The Polar Dual of a Convex Polyhedral Set in Hyperbolic Space

Ruth Charney and Michael Davis

Introduction

Let X be a convex polyhedral set in a space of constant curvature. Its "polar dual" $P(X)$ is the set of outward pointing unit normal vectors to the supporting hyperplanes of X. In the spherical and Euclidean cases, $P(X)$ is a subset of the unit sphere. In the hyperbolic case, $P(X)$ is a subset of the unit pseudo-sphere in Minkowski space (sometimes called the "de Sitter sphere"). In all three cases $P(X)$ naturally has the structure of a piecewise spherical cell complex. The spherical cells of $P(X)$ correspond bijectively to the faces of X.

A piecewise spherical cell complex, with its intrinsic metric, is *large* if there is a unique geodesic between any two points of distance less than π. Equivalently, it is large if it satisfies Gromov's "CAT(1)-inequality" (cf. [G]).

Piecewise spherical cell complexes play an important role in the study of certain singular metric spaces: the link (or "space of directions") of a point in such a singular space often has a piecewise spherical structure. The largeness condition is closely related to the notion of nonpositive curvature in the sense of Alexandrov and Gromov. For example, a polyhedron of piecewise constant curvature has curvature bounded from above if and only if the link of each point is large (cf. [G], [B]). A similar result holds for the induced singular metric on the branched cover of a Riemannian manifold (cf. [CD1]).

In his 1986 Princeton Ph.D. thesis [R1], Igor Rivin proved that the polar dual of a convex polytope in hyperbolic 3-space is large. (This was published as the paper [RH] of Rivin and Hodgson.) The proof is a simple geometric argument. (The main result of [RH] is in the converse direction.) The referee has pointed out that some related results, in the smooth category, are proved in [S].

In his 1988 Ohio State Ph.D. thesis [M], Gabor Moussong considered a related situation: he gave a simple, necessary and sufficient condition for a piecewise spherical, simplicial complex with all edge lengths $\geq \frac{\pi}{2}$ to be large. The results of both Rivin and Moussong are generalizations of Andreev's Theorem, [A].

In this paper we use Moussong's ideas to extend Rivin's argument to any hyperbolic convex polyhedral set (not necessarily compact) of any dimension. The main result is the following.

Theorem. *The polar dual of a convex polyhedral set in hyperbolic space is large.*

This theorem is stated, in a sharper form, and proved in Section 4 as Theorem 4.1.1.

A geodesic in $P(X)$ is a special case of a "broken geodesic", $\gamma = (\gamma_1, \ldots, \gamma_k)$. This means that each γ_i is the arc of a great circle in some cell of $P(X)$ and that the terminal point x_i of γ_i coincides with the initial point of γ_{i+1}. As in [RH], one associates to each point x_i a supporting hyperplane H_i of X and to each γ_i the codimension two subspace $G_i = H_i \cap H_{i-1}$. (This is explained in subsection 3.1.) Thus, G_i and G_{i+1} are hyperplanes in H_i. The argument of [RH] then has two steps:

(1)   When γ is a geodesic, one shows that G_i and G_{i+1} are either parallel or ultraparallel. Hence, there is a region S_i in H_i bounded by G_i and G_{i+1}.

(2)   When γ is a closed geodesic the regions S_i fit together to give a piecewise isometrically immersed "cylinder" in the ambient hyperbolic space. The exterior dihedral angle along G_i is the length of γ_i. From this one concludes that the length of γ must be $\geq 2\pi$.

It follows from this last fact that $P(X)$ is large. The proof of (2) generalizes easily to higher dimensions; however, the proof of (1) does not. As we explain in Section 3 (Remark 3.2.3), the argument for (1) in [RH] comes down to the Gauss-Bonnet Theorem in dimension 2 and there is no obvious generalization of this argument to higher dimensions.

Our proof of (1) in higher dimensions occupies all of Section 3. It is based on ideas of [M]. The basic result, Theorem 3.2.2, relates the metric on $P(X)$ to the inner product on the ambient Minkowski space.

In Section 5 we relate the main result to a conjecture of [CD2]. Suppose K^{2m-1} is a piecewise spherical cell complex homeomorphic to the $(2m-1)$-sphere. In [CD2] we considered a number $\kappa(K^{2m-1})$, defined as a certain alternating sum of normalized volumes of the dual cells to the cells in K^{2m-1}. Hopf's Conjecture on the sign of the Euler characteristic of a nonpositively curved $2m$-manifold leads to the conjecture that $(-1)^m \kappa(K^{2m-1}) \geq 0$, whenever K^{2m-1} is large. In subsection 5.2 we recall the following formula of Hopf for the normalized volume v of a $2m$-dimensional hyperbolic polytope X^{2m}: $(-1)^m 2v(X^{2m}) = \kappa(P(X))$. (This is a special case of the Gauss-Bonnet Theorem of [AW].) It follows that the sign of $\kappa(K^{2m-1})$ is correct when K^{2m-1} is the polar dual of a hyperbolic polytope.

In Section 6, we relate the main result to the study of the space of large piecewise spherical structures on a given simplicial complex.

1.   Piecewise spherical polyhedra.

In this section we collect some well-known facts about piecewise spherical polyhedra.

1.1 Convex polyhedral sets. Let V be a finite dimensional real vector space. A subset X of V is a *convex polyhedral set* if it is defined by a finite number of affine

inequalities. In other words, there are linear functions $\lambda_i \in V^*$, and constants $c_i \in \mathbb{R}$, $i = 1, \cdots, k$, such that
$$X = \{v \in V | \lambda_i(v) \le c_i, i = 1, \cdots, k\}.$$

A *face* of X is a (nonempty) subset F of X obtained by changing the inequalities $\lambda_i(v) \le c_i$ into equalities $\lambda_i(v) = c_i$ for i in some subset of $\{1, \cdots, k\}$. (Equivalently , F is a face of X if for any two distinct points x, y in X such that $(x, y) \cap F \ne \emptyset$, we have $[x, y] \subset F$, where (x, y) and $[x, y]$ denote, respectively, the open and closed line segment from x to y.) Let $\mathcal{F}(X)$ denote the set of faces of X partially ordered by conclusion and let $\mathcal{F}_0(X)$ be the set of proper faces of X.

A convex polyhedral set X is called a *convex polytope* (or a *convex cell*) if it is compact. Equivalently, X is a convex polytope if it is the convex hull of a finite set.

A *polyhedral cone C* in V is a convex polyhedral subset defined by a finite number of linear inequalities, i.e.,

$$C = \{v \in V | \lambda_i(v) \le 0, i = 1, \cdots, k\}.$$

The *linear part* of C, noted $L(C)$, is the largest linear subspace contained in C. Thus, $L(C)$ is the minimum element in $\mathcal{F}(C)$. The cone C is *nondegenerate* if $L(C) = \{0\}$.

For any subset S of V, let $\mathrm{Span}(S)$ denote the linear subspace spanned by S. A polyhedral cone C is of *full dimension* if $\mathrm{Span}(C) = V$

Suppose X is a convex polyhedral set in V and $x \in X$. The *tangent cone* at x, denoted $T_x(X)$, is the polyhedral cone in $V(= T_xV)$ consisting of all inward pointing vectors at x, i.e.,
$$T_x(X) = \{v \in V | x + tv \in X \text{ for some } t > 0\}.$$

If F is a face of X, then this cone is independent of the choice of x in the relative interior of F and is denoted by $T(F, X)$. Explicitly, $T(F, X)$ is the cone determined by the codimension 1 faces of X containing F. The linear part of $T(F, X)$ is given by

$$L(T(F, X)) = T(F, F) = \mathrm{Span}(F - x)$$

(where $F - x$ is a translate of F to the origin) and

$$\mathrm{Span}\ (T(F, X)) = \mathrm{Span}(X - x).$$

1.2 Intrinsic metrics. Now identify V with some Euclidean space and let $\mathbb{S}(V)$ denote the unit sphere in V. A *spherical convex polyhedral set σ* is the intersection of a polyhedral cone C with $\mathbb{S}(V)$. If C is nondegenerate (or equivalently, if σ contains no pairs of antipodal points), then σ is called a *spherical convex polytope* or a *spherical cell*. A *piecewise spherical cell complex K* is a cell complex together with an identification of each cell with a spherical convex polytope. All such cell complexes will be assumed to be locally finite and to have a lower bound on the height of their cells. (The height of a cell σ is the minimum distance between disjoint faces of σ.)

3

A *broken geodesic* γ in a K is a sequence $(\gamma_1, \ldots, \gamma_k)$ where each γ_i is an oriented geodesic segment (= segment of a great circle) in some cell of K and where the terminal point of γ_i is the initial point of γ_{i+1}. The *length* of γ, denoted by $\ell(\gamma)$, is the sum of the lengths of the γ_i. The *intrinsic metric* on K is defined as follows: the distance $d(x, y)$ between points x and y in K is the infimum of the lengths of all broken geodesics from x to y. It follows from the assumptions above that, if x and y belong to the same component of K, then they can be joined by a broken geodesic whose length realizes this infimum (see [Br, Section 1] or [Pa]). Parameterizing by arc length gives a path $\gamma\colon [a, b] \to X$ such that $d(\gamma(s), \gamma(t)) = |s - t|$, for all $s, t \in [a, b]$. This means that K is a *geodesic space* (also called a *length space*). Such an isometric path is a *geodesic*.

The underlying metric space of a piecewise spherical cell complex with its intrinsic metric is called a *piecewise spherical polyhedron*.

1.3 Orthogonal joins. Suppose that V and W are disjoint Euclidean spaces, that $\sigma \subset \mathbb{S}(V)$ and $\tau \subset \mathbb{S}(W)$ are spherical cells, and that $C(\sigma)$ and $C(\tau)$ are the associated polyhedral cones. Then $C(\sigma) \times C(\tau)$ is a polyhedral cone in $V \oplus W$. The *orthogonal join* of σ and τ, denoted $\sigma * \tau$, is the spherical cell defined by

$$\sigma * \tau = C(\sigma) \times C(\tau) \cap \mathbb{S}(V \oplus W)$$

It is obvious how to extend this to a definition of the orthogonal join of two piecewise spherical complexes K_1 and K_2. The resulting intrinsic metric on $K_1 * K_2$ depends only on the intrinsic metrics on K_1 and K_2 and not on the particular cell structure.

Suppose K is a piecewise spherical cell complex. The *spherical cone* on K (or simply the "cone") is the orthogonal join of K with a point. The *k-fold suspension of K* is the orthogonal join $\mathbb{S}^{k-1} * K$. We note that there is no canonical cell structure on $\mathbb{S}^{k-1} * K$, rather it is partitioned into spherical convex sets of the form $\mathbb{S}^{k-1} * \sigma$, where σ is a cell of K or $\sigma = \emptyset$. (For more information about orthogonal joins, see [CD1, Appendix].)

1.4 Links. Suppose that σ is a spherical cell in $\mathbb{S}^n$ and that $C(= C(\sigma))$ is the associated polyhedral cone in $\mathbb{R}^{n+1}$. Let x be a point in σ and F the face of C which contains x in its relative interior. The *tangent cone* of σ at x is defined by

$$T_x(\sigma) = T(F, C) \cap x^\perp.$$

The *link of x in σ* is the set of unit vectors in this tangent cone,

$$Lk(x, \sigma) = T(F, C) \cap \mathbb{S}(x^\perp).$$

If $\tau = F \cap \mathbb{S}^n$, then the *link of τ in σ* is defined by

$$Lk(\tau, \sigma) = T(F, C) \cap \mathbb{S}(F^\perp)$$

and we have a natural identification

$$Lk(x, \sigma) = Lk(\tau, \sigma) * \mathbb{S}(x^\perp \cap F)$$
$$= Lk(\tau, \sigma) * \mathbb{S}^{k-1}$$

where $k = \dim \tau$ and x is in the relative interior of τ.

If x is a point in a piecewise spherical complex K, then define

$$Lk(x, K) = \bigcup_{x \in \sigma} Lk(x, \sigma)$$

where the union is taken over all cells σ containing x. Similarly, for a cell τ in K, define

$$Lk(\tau, K) = \bigcup_{\tau \subset \sigma} Lk(\tau, \sigma)$$

where the union is taken over all cells σ properly containing τ. Thus, $Lk(\tau, K)$ is naturally a piecewise spherical cell complex. If $\dim \tau = k$ and x is in the relative interior of τ, then

$$Lk(x, K) = Lk(\tau, K) * \mathbb{S}^{k-1}$$

so $Lk(x, K)$ is also a piecewise spherical polyhedron.

1.5 Local geodesics. Suppose K is a piecewise spherical cell complex. A *closed geodesic* in K is an isometric embedding of a circle.

