WEAKLY INSCRIBED POLYHEDRA

HAO CHEN AND JEAN-MARC SCHLENKER

Abstract. We study convex polyhedra in $\mathbb{R}P^3$ with all their vertices on a sphere. In particular, we do not require that the polyhedra lie in the interior of the sphere, hence the term "weakly inscribed". With a purely combinatorial characterization of the 1-skeleta of weakly inscribed polyhedra, we advance one step towards a complete answer to Steiner's inscribability problem, which dates back to 1832.

The key idea of this paper is to regard $\mathbb{R}P^3$ as a combination of the hyperbolic space and the de Sitter space, with the sphere as the common ideal boundary. Then weakly inscribed polyhedra can be interpreted as ideal polyhedra, and our combinatorial characterization follows from a linear programming characterization of their exterior dihedral angles. We also describe the hyperbolic-de Sitter structures induced on the boundaries of weakly ideal polyhedra. They are obtained by properly gluing hyperbolic planes to a de Sitter surface along their ideal boundaries.

Contents

Date: May 20, 2018.

Hao Chen was partly supported by the National Science Foundation under Grant No. DMS-1440140 while the author was in residence at the Mathematical Sciences Research Institute in Berkeley, California, during the Fall 2017 semester.

Jean-Marc Schlenker's research was partially supported by University of Luxembourg research projects GeoLoDim (9/2014-8/2017) and NeoGeo (9/2017-8/2019), and by the Luxembourg National Research Fund projects INTER/ANR/15/11211745, INTER/ANR/16/11554412/SoS and OPEN/16/11405402.

1. Introduction

1.1. **Steiner's inscribability problem.** In 1832, Steiner [Ste32] asked the following question:

> 77) Wenn irgend ein convexes Polyëder gegeben ist, lässt sich dann immer (oder in welchen Fällen nur) irgend ein anderes, welches mit ihm in Hinsicht der Art und der Zusammensetzung der Grenzflächen übereinstimmt (oder von gleicher Gattung ist), in oder um eine Kugelfläche, oder in oder um irgend eine andere Fläche zweiten Grades beschreiben (d. h. dass seine Ecken alle in dieser Fläche liegen oder seine Grenzflächen alle diese Fläche berühren)?

In English, the question asks: Does every polyhedron have a combinatorially equivalent realization that is inscribed or circumscribed to a sphere, or to another quadratic surface? If not, which polyhedra have such realizations? In the parenthesis, Steiner made the following definition:

Definition 1.1. A polyhedron P is *inscribed to* a quadratic surface S if all the vertices of P lie on S.

Steiner also defined that P is *circumscribed* to S if all its facets are tangent to the surface. We will see that circumscription and inscription are closely related through polarity, hence we only need to focus on one of them. A polytope is called *inscribable to* S if it has a combinatorially equivalent realization that is inscribed to S.

Note that we say P is inscribed "to" S instead of "in" S. This apparent grammar mistake makes a significant distinction here.

In elementary geometry, when we say that a figure A is inscribed in another figure B, we usually require that A is contained inside B, hence the preposition "in". This is also assumed in many advanced studies on inscribed polytopes. Indeed, Steiner's problem were usually considered in Euclidean space. In this case, a polyhedron inscribed to a sphere must be inside the sphere. But this requirement is not present in Steiner's definition. Actually, since Steiner's problem is obviously projectively invariant, it is quite natural to consider it in projective space. It is then possible for a polyhedra to be inscribed *to* but not *in* the sphere.

We now define the mathematical terminology to avoid the subtlety of prepositions in the informal language.

Definition 1.2. In the projective space $\mathbb{R}\mathrm{P}^3$, a polyhedron P inscribed to a quadratic surface S is *strongly inscribed in* S if the interior of P is disjoint from S, or *weakly inscribed to* S otherwise.

Hence "strongly inscribed in" is the synonym of "inscribed *in*", and "weakly inscribed to" means "inscribed *to* but not *in*".

1.2. **History.** Steiner's problem stayed open for nearly a century. There were even beliefs [Brü00] that every simplicial polyhedron is strongly inscribable in a sphere. The first polyhedra without any strongly inscribed realization were discovered in 1928 by Steinitz [Ste28].

Hodgson, Rivin and Smith [HRS92] gave a characterization for the combinatorial types of polyhedra that are strongly inscribable in a sphere. They view the sphere as the ideal boundary of the projective model of the hyperbolic space, so that strongly inscribed polyhedra become hyperbolic ideal polyhedra. The characterization is then formulated in terms of hyperbolic exterior dihedral angles. More specifically, a polyhedron is strongly inscribable in a sphere if and only if one can assign weights (angles) to the edges subject to a family of equalities and inequalities. Hence strong inscribability in the sphere can be determined by solving a feasibility problem.

Rivin [Riv94] also gave another characterization in terms of the metric induced on the boundaries of ideal hyperbolic polyhedra. More specifically, every complete hyperbolic metric of finite area on an n-times punctured sphere can be isometrically embedded as the boundary of a n-vertices

polyhedron strongly inscribed to a sphere, viewed as an ideal hyperbolic polyhedron (possibly degenerate and contained in a plane).

Other quadratic surfaces were mostly ignored until very recently. There are three quadratic surfaces in $\mathbb{R}\mathrm{P}^3$ up to projective transformation: the sphere, the one-sheeted hyperboloid, and the cylinder. Danciger, Maloni and the second author [DMS14] characterized the combinatorial types of polyhedra that are strongly inscribable in a one-sheeted hyperboloid or in a cylinder. This characterization is based on the possible dihedral angles of ideal polyhedra in the Anti-de Sitter space (for the one-sheeted hyperboloid) or in the "half-pipe" space (for the cylinder; see [Dan14, Dan13]). For the one-sheeted hyperboloid, they also describe the possible induced metrics on the boundary of those ideal polyhedra in the Anti-de Sitter space.

Remark 1.3. The definition in [DMS14] is slightly stronger than ours. They require that $P \cap S$ consists of exactly the vertices of P. The two definitions are equivalent only when S is the sphere. Otherwise, it is possible that some edges of P are contained in S.

All these works, like many others, are about strong inscription. Steiner's original definition seems to be overlooked. For the sphere, Steiner's definition recently reappeared after a long detour.

Schulte [Sch87] considered higher dimensional generalizations of Steiner's problem in Euclidean space, namely d-dimensional polytopes with all the k-dimensional faces tangent to the sphere. It reduces to the strong inscription when $k = 0$. He then followed an idea of Grünbaum and Shephard [GS87], and defined a weaker version requiring only affine hulls of the k-faces to be tangent to the sphere. But since he worked in Euclidean space, the weak version with $k = 0$ coincides with the strong inscription.

Padrol and the first author [CP17] considered Schulte's generalizations in projective space, where they noticed that the weak version with $k = 0$ reduces to the weak inscription in the current paper. The projective setting is weaker then Schulte's Euclidean setting: there are polyhedra (e.g. the triakis tetrahedron) weakly inscribed to the sphere, but not strongly inscribable. However, the projective setting behaves much better under polarity, which led to the solution [CP17] of some Euclidean problems left open by Schulte.

1.3. **Main results.** The current paper answers Steiner's question for polyhedra weakly inscribed to a sphere. We also plan, in the future, to study polyhedra weakly inscribed to a one-sheeted hyperboloid or to a cylinder, in the hope of completely solving Steiner's problem.

Note that for any quadratic surface S, there is a polyhedron P that has no realization inscribed to S. A universal example is provided by the cube with one vertex truncated. This follows from the well-known fact that if seven vertices of a cube lie on a quadratic surface, so does the eighth one [BS08, Section 3.2]; see [CP17, Example 4.1] for a complete argument. Hence we seek to characterize polyhedra inscribed to quadratic surfaces in $\mathbb{R}\mathrm{P}^3$.

As we have mentioned, strong inscriptions in the sphere, the one-sheeted hyperboloid and the cylinder are essentially characterized in previous works [HRS92, Riv94, DMS14]. The current paper focuses on polyhedra weakly inscribed to a sphere.

In the following theorem, we consider a single vertex as a 1-cycle, and a single edge as a 2-cycle.

Theorem 1.4 (Combinatorial characterization)**.** *A 3-connected planar graph $\Gamma = (V, E)$ is the 1-skeleton of a polyhedron $P \subset \mathbb{R}\mathrm{P}^3$ weakly inscribed to a sphere if and only if*

(C1) *Γ admits a vertex-disjoint cycle cover consisting of two cycles*

and, if we color the edges connecting vertices on the same cycle by red (r), and those connecting vertices from different cycles by blue (b), then

(C2) *there is a cycle visiting all the edges (repetition allowed) along which the edge color has the pattern*

- *... bbrbbr... if the cycle cover contains a 1-cycle, or*
- *... brbr... otherwise.*

If the cycle cover contains a 1-cycle with a single vertex v, Condition **(C2)** has a much simpler formulation, namely that v is connected to every other vertex. We decide to adopt the current formulation for comparison with the other case.

This characterization is purely combinatorial in the sense that, unlike the characterization of strong inscription in [HRS92], it does not involve any feasibility problem. For strong inscription, efforts were also made [DS96, Che03] towards a purely combinatorial characterization without linear programming.

In fact, Theorem 1.4 is a consequence of the following linear programming characterization.

Theorem 1.5 (Linear programming characterization)**.** *A 3-connected planar graph $\Gamma = (V, E)$ is the 1-skeleton of a polyhedron $P \subset \mathbb{R}\mathrm{P}^3$ weakly inscribed to a sphere if and only if*

(C1) *Γ admits a vertex-disjoint cycle cover consisting of two cycles*

and, if we color the edges connecting vertices on the same cycle by red, and those connecting vertices from different cycles by blue, there is a weight function $w : E \to \mathbb{R}$ such that

(W1) *$w > 0$ on red edges, and $w < 0$ on blue edges;*

(W2) *w sums up to 0 over the edges adjacent to a vertex v, except when v is the only vertex in a 1-cycle, in which case w sums up to -2π over the edges adjacent to v.*

Recall that we consider a single vertex as a 1-cycle, and a single edge as a 2-cycle. The feasibility problem involved here is visibly much simpler than that in [HRS92]. In particular, there is no inequality on the non-trivial cuts.

Given a sphere $S \subset \mathbb{R}\mathrm{P}^3$, its interior (resp. exterior) can be seen as the projective model for the 3-dimensional hyperbolic space $\mathbf{H}^3$ (resp. de Sitter space $\mathrm{d}\mathbf{S}^3$). In [Sch98, Sch01], $\mathbf{H}^3$ and $\mathrm{d}\mathbf{S}^3$ together make up the *hypebolic-de Sitter space* (HS space for short) which is denoted by $\mathbf{HS}^3$. Then a polyhedron P ideal to S can be considered as an ideal polyhedron in $\mathbf{HS}^3$. We say that P is strongly ideal if P is contained in $\mathbf{H}^3$, or weakly ideal otherwise.

We will see in Section 8 that Theorem 1.5 follows from Theorem 3.2 below, which describes the possible dihedral angles of convex polyhedra in $\mathbf{HS}^3$. More specifically, the dihedral angles at the edges of P form a weight function θ satisfying all the conditions of Theorem 1.5 and, additionally, that $|\theta| < \pi$ and the sum of θ over the blue edges is bigger than -2π. These additional conditions are, however, redundant in the linear programming characterization, as we will prove in Section 8.

In view of Rivin's work [Riv94], we also provide a characterization for the geometric structure induced on the boundary of a weakly ideal polyhedron in $\mathbf{HS}^3$. This geometric structure, as distinguished from that induced on a strongly ideal polyhedron, contains a de Sitter part; that is, a part locally modeled on the de Sitter plane. We call this induced data an "HS-structure", since it is locally modeled on $\mathbb{HS}^2$, a natural extension of the hyperbolic plane by the de Sitter plane. Relevant definitions in the following statement will be recalled in the next section.

Theorem 1.6 (Metric characterization)**.** *Let P be a weakly ideal polyhedron in $\mathbf{HS}^3$ with n vertices. Then the induced HS-structure on ∂P is a complete, maximal HS structure on the punctured sphere, obtained by gluing copies of $\mathbf{H}^2$ to a de Sitter surface along their ideal boundaries by C^1 piecewise projective maps such that, at the "break points" where the maps fail to be projective, the second derivative has a positive jump. Conversely, each HS structure of this type is induced on a unique weakly ideal polyhedron in $\mathbf{HS}^3$.*

Note that both the hyperbolic and de Sitter parts of the metric have a well-defined real projective structure at infinity, so it is meaningful to ask for a piecewise projective gluing map. More explanations on the statement of Theorem 1.6 can be found in Section 3.3.

Remark 1.7. For the interest of physics audience, weakly ideal polyhedra can be interpreted as a description of interactions of "photons"; see [BBS11]. More specifically, $\mathbb{HS}^3$ models the link of an event in the space-time. The vertices on different boundary components of the de Sitter surface in Theorem 1.6, or, combinatorially, on different cycles in Condition **(C1)**, correspond to incoming and outgoing photons (depending on the direction of time) involved in an interaction. A special case is the single vertex in a degenerate boundary component, or, combinatorially, in a 1-cycle, which corresponds to an extreme BTZ-like singularity.

The paper is organized as follows. The essential definitions are made in Section 2 in a general setting. Then we can view polyhedra weakly inscribed to the sphere as ideal polyhedra in $\mathbf{HS}^3$. In

Section 3, we announce characterizations for the dihedral angles and induced metrics of weakly ideal polyhedra in $\mathbf{HS}^3$, which are actually reformulations of Theorems 1.5 and 1.6. We also outline the proof strategy, which is carried out in the following sections. In particular, there are some technical challenges which were not encountered in the previous works on strong inscription. For instance, the space of weakly inscribed polyhedra is not simply connected. Finally, in Section 8, we deduce Theorem 1.4 from the linear programming characterization of dihedral angles.

2. Definitions

We are mainly interested in three dimensional polyhedra. However, the definitions in this section are more general than strictly necessary, and cover the anti-de Sitter and half-pipe spaces that we hope to study in a further work. Lower dimensional cases are used as examples.

2.1. **The hyperbolic, anti-de Sitter and half-pipe spaces.** The *projective space* $\mathbb{R}\mathrm{P}^d$ is the set of linear 1-subspaces of $\mathbb{R}^{d+1}$. An *affine chart* of $\mathbb{R}\mathrm{P}^d$ is an affine hyperplane $H \subset \mathbb{R}^{d+1}$ which is identified to the set of linear 1-dimensional subspaces intersecting H. The linear hyperplane parallel to H is projectivized as the *hyperplane at infinity*; linear 1-dimensional subspaces contained in this hyperplane have no representation in the affine chart H.

Let $\mathbb{R}^{d+1}_{p,q}$ denotes $\mathbb{R}^{d+1}$ equipped with an inner product $\langle\cdot,\cdot\rangle$ of signature (p,q), $p+q \leq d+1$. We say that $\mathbb{R}^{d+1}_{p,q}$ is *non-degenerate* if $p+q = d+1$. For convenience, we will assume that

$$\langle \mathbf{x}, \mathbf{y}\rangle := \sum_{i=0}^{p-1} x_i y_i - \sum_{i=0}^{q-1} x_{d-i} y_{d-i}.$$

For $p+q \leq d+1$, we define $\mathbb{H}^d_{p,q} = \{\mathbf{x} \in \mathbb{R}^{d+1}_{p,q} \mid \langle \mathbf{x}, \mathbf{x}\rangle = -1\}$ and $\mathbb{S}^d_{p,q} = \{\mathbf{x} \in \mathbb{R}^{d+1}_{p,q} \mid \langle \mathbf{x}, \mathbf{x}\rangle = 1\}$. Both $\mathbb{H}^d_{p,q}$ and $\mathbb{S}^d_{p,q}$ are equipped with the metric induced by the inner product. The metrics of $\mathbb{H}^d_{p,q}$ and $\mathbb{S}^d_{q,p}$ differ only by a sign.

