 0 & \langle \mathbf{v_1}, \mathbf{v_3} \rangle & 0 \\ 0 & 1 & 0 & \langle \mathbf{v_2}, \mathbf{v_4} \rangle \\ \langle \mathbf{v_3}, \mathbf{v_1} \rangle & 0 & 1 & 0 \\ 0 & \langle \mathbf{v_4}, \mathbf{v_2} \rangle & 0 & 1 \end{bmatrix}$$

has determinant 0 if the $\mathbf{v}$'s are linearly dependent, and otherwise represents the inner product of $E^{3,1}$ and hence has negative determinant. In both cases we have

$$detQ = (1 - \langle \mathbf{v_1}, \mathbf{v_3} \rangle^2)(1 - \langle \mathbf{v_2}, \mathbf{v_4} \rangle^2) \leq 0.$$

So $\langle \mathbf{v_1}, \mathbf{v_3} \rangle^2 \leq 1$ and $\langle \mathbf{v_2}, \mathbf{v_4} \rangle^2 \geq 1$ or vice versa (perhaps one or both are equalities). This means that one of the opposite pairs of faces of the 4-circuit intersect, perhaps at a point at infinity. We can suppose that this pair is $H_{\mathbf{v_1}}$ and $H_{\mathbf{v_3}}$.

If $H_{\mathbf{v_1}}$ and $H_{\mathbf{v_3}}$ intersect in $\mathbb{H}^3$, they do so with positive dihedral angle. Since $H_{\mathbf{v_2}}$ intersects each $H_{\mathbf{v_1}}$ and $H_{\mathbf{v_3}}$ with dihedral angle $\pi/2$ the three faces pairwise intersect and have dihedral angle sum $> \pi$. By Lemmas 3.2 and 3.3 these three faces intersect at a point in $\mathbb{H}^3$ which is a vertex of P. In this case, the 4-circuit $H_{\mathbf{v_1}}, H_{\mathbf{v_2}}, H_{\mathbf{v_3}}, H_{\mathbf{v_4}}$ is not prismatic.

Otherwise, $H_{\mathbf{v_1}}$ and $H_{\mathbf{v_3}}$ intersect at a point at infinity. In this case, since $H_{\mathbf{v_2}}$ intersects each $H_{\mathbf{v_1}}$ and $H_{\mathbf{v_3}}$ with dihedral angle $\pi/2$ the three faces intersect at this point at infinity by Lemma 3.2 and then by Lemma 3.3 P is not compact, contrary to assumption.

Hence, if $H_{\mathbf{v_1}}, H_{\mathbf{v_2}}, H_{\mathbf{v_3}}, H_{\mathbf{v_4}}$ forms a prismatic 4-circuit, the sum of the dihedral angles cannot be 2π.

For condition (5), suppose that the quadrilateral is formed by edges e_1, e_2, e_3, e_4. Violation of one of the inequalities would give that the dihedral angles at each of the edges e_{ij} leading to the quadrilateral is $\pi/2$ and that the dihedral angles at two of the opposite edges of the quadrilateral are $\pi/2$. See the diagram below:

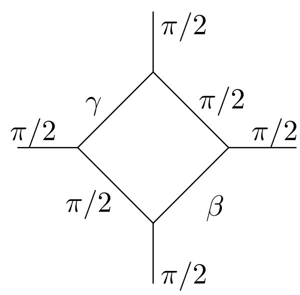

Consider a vertex of the quadrilateral formed by edges e_i, e_j, and e_{ij}. Violation of the inequality gives that $\alpha_{ij} = \pi/2$ and either $\alpha_i = \pi/2$ or $\alpha_j = \pi/2$. Using Equation (1), we see that the face angle in the quadrilateral at this vertex must be $\pi/2$. So, we have that each of the face angles of the quadrilateral is $\pi/2$, which is a contradiction to the Gauss-Bonnet Theorem. Hence both of the inequalities in condition (5) must be satisfied.

This was the last step in proving Proposition 3.1. □

4 The mapping α is injective.

Proposition 4.1 *The mapping $\alpha : \mathcal{P}_C \to \mathbb{R}^E$ is injective.*

This proposition depends on a famous lemma, due to Cauchy, who also used it to prove a rigidity result: a convex Euclidean polyhedron with given faces is rigid.

The reader is highly encouraged to draw the diagrams corresponding to the combinatorial configurations described in this section. It clarifies matters greatly.
Proof of Proposition 4.1 The proof of this proposition consists of the following two propositions.

Proposition 4.2 *Let P, P' be two convex hyperbolic polyhedra, and $f : P \to P'$ a homeomorphism which maps vertices to vertices and edges to edges. Suppose that the dihedral angles which correspond under f are equal. Label each edge e of C with $-, 0, +$ if the edge e in P' is smaller, equal, or greater in length than e is in P Then: for any face F of C not marked entirely with 0's, the edges labeled $+$ cannot all be in the same component of ∂P with all of the edges labeled $-$ removed.*

The reader should notice that the exact same lemma is obtained switching $+$ with $-$ by switching P with P'. Hence, from this proposition, any face of C which is not labeled with all 0s must have both edges labeled $+$ and edges labeled $-$ and these labels cannot "have all $+$'s and 0's on one side and all $-$'s and $0's$ on the other".

Proof. Let us define the *dist-angle* of two half-planes in $\mathbb{H}^2$ to be the distance between the two boundary lines if these do not intersect, and the negative of the angle between them if they do. The dist-angle is a number in $(-\pi, \infty)$.

Now suppose that we have a polygon $Q \subset \mathbb{H}^2$, perhaps non-compact, but with connected boundary. That is, if Q is non-compact, it has exactly two non-compact sides. Let s, s', s'' be three distinct sides of Q, with s' and s'' either consecutive, or non-compact. We will consider deformations of Q in which all sides are kept a constant length except s, and all angles are kept constant except that between s' and s'' (if it exists.)

Lemma 4.3 *If Q is deformed as above, the dist-angle between sides s' and s'' is a monotone increasing function of the length of s.*

Proof. Embed Q in the band model $\mathbb{B}$ of the hyperbolic plane so that the side s is on the axis. Let $\partial'(Q)$ be the part of ∂Q which contains s', and $\partial''(Q)$ the part which contains s''. Then horizontal translation of $\partial''(Q)$, keeping $\partial'(Q)$ fixed realized the deformations above. The following two figures now prove the lemma.

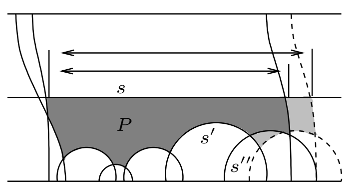

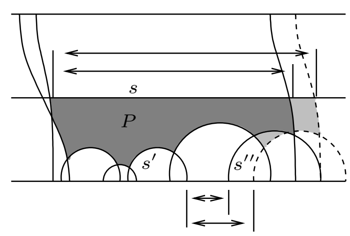

□

Continuing the proof of the Proposition 4.2. Suppose for contradiction that at least some side s of P is labeled $+$, and that the sides labeled $+$ are all in one component of ∂P with the sides labeled $-$ removed. It is then possible to choose sides s' and s'', distinct from each-other and from s, such that the component of $\partial P - (s' \cup s'')$ containing s contains only sides labeled $+$ or 0 and the other component of $\partial P - (s' \cup s'')$ (possibly empty) contains no edges labeled $+$. We do not care how s and s' are labeled.

Looking at the polygon $P+$ bounded by $s' \cup s''$ and the component of $\partial P - (s' \cup s'')$ containing s, the above lemma guarantees that the dist-angle between sides s' and s'' in P' is strictly greater than in P. But looking at the polygon bounding s' and s'' and the other component of $\partial P - (s' \cup s'')$, either s and s' meet at a vertex, so the distangle remains the same or, if there are other sides between s and s' in this component all of which are labeled 0 or $-$. So the lemma guarantees that the dist-angle between sides s' and s'' in P' is at most equal to that in P. This is a contradiction. □

Now suppose that $P_1, P_2 \in \mathcal{P}_C^0$ are two polyhedra such that $\alpha(P_1) = \alpha(P_2)$. We can label each edge e of C by $-, 0$, or $+$ depending on whether the length of e in P_1 is less than, equal to, or greater than the length of e in P_2.

Now the injectivity of the mapping α follows from Cauchy's result.

Proposition 4.4 Cauchy's Lemma *If we mark all the edges of C with $+, 0, -$ so as to satisfy the conditions of the above proposition, then all edges are marked 0.*

Proof.

If there are any edges labeled $+$ and $-$, we will derive a contradiction from the hypothesis of how these edges of C must be distributed, by defining a line field on $\mathbb{S}^2$ having index sum of the singular points ≤ 0, contradicting the well-known Poincaré-Hopf Theorem. So, assume that there are some edges labeled $+$ and $-$.

We will call each face whose edges are not all labeled 0 *essential*. We first define the line field in a neighborhood of the edges of essential faces as follows. The line field will be transverse to the edges labeled $+$, and tangent to the edges labeled $-$. For edges labeled 0, it will be tangent to all of these except if along the boundary of an essential face we see three consecutive edges labeled $+, 0, +$, then we put a singularity of the line field on this edge labeled 0, as in the figure. We will call such edges "singular 0-edges". The singular edges will be essential to help to cancel out any possible index from singularities with positive index.

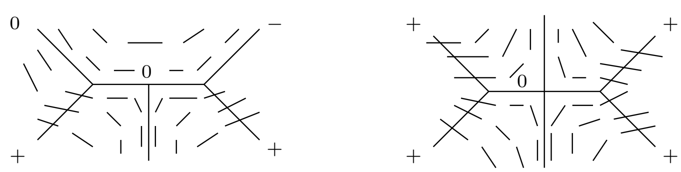

Singular 0-edges

The first thing to check is that this line field can be extended to each face whose edges are not all labeled 0 (called *essential faces*), with at most a singularity of negative index in the interior. (See the following figure.) Let F be an essential face. Choose some point p "in the middle" of F. Consider the edges labeled $+$ and the singular 0-edges, with all other edges removed. We will draw smooth curves from points on each component to p. If the component contains any singular 0-edges, draw the curves from these points, otherwise draw a single curve from a random point in the interior of one of the $+$ edges. These curves can be made to not intersect. On should check that Proposition 4 guarantees that there will be at least $n \geq 2$ paths from ∂F to p in total. These paths divide F into "sectors" with each sector containing exactly two components labeled $+$. Connect smooth curves between these two components, completely filling out the sector. The line field on F will be the collection of lines tangent to each of these curves. One can check that there is either no singular point, corresponding to $n = 2$, or that the index of the singular point is $-(n-2)/2$, for $n \geq 2$

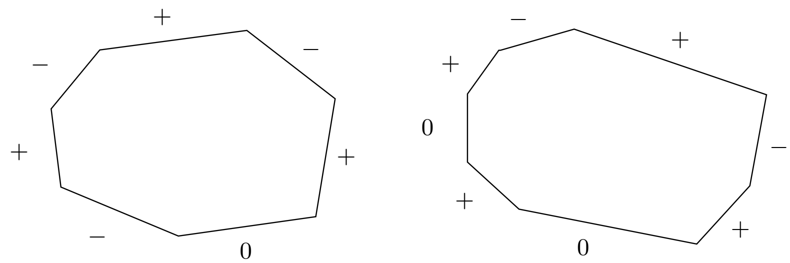

Two different examples of essential faces

Now consider the connected components consisting of inessential faces. The line field is already defined on the boundary of these components and tangent to it, so the line field can be extended to the interior. The total index inside the component is equal to its Euler characteristic. This will be 0 or negative if the component is not homeomorphic to the disc, but in the case of the disc, we will be forced to put a singularity of index $+1$ in the interior.

We have created a line field everywhere on $\mathbb{S}^2$ with finitely many singularities consisting of:

- At most one singularity of negative index inside of each essential face.
- Possible singularities at each vertex of C.
- A singularity in each singular 0-edge.
- Some finite number of singularities in each connected component of non-essential faces with index sum at most 1.

There are exactly four types of line field locally at each vertex of C, depending on the number of + edges which enter the vertex. The reader should check that if there are 3 or 0 + edges at a vertex, the vertex will have singularity with index $-1/2$. Otherwise, the line field can be extended continuously in a neighborhood of the vertex.

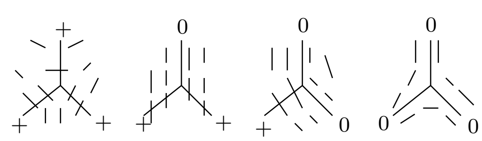

The four types of singular vertices

By construction, the singularities on singular 0-edges have index $-1/2$ or -1, depending on whether they are singular on each side.

So we see that the only singularities of positive index must be on the interior of the connected components of non-essential faces. Since each such component D has total index sum ≤ 1 it will be sufficient to find enough singularities on ∂D of negative index to cancel out at least $+1$.

Such a component D must have non-empty boundary, since we assume that there are some edges that are not labeled 0. Its boundary has at least three vertices that have edges leading from from them, out of D. (Two or fewer would be geometrically impossible, since C corresponds to a convex polyhedron.) If at least two such edges are labeled 0 or $-$, then these vertices have index $\leq -1/2$, and cancel the $+1$ from the interior of D. Otherwise there is either exactly one such edge labeled 0 or $-$, and hence at least two consecutive edges labeled $+$, or all the edges leaving ∂D are labeled $+$. In the first case, the edge labeled $-$ or 0 contributes a singularity of index $\leq -1/2$ and the edge of ∂D between two consecutive edges labeled $+$ contributes a singularity of $-1/2$ also. (Since it is a "singular 0-edge".) In the second case, with all of the edges labeled $+$, each edge of ∂D is a "singular 0-edge", hence contributes a singularity of index $-1/2$. Since there must be at least three of these edges, the indices of these singularities cancels out the index $+1$ from the interior of D.