A path $\gamma \colon [a, b] \to K$ is a *local geodesic* if for each $t \in [a, b]$ there is a neighborhood J of t such that $\gamma|_J$ is geodesic. (A *closed local geodesic* is defined similarly.)

Suppose $\gamma = (\gamma_1, \ldots, \gamma_k)$ is a broken geodesic in K, where each γ_i is an oriented arc of a great circle from x_{i-1} to x_i. If γ is closed (i.e., if $x_0 = x_k$), then the indices i are interpreted as integers modulo k. At each break point x_i, define points γ_i' and γ_{i+1}' in $Lk(x_i, K)$ as the unit tangent vectors of γ_i and γ_{i+1} at x_i, oriented outward from x_i. Define an "angle" θ_i by

$$(1.5.1) \qquad\qquad \theta_i = d(\gamma_i', \gamma_{i+1}')$$

where d is the intrinsic metric in $Lk(x_i, K)$. Then γ is a local geodesic if and only if it satisfies the following "angle condition" (cf. [M, Section 4] or [Br, p. 385]):

$$(1.5.2) \qquad\qquad \theta_i \geq \pi, \qquad i = 0, 1, \ldots, k-1.$$

1.5 Large piecewise spherical polyhedra. Any two points in the same component of a piecewise spherical polyhedron K are connected by a geodesic segment. We say that K is *large* if for any $x, y \in K$ with $d(x, y) < \pi$, the geodesic segment from x to y is unique. The condition that K is large is equivalent to the condition that it satisfies the "CAT(1)-inequality" of [G, p. 106]. (See [CD1, Theorem 3.1].)

The *systole* of K, denoted sys(K), is the infimum of the lengths of all closed geodesics in K.

The following lemma is a result of Gromov, [G, §4.2.A and 4.2.B].

5

Lemma 1.6.1 (Gromov). *A piecewise spherical polyhedron K is large if and only if the following two conditions hold:*

(i) $sys(K) \geq 2\pi$

(ii) $sys(Lk(\sigma, K)) \geq 2\pi$ *for all cells σ in K.*

It is fairly obvious that if K is large then (i) and (ii) hold; the proof of the converse uses the angle condition (1.5.2). For the details of the proof, see [B, Theorem 15] and [CD1, Theorem 3.1].

Lemma 1.6.2. *Let K be a large piecewise spherical polyhedron. Then any local geodesic of length $\leq \pi$ is actually a geodesic.*

Proof: The proof uses Alexsandrov's version of the largeness condition (see [Tr, Théorème 4]). This asserts that, given any geodesic triangle in K of perimeter $< 2\pi$, its "angles" must be no greater than those of its comparison triangle in $\mathbb{S}^2$. Let $\gamma = (\gamma_1, \ldots, \gamma_k)$ be a local geodesic and let x_{i-1} and x_i be the endpoints of γ_i. We may assume (by subdividing the γ_i) that each γ_i is a geodesic. The argument proceeds by induction on k.

If $k = 1$, then $\gamma = (\gamma_1)$ is geodesic by definition. Let $k > 1$ and assume by induction that $\alpha = (\gamma_1, \ldots, \gamma_{k-1})$ is an actual geodesic from x_0 to x_{k-1}. Suppose γ is not an actual geodesic. Then the actual geodesic β from x_0 to x_k has length $< \pi$. Consider the geodesic triangle $(\alpha, \gamma_k, \beta)$. The "angle", between α and γ_k at x_{k-1} (in the sense of Alexsandrov), coincides with $\min\{\pi, \theta_{k-1}\}$, where θ_{k-1} is defined by (1.5.1). By (1.5.2), $\theta_{k-1} \geq \pi$. It follows that the comparison triangle in $\mathbb{S}^2$ (which exists since perimeter $(\alpha, \gamma_k, \beta) < 2\pi$) must degenerate to a segment. The CAT(1)-inequality then implies that the same holds for $(\alpha, \gamma_k, \beta)$, that is, $\gamma = \beta$. $\qquad\qquad\square$

Definition 1.6.3. Suppose K_1 is a subpolyhedron of a piecewise spherical polyhedron K. Then K_1 is a *locally convex* subset of K (or a "totally geodesic subset") if any local geodesic in K_1 is actually a local geodesic in K.

The next lemma is an immediate consequence of Lemma 1.6.2.

Lemma 1.6.4. *Suppose that K is large and that K_1 is a locally convex subpolyhedron. Let d_1 and d denote the intrinsic metrics on K_1 and K, respectively. Then any local geodesic in K_1 of length $\leq \pi$ is actually a geodesic in K. Thus, if $x, y \in K_1$ and $d_1(x, y) \leq \pi$, then $d(x, y) = d_1(x, y)$.*

There is also the following simple formulation of local convexity in terms of links. It is an immediate consequence of (1.5.2).

Lemma 1.6.5. *Let K_1 be a subpolyhedron of K. Then K_1 is locally convex in K if and only if the following condition holds at each point u of K_1:*

6

(∗)   *If $x, x' \in Lk(u, K_1)$ and $d_1(x, x') \geq \pi$, then $d(x, x') \geq \pi$, where d_1 and d denote the intrinsic metrics in $Lk(u, K_1)$ and $Lk(u, K)$, respectively.*

2.   Polar duality.

In this section we investigate convex polyhedral sets in $\mathbb{S}^n, \mathbb{E}^n$, and $\mathbb{H}^n$. To such a set X, we will associate a piecewise spherical polyhedron $P(X)$, called the "polar dual" of X. In essence, $P(X)$, consists of the outward pointing unit normal vectors to X. The cell structure on $P(X)$ is induced by the face structure on X: a cell of $P(X)$ corresponds to the outward pointing normals to a face of X.

2.1 Dual Cones. Let V be a finite dimensional real vector space as above. In this subsection, V is equipped with a nondegenerate, symmetric, bilinear form $\langle \ , \ \rangle$. If S is a subset of V, then $S^\perp$ denotes the linear subspace of V orthogonal to Span(S).

Let C be a polyhedral cone in V. Its *dual cone* C^* is the polyhedral cone defined by

(2.1.1) $$C^* = \{w \in V \,|\, \langle v, w \rangle \leq 0 \text{ for all } v \in C\}.$$

Lemma 2.1.2 *Suppose $u_1, \cdots, u_k \in V$ and*

$$C = \{v \in V \,|\, \langle u_i, v \rangle \leq 0, i = 1, \cdot s, k\}.$$

Then

  (i) $C^* = \{$nonnegative linear combinations of $u_1, \cdots, u_k\}$
  (ii) $C^{**} = C$
  (iii) $L(C^*) = C^\perp$
  (iv) $Span(C^*) = L(C)^\perp$

Proof: Let $C_{\bar{u}}$ denote the cone of positive linear combinations of $u_1, \cdots, u_k$. It is easy to see that $C_{\bar{u}} \subset C^*$ and hence $C^{**} \subset C_{\bar{u}}^* = C$. On the other hand, it is clear from the definition of C^* that $C \subset C^{**}$. This proves (i) and (ii). Part (iii) is clear since $w \in L(C^*)$ if and only if $w, -w \in C^*$. Part (iv) follows from (iii) since $L(C) = (C^*)^\perp$ implies $L(C)^\perp = \text{Span}(C^*)$. □

In particular, C is of full dimension if and only if C^* is nondegenerate and vice versa.

If F is a face of C, then its *dual face* $\check{F}$ is the face of C^* defined by

$$\check{F} = T(F, C)^* = F^\perp \cap C^*$$

The correspondence $F \mapsto \check{F}$ is an order-reversing bijection from $\mathcal{F}(C)$ to $\mathcal{F}(C^*)$. It follows from Lemma 2.1.2 that

$$\text{Span}(\check{F}) = F^\perp \cap L(C)^\perp \text{ and } L(\check{F}) = C^\perp.$$

2.2 Polar duals in $\mathbb{E}^n$. Suppose $\langle \, , \, \rangle$ is positive definite and identity V with Euclidean space $\mathbb{E}^n$. Let X be a convex polyhedral set in V defined by

$$X = \{v \in V \mid \langle u_i, v \rangle \le c_i, i = 1, \cdots, k\}$$

and assume without loss of generality that X contains the origin. (The definition of polar dual is translation invariant.) Let $W = \mathrm{Span}(X)$. Let $C^*(X)$ denote the polyhedral cone in V spanned by $u_i, \cdots u_k$ and $W^\perp$ (i.e. $C^*(X)$ is the set of all nonnegative linear combinations of the u_i plus a vector in $W^\perp$).

Define the *full polar dual* of X to be

$$\widetilde{P}(X, V) = C^*(X) \cap \mathbb{S}(V)$$

and the *polar dual* of X to be

$$P(X) = \widetilde{P}(X) \cap W = C^*(X) \cap \mathbb{S}(W)$$

Since $\mathbb{S}(W^\perp) \subset \widetilde{P}(X, V)$, it follows that $\widetilde{P}(X, V)$ is the orthogonal join

$$\widetilde{P}(X, V) = \mathbb{S}(W^\perp) * P(X).$$

Let $C(X)$ denote the dual cone to $C^*(X)$. Then $C(X)$ is the largest polyhedral cone contained in X.

Lemma 2.2.1 *Let L_0 be the linear part of $C(X)$ and let L_1 be the orthogonal complement of L_0 in W. Then $P(X)$ is a convex polyhedral set in $\mathbb{S}(L_1)$. X is a Euclidean cell (i.e. compact) if and only if $L_1 = W$ and $P(X) = \mathbb{S}(W)$.*

Proof: By Lemma 2.1.1, $C^*(X) \subset L(C(X))^\perp = L_0^\perp$. The first statement of the theorem follows. For the second statement, note that X is compact if and only if $C(X) = \{0\}$, or equivalently, $C^*(X) = V$. $\qquad\square$

For each proper face F of X, define

$$\check{F} = T(F, X)^*$$
$$\sigma_F = \check{F} \cap P(X)$$

Note that $\check{F} \subset C^*(X)$, but since F is not necessarily a face of $C(X)$, $\check{F}$ need not be a face of $C^*(X)$. (Figure 1)

Figure 1

Lemma 2.2.2 σ_F *is a spherical cell. These cells partition* $P(X)$ *into a piecewise spherical complex.*

Proof: It is easy to see that $C(X) \subset T(F,X)$ and hence $\check{F} \subset C^*(X)$. The linear part of $\check{F}$ is $\mathrm{Span}(T(F,X))^\perp = \mathrm{Span}(X)^\perp = W^\perp$. Thus, $\check{F} \cap W$ is a nondegenerate polyhedral cone, so $\check{F} \cap P(X) = \check{F} \cap \mathbb{S}(W)$ is a spherical cell.

To see that these cells partition $P(X)$, it suffices to observe that the cones $T(F,X)^*$ partition $C^*(X)$, that is, every $v \in C^*(X)$ lies in the relative interior of $T(F,X)^*$ for a unique face F of X. Let $v \in C^*(X)$ and consider the linear function $x \to \langle x, v \rangle$ on X. The set of maxima of this function is necessarily a face of X which we denote by F_v. Then, for an arbitrary face F of X, we have

$$\begin{aligned} v \in T(F,X)^* &\Leftrightarrow \langle w, v \rangle \leq 0, \text{for all } w \in T(F,X) \\ &\Leftrightarrow \langle x + w, v \rangle \leq \langle x, v \rangle, \text{for all } x \in F, w \in T(F,X) \\ &\Leftrightarrow x \in F_v, \text{for all } x \in F \\ &\Leftrightarrow F \subset F_v \end{aligned}$$

It follows that v lies in the relative interior of a unique $T(F,X)^*$, namely $T(F_v,X)^*$. $\quad\square$

2.3 Polar duals in $\mathbb{S}^n$. In this subsection, $V = \mathbb{R}^{n+1}, \langle\,,\,\rangle$ is positive definite, and $\mathbb{S}(V)$ is the unit sphere in V. Let X be a convex polyhedral set in $\mathbb{S}(V)$. Then there is a unique polyhedral cone $C = C(X)$ in V such that $X = C \cap \mathbb{S}(V)$.

Let C^* be the dual cone.