Here and through out this paper, if a space in $\mathbb{R}^{d+1}_{p,q}$ is denoted by a blackboard boldface letter, we use the corresponding simple boldface letter to denote its projectivization in $\mathbb{R}\mathrm{P}^d_{p,q}$. For example, $\mathbf{H}^d_{p,q}$ and $\mathbf{S}^d_{p,q}$ are the quotient of $\mathbb{H}^d_{p,q}$ and $\mathbb{S}^d_{p,q}$ by the antipodal map.

Example 2.1.

- $\mathbf{H}^d := \mathbf{H}^d_{d,1}$ is the projective model of the hyperbolic space;
- $\mathbb{HP}^d := \mathbb{H}^d_{d-1,1}$ is the half-pipe space, see [Dan13];
- $\mathrm{Ad}\mathbb{S}^d := \mathbb{H}^d_{d-1,2}$ is the Anti-de Sitter space;
- $\mathbb{S}^d := \mathbb{S}^d_{d+1,0}$ is the spherical space;
- $\mathrm{d}\mathbb{S}^d := \mathbb{S}^d_{d,1}$ is the de Sitter space.

We define $\mathbb{HS}^d_{p,q} = \{\mathbf{x} \in \mathbb{R}^{d+1}_{p,q} \mid |\langle \mathbf{x}, \mathbf{x}\rangle| = 1\}$. It is equipped with a complex-valued "distance", which restricts to each connected component as the natural constant curvature metric, and can be defined in terms of the Hilbert metric of the boundary quadric, see [Sch98]. If p and q are both non-zero, $\mathbb{HS}^d_{p,q}$ consists of a copy of $\mathbb{H}^d_{p,q}$ and a copy of $\mathbb{S}^d_{p,q}$ identified along their ideal boundaries.

Example 2.2.

- $\mathbb{HS}^d_{d,1}$ consists of two copies of the hyperbolic space $\mathbf{H}^d$ and a copy of the de Sitter space $\mathrm{d}\mathbb{S}^d$; we call it the "hyperbolic-de Sitter space", and simplify the notation to $\mathbb{HS}^d$.
- Another situation that concerns us in the future is $\mathbb{HS}^d_{p,p}$, $2p < d$ consisting of two copies of $\mathbb{H}^d_{p,p}$ differing by the sign of the metric. We denote it by $2\mathbb{H}^d_{p,p}$. In particular, $\mathbb{HS}^3_{2,2} = 2\mathrm{Ad}\mathbb{S}^3$.
- Up to a sign of metric, there are five possible $\mathbb{HS}^2_{p,q}$ metrics, namely $\mathbb{HS}^2_{3,0}(=\mathbb{S}^2)$, $\mathbb{HS}^2_{2,1}(=\mathbb{HS}^2)$, $\mathbb{HS}^2_{2,0}$, $\mathbb{HS}^2_{1,1}(=2\mathbb{HP}^2)$ and $\mathbb{HS}^2_{1,0}$.
- Up to a sign of metric, there are three possible $\mathbb{HS}^1_{p,q}$ metrics. We call a 1-subspace *space-*, *light-* or *time-*like if it is isometric to $\mathbb{HS}^1_{1,1}$, $\mathbb{HS}^1_{1,0}$ or $\mathbb{HS}^1_{2,0}$ respectively.

In an affine chart of $\mathbb{R}\mathrm{P}^d_{p,q}$, the boundary $\partial\mathbf{HS}^d_{p,q} := \partial\mathbf{H}^d_{p,q} = \partial\mathbf{S}^d_{p,q}$ appears as a quadric in $\mathbb{R}^d$.

Example 2.3.

- In the affine chart $x_3 = 1$: $\mathbf{H}^3$ appears as a unit open ball; $\mathbf{HP}^3$ appears as the interior of a circular cylinder; Ad$\mathbf{S}^3$ appears as the simply connected side of a one-sheeted hyperboloid.
- In the affine chart $x_2 = 1$: $\mathbf{H}^3$ appears as the two components of the complement of a two-sheeted hyperboloid that do not share a boundary; $\mathbf{HP}^3$ appears as two circular cones; Ad$\mathbf{S}^3$ appears as the non-simply connected side of a one-sheeted hyperboloid.

A totally geodesic subspace in $\mathbf{HS}^d_{p,q}$ is given by a projective subspace $L \subset \mathbb{R}\mathrm{P}^d$. If L is of codimension k, then the induced metric on L is isometric to $\mathbf{HS}^d_{p',q'}$ for some $p'+q' \leq d-k$ and, by Cauchy's interlacing theorem, we have $0 \leq p-p' \leq k$ and $0 \leq q-q' \leq k$. If $\mathbf{HS}^d_{p,q}$ is non-degenerate, then there are three possible metrics on a totally geodesic hyperplane H (codimension 1). We say that H is *space-*, *time-* or *light-like* if it is isometric to $\mathbf{HS}^{d-1}_{p,q-1}$, $\mathbf{HS}^{d-1}_{p-1,q}$ or $\mathbf{HS}^{d-1}_{p-1,q-1}$, respectively.

Example 2.4. In $\mathbf{HS}^d$, a hyperplane H is space-like if it is disjoint from the closure of $\mathbf{H}^d$, time-like if it intersects $\mathbf{H}^d$, or light-like if it is tangent to the boundary of $\mathbf{H}^d$.

The polar of a set $X \subset \mathbb{R}\mathrm{P}^d$ is defined by

$$X^* = \{[\mathbf{x}] : \langle \mathbf{x}, \mathbf{y}\rangle \leq 0 \text{ for all } [\mathbf{y}] \in X\}.$$

The polar of a subspace $L \subset \mathbb{R}\mathrm{P}^d$ is its orthogonal companion, i.e.

$$L^* = L^\perp = \{[\mathbf{x}] : \langle \mathbf{x}, \mathbf{y}\rangle = 0 \text{ for all } [\mathbf{y}] \in L\}.$$

If $\mathbf{HS}^d_{p,q}$ is non-degenerate, L is isometric to $\mathbf{HS}^k_{r,s}$ and $L^\perp$ to $\mathbf{HS}^{k'}_{r',s'}$, then we have $k+k' = d-1$ and $k-r-s = k'-r'-s' = p-r-r' = q-s-s'$. In particular, the polar of a hyperplane H is a point in $\mathbf{H}^d_{p,q}$ if H is space-like, in $\mathbf{S}^d_{q,p}$ if H is time-like, or on $\partial\mathbf{HS}^d_{p,q}$ if H is light-like.

2.2. **Ideal polytopes.** A set $X \subset \mathbb{R}\mathrm{P}^d$ is *convex* if it is convex in some affine chart that contains it. Equivalently [dGdV58], $X \subset \mathbb{R}\mathrm{P}^d$ is convex if for any two points $p, q \in X$, exactly one of the two segments joining p and q is contained in X.

Example 2.5.

- $\mathbf{H}^d$ is convex;
- Ad$\mathbf{S}^d$ is not convex;
- $\mathbf{HP}^d$ is convex, but its closure is not.

Two convex sets are *consistent* if some affine chart contains both of them, or *inconsistent* otherwise.

A *convex hull* of a set X is a minimal convex set containing X. Note that there is usually more than one convex hull. A *convex polytope* P is a convex hull of finitely many points. A (closed) *face* of P is the intersection of P with a *supporting hyperplane*, i.e. a hyperplane that intersects the boundary of P but disjoint from the interior of P. The faces of P decompose the boundary ∂P into a cell complex, giving a face lattice. Two polytopes are *combinatorially equivalent* if they have the same face lattice. The polar P^* is combinatorially dual to P, i.e. the face lattice of P^* is obtained from P by reversing the inclusion relations. We recommend the books [Grü03, Zie95] as general references for polytope theory.

Definition 2.6. A convex polytope $P \subset \mathbb{R}\mathrm{P}^d$ is *ideal* to $\mathbf{HS}^d_{p,q}$ if all its vertices are on the boundary of $\mathbf{HS}^d_{p,q}$. An ideal polytope P is *strongly ideal* if the interior of P is disjoint from $\partial\mathbf{HS}^d_{p,q}$, or *weakly ideal* otherwise.

In the case that P is (strongly) ideal to $\mathbf{HS}^d_{p,q}$, we also say that P is (strongly) ideal to $\mathbf{H}^d_{p,q}$ or to $\mathbf{S}^d_{p,q}$.

A *(polyhedral) HS structure* of a k-dimensional manifold is a triangulation the manifold together with an isometric embedding of each k-simplex into $\mathbf{HS}^k_{p,q}$, such that the simplices are isometrically identified on their common faces. If $\mathbb{R}\mathrm{P}^d$ is equipped with a $\mathbf{HS}^d_{p,q}$ metric, then a convex polytope

$P \subset \mathbb{R}\mathrm{P}^d$ with n vertices naturally induces an HS structure on the n-times punctured $\mathbb{S}^{d-1}$. If P is ideal to $\mathbf{HS}^d_{p,q}$, then this metric is geodesically complete.

In any affine chart, $\partial\mathbf{HS}^d_{p,q}$ appears as a quadratic surface, and an ideal polytope appears inscribed to this surface.

For $\mathbf{H}^d$, an ideal polytope P is strongly ideal if and only if it is consistent with $\mathbf{H}^d$ [CP17]. Polyhedra strongly ideal to $\mathbf{H}^3$ are then inscribed to a sphere. Their combinatorics was characterized by Hodgson, Rivin and Smith [HRS92]. Polyhedra strongly ideal to $\mathbf{HP}^3$ are inscribed to a circular cylinder. Polyhedra strongly ideal to $\mathrm{Ad}\mathbf{S}^3$ are inscribed to and contained in a one-sheeted hyperboloid. Danciger, Maloni and the second author [DMS14] have *essentially* provided characterizations of the combinatoric types of these polytopes.

We will focus on *weakly ideal* polyhedra, i.e. ideal polyhedra that are not strongly ideal. We prefer affine charts that contains the polytope P; such an affine charts cannot contain $\mathbf{H}^3$, $\mathbf{HP}^3$ or $\mathrm{Ad}\mathbf{S}^3$ by the discussion above. Polyhedra weakly ideal to $\mathbf{H}^3$ are then inscribed to a two-sheeted hyperboloid. Polyhedra weakly ideal to $\mathbf{HP}^3$ are inscribed to a circular cone. And finally, polyhedra weakly ideal to $\mathrm{Ad}\mathbf{S}^3$ are inscribed to, but not contained in, a one-sheeted hyperboloid. This covers all the quadratic surfaces, and characterizing the weakly ideal polyhedra in $\mathbf{H}^3, \mathbf{HP}^3$ and $\mathrm{Ad}\mathbf{S}^3$ would provides a complete answer to Steiner's problem.

3. Overview

From now on, we will focus on projective polyhedra weakly inscribed to the sphere, which is equivalent to projective polyhedra weakly ideal to $\mathbf{H}^3$, or Euclidean polyhedra inscribed to the two-sheeted hyperboloid.

Recall that a polyhedron P weakly ideal to $\mathbf{H}^3$ is not consistent with $\mathbf{H}^3$. Since we prefer affine charts containing P, $\mathbf{H}^3$ would appear, up to a projective transformation, as the set $x_0^2+x_1^2-x_2^2<-1$ in such charts. This is projectively equivalent to the Klein model. We use $\mathbf{H}^3_+$ and $\mathbf{H}^3_-$ to denote the parts of $\mathbf{H}^3$ with $x_2>0$ and $x_2<0$, respectively. Moreover, the boundary $\partial\mathbf{H}^3$ appears as a two-sheeted hyperboloid.

3.1. **Ideal polyhedra.** For a polyhedron P weakly ideal to $\mathbf{H}^3$, let V denotes the set of its vertices; then $V\subset\partial\mathbf{H}^3$ by definition. We write $V^+=V\cap\partial\mathbf{H}^3_+$ and $V^-=V\cap\partial\mathbf{H}^3_-$, and say that P is (p,q)-ideal if $|V^+|=p$ and $|V^-|=q$. P is strongly ideal if $p=0$ or $q=0$; we only consider weakly ideal polyhedra, hence $p>0$ and $q>0$. Following the curves $P\cap\partial\mathbf{H}^3$, we label vertices of V^+ by $1^+,\dots,p^+$, and vertices of V^- by $1^-,\dots,q^-$, in the order compatible with the right-hand rule.

Let $\mathscr{P}_n$ denote the space of labeled polyhedra with $n\geq 4$ vertices that are weakly ideal to $\mathbf{H}^3$, considered up to hyperbolic isometries, and $\mathscr{P}_{p,q}$ denote the space of labeled (p,q)-ideal polyhedra, $p+q\geq 4$. Then $\mathscr{P}_n$ is the disjoint union of $\mathscr{P}_{p,q}$ with $p+q=n$. We only need to study connected components $\mathscr{P}_{p,q}$, and may assume $p\leq q$ without loss of generality. We usually distinguish two cases, namely $p<2<q$ and $2\leq p\leq q$. Note that we always assume that $p\geq 1$ since we only consider only *weakly* ideal polyhedra, so $p<2$ below always means $p=1$.

3.2. **Admissible graphs.** We define a *weighted graph* (or simply *graph*) on a set of vertices V as a real valued function w defined on the unordered pairs $\binom{|V|}{2}$. The *weight* w_v at a vertex $v\in V$ is defined as the sum $\sum_u\theta(u,v)$ over all $u\neq v$.

Unless stated otherwise, the support of w is understood as the set E of edges. We can treat w as a usual graph with edge weight θ, and talk about notions such as subgraph, planarity and connectedness. But we will also take the liberty to include edges of zero weight, as long as it does not destroy the property in the center of our interest. For example, graphs in this paper are used to describe the 1-skeleta of polyhedra, i.e. 3-connected planar graphs. Hence whenever convenient, we will consider maximal planar triangulations. If this is not the case with the support of θ, we just triangulate the non-triangle faces by including edges of zero weight.

The advantage of this unconventional definition is that graphs can be treated as vectors in $\mathbb{R}^{\binom{|V|}{2}}$. Weighted graphs of a fixed combinatorics, together with their subgraphs, then form a linear subspace. Graphs with a common subgraph correspond to subspaces with nontrivial intersection. This makes it convenient to talk about neighborhood, convergence, etc. For a fixed polyhedral

combinatorics, our main result implies that the set of weighted graphs form a $(|E|-|V|)$-dimensional cell. Weighted graphs of a fixed number of vertices then form a cell complex of dimension $2|V|-6$ in $\mathbb{R}^{\binom{|V|}{2}}$: The maximal cells correspond to triangulated (maximal) planar graphs, and they are glued along their faces corresponding to common subgraphs.

Consider an edge e of an ideal polyhedron P. Then e is either a geodesic in $\mathbf{H}^3$, or a time-like geodesic in $\mathrm{d}\mathbf{S}^3$. In both cases, the faces bounded by e expand to half-planes forming a hyperbolic exterior dihedral angle, denoted by ϑ. We assign to e the *HS exterior dihedral angle* θ, which equals ϑ if $e \subset \mathbf{H}^3$, or $-\vartheta$ if $e \subset \mathrm{d}\mathbf{S}^3$. We will refer to θ as exterior angles, dihedral angles, or simply angles, and should not cause any confusion.