Hence, if there are any edges not labeled 0, one can create a linefield on $\mathbb{S}^2$ having singularities with index sum ≤ 0, whereas the Poincaré-Hopf Theorem states that this index sum must equal exactly 2. □

The construction of such a linefield with non-positive index sum fails if all of the edges are labeled zero since we essentially used that each component of inessential faces had boundary. If all of the edges were labeled 0, there would be one inessential component, consisting of the whole sphere.

So, we see that if $\alpha(P_1) = \alpha(P_2)$ then each pair of corresponding edges has the same length. But then P_1 and P_2 are congruent, since the faces are congruent. Notice that we have not used in this section any restriction on the dihedral angles, so in fact we have shown that $\alpha : \mathcal{P}_C \to \mathbb{R}^E$ is injective, proving Proposition 4.1. □

5 The mapping α is proper.

Proposition 5.1 *The mapping $\alpha : \mathcal{P}_C^0 \to A_C$ is proper.*

Proof: Let P_i be a sequence in $\mathcal{P}_C^0$, such that $\alpha(P_i) = \mathbf{a}_i$ converges in A_C. Then, by sequential compactness, we must show that a subsequence of the P_i converges in $\mathcal{P}_C^0$. We will first prove the following two lemmas:

Lemma 5.2 *Let $\mathcal{F}$ be a face of a hyperbolic polyhedron P with non-obtuse dihedral angles. If a face angle of $\mathcal{F}$ equals $\pi/2$ at the vertex v, then the dihedral angle of the edge opposite the face angle is $\pi/2$ and the dihedral angle of one of the two edges in $\mathcal{F}$ that enters v is $\pi/2$.*

Proof: This will follow from Equation (1) in Lemma 1.1 which one can use to calculate face angles from the dihedral angles at a vertex. In Equation (1), if $\beta_i = \pi/2$ we have:

$$0 = \frac{\cos(\alpha_i) + \cos(\alpha_j)\cos(\alpha_k)}{\sin(\alpha_j)\sin(\alpha_k)},$$

where α_i is the dihedral angle opposite of the face angle β_i and α_j, α_k are the dihedral angles of the other two edges entering v. Both $\cos(\alpha_i) \geq 0$ and $\cos(\alpha_j)\cos(\alpha_k) \geq 0$ for non-obtuse α_i, α_j, and α_k, so that $\cos(\alpha_i) = 0$ and $\cos(\alpha_j)\cos(\alpha_k) = 0$. Hence $\alpha_i = \pi/2$ and either $\alpha_j = \pi/2$ or $\alpha_k = \pi/2$. $\square$

Lemma 5.3 *If P_i is a sequence in $\mathcal{P}_C^0$ such that $\alpha(P_i) = \mathbf{a}_i$ converges in A_C then each edge length of the P_i remain bounded in $[\epsilon, \kappa]$ for some positive numbers ϵ and κ.*

Proof:

To establish the upper bound κ, we must show that no subsequence of the P_i can have diameters tending to infinity. So, re-indexing, if necessary, suppose that the diameters of the P_i tend to infinity; let p_i, q_i be two vertices realizing the diameter, and denote by $r_{1,i}, ...r_{k,i}$ the projections of the other vertices onto the line through p_i, q_i. (In fact, $k = 2N - 6$ if N is the number of faces.) As i tends to infinity, the largest gap in the $p_i, r_{1,i}, ...r_{k,i}, q_i$ must also tend to infinity. Using the ball model for $\mathbb{H}^3$, put the center of this gap at the origin, and the line p_i, q_i on some fixed line, say the vertical axis. Then all of the vertices will be in some neighborhoods N of the north pole and S of the south pole, and connecting groups of vertices there will be some edges $e_{1,i}, ..., e_{l,i}$ which intersect the equatorial plane H almost orthogonally, and close to the origin.

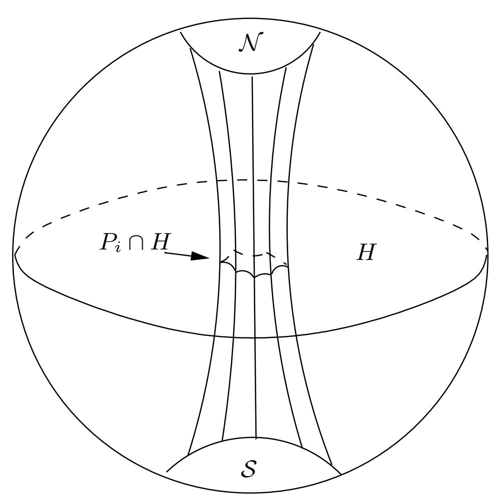

Thus the intersection $P_i \cap H$ will be almost a Euclidean polygon and its angles will be almost the dihedral angles $\alpha(e_{1,i}), ...\alpha(e_{l,i})$; in particular, for i sufficiently large they will be at most only slightly larger than $\pi/2$. This implies that l_i, the number of such edges, is 3 or 4 for i sufficiently large, as a Euclidean polygon with at least 5 faces has at least one angle $\geq 3\pi/5$.

If the faces $f_{1,i}, ..., f_{l,i}$ intersecting H are a prismatic circuit, and $l = 3$, the sum $\alpha(e_{1,i}) + \alpha(e_{2,i}) + \alpha(e_{3,i})$ tends to π, hence $\mathbf{a}_i$ cannot converge to a point in A_C by condition (3). Similarly, if $l_i = 4$, the corresponding sum tends to 2π violating condition (4).

Thus, the circuit $e_1, ..., e_l$ is not prismatic. If $l = 3$, this means that near one of the poles there is exactly one vertex and that the sum of the dihedral angles at that vertex tends to π, violating condition (2).

So we are left with the possibility that $l = 4$, and that f_1, f_2, f_3, f_4 does not form a prismatic 4-circuit. Therefore, a pair of opposite faces, say f_1 and f_3 intersect. Since the sum of the dihedral angles along this four circuit limits to 2π, and each dihedral angle is non-obtuse, each dihedral angle in this four circuit limits to $\pi/2$. If either f_1, f_2, f_3 don't meet at a vertex, or f_1, f_3, f_4 don't meet at a vertex, then this triple of faces will form a prismatic 3-circuit with dihedral angle sum limiting to $\geq \pi$ because the dihedral angles at two of the edges of the circuit limit to $\pi/2$. So, we assume that f_1, f_2, f_3 meet at a vertex r and f_1, f_3, f_4 meet at another vertex s, both in a neighborhood of a point at infinity, say the north pole. This situation is shown in the diagram below. At the vertex r, the dihedral angles between f_1 and f_2 and between f_2 and f_3 will converge to $\pi/2$, as mentioned before. In this situation, one can use Equation (1) to check that the face angle in the face f_2 will converge to the dihedral angle between faces f_1 and f_3. This is because the right hand side of the equation limits to $\cos(\alpha_{13})$, where α_{13} is the dihedral angle between faces f_1 and f_3. Then, as r diverges to infinity, the face angle in f_2 must limit to 0, and hence the dihedral angle between f_1 and f_3 must as well, contrary to condition (1).

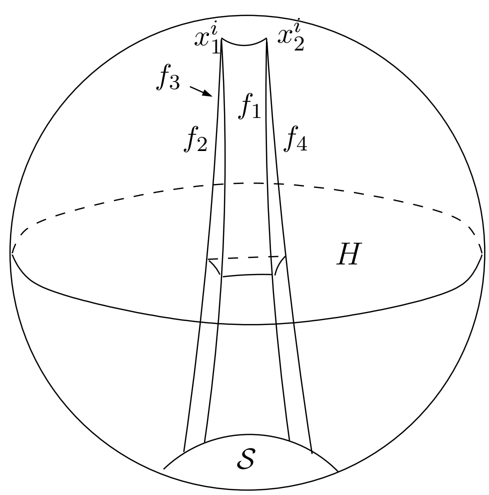

In conclusion, if the diameters of the P_i diverge the sequence $\mathbf{a_i}$ must diverge as well. This is contrary to the assumption that $\mathbf{a_i}$ converges.

Now, we establish the bound ϵ. We suppose that for some subsequence of the P_i some of the edges shrink to zero length. Re-indexing if necessary, we suppose that this happens for the P_i.

First, we check that none of the faces of the P_i can degenerate to either a point or a line segment. Any face, $\mathcal{F}$, that degenerates to a point of a line segment would become almost Euclidean. By Lemma 1.1, the face angles are all non-obtuse; this restricts $\mathcal{F}$ to either a triangle or a quadrilateral, the only Euclidean polygons having non-obtuse angles. If $\mathcal{F}$ is a triangle, the three edges leading to $\mathcal{F}$ form a prismatic 3-circuit because we assume that C is not the simplex. If $\mathcal{F}$ degenerates to a point, the three faces adjacent to $\mathcal{F}$ would meet at a finite vertex, in the limit. Therefore, by Lemma 3.2, the sum of the dihedral angles at the edges leading to $\mathcal{F}$ would limit to a value $> \pi$, contrary to condition (3). Otherwise, if $\mathcal{F}$ is a triangle and $\mathcal{F}$ degenerates to a line segment, in the limit, two of the face angles become $\pi/2$. Then, by Lemma 5.2, the dihedral angles at the edges opposite of these face angles become $\pi/2$. However, these edges are part of the prismatic 3-circuit of edges leading to $\mathcal{F}$, resulting in an angle sum $\geq \pi$, contrary to condition (3).

In the case that $\mathcal{F}$ is a quadrilateral, each of the face angles would have to limit to $\pi/2$. By Lemma 5.2, the dihedral angles at each of the edges leading from $\mathcal{F}$ to the rest of P would limit to $\pi/2$, as well as at least one edge of $\mathcal{F}$ leading to each vertex of $\mathcal{F}$. Therefore, the dihedral angles at each of the edges leading from $\mathcal{F}$ to the rest of P_i and at at least one opposite pair of edges of $\mathcal{F}$ limit to $\pi/2$, in violation of condition (5).

Since none of the faces of the P_i can degenerate to a point or a line segment, neither can the P_i. Suppose that the P_i degenerate to a polygon, $\mathcal{G}$. Because the dihedral angles are non-obtuse, only two of the faces of the P_i can limit to the polygon $\mathcal{G}$. Therefore the rest of the faces of the P_i must limit to points or line segments, contrary to our reasoning above.

We now suppose that some of the edges of the P_i shrink to length 0 in such a way that the P_i do not shrink to a point, a line segment, or a polygon. Let $v_1, \cdots, v_k$ be a subset of the vertices that converge to some point p with $k > 1$. Then, since the P_i do not shrink to a point, a line segment, or a polygon, there are at least three vertices η, κ and γ that don't converge to p and that don't converge to each-other. Perform the appropriate isometry taking p to the origin in the ball model. Place a small sphere S centered at the origin, so that η, κ and γ never enter S. For large enough i, the intersection $P_i \cap S$ approximates a spherical polygon whose angles approximate the dihedral angles between the faces of P_i that enter S. These spherical polygons cannot degenerate to a point or a line segment because the polyhedra P_i do not degenerate to a line segment or a polygon. By reasoning similar to that of Proposition 1.1, one can check that this polygon must have only three sides and angle sum $> \pi$. The edges of this triangle form a prismatic 3-circuit in C^*, since for each i, P_i has more than one vertex inside the sphere ($k > 1$) and at least the three vertices η, κ and γ outside of the sphere. So, the P_i would have a prismatic three circuit whose angle sum limits to a value $> \pi$. However, this contradicts condition (3) of our supposition that $\mathbf{a_i}$ converges in A_C.

So none of the edges of P_i can limit to zero length hence we have some lower bound ϵ for which each edge of each P_i has length greater than ϵ. $\square$

We now finish the proof of Proposition 5.1. By the above lemma, the vector of edge lengths remains confined in $[\epsilon, \kappa]^E$. We normalize the P_i so that the center of mass of each is at the origin in the conformal ball model, v_1 is on the x-axis, and v_2 is in the xy-plane. Since the edge lengths are confined to $[\epsilon, \kappa]^E$, the vertices of each P_i are confined to a closed ball of finite (hyperbolic) radius. Therefore, we can choose a subsequence of the P_i so that these vertices converge. Because the edge lengths are bounded from below, each of the vertices converges to a distinct point in $\mathbb{H}^3$. Therefore, this subsequence of the P_i converge in $\mathcal{P}'_C$ $\square$

There are two Propositions that will be needed in the next section that use sequences of polyhedra as we used above. Both of them suppose that Andreev's Theorem is satisfied for a specific abstract polyhedron C and they consider two ways of making polyhedra with infinite vertices as limits of sequences of compact polyhedra realizing C. The proofs of both propositions are so similar to the proof above that we include them here.

Proposition 5.4 *Given a abstract polyhedron C having no prismatic 3-circuits for which Andreev's Theorem is satisfied. For any edge e_0 of C, let C_0 be the complex obtained by contracting e_0 to a point. Then, there exists a non-compact polyhedron P_0 realizing C_0 with the edge e_0 contracted to a single vertex at infinity and the rest of the vertices at finite points in $\mathbb{H}^3$.*

Proof: Let v_1 and v_2 be the vertices at the ends of e_0, let e_1, e_2, e_3, e_4 the edges emanating from the ends of e_0, and f_1, f_2, f_3, f_4 be the four faces meeting at any of these edges. See the diagram below.

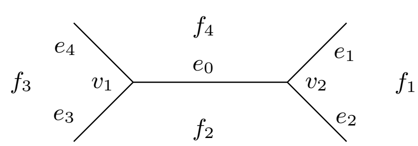

The angles: $\alpha(e_0) = \epsilon$, $\alpha(e_1) = \alpha(e_2) = \alpha(e_3) = \alpha(e_4) = \pi/2$, and $\alpha(e) = 2\pi/5$ for all other edges e, are in A_C since C has no prismatic 3-circuits. Therefore, because we assume that Andreev's Theorem holds for C, there is a polyhedron $P_\epsilon \in \mathcal{P}^0_C$ realizing these angles. Choose a sequence $\epsilon_n > 0$ converging to 0. We will show that the sequence of polyhedra P_{ϵ_n} has a subsequence converging to a non-compact polyhedron P_0 realizing C_0, and hence has the entire edge e_0 shrunken down to a single point at infinity.