Points in C^* may be viewed as outward pointing normal vectors to faces of C. The interior points of C^* correspond to normal vectors at the cone point and hence should not be viewed as normals to X. To define the polar dual of X, therefore, we consider ∂C^*. Here, ∂C^* denotes the topological boundary of C^* in V so that if C^* is of dimension less than $n+1$, then $\partial C^* = C^*$. The *full polar dual* of X is then defined by

$$\widetilde{P}(X,V) = \partial C^* \cap \mathbb{S}(V).$$

Letting $W = \mathrm{Span}(C)$, the *polar dual* of X is defined by

$$P(X) = \widetilde{P}(X,V) \cap W = \partial C^* \cap \mathbb{S}(W).$$

As in the Euclidean case,
$$\widetilde{P}(X,V) = \mathbb{S}(W^\perp) * P(X).$$

Lemma 2.3.1 *Let L be the linear part of C and let $k = \dim X = \dim W - 1$. Then $P(X)$ is the boundary of a spherical cell X^* of dimension $k - \dim L$. In particular, X is a spherical cell (i.e. $L = \{0\}$) if and only if $P(X)$ is homeomorphic to a $(k-1)$-sphere.*

Proof: Let $X^* = C^* \cap \mathbb{S}(W)$. The cone $C^* \cap W$ is nondegenerate, so X^* is a spherical cell with $\partial X^* = P(X)$. $\quad\square$

The boundary of a spherical cell is naturally a piecewise spherical complex. The cells of $P(X)$ are the dual faces, $\sigma_F = \check{F} \cap P(X) = F^\perp \cap P(X)$, where F is a proper, nonzero face of $C(X)$.

It is easy to see that these partition $P(X)$ into spherical cells.

2.4 Polar duals in $\mathbb{H}^n$. In this subsection, V is $(n+1)$-dimensional and the form $\langle \, , \, \rangle$ on V is indefinite of type $(n, 1)$. In other words, $\langle \, , \, \rangle$ has 1 negative eigenvalue and the rest are positive. (Thus, V could be identified with Minkowski space $\mathbb{R}^{n,1}$.) Let q be the associated quadratic form

$$q(v) = \langle v, v \rangle \, .$$

We consider the hyperquadrics $q^{-1}(-1)$, $q^{-1}(0)$ and $q^{-1}(1)$ in V. The first, $q^{-1}(-1)$ is a 2-sheeted hyperboloid; choose a component, denote it by $\mathbb{H}(V)$ and call it the *hyperbolic space* of V. The hyperquadric $q^{-1}(0)$ is denoted by $\mathbb{L}$ and called the *light-cone*, while $q^{-1}(1)$ is denoted $\mathbb{S}_1(V)$ and called the *unit pseudo-sphere* in V. ($\mathbb{S}_1(V)$ is sometimes called the "de Sitter sphere.")

The form $\langle \, , \, \rangle$ induces a Riemannian metric on $\mathbb{H}(V)$ of constant sectional curvature -1 and a Lorentzian metric on $\mathbb{S}_1(V)$ of constant sectional curvature $+1$. (See [O, pp. 108-114].)

A subspace W of V is *spacelike* if the restriction of $\langle \, , \, \rangle$ to W is positive definite, *lightlike* if $\langle \, , \, \rangle$ restricted to W is positive semidefinite but not definite, and *timelike* if W contains a vector w with $q(w) < 0$.

If C is a polyhedral cone in V, then

$$X = C \cap \mathbb{H}(V)$$

is called a *convex polyhedral subset* of $\mathbb{H}(V)$; X is called a *hyperbolic cell* (or a *hyperbolic convex polytope*) if it is compact. This is the case if and only if C is contained entirely inside the light-cone, i.e., if $q(v) < 0$ for all nonzero v in C.

Unlike the spherical case, the set X does not uniquely determine the cone C. However, the cone $C = C(X)$ is unique if we require that all codimension 1 faces of C be timelike, i.e.

$$C = \{v \in V \mid \langle u_i, v \rangle \leq 0, i = 1, \cdots, k\}$$

with $q(u_i) > 0$. (Lightlike and spacelike subspaces do not intersect $\mathbb{H}^n$ and hence do not affect X.) We will assume this is the case. Even under this hypothesis, however, C may have lower dimensional faces which are not timelike (since, for example, the intersection of two timelike hyperplanes may be a lightlike or spacelike subspace). For the polar dual, we are interested only in outward pointing normals to faces of X or, in other words, to timelike faces of C.

Let C^* be the dual cone to C. The timelike faces F of C correspond to the spacelike faces $\check{F}$ of C^*. We therefore let

$$C^*_{>0} = \text{ union of the spacelike faces of } C^*.$$

Figure 2

and define the *full polar dual* of X to be

$$\widetilde{P}(X, V) = C^*_{>0} \cap \mathbb{S}_1(V)$$

Letting $W = \mathrm{Span}(C)$, the *polar dual* of X is defined as

$$P(X) = \widetilde{P}(X, V) \cap W = C^*_{>0} \cap \mathbb{S}_1(W).$$

Unlike the spherical and Euclidean cases where $P(X)$ is always homeomorphic to a disk or a sphere, the polar dual of a convex hyperbolic set can have a much more arbitrary topology. However, in the case that X is a hyperbolic cell (i.e. compact), the following lemma states that $P(X)$ is indeed a sphere.

Lemma 2.4.1 *Suppose $k = \dim X = \dim C - 1$. If X is a hyperbolic cell, then $C^*_{>0} = \partial C^*$ and $P(X)$ is homeomorphic to a $(k-1)$-sphere*

Proof: Let $w \in \mathbb{H}^n$ and let $C_0 = \{v \in V \mid q(v) \le 0, \langle v, w \rangle \le 0 \text{ for all } w \in \mathbb{H}^n\}$. Then C_0 is the cone bounded by the (positive) light cone and $C_0^* = C_0$. (C_0 is not polyhedral, but the definition (2.1.1) of the dual cone still makes sense.) If $X = C \cap \mathbb{H}^n$ is a hyperbolic cell, then $C \subset \mathrm{int}\, C_0$; hence, $C_0 = C_0^* \subset \mathrm{int}\, C^*$. It follows that the spacelike faces of C^* are precisely those on its boundary, i.e., $C^*_{>0} = \partial C^*$.

For $w \in \mathbb{H}^n$, the restriction of $\langle\,,\,\rangle$ to $w^\perp$ is positive definite and $\mathbb{S}_1(V)$ is homeomorphic to $\mathbb{S}(w^\perp) \times \mathbb{R}$ ($\cong \mathbb{S}^{n-1} \times \mathbb{R}$) via the map

$$\mathbb{S}(w^\perp) \times \mathbb{R} \longrightarrow \mathbb{S}_1(V)$$
$$(v, r) \longmapsto (1 + r^2)^{\frac{1}{2}} v + rw$$

Thus, to prove the second statement of the lemma, it suffices to prove that for each $v \in \mathbb{S}(w^\perp)$, the curve $\alpha_v(r) = (1 + r^2)^{\frac{1}{2}} v + rw$ on $\mathbb{S}_1(V)$ intersects ∂C^* precisely once. Restricting to the 2-plane spanned by v and w reduces the problem to the 2-dimensional case, which follows easily from the fact that $C_0 \subset \mathrm{int}\, C^*$. (Figure 2) $\square$

For any proper face F of $C(X)$, define

$$P_F = \check{F} \cap P(X) = F^\perp \cap P(X).$$

If F is a timelike face (so $F \cap X$ is a *face* of X), then P_F turns out to be a spherical cell which we denote by $\sigma_F (= P_F)$.

Lemma 2.4.2 *If F is timelike, then σ_F is a spherical cell. The set $\{\sigma_F\}_{F \text{ timelike}}$ partitions $P(X)$ into a piecewise spherical cell complex.*

Proof: If F is a timelike face of C then $F^\perp$ is a spacelike subspace of V, so $F^\perp \cap \mathbb{S}_1(V) = \mathbb{S}(F^\perp)$ is a sphere. Also, $\check{F} \subset C^*_{>0}$ and the linear part of $\check{F}$ is $W^\perp$. Thus, $\sigma_F = \check{F} \cap P(X) = \check{F} \cap W \cap \mathbb{S}_1(V)$ is a spherical cell in $\mathbb{S}(F^\perp)$. Since the faces of $C^*_{>0}$ are precisely the duals of the timelike faces of C, these spherical cells partition $P(X)$ into a piecewise spherical complex. $\square$

Lemma 2.4.3 *Suppose F is spacelike. Then $F^\perp \cap X$ is a convex polyhedral set in the hyperbolic space $F^\perp \cap \mathbb{H}^n$, and $P_F = P(F^\perp \cap X)$. Thus, P_F is the polar dual of a convex set of smaller dimension.*

Proof: Since F is spacelike, $F^\perp$ is timelike. $F^\perp \cap C$ is a polyhedral cone in $F^\perp$ whose span is $F^\perp \cap W$. It is easy to see that the full polar dual of $F^\perp \cap X$ in $F^\perp$ is

$$\widetilde{P}(F^\perp \cap X, F^\perp) = C(X)^*_{>0} \cap \mathbb{S}_1(F^\perp)$$

and hence

$$P(F^\perp \cap X) = C(X)^*_{>0} \cap \mathbb{S}_1(F^\perp \cap W)$$
$$= F^\perp \cap P(X). \qquad\qquad \square$$

We shall analyze what happens when F is lightlike in the next subsection.

2.5 Cusps. Let $\mathbb{L}_+$ denote the positive light cone in $\mathbb{R}^{n,1}$, i.e., $\mathbb{L}_+ = \{v \in \mathbb{R}^{n,1} \mid q(v) = 0 \text{ and } \langle v, x \rangle \leq 0 \text{ for all } x \in \mathbb{H}^n\}$. The *sphere at infinity* of $\mathbb{H}^n$, denoted by S_∞, is the projective image of $\mathbb{L}_+ - \{0\}$. Thus, a point in S_∞ is an open ray in $\mathbb{L}_+ - \{0\}$. Geometrically, a point in S_∞ can be thought of as the limit point of a geodesic ray in $\mathbb{H}^n$.

If $y \in S_\infty$, then $y^\perp$ (a linear subspace of $\mathbb{R}^{n,1}$) is the tangent space along y to $\mathbb{L}_+$. Pick a point x in $\mathbb{H}^n$, and consider the affine subspace

$$A_y = x + y^\perp.$$

If $v \in y$, then A_y is the locus of those $w \in \mathbb{R}^{n,1}$ such that $\langle v, w \rangle = $ constant, where the constant is, of course, $\langle v, x \rangle$. The *horosphere* at y (passing through x) is the subset E_y of $\mathbb{H}^n$ defined by

$$E_y = A_y \cap \mathbb{H}^n.$$

12

With the induced metric, E_y is isometric to Euclidean $(n-1)$-space.

We return to the situation where X is a convex polyhedral set in $\mathbb{H}^n$. Let $S_\infty(X)$ denote the image of $(C(X) \cap \mathbb{L}_+) - \{0\}$ in S_∞. The points in $S_\infty(X)$ are the *points at infinity* of X.

For any $y \in S_\infty(X)$, let F_y be the unique face of $C(X)$ which contains y in its relative interior. Thus, F_y is either timelike or lightlike. The point y is called a *cusp point* of X if F_y is lightlike. If F is any lightlike face of $C(X)$, then there is a unique cusp point y with $F_y = F$; in fact, $y = F \cap \check{F}$.

Suppose $y \in S_\infty(X)$ and that E_y is any sufficiently small horosphere at y. ("Sufficiently small" means that E_y intersects only those faces of $C(X)$ which contain F_y.) Put

$$Z_y = E_y \cap X.$$

Then Z_y is an $(n-1)$-dimensional Euclidean convex polyhedral set, well-defined up to similarity.

Lemma 2.5.1. *Let $y \in S_\infty(X)$. F_y is timelike if and only if Z_y is a polyhedral cone. Thus, y is a cusp point if and only if Z_y is not a polyhedral cone.*

Proof: Let $\mathcal{E}_y$ be the poset of timelike faces of $C(X)$ which contain y. Then $\mathcal{E}_y$ indexes the poset of faces of Z_y. Thus, Z_y is a polyhedral cone if and only if $\mathcal{E}_y$ contains a minimum (= lower bound). But $\mathcal{E}_y$ contains such a minimum if and only if F_y is timelike (in which case F_y is the minimum). $\qquad\square$

If y is a cusp point of X, then put

$$P_y = P_{F_y} = F_y^\perp \cap P(X).$$

Lemma 2.5.2. *If y is a cusp point, then P_y can be canonically identified with $P(Z_y)$. Thus, the subcomplex P_y corresponding to a cusp is isometric to the polar dual of a convex set in $\mathbb{E}^{n-1}$.*

Proof: There is a canonical identification of the horosphere E_y with the Euclidean space $y^\perp / \langle y \rangle = \mathbb{E}^{n-1}$, well-defined up to translation. (Here, $\langle y \rangle$ denotes the line determined by y.) On tangent spaces, this identification corresponds to the isomorphism

$$T_x E_y = x^\perp \cap y^\perp \overset{\cong}{\longrightarrow} y^\perp / \langle y \rangle$$

induced by the projection $y^\perp \to y^\perp / \langle y \rangle$. If F is a timelike face of $C(X)$ containing y, and $x \in F \cap E_y$, then $F^\perp \subset F^\perp \cap x^\perp \subset y^\perp \cap x^\perp = T_x E_y$. It follows that unit normals to F in $\mathbb{R}^{n,1}$ can be identified with unit normals to $T_x(F \cap E_y)$ in $T_x E_y$, and hence with unit normals to $F \cap E_y$ in E_y, viewed as $\mathbb{E}^{n-1}$. $\qquad\square$

2.6 Links in $P(X)$. Let us summarize the results of sections 2.3-2.5. Let X be a convex polyhedral set in $\mathbb{E}^n, \mathbb{S}^n$ or $\mathbb{H}^n$ and let $C = C(X)$ be the cone determined by X

(as in 2.3-2.5). If Y is a face of X, then $F_Y = C(Y)$ is a face of C. (In the spherical and hyperbolic cases, F_Y is the unique face such that $Y = X \cap F_Y$.) Associated to Y is a spherical cell $\sigma_Y = \check{Y} \cap P(X)$. (In the spherical and hyperbolic cases, $\check{Y} \equiv \check{F}_Y, \sigma_Y \equiv \sigma_{F_Y}$; in the Euclidean case, $\check{Y} \subset \check{F}_Y$, but they need not be equal.) Let $\mathcal{F}_0(X)$ denote the set of proper faces of X ordered by inclusion. Let $\mathcal{F}_0(X)^{op}$ denote this set with the opposite ordering.