This angle assignment induces a graph on V, also denoted by θ, supported by the edges of P. We have thus obtained a function Θ that maps an ideal polyhedron P to the graph θ of its angles. Obviously, $\Theta(P)$ is polyhedral, i.e. 3-connected planar. We will see that, if P is (p,q)-ideal, then

(C1) $\theta = \Theta(P)$ admits a vertex-disjoint cycle cover consisting of a p-cycle and a q-cycle

and, if we color the edges connecting vertices on the same cycle by red, and those connecting vertices from different cycles by blue, then

(A1) $0 < \theta < \pi$ on red edges, and $-\pi < \theta < 0$ on blue edges;

(A2) $\theta_v = \sum_u \theta(u,v) = 0$, with the exception when v is the only vertex in a 1-cycle, in which case $\theta_v = -2\pi$;

(A3) The sum of θ over blue edges is $\leq -2\pi$, and the equality only happens when $p < 2 < q$.

The exception in Condition **(A2)** only happens when $p < 2 < q$.

Definition 3.1. A (p,q)-*admissible* graph is a weighted polyhedral graph satisfying Conditions **(C1)** and **(A1)**–**(A3)**.

Given a (p,q)-admissible graph drawn on the plane, we may label the vertices on the p-cycle by $1^+, \dots, p^+$, and vertices on the q-cycle by $1^-, \dots, q^-$, both in the clockwise order. Let $\mathscr{A}_{p,q}$ denote the space of labeled (p,q)-admissible graphs with $p+q \geq 4$. We use $\mathscr{A}_n$, $n \geq 4$, to denote the disjoint union of $\mathscr{A}_{p,q}$ with $p+q=n$. Our main Theorem 1.5 is the consequence of the following theorem:

Theorem 3.2. *Θ is a homeomorphism from $\mathscr{P}_{p,q}$ to $\mathscr{A}_{p,q}$.*

Note that Condition **(A3)** and some inequalities in Condition **(A1)** are not present in Theorem 1.5. We will see that they are indeed redundant when formulating a feasibility problem.

3.3. **Admissible HS structures.** Let Δ denote the function that maps an ideal polyhedron to its induced HS structure. If P is (p,q)-ideal, it follows from the definition that $\Delta(P)$ is geodesically complete, and it is maximal in the sense that it does not embed isometrically as a proper subset of another HS structure.

The part of ∂P in $\mathbf{H}^3$ has no interior vertex, hence is isometric to a disjoint union of copies of $\mathbf{H}^2$. In the case $2 \leq p \leq q$, we use $\mathbf{H}^2_+$ to denote the copies induced by $\partial P \cap \mathbf{H}^3_+$. If $p < 2 < q$, we have only $\mathbf{H}^2_- = \partial P \cap \mathbf{H}^3_-$. The part of ∂P in $\mathrm{d}\mathbf{S}^3$ has no interior vertex, neither, hence $\partial P \cap \mathrm{d}\mathbf{S}^3$ is isometric to a complete de Sitter surface.

The intersection of ∂P with a space-like plane in $\mathrm{d}\mathbf{S}^3$ is a simple polygonal closed space-like curve in $\partial P \cap \mathrm{d}\mathbf{S}^3$. If $2 \leq p \leq q$, this polygonal curve can be deformed to one of maximal length, say γ_0, which is therefore geodesic in $\partial P \cap \mathrm{d}\mathbf{S}^3$. Considered as a polygonal curve in $\mathrm{d}\mathbf{S}^3$, γ_0 is then E-convex in the sense of [Sch98, Def 7.13], and it follows that its length ℓ is less than 2π, see [Sch98, Prop 7.14]. As a consequence, γ_0 is the unique simple closed space-like geodesic in $\partial P \cap \mathrm{d}\mathbf{S}^3$, because any other simple closed space-like geodesic would need to cross γ_0 at least twice (there is no de Sitter annulus with space-like, geodesic boundary by the Gauss-Bonnet formula), and two successive intersection points would be separated by a distance π, leading to a contradiction. We denote the metric space $\partial P \cap \mathrm{d}\mathbf{S}^3$ by $\mathrm{d}\mathbb{S}^2_\ell$. $\mathrm{d}\mathbb{S}^2_\ell$ has two boundary components, both homeomorphic to a circle.

If $p < 2 < q$, then one boundary component of S degenerates to a point. In this case, the metric space $\partial P \cap \mathrm{d}\mathbf{S}^3$ does not contain any closed space-like geodesic, and we denote it by $\mathrm{d}\mathbb{S}^2_0$.

Hence $\Delta(P)$ is obtained by gluing one or two copies of $\mathbf{H}^2$ to the non-degenerate boundary components of a de Sitter surface. Let $\gamma_\pm$ be the map that glues $\partial\mathbf{H}^2_\pm$ to $\partial\mathrm{d}\mathbb{S}^2_\ell$. We will see that $\gamma_\pm$ are C^1 piecewise projective maps (CPP maps for short). More specifically, they are projective except at the vertices of P. The points where the map is not projective are called *break points*. A break point is said to be positive (resp. negative) if the jump in the second derivative at this point is positive (resp. negative). We will see that the break points of $\gamma_\pm$ are all positive.

Definition 3.3. A (p,q)-*admissible* HS structure, $p+q\geq 4$, is obtained

In the case $p<2<q$**:** by gluing a copy of $\mathbf{H}^2$ to $\mathrm{d}\mathbb{S}^2_0$ along the non-degenerate ideal boundary by a CPP map with q positive break points.

In the case $2\leq p\leq q$**:** by gluing two copies of $\mathbf{H}^2$ to $\mathrm{d}\mathbb{S}^2_\ell$, $0<\ell<2\pi$, along the ideal boundaries through CPP maps with, respectively, p and q positive break points.

In Figure 1 we sketch the situation of $p=2$ and $q=3$ is sketched in Figure 1.

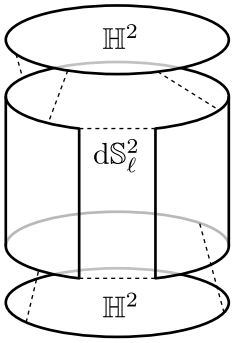

FIGURE 1. A sketch of a $(2,3)$-admissible HS structure. The dashed segments indicates the gluing maps, including one that produces $\mathrm{d}\mathbb{S}^2_\ell$, and two CPP maps with two and three break points, respectively.

Given a (p,q)-admissible HS structure, we may label the break points in the two boundary components of S by $1^+,\dots,p^+$ and $1^-,\dots,q^-$, respectively. Let $\mathscr{M}_{p,q}$ denote the space of (p,q)-ideal HS structures up to isometries. Our main Theorem 1.6 is the consequence of the following theorem.

Theorem 3.4. *Δ is a homeomorphism from $\mathscr{P}_{p,q}$ to $\mathscr{M}_{p,q}$.*

3.4. **Outline of proofs.** We will prove that Θ and Δ are local immersions (Section 6) with images in $\mathscr{A}_{p,q}$ (Section 4.2) and $\mathscr{M}_{p,q}$ (Section 4.3), respectively. They are then local homeomorphisms because $\mathscr{P}_{p,q}$, $\mathscr{A}_{p,q}$ and $\mathscr{M}_{p,q}$ have the same dimension $2(p+q-3)$ (Section 7). Moreover, they are proper maps (Section 5), hence are covering maps. A difference from the previous works lies in the fact that $\mathscr{P}_{p,q}$, $\mathscr{A}_{p,q}$ and $\mathscr{M}_{p,q}$ are not simply connected if $2\leq p\leq q$. We will use open covers and universal covers to conclude that the covering numbers of Θ and Δ are one (Section 7).

4. Necessity

4.1. **Combinatorial conditions.** We first verify combinatorial Condition **(C1)**. For this we will need some lemmata about convex sets in $\mathbb{R}\mathrm{P}^d$.

Lemma 4.1. *Let A and B be two convex sets in $\mathbb{R}\mathrm{P}^d$. If A and B are consistent, then $A\cap B$ consists of at most one connected component. If A and B are inconsistent, then $A\cap B$ consists of exactly two connected components.*

Proof. If A and B are consistent, we can regard them as convex sets in Euclidean space. Hence they are either disjoint, or their intersection is convex, hence connected.

Conversely, if $A\cap B$ consists of at most one connected component, they can be lifted to two convex sets A' and B' of $\mathbb{S}^d$. We may assume that $A'\cap B'=\emptyset$. If this is not the case, it suffices

to replace B' by its antipodal image. By the spherical hyperplane separation theorem, A' and B' is separated by a spherical hyperplane. We then project back to $\mathbb{R}\mathrm{P}^d$, taking this separating hyperplane at infinity. This gives us an affine chart that contains both A and B.

Consequently, if A and B are inconsistent, $A \cap B$ will have at least two connected components. To prove that there are exactly two components, we only need to work in dimension $d = 2$. Higher dimensional cases follow by restricting to a 2-dimensional subspace.

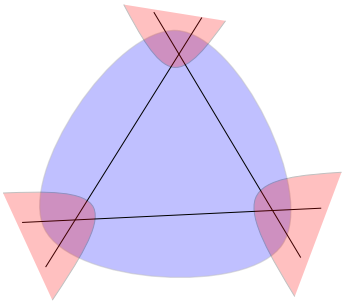

FIGURE 2.

For the sake of contradiction, assume that $A \cap B$ has three components, and take a point from each of them. The three points determine six segments. In an affine chart that contains A, the three bounded segments are contained in A, the three unbounded segments are contained in B. The situation is illustrated in Figure 2, where A is blue and B is red. Since B is convex, there should be a projective line avoiding B. To avoid finite points of B in the chosen affine chart, such a line must be parallel to one of the three lines spanned by the points. But this means that this line intersects B at infinity. Hence such a projective line does not exist, contradicting the convexity of B. □

If A and B are two inconsistent convex regions in $\mathbb{R}\mathrm{P}^2$, then $\partial A \cap \partial B$ consists of at most four connected components, at most two on the boundary of each connected component of $A \cap B$. Otherwise, either the interior or the closure of $A \cap B$ would consist of more than two connected components, contradicting the lemma above.

A particular case is when $A = \mathbf{H}^2$ and B is a polygon, and their boundary intersect at the vertices of B, i.e. B is weakly ideal to $\mathbf{HS}^2$. In this case, Lemma 4.1 implies

Corollary 4.2.

- *Any polygon strongly ideal to* $\mathbf{HS}^2$ *with at least three vertices is disjoint from* $\mathrm{d}\mathbf{S}^2$.
- *Any polygon ideal to* $\mathbf{HS}^2$ *with at least five vertices is strongly ideal in* $\mathbf{HS}^2$.
- *A weakly ideal polygon P has three or four vertices, at most two in each connected component of* $\partial P \cap \mathbf{H}^2$.

A dual version of this corollary was proved in [CP17]. Notice that there is only one possibility for a weakly ideal triangle; see Figure 3.

Moreover, it is known that every connected component of $A \cap B$ is convex [Tod10].

Lemma 4.3. *Let A and B be two inconsistent convex sets in $\mathbb{R}\mathrm{P}^d$, and C_1, C_2 be the connected components of $A \cap B$. Then $\partial C_i \cap \partial A$ and $\partial C_i \cap \partial B$, $i = 1, 2$, are all contractible.*

Proof. We only need to argue for $\partial C_2 \cap \partial B$. The other cases follow similarly.

We work in an affine chart containing A; thus it does not contain B. Let $p \in \partial C_1 \cap \partial B$. Then for any $q \in \partial C_2 \cap \partial B$, the bounded closed segment $[pq]$ is disjoint from the interior of B, hence also from the interior of C_2. On the other hand, any $q' \in \partial C_2 \setminus \partial B$ (if not empty) is in the interior of B, hence $[pq']$ must intersect the interior of C_2. In other words, $\partial C_2 \cap \partial B$ is the part of ∂C_2 "visible" from p, which must be contractible as C_2 is.

The proof for $d = 2$ is illustrated in Figure 4. □

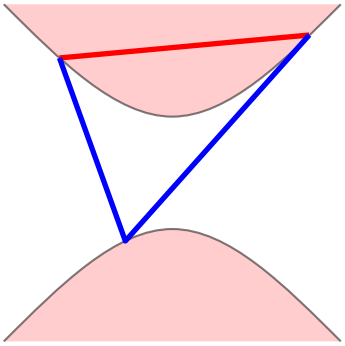

FIGURE 3.

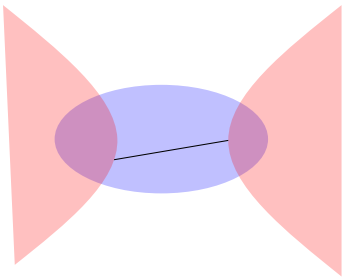

FIGURE 4.

We call an (open) face (that is vertex, edge or facet) F of P *interior* if $F \subset \overline{\mathbf{H}^3}$, or *exterior* if $F \subset \mathrm{d}\mathbb{S}^3$. For example, every vertex of an ideal polyhedron is interior, and every face of a strongly ideal polyhedron is interior. An edge of an ideal polyhedron is either interior or exterior. Let $\mathcal{I}(P)$ be the union of interior faces, and $\mathcal{E}(P)$ be the union of exterior faces.

Proposition 4.4 (Condition **(C1)** and more)**.** *Let P be a polyhedron weakly ideal to $\mathbf{H}^3$. Then $\mathcal{I}(P)$ consists of two connected components, both contractible. A component is homeomorphic to a closed disk if it contains at least three vertices. Vertices in each component induce an outerplanar graph. Moreover, $\mathcal{E}(P)$ consists of disjoint open segments; there is no exterior facet.*

Proof. Assume a triangular facet F that is not interior, then it would be a weakly ideal in $\mathrm{span}(F)$. Hence P has no exterior facet. The only exterior faces are edges. Then we observe from Figure 3 that, whenever F is not interior, $F \cap \partial\mathbf{H}^3$ is an arc which is homeomorphic to the unique interior edge of F. This homeomorphism induces an homotopy from $\partial P \cap \overline{\mathbf{H}^3}$ to $\mathcal{I}(P)$. The former is contractible by Lemma 4.3, hence so is $\mathcal{I}(P)$.

The vertices of P are all on the boundaries of $\partial P \cap \overline{\mathbf{H}^3}$ (and of $\mathcal{I}(P)$). Otherwise, through a vertex in the interior of $\partial P \cap \overline{\mathbf{H}^3}$, we can find an hyperplane H such that the intersection of $H \cap \partial P$ and $H \cap \partial\mathbf{H}^3$ consists of more than two connected components, contracting Lemma 4.1. Hence the vertices of each component induce an outerplanar graph. If there are at least three vertices in this component, the boundary edges form a Hamiltonian cycle (of the induced graph), so the outerplanar graph is 2-connected. We then conclude that the component of $\mathcal{I}(P)$ is homeomorphic to a disk. □

This proves the necessity of Condition **(C1)** since the boundary edges of the 2-connected outerplanar graphs form a cycle cover consisting of two cycles.

Induction in higher dimensions yields that the union of interior faces consists of two contractible components, and vertices and edges in each component form a 2-connected graph. However, it is possible that a component is not homeomorphic to the $(d-1)$-ball, as shown in the following example.

Example 4.5. Consider the eight points $[\pm\sqrt{2}, \pm 1, 0, 0]$, $[2(1-\epsilon+\epsilon^2), 0, \pm(1-\epsilon), \pm\epsilon]$, taking all possible combinations of $\pm$. Their convex hull is a 4-dimensional polytope weakly inscribed to the two-sheeted hyperboloid defined by the equation $-x_0^2 + x_1^2 + x_2^2 + x_3^2 = -1$. If ϵ is sufficiently small, the polytope has two interior facets in $\mathbf{H}^4_+$, whose intersection is a single edge connecting $[\sqrt{2}, \pm 1, 0, 0]$. Hence the union of interior faces is not homeomorphic to a 3-ball.

4.2. **Angle conditions.** We now prove that the conditions in Theorem 3.2 are necessary.

