RIGIDITY OF CIRCLE POLYHEDRA IN THE 2-SPHERE AND OF HYPERIDEAL POLYHEDRA IN HYPERBOLIC 3-SPACE

JOHN C. BOWERS, PHILIP L. BOWERS, AND KEVIN PRATT

ABSTRACT. We generalize Cauchy's celebrated theorem on the global rigidity of convex polyhedra in Euclidean 3-space $\mathbb{E}^3$ to the context of circle polyhedra in the 2-sphere $\mathbb{S}^2$. We prove that any two convex and proper non-unitary c-polyhedra with Möbius-congruent faces that are consistently oriented are Möbius-congruent. Our result implies the global rigidity of convex inversive distance circle packings in the Riemann sphere as well as that of certain hyperideal hyperbolic polyhedra in $\mathbb{H}^3$.

INTRODUCTION

Start with an abstract oriented triangulation of the 2-sphere $\mathbb{S}^2$ in which each edge is labeled with a non-negative real number. A *circle realization* of this triangulation is a drawing of circles on the sphere, one circle for each vertex, that respects orientation, and with the inversive distance between adjacent circles equal to the label for that edge. When connecting adjacent circle centers by geodesic arcs results in a geodesic triangulation of the 2-sphere, this circle realization is called an *inversive distance circle packing*.[1]

The question addressed in this paper is whether these inversive distance circle packings of $\mathbb{S}^2$, and more generally these circle realizations, are uniquely determined by the underlying triangulation and the inversive distances between adjacent circles. Of course, uniqueness is with respect to the group of orientation-preserving, circle-preserving transformations of the 2-sphere, the group of Möbius transformations Möb($\mathbb{S}^2$). In [8], Bowers and Stephenson questioned this uniqueness for general surfaces. In the case of packings on closed, orientable Euclidean or hyperbolic surfaces, Guo [12] verified their local rigidity, and later Luo [14] verified their global rigidity, as long as adjacent circles overlap in angles at most $\pi/2$.

[1]Perhaps a better term would be *circle pattern* instead of *packing* since adjacent circles may in fact be disjoint. We very soon will dispense with packings altogether and consider more general *circle frameworks* and, specifically, *circle polyhedra.*

Surprisingly, such packings are not unique on $\mathbb{S}^2$. Ma and Schlenker [15] produced counter-examples by constructing pairs of circle packings on $\mathbb{S}^2$ for the octahedral graph realizing the same inversive distance edge data, but for which there is no Möbius transformation taking one pattern onto the other. Their construction used an infinitesimally flexible Euclidean polyhedron, embeddings in de Sitter space $\mathbb{S}_1^3$, and the Pogorelov map between different geometries. In [4], the first two authors of the present work produced families of counter-examples inspired by Ma and Schlenker using only the inversive geometry of the sphere.

Ma and Schlenker's result was surprising particularly because the famous Koebe-Andre'ev-Thurston Circle Packing Theorem implies the uniqueness of such packings when all inversive distance edge labels are in the closed unit interval $[0, 1]$. In this case, adjacent circles always overlap, and in angles between 0 (tangent circles) and $\pi/2$ (orthogonal overlap). In contrast, our interest is in studying packings where adjacent circles may be disjoint, where inversive distances are greater than unity. Of particular interest to us are *edge-separated* packings in which all adjacent circles are disjoint. These packings always have *ortho-circles*, meaning that any three mutually adjacent circles have a unique ortho-circle, a circle orthogonal to all three. In this paper we will allow adjacent circles to overlap as long as this ortho-circle property is preserved. We even will allow adjacent circles to overlap in an angle greater than $\pi/2$ but we will not allow tangencies. Our study then is of *non-unitary circle packings with ortho-circles*, inversive distance circle packings for which all inversive distance edge labels take their values in $(-1, 1) \cup (1, \infty)$ and for which each face admits an ortho-circle. The constructions given in [4] show that in general such packings are not unique. In contrast to this, our main result shows that if we restrict ourselves to *convex* non-unitary circle packings with ortho-circles, then the Bowers–Stephenson question is once more answered in the affirmative.

Main Theorem. *Let $\mathcal{C}$ and $\mathcal{C}'$ be two non-unitary, inversive distance circle packings with ortho-circles for the same oriented edge-labeled triangulation of the 2-sphere $\mathbb{S}^2$. If $\mathcal{C}$ and $\mathcal{C}'$ are convex and proper, then there is a Möbius transformation $\mathsf{T} : \mathbb