Using reasoning very similar to that in Proposition 5.1, one can show that for each edge e that is not one of the e_0, e_1, e_2, e_3, e_4, the edge length must remain bounded away from 0 and ∞. So, one can choose an appropriate subsequence of the P_{ϵ_n} so that each of these edge lengths converges to a non-zero edge length.

Since the dihedral angle at edge e_0 decreases to 0, in the limit, the two faces adjacent along e_0 intersect at a single point at infinity. The vertices v_1 and v_2 are in this intersection, therefore they limit to the same point at infinity. In particular, the edges e_1, e_2, e_3, e_4 connecting v_1 and v_2 to the rest of the polyhedron limit to infinite length.

So, we have a subsequence of the P_{ϵ_n} for which each edge e_1, e_2, e_3, e_4 limits to infinite length, edge e_0 shrinks to a single point at infinity, and all of the other edges limit to finite lengths. As in Proposition 5.1, each of the face angles must converge. So each of these faces converges to a given non-empty (possibly infinite) hyperbolic polygon, and, as in Proposition 5.1, these polygons must fit together to form a polyhedron P_0. P_0 realizes C_0 since it realizes the same cell complex on the sphere as the P_{ϵ_n}, except with the edge e_0 shrunk to a point at infinity. $\square$

Proposition 5.5 *Suppose that Andreev's Theorem is true for C and let $\mathbf{a}_i \in A_C$ be a sequence that converges to $\mathbf{a} \in \partial A_C$, satisfying conditions (1,3-5). If condition (2) is satisfied for vertices $v_1, \cdots, v_k$ of C, but not for $v_{k+1}, \cdots, v_n$ for which the dihedral angle sum is exactly π, then there exists a non-compact polyhedron P_0 realizing C with dihedral angles $\mathbf{a}$. Furthermore, P_0 has vertices $v_1, \cdots, v_k$ at distinct finite points and the $v_{k+1}, \cdots, v_n$ at distinct points at infinity.*

Proof: Because we assume that Andreev's theorem is true for the abstract polyhedron C, there exist unique hyperbolic polyhedra P_i realizing C with dihedral angles given by $\mathbf{a}_i$. In the proof of Lemma 5.3, we did not use condition (2) to establish the lower bound on the lengths of edges of the P_i, so the same proof applies here.

Using the the same methods as from Lemma 5.3 one can situate the P_i in the conformal ball model in such a way that all of the vertices are in neighborhoods N of the north pole and S of the south pole. We will assume, without any loss in generality, that S contains more of the vertices than N does. With exactly the same reasoning as in Lemma 5.3, we can use that the $\mathbf{a}_i$ satisfy conditions (1,3-5) to eliminate any case in which there is more than one vertex in N.

Now, consider each P_i in the conformal ball model normalized so that each is situated with the center of mass at the origin, v_1 on the x-axis, and v_2 in the x-y-plane. (Notice that we do not eliminate the case $k = 0$.) In this way we represent each isometry class of P_i with a fixed polyhedron.

Then, for each P_i the locations of the vertices are defined and each vertex is bounded away from the others in $\overline{\mathbb{H}^3}$. We will denote the locations of the vertices of P_i by $v_1^i, \cdots, v_n^i$. By compactness of the unit ball, we can choose a subsequence of the P_i so that each vertex converges in $\overline{\mathbb{H}^3}$. Because the $v_1^i, \cdots, v_n^i$ are bounded away from each-other, the limit points $v_1, \cdots, v_n$ are distinct. Let P_0 be the polyhedron spanned by $v_1, \cdots, v_n$. By Lemma 3.2, the vertices $v_1, \cdots, v_k$ must be vertices at infinity because the three edges meeting at them have dihedral angle sum π and the vertices $v_k, \cdots, v_n$ are finite vertices because the three edges meeting at them have dihedral angle sum $> \pi$. $\square$

6 $A_C \neq \emptyset$ implies $\mathcal{P}_{\mathcal{C}} \neq \emptyset$

At this point, we know the following result:

Proposition 6.1 *If $\mathcal{P}_C^0 \neq \emptyset$, then $\alpha : \mathcal{P}_C^0 \to A_C$ is a homeomorphism.*

Proof: Indeed, since $\alpha : \mathcal{P}_C \to \mathbb{R}^E$ is continuous and injective, invariance of domain gives that it is a homeomorphism onto its image. Therefore α restricted to $\mathcal{P}_C^0$ is a homeomorphism onto its image. We now use the fact that a local homeomorphism (between metric spaces) that is proper is a finite sheeted covering map. Therefore α is an injective covering map to the connected set A_C. Hence, if $\mathcal{P}_C^0 \neq \emptyset$, $\alpha : \mathcal{P}_C^0 \to A_C$ is a homeomorphism onto A_C. $\Box$

But what is left is absolutely not obvious, and is the hardest part of the whole proof: proving that if $A_C \neq \emptyset$, then $\mathcal{P}_C^0 \neq \emptyset$. We have no tools to approach it and must use bare hands. We follow the proof of Andreev, although the proof of his key lemma was wrong. We provide our own correction.

First recall that in Corollary 1.4, we saw that if C has no prismatic 3-circuits, $A_C \neq \emptyset$. We first prove that $\mathcal{P}_C^0 \neq \emptyset$ for these these *simple polyhedra*, and hence by Proposition 6.1 that Andreev's Theorem is true for simple polyhedra. We then show that for any C having prismatic 3-circuits, if $A_C \neq \emptyset$, then $\mathcal{P}_C^0 \neq \emptyset$ by making a polyhedron realizing C from (possibly many) simple polyhedra. By Proposition 6.1, this final step will finish the proof of Andreev's theorem.

Proof of Andreev's Theorem for Simple Polyhedra

Proposition 6.2 *If C is simple and has $N > 5$ faces, $\mathcal{P}_C^0 \neq \emptyset$. In words: every simple polyhedron is realizable.*

Proof. The proof comprises three lemmas. We will first state the lemmas and prove this proposition using them. Then we will prove the lemmas.

Lemma 6.3 *Let Pr_N and D_N be the abstract polyhedron corresponding to the N faced prism and the N faced "split prism", as illustrated below. If $N > 4$, $\mathcal{P}_{Pr_N}$ is nonempty and if $N > 7$, $\mathcal{P}_{D_N}$ is nonempty.*

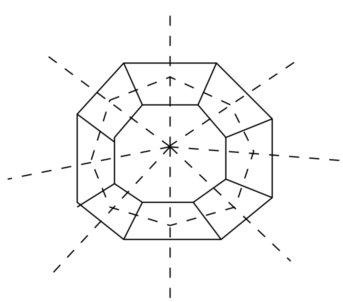

Prism with 10 faces

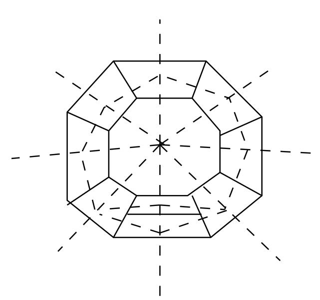

Splitprism with 11 faces

The reader should recall that a Whitehead move on an edge e of an abstract polyhedron is given the local change described the following diagram. The Whitehead move in the dual complex is dashed.

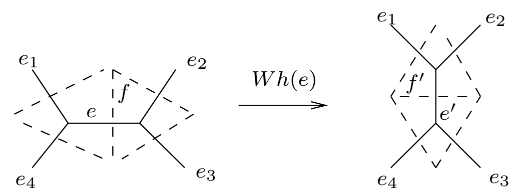

Whitehead move on edge e

Lemma 6.4 *Let the abstract polyhedron C' be obtained from the simple abstract polyhedron C by a Whitehead move Wh_e. Then if $\mathcal{P}^0_C$ is non-empty, so is $\mathcal{P}^0_{C'}$.*

Lemma 6.5 Whitehead Sequence

Let C be a simple abstract polyhedron on $\mathbb{S}^2$ which is not a prism. If C has $N > 7$ faces, one can simplify C by a finite sequence of Whitehead moves to D_N with all of the intermediate abstract polyhedra C_i simple.

Proof of Proposition 6.2, assuming these three lemmas:

Given simple C with $N > 5$ faces; if C is the prism, the statement is proven by Lemma 6.3. One can check that if C has 7 or fewer faces (and is not the tetrahedron) it is a prism. So, if C is not the prism, we have $N > 7$. Then, according to Lemma 6.5, one finds a reduction by (say n) Whitehead moves to D_N, with each intermediate abstract polyhedron simple. Applying Lemma 6.4 n times, one sees that $\mathcal{P}^0_C$ is non-empty if and only if $\mathcal{P}^0_{D_N}$ is non-empty. However $\mathcal{P}^0_{D_N}$ is non-empty by Lemma 6.3. $\square$

The hard technical part of this is the proof of Lemma 6.5. Andreev's original proof [3] provides an algorithm giving the Whitehead moves needed for this lemma but the algorithm *just doesn't work*. It was implemented as a computer program by the author and failed on the first test case, C being the dodecahedron. On one of the final steps, it produced an abstract polyhedron which had a prismatic 3-circuit. This error was then traced back to a false statement in Andreev's proof of the lemma. We will explain the details of this error in the proof of Lemma 6.5.

We proceed to prove the lemmas:

Proof of Lemma 6.3

We construct the N faced prism by explicit construction. First, construct a regular polygon with $N-2$ sides in the disc model for $\mathbb{H}^2$. ($N-2 \geq 3$, since $N \geq 5$.) We can do this with the angles arbitrarily small. Now view $\mathbb{H}^2$ as the equatorial plane of $\mathbb{H}^3$; and consider the hyperbolic planes perpendicular to the equatorial plane containing the sides of the polygon. In Euclidean geometry these are hemispheres with centers on the boundary of the equatorial disc. The dihedral angles of these planes are the angles of the polygon.

Now consider two hyperbolic planes close to the equatorial plane, one slightly above and one slightly beneath, both perpendicular to the z-axis. These will intersect the previous planes at angles slightly smaller than $\pi/2$. The region defined by these N planes makes a hyperbolic polyhedron realizing the cell structure of the prism.

Now for D_N with $N > 7$ faces: we will get this by cutting D_N into two prisms and using Prop. 6.1, to realize these prisms with appropriate angles so that they fit together to give D_N. Consider the prism having $N - 1$ faces with dihedral angles as labeled below.

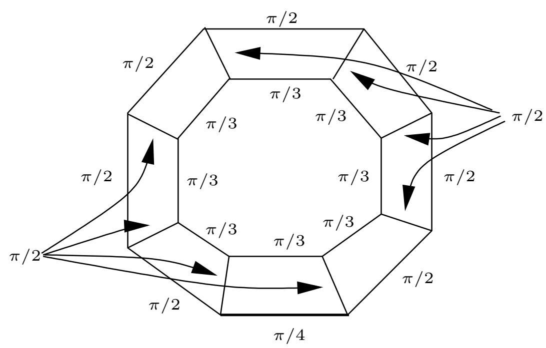

These angles satisfy Andreev's conditions (1 - 5), and since for the $N - 1 > 6$ prism we have that $\mathcal{P}^0_{Pr_{N-1}}$ nonempty, this prism exists. When two are glued together, the edges labeled $\pi/2$ on the outside disappear as edges, and the edges labeled on the outside by $\pi/4$ glue together becoming an edge with dihedral angle $\pi/2$. Hence, we have constructed a polyhedron realizing with D_N (assuming $N >$ 7.) (Note: one can apply this construction when $N = 7$, but the result is in fact Pr_7. D_7 is combinatorially equivalent to Pr_7.)

Proof of Lemma 6.4

We are given C and C' simple with C' obtained by a Whitehead move on the edge e_0 and we are given that $P_C \neq \emptyset$. Since $P_C \neq \emptyset$, we conclude that Andreev's Theorem is satisfied for the abstract polyhedron C, by Proposition 6.1. Let C_0 be the complex obtained from C by shrinking the edge e_0 down to a point. By Proposition 5.4, there exists a non-compact polyhedron P^0 realizing C^0 since Andreev's Theorem holds for C.

Having established the existence of P_0 it is easy to prove the lemma. Let us use the upper half-space model of $\mathbb{H}^3$, and normalize so that e_0 has collapsed to the origin of $\mathbb{C} \subset \partial\mathbb{H}^3$. The faces incident to e_0 are carried by 4 planes $H_1, ..., H_4$ each intersecting the adjacent ones at right angles, and all meeting at the origin. Their configuration will look like the center of the following figure. (Recall that planes in the upper half-space model of $\mathbb{H}^3$ are hemispheres which intersect $\partial\mathbb{H}^3$ in their boundary circles. The dihedral angle between a pair of planes is the angle between the corresponding pair of circles in $\partial\mathbb{H}^3$.)

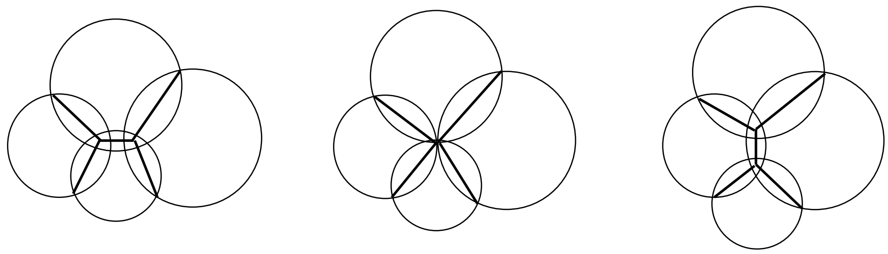

The pattern of circles in the center of the figure can by modified forming the figures on the left and the right with each of the four circles intersecting the adjacent two circles orthogonally. If we leave the other faces of P^0 fixed we can make a small enough modification that the edges e_1, e_2, e_3, e_4 still have positive length and the vertices at the far ends of these edges remain finite. Since each of the dihedral angles corresponding to edges other than e_0, e_1, e_2, e_3, and e_4 were chosen to be $2\pi/5$, this small modification will not increase any of these angles past $\pi/2$.