Proposition 2.6.1 *Let X be a convex polyhedral set in $\mathbb{E}^n, \mathbb{S}^n$, or $\mathbb{H}^n$. Then the polar dual of X, $P(X)$ is a piecewise spherical cell complex whose poset of cells $\{\sigma_Y\}$ is naturally identified with $\mathcal{F}_0(X)^{op}$. If X is a cell of dimension k in $\mathbb{E}^n, \mathbb{S}^n$, or $\mathbb{H}^n$, then $P(X)$ is homeomorphic to a $(k-1)$-sphere.*

Lemma 2.6.2 *Let X be a convex polyhedral set in $\mathbb{E}^n, \mathbb{S}^n$, or $\mathbb{H}^n$ and let $W = Span(C(X))$. Let $Y \in \mathcal{F}_0(X)$ and let y be a point in the relative interior of σ_Y. Then*

(1) $Lk(y, P(X)) = \widetilde{P}(Y, y^\perp \cap W)$
(2) $Lk(\sigma_Y, P(X)) = P(Y)$

Proof: The correspondence $F \mapsto \check{F}$ gives a bijection between faces of C and faces of C^* such that
$$F = \check{\check{F}} = T(\check{F}, C^*)^*.$$
Thus, $F^* = T(\check{F}, C^*)$ (where F^* is the dual cone in V to F, viewed as a polyhedral cone).

Let $F = F_Y = C(Y)$. Then y lies in the relative interior of $\check{F}$, so
$$F^* = T_y(C^*).$$

The faces of F^* are of the form $T(\check{F}, \check{G})$ where G is a face of F. In the hyperbolic case, $T(\check{F}, \check{G})$ is spacelike if and only if $\check{G}$ is spacelike (since $\text{Span}(T(\check{F}, \check{G})) = \text{Span}(\check{G})$), so
$$F^*_{>0} = T_y(C^*_{>0}).$$

Thus, in the spherical and Euclidean cases,
$$\begin{aligned} Lk(y, \widetilde{P}(X)) &= \mathbb{S}(y^\perp) \cap T_y(C^*) \\ &= \mathbb{S}(y^\perp) \cap F^* \\ &= \widetilde{P}(Y, y^\perp). \end{aligned}$$

The same holds in the hyperbolic case if we replace C^* by $C^*_{>0}$ and F^* by $F^*_{>0}$. Intersecting both sides of the equation with W gives part (1) of the lemma; intersecting with $\text{Span}(Y)$ $(= (\check{Y})^\perp)$ gives part (2). $\square$

3. The intrinsic metric on $P(X)$.

3.1 A geometric interpretation of local geodesics in $P(X)$. There is an interesting interpretation of a local geodesic in $P(X)$ in terms of the hyperbolic geometry of X. This relationship is explained in [RH]. We shall now recall it.

14

Each point v in $\mathbb{S}_1^n$ corresponds to an oriented hyperbolic hyperplane

$$H_v = v^\perp \cap \mathbb{H}^n .$$

If $v \in P(X)$, H_v is called a *supporting hyperplane* of X. This terminology is appropriate since v lies in $P(X)$ if and only if X lies in the *half-space*

$$D_v = \{ u \in \mathbb{H}^n \mid \langle u, v \rangle \le 0 \}$$

and $X \cap H_v$ is nonempty. (If v belongs to the relative interior of the spherical cell σ_Y, then $H_v \cap X = Y$.)

Suppose that σ_Y is a cell of $P(X)$, that v_0 and v_1 are points in distinct faces of $\partial \sigma_Y$ and that γ is the geodesic segment in σ_Y from v_0 to v_1. Put

$$G = \gamma^\perp \cap \mathbb{H}^n = H_{v_0} \cap H_{v_1} .$$

Then G is a totally geodesic codimension two subspace of $\mathbb{H}^n$. The length $\ell(\gamma)$ of γ is $\cos^{-1} \langle v_0, v_1 \rangle$. Hence, $\ell(\gamma)$ measures the exterior dihedral angle between H_{v_0} and H_{v_1}. Corresponding to the points of γ, there is a 1-parameter family of supporting hyperplanes obtained by rotating H_{v_0} to H_{v_1} about G through an angle of $\ell(\gamma)$.

Next suppose that $\gamma = (\gamma_1, \ldots, \gamma_k)$, is a broken geodesic in $P(X)$, where γ_i is a geodesic segment in some cell from v_{i-1} to v_i, as in the previous paragraph. Put $H_i = H_{v_i}$ and $G_i = (\gamma_i)^\perp \cap \mathbb{H}^n = H_{i-1} \cap H_i$. Thus, γ gives us a 1-parameter family of supporting hyperplanes "rolling" along the boundary of the convex body X and containing $H_0, H_1, \ldots, H_{k-1}$ in succession.

Consider the $(n-1)$-dimensional hyperbolic space H_i. Then G_i and G_{i+1} are oriented hyperplanes in H_i with corresponding unit normal vectors γ_i' and γ_{i+1}' in $Lk(v_i, P(X))$. Suppose that the broken geodesic γ is actually a local geodesic, i.e., $d(\gamma_i', \gamma_{i+1}') \ge \pi$, (cf. (1.5.2)). We shall prove in Proposition 3.3.3 below that this implies that the hyperplanes G_i and G_{i+1} do not intersect in H_i. Hence, G_i and G_{i+1} bound an infinite $(n-1)$-dimensional "strip" S_i in H_i.

Finally, we suppose that $\gamma = (\gamma_1, \ldots, \gamma_k)$ is a closed local geodesic. Then we get a sequence of strips $(S_0, \ldots, S_{k-1})$ such that $S_i \subset H_i$ and $S_{i-1} \cap S_i = G_i$. Let $S_\gamma = S_0 \cup \cdots \cup S_{k-1}$ be the union of these strips in $\mathbb{H}^n$. We note that, even when the indices i and j are not consecutive (modulo k), the strips S_i and S_j might have nonempty intersection. To eliminate these extraneous intersections, one introduces an abstract $(n-1)$-dimensional hyperbolic manifold $\widehat{S}_\gamma$ formed by taking the quotient of the disjoint union $\coprod S_i$ by the equivalence relation which identifies the boundary components of S_{i-1} and S_i which correspond to G_i. The natural map $\widehat{S}_\gamma \to S_\gamma \subset \mathbb{H}^n$ is then a piecewise totally geodesic immersion.

The $(n-1)$-manifold $\widehat{S}_\gamma$ is a "hyperbolic cylinder." That is to say, it is isometric to $\mathbb{H}^{n-1}/\langle \rho_\gamma \rangle$, where ρ_γ is either a parabolic or hyperbolic isometry of $\mathbb{H}^{n-1}$ and where $\langle \rho_\gamma \rangle$ denotes the infinite cyclic group generated by ρ_γ. Thus, to each closed local geodesic γ in $P(X)$ we have associated an immersed hyperbolic cylinder S_γ in $\mathbb{H}^n$ such that the length

of γ is the sum of the exterior dihedral angles in S_γ. This cylinder will be used in Section 4 to prove the main theorem (Theorem 4.1.1).

3.2 A version of a result of Moussong.

Our goal in this subsection is to prove Theorem 3.2.2, below, which relates the intrinsic metric on $P(X)$ to the Lorentzian inner product on $\mathbb{R}^{n,1}$. We begin with a preliminary lemma which clarifies the statement of Theorem 3.2.2.

Lemma 3.2.1. *If $v, w \in P(X)$, then $\langle v, w \rangle \leq 1$, with equality if and only if $v = w$.*

Proof: Suppose that v, w are distinct vectors in $P(X)$ with $\langle v, w \rangle \geq 1$. Then the supporting hyperplanes H_v and H_w do not intersect in $\mathbb{H}^n$. Thus, for any x in H_v, the quantity $\langle w, x \rangle$ is never zero. If it is negative, then the half-space D_v is contained in D_w; if positive, then $D_w \subset D_v$. In either case, H_v and H_w cannot both be supporting hyperplanes, a contradiction. $\square$

Theorem 3.2.2. *Suppose v, w are points in $P(X)$ with $\langle v, w \rangle > -1$. Then*

$$d(v, w) \leq \cos^{-1} \langle v, w \rangle,$$

where d is the intrinsic metric on $P(X)$. In fact, suppose neither of the following conditions hold:

(i)   *v and w belong to some cell of $P(X)$,*

(ii)   *v and w belong to P_y for some cusp point y of X.*

Then $d(v, w) < \cos^{-1} \langle v, w \rangle$.

The above theorem is a modification of Lemma 9.7 in the thesis of G. Moussong [M]. Although Moussong was working in a slightly different context, both the statement and proof of Theorem 3.2.2 are essentially the same as in [M] (our case is somewhat easier).

Remark 3.2.3. When X is a 2-dimensional convex polyhedral set in $\mathbb{H}^2$ there is an easy geometric proof of Theorem 3.2.2, along the lines of the proof of Lemma 3.4 in [RH], which we now recall. Suppose $v, w \in P(X)$ and $\langle v, w \rangle \geq -1$. Then the supporting lines H_v and H_w intersect at a point q_0 which either lies in $\mathbb{H}^2$ or on the circle at infinity (when $\langle v, w \rangle = -1$). Let Q be the polygon, exterior to X, which is bounded by H_v, H_w and some of the edges of X as indicated in Figure 3, below. Let $q_1, \ldots, q_k$ be the other vertices of Q and let $\alpha, \theta_1, \ldots, \theta_k$ be the indicated angles. Then $\alpha = \cos^{-1} \langle v, w \rangle$ and $d(v, w) = \theta_1 + \cdots + \theta_k$.

Figure 3

For any finite area polygon Q in $\mathbb{H}^2$ (convex or not), the Gauss-Bonnet Theorem states that the sum of the exterior angles of Q is equal to $2\pi + \mathrm{Area}(Q)$. In our case, the exterior angle at q_0 is α, at q_1 it is $\pi - \theta_1$ and at q_k it is $\pi - \theta_k$, while, for $2 \le i \le k - 1$ the exterior angle at q_i is $-\theta_i$ (a negative number). So,

$$\alpha + (\pi - \theta_1) - (\theta_2 + \cdots + \theta_{k-1}) + (\pi - \theta_k) = 2\pi + \mathrm{Area}(Q)$$

Hence, $\alpha = \sum \theta_i + \mathrm{Area}(Q)$, that is,

$$\cos^{-1} \langle v, w \rangle = d(v, w) + \mathrm{Area}(Q).$$

Thus, $d(v, w) \le \cos^{-1} \langle v, w \rangle$ and the inequality is strict unless $\mathrm{Area}(Q) = 0$. If $\mathrm{Area}(Q) = 0$, then H_v and H_w are actually supported by edges of X, and q_0 is a vertex of X (possibly at ∞) and so we are in situation (i) or (ii) of the theorem.

Before giving the details of the proof of Theorem 3.2.2, let us sketch the rough idea. Suppose v and w are distinct points in $P(X)$. By Lemma 3.2.1 we always have $\langle v, w \rangle < 1$. Hence, the condition that $\langle v, w \rangle > -1$ means that $|\langle v, w \rangle| < 1$, i.e., that v and w span a spacelike subspace of $\mathbb{R}^{n,1}$. The cone generated by v and w then intersects $\mathbb{S}_1^n$ in a circular arc δ of length $\cos^{-1} \langle v, w \rangle$. In general, δ will not lie in $P(X)$. Let u denote the unit tangent vector to δ at v. In a sequence of lemmas, stated below (in particular, in Lemma 3.2.8), we show that one can always find a vector $z \in Lk(v, P(X))$ so that as x moves along the geodesic γ in direction z, the quantity $\cos^{-1} \langle x, w \rangle$ decreases at least as fast as it does when x moves along δ. (The condition that z have this property is just that $\langle u, z \rangle \ge 1$.) Using a compactness argument, we show that we can extend γ to a path from v to w of length at most that of δ.