Proposition 4.6.

$$\Theta(\mathscr{P}_{p,q}) \subseteq \mathscr{A}_{p,q}.$$

For that, we need to verify all the conditions defining an admissible graph.

We color interior edges by red, and exterior edges by blue, and prove the angle conditions with respect to this coloring. In the previous part we have seen that this coloring coincides with the combinatorial description in Theorem 3.2.

Among the conditions that involve angles, Condition **(A1)** comes from the definition of angle. To see Condition **(A2)** we need the vertex figures.

Recall that a *horosphere* in $\mathbf{H}^d$ based at an ideal point $x \in \partial\mathbf{H}^d$ is a hypersurface that intersects orthogonally all the geodesics emerging from x. In the projective model of $\mathbf{H}^3$, horospheres appear as flattened spheres tangent to the hyperbolic boundary. Similarly, a horosphere in $\mathrm{d}\mathbf{S}^d$ based at $x \in \partial\mathrm{d}\mathbf{S}^d$ is a hypersurface that intersects orthogonally all the geodesics emerging from x. Horospheres in $\mathbf{H}^d$ and $\mathrm{d}\mathbf{S}^d$ are paired through polarity: the set in $\mathrm{d}\mathbf{S}^d$ polar to a horosphere in $\mathbf{H}^d$ is a horosphere with the same base point, and vice versa. See Figure 5. In the following, by a horosphere in $\mathbf{HS}^d$, we mean a horosphere in $\mathbf{H}^d$ or in $\mathrm{d}\mathbf{S}^d$.

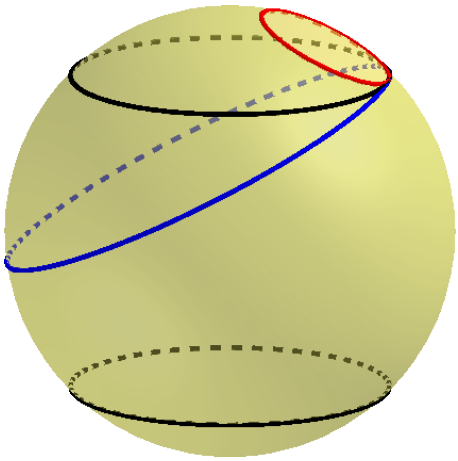

FIGURE 5. Horocycles in $\mathbb{H}^2$ (red) and in $\mathrm{d}\mathbb{S}^2$ (blue). The black circles are the boundaries of $\mathbf{HS}^2$.

Now consider an ideal polyhedron P. The *vertex figure* of P at a vertex v, denoted by P/v, is the projection of P with respect to v. P/v is therefore a polygon in $\mathbb{R}\mathrm{P}^2$. A chart is provided by a horosphere in $\mathbf{HS}^3$ based at v. This chart sends $v^\perp$ to the line at infinity, hence does not contain P/v unless the neighborhood of v in ∂P is contained in $\mathbf{H}^3$ or in $\mathrm{d}\mathbf{S}^3$.

Figure 6 shows a typical situation of P/v not contained in the horosphere. The vertices of the polygon P/v correspond to the edges of P adjacent to v. If one walks along the polygon, it can be arranged (as in Figure 6) that he turns anti-clockwise at the vertices corresponding to red edges, and clockwise at the vertices corresponding to blue edges. Indeed, the turning direction switches when the walker passes through infinity. The turning angles are then the dihedral angles θ at the corresponding edges, taking anti-clockwise turns as positive, and clockwise turns as negative. Condition **(A2)** then follows immediately.

We now verify the last angle condition.

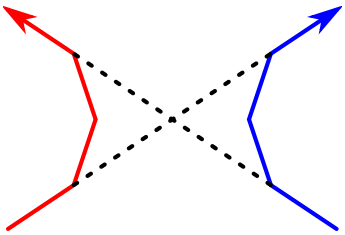

FIGURE 6. A typical vertex figure.

Proposition 4.7 (Condition **(A3)**). *If P is weakly ideal, then $\Theta(P)$ sum up to at least -2π over the blue edges, and -2π is achieved only when $p < 2 < q$.*

Proof. It is clear from the vertex figure that the sum over blue edges is -2π when $p = 1$. Hence we focus on the case $2 \leq p \leq q$ and prove that the sum over blue edges is strictly larger than -2π.

If $2 \leq p \leq q$, the polar P^* of P is a compact polyhedron in $\mathrm{d}\mathbf{S}^3$. All its faces are light-like (isotropic). The edges of P^* polar to the blue edges of P form a closed space-like polygonal curve γ, whose vertices are polar to the non-interior facets of P.

The polygonal curve γ is a $\mathcal{T}$-geodesic for the induced HS-structure on P^*; see [Sch01, Definition 3.4]. It then follows from point C. in [Sch01, Theorem 1.5] that γ has length strictly less than 2π. The proposition follows by polarity. □

4.3. **Metric conditions.** We show in this part that the conditions of Theorem 1.6 are necessary.

Theorem 4.8.

$$\Delta(\mathscr{P}_{p,q}) \subseteq \mathscr{M}_{p,q}.$$

If P is (p, q)-ideal, we have argued that $\Delta(P)$ consists of one or two copies of $\mathbf{H}^2$ and a de Sitter surface. It remains to verify that the pieces are glued along their ideal boundaries by CPP (C^1 piecewise projective) maps with positive break points at the vertices of P. We may focus on the boundary component of $\mathrm{d}\mathbb{S}^2_\ell$, $\ell \geq 0$, consisting of $q > 1$ vertices of P. Let γ be the map that sends $\partial\mathbf{H}^2$ to this boundary.

We can identify $\partial\mathbf{H}^2$ to the real projective line $\mathbb{R}\mathrm{P}^1$. Let $v_0 < v_1 < \cdots < v_p = v_0$, in this order, be the p vertices of P. They divide $\mathbb{R}\mathrm{P}^1$ into p segments $[v_i, v_{i+1}]$. Each segment corresponds to a segment of $\partial\mathbf{HS}^2$ in the interior of a face triangle weakly ideal to $\mathbf{HS}^2$. Hence for each i, $0 \leq i < p$, the restriction $\gamma_i = \gamma|_{[v_i, v_{i+1}]}$ coincides with the restriction of an element of $\mathrm{PSL}(2, \mathbb{R})$. This proves that γ is piecewise projective.

To study differentiabilities, we will follow the work of Martin [Mar05] on CPP homeomorphisms of $\mathbb{R}\mathrm{P}^1$. Note that $\partial\mathrm{d}\mathbb{S}^2_\ell$ is not projectively equivalent to $\mathbb{R}\mathrm{P}^1$. But our CPP maps are local homeomorphisms. Hence the definitions and many results from [Mar05] remain valid for our case.

Definition 4.9 ([Mar05]). For $x \in \mathbb{R}\mathrm{P}^1$, let $\gamma_x^{\leftarrow} \in \mathrm{PSL}(2, \mathbb{R})$ (resp. $\gamma_x^{\rightarrow} \in \mathrm{PSL}(2, \mathbb{R})$) be the left (resp. right) germ of a piece-wise projective map γ. The projective transformation $D_x\gamma = (\gamma_x^{\rightarrow})^{-1} \circ \gamma_x^{\leftarrow}$ is called the *shift* of γ at x.

The projective transformations $\gamma_x^{\leftarrow}$, $\gamma_x^{\rightarrow}$ and $D_x\gamma$ extend uniquely to isometries of $\mathbf{HS}^2$. We abuse the same notations for these extensions. The shift $D_x\gamma$ measures how much γ fails to be projective at x. It can be understood as the holonomy of the HS structure along a curve going around x. The following lemma reveals the relation between the shift and the differentiability of γ.

Lemma 4.10. *γ is C^1 at x if and only if $D_x\gamma$ preserves the horocycles based at x.*

Proof. The proof of [Mar05, Proposition 2.3] can be used here, word by word, to prove the "only if". For the "if" part, assume $x = 0$. Note that $u := D_0\gamma : \mathbb{R}\mathrm{P}^1 \to \mathbb{R}\mathrm{P}^1$ fixes 0. If (the extension of) u preserves the horocycles based at 0, it must be of the form $u(t) = t/(ct + 1)$. Hence $u'(0) = 1$, therefore γ is C^1. □

Then γ being C^1 follows from the following proposition.

Proposition 4.11. *For any $x \in \partial\mathbf{H}^2$, $D_x\gamma$ preserves the horocycles based at x.*

Proof. As a measure of how much γ fails to be projective at x, $D_x\gamma$ must be trivial except at vertices of of P. Let $v \in V^+$ be a vertex of P.

We first note that in a neighborhood of v, exactly two faces of ∂P, say $F^{\leftarrow}$ and $F^{\rightarrow}$, have non-empty intersections with $\partial\mathbf{H}^3$. This can be seen from the vertex figure P/v. Recall that vertices (resp. edges) of P/v correspond to edges (resp. faces) of P adjacent to v. In the affine chart of $\mathbb{R}\mathrm{P}^2$ provided by a horosphere at v, the tangent plane of $\partial\mathbf{H}^3$ at v is sent to infinity, and intersects P/v in exactly two edges, corresponding to $F^{\leftarrow}$ and $F^{\rightarrow}$. In particular, no vertex of P/v is at infinity; otherwise such a vertex would correspond to an edge e of P tangent to the quadratic $\partial\mathbf{H}^3$, and the other end of e can not lie on the same quadratic.

Our gluing map γ coincides with $\gamma_v^{\leftarrow}$ on $\partial\mathbf{HS}^3 \cap F^{\leftarrow}$, and with $\gamma_v^{\rightarrow}$ on $\partial\mathbf{HS}^3 \cap F^{\rightarrow}$.

$F^{\leftarrow}$ and $F^{\rightarrow}$ extend to half-planes bounding a dihedral angle Φ containing P. The boundary of Φ is isometric to $\mathbf{HS}^2$. Let $\tilde{h}$ be a horocycle in $\mathbf{HS}^3$ based at v. Then the intersection $h = \partial\Phi \cap \tilde{h}$ give a horocycle in $\mathbf{HS}^2$. On the other hand, $h' = \partial P \cap \tilde{h}$ give a horocycle in ∂P based at v. In a neighborhood of v, h' lies within $F^{\leftarrow}$ and $F^{\rightarrow}$.

We then conclude that, for any horocycle h in $\mathbf{HS}^2$, we have $h' = \gamma_x^{\leftarrow}(h) = \gamma_x^{\rightarrow}(h)$, i.e. $D_x\gamma(h) = h$. □

Proposition 4.11 can be interpreted as horocycles "closing up" after going around a vertex. In Rivin's characterization of polyhedra strongly inscribed in the sphere, the same phenomenon is reflected by the shearing coordinates summing up to 0 around each vertex. We can do the same with a proper definition of shearing coordinates.

Note that three vertices on $\partial\mathbf{HS}^2$ determine a unique strongly ideal (hyperbolic) triangle. Hence for a given ideal HS structure $\delta \in \mathscr{M}_{p,q}$, we can replace every triangle in δ, if not already strongly ideal, by the unique hyperbolic triangle with the same vertices. The result is a hyperbolic structure η (a triangulation with hyperbolic simplices) of the n-times punctured $\mathbb{S}^2$ (not embedded). We define the *HS shearing* (or simply shearing) along an edge of δ as the hyperbolic shearing (see [Pen87, Pen12]) along the corresponding edge of η.

Shearing can be easily read from the vertex figures. First note that the vertex figure of η at a vertex v can be obtained from that of δ by replacing the segments through infinity by the unique other segments with the same vertices. For example, Figure 7 is obtained from Figure 6. Let e be an edge of P adjacent to a vertex v. The shearing along e then equals to the logarithm of the length ratio of the segments adjacent to e in the vertex figure of η at v.

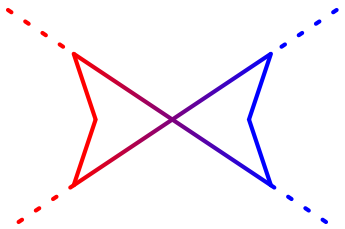

FIGURE 7. The vertex figure of the hyperbolic structure corresponding to the vertex figure of the HS structure shown in Figure 6.

Horocycles in η close up if and only if the hyperbolic shearings sum up to 0 over the edges adjacent to v. Then we see from the vertex figure that the horocycles in δ also "close up". And by definition, the HS shearings of η must also sum up to 0. Different triangulation of δ would yield a different hyperbolic metric η. But for a fixed triangulation, it is well-known (see [Pen87, Pen12]) that the hyperbolic shearing on the edges of η provide a coordinate system for the hyperbolic structure. Hence the HS shearing on the edges of δ provide a coordinate system for the ideal HS metrics.

Now back to the proof of necessity. Proposition 4.11 asserts that the $D_x\gamma$ are parabolic transformations for every $x \in \mathbb{R}\mathrm{P}^1$. Consider the projective transformation $u_x : t \mapsto 1/(t-x)$

sending x to infinity. Then the conjugate $u_x D_x(\gamma) u_x^{-1}$ is a translation of the form $t \mapsto t + d_x(\gamma)$. We have $d_x(\gamma) = 0$ at projective points. At break points:

Lemma 4.12. *$d_x(\gamma) < 0$ (resp. > 0) if x is a positive (resp. negative) break points.*

Proof. We may assume $x = 0$, then $u_0 = 1/t$. We already figured that $D_0(\gamma)$ is of the form $D_0(\gamma) : t \mapsto t/(ct+1)$. So the conjugate $u_0 D_0(\gamma) u_0^{-1}$ has the form $t \mapsto t + c$, i.e. $d_0(\gamma) = c$. On the other hand, an elementary computation shows that the second derivative $\frac{d^2}{dt^2}|_{t=0} D_0(\gamma) = -2c$. □

In the half-space model of $\mathbf{H}^2$, let h be a horocycle based at x. Then $d_x(\gamma) < 0$ (resp. > 0) if and only if $D_x\gamma$ moves points on h in the clockwise (resp. anti-clockwise) direction. We are now ready to prove that every break point of γ is positive.

Proposition 4.13. *$d_v(\gamma) < 0$ at every vertex v of P.*

Proof. We keep the definitions and notations in the proof of Proposition 4.11. We can recover P by truncating the dihedral angle Φ with planes through v. From the vertex figure, we observe the effect of a truncation on a horocycle h based at v: it replaces a segment of h with a shorter one. See Figure 8. Hence $D_v\gamma = (\gamma_v^{\rightarrow})^{-1} \circ \gamma_v^{\leftarrow}$ moves points on h in the clockwise direction, i.e. $d_v(\gamma) < 0$. □

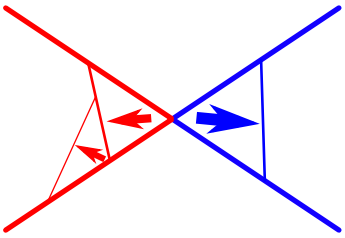

FIGURE 8. Each truncation replaces a horocyclic segment with a shorter one.

Remark 4.14. Lemma 6.5 of [Mar05] asserts that $d_x(\gamma)$ equals the change of length of a well chosen segment of horocycle based at x. Up to a scaling, this also suffices for us to conclude that $d_v(\gamma) < 0$.

5. PROPERNESS

The following theorem states that the maps Θ and Δ are proper.

Theorem 5.1. *Consider a sequence of polyhedra $(P_k)_{k\in\mathbb{N}}$ diverging in $\mathscr{P}_{p,q}$, then $\theta_k = \Theta(P_k)$ diverges in $\mathscr{A}_{p,q}$ and $\delta_k = \Delta(P_k)$ diverges in $\mathscr{M}_{p,q}$.*

Up to hyperbolic isometries, we may fix three vertices for every polyhedron in (P_k). As $\partial\mathbf{H}^3$ is compact, we may assume that vertices of P_k have well defined limits by taking a subsequence. But the limit of P_k, denoted by P_∞, is not a (p,q)-ideal polyhedron, since the sequence is diverging.