Therefore, these modified patterns of intersecting circles correspond to polyhedra are elements of $\mathcal{P}^0_C$ and $\mathcal{P}^0_{C'}$. Therefore since $\mathcal{P}^0_C \neq \emptyset$, we have also $\mathcal{P}^0_{C'} \neq \emptyset$. $\square$

Proof of Lemma 6.5 We assume that C is a simple abstract polyhedron having $N > 7$ faces which is not Pr_N. We will construct a sequence of Whitehead moves which change C to D_N with each intermediate complex having no prismatic 3-circuits.

Find a vertex v_∞ of C^* which is connected to the greatest number of other vertices, k. We will call its link, a cycle of k vertices and k edges the *outer-polygon*. Most of the work is to show that we can do Whitehead moves to increase k to $N-3$ without introducing any prismatic 3-circuits during the process. Once this is completed, it will be easy to change the resulting complex to D^*_N by additional Whitehead moves.

Let's set up some notation. We will draw the dual complex C^* in the plane with the vertex v_∞ at infinity and the outer polygon, P, surrounding the remaining vertices, and triangles. We will call the vertices inside of P *interior vertices*. We will call all of the edges inside of P which don't have an endpoint on P *interior edges*. The graph of interior vertices and edges is connected, since C^* is simple. An interior vertex which is connected to only one other interior vertex will be called an *endpoint*. See below:

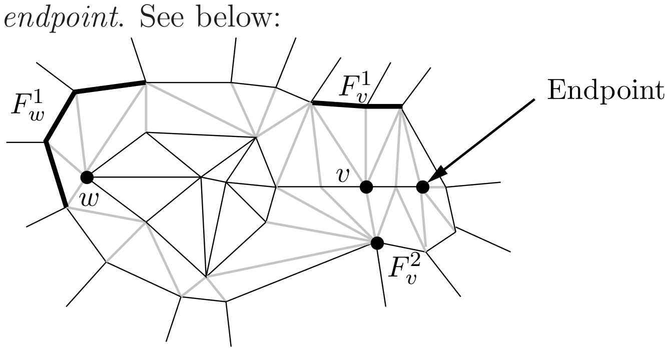

Throughout this proof we will draw P in black and we draw interior edges and vertices of C^* in black, as well. The connections between P and the interior vertices will be in grey. Connections between P and v_∞ will be black, if shown at all.

The link of an interior vertex v will intersect P in a number of components $F_v^1, \cdots, F_v^n$. (Possibly $n = 0$.) See the above figure. We will say that v is *connected to* P *in these components.* Notice that since C^* is simple, an endpoint is always connected to P in exactly one such component.

Sub-lemma 6.6 *If a Whitehead move on* C_1^* *results in the complex* C_2^* *(replacing* e *by* e'*), and if* δ *is a simple closed path in* C_1^**, which separates one endpoint of* e' *from the other, then any newly-created 3-circuit will contain some vertex of* δ *which shares edges with both endpoints of* e'*.*

Proof: A newly created prismatic 3-circuit γ would consist of the new edge e' as well as two additional edges e_1 and e_2 connecting from a single vertex V to the two endpoints of e'. By the Jordan Curve Theorem, the path $e_1 e_2$ connecting the two endpoints of e' must intersect δ since δ separates these endpoints. The edges e_1 and e_2 would have to have been edges in C_1 since e_0' is the only new edge in C_2. Since each of these paths is made from edges in C_1, the vertex V must be a vertex on δ. $\square$

We will now use this sub-lemma to prove three additional sub-lemmas which specify certain operations by Whitehead moves that can be done without introducing prismatic 3-circuits.

Sub-lemma 6.7 *Suppose that in the dual of some simple polyhedron, there exists a vertex* A *which is connected to* P *in exactly one component consisting of exactly two consecutive vertices* Q *and* R*. The Whitehead move Wh(QR) results in a complex outer-polygon has been increased in length by one, which is also the dual of a simple polyhedron.*

Proof:

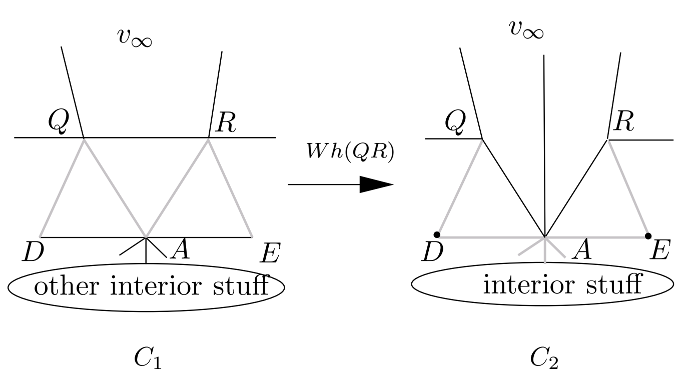

Clearly this Whitehead move increases the length of P by one. We apply Sub-Lemma 6.6 to see that this move introduces no prismatic 3-circuits. We let

$\delta = P$, the outer polygon, which clearly separates the interior vertex A from v_∞ in C_1. Any new prismatic 3-circuit would consist of a point on P connected to both A and v_∞. However, by hypothesis, there were only the two points Q and R on P connected to A. These two points result in the new triangles AQv_∞ and ARv_∞ therefore they don't result in prismatic 3-circuits. We conclude that this Whitehead move introduces no prismatic 3-circuits. □

In the above sub-lemma, the condition that A is connected to exactly two consecutive vertices of P prevents A from being an endpoint. If A were an endpoint, let B be the unique interior vertex that A is connected to. In this case the three circuit BQR would surround A and hence would be a prismatic 3-circuit in C_1. Therefore any endpoint must be connected to P in a single component having three or more vertices.

Sub-lemma 6.8 *Given an interior vertex A which is connected to P in a component consisting of M consecutive vertices $Q_1, \cdots, Q_M$ of P (and possibly other components).*

1. *If A is not an endpoint and $M > 2$, the sequence of Whitehead moves $Wh(AQ_M), \cdots, Wh(AQ_3)$ result in a complex in which A is connected to this component of P in only Q_1 and Q_2. These moves leave P unchanged, and introduce no prismatic 3-circuits.*

2. *If A is an endpoint and $M > 3$, the sequence of Whitehead moves $Wh(AQ_M), \cdots, Wh(AQ_4)$ result in a complex in which A is connected to this component of P in only Q_1, Q_2, and Q_3. These moves leave P unchanged and introduce no prismatic 3-circuits.*

Proof: If A is not an endpoint we have the following setup:

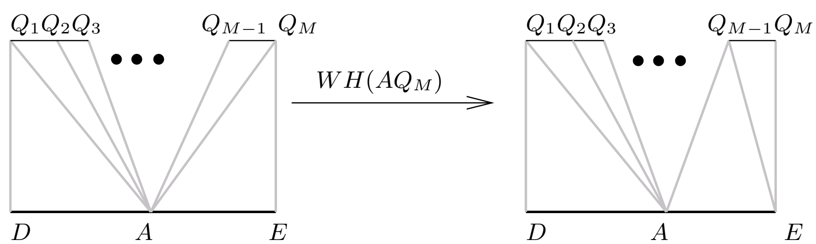

Part of C_1 Part of C_2

Clearly the move $Wh(AQ_M)$ decreases M by one. We check that if $M > 2$, this move introduces no prismatic 3-circuits. We let δ be that path $v_\infty Q_{M-2} A Q_M$ which separates Q_{M-1} and E in C_1. By Sub-lemma 6.6, any new prismatic 3-circuit would contain a vertex on δ that connects to both E and Q_{M-1}. Clearly v_∞ is not connected to the interior vertex E. Also, a connection of Q_{M-2} to E would correspond to a pre-existing prismatic 3-circuit $EQ_{M-2}A$ in C_1, contrary to the assumption that C_1 was simple. So the only two vertices on δ that are connected to both E and Q_{M-1} are A and Q_M, but these connections form the two triangles $AQ_{M-1}E$ and $Q_{M-1}Q_ME$ in C_2. Hence there are no new prismatic 3-circuits. Since we can always reduce M by one, when $M > 2$, we can reduce M to 2.

If A is an endpoint we have the following setup:

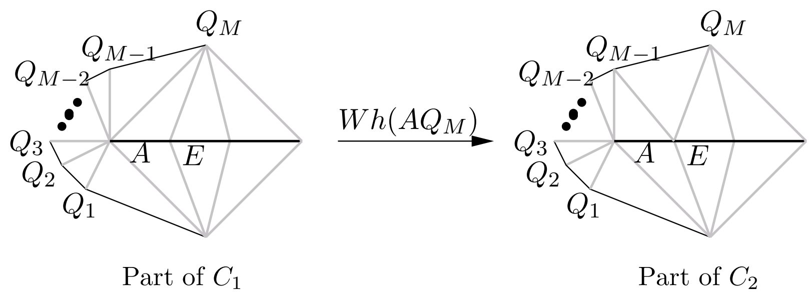

So the move $Wh(AQ_M)$ decreases M by one. We check that as long as $M > 3$, this move introduces no prismatic 3-circuits. Let δ be the curve $v_\infty Q_M A Q_{M-2}$ which separates Q_{M-1} from E in C_1 which are the ends of the new edge in C_2. By Sub-lemma 6.6, any new prismatic 3-circuit would contain of a vertex on δ that is connected to both Q_{M-1} from E. Clearly v_∞ is not connected to E since E is interior. Also, since $M > 3$, Q_{M-2} is not connected to E. So the only vertices on δ that are connected to E are Q_M and A. However these connections form the triangles $AQ_{M-1}E$ and $Q_{M-1}Q_ME$, hence they do not correspond to prismatic 3-circuits.

So, as long as $M > 3$ we can reduce M by one without introducing prismatic 3-circuits. Hence, we can reduce M to 3. Recall that an endpoint of a simple complex cannot be connected to fewer than three points of P, so this is the best that we can hope to do. $\square$

Note: In both parts (1) and (2), each of these Whitehead moves $Wh(AQ_M)$ transfers the connection between A and Q_M to a connection between the neighboring interior vertex E and Q_M. This will be essential later in the proof (Case 2.) where we will need to increase the number of vertices in a component where E is connected to P by decreasing the number of vertices in a component where A is connected to P.

Sub-lemma 6.9 *If an interior vertex A is connected to P in two points X and Y, one can do Whitehead moves to eliminate any components of connections of A to P that do not contain X or Y without introducing any new prismatic 3-circuits.*

Example: Here A is connected to P in four components containing six points. We can eliminate connections of A to all of the components except for the single-point components X and Y. Note: It is essential later in the proof of Lemma 6.5 that the hypotheses of this Sub-lemma do not require that X and Y be in the same component. (However, we often use this Sub-lemma in the simpler case where X and Y are in the same component.)

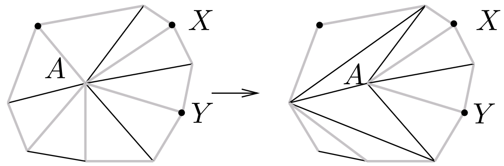

Proof: Notice that since A is an interior vertex connected to more than one component on P, it is not an endpoint. In particular, there will be unique interior vertices D and E forming triangles ADW and AEV as in the figure below.

Let O be a component that doesn't contain X or Y which we wish to eliminate. If O contains more than two vertices, we can reduce it to two vertices V and W by Sub-lemma 6.8, part (1). Having done so, we can do Whitehead moves $Wh(AW)$ and $Wh(AV)$ to eliminate these final connections. This is slightly delicate, and requires in an essential way that A is connected to P in at least two other vertices.

First, one does $Wh(AW)$ resulting in the elimination of edge AW and the creation of the new edge DV as in the diagram below. We will use Sub-lemma 6.6 to check that this move introduces no new prismatic 3-circuits. Let δ be the loop $v_\infty YAW$ which separates D from V in C_1. See the dashed curve in the figure below. Any new prismatic 3-circuit would contain a point on δ that is connected to both D and V. Clearly v_∞ is not connected to D since D is interior. Also, Y is not connected to W, since Y and W are in different components of connection between A and P, by hypothesis. So, only A and W are connected to D and to V, but they form the triangles ADV and WVD in C_2, hence do not correspond to prismatic 3-circuits. Therefore, $Wh(AW)$ results in no prismatic 3-circuits.

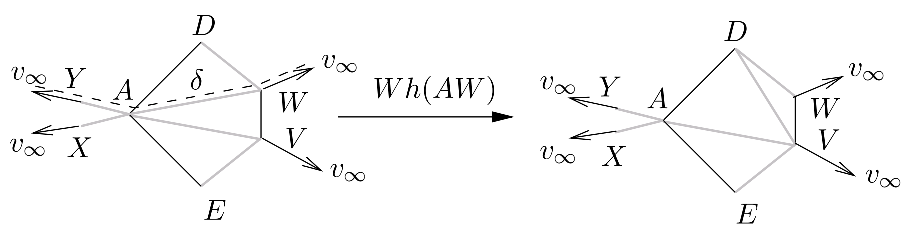

Part of C_1 Part of C_2

Second, one does $Wh(AV)$, as depicted below. Let δ_1 be the curve $v_\infty YAV$ and δ_2 be the curve $v_\infty XAV$ in C_1. See the two dashed curves in the figure below. Both of these curves separate D and E in C_1. So, applying sub-lemma 6.6 twice, we conclude that any newly created prismatic 3-circuit would contain a point that is *both on* δ_1 *and on* δ_2 and that connects to both D and E. The only points on both δ_1 and δ_2 are v_∞, A, and V. Since D and E are interior, v_∞ cannot connect to either of them. The connections from A and from V to D and E result in the triangles ADE and VDE, hence do not result in prismatic 3-circuits. Therefore, we conclude that $Wh(AV)$ results in no prismatic 3-circuits. □

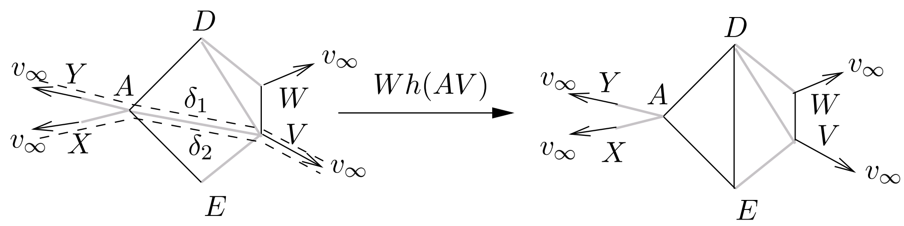

Part of C_1 Part of C_2

In his paper [3], on pages 433 and 434, Andreev describes a nearly identical process to Sub-lemma 6.9 for decreasing the number of components in which an interior point A is connected to P. However, he merely assumes that A is connected to P in at lease one component in addition to the components being eliminated. He does not require, as we have, that A is connected to P in at least *two points* outside of the components being eliminated. Andreev asserts: "It is readily seen that all of the polyhedra obtained in this way are simple..." However this assertion is incorrect. Consider the case below where A is connected to P in two components, the two points M and N. Doing the Whitehead move $Wh(AN)$ eliminates the connection of A to P at N but also creates the prismatic 3-circuit DEM surrounding A, which is shown as the dashed curve in the right hand side of the figure.