Suppose G is a subspace of $\mathbb{H}^n$ and $y \in S_\infty$. We say that G *meets* y if there is a geodesic ray in G which limits at y.

Lemma 3.2.4. *Suppose that $v, w \in P(X)$ are such that $H_v \cap H_w \cap X = \emptyset$. If y is a point in $S_\infty(X)$ which meets H_v and H_w, then y is a cusp point of X.*

Proof: Let E_y be a small horosphere at y and $Z_y = E_y \cap X$. Then Z_y is a Euclidean convex polyhedral set and $v^\perp \cap E_y$ and $w^\perp \cap E_y$ are supporting hyperplanes of Z_y which do not intersect. It follows that Z_y is not a polyhedral cone (if it were, the intersection of any two supporting hyperplanes would contain the minimum face). By Lemma 2.5.1, y is a cusp point.

Lemma 3.2.5. *Suppose that u is a unit vector in $C^*(X)$. Let $d(H_u, X)$ denote the distance (in $\mathbb{H}^n$) from H_u to X.*

(i)   *If $d(H_u, X) > 0$, then there is a unique point x_0 in X such that $d(H_u, X) = d(H_u, x_0)$.*

(ii) *If $d(H_u, X) = 0$, then either*

17

(a) H_u is a supporting hyperplane for X (i.e., $u \in P(X)$) or

(b) H_u meets $S_\infty(X)$ in a single point.

(iii)   There is a vector z in $P(X)$ such that $\langle u, z \rangle \geq 1$. Moreover, if $d(H_u, X) > 0$, then z can be chosen so that $\langle u, z \rangle > 1$.

Proof: (i) If $\alpha(t)$ is any infinite geodesic ray in X, then

$$\lim_{t \to \infty} d(H_u, \alpha(t)) = \begin{cases} 0; & \text{if } \alpha(t) \text{ and } H_u \text{ meet at infinity} \\ \infty; & \text{otherwise} \end{cases}$$

Since X is geodesically convex and since the restriction of the function $x \to d(H_u, x)$ to any geodesic segment is strictly convex, it follows that if $d(H_u, x)$ is bounded away from 0, then there is a unique point x_0 in X where the minimum value is obtained.

(ii) Suppose $\alpha_1(t)$ and $\alpha_2(t)$ are two geodesic rays in X which both meet H_u at infinity. Let y_1 and y_2 be the corresponding limit points in $S_\infty(X)$. If $y_1 \neq y_2$ then, by convexity, the infinite geodesic from y_1 to y_2 is contained in $H_u \cap X$ and hence, H_u is a supporting hyperplane for X. Therefore, if H_u is not a supporting hyperplane, then $y_1 = y_2$, which proves (ii).

(iii) First, suppose that $d(H_u, X) > 0$. Then we choose z to be the parallel transport of u to the closest point x_0. It is then geometrically clear that H_z is a supporting hyperplane. This can also be checked by the following algebraic calculation. The vector z is the orthogonal projection of u onto $(x_0)^\perp$, normalized to have unit length, i.e.,

$$z = \frac{u + \langle u, x_0 \rangle x_0}{\left(1 + \langle u, x_0 \rangle^2\right)^{\frac{1}{2}}}$$

Since $d(H_u, x) = -\sinh \langle u, x \rangle$ (see [T, Section 2]), the function $x \to \langle u, x \rangle$ takes its maximum value on X at x_0, i.e.,

$$\langle u, x_0 \rangle \geq \langle u, x \rangle, \qquad \text{for all } x \in X.$$

Since the inner product of any two points in $\mathbb{H}^n$ is ≤ -1,

$$\langle x_0, x \rangle \leq -1$$

Hence, for any $x \in X$,

$$\langle u + \langle u, x_0 \rangle x_0, x \rangle = \langle u, x \rangle + \langle u, x_0 \rangle \langle x_0, x \rangle \leq \langle u, x \rangle - \langle u, x \rangle = 0.$$

So, $u + \langle u, x_0 \rangle x_0$ belongs to $C^*(X)$ and therefore, so does its positive multiple z. Since $\langle z, x_0 \rangle = 0$ and $\langle z, z \rangle = 1$, we see that $z \in P(X)$. Moreover,

$$\langle u, z \rangle = \left(1 + \langle u, x_0 \rangle^2\right)^{\frac{1}{2}} > 1,$$

as claimed.

Finally, suppose that $d(H_u, X) = 0$. By (ii) there are two cases to consider. If H_u is a supporting hyperplane, then put $z = u$. Otherwise, let y be the unique point in $S_\infty(X)$ which meets H_u. Let E_y be a sufficiently small horosphere at y and $Z_y = E_y \cap X$. Then Z_y is a Euclidean convex polyhedral set and $u^\perp \cap E_y$ is an affine hyperplane which bounds a halfspace containing Z_y. let H' be the oriented affine hyperplane in E_y obtained by translating $u^\perp \cap E_y$ to a closest point in Z_y. Then there is a unique $z \in \mathbb{S}_1^n$ such H_z and H_u are parallel and meet at infinity at y and such that $H_z \cap E_y = H'$. Clearly, $z \in P(X)$ and $\langle u, z \rangle = 1$, as desired.

In the next two lemmas and in subsection 3.3 we must deal with the case where the convex polyhedral set might be of smaller dimension than the ambient hyperbolic space. In other words, the full polar dual might not coincide with the polar dual. We set up some notation to deal with this situation below.

Notation 3.2.6. Suppose Y is a convex polyhedral set in $\mathbb{H}^{n-1}$ (possibly of dimension $< n-1$). Let $u \in \mathbb{R}^{n-1,1}$. If u is spacelike, then $\bar{u} = q(u)^{-\frac{1}{2}} u$ denotes its normalization to a unit vector. Denote by u_1 and u_2 the orthogonal projections of u onto $Y^\perp$ and $\mathrm{Span}(Y)$, respectively. Identify $Y^\perp$ with $\mathbb{R}^\ell$ and $\mathrm{Span}(Y)$ with $\mathbb{R}^{n-\ell-1,1}$. If u_2 is any spacelike vector in $\mathbb{R}^{n-\ell-1,1}$, then G_{u_2} denotes the hyperbolic hyperplane $(u_2)^\perp \cap \mathbb{H}^{n-\ell-1}$. As in 2.5, $\widetilde{P}(Y) = \mathbb{S}^{\ell-1} * P(Y)$ (where the ambient space $V = \mathbb{R}^{n-1,1}$ is omitted from the notation).

Lemma 3.2.7. *Let u be a unit vector in $\mathbb{R}^{n-1,1}$ such that $u_2 \in C^*(Y)$. Then there is a vector $z \in \widetilde{P}(Y)$ such that $\langle u, z \rangle \geq 1$. Moreover, z can be chosen so that the inequality is strict except possibly in the following three cases:*

(a) $u_2 = 0$,

(b) u_2 *is spacelike and* $d(G_{u_2}, Y) = 0$,

(c) u_2 *is lightlike and* $u_2 \in S_\infty(Y)$.

Proof: There are four cases to consider as the vector u_2 is zero, timelike, spacelike or lightlike.

Case 1: $u_2 = 0$. Put $z = u$. Then z lies in the suspension sphere and hence, in $\widetilde{P}(Y)$. Furthermore, $\langle u, z \rangle = 1$.

Case 2: u_2 *is timelike.* Then $q(u_1) = 1 - q(u_2) > 1$. In particular, $u_1 \neq 0$. Put $z = \bar{u}_1$. Again, z lies in the suspension sphere and $\langle u, z \rangle = q(u_1)^{\frac{1}{2}} > 1$.

Case 3: u_2 *is spacelike.* Apply Lemma 3.2.5 to $\bar{u}_2$ and Y to find $z_2 \in P(Y)$ so that $\langle \bar{u}_2, z_2 \rangle \geq 1$. Put $z = u_1 + q(u_2)^{\frac{1}{2}} z_2$. Then $q(z) = q(u_1) + q(u_2) = 1$, so $z \in \mathbb{S}^{\ell-1} * P(Y)$. Furthermore, $\langle u, z \rangle = q(u_1) + q(u_2) \langle \bar{u}_2, z_2 \rangle \geq q(u_1) + q(u_2) = 1$, and the inequality is strict if and only if $\langle \bar{u}_2, z_2 \rangle > 1$. By Lemma 3.2.5 (iii), we can choose z_2 so that $\langle \bar{u}_2, z_2 \rangle > 1$ unless $d(G_{u_2}, Y) = 0$.

19

Case 4: u_2 *is lightlike.* In this case one might try the approach in Case 2 or the one in Case 3. Both methods work; however, only the second method yields a strict inequality. Since $q(u_2) = 0$, $q(u_1) = 1 - q(u_2) = 1$. The first method is to choose $z = u_1$. Then z is in the suspension sphere and $\langle u, z \rangle = 1$. Let y be the point in S_∞ determined by u_2. The second method can be applied whenever $y \notin S_\infty(Y)$. As in Lemma 3.2.5, we find a point x_0 in Y which is closest to y. Explicitly, consider the smallest horosphere at y which intersects Y. The intersection must be a single point x_0. Let G be the hyperbolic hyperplane in $\mathbb{H}^{n-\ell-1}$ which is tangent to the horosphere at x_0 and let z_2 be the unit normal to G which points into the horoball. Clearly, $z_2 \in P(Y)$. We will choose z to be of the form

$$z = (\cos \theta) u_1 + (\sin \theta) z_2$$

for an appropriate choice of $\theta \in \left(0, \frac{\pi}{2} \right]$. Any such z will lie in $\mathbb{S}^{\ell-1} * P(Y)$. Let $\varepsilon = \langle u_2, z_2 \rangle$. Then $\varepsilon > 0$ and $\langle u, z \rangle = \langle u_1, (\cos \theta) u_1 \rangle + \langle u_2, (\sin \theta) z_2 \rangle = \cos \theta + \varepsilon \sin \theta$. So, pick $\theta \in \left(0, \frac{\pi}{2} \right]$ so that $\sin \theta < \varepsilon$. Then $\cos \theta + \varepsilon \sin \theta > \cos \theta + \sin^2 \theta > \cos^2 \theta + \sin^2 \theta = 1$ and hence, $\langle u, z \rangle > 1$. $\square$

Lemma 3.2.8. *Suppose v, w are points in $P(X)$ with $\langle v, w \rangle > -1$. Let u denote the orthogonal projection of w onto $v^\perp$, normalized to have unit length:*

$$u = \frac{2 - \langle v, w \rangle \, v}{(1 - \langle v, w \rangle^2)^{\frac{1}{2}}}.$$

(So, u is the unit tangent vector to the circular arc in $\mathbb{S}_1^n$ from v to w.) Let Y be the face of X such that v belongs to the relative interior of σ_Y. Then there is a vector z in $Lk(v, P(X))$ $(= \widetilde{P}(Y, v^\perp))$ so that $\langle u, z \rangle \geq 1$. Moreover, if neither condition (i) nor (ii) of Theorem 3.2.2 holds, then z can be chosen so that $\langle u, z \rangle > 1$.

Proof: Denote the orthogonal projections of v and w onto $\mathrm{Span}(Y)$ by v_2 and w_2, respectively. Since $v_2 = 0$, $u_2 = (1 - \langle v, w \rangle^2)^{-\frac{1}{2}} w_2$. For any $x \in Y, 0 \geq \langle w, x \rangle = \langle w_1, x \rangle + \langle w_2, x \rangle = \langle w_2, x \rangle$. Hence, w_2 and its positive multiple u_2 lie in $C^*(Y)$. Thus, we can apply Lemma 3.2.7 to find $z \in \widetilde{P}(Y)$ with $\langle u, z \rangle \geq 1$.

As for the question of strict inequality, it remains to show that cases (a), (b) and (c) of Lemma 3.2.7 correspond to situations (i) and (ii) of Theorem 3.2.2.

Case (a): $u_2 = 0$. Then $w_2 = 0$ and hence, $w \in \check{Y}$. Thus, both v and w belong to σ_Y and we are in situation (i).