Hence in the limit, P_∞ must fail to be *strictly* convex at some vertex v. Let P'_∞ be the convex hull of all the other vertices. There are three possibilities, namely that v is in the relative interior of a vertex, an edge or a facet of P'_∞. But every straight line intersect a quadratic surface in at most two points, hence an ideal vertex can not be in the relative interior of an edge. Thus we only need to consider the remaining two possibilities.

Remark 5.2. For strongly ideal polyhedra, an ideal vertex can not lie in the interior of a facet. Hence there is only one possibility to consider. See [DMS14].

Proposition 5.3. *If some vertices of $(P_k)_{k\in\mathbb{N}}$ converge to the same vertex of P_∞, then the limit graph θ_∞ violates Condition* **(A2)** *or* **(A3)**, *and some break points in δ_k merge into a single break point in the limit metric δ_∞.*

Proof. The divergence of the induced metrics follows immediately from the correspondence between vertices of P and break points in $\Delta(P)$. Hence we will focus on the divergence of the admissible graphs.

Note that the ideal boundary of $\mathbf{H}^3$ can be seen as two copies of $\mathbf{H}^2$ identified along their ideal boundaries. With our choice of affine chart for the projective model, the two copies of $\mathbf{H}^2$ appear as the two sheets of the hyperboloid $x_0^2 + x_1^2 - x_2^2 = -1$.

The vertex set of each $P_k \in \mathscr{P}_{p,q}$ then corresponds to two point sets V_k^+ and V_k^- in $\mathbf{H}^2$ of cardinality p and q respectively. They converge to two point sets V_∞^+ and V_∞^- of cardinality p' and q' respectively, corresponding to the vertices of P_∞. If some vertices of (P_k) converges to the same vertex of P_∞, we must have $p' + q' < p + q$. Up to hyperbolic isometries, we may assume three fixed vertices shared by all P_k, hence $3 \leq p' + q'$.

Consequently, the graph $\Theta(P_\infty)$ has at least three vertices, but stricly less vertices than θ_∞. In fact, it is obtained by contracting vertices of θ_∞. Recall that $\theta_k = \Theta(P_k)$ and their limit θ_∞ are vectors of dimension $\binom{|V|}{2} = \binom{p+q}{2}$.

Assume that a set of vertices $S \subset V$ is merged into a vertex of P_∞. If $S \neq V^+$ or V^-, we have on the one hand

$$\sum_{u \notin S, v \in S} \theta_\infty(u, v) = 0.$$

On the other hand, as the limit of θ_k, Condition **(A2)** asserts that

$$\sum_{u \in V} \theta_\infty(u, v) = 0$$

for any $v \in S$. Comparing the two sums, we conclude that

$$\sum_{u, v \in S} \theta_\infty(u, v) = 0.$$

Now assume that θ_∞ is (p, q)-admissible. Since S are vertices on the same polar circle, $\theta_\infty(u, v)$ is non-negative, hence must be 0, if $u, v \in S$. In other words, S induce an empty graph, contradicting the fact that vertices in S are consecutive in a cycle.

If $S = V^+$ or V^-, we must have $p > 1$. But it is easy to conclude that the sum over negative weights in θ_∞ is -2π, contradicting Condition **(A3)**. □

Proposition 5.4. *If a vertex v of P_k converges to a vertex of P_∞ that is contained in a unique supporting plane, then v is an isolated vertex in the limit graph θ_∞, and v is not a break point in the HS structure induced by δ_∞.*

Proof. Under the assumption of the proposition, every face of P_∞ adjacent to v must lie in this unique supporting plane. Otherwise, the supporting plane of the face would be another supporting plane containing v. Then the dihedral angles vanish on all the edges incident to v. Since v is in the interior of a weakly ideal HS triangle, the gluing map is projective at v. □

A special case is of particular importance for us: If ∂P_k converges to a double cover of a plane, then the limit polyhedron P_∞ is equal to $\mathbb{R}\mathrm{P}^3$. In this case, every vertex is "flat": θ_∞ is identically 0 (empty graph), and δ_∞ is the double cover of $\mathbf{HS}^2$. We call this polyhedron a *flat polyhedron.*

Proposition 5.5. *If two faces of P_∞ intersect in their relative interiors and span a plane, then P_∞ is flat.*

Proof. Under the assumption of the proposition, the only plane that avoids the interior of P_∞ is the plane spanned by the two intersecting faces. Hence every face admits this plane as the supporting plane. In other words, every face lies in this plane. □

6. Rigidity

6.1. **The infinitesimal Pogorelov map.** We recall here the definition of the infinitesimal Pogorelov map, as well as its key properties. We refer to [Sch98] for the proofs, see in particular Définition 5.6 and Proposition 5.7 in [Sch98]. Other relevant references are [Fil07, Izm09, LS00, Sch05].

With affine charts containing weakly ideal polytopes, the hyperplane at infinity H_∞ is space-like. Apart from the HS metric and the usual Euclidean metric, the affine charts can also carry the Minkowski metric. Then the point $x_0 = H_\infty^\perp$ is the "center" of the Minkowski space $\mathbb{R}^{2,1}$. The set of light-like geodesics passing through x_0 is called the *light cone* at x_0, denoted by $C(x_0)$.

Let $U = \mathbb{R}\mathrm{P}^3 \setminus H_\infty$ be an affine chart, and $\iota : U \to \mathbb{R}^{2,1}$ be the projective embedding into the Minkowski space. The infinitesimal Pogorelov map Υ is then defined as the bundle map $\Upsilon : TU \to T\mathbb{R}^{2,1}$ over the inclusion $\iota : U \hookrightarrow \mathbb{R}^{2,1}$ as follows: Υ agrees with $d\iota$ on $T_{x_0}U$. For any $x \in U \setminus C(x_0)$, and any vector $v \in T_xU$, write $v = v_r + v_\perp$, where v_r is tangent to the radial geodesic passing through x_0 and x, and $v_\perp$ is orthogonal to this radial geodesic, and define

$$\Upsilon(v) = \sqrt{\frac{\|\hat{x}\|^2_{\mathbf{HS}}}{\|d\iota(\hat{x})\|^2_{2,1}}}\, d\iota(v_r) + d\iota(v_\perp),$$

where the norm $\|\cdot\|_{\mathbf{HS}}$ in the numerator of the first term is the HS metric, the norm $\|\cdot\|_{2,1}$ in the denominator is the Minkowski metric and $\hat{x}$ is the normalized radial vector (so $\|\hat{x}\|^2_{2,1} = \pm 1$).

The key property of the infinitesimal Pogorelov map is the following (the proof is an easy computation in coordinates, that can be adapted from [Fil11, Lemma 3.4]).

Lemma 6.1. *Let Z be a vector field on $U \setminus C(x_0) \subset \mathbf{HS}^3$. Then Z is a Killing vector field if and only if $\Upsilon(Z)$ (wherever defined) is a Killing vector field for the Minkowski metric on $\mathbb{R}^{2,1}$.*

In fact, the lemma implies that the bundle map Υ, which so far has only been defined over $U \setminus C(x_0)$, has a continuous extension to all of U. The bundle map Υ is called an *infinitesimal Pogorelov map*.

Next, the bundle map $\Xi : T\mathbb{R}^{2,1} \to T\mathbb{R}^3$ over the identity, which simply changes the sign of the n-th coordinate of a given tangent vector, has the same property: it sends Killing vector fields in $\mathbb{R}^{2,1}$ to Killing vector fields for the Euclidean metric on $\mathbb{R}^3$. Hence the map $\Pi = \Xi \circ \Upsilon$ is a bundle map over the inclusion $U \hookrightarrow \mathbb{R}^3$ with the following property:

Lemma 6.2. *Let Z be a vector field on $U \subset \mathbf{HS}^3$. Then Z is a Killing vector field if and only if $\Pi(Z)$ is a Killing vector field for the Euclidean metric on $\mathbb{R}^3$.*

The bundle map Π is also called an infinitesimal Pogorelov map, since it is an infinitesimal version of a remarkable map introduced by Pogorelov [Pog73] to handle rigidity questions in spaces of constant curvature.

6.2. **Rigidity with respect to HS structures.** Here, an *infinitesimal deformation* of P associates a vector tangent to $\partial\mathbf{H}^3$ to each vertex of P; the infinitesimal deformation is trivial if it is the restriction of a global Killing field of $\mathbf{HS}^3$.

Proposition 6.3. *Let $P \in \mathscr{P}_{p,q}$ and $\dot{P}$ be an infinitesimal deformation of P within $\mathscr{P}_{p,q}$. If $\dot{P}$ does not change the HS structure $\Delta(P)$ at first order, then $\dot{P}$ is trivial.*

Proof. As always, we work in an affine chart containing P. Suppose that $\dot{P}$ is a non-trivial infinitesimal deformation of P that does not change, at first order, the HS metric induced on P. Then the induced HS metric on each facet is constant at first order. Hence for each facet F, there is a Killing field κ_F such that the restriction of κ_F to the vertices of F is equal to the restriction of $\dot{P}$, and for two facets F and G, κ_F and κ_G agree on the common edge of F and G.

Lemma 6.2 shows that $\bar{\kappa}_F = \Pi(\kappa_F)$ is the restriction of a Killing field of $\mathbb{R}^3$, while it is clear that if F and G share an edge, then $\bar{\kappa}_F$ and $\bar{\kappa}_G$ agree on this edge. Therefore the restriction of $\bar{\kappa}_F$ to the vertices of P define an isometric first-order Euclidean deformation of P.

However, Alexandrov [Ale05] proved that convex polyhedra in $\mathbb{R}^3$ are infinitesimally rigid: any first-order Euclidean isometric deformation must be the restriction of a global Killing vector field of $\mathbb{R}^3$. So $\bar{\kappa}_F$ must be the restriction of a global Killing vector field $\bar{\kappa}$. Lemma 6.2 therefore implies that κ_F are the restriction to the faces of P of a global Killing vector field $\kappa = \Pi^{-1}(\bar{\kappa})$, which contradicts our hypothesis. □

6.3. **Shape parameters associated to edges.** We can identify $\partial\mathbf{H}^3$ with the extended complex plane $\mathbb{C}\mathrm{P}^1$, then vertices of an ideal polytope P can be described by complex numbers. By subdividing non-triangular facets if necessary, we may assume that facets of P are all triangles. Let $z_1z_2z_3$ and $z_2z_1z_4$ be two facets of P oriented with outward pointing normal vectors. The *shape parameter* on the common (oriented) edge z_1z_2 is the cross ratio

$$\tau = [z_1, z_2; z_3, z_4] = \frac{(z_1 - z_3)(z_2 - z_4)}{(z_2 - z_3)(z_1 - z_4)}.$$

Recall that each triangular facet of P determines a strongly ideal triangles with the same vertices. The two (oriented) hyperbolic triangles corresponding to $z_1z_2z_3$ and $z_2z_1z_4$ form a hyperbolic dihedral angle at their common edge z_1z_2. Let ϕ denote the *hyperbolic* exterior angle at z_1z_2. Then the shape parameter τ has a geometric interpretation: it can be written in the form of $\tau = \exp(\sigma + i\phi)$, where σ is the shearing between the two hyperbolic triangles.

The angle ϕ can be read from the hyperbolic vertex figure (see Section 4.3). If one walks along the polygonal curve, in the same direction as we specified for reading θ (see Section 3), then ϕ is nothing but the turning angle at every vertex, taking anti-clockwise turns as positive, and clockwise turns as negative.

Let v be a vertex of P, and let $\tau_1, \tau_2, \cdots, \tau_k$ be the shape parameters associated to the edges of P adjacent to v, in this cyclic order. The following relations, which holds for strongly inscribed polyhedra, also holds for the weakly inscribed P.

$$\prod_{i=1}^{k} \tau_i = 1, \tag{1}$$

$$\sum_{j=1}^{k} \prod_{i=1}^{j} \tau_i = 0. \tag{2}$$

Both equations can be easily understood by considering the hyperbolic vertex figure at v: (1) follows from the fact that $\sum \sigma_i = 0$ while $\sum \phi_i$ is a multiple of 2π. (2) is just saying that the vertex figure, considered as a polygonal curve in the Euclidean plane, closes up.

The shape parameters determine the local geometry (angle and shearing) at every edge, hence completely describe the polyhedron. A small perturbation in the shape parameter subject to (1) and (2) corresponds to a deformation of P into another weakly ideal polyhedron. Indeed, the convexity is stable under a small perturbation, then (1) and (2) guarantee that the hyperbolic vertex figures are closed polygonal curves, hence they are vertex figures of a weakly inscribed polyhedron.

6.4. **Rigidity with respect to dihedral angles.** We now have the necessary tools to prove the infinitesimal rigidity of weakly ideal polyhedra with respect to their dihedral angles.

Proposition 6.4. *Let $P \in \mathscr{P}_{p,q}$ and $\dot{P}$ be an infinitesimal deformation of P within $\mathscr{P}_{p,q}$. If $\dot{P}$ does not change the dihedral angles $\Theta(P)$ at first order, then $\dot{P}$ is trivial.*

Proof. Let $\dot{P} \in \mathbb{C}^{p+q}$ be an infinitesimal deformation of P, represented by the velocity of the vertices in $\mathbb{C}\mathrm{P}^1$. Let $\dot{\tau} = (\dot{\tau}_e)_{e\in E} \in \mathbb{C}^{|E|}$ be the corresponding first-order variation of the shape parameters associated to the edges. Suppose that $\dot{P}$ does not change the dihedral angles of P (at first order). This means that for all $e \in E$, $\dot{\tau}_e/\tau_e$ is real, because the argument of τ_e is equal to the dihedral angle at the corresponding edge.

Now consider the first-order variation $i\dot{\tau} = (i\dot{\tau}_e)_{e\in E}$ of the shape parameters. A crucial observation is that, since the conditions (1) and (2) above are polynomial, $i\dot{\tau}$ again corresponds to a first-order deformation of P, which we can call $i\dot{P}$. Now for all $e \in E$, $i\dot{\tau}_e/\tau_e$ is imaginary. This means that in the first-order deformation $i\dot{P}$, the shear along the edges remains fixed (at first order). So $i\dot{P}$ does not change, at first order, the HS-structure induced on P.

By Proposition 6.3, $i\dot{P}$ is trivial, and it follows that the infinitesimal deformation $\dot{P}$ is also trivial. The result follows. □

7. Topology

7.1. **Ideal polyhedra.** In this section we will conclude that Θ and Δ are homeomorphisms. The first step is to prove that the domain $\mathscr{P}_{p,q}$ is connected.

We work in a projective chart inconsistent with $\mathbb{H}^3$, in which $\partial\mathbb{H}^3$ is the quadric of equation $x_2 = f(x_0, x_1) = \pm\sqrt{x_0^2 + x_1^2 + 1}$. The following lemma allows us to place any weakly (p, q)-ideal polyhedron in a convenient position.

Lemma 7.1. *For any $P \in \mathscr{P}_{p,q}$, there is an isometry T of $\mathbb{H}^3$ such that $T(P)$ contains the origin and the points $(0, 0, \pm 1)$.*

Proof. The proof uses the Hyperplane Separation Theorem for hyperbolic space. We sketch here a quite standard proof in the spirit of [BV04], as some ingredient in this proof would be useful for us.