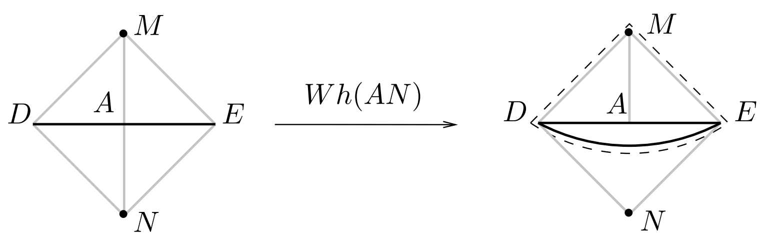

Having assumed this assertion, the remainder of Andreev's proof is relatively easy. Unfortunately, the situation pictured above is not uncommon (as we will see in Case 3 below)! Using Sub-lemma 6.9 we will have to work a little bit harder.

We will now use these three sub-lemmas to show that if the length of P is less than $N-3$ (so that there are at least 3 interior vertices), then we can do Whitehead moves to increase the length of P by one, without introducing any prismatic 3-circuits.

Case 1: An interior point which isn't an endpoint connects to P in a component with two or more vertices, and possibly in other components, as well.

Apply Sub-lemma 6.8 decreasing this component to two vertices. We can then apply Sub-lemma 6.9, eliminating any other components since this component contains two vertices. Finally, apply Sub-lemma 6.7 to increase the length of the outer polygon by 1.

Case 2: An interior vertex that is an endpoint is connected to more than three vertices of P.

We assume that each of the interior points that are not endpoints are connected to P in components consisting of single points, otherwise we are in Case 1.

Let A be the endpoint which is connected to more than three vertices of P. By Sub-lemma 6.8, part (2), we can do a Whitehead move to transfer one of these connections to the interior vertex E that is next to A. Now, one of the components in which E is connected to P has exactly two vertices. The point E is not an endpoint since $k < N - 3$ implies that there are at least three interior vertices. Once this is done, we can apply Case 1.

Case 3: Each interior point which is an endpoint is connected to exactly 3 points of P and each interior point which is not an endpoint is connected to P in components consisting of single points.

First, notice that if the interior vertices and edges form a line, this restriction on how interior points are connected to P results in the following complex, which is the prism:

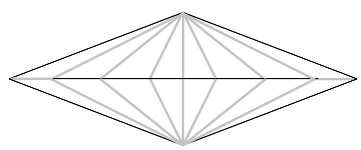

This case is ruled out by our assumption that C is not the prism. However, there are plenty of complexes satisfying the hypotheses of this case which have interior vertices and edges forming a graph more complicated than a line:

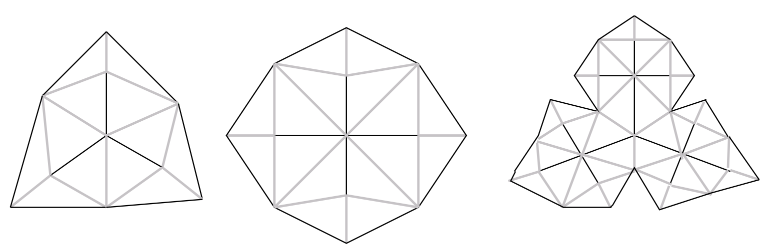

For such complexes we need a very special sequence of Whitehead moves to increase the length of P.

Pick an interior vertex which is an endpoint and label it I_1. Denote by P_1, P_2, and P_3 the three vertices of P to which I_1 connects. I_1 will be connected to a linear sequence of interior vertices $I_2, I_3, \cdots I_m, m \geq 2$, with I_m the first interior vertex in the sequence that is connected to more than two other interior vertices. Vertex I_m must exist since we are assuming that the interior vertices don't form a line, a configuration that we ruled out above. By hypothesis, $I_2, \cdots, I_m$ can only connect to P in components which each consist of a vertex, hence each must be connected to P_1 and to P_3. Similarly, there is an interior vertex (call it X) which connects both to I_m and to P_1 and another vertex Y which connects to I_m and P_3. Vertex I_m may connect to other vertices of P and other interior vertices, as shown on the left side of the following diagram, which depicts the general situation in a neighborhood of an endpoint, in Case 3.

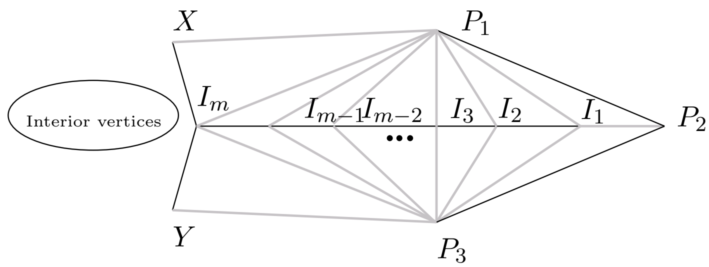

Now we describe a sequence of Whitehead moves that can be used to connect I_m to P in only P_1 and P_2, which will allow us to use Sub-lemma 6.7 to increase the length of P by one.

First, using Sub-lemma 6.9, one can eliminate all possible connections of I_m to P in places other than P_1 and P_3. Next, one does the move $Wh(I_mP_3)$ so that I_m connects to P in only one vertex, P_1. We check that this Whitehead move does not create any prismatic 3-circuits. Let δ be the curve $v_\infty P_1 I_m P_3$ separating I_{m-1} from Y. By Sub-lemma 6.6, any newly created prismatic 3-circuit would contain a point on δ connected to both I_{m-1} and Y. Since Y and I_{m-1} are interior, v_∞ does not connect to them. Also, P_1 is not connected to Y as this would correspond to a pre-existing prismatic 3-circuit P_1I_mY, contrary to assumption. So, the only vertices of δ that are connected to both I_{m-1} and Y are I_m and P_3, which result in the triangles $I_mI_{m-1}Y$ and P_3YI_{m-1}, hence do not correspond to newly created prismatic 3-circuits. We conclude that $Wh(I_mP_3)$ introduces no prismatic 3-circuits.

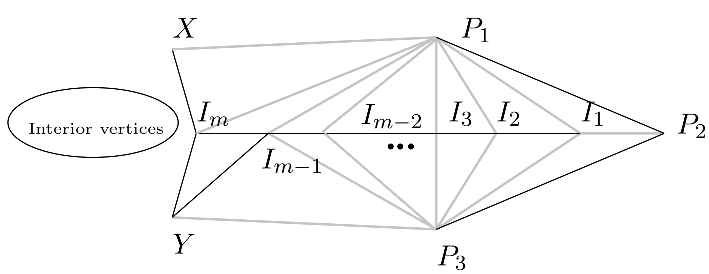

Next, one must do the moves $Wh(I_{m-1}P_1)$,...,$Wh(I_1P_1)$, in that order. We check that each of these moves creates no prismatic 3-circuits (see the figure below). Fix $1 \leq l \leq m-1$, and let δ be the loop $v_\infty P_1 I_1 P_3$. $Wh(I_1P_1)$ creates a new edge $I_{l-1}I_m$ if $l > 1$, or P_2I_m if $l = 1$, the vertices of which are separated by δ. Since I_m is interior, v_∞ does not connect to I_m. Also, I_m is no longer connected to P_3. Therefore the only points of δ that are both connected to I_m and I_{l-1} are I_l and P_1. The connections form the new triangles $P_1I_mI_{l-1}$ and $I_lI_{l-1}I_m$, hence no prismatic 3-circuits (when $l = 1$, the above is true with P_2 in place of I_{l-1}). So the move $Wh(I_lP_1)$ introduces no prismatic 3-circuits. Hence we can do each of the moves listed above.

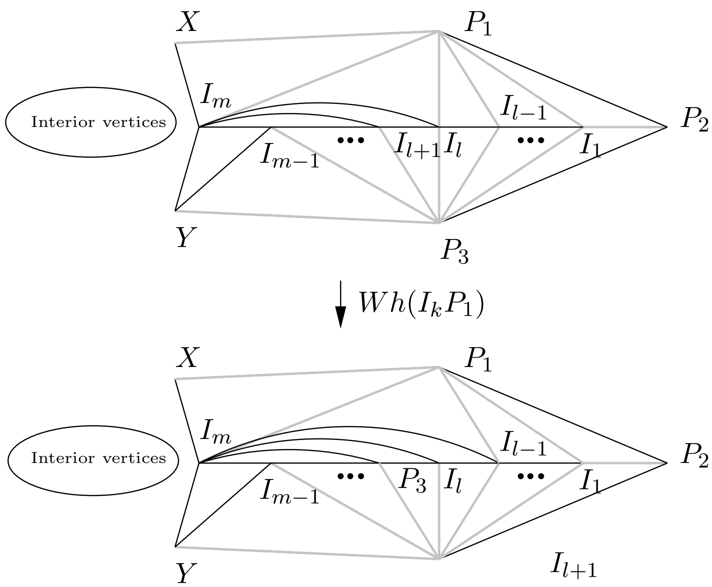

When we finish this sequence of Whitehead moves, we obtain the diagram shown below, with I_m connected to P exactly at P_1 and P_2, so that one can apply Sub-lemma 6.7 to increase the length of P by the move $Wh(P_1P_2)$, also shown below.

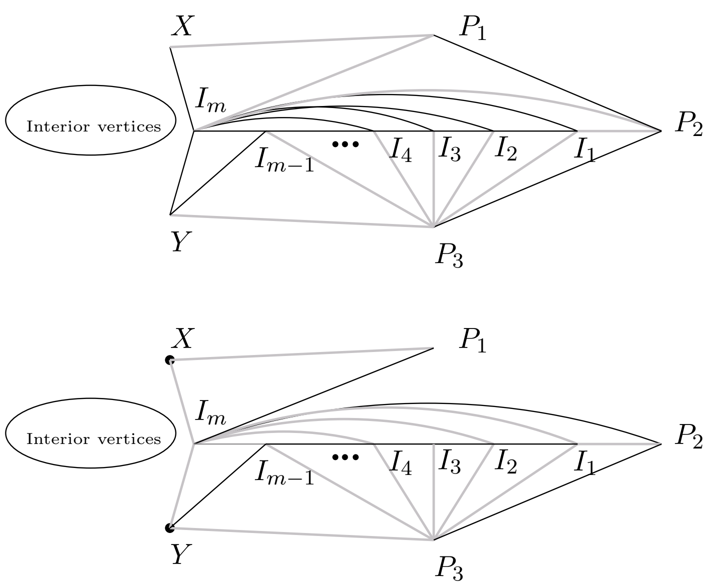

This concludes Case 3.

Since C^* must belong to one of these cases, we have seen that if the length of P is less than $N-3$ we can do Whitehead moves to increase its length to $N-3$ without creating prismatic 3-circuits. Hence we can reduce to the case of two interior vertices, as shown in the diagram below on the left. Since $N > 7$ such a diagram exists, without prismatic 3-circuits.

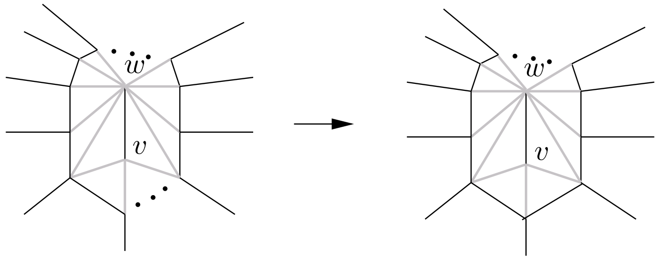

Each of these interior vertices v and w will be endpoints and hence each connected to P in a single component, F_v and F_w, with each of these components containing three or more vertices. One can reduce to the desired complex D_N^* by choosing one of the interior vertices, say v and

This completes the proof of Lemma 6.5, and hence of Proposition 6.2. applying Sub-lemma 6.8, part (2), to reduce F_v until it contains only three vertices. See the right side of the above diagram. $\square$

Proof of Andreev's Theorem for general polyhedra

So, we have seen that Andreev's Theorem is true for every simple (3-aprismatic) abstract polyhedron C. Now we consider the case of C having prismatic 3-circuits. (Of these, at this point we only know that the triangular prism exists.) Recall that there are some such C for which $A_C = \emptyset$, so we can only hope to prove that $\mathcal{P}_C^0 \neq \emptyset$ when $A_C \neq \emptyset$. The following lemma may help one check whether $A_C \neq \emptyset$.