Case (b): u_2 *is spacelike and* $d(G_{u_2}, Y) = 0$. By Lemma 3.2.5 (ii), either G_{u_2} is a supporting hyperplane of Y, or G_{u_2} meets $S_\infty(Y)$ in a single point y. If G_{u_2} is a supporting hyperplane, then it intersects Y in a face Y'. For any $x \in Y'$, $\langle w, x \rangle = \langle w_2, x \rangle = \langle u_2, x \rangle = 0$, so $Y' \subset H_v \cap H_w$, i.e., v, w both belong to $\sigma_{Y'}$. This is situation (i). Otherwise, $H_v \cap H_w \cap X = \emptyset$ and $H_v \cap H_w$ meets $S_\infty(X)$ at y. By Lemma 3.2.4, y is a cusp point of X. Hence, $v, w \in P_y$ and we are in situation (ii).

20

Case (c): u_2 *is lightlike and represents a point y in $S_\infty(Y)$.* We then argue exactly as in the end of the previous paragraph, to conclude that we are again in situation (ii). □

Proof of Theorem 3.2.2: Consider the function

$$f(v) = \cos^{-1}\langle v, w\rangle$$

defined on the set

$$U = \{v \in P(X) \mid \langle v, w\rangle > -1\}.$$

By Lemma 3.2.1, for any $c > -1$,

$$U_c = \{v \in P(X) \mid \langle v, w\rangle \geq c\}$$

is a compact subset of U.

Claim: *For any $v \in U$ there is a point $r(v) \in U$ such that*

$$d(v, r(v)) + f(r(v)) \leq f(v)$$

and $r(v) \neq v$ if $v \neq w$.

To prove the claim, suppose $v \in U$, $v \neq w$. Let u and z be as in Lemma 3.2.8, i.e.,

$$u = \frac{w - \langle v, w\rangle v}{(1 - \langle v, w\rangle^2)^{\frac{1}{2}}}, \quad \text{and}$$
$$\langle u, z\rangle \geq 1, \qquad z \in Lk(v, P(X)).$$

For sufficiently small $\varepsilon > 0$, put

$$r(v) = (\cos\varepsilon)v + (\sin\varepsilon)z.$$

Let $\theta = \cos^{-1}\langle v, w\rangle$ so that $w = (\cos\theta)v + (\sin\theta)u$. Then

$$\langle r(v), w\rangle = \cos\varepsilon\cos\theta + \langle u, z\rangle\sin\varepsilon\sin\theta$$
$$\geq \cos\varepsilon\cos\theta + \sin\varepsilon\sin\theta$$
$$= \cos(\theta - \varepsilon).$$

Hence,

$$d(v, r(v)) + f(r(v)) \leq \varepsilon + \cos^{-1}\langle r(v), w\rangle$$
$$\leq \varepsilon + \theta - \varepsilon$$
$$= f(v).$$

If neither condition (i) nor (ii) of the theorem holds, then we can choose z so that $\langle u, z\rangle > 1$ (Lemma 3.2.8). Then $\langle r(v), w\rangle > \cos(\theta - \varepsilon)$, which gives a strict inequality, $d(v, r(v)) + f(r(v)) < f(v)$.

21

Now fix $v \in U$ and let V be the set

$$V = \{x \in U \mid d(v,x) + f(x) \le d(v,r(v)) + f(r(v))\}.$$

Note that for any $x \in V$, $r(x)$ also lies in V since

$$d(v,r(x)) + f(r(x)) \le d(v,x) + d(x,r(x)) + f(r(x))$$
$$\le d(v,x) + f(x).$$

Note also that V is contained in $U_{\langle v,w \rangle}$ since for $x \in V$

$$f(x) \le d(v,x) + f(x) \le d(v,r(v)) + f(r(v)) \le f(v).$$

Thus, V is nonempty (since $r(v) \in V$) and compact, so f attains its minimum on V, say at $p \in V$. We claim that $p = w$. For if not, then $r(p)$ lies in V and since $r(p) \ne p$,

$$0 < d(p,r(p)) \le f(p) - f(r(p)),$$

which contradicts the minimality of f at p. Hence, $p = w$. In particular, $w \in V$. Thus,

$$d(v,w) + f(w) \le d(v,r(v)) + f(r(v)) \le f(v)$$

with the last inequality strict except under conditions (i) or (ii). Since $\langle w,w \rangle = 1$, $f(w) = 0$, so this proves the theorem.

3.3 Some consequences of Theorem 3.2.2. The goal of this subsection is to prove Proposition 3.3.3, below, which is the key technical lemma in the proof of the main result. We begin by recalling a general result concerning the orthogonal join of two piecewise spherical polyhedra.

Suppose that K_1 and K_2 are piecewise spherical polyhedra and that $K_1 * K_2$ denotes their orthogonal join (cf. 1.3). If $x_1 \in K_1$, $x_2 \in K_2$, then there is a geodesic segment (= arc of a great circle) of length $\frac{\pi}{2}$ in $K_1 * K_2$ from x_1 to x_2. For $t \in \left[0, \frac{\pi}{2}\right]$, let $[x_1, x_2, t]$ denote the point on this segment of distance t from x_1.

The following is proved in [CD1] as Lemma A2 and A6 of the Appendix.

Lemma 3.3.1. *Suppose that K_1, K_2 are nonempty piecewise spherical polyhedra and that d_1, d_2 and d denote the intrinsic metrics on K_1, K_2 and $K_1 * K_2$, respectively. Let $x = [x_1, x_2, s]$ and $y = [y_1, y_2, t]$ be points in $K_1 * K_2$. Then $d(x,y) \le \pi$. Moreover, $d(x,y) = \pi$ if and only if $s = t$ and one of the following conditions holds:*

(a) $s = 0$ *and* $d_1(x_1, y_1) \ge \pi$,

(b) $s = \frac{\pi}{2}$ *and* $d_2(x_2, y_2) \ge \pi$,

(c) $s \ne 0$, $s \ne \frac{\pi}{2}$, $d_1(x_1, y_1) \ge \pi$ *and* $d_2(x_2, y_2) \ge \pi$.

In the next lemma and its proof we return to the notation of 3.2.6.

Lemma 3.3.2. *Let Y be a convex polyhedral set in $\mathbb{H}^{n-1}$ and let $v, w \in \widetilde{P}(Y)$ be such that $\tilde{d}(v, w) \geq \pi$, where $\tilde{d}$ is the intrinsic metric on $\widetilde{P}(Y)$. Then*

$$\langle v, w \rangle \leq -1.$$

In other words, the hyperbolic hyperplanes $G_v = v^{\perp} \cap \mathbb{H}^{n-1}$ and $G_w = w^{\perp} \cap \mathbb{H}^{n-1}$ are either parallel or ultraparallel.

Proof: If Y is of dimension $n-1$ (the maximum possible), then $\widetilde{P}(Y) = P(Y)$ and the lemma follows immediately from Theorem 3.2.2. So, suppose $\dim Y < n - 1$. For any vector $v \in \widetilde{P}(Y)$ let v_1 and v_2 denote its orthogonal projections onto $Y^{\perp}$ ($= \mathbb{R}^{\ell}$) and $\mathrm{Span}(Y)$, respectively. If $q(v_i) \neq 0$, for $i = 1, 2$, let $\bar{v}_i = q(v_i)^{-\frac{1}{2}} v_i$ be the corresponding unit vector. Thus, $\bar{v}_1 \in \mathbb{S}^{\ell-1}$ and $\bar{v}_2 \in P(Y)$. We want to apply the previous lemma to $\widetilde{P}(Y) = \mathbb{S}^{\ell-1} * P(Y)$. In the notation of that lemma, $v = [\bar{v}_1, \bar{v}_2, s]$ and $w = [\bar{w}_1, \bar{w}_2, t]$, where $s = q(v_1)^{\frac{1}{2}}$ and $t = q(w_1)^{\frac{1}{2}}$. Hence, Lemma 3.3.1 gives that $\tilde{d}(v, w) = \pi$, that $s = t$, and that one of the following conditions holds:

(a) v and w are antipodal points on the suspension sphere $\mathbb{S}^{\ell-1}$,

(b) v and w lie in $P(Y)$ and $d(v, w) \geq \pi$, or

(c) $q(v_1) \neq 0$, $q(v_2) \neq 0$, $v_1 = -w_1$, and $d(\bar{v}_2, \bar{w}_2) \geq \pi$.

In case (a), $\langle v, w \rangle = -1$. In case (b), $\langle v, w \rangle \leq -1$, by Theorem 3.2.2. So, suppose (c) holds. Then $\langle \bar{v}_1, \bar{w}_1 \rangle = -1$ and (by Theorem 3.2.2) $\langle \bar{v}_2, \bar{w}_2 \rangle \leq -1$. Since $s = t$, $q(v_1) = q(w_1)$ and $q(v_2) = q(w_2)$. Hence, $\langle v_1, w_1 \rangle = -q(v_1)$ and $\langle v_2, w_2 \rangle \leq -q(v_2)$. Thus, $\langle v, w \rangle = \langle v_1, w_1 \rangle + \langle v_2, w_2 \rangle \leq -(q(v_1) + q(v_2)) = -1$, so the lemma also holds in this case. $\qquad\square$

In the next proposition we return to the situation considered in 3.1.

Proposition 3.3.3. *Let $\gamma = (\gamma_1, \ldots, \gamma_k)$ be a closed local geodesic in $P(X)$ with break points $v_0, \ldots, v_{k-1}$, $v_k = v_0$. (Interpret the indices as integers modulo k.) Let $H_i = H_{v_i}$ and $G_i = H_{i-1} \cap H_i$ be the subspaces of $\mathbb{H}^n$ considered in 3.1.*

(i) The subspaces G_i and G_{i+1} are either parallel or ultraparallel hyperplanes in H_i.

(ii) For $i = 0, \ldots, k-1$, suppose that G_i and G_{i+1} are parallel. Let $y_i \in S_{\infty}(H_i)$ be the point at infinity where G_i and G_{i+1} meet. Further, suppose that $y_0 = y_1 = \cdots = y_{k-1}$. Then $\ell(\gamma) = 2\pi$ and, if y denotes the common value of the y_i, then y is a cusp point of X.

Proof: (i) As in 1.5, from the broken geodesic γ we get unit tangent vectors, γ_i' and γ_{i+1}' in $Lk(v_i, P(X))$, where, in fact, γ_i' and γ_{i+1}' are the unit outward normals to G_i and G_{i+1} in H_i. Since γ is a local geodesic, condition (1.5.2) gives

$$d_i(\gamma_i', \gamma_{i+1}') \geq \pi,$$

where d_i is the intrinsic metric on $Lk(v_i, P(X))$. Let Y_i be the face of X such that v_i belongs to the relative interior of σ_{Y_i}, i.e., $Y_i = H_i \cap X$. By Lemma 2.6.2 (1), $Lk(v_i, P(X)) = \widetilde{P}(Y_i, (v_i)^\perp)$. Applying the previous lemma with $Y = Y_i$, $v = \gamma_i'$ and $w = \gamma_{i+1}'$, we conclude that (i) holds.

(ii) Suppose the hypotheses of (ii) hold. Let Z_y denote the intersection of the half-spaces bounded by the H_i with a horosphere E_y at y. Then Z_y is an $(n-1)$-dimensional Euclidean prism (i.e., it is isometric to the Cartesian product of a polygon and an $(n-3)$-dimensional Euclidean space), and $\ell(\gamma)$ is the sum of the exterior angles. Thus, $\ell(\gamma) = 2\pi$. Moreover, if v, w are any two points on γ, then the length of the segment of γ between them is $\cos^{-1}\langle v, w \rangle$. Choose such v, w which are arbitrarily close and which do not lie in the same cell of $P(X)$. If y is not a cusp point, then by Theorem 3.2.2, $d(v, w) < \cos^{-1}\langle v, w \rangle$, which contradicts the hypothesis that γ is a local geodesic. Hence, y is a cusp point and the lemma is proved.

4. The main result.

4.1 $P(X)$ is large.

Theorem 4.1.1. *Suppose X is a hyperbolic convex polyhedral set of dimension n. Then*

1) *Its polar dual $P(X)$ is large,*

2) *If γ is any closed local geodesic of length 2π, then γ must lie in the subcomplex P_y for some cusp point y of X.*

The proof uses Proposition 3.3.3, together with an argument of Rivin and Hodgson [RH].

Proof: According to Lemma 1.6.1, $P(X)$ is large if and only if

(i) $\text{sys}(P(X)) \geq 2\pi$,

(ii) $\text{sys}(Lk(\sigma, P(X))) \geq 2\pi$ for all cells σ in $P(X)$.

Each cell σ of $P(X)$ is of the form $\sigma = \sigma_Y$ for some proper face Y of X. By Lemma 2.6.2 (2), $Lk(\sigma_Y, P(X)) = P(Y)$. We can assume, by induction on the dimension, that the theorem holds for Y; in particular, $\text{sys}(P(Y)) \geq 2\pi$, so condition (ii) holds. Thus, it suffices to check (i).