Note that $\mathbb{H}^3 \cap P$ consists of two disjoint components, denoted by P^+ and P^-, both are convex subsets of $\mathbb{H}^3$. Let $u \in P^+$ and $v \in P^-$ such that the hyperbolic distance between u and v achieves the minimum hyperbolic distance between P^+ and P^-. This distance is necessarily finite, hence u and v are necessarily on the boundary ∂P. We claim that the hyperbolic plane H that perpendicularly bisects the segment $\overline{uv}$ separates P^+ and P^-, i.e. $P^\pm$ are on different sides of H. To see this, assume $u' \in P^+$ is on the same side of H as P^-. Then a perturbation of u towards u' would be closer to v, contradicting our choice of u and v.

Now let T be the isometry of $\mathbb{H}^3$ that sends the separating plane H to infinity. Then the polar point of H, which is contained in P, is sent to the origin in the interior of $T(P)$. Moreover, the line through u and v is sent to the x_2-axis. In particular, the points $(0, 0, \pm 1)$ are also in the interior of $T(P)$. □

Proposition 7.2. *$\mathscr{P}_{p,q}$ is connected.*

Proof. Let $P \in \mathscr{P}_{p,q}$. Thanks to the previous lemma, we may assume that P contains the origin and the points $(0, 0, \pm 1)$. We now define a deformation of $P \in \mathscr{P}_{p,q}$.

If $v \in V_\pm$, define v_t, $t > 1$, as follows: $v_t = v$ if the x_2-coordinate of v is smaller than t; otherwise, v_t is a point of height $x_2 = \pm t$ obtained by moving v along the gradient of f (for metric induced on the quadric by the Euclidean metric $dx_0^2 + dx_1^2 + dx_2^2$ in a chart) towards $(0, 0, \pm 1)$. We claim that the point set $V_t = \{v_t \mid v \in V_+ \cup V_-\}$ remains in convex position for all $t > 1$. If $v_t \in V_t$ is at height $x_2 = \pm t$, v_t is on the circle $x_0^2 + x_1^2 = t^2 - 1$; the convexity at v_t is then immediate. Otherwise, v_t coincides with a vertex v of P, and we claim that a supporting hyperplane H_v of P at v is also a supporting hyperplane of P_t. This can be seen by noting that, for any $u \in \partial\mathbb{H}^3_\pm$ on the same side of H_v as $(0, 0, \pm 1)$, as long as u is sufficiently close to H_v, the gradient of f at u points away from H_v. Because $(0, 0, \pm 1) \in P$ and H_v is supporting, no vertex v_t would move across H_v as t decreases.

Define P_t as the convex hull of V_t. We see that $P_t = P$ for t sufficiently large. As t approaches 1, the vertices of P_t lie, eventually, on two horizontal planes $x_2 = \pm t$.

Now assume another polyhedra $P' \in \mathscr{P}_{p,q}$, which also contains the origin and the points $(0, 0, \pm 1)$. For t sufficiently close to 1, both P_t and P'_t have vertices on the horizontal planes $x_2 = \pm t$. Polyhedra with vertices on these two planes form a connected subset of $\mathscr{P}_{p,q}$; indeed, any choice of p and q ideal points on these two planes determines uniquely such a polyhedra. Hence we find a continuous path from between P and P', which proves the connectedness of $\mathscr{P}_{p,q}$. □

Let $\mathscr{P}^i_{p,q}$ denote the open subset of $\mathscr{P}_{p,q}$ consisting of polyhedra with an edge 1^+i^-. $\{\mathscr{P}^i_{p,q} \mid 1 \le i \le q\}$ form an open cover of $\mathscr{P}_{p,q}$. During the deformation P_t in the previous proof, we may rotate V_- around the x_2-axis so that 1^+i^- remains an edge. This shows that

Proposition 7.3. *$\mathscr{P}^i_{p,q}$ is connected.*

7.2. **Admissible graphs.** Correspondingly, let $\mathscr{A}^i_{p,q}$ denote the open subset of $\mathscr{A}_{p,q}$ consisting of graphs θ with 1^+i^- as an edge (recall the vertex labeling from Section 2). That is, either $\theta(1^+i^-) < 0$, or $\theta(1^+i^-) = 0$ but the graph remains polyhedral if we include 1^+i^- as an edge. Then $\{\mathscr{A}^i_{p,q} \mid 1 \le i \le q\}$ form an open cover of $\mathscr{A}_{p,q}$.

Proposition 7.4. *For $1 \leq i \leq q$, $\mathscr{A}^i_{p,q}$ is homeomorphic to $\mathbb{R}^{2(p+q-3)}$.*

We treat two cases separately.

Proof for the case $p < 2 < q$. We may take $i = 1$ without loss of generality. To construct a graph $\theta \in \mathscr{A}^1_{p,q}$, we first assign positive weights then negative weights.

For the positive weights, our task is to find an outerplanar graph θ^+ on the vertices of the q-cycle with only positive edge weights. For the final result to be an element of $\mathscr{A}^1_{p,q}$, we need the support of θ^+ to contain the q-cycle, and the sum of the weights of θ^+ to be π. This condition on the sum can be seen by adding up the weight sum of θ around all the vertices of the q-cycle, and noticing that the edges of θ^+ are double counted. Let $\mathscr{A}^+$ be the set of all θ^+ satisfying those conditions.

Note that θ^+ can be embedded in the plane in such a way that the q-cycle is embedded as a q-gon, and the other edges are embedded as non-crossing diagonals of the q-gon. Let $\mathscr{A}^+_0$ denote the set of positive graphs that are only supported on the q-cycle, and $\mathscr{A}^+_1$ the set of positive graphs that are only supported on non-crossing diagonals. Then any $\theta^+ \in \mathscr{A}^+$ can be written in the form $\theta^+ = (1-t)\theta^+_0 + t\theta^+_1$ where $\theta^+_0 \in \mathscr{A}_0$, $\theta^+_1 \in \mathscr{A}_1$, for some $0 \leq t < 1$ (note the strict inequality!).

It is easy to see that $\mathscr{A}^+_0$ is a $(q-1)$-dimensional open simplex. Graphs in $\mathscr{A}^+_1$ with the same combinatorics (that is, supported on the same diagonals) also form an open simplex, whose dimension is the cardinality of their support minus 1. $\mathscr{A}^+_1$ is therefore a simplicial complex: The maximal simplices are of dimension $q-4$, corresponding to the maximal set of non-crossing diagonals, and they are glued along their faces corresponding to common subsets. This simplicial complex is well-known as the boundary of a polytope, namely the *associahedron* [Lee89].

Therefore, the closure of $\mathscr{A}^+ = (1-t)\mathscr{A}^+_0 + t\mathscr{A}^+_1$, $0 \leq t < 1$, is topologically the join of a $(q-1)$-ball and a $(q-4)$-sphere, hence homeomorphic to a $(2q-4)$-ball. The openness of $\mathscr{A}^+$ follows from the openness of $\mathscr{A}^+_0$ and the strict inequality $t < 1$.

Once positive weights are assigned, the negative weights are uniquely determined since $p = 1$ and all vertices of the q-cycle are connected to only one negatively weighted edge. Hence $\mathscr{A}^i_{p,q}$ is homeomorphic to $\mathbb{R}^{2(p+q-3)}$. □

We need more ingredients for the case $2 \leq p \leq q$.

First, we claim that if a graph θ is admissible, and the negative weights of θ sum up to $-2\omega > -2\pi$, then for any $0 < t < \pi/\omega$, the scaled graph $t\theta$ is also admissible. The claim follows from the following lemma, which guarantees that Condition **(A1)** is not violated after the scaling:

Lemma 7.5. *Let $-2\omega > -2\pi$ be the sum of negative weights of θ. Then any negative weight of θ is at least $-\omega$.*

Proof. We argue by contradiction and suppose that there is an edge e with negative weight strictly less than $-\omega$. Let v be an endpoint of e. The sum of the positive weights on the red edges $e_1, \cdots, e_k$ adjacent to v is at least ω. Let $v_1, \cdots, v_k$ be their endpoints different from v; e is not adjacent to any of them. Then the sum of the negative weights over the blue edges adjacent to $v_1, \cdots, v_k$ must be strictly less than $-\omega$. We then conclude that the sum of negative weights is not -2ω as assumed, but strictly less. □

This lemma also proves the redundancy of Condition **(A3)** and part of Condition **(A1)**. Any weight function that satisfy Conditions **(W1)** and **(W2)** can be normalized to satisfy Conditions **(A1)**–**(A3)**. Hence these conditions are not present in the main Theorem 1.5.

Proof for the case $2 \leq p \leq q$. We may take $i = 1$ without loss of generality. Fix a number $0 < \omega < \pi$. We will prove that the set of admissible graphs in $\mathscr{A}^1_{p,q}$ with negative weights summing up to -2ω is homeomorphic to $\mathbb{R}^{2(p+q)-7}$. To construct such a graph θ, we follow the same strategy as before. That is, first assign positive weights then negative weights.

For the positive weights, we need to find θ^+ that is the disjoint union of two outerplanar graphs, one on the vertices of the p-cycle, and the other on the vertices of the q-cycle, with only positive edge weights. Moreover, we need the sum of the weights of each disjoint component to be ω. Hence each component can be obtained by taking the θ^+ constructed for the case $p < 2 < q$,

and multiplying its weights by a constant ω/π. The space of such θ^+ is then homeomorphic to $\mathbb{R}^{2(p+q-4)}$.

We then propose an algorithm that assigns weights to negative edges and outputs an admissible graph in $\mathscr{A}^1_{p,q}$. This algorithm depends on one parameter taken in a segment, hence proves the proposition.

Recall that vertices are labeled by $1^+, \dots, p^+$ and $1^-, \dots, q^-$ respectively, following the boundary of the outerplanar subgraphs in a compatible direction. Also recall that the weight w_v of a vertex v is the sum of weights $\theta(u, v)$ over all other vertices u. The vertex weights change as we update the edge weights. Before we proceed, the weights are positive on every vertex, because only positive weights are assigned. Our goal is to cancel them with negative weights. We also keep track of two indices i and j, initially both 1. At each step, we assign a negative weight to the edge i^+j^-. Moreover, the graph will be embedded in $\mathbb{S}^2$ throughout the algorithm.

We start with an embedding of θ^+ in $\mathbb{S}^2$, such that the two outerplanar components are embedded as two disjoint polygons with non-crossing diagonals. Interiors of the two polygons are declared as forbidden area: during our construction, no new edge is allowed to intersect this area. In other words, we are only allowed to draw edges within a belt.

For the first step, we draw a curve connecting 1^+ and 1^-, to which we are free to assign any non-positive weight ranging from $-\min(w_{1^+}, w_{1^-})$ to 0. We also have the freedom to choose a sign $\sigma = \pm$, and increment i if $\sigma = +$, increment j if $\sigma = -$.

In the following steps, we adopt a greedy strategy. We draw a curve connecting i^+ and j^-, to which we assign the weight $-\min(w_{i^+}, w_{j^-})$. After this assignment, the face bounded by i^+j^- and the previously assigned edge is considered as a forbidden area for later construction: no new edge is allowed to intersect this area. Then we increment i if $w_{i^+} = 0$, and increment j if $w_{j^-} = 0$. If both weights vanish, we increment both indices.

Eventually, we will have $i = p + 1$ and $j = q + 1$, and the weights vanish at all vertices. Then the algorithm stops. The result is by construction the embedding of a (p, q)-admissible graph.

In this algorithm, being greedy is not only good, but also necessary. Note that the weights between vertices of smaller indices are already fixed. If we choose any bigger negative weight for the edge i^+j^-, then both w_{i^+} and w_{j^-} remain positive. They both need to be connected to a vertex with larger index to cancel the weight. This is however not possible, since neighborhoods of these vertices have been declared as forbidden area.

The algorithm is parametrized only by the two choices at the first step: a non-positive weight and a sign. The space of parameters is therefore homeomorphic to a segment. □

We have $\Theta^{-1}(\mathscr{A}^i_{p,q}) \subseteq \mathscr{P}^i_{p,q}$. Let Θ_i denote the restriction of Θ on $\mathscr{P}^i_{p,q}$; it is a covering map with images in $\mathscr{A}^i_{p,q}$. Since $\mathscr{P}^i_{p,q}$ is connected and $\mathscr{A}^i_{p,q}$ simply connected, we conclude that $\Theta^i_{p,q}$ is a homeomorphism. This proves that the covering number of Θ is 1, i.e. Θ is a homeomorphism.

7.3. **Admissible HS structures.** Let $\mathbf{H}^2_\ell$ denote the complete, simply connected hyperbolic surface with one cone singularity of angle ℓ. We extend this notation, and use $\mathbf{H}^2_0$ for the hyperbolic surface with one cusp. A non-degenerate boundary component of $\mathrm{d}\mathbb{S}^2_\ell$ can be identified to the boundary of $\mathbf{H}^2_\ell$.

Let B be a subset of p points on $\partial\mathbf{H}^2$, considered up to isometries of $\mathbf{H}^2$. Let $\mathscr{G}_{B,\ell}$ be the space of CPP maps from $\partial\mathbf{H}^2$ to $\partial\mathbf{H}^2_\ell$, $0 < \ell < 2\pi$, up to isometries of both $\mathbf{H}^2$ and $\mathbf{H}^2_\ell$, with positive break point set B. We denote by $\mathscr{G}_B$ the union of $\mathscr{G}_{B,\ell}$ for $0 < \ell < 2\pi$, i.e. set of all CPP maps on $\partial\mathbf{H}^2$ with positive break point set B, up to isometries.

A *horocyclic p-gon* is the intersection of p horodisks in $\mathbf{H}^2$. Figure 9 shows a horocyclic 3-gon and a horocyclic 4-gon.

The key observation is the following homeomorphism η from $\mathscr{G}_B$ to the space of horocyclic p-gons bounded by horocycles based at B.

Label the elements of B as $b_1, \dots, b_p = b_0$ in the clockwise order in the disk model. They are the vertices of an ideal p-gon $P \subset \mathbf{H}^2$. Consider a map $\gamma \in \mathscr{G}_{B,\ell}$. It maps P to an ideal p-gon $P' \subset \mathbf{H}^2_\ell$ with a cone singularity of angle ℓ in its interior. The vertices of P' are $\gamma(b_i)$.

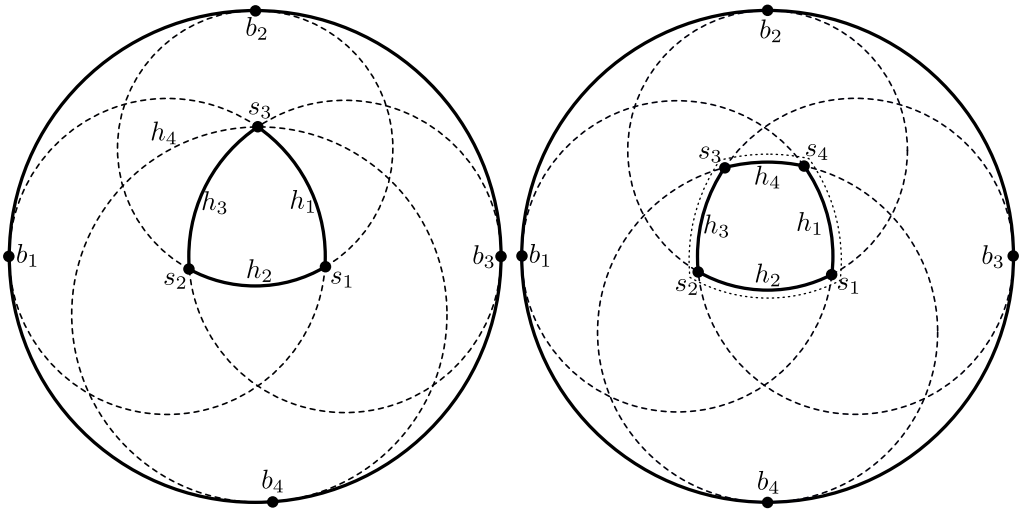

FIGURE 9. The horocyclic 3-gon on the left, together with the horocycle h_4, represent a point in $\partial\mathscr{G}_B$ with $B = \{b_1, b_2, b_3, b_4\}$. Shrinking h_4 truncates the 3-gon into the 4-gon on the right. We also sketch a scaling of the 4-gon, which is the key in the proof of Proposition 7.10.