Lemma 6.10 *If $A_C \neq \emptyset$, then there are points in A_C arbitrarily close to $(\pi/3, \pi/3, \cdots, \pi/3)$.*

Proof: Let $\mathbf{a} \in A_C$ and let $\mathbf{a_t} = \mathbf{a}(1-t) + (\pi/3, \pi/3, \cdots, \pi/3)t$. For each $t \in [0, 1)$, $\mathbf{a_t} \in A_C$. We check conditions (1-5): Each component is clearly in $(0, \pi/2]$, so condition (1) is satisfied. Given edges e_i, e_j, e_k meeting at a vertex we have $(\mathbf{a_i} + \mathbf{a_j} + \mathbf{a_k})(1-t) + \pi t > \pi(1-t) + \pi t = \pi$ for $t < 1$, since $(\mathbf{a_i} + \mathbf{a_j} + \mathbf{a_k}) > \pi$. So, condition (2) is satisfied. Similarly, given a prismatic 3-circuit intersecting edges e_i, e_j, e_k we have $(\mathbf{a_i} + \mathbf{a_j} + \mathbf{a_k})(1-t) + \pi t < \pi(1-t) + \pi t = \pi$ for $t < 1$, so condition (3) is satisfied. Conditions (4) and (5) are satisfied since each component of $\mathbf{a_t}$ is $< \pi/2$ for $t > 0$ and since $\mathbf{a}$ satisfies these conditions for $t = 0$. $\square$

We will distinguish two types of prismatic 3-circuits. If a prismatic three circuit in C^* separates one point from the rest of C^*, we will call it a *truncated triangle*, otherwise we will call it an *essential 3-circuit*. The name truncated triangle comes from the fact that such a 3-circuit in C^* corresponds geometrically to the truncation of a vertex, forming a triangular face. We will first prove the following sub-case:

Proposition 6.11 *Let C be an abstract polyhedron in which every prismatic 3-circuit in C^* is a truncated triangle. If A_C is non-empty, then $\mathcal{P}_C^0$ is non-empty.*

We will need the following three elementary lemmas in the proof:

Lemma 6.12 *Given three planes in $\mathbb{H}^3$ that intersect pairwise, but which do not intersect at a point in $\overline{\mathbb{H}^3}$, there is a fourth plane that intersects each of these planes at right angles.*

Proof: Suppose that the three planes are given by $P_{\mathbf{v_1}}, P_{\mathbf{v_2}}$, and $P_{\mathbf{v_3}}$. The line $P_{\mathbf{v_1}} \cap P_{\mathbf{v_2}} \cap P_{\mathbf{v_3}}$ in $E^{3,1}$ is outside of the light-cone, so the hyper-plane $(P_{\mathbf{v_1}} \cap P_{\mathbf{v_2}} \cap P_{\mathbf{v_3}})^{\perp}$ intersects $\mathbb{H}^3$ and hence defines a plane orthogonal to $P_{\mathbf{v_1}}, P_{\mathbf{v_2}}$, and $P_{\mathbf{v_3}}$. $\square$

Lemma 6.13 *Given two circles C_1 and C_2 in the Euclidean plane that intersect with a non-obtuse exterior angle. Decreasing either or both radii of C_1 and C_2 while keeping their centers fixed decreases the angle of intersection between C_1 and C_2.*

Proof: See the diagram below:

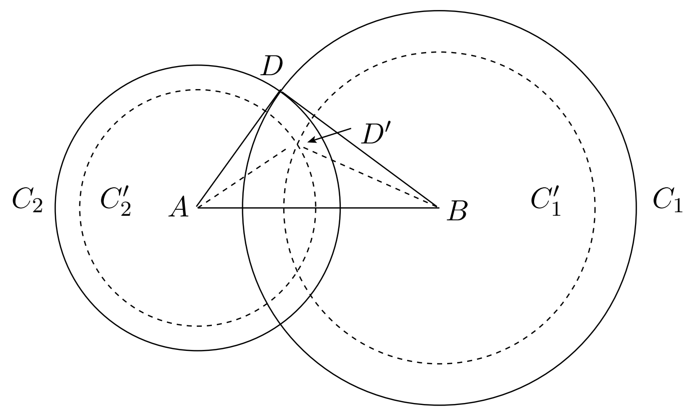

We wish to show that the exterior angle α between lines tangent to C_1 and C_2 at D is larger than the exterior angle α' tangent to C_1' and C_2' at D'. Let β be the angle at vertex D of triangle ADB and β' be the angle at vertex D' of triangle $AD'B$. Elementary geometry shows that $\alpha = \pi - \beta$ and $\alpha' = \pi - \beta'$. So, we must only check that decreasing the radius of either C_1 or C_2 or both, as in the diagram above, results in $\beta' > \beta$. Notice that if the initial exterior angle $\alpha \leq \pi/2$ gives that $\beta \geq \pi$ which gives that the region inside of both circles C_1 and C_2 and above line segment AB is exactly inside of triangle ADB. The vertex D' must be strictly inside of this region and hence inside of the triangle ADB. However, this gives that $\beta' > \beta$. $\square$

Lemma 6.14 *Given a non-compact hyperbolic polyhedron P with dihedral angles in $(0, \pi/2]$ and with trivalent vertices. Suppose that the vertices $v_1, ..., v_n$ are at distinct points at infinity and the rest of the vertices are at finite points in $\mathbb{H}^3$. Then there exists a polyhedron P' with $v_1, ..., v_n$ truncated by triangular faces orthogonal to their adjacent faces and the remaining vertices at finite points in $\mathbb{H}^3$. Each of the dihedral angles of P' will be in $(0, \pi/2]$.*

Proof: Suppose that P is in the upper half-space model of $\mathbb{H}^3$ so that it is completely determined by the pattern of circles where the planes carrying its faces intersect the plane at $z = 0$. (As in the proof of Lemma 6.4.) One can do this with one of the planes corresponding to a large circle, K_1, with the the centers of the circles, $K_2, \cdots, K_N$, which correspond to the other faces, entirely within K_1. Such a configuration is drawn on the left hand side of the figure below. The infinite vertices of P will correspond to three circles intersecting in a single point because we have specified that each infinite vertex is at a distinct point and trivalent. Consider the effect of decreasing the radius of each of the circles $K_2, \cdots, K_N$ by a small number ϵ while keeping the centers of each of the circles and the radius of K_1 fixed. Because the circles intersect with non-obtuse dihedral angles, Lemma 6.13 gives that the result is a small decrease of the dihedral angles between each intersecting pair of planes. Clearly one can choose ϵ small enough to achieve an arbitrarily small decrease in the dihedral angles between pairs of intersecting planes.

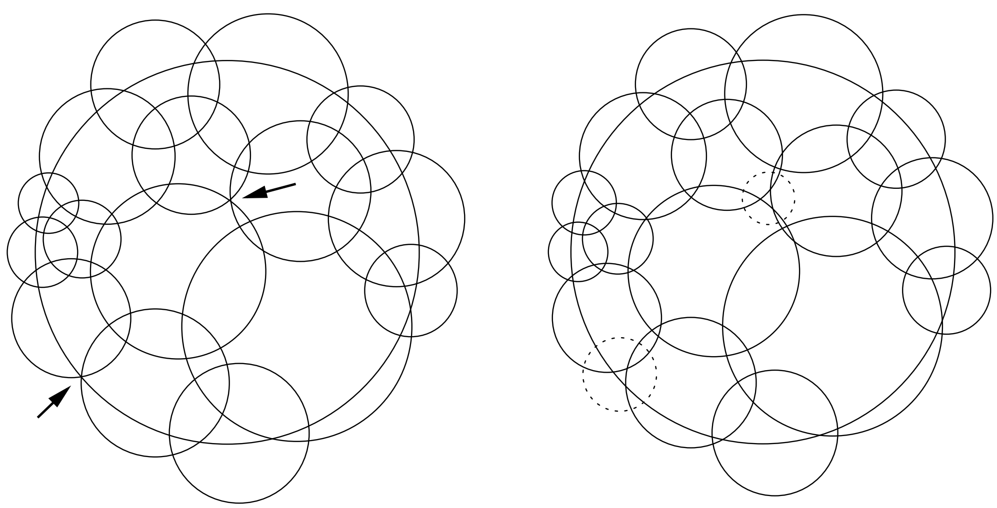

By Lemma 3.2, the sum of the dihedral angles at each of the finite vertices of P is $> \pi$, whereas the sum of the dihedral angles at each of the infinite vertices $= \pi$. We can choose ϵ small enough that each of the dihedral angles is decreased by a amount small enough that the sum of the dihedral angles between triples of faces that correspond to finite vertices of P remains $> \pi$ and so that the dihedral angles between each pair of intersecting circles remains > 0. The sum of dihedral angles between triples of faces that intersect at infinite vertices in P becomes $< \pi$. Therefore, by Lemma 3.2, the triples of faces that intersected at finite vertices in P intersect at finite vertices in the resulting pattern of circles and the triples of faces that intersected at infinite vertices in P do not intersect at a vertex, finite or infinite, but still intersect pairwise. By lemma 6.12, for each triple of planes that previously intersected at an infinite vertex there exists a plane that is perpendicular to each of them. The circles corresponding to these perpendicular planes are drawn in dashed lines in the pattern of circles on the right hand side of the figure above.

This resulting pattern of circles corresponds to the polyhedron P' which has the same combinatorics as P, except that each of the infinite vertices of P is replaced by a triangular plane perpendicular its three adjacent faces. By construction we have that the dihedral angles of P' are in $(0, \pi/2]$. $\square$

Proof of Proposition 6.11.

The abstract polyhedron C cannot correspond to the tetrahedron since the triangular faces of the tetrahedron do not correspond to prismatic circuits. If C corresponds to the triangular prism, $\mathcal{P}_C^0$ is non empty by Lemma 6.3. So, we can restrict to the case where C has more than 5 faces. In this case, one can replace all (or all but one) of the truncated triangles by single vertices, as shown in the diagram below, to reduce C to either Pr_5 or a simple abstract polyhedron. In either case we call the resulting abstract polyhedron C^0. (It can happen that replacing all of the truncated triangles of C by single vertices results in the tetrahedron, this is exactly the case where one must instead replace all but one of the truncated triangles by vertices which results in $C^0 = Pr_5$, the once truncated tetrahedron.)

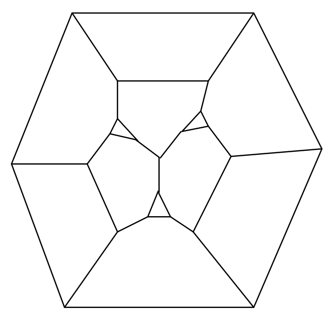

C

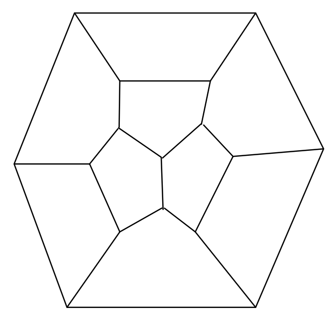

C^0

Using that $A_C \neq \emptyset$ and Lemma 6.10, choose a point $\beta \in A_C$ so that each component of β is within δ of $\pi/3$, with $\delta < \pi/18$. It will be convenient to number the edges of C and C^0 in the following way: If there is a prismatic 3-circuit in C^0 (i.e. $C^0 = Pr_5$), we number these edges $1, 2,$ and 3 in C and C^0. (Otherwise, we just pick three edges of C^0 which do not meet at a vertex.) Next, we number the remaining edges common to C and C^0 by $4, 5, \cdots, k$. Finally, we number the edges of C that were removed to form C^0 by $k+1, \cdots, n$ so that the edges surrounded by prismatic three circuits of C with smaller angle sum (given by β) come before those surrounded by prismatic three circuits with larger angle sum.

The point $\gamma = (\beta_1, \beta_2, \beta_3, \beta_4 + 2\delta, \beta_5 + 2\delta, ..., \beta_k + 2\delta)$ satisfies $\gamma \in A_{C^0}$. We check conditions (1-5). Each of the dihedral angles specified by γ is in $(0, \pi/2)$ because $0 < \beta_i + 2\delta < \pi/3 + 3\delta < \pi/3 + \pi/6 = \pi/2$. Therefore, condition (1) is satisfied, as well as conditions (4) and (5) because the angles are acute. Two of the edges labeled 4 and higher will enter any vertex of C^0 so the sum of the three dihedral angles at each vertex is at least 4δ greater than the sum of the three dihedral angles given by β, which is $> \pi - 3\delta$. Therefore condition (2) is satisfied. If there is a prismatic 3-circuit in C^0, it crosses the first three edges of C^0 and is also a prismatic 3-circuit in C. By our definition of γ, $\gamma_1 + \gamma_2 + \gamma_3 = \beta_1 + \beta_2 + \beta_3 < \pi$, if this is a prismatic 3-circuit, so condition (3) is satisfied.

Now define $\alpha(t) = (1-t)\gamma + t(\beta_1, ...\beta_k)$. Let $T_0, ..., T_{l-1} \in (0,1)$ be the values of t at which there is a vertex of C^0 for which $\alpha(t)$ gives an angle sum of π. (For pedantic reasons, we set $T_{-1} = 0, T_l = 1$.) In fact, this may happen for many vertices simultaneously. We will label the vertices that have angle sum π at T_i by $v_1^i, \cdots, v_{m(i)}^i$. Let C^{i+1} be C^i with $v_1^i, \cdots, v_{m(i)}^i$ truncated for $i = 0, \cdots, l-1$. (This way, C^l is C.)

Suppose that C^i has n_i more edges than C^0. We will conveniently miss-use notation and set $\alpha(t) = (\alpha_1(t), \cdots, \alpha_k(t), \pi/2, \cdots, \pi/2)$, with the last n_i components being $\pi/2$.

We know that Andreev's Theorem is true for C^0 because C^0 is either simple, or the 3-prism. So, it will be sufficient to show that if Andreev's Theorem is satisfied for C^i then it is satisfied for C^{i+1}, for each $i = 0, \cdots, l-1$. To do this, we must generate a polyhedron realizing C^i with the vertices $v_1^i, ..., v_{m(i)}^i$ at infinity and the other vertices a finite points in $\mathbb{H}^3$. This will be easy with our definition of $\alpha(t)$ and Proposition 5.5. We will then use Lemma 6.14 to truncate these vertices.