Let γ be a closed local geodesic in $P(X)$. As in subsection 3.1, we construct an $(n-1)$-dimensional hyperbolic cylinder $\widehat{S}_\gamma$ and its image S_γ in $\mathbb{H}^n$ so that $\ell(\gamma)$ is the sum of the exterior dihedral angles in S_γ. thus, $\widehat{S}_\gamma$ is isometric to $\mathbb{H}^n / \langle \rho_\gamma \rangle$ where the isometry ρ_γ is either hyperbolic or parabolic. So, there are two cases to consider.

Case 1: *ρ_γ is hyperbolic.* In this case the argument of [RH] shows that $\ell(\gamma) > 2\pi$. (Actually, [RH] is concerned only with dimension 3, but this part of the argument works in any dimension.) The argument goes as follows. Since ρ_γ is hyperbolic, there is a unique

closed geodesic $\widehat{\alpha}$ on $\widehat{S}_\gamma$ of shortest length ($\widehat{\alpha}$ is the image of the axis of ρ_γ). Let α denote its image in S_γ. Then α is a piecewise geodesic in $\mathbb{H}^n$. The break points of α lie on the boundaries of the "strips" S_i, i.e., such a break point lies on the intersection of two hyperplanes H_i and H_{i+1}. By Lemma 3.2 in [RH], the "turning angle" of α at a break point is $\leq$ the corresponding exterior dihedral angle between H_i and H_{i+1}. Thus, the sum of the turning angles of α is $\leq \ell(\gamma)$. By Theorem 3.1 in [RH], this sum of turning angles is $> 2\pi$. hence, $\ell(\gamma) > 2\pi$.

Case 2: ρ_γ *is parabolic.* In this case we are in the situation of part (ii) of Proposition 3.3.3. The proposition implies that γ lies in some P_y for y a cusp point of X and its proof shows that $\ell(\gamma) = 2\pi$.

Thus, $\ell(\gamma) \geq 2\pi$ in both cases, which verifies condition (i) and hence, statement (1) of the theorem. The analysis in Case 2 also shows that if $\ell(\gamma) = 2\pi$, then γ lies in some P_y, which is statement (2) of the theorem.

4.2 The induced metric at a cusp. We have the following converse to the last part of Theorem 3.2.2.

Proposition 4.2.1. *Suppose X is a hyperbolic convex polyhedral set of dimension n and that d is the intrinsic metric on $P(X)$. Suppose further that v, w are points in $P(X)$ and that either (i) v and w belong to the same cell of $P(X)$ or (ii) v and w belong to P_y for some cusp point y. Then*

$$d(v, w) = \cos^{-1} \langle v, w \rangle.$$

Proof: We shall only verify this in case (ii), the proof in case (i) being similar and easier. Let d_1 denote the intrinsic metric on P_y. We first note that, since P_y is the polar dual of a Euclidean convex set, if $v, w \in P_y$, then $\langle v, w \rangle \geq -1$ and $d_1(v, w) = \cos^{-1} \langle v, w \rangle$ (by Lemma 2.2.1 (2)). The proof is by induction on n. So, suppose the proposition holds in dimensions $< n$. We first claim that P_y is locally convex in $P(X)$. To see this, we must check condition (∗) of Lemma 1.6.5. Let $u \in P_y$ and Y the face of X such that u belongs to the relative interior of σ_Y. Then

$$Lk(u, P(X)) = \widetilde{P}(Y, u^\perp) = \mathbb{S}^{\ell-1} * P(Y)$$
$$Lk(u, P_y) = \mathbb{S}^{\ell-1} * P_y(Y)$$

where $P_y(Y) = P_y \cap Y$ and $\mathbb{S}^{\ell-1}$ is the unit sphere in $Y^\perp \cap u^\perp$. Let d' and d_1' denote the intrinsic metrics on $P(Y)$ and $P_y(Y)$, respectively. Suppose we are given $x, x' \in P_y(Y)$ such that $d_1'(x, x') \geq \pi$. Then $d_1'(x, x') = \pi$ and $\langle x, x' \rangle = -1$. By the inductive hypothesis, $d'(x, x') = \pi$. By Lemma 3.3.1, the same result holds after taking ℓ-fold suspensions. This verifies condition (∗) of Lemma 1.6.5 and hence, shows that P_y is locally convex in $P(X)$. By Lemma 1.6.4, any geodesic in P_y is actually geodesic in $P(X)$. Hence,

$$\cos^{-1} \langle v, w \rangle = d_1(v, w) = d(v, w).$$

Definition 4.2.2. Let $\text{Cone}(P_y)$ denote the orthogonal join of P_y with a point. The *completed polar dual* of X, denoted $\widehat{P}(X)$, is the piecewise spherical complex formed by gluing $\text{Cone}(P_y)$ to $P(X)$ along P_y for each cusp point y.

For example, if S is n-dimensional and of finite volume, then each P_y is isometric to the round sphere $\mathbb{S}^{n-2}$ (by Lemma 2.2.1 (3)), so $\text{Cone}(P_y)$ is a hemisphere. Thus, $\widehat{P}(X)$ is obtained from $P(X)$ by "capping off" the P_y with hemispheres. It is then clear that $\widehat{P}(X)$ is homeomorphic to the $(n-1)$-sphere.

In general, if K_1, K_2 are large piecewise spherical subcomplexes, and K_0 is a common subcomplex which is locally convex in both, then the result of gluing K_1 to K_2 along K_0 is large. (After taking Euclidean cones on K_1, K_2 and K_0, this is a special case of Gromov's Gluing Lemma; see [Pa, Lemma 4.3].) As a corollary of Proposition 4.2.1 and Theorem 4.1.1, we therefore have the following.

Corollary 4.2.3. $\widehat{P}(X)$ *is large.*

4.3 Convex polytopes. In this subsection we state two special cases of the above results.

Corollary 4.3.1. *Let X be an n-dimensional hyperbolic convex polytope. Then $P(X)$ is a large piecewise spherical structure on $\mathbb{S}^{n-1}$. Moreover, the systole of $P(X)$ is greater than 2π, as is the systole of the link of any cell in $P(X)$.*

Corollary 4.3.2. *Let X be an n-dimensional convex polyhedral set of finite volume in $\mathbb{H}^n$. Then $\widehat{P}(X)$ is a large, piecewise spherical structure on $\mathbb{S}^{n-1}$.*

5. The Gauss-Bonnet Formula for hyperbolic polytopes.

5.1 The quantity $\kappa(K)$. For σ a p-dimensional spherical cell, let $a(\sigma)$ denote its p-dimensional volume, normalized so that the volume of $\mathbb{S}^p$ is 1, i.e.,

$$a(\sigma) = \frac{\text{vol}(\sigma)}{\text{vol}(\mathbb{S}^p)}.$$

Also, let $a^*(\sigma) = a(\sigma^*)$, where $\sigma^* = C^*(\sigma) \cap S^p$ is the dual cell to σ (cf. 2.3.1).

Given a finite, piecewise spherical cell complex K, consider the following quantity:

$$(5.1.1) \qquad \kappa(K) = 1 + \sum_{\sigma} (-1)^{\dim \sigma + 1} a^*(\sigma),$$

where the summation is over all cells σ of K. In [CMS] it is shown that the value of κ on the link of a vertex in a piecewise Euclidean complex plays the role of the Gauss-Bonnet integrand. For this reason, in analogy with a well-known conjecture of Hopf, we made the following conjecture in [CD2].

Conjecture 5.1.2. *Suppose that K is a large, piecewise spherical structure on the $(2m-1)$-sphere. Then $(-1)^m \kappa(K) \geq 0$.*

5.2 Hopf's formula. Suppose X^n is an n-dimensional hyperbolic convex polytope. Its normalized volume $v(X^n)$ is defined by

$$v(X^n) = \frac{\text{vol}(X^n)}{\text{vol}(\mathbb{S}^n)}.$$

Thus, $v(X^n)$ is a positive constant times the usual volume of X^n.

The Gauss-Bonnet Formula for hyperbolic convex polytopes is due to H. Hopf [H]. (Of course, Hopf's result predates the general Gauss-Bonnet Formula due to Allendoerfer and Weil [AW].) Hopf's formula is the following.

Theorem 5.2.1 (Hopf). *Suppose X^{2m} is a hyperbolic convex polytope of dimension $2m$. Then*

$$(-1)^m 2v(X) = \kappa(P(X)).$$

Remarks 5.2.3. (1) The analogous formula for X^{2m} spherical or Euclidean was proved by Poincaré in 1905, [P]. It states that if X is Euclidean then $\kappa(P(X)) = 0$ (this also holds in odd dimensions), while if X^{2m} is a spherical cell, then $2a(X) = \kappa(P(X))$.

(2) In writing the quantity $\kappa(P(X))$, we have dualized twice: once in forming $P(X)$ and a second time in taking the normalized "exterior angles" $a^*(\sigma)$. Thus, $\kappa(P(X))$ could have been written as an alternating sum of interior angles of X (and this is the actual form of the formulas of Poincaré and Hopf).

(3) If we subdivide X^{2m} into $2m$-dimensional cells, $X_1, \ldots, X_k$, then a formal calculation shows that $\kappa(P(X)) = \sum \kappa(P(X_i))$. (Compare Theorem III in [AW].) The left-hand side of Hopf's formula is also clearly additive with respect to subdivisions. Thus, it suffices to prove the formula in the case of a simplex.

(4) The formula also holds when X is a convex polyhedral set of finite volume (use a limit argument). In this case it does not matter whether we write the right-hand side as $\kappa(P(X))$ or as $\kappa(\widehat{P}(X))$. The reason is that for any cusp point y, $\kappa(P_y) = 0$ (by Poincaré's result) and hence, $\kappa(\text{Cone}(P_y)) = 0$ (since κ is multiplicative for orthogonal joins, by [CMS], formula (3.2.9). Thus, $\kappa(P(X)) = \kappa(\widehat{P}(X))$.

Since $v(X^{2m}) > 0$, we have the following corollary.

Corollary 5.2.3. *Let X^{2m} be a convex polyhedral set of finite volume in $\mathbb{H}^{2m}$ (e.g., a convex polytope). Then Conjecture 5.1.2 holds for $\widehat{P}(X)$. In fact, we have a strict inequality $(-1)^m \kappa(\widehat{P}(X)) > 0$.*

6. Spaces of geometric structures.

6.1 The space of piecewise spherical structures. Let K be a finite simplicial complex. Denote by $K^{(i)}$ the set of i-simplices in K. A piecewise spherical structure on K is determined by a function $\ell : K^{(1)} \to (0, \pi)$, written as $e \to \ell_e$, giving the edge lengths. Conversely, such a function comes from a piecewise spherical structure on K if for each simplex σ in K of dimension > 1 there is a spherical simplex with edge lengths as prescribed by ℓ.

This last condition is equivalent to a simple condition using linear algebra, which we shall now describe. Given $\ell : K^{(1)} \to (0, \pi)$ define a new function $x : K^{(i)} \to (-1, 1)$ by $x_e = \cos(\ell_e)$. Clearly, ℓ and x determine one another. Suppose that $\sigma \in K^{(m)}, m > 1$, and that $\{v_1, \cdots, v_{m+1}\}$ is the vertex set of σ. Given an arbitrary function $x : K^{(1)} \to (-1, 1)$, define a $(m + 1) \times (m + 1)$ symmetric matrix $c_\sigma(x)$ by

$$
c_\sigma(x)_{ij} = \begin{cases} 1 & , \ i = j \\ x_{\overline{v_i v_j}} & , \ i \neq j \end{cases}
$$

where $\overline{v_i v_j}$ denotes the edge from v_i to v_j. Then there exists a spherical m-simplex with edge lengths as prescribed by x if and only if $c_\sigma(x)$ is positive definite. (If σ can actually be realized as a m-simplex in $\mathbb{S}^m$, then $v_1, \ldots, v_{m+1}$ are unit vectors in $\mathbb{R}^{m+1}$ and $c_\sigma(x)$ is the matrix of inner products (v_i, v_j).)

The space $\mathcal{PS}_K$ of piecewise spherical structures on K is defined to be the subset of the Euclidean space $\mathbb{R}^{K^{(1)}}$ consisting of all functions $x : K^{(1)} \to (-1, 1)$ such that $c_\sigma(x)$ is positive definite for all simplices σ in K (of dimension > 1).

Lemma 6.1.1 $\mathcal{PS}_K$ *is a convex open subset of* $\mathbb{R}^{K^{(1)}}$.

Proof: That the $c_\sigma(x)$ be positive definite is clearly an open condition. A convex linear combination of positive definite matrices is positive definite; hence, $\mathcal{PS}_K$ is a convex subset of $\mathbb{R}^{K^{(1)}}$. $\qquad \square$

Given $x \in \mathcal{PS}_K$, let K_x denote the space K with the piecewise spherical structure determined by x. Similarly, if σ is a simplex in K, then σ_x denotes the corresponding spherical simplex determined by x.