We then obtain a horocyclic p-gon $\eta(\gamma)$ as follows. Cut P' into p triangles by the geodesics from $\gamma(b_i)$ to s, the cone singularity of $\mathbf{H}^2_\ell$. Then fold each triangle $\gamma(b_i)s\gamma(b_{i+1})$ inward, isometrically, into the triangle $b_is_ib_{i+1}$ in $\mathbf{H}^2$. Let h_i be the horocycle based at b_i passing through s_i. Then $\eta(\gamma)$ is the intersection of the horodisks bounded by the h_i's.

Since γ preserves horocycles, s_{i-1} must also lie on h_i. Hence the horocyclic segments s_is_{i+1} form a piecewise horocyclic closed curve, denoted by h.

Lemma 7.6. *h is embedded as the boundary of $\eta(\gamma)$.*

Proof. Since b_i are positive break points of γ, the triangles $b_is_ib_{i+1}$ and $b_{i-1}s_{i-1}b_i$ must overlap; see Proposition 4.13 and Remark 4.14. In other words, the points b_i, s_{i-1} and s_i are placed on h_i in the clockwise order. See Figure 9 for examples.

For $x \in \partial\mathbf{H}^2$ and $y \in \mathbf{H}^2$, we use $\xi(x, y)$ to denote the other ideal end of the geodesic that emerges from x and passes through y. Define a map $g : \partial\mathbf{H}^2 \to \partial\mathbf{H}^2$ such that $g^{-1}(b_i) = \{\xi(b_i, y) \mid y \in h_i \cap h\}$ and, for some x between b_i and b_{i+1}, $g^{-1}(x) = \xi(x, s_i)$. The map g is continuous and monotone, and has the property that $x \notin g^{-1}(x)$, hence its degree must be 1. This proves that h is embedded, hence the boundary of $\eta(\gamma)$. □

Conversely, given a horocyclic polygon bounded by horocycles h_i based at $b_i \in B$, let $s_i \in h_i \cap h_{i+1}$ be its vertices. Then we can glue the triangles $b_is_ib_{i+1}$ into an ideal p-gon with a cone singularity of angle ℓ. More specifically, ℓ is the sum of the angles $\angle b_is_ib_{i+1}$. This proves that η is a homeomorphism.

Remark 7.7. It is interesting to note from the horocyclic polygons that $\ell < 2\pi$. Let s be a point in the interior of the horocyclic polygon. We then have $\angle b_is_ib_{i+1} < \angle b_isb_{i+1}$ for all i. Yet the sum of $\angle b_isb_{i+1}$ is equal to 2π.

Proposition 7.8. *$\mathscr{G}_B$ is homeomorphic to $\mathbb{R}^{|B|}$.*

Proof. The proof is by induction on the cardinality $p = |B|$. Up to isometries of $\mathbf{H}^2$, we may consider B fixed.

For $p = 2$, a horocyclic 2-gon P_2 is bounded by two horocycles. It is determined by the position of an intersection point of the two horocycles. This point can be chosen arbitrarily in $\mathbf{H}^2$, proving the statement for $p = 2$.

Now consider a horocyclic $(p-1)$-gon P_{p-1} bounded by horocycles $h_1, \dots, h_{p-1}$ based at $b_1, \dots, b_{p-1} \in B$, and let $s_i \in h_i \cap h_{i+1}$ and $s_{p-1} \in h_{p-1} \cap h_1$ denote the vertices of P_{p-1}. We now construct a horocyclic p-gon P bounded by horocycles with bases in B. For this, it suffices to choose a horocycle h_p that truncates the vertex s_{p-1} of P_{p-1}. This can be done by taking the horocycle h_p at b_p passing through s_{p-1}, then shrinking it. On the left of Figure 9 we illustrate a truncation of a 3-gon into a 4-gon. We can continue to shrink h_p until it hits another vertex of P_{p-1}.

Hence $\mathscr{G}_B$ is homeomorphic to $\mathscr{G}_{B\setminus\{b_p\}} \times \mathbb{R}$, which is $\mathbb{R}^{|B|}$ by induction. □

In the following we will consider the closures of $\mathscr{G}_B$. The boundary $\partial\mathscr{G}_B$ consists of three parts, namely $\mathscr{G}_{B,0}$, $\mathscr{G}_{B,2\pi}$, and $\partial\mathscr{G}_{B,\ell}$, $0 < \ell < 2\pi$. We now define and describe these spaces.

We use the notation $\mathscr{G}_{B,0}$ for the space of CPP maps from $\partial\mathbf{H}^2$ to $\partial\mathbf{H}^2_0$, up to isometries of both $\mathbf{H}^2$ and $\mathbf{H}^2_0$, with positive break point set B. As before, we can interpret a map $\gamma \in \mathscr{G}_{p,0}$ as gluing the boundary of an ideal p-gon $P \subset \mathbf{H}^2$ to the boundary of an ideal p-gon $P' \subset \mathbf{H}^2_0$, where P' contains a cusp s. We triangulate P' by connecting its vertices to s, and triangulate P arbitrarily, and glue them through γ to obtain a triangulation T of the 2-sphere. The shearing coordinates on the $p-3$ interior edges of P are determined by the positions of the break points. The shearing coordinates on the remaining $2p$ edges of T are governed by the condition that the shearing around each vertex of T should sum up to 0. Since T has $p+1$ vertices, we conclude that $\mathscr{G}_{B,0}$ is homeomorphic to $\mathbb{R}^{p-1} = \mathbb{R}^{|B|-1}$.

The part $\partial\mathscr{G}_{B,\ell}$, $0 < \ell < 2\pi$, consists of CPP maps from $\partial\mathbf{H}^2$ to $\partial\mathbf{H}^2_\ell$ up to isometries of both $\mathbf{H}^2$ and $\mathbf{H}^2_\ell$, with break point set $B' \subset B$ and marked projective point set $B \setminus B'$. The homeomorphism η extends continuously to this boundary. More specifically, let $p = |B|$ and $p' = |B'|$. Then for $\gamma \in \partial\mathscr{G}_{B,\ell}$, $\eta(\gamma)$ is a horocyclic p'-gons P' bounded by horocycles with bases in B', together with $p - p'$ additional horocycles with bases in $B \setminus B'$ "supporting" P' (that is, they intersect $\partial P'$ but disjoint from the interior of P'). The left side of Figure 9 is an example with $p = 4$ and $p' = 3$.

We also abuse the notation $\mathscr{G}_{p,2\pi}$ for the space of projective homeomorphisms from $\partial\mathbf{H}^2$ to itself up to isometries of $\mathbf{H}^2$, with a set of marked points B. In fact, the marking here is superficial; hence $\mathscr{G}_{p,2\pi}$ consists of a single element.

Let δ_i denote the distance from s_i to the geodesic b_ib_{i+1}. The hyperbolic triangle formula shows that α_i and δ_i are related by the formula $\cosh\delta_i \sin\alpha_i/2 = 1$. We now deform the horocyclic p-gon by moving every s_i, simultaneously, along the geodesic perpendicular to b_ib_{i+1}, to a new position s'_i. The following lemma is the crucial observation for the proof that follows later.

Lemma 7.9. *If the deformation is performed in such a way that the ratio* $\cosh\delta_i/\cosh\delta'_i$ *is the same for every* i*, then there is a horocycle* h'_i *passing through every adjacent pair* s'_{i-1} *and* s'_i*.*

Proof. Let k be the common ratio of $\cosh\delta_i/\cosh\delta'_i$. We use the half-plane model of $\mathbf{H}^2$, and assume that $b_i = \infty$. The situation is illustrated in Figure 10. Let ψ_i be the (Euclidean) angle $\angle s_ib_{i+1}b_{i-1}$ and ψ_{i-1} be the angle $\angle s_{i-1}b_{i-1}b_{i+1}$. Then for $j = i$ or $i-1$, we can calculate the distances (see for instance [And05, §3.5])

$$\delta_j = \ln|\csc\psi_j + \cot\psi_j| = \operatorname{arccosh}\csc\psi_j.$$

This is particularly convenient because we have $\cosh\delta_j \sin\psi_j = 1$. Hence $\cosh\delta_j/\cosh\delta'_j = k$ implies that $\sin\psi_j/\sin\psi'_j = 1/k$ for every i. In the half-plane model, s_i and s_{i-1} are moving along the circles centered at b_{i+1} and b_{i-1}, respectively, such that their heights are both scaled by $1/k$. Therefore, their new positions s'_{i-1} and s'_i are again at the same height, hence on the same horocycle h'_i based at $b_i = \infty$. □

This deformation is sketched on the right side of Figure 9. We see from Figure 10 that, by moving s_j's outwards the horocyclic polygon, one can multiply $\cosh(\delta_i)$ by an arbitrarily large constant. However, if we move s_j's inwards the horocyclic polygon, s_i and s_{i-1} would eventually merge.

Proposition 7.10. $\mathscr{G}_{B,\ell}$, $0 \le \ell \le 2\pi$*, are contractible, and homeomorphic to* $\mathbb{R}^{|B|-1}$ *if* $\ell < 2\pi$*.*

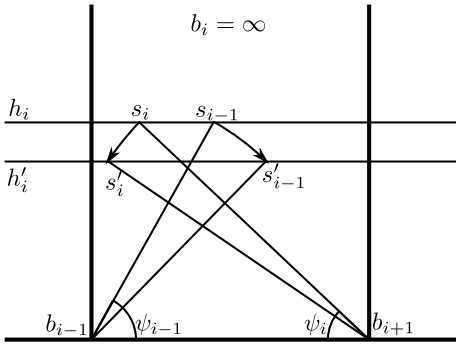

FIGURE 10. Proof of Lemma 7.9.

Proof. Through any given point $\gamma \in \mathscr{G}_B$, the deformation described above defines a continuous path with monotonically changing cone angle ℓ. We use ℓ as the parameter and denote this path by $c_\gamma(\ell)$. See Figure 11 for an illustration.

To decrease ℓ, one needs to move the vertices outwards the horocyclic polygon. We have seen that, in this direction, one can travel along $c_\gamma(\ell)$ until hitting $\mathscr{G}_{B,0}$.

To increase ℓ, one needs to move the vertices inwards the horocyclic polygon. In this direction, however, the path $c_\gamma(\ell)$ would in general hit some $\gamma' \in \partial\mathscr{G}_{B,\ell}$ for some $\ell < 2\pi$, as shown in Figure 11. Further movement of the vertices in the same direction would destroy the p-gon. However, we can continue deforming γ' as a gluing map in $\mathscr{G}_{B'}$ for some $B' \subset B$. Hence $c_\gamma(\ell)$ is extended into the closure of $\mathscr{G}_B$, along which one can increase ℓ up to 2π; see Figure 11.

This path is uniquely defined through every $\gamma \in \overline{\mathscr{G}}_B$, and two path do not intersect inside $\mathscr{G}_B$; intersection is only possible on the boundary. For $0 \leq \ell \neq \ell' \leq 2\pi$, let $f_{\ell,\ell'} \colon \mathscr{G}_{B,\ell} \to \mathscr{G}_{B,\ell'}$ be the continuous map that sends $\gamma \in \mathscr{G}_{B,\ell}$ to $\gamma' = c_\gamma(\ell') \in \mathscr{G}_{B,\ell'}$. Then $f_{\ell,\ell'}$ and $f_{\ell',\ell}$ define a homotopy equivalence between $\overline{\mathscr{G}}_{B,\ell}$ and $\overline{\mathscr{G}}_{B,\ell'}$. Consequently, $\overline{\mathscr{G}}_{B,\ell}$ are all of the same homotopy type. We have seen that $\mathscr{G}_{B,0}$ and $\mathscr{G}_{B,2\pi}$ are contractible, and so must be $\mathscr{G}_{B,\ell}$ for $0 < \ell < 2\pi$.

In general, $f_{\ell,\ell'}$ and $f_{\ell',\ell}$ are not inverse to each other. However, if the path emerging from $\gamma \in \mathscr{G}_{B,\ell}$ arrives at $\gamma' = f_{\ell,\ell'}(\gamma) \in \mathscr{G}_{B,\ell'}$ without hitting the boundary of $\mathscr{G}_B$, then one can travel backwards along the same path. The reversed path defines $f_{\ell',\ell}$, hence we have $f_{\ell',\ell} \circ f_{\ell,\ell'}(\gamma) = f_{\ell',\ell}(\gamma') = \gamma$.

This is the case, in particular, when $\ell > \ell'$ and $\gamma \in \mathscr{G}_{B,\ell}$. Then $f_{\ell,\ell'}$ defines a homeomorphism between $\mathscr{G}_{B,\ell}$ and its image $f_{\ell,\ell'}(\mathscr{G}_{B,\ell}) \subset \mathscr{G}_{B,\ell'}$. We finally conclude that $\mathscr{G}_{B,\ell}$, $0 < \ell < 2\pi$, by its continuity in ℓ, are all homeomorphic to $\mathscr{G}_{B,0}$, thus to $\mathbb{R}^{|B|-1}$. □

It is now clear that $\mathscr{G}_B$ is foliated by $\mathscr{G}_{B,\ell}$, $0 < \ell < 2\pi$, as illustrated in Figure 11.

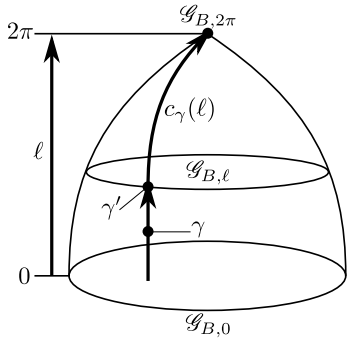

FIGURE 11. The structure of $\mathscr{G}_B$ showing a path c_γ.

An admissible HS structure is obtained by gluing copies of $\mathbf{H}^2$ to the ideal boundary components of $\mathrm{d}\mathbb{S}^2_\ell$. We now conclude on the topology of $\mathscr{M}_{p,q}$, and prove Theorem 3.4.

If $p < 2 < q$, an element of $\mathscr{M}_{1,q}$ can be constructed by first choosing a set of q points $B \subset \partial\mathbf{H}^2$ up to isometries, and then a gluing map $\gamma \in \mathscr{G}_{B,0}$. For the first step, we may fix three points up to hyperbolic isometry, hence the space of B is homeomorphic $\mathbb{R}^{q-3}$. In the second step, we have seen that $\mathscr{G}_{B,0}$ is homeomorphic to $\mathbb{R}^{q-1}$. Hence $\mathscr{M}_{1,q}$ is homeomorphic to $\mathbb{R}^{2q-4} = \mathbb{R}^{2(p+q-3)}$. Theorem 3.4 follows since both $\mathscr{G}_{p,0}$ and $\mathscr{P}_{p,0}$ are contractible.

If $2 \leq p \leq q$, we first count the dimension. For the gluing map on one boundary of $\mathrm{d}\mathbb{S}_\ell$, we need to choose a set B of p break points, then pick a gluing map from $\mathscr{G}_B$. Up to isometries, the space of this gluing map is of dimension $2p-3$. Similarly, the gluing map on the other boundary contributes $2q-3$ dimensions. The two gluing maps have the same cone angle, removing one degree of freedom. But we can also rotate the break points on $\partial\mathbf{H}^2$, corresponding to translations in $\mathbb{R}\mathrm{P}^1$. This contributes another dimension, hence the dimension of $\mathscr{M}_{p,q}$ is $2(p+q-3)$.