To use Proposition 5.5, we must check that $\alpha(t) \in A_{C^i}$ when $t \in (T_{i-1}, T_i)$. This follows almost directly from how $\alpha(t)$ is defined. To check condition (1), notice that both β_j and γ_j are non-zero and non-obtuse, for each j, so $\alpha_j(t)$ must be as well. Any vertex of C^i is either a vertex of both C and C^0, such a vertex has dihedral angle sum $> \pi$ because it does for β and γ, or it is a truncated triangle of C but not of C^i and therefore has dihedral angle sum $> \pi$ by definition of T_i, or it is a vertex of C^i on one of the truncated triangles, in which case two of the edges entering this vertex have dihedral angles $= \pi/2$. In each of these cases, condition (2) is satisfied.

Notice that any prismatic 3-circuit in C^i is is either a prismatic 3-circuit in both C^0 and C (the special case where C^0 is the triangular prism) or is a prismatic circuit of C^i which wasn't a prismatic circuit of C^0. In the first case, the dihedral angle sum is $< \pi$ because condition (3) is satisfied by both β and γ and in the second case, the angle sum is $< \pi$ by definition of the T_i.

For each $j = 1, \cdots, k$ we have $\beta_j, \gamma_j \in (0, \pi/2)$, so $\alpha_j(t) \in (0, \pi/2)$. Since C^0 is the only abstract polyhedron that could be the 3-prism, condition (5) is satisfied because C^0 has only the edges $e_1, \cdots, e_k$, each of which is assigned an acute dihedral angle as seen above. However, $\alpha_j(t) = \pi/2$ for $j > k$, corresponding to the edges of the added triangular faces. Fortunately, a prismatic 4-circuit cannot cross edges of these triangular faces, since it would have to cross two edges from the same triangle, which meet at a vertex. This would be contrary to the definition of a prismatic circuit. So a 4-prismatic circuit can only cross edges numbered less than or equal to k, each of which is assigned an acute dihedral angle, and hence has dihedral angles sum $< 2\pi$, satisfying condition (4).

Consider the sequence of dihedral angles $\alpha_{n,i} = \alpha(T_{i-1} + (1-1/n)(T_i - T_{i-1}))$. By our above analysis, $\alpha_{n,i} \in A_{C^i}$ for each n, i. In fact, by definition $\alpha_{n,i}$ limits to the point $\alpha(T_i) \in \partial A_{C^i}$, which satisfies conditions (1-5) to be in A_{C^i}, except that the sum of the dihedral angles at each vertex $v_1^i, \cdots, v_{m(i)}^i$ is exactly π. We assume that Andreev's Theorem holds for C^i, so by Proposition 5.5, there exists a non-compact polyhedron P^i realizing C^i with each of the vertices $v_1^i, \cdots, v_{m(i)}^i$ at infinity and the rest of the vertices at finite points.

Now, by Lemma 6.14, the existence of P^i gives that there is a polyhedron realizing C^{i+1}, therefore, by Proposition 6.1, that Andreev's Theorem is satisfied for the abstract polyhedron C^{i+1}. Repeating this process until $i+1 = l$ gives that Andreev's Theorem is true for C^l, which is our original abstract polyhedron C. □

Proposition 6.15 *If $A_C \neq \emptyset$, then $\mathcal{P}_C^0 \neq \emptyset$.*

This proposition, combined with Proposition 6.1, concludes the proof of Andreev's Theorem.

Proof.

By Proposition 6.2 and Proposition 6.11 we know that Proposition 6.15 is true for every simple abstract polyhedron C and for every C whose only prismatic 3-circuits are truncated triangles. So we assume that there are $k > 0$ essential 3-circuits. From this point on, we will work entirely within the dual complex C^*.

Label the essential 3-circuits $\gamma_1, ..., \gamma_k$. The idea will be to replace C^* with $k+1$ separate abstract polyhedra $C_1^*, ..., C_{k+1}^*$ each of which has no essential 3-circuits. The γ_i separate the sphere into exactly $k+1$ components. Let C_i^* be the i-th of these components. To make C_i^* a simplicial complex on the sphere we must fill in the holes. Each of the holes is bounded by 3 edges (some γ_l). Fill in such a hole with the following figure (the dark outer edge is γ_l). The idea is that we are filling in each hole with a truncated triangle:

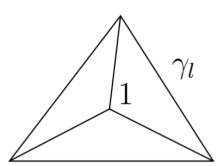

In C_i, we will call each vertex, edge, or face obtained by such filling in a *new vertex, new edge,* or *new face* respectively. We will call all of the other edges *old edges.* We label each such new vertex with the number l corresponding to the 3-circuit γ_l that was filled in. Clearly for each l, there will be exactly 2 new vertices labeled l which are in two different C_i, C_j. We must keep track of this pairing for a later gluing. See the following diagram, although the general case is obviously more complicated.

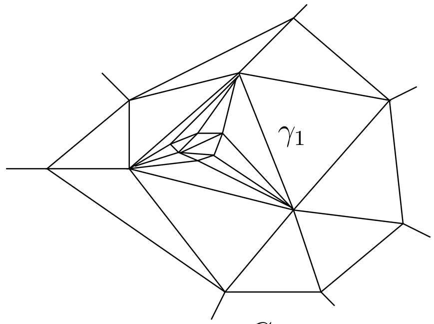

C

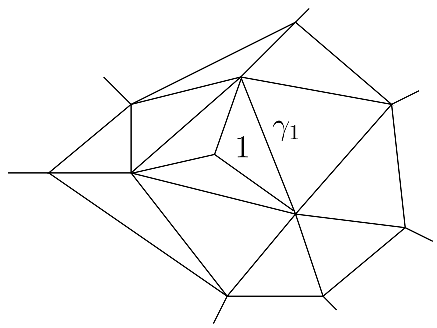

C_1

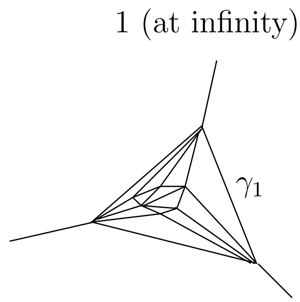

C_2

Notice that none of the C_i is a triangular prism, since we have divided up C along essential prismatic 3-circuits.

The choice of angles $\mathbf{a} \in A_C$ gives dihedral angles assigned to each old edge in each C_i^*. Assign to each of the new edges $\pi/2$. This gives a choice of angles $\mathbf{a_i}$ for each C_i^*. One must now check that $\mathbf{a_i} \in A_{C_i}$ for each i. Luckily, this is easy.

Clearly condition (1) is satisfied since these angles are non-zero and none of them obtuse.

The angles along each triangle of old edges in C_i^* already satisfy condition (2) since $\mathbf{a} \in A_C$. For each of the new triangles added, two of the edges are assigned $\pi/2$ and the third was already assigned a non-zero angle, according to $\mathbf{a}$, so condition (2) is satisfied for these triangles, too.

None of the new edges in C_i^* can be in a prismatic 3-circuit or a prismatic 4-circuit since such a circuit would have to involve two such new edges, which form two sides of a triangle. Therefore, each prismatic 3 or 4-circuit has come from such a circuit in C^*, so the choice of angles $\mathbf{a_i}$ will satisfy (3) and (4).

Since none of the C_i^* is a triangular prism, condition (5) is a consequence of condition (4), and hence is satisfied.

Therefore $\mathbf{a_i} \in A_{C_i}$ for each i. If for each i there exists a polyhedron P_i realizing the data $(C_i, \mathbf{a_i})$, we are done. To see this, notice that each of the new vertices corresponds to triangular face in P_i which is perpendicular to each of its neighboring faces. Also note that for each pair of new vertices labeled l the two faces dual to them are isomorphic. (By Proposition 1.1 the face angles are the same.) So one can glue all of the P_i together by gluing the pair of triangles labeled l together for each l. Since the edges of these triangles were assigned dihedral angles of $\pi/2$, the faces coming together from opposite sides of such a glued pair fit together smoothly. The result is a polyhedron P realizing C and angles $\mathbf{a}$. See the following diagram.

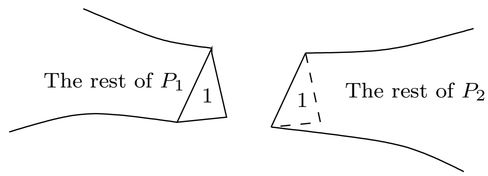

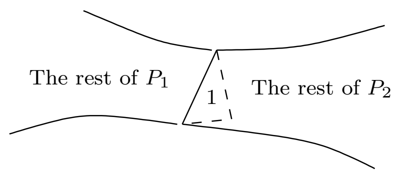

So, we have reduced the problem to proving that if C has no essential 3-circuit and if $A_C \neq \emptyset$, then $\mathcal{P}_C^* \neq \emptyset$. However, this is proven in Proposition 6.11. $\square$

That concludes the proof of Andreev's Theorem.

7 Hyperbolic tetrahedra

For the sake of completeness, we include this final section on hyperbolic tetrahedra having non-obtuse dihedral angles. The results of this section will make it clear why hyperbolic tetrahedra are a special case, distinct from Andreev's Theorem and they will also give insight into the reasons why an extension of Andreev's Theorem to include obtuse dihedral angles would be rather hard.

Theorem 7.1 *Suppose that the cell complex Δ on $\mathbb{S}^2$ gives the face structure of the tetrahedron. Given non-obtuse dihedral angles $\alpha_1, \cdots, \alpha_6$, determining non-obtuse face angles $\beta_1(\alpha_1, \cdots, \alpha_6), \cdots, \beta_{12}(\alpha_1, \cdots, \alpha_6)$, then, there is a compact hyperbolic polyhedron realizing Δ with non-obtuse dihedral angles $\alpha_1, \cdots, \alpha_6$ if and only if:*

1. *For each edge e_i, $0 < \alpha_i \leq \pi/2$.*
2. *Whenever 3 distinct edges e_i, e_j, e_k meet at a vertex, $\alpha_i + \alpha_j + \alpha_k > \pi$.*
3. *For each face the sum of the face angles satisfies $\beta_i + \beta_j + \beta_k < \pi$.*

Recall from Lemma 1.1 that the face angles β_i are calculable from the dihedral angles α_i and are themselves non-obtuse so that condition (3) is a condition in terms of the dihedral angles. We will see an explicit calculation of the face angles from the dihedral angles later in this section and use this to express the conditions (1-3) from Theorem 7.1 entirely in terms of the dihedral angles. We will denote the subset of $\mathbb{R}^6$ of dihedral angles satisfying conditions (1-3) by A_Δ.

Proof of Theorem 7.1: The method of proof is the same as in the proof of Andreev's Theorem, we consider the mapping $\alpha_\Delta : \mathcal{P}^0_\Delta \to A_\Delta$ and show that it is an injective covering map. We then show that $\mathcal{P}^0_\Delta$ is non-empty and that A_Δ is connected so that α_Δ is a homeomorphism.

The first step is to make sure that the dihedral angles of a tetrahedron satisfy conditions (1-3). The proofs of conditions (1) and (2) are identical to those in Section 3. Each face of a hyperbolic tetrahedron is a hyperbolic triangle of non-zero area so the Gauss-Bonnet formula gives condition (3).

The proofs that the mapping α_C is injective in Section 4 did not use anywhere that the polyhedron is not the tetrahedron, so we conclude that α_Δ is injective.

However, there is an elementary proof that α_C is injective: Since the face angles are uniquely determined by the dihedral angles and each face is a hyperbolic triangle, one can calculate the length of each edge using the hyperbolic law of cosines.

The only place in the proof that α_C is proper from Section 5 where it was used that C is not the tetrahedron was to show that the sequence of polyhedra P_i did not have any faces that degenerate to a points or a line segments. However, if they did, each of these faces would become Euclidean and would then violate condition (3). Notice that the conditions for prismatic three and four circuits as well as quadrilateral faces that were used in the proof in Section 5 are satisfied vacuously, since the tetrahedron has none of these combinatorial features. Therefore we conclude that α_Δ is proper.

Of course, as in the proof on Andreev's Theorem, this gives that α_Δ is an injective covering map from $\mathcal{P}^0_\Delta$ to A_Δ and it remains to show that $\mathcal{P}^0_\Delta$ is nonempty and that A_Δ is connected.

The easiest way to see that $\mathcal{P}^0_\Delta \neq \emptyset$ is by explicit construction. Let $\mathbf{v}_1 = (0,1,0,0), \mathbf{v}_2 = (0,0,1,0), \mathbf{v}_3 = (0,0,0,1)$, and $\mathbf{v}_4 = \frac{1}{\sqrt{3}}(0,1,1,1)$. Then the intersection of the half spaces $H_{\mathbf{v}_1} \cap H_{\mathbf{v}_2} \cap H_{\mathbf{v}_3} \cap H_{\mathbf{v}_4}$ is a hyperbolic tetrahedron with dihedral angles $\alpha_{1,2} = \pi/2, \alpha_{1,3} = \pi/2, \alpha_{2,3} = \pi/2, \alpha_{1,4} = \alpha_{2,4} = \alpha_{3,4} = \arccos(1/\sqrt{3}) < \pi/2$. Hence, we conclude that $\mathcal{P}^0_\Delta \neq \emptyset$.

To see that A_Δ is connected is significantly harder than for A_C with C not the tetrahedron because the inequalities are not linear. We will have to do detailed analysis of the equation that expresses a face's angles in terms of the dihedral angles.