A point $x \in \mathcal{PS}_K$ is *extra large* if the following two conditions hold:

a) sys $(K_x) > 2\pi$

b) sys $(Lk(\sigma_x, K_x)) > 2\pi$ for all simplices σ in K.

(See §1.6 for the definition of "sys".) Let EL_K denote the set of those x in $\mathcal{PS}_K$ which are extra large.

Moussong showed in [M, Lemma 5.11] that the function sys : $\mathcal{PS}_K \to \mathbb{R}$ defined by $x \mapsto$ sys (K_x) is lower semicontinuous. This gives the following result.

Lemma 6.1.2 (Moussong) EL_K *is an open subset of* $\mathcal{PS}_K$.

If X is a convex polyhedral set in hyperbolic space, without cusp points, then our main result (Theorem 4.1.1) asserts that $P(X)$ is extra large. Hence, whenever K is combinatorially equivalent to the polar dual of such an X, the space EL_K is nonempty.

6.2 Spaces of polytopes. Suppose that Q is some n-dimensional convex simplicial polytope. (*Simplicial* means that each proper face of Q is a simplex). Since Q is simplicial, its dual polytope Q^* is called *simple*. We are only interested in Q and Q^* up to combinatorial equivalence. The boundary complex of Q is a simplicial complex which we shall denote by K. (Thus, K is PL-homeomorphic to $\mathbb{S}^{n-1}$.)

Let H_K (respectively, S_K or E_K) denote the space of isometry classes of hyperbolic (respectively, spherical or Euclidean) polytopes which are combinatorially equivalent to Q^*. The topology on these spaces will be described in the course of the proof of the next lemma.

Lemma 6.2.1 *The spaces H_K, S_K, and E_K are naturally smooth manifolds of dimension $nf_0 - \binom{n+1}{2}$, where f_0 denotes the number of vertices in K.*

Proof: Suppose that X is a convex n-cell in $\mathbb{H}^n$ combinatorially isomorphic to Q^*. For each vertex v_i of K we have a codimension one face of X and a outward pointing unit normal u_i in the unit pseudo-sphere $\mathbb{S}_1(\mathbb{R}^{n,1})$. Thus, X determines an f_0-tuple $u : K^{(0)} \to \mathbb{S}_1(\mathbb{R}^{n,1})$ of vectors in $\mathbb{S}_1(\mathbb{R}^{n,1})$.

A small neighborhood of u in $[S_1(\mathbb{R}^{n,1})]^{f_0}$ will determine a polytope of the same combinatorial type (since X is simple). Since $\dim \mathbb{S}_1(\mathbb{R}^{n,1}) = n$, this shows that the space of such X is a manifold of dimension nf_0. The isometry group $O(n,1)$ of $\mathbb{H}^n$ acts properly and freely on this manifold and H_K is the quotient manifold. Since $\dim O(n,1) = \binom{n+1}{2}$, we conclude that $\dim H_K = nf_0 - \binom{n+1}{2}$. The analysis for S_K and E_K is entirely similar. $\square$

Remark Actually H_K should be called the space of "marked" hyperbolic polytopes of the same combinatorial type as Q^* (since we have prescribed an identification of the codimension one faces of X with the vertices of K). The finite group G of combinatorial symmetries of Q^* acts on H_K and H_K/G is the "unmarked" space. Similar remarks apply to S_K and E_K.

If X represents an element of H_K (respectively, S_K or E_K), then $P(X)$ represents an element of $\mathcal{PS}_K$. This defines a map $P : H_K \to \mathcal{PS}_K$ (respectively, $P : S_K \to \mathcal{PS}_K$ or $P : E_K \to \mathcal{PS}_K$). The following result for $n = 3$ is proved as Corollary 4.6 of [RH]. Another 3-dimensional result, Theorem 4.1 of [RH], immediately implies that both results hold in all dimensions ≥ 3.

Proposition 6.2.2 ([RH]) *For $n \geq 3$, the maps $P : H_K \to \mathcal{PS}_K$ and $P : S_K \to \mathcal{PS}_K$ are embeddings.*

The map $P : E_K \to \mathcal{PS}_K$ is not an embedding. To see this, first note that the elements of $P(E_K)$ are all isometric to the round sphere $\mathbb{S}^{n-1}$. Thus, the points in $P(E_K)$ may be regarded as geodesic triangulations of $\mathbb{S}^{n-1}$ of the combinatorial type of K. Conversely, any such geodesic triangulation is clearly in $P(E_K)$. (If $\{v_1, \ldots, v_k\} \subset \mathbb{S}^{n-1}$ is the vertex set of such a triangulation, then the Euclidean n-cell X defined as $\{z \in \mathbb{R}^n \mid \langle z, v_i \rangle \leq 1, 1 \leq i \leq k\}$ gives back the geodesic triangulation as its polar dual.) Arguing as in Lemma 6.2.1, we see that $P(E_K)$ is a submanifold of dimension $(n-1)f_0 - \binom{n}{2}$. Hence, the fiber of $P : E_K \to \mathcal{PS}_K$ has dimension $f_0 - n$.

Lemma 6.2.3 $P(E_K) \subset \overline{P(H_K)} \cap \overline{P(S_K)}$.

Proof: Realize a point in $P(E_K)$ as a geodesic triangulation of $\mathbb{S}^{n-1}$. Identify $\mathbb{S}^{n-1}$ with a codimension one sphere in the unit pseudo-sphere $\mathbb{S}_1^n(\mathbb{R}^{n,1})$, e.g., let $z \in \mathbb{H}^n$ and identify $\mathbb{S}^{n-1}$ with the intersection of $z^\perp$ and $\mathbb{S}_1^n(\mathbb{R}^{n,1})$. Push the vertices of the triangulation slightly into the "half-pseudo-sphere" defined by $v \cdot z < 0$. If $u_1, \ldots, u_k$ are the resulting unit vectors in $\mathbb{S}_1^n(\mathbb{R}^{n,1})$, then the hyperbolic polytope X defined by $X = \{w \in \mathbb{H}^n \mid u_i \cdot w \leq 0\}$ has $P(X)$ combinatorially equivalent to K and close to the original triangulation of $\mathbb{S}^{n-1}$. Thus, $P(E_K) \subset \overline{P(H_K)}$. The proof that $P(E_K) \subset \overline{P(S_K)}$ is similar, with the pseudo-sphere $\mathbb{S}_1^n(\mathbb{R}^{n,1})$ replaced by the unit sphere $\mathbb{S}^n$ in $\mathbb{R}^{n+1}$. $\qquad \square$

Thus, when K is the boundary complex of a simplicial polytope, the following picture has emerged. The space $\mathcal{PS}_K$ of piecewise spherical structures on K is diffeomorphic to Euclidean space of dimension f_1, where f_1 denotes the number of edges in K. $\mathcal{PS}_K$ contains submanifolds $P(H_K)$ and $P(S_K)$ each of dimension $nf_0 - \binom{n+1}{2}$. The fact that $P(H_K)$ is nonempty means that the open subset EL_K of extra large structures is nonempty. The fact that $P(E_K) \subset \overline{P(H_K)}$ means that any geodesic triangulation of the round sphere can be deformed to an extra large piecewise spherical structure.

An interesting problem is to understand the topology of EL_K. For example, when is EL_K connected?

6.3 The Lower Bound Theorem. Since $P(H_K)$ is a submanifold of $\mathcal{PS}_K$, we have

$$\dim P(H_K) \leq \dim \mathcal{PS}_K.$$

In the combinatorial theory of convex polytopes, this fact is known as "the Lower Bound Theorem" which we state below.

Theorem 6.3.1 (Barnette [Ba]) *Let Q be an n-dimensional simplicial polytope, $n \geq 3$, with f_0 vertices and f_1 edges. Then*

$$f_1 \geq nf_0 - \binom{n+1}{2}.$$

For $n > 3$, the above inequality is strict unless Q is a so called "stacked polytope". This means that $K (= \partial Q)$ is obtained from the boundary of an n-simplex via subdivisions which involve adjoining barycenters of (n-1)-simplicies. Thus, if $n > 3$ and Q is not stacked, then $P(H_K)$ is a proper subset of EL_K.

Some explicit examples of elements in $EL_K - P(H_K)$ are given in [M]. One such class of examples occurs when Q^* is the product of two 2-simplices (so that K is the join of two triangles). Realize Q^* as the Cartesian product of two 2-simplices in $\mathbb{H}^2$ with acute angles. Its polar dual is a piecewise spherical structure x on K which is the orthogonal join of two large triangles. It follows that x is large. Consider the 6×6 matrix $c(x)$ obtained by taking inner products of the 6 unit vectors in $\mathbb{R}^{2,1} \times \mathbb{R}^{2,1}$ which are normal to the codimension-one faces of Q^*. Thus, $c(x)$ has signature (4,2). (Alternately described, $c(x)$ is the matrix of cosines of edge lengths in the piecewise spherical structure x.) There are

9 edges of K of length $\pi/2$. If we deform x to x' by slightly increasing these edge lengths, then it follows from Moussongs Theorem that x' is extra large. On the other hand, for small deformations, the matrix $c(x')$ will still have signature (4,2). Hence, x' cannot arise as a polar dual of a hyperbolic polyhedra (since this would require $c(x')$ to be of signature (5,1)).

References

[AW]   C.B. Allendoerfer and A. Weil, *The Gauss-Bonnet Theorem for Riemannian polyhedra*, Trans. Am. Math. Soc. 53 (1943), 101–129.

[A]   E. M. Andreev, *On convex polyhedra in Lobačevskiĭ space*, Math. USSR Sb. 10 (1970), 413–440.

[B]   W. Ballman, *Singular spaces of non-positive curvature*, in "Sur les Groupes Hyperboliques d'apres Mikhael Gromov", edited by E.Ghys and P. de la Harpe, Progress in Math. 83, Birkhaüser, Boston, Basel, Berlin, 1990, 189–202.

[Ba]   D. Barnette, *A proof of the lower bound conjecture for convex polytopes*, Pacific J. Math. 46,(1973), 349-345.

[Br]   M. R. Bridson, *Geodesics and curvature in metric simplicial complexes*, in "Group Theory from a Geometrical Viewpoint," edited by E. Ghys, A. Haefliger and A. Verjovsky, Word Scientific, Singapore, 1991, 373–463.

[CD1]   R. Charney and M. W. Davis, *Singular metrics of nonpositive curvature on branched covers of Riemannian manifolds*, Amer. J. Math. 115 (1993), 929-1009.

[CD2]   _____, *The Euler characteristic of a nonpositively curved, piecewise Euclidean manifold*, to appear in Pacific J. Math.

[CMS]   J. Cheeger, W. Müller and R. Schrader, *On the curvature on piecewise flat spaces*, Commun. Math. Phys. 92 (1984), 405–454.

[G]   M. Gromov, *Hyperbolic groups*, in "Essays in Group Theory," edited by S. M. Gersten, M.S.R.I. Publ. 8, Springer-Verlag, New York and Berlin, 1987, 75–264.

[H]   H. Hopf, *Die Curvatura integra Clifford-Kleinscher Raumformen*, Nachr. Ges. Wiss. Gottingen, Math. Phys. (1925), 131–141.

[M]   G. Moussong, *Hyperbolic Coxeter groups*, Ph.D. thesis, The Ohio State Univ., 1988.

[O]   B. O'Neill, "Semi-Riemannian Geometry with Applications to Relativity," Academic Press, New York, 1983.

[Pa]   F. Paulin, *Constructions of hyperbolic groups via hyperbolizations of polyhedra*, in "Group theory from a Geometrical Viewpoint," edited by E. Ghys, A. Haefliger and A. Verjovsky, World Scientific, Singapore, 1991, 313–372.

[P]   H. Poincaré, *Sur la généralisation d'un thèoréme elémentaire de Géométrie*, C. R. Acad. Sc. Paris (1905), I, 113–117.

[R1]   I. Rivin, *On the geometry of convex polyhedra in hyperbolic 3-space*, Ph.D. thesis, Princeton Univ., 1986.

[R2]   ⸻, *On the geometry of ideal polyhedra in hyperbolic 3-space*, Topology 32 (1993).

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

[S]   J.-M. Schlenker, *Surfaces convex dans des espaces Lorentziens à courbure constante* doctoral thesis, Ecole Polytechnique, 1994.

[T]   W. Thurston, *The geometry and topology of 3-manifolds*, reproduced lecture notes, Princeton Univ., 1979.

[Tr]   M. Troyanov, *Espaces à courbure négative et groupes hyperboliques*, in "Sur les Groupes Hyperboliques d'apres Mikhael Gromov," edited by E. Ghys and P. de la Harpe, Progress in Math. 83, Birkhaüser, Boston, Basel, Berlin, 1990, 47–66.