The rotation of $\partial\mathbf{H}^2$ generates the non-trivial fundamental group of $\mathscr{M}_{p,q}$. To prove that the map Δ is a homeomorphism, we lift it to a map $\tilde{\Delta}$ between the universal covers $\tilde{\mathscr{P}}_{p,q}$ and $\tilde{\mathscr{M}}_{p,q}$. A point in $\tilde{\mathscr{P}}_{p,q}$ corresponds to a (p,q)-ideal polyhedron equipped with a path (defined up to homotopy) connecting vertex 1^+ and 1^-. A point in $\tilde{\mathscr{M}}_{p,q}$ corresponds to a (p,q)-admissible HS structure with a path (up to homotopy) connecting 1^+ and 1^- in $\mathrm{d}\mathbb{S}^2_\ell$. Hence $\tilde{\Delta}$ is a homeomorphism between marked (p,q)-ideal polyhedra and marked (p,q)-admissible HS structures. This proves that the covering number of Δ is 1.

8. Combinatorics

It remains to prove Theorem 1.4 from Theorem 1.5. In other terms, assume that a graph $\Gamma = (V, E)$ satisfies Condition **(C1)** and the edges are colored as specified in Theorem 1.5. We need to prove that Condition **(C2)** implies the existence of a weight function $w : E \to \mathbb{R}$ satisfying Conditions **(W1)** and **(W2)** and, conversely, existence of such a weight function implies Condition **(C2)**.

We consider two cases.

Case $p < 2 < q$. In this case, the cycle cover contains a 1-cycle, say with vertex set $V^+ = \{v_0\}$. The other vertices V^- induce a 2-connected outerplanar subgraph. We color the edges adjacent to v_0 by blue, and other edges by red.

Theorem 1.4 requires a cycle visiting all the edges along which the edge color has the pattern ...-blue-blue-red-.... This is actually equivalent to a much simpler condition:

Lemma 8.1. *In the case of $p < 2 < q$, Condition* **(C2)** *is equivalent to*

(C'2) *v_0 is connected to every vertex in V^-.*

Proof.

(C'2) $\implies$ (C2): Immediate.

(C2) $\implies$ (C'2): If some $v \in V^-$ is not connected to v_0, then the edges adjacent to v can not belong to a cycle as required in Condition **(C2)**.

□

Proof of Theorem 1.4 when $p < 2 < q$.

(C2) $\implies$ (W1) $\wedge$ (W2): Suppose a cycle c specified by Condition **(C2)**. Let n be its length; n is necessarily a multiple of 3. Let n_e be the number of times that c visits e. Assign to e the weight n_e if e is red, or the weight $-n_e$ if e is blue. After normalization by a factor $3\pi/n$, we obtain a graph that satisfies Conditions **(W1)** and **(W2)**.

(W1) $\wedge$ (W2) $\implies$ (C'2): Assume a graph function satisfying Condition **(W1)**. If some vertex $v \in V^-$ is not connected to v_0, then the edges adjacent to v are all of positive weight. This violates Condition **(W2)**.

□

Case $2 \leq p \leq q$. Condition **(C2)** requires a cycle with alternating colors; we call such a cycle an *alternating cycle*. As in the previous case, the existence of an alternating cycle that visits every edge implies that every vertex is adjacent to a blue edge, but the converse is not true. However, we have the following equivalence, which does not depend on Condition **(C1)**.

Proposition 8.2. *If the edges of a graph are colored in blue and red, then Condition* **(C2)** *is equivalent to*

(C"2) *Every edge belongs to an alternating cycle (which does not necessarily visits every edge).*

The proof is immediate from the composition and decomposition of cycles.

Proof of Theorem 1.4 when $2 \leq p \leq q$.

(C2) $\Longrightarrow$ **(W1)** $\wedge$ **(W2):** Suppose a cycle c specified by Condition **(C2)**. Let n_e be the number of times that c visits e. Assign to e the weight n_e if e is red, or the weight $-n_e$ if e is blue. We obtain a graph that satisfies Conditions **(W1)** and **(W2)**.

(W1) $\wedge$ **(W2)** $\Longrightarrow$ **(C"2):** Let w be a graph satisfying Conditions **(W1)** and **(W2)**.

If w has an alternating cycle c, the number of visits defines a graph supported on the edges of c, which we denote by w_c. We can choose a positive number α such that $w' = w - \alpha w_c$ is supported on a proper subgraph of w, but still satisfy Condition **(W1)** on other edges. Most importantly, w' satisfy Condition **(W2)** because both w and w_c do. The cycle c no longer exist in w'. We repeat this operation if there are other alternating cycles. Since the graph is finite, we will obtain a graph $\tilde{w}$ with no alternating cycle in finitely many steps.

If an edge e_0 of w does not belong to any alternating cycle, it must also be an edge of $\tilde{w}$. But we prove in the following that this is not possible.

Assume that e_0 is red and let v_0^+ and v_0^- be its vertices. Condition **(W2)** guarantees that v_0^+ is adjacent to a blue edge, e_1, whose other vertex is denoted by v_1^+. The same argument continues and we obtain an alternating path $e_0, e_1, e_2, \ldots$.

Since the graph is finite, this path will eventually intersect itself. That is, $v_i^+ = v_{i'}^+$ for some $0 \leq i' < i$ (note that we don't consider v_0^- as visited by e_0). If e_i and $e_{i'}$ are of the same color, $e_{i'+1}, \ldots, e_i$ form an alternating cycle, contradicting our assumption. Hence e_i and $e_{i'}$ must have different colors.

The same argument applies on the other vertex v_0^- of e_0. We obtain an alternating path $e_0, e_{-1}, e_{-2}, \ldots$. Let v_j^- denote the common vertex of e_{-j} and e_{-j-1}. This path eventually intersect itself, i.e. $v_j^- = v_{j'}^-$ for some $0 \leq j' < j$ (this time we don't consider v_0^+ as visited by e_0). Again, e_{-j} and $e_{-j'}$ must have different colors.

But then, $e_0, \ldots, e_i, e_{i'}, \ldots, e_0, \ldots, e_{-j}, e_{-j'}, \ldots, e_0$ form an alternating cycle; see Figure 12. This contradicts our assumption.

□

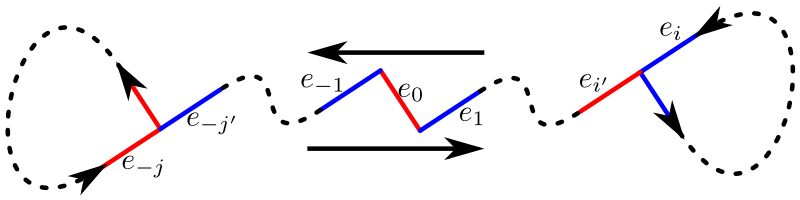

FIGURE 12. The alternating cycle in the last step in the proof of Theorem 1.4.

Remark 8.3. The feasibility region specified in Theorem 1.5 is a polyhedral cone. The proof above shows that the extreme rays of this cone correspond to the minimal alternating cycles in the graph.

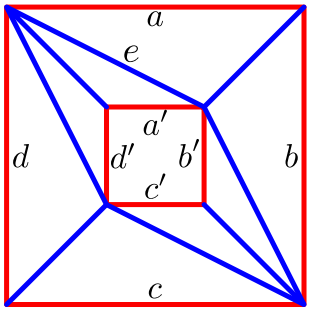

FIGURE 13. This graph is not the 1-skeleton of any weakly ideal polyhedron.

Example 8.4. The example in Figure 13 shows that Condition **(C2)** is essential. This graph is not the 1-skeleton of any weakly ideal polyhedron with the inner square in $\mathbf{H}^3_+$ and the outer square in $\mathbf{H}^3_-$. A fairly elementary argument, left to the reader, shows that there is no alternating cycle containing edge e. This can also be shown using Theorem 1.5, since if a graph w satisfies Conditions **(W1)** and **(W2)**, we would have

$$\begin{aligned} w(a) + w(b) &< w(a') + w(b') \ , \\ w(c) + w(d) &< w(c') + w(d') \ , \\ w(b) + w(c) &> w(b') + w(c') \ , \\ w(d) + w(a) &> w(d') + w(a') \ , \end{aligned}$$

from which a contradiction immediately follows.

Remark 8.5. Given a graph G with edges colored in blue and red, we define a directed graph $\tilde{G}$ as follows:

- Each vertex v of G lifts to two vertices v_+ and v_- in $\tilde{G}$.
- Each red edge uv of G lifts to two oriented edges u_-v_+ and v_-u_+ in $\tilde{G}$.
- Each blue edge uv of G lifts to two oriented edges u_+v_- and v_+u_- in $\tilde{G}$.

It is quite clear from the definition that an alternating cycle in G lifts to two oriented cycle in $\tilde{G}$, and any oriented cycle in $\tilde{G}$ projects to an alternating cycle in G. Hence finding an alternating cycle in G is equivalent to finding an oriented cycle in $\tilde{G}$. The latter can be solved by a simple depth- or breath-first search.

REFERENCES

[Ale05] A. D. Alexandrov. *Convex polyhedra*. Springer Monographs in Mathematics. Springer-Verlag, Berlin, 2005. Translated from the 1950 Russian edition by N. S. Dairbekov, S. S. Kutateladze and A. B. Sossinsky, With comments and bibliography by V. A. Zalgaller and appendices by L. A. Shor and Yu. A. Volkov.

[And05] James W. Anderson. *Hyperbolic geometry. 2nd ed.* London: Springer, 2nd ed. edition, 2005.

[BBS11] Thierry Barbot, Francesco Bonsante, and Jean-Marc Schlenker. Collisions of particles in locally AdS spacetimes I. Local description and global examples. *Comm. Math. Phys.*, 308(1):147–200, 2011.

[Brü00] Max Brückner. Vielecke und Vielflache; Theorie und Geschichte., 1900.

[BS08] Alexander I. Bobenko and Yuri B. Suris. *Discrete differential geometry*, volume 98 of *Graduate Studies in Mathematics*. American Mathematical Society, Providence, RI, 2008.

[BV04] Stephen Boyd and Lieven Vandenberghe. *Convex optimization*. Cambridge University Press, 2004.

[Che03] Kevin King Hin Cheung. *Subtour elimination polytopes and graphs of inscribable type*. PhD thesis, University of Waterloo (Canada), 2003.

[CP17] Hao Chen and Arnau Padrol. Scribability problems for polytopes. *European J. Combin.*, 64:1–26, 2017.

[Dan13] Jeffrey Danciger. A geometric transition from hyperbolic to anti-de Sitter geometry. *Geom. Topol.*, 17(5):3077–3134, 2013.

[Dan14] Jeffrey Danciger. Ideal triangulations and geometric transitions. *J. Topol.*, 7(4):1118–1154, 2014.

[dGdV58] J. de Groot and H. de Vries. Convex sets in projective space. *Compositio Math.*, 13:113–118, 1958.

[DMS14] Jeffrey Danciger, Sara Maloni, and Jean-Marc Schlenker. Polyhedra inscribed in a quadric. 2014. Preprint, 42 pp., `arXiv:1410.3774`.

[DS96] Michael B. Dillencourt and Warren D. Smith. Graph-theoretical conditions for inscribability and Delaunay realizability. *Discrete Math.*, 161(1-3):63–77, 1996.

[Fil07] François Fillastre. Polyhedral realisation of hyperbolic metrics with conical singularities on compact surfaces. *Ann. Inst. Fourier (Grenoble)*, 57(1):163–195, 2007.

[Fil11] François Fillastre. Fuchsian polyhedra in Lorentzian space-forms. *Math. Ann.*, 350(2):417–453, 2011.

[Grü03] Branko Grünbaum. *Convex polytopes*, volume 221 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, second edition, 2003. Prepared and with a preface by Volker Kaibel, Victor Klee and Günter M. Ziegler.

[GS87] Branko Grünbaum and Geoffrey C. Shephard. Some problems on polyhedra. *J. Geom.*, 29(2):182–190, 1987.

[HRS92] Craig D. Hodgson, Igor Rivin, and Warren D. Smith. A characterization of convex hyperbolic polyhedra and of convex polyhedra inscribed in the sphere. *Bull. Amer. Math. Soc. (N.S.)*, 27(2):246–251, 1992.

[Izm09] Ivan Izmestiev. Projective background of the infinitesimal rigidity of frameworks. *Geom. Dedicata*, 140(1):183–203, 2009.

[Lee89] Carl W. Lee. The associahedron and triangulations of the n-gon. *Eur. J. Comb.*, 10(6):551–560, 1989.

[LS00] François Labourie and Jean-Marc Schlenker. Surfaces convexes fuchsiennes dans les espaces lorentziens à courbure constante. *Math. Ann.*, 316(3):465–483, 2000.

[Mar05] Xavier Martin. Piecewise-projective representation of Thompson's group T. In *Geometric methods in group theory*, volume 372 of *Contemp. Math.*, pages 161–180. Amer. Math. Soc., Providence, RI, 2005.

[Pen87] R. C. Penner. The decorated Teichmüller space of punctured surfaces. *Comm. Math. Phys.*, 113(2):299–339, 1987.

[Pen12] Robert C. Penner. *Decorated Teichmüller theory*. QGM Master Class Series. European Mathematical Society (EMS), Zürich, 2012. With a foreword by Yuri I. Manin.

[Pog73] Aleksei V. Pogorelov. *Extrinsic Geometry of Convex Surfaces*. American Mathematical Society, 1973. Translations of Mathematical Monographs. Vol. 35.

[Riv94] Igor Rivin. Intrinsic geometry of convex ideal polyhedra in hyperbolic 3-space. In *Analysis, algebra, and computers in mathematical research (Luleå, 1992)*, volume 156 of *Lecture Notes in Pure and Appl. Math.*, pages 275–291. Dekker, New York, 1994.

[Sch87] E. Schulte. Analogues of Steinitz's theorem about non-inscribable polytopes. In *Intuitive geometry (Siófok, 1985)*, volume 48 of *Colloq. Math. Soc. János Bolyai*, pages 503–516. North-Holland, Amsterdam, 1987.

[Sch98] Jean-Marc Schlenker. Métriques sur les polyèdres hyperboliques convexes. *J. Differential Geom.*, 48(2):323–405, 1998.

[Sch01] Jean-Marc Schlenker. Convex polyhedra in Lorentzian space-forms. *Asian J. Math.*, 5(2):327–363, 2001.

[Sch05] Jean-Marc Schlenker. A rigidity criterion for non-convex polyhedra. *Discr. Comput. Geom.*, 33(2):207–221, 2005. math.DG/0301333.

[Ste32] Jakob Steiner. *Systematische Entwicklung der Abhängigkeit geometrischer Gestalten von einander: mit Berücksichtigung der Arbeiten alter und neuer Geometer uber Porismen, Projections-Methoden, Geometrie der Lage, Transversalen, Dualität und Reciprocität, etc.* Fincke, 1832.

[Ste28] E. Steinitz. Über isoperimetrische Probleme bei konvexen Polyedern. *J. Reine Angew. Math.*, 159:133–143, 1928.

[Tod10] Takahisa Toda. Multi-convex sets in real projective spaces and their duality. 2010. Preprint, 21 pp., `arXiv:1005.1852`.

[Zie95] Günter M. Ziegler. *Lectures on polytopes*, volume 152 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 1995.

Georg-August-Universität Göttingen, Institut für Numerische und Angewandte Mathematik, Lotzestrasse 16–18, D-37083 Göttingen, Germany

Email address: `hao.chen.math@gmail.com`

University of Luxembourg, UR en Mathématiques, Maison du nombre, 6 avenue de la Fonte, L-4364 Esch-sur-Alzette, Luxembourg.

Email address: `jean-marc.schlenker@uni.lu`