Lemma 7.2 A_Δ *is path connected.*

Proof: Recall that the face angle β_i at a vertex (e_i, e_j, e_k) in the face containing e_j and e_k is the length of the edge of the unique spherical triangle with angles

$(\alpha_i, \alpha_j, \alpha_k)$ that is opposite from the angle α_i. This can be calculated explicitly using the Law of Cosines rule from spherical geometry:

$$\cos(\beta_i) = \frac{\cos(\alpha_i) + \cos(\alpha_j)\cos(\alpha_k)}{\sin(\alpha_j)\sin(\alpha_k)}$$

First, notice that given $\alpha \in A_\Delta$, decreasing any of the components of α does not increase any of the β_i. One can check that if:

$$F(x, y, z) = \frac{\cos(x) + \cos(y)\cos(z)}{\sin(y)\sin(z)}$$

Then we have:

$$\frac{\partial F}{\partial x} = -\frac{\sin(x)}{\sin(y)\sin(z)},$$
$$\frac{\partial F}{\partial y} = \frac{-\sin(y)\sin(z)\sin(y)\cos(z) - \cos(y)\cos(z)\cos(y)\sin(z)}{\sin^2(y)\sin^2(z)},$$
$$\frac{\partial F}{\partial z} = \frac{-\sin(y)\sin(z)\cos(y)\sin(z) - \cos(y)\cos(z)\sin(y)\cos(z)}{\sin^2(y)\sin^2(z)}$$

These have the nice property that for all $(x, y, z) \in (0, \pi/2]^3$ we have $\frac{\partial F}{\partial x} < 0, \frac{\partial F}{\partial y} < 0$, and $\frac{\partial F}{\partial z} < 0$. Because arccos is a decreasing function, this gives that $\beta(\alpha_i, \alpha_j, \alpha_k) \leq \beta(\gamma_i, \gamma_j, \gamma_k)$ when $\alpha_i \leq \gamma_i, \alpha_j \leq \gamma_j$, and $\alpha_k \leq \gamma_k$.

Therefore, given $\alpha \in A_\Delta$, decreasing the angles of α cannot result in a violation of condition (3).

Let $A_i \subset \partial A_\Delta$ be the subset obtained by restricting the dihedral angle sum at each of the vertices, expect v_i, to equal π. Using the formula for the β_j, one can check that at each vertex with dihedral angle sum exactly π, each of the face angles is 0. One can also check that each of the face angles at v_i is non-obtuse, since each of the dihedral angles is non-obtuse. Therefore, for any point in A_i, for each $i = 1, \cdots, 4$, each of the face angle sums is $\leq \pi/2$. Since the formula for face angles in terms of dihedral angles is continuous, there exists a neighborhood NA_i of each A_i in A_Δ which is connected, since A_i is convex.

For $i = 1, \cdots, 4$, each A_i contains $(\pi/3, \cdots, \pi/3)$ hence $NA_1 \cap NA_2 \cap NA_3 \cap NA_4 \neq \emptyset$. Therefore $NA_1 \cup NA_2 \cup NA_3 \cup NA_4$ is path connected. Denote this set by $\mathcal{N}$

Given any $\alpha \in A_\Delta$, we will create a path from α to a point in $\mathcal{N}$. This will be sufficient to prove that A_Δ is connected.

First, consider $t \cdot \alpha$ decreasing t from 1 to 0. For some first value of t, the sum of dihedral angles at one of the vertices, say v_1, will be π. Next, decrease only the dihedral angles of edges not entering v_1 in the same uniform way until the sum of the dihedral angles at another of the vertices, say v_2 equals π. Finally, decrease the dihedral angle on the edge that does not enter v_1 or v_2 until one the two remaining vertices has dihedral angle sum π, call this vertex v_3.

Since we have decreased the dihedral angles during the duration of this path, condition (3) was satisfied throughout. Condition (1) was satisfied throughout because we decreased the dihedral angles, so none exceded $\pi/2$ and since we decreased them by scaling, so that none reached 0.

This path must have entered $\mathcal{N}$ because it connected a point in A to A_1. $\square$

Therefore, since $\alpha_\Delta : \mathcal{P}^0_\Delta \to A_\Delta$ is an injective covering map with $\mathcal{P}^0_\Delta \neq \emptyset$ and A_Δ path connected, we conclude that α_Δ is a homeomorphism. This proves Theorem 7.1. $\square$

Using the explicit calculation in the above proof, we can re-express Theorem 7.1 entirely in terms of the dihedral angles.

Theorem 7.3 *Suppose that the cell complex Δ on $\mathbb{S}^2$ gives the face structure of the tetrahedron. There is a compact hyperbolic polyhedron realizing Δ with non-obtuse dihedral angles $\alpha_1, \cdots, \alpha_6$ if and only if:*

1. *For each edge e_i, $0 < \alpha_i \leq \pi/2$.*
2. *Whenever 3 (distinct) edges e_i, e_j, e_k meet at a vertex, $\alpha_i + \alpha_j + \alpha_k > \pi$.*
3. *For each face F bounded by edges e_i, e_j, e_k with edges $e_{i,j}, e_{j,k}, e_{k,i}$ emanating from the vertices, we have:*

$$\arccos\left(\frac{\cos(\alpha_{i,j}) + \cos(\alpha_i)\cos(\alpha_j)}{\sin(\alpha_i)\sin(\alpha_j)}\right) +$$

$$\arccos\left(\frac{\cos(\alpha_{j,k}) + \cos(\alpha_j)\cos(\alpha_k)}{\sin(\alpha_j)\sin(\alpha_k)}\right) +$$

$$\arccos\left(\frac{\cos(\alpha_{k,i}) + \cos(\alpha_k)\cos(\alpha_i)}{\sin(\alpha_k)\sin(\alpha_i)}\right) < \pi$$

The proof is evidently a direct consequence of Theorem 7.1 and the formula for the face angles.

In terms of the dihedral angles, condition (3) is reasonably nasty. In fact, it results in A_C being non-convex! Consider the hyperbolic tetrahedron with dihedral angles x and y assigned to two edges that meet at a vertex and dihedral angles α assigned to the remaining edges. The following figure was computed in Maple [1] and shows the cross section of A_C when $\alpha = 1.3$.

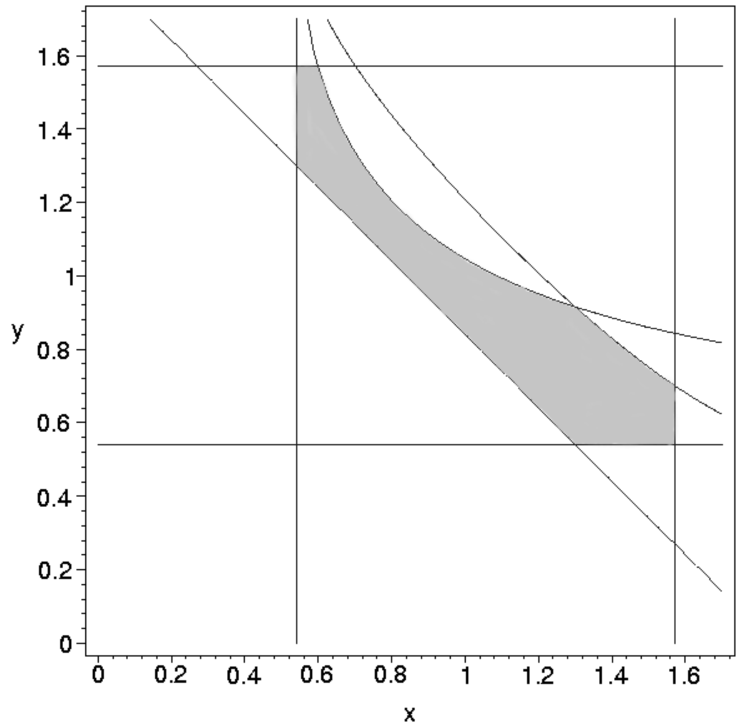

This classification of hyperbolic tetrahedra in terms of their dihedral angles gives us some understanding of how a generalization of Andreev's Theorem to include obtuse dihedral angles would be significantly more complicated than Andreev's Theorem. For instance, as in the case of the tetrahedron, one could have a triangular face truncating a finite vertex. With non-obtuse dihedral angles (for C not the tetrahedron) this cannot happen, by Lemma 3.3. However, this lemma used in an essential way that the dihedral angles are non-obtuse. In the general case, as for the tetrahedron, one would have to write a condition guaranteeing that the triangular faces do not degenerate in terms of the face angles, which depend in a non-linear way on the dihedral angles. One can only expect that other conditions necessary to prevent more exotic types of degeneracies.

One must also realize that in the case of arbitrary dihedral angles there is the additional difficulty that one cannot restrict to studying polyhedra realizing trivalent abstract polyhedra.

8 Example of the combinatorial algorithm from Lemma 6.5

In this final section we include an example of the combinatorial algorithm described in Lemma 6.5, which gives a sequence of Whitehead moves to reduce the dual complex of a simple abstract polyhedron, C^*, to the dual complex D_N^*, where N is the number of faces in C.

The Whitehead Sequence for a complex C for which Andreev's proof does not work:

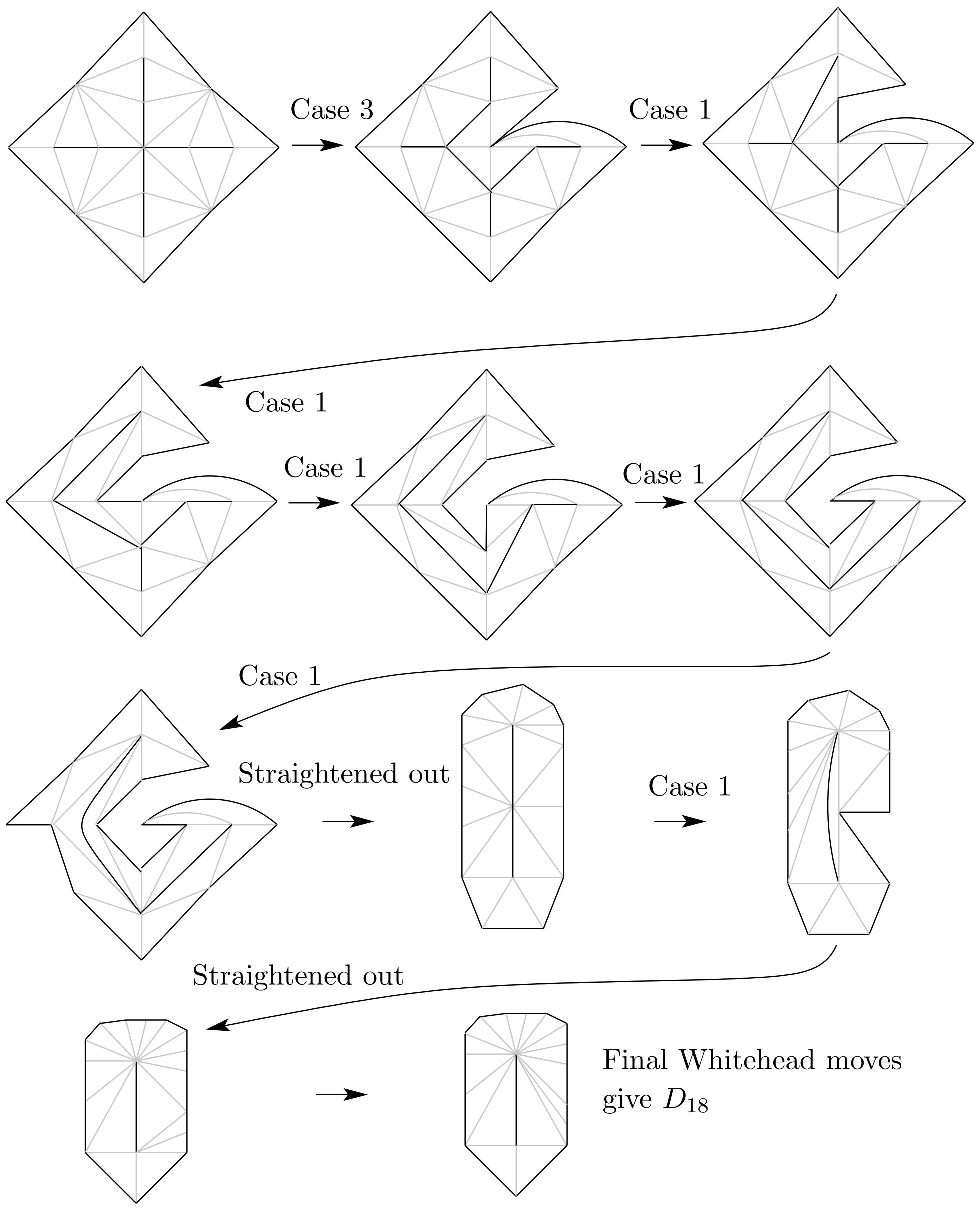

Bibliographie ∼ Bibliography

[1] www.maplesoft.com.

[2] www.geomview.org, Developed by The Geometry Center at the University of Minnesota in the late 1990's.

[3] E. M. Andreev. On convex polyhedra in Lobacevskii spaces (English Translation). *Math. USSR Sbornik*, 10:413–440, 1970.

[4] Michel Boileau. *Uniformisation en dimension trois*, pages 137–174. Number 266. Astérisque, 2000.

[5] C. D. Hodgson. Deduction of Andreev's theorem from Rivin's characterization of convex hyperbolic polyhedra. In *Topology 90*, pages 185–193. de Gruyter, 1992.

[6] A. Marden and B. Rodin. On Thurston's formulation and proof of Andreev's Theorem. In *Computational Methods and Function Theory*, volume 1435 of *Lecture Notes in Mathematics*, pages 103–115. Springer-Verlag, 1990.

[7] Jean-Pierre Otal. Thurston's hyperbolization of haken manifolds. In *Surveys in differential geometry, Vol. III (Cambridge, MA, 1996)*, pages 77–194. Int. Press, 1998.

[8] I. Rivin and C. D. Hodgson. A characterization of compact convex polyhedra in hyperbolic 3-space. *Invent. Math.*, 111:77–111, 1993.

[9] W. P. Thurston. *The Geometry and Topology of 3-manifolds.* Princeton University Notes, Princeton, New Jersey, 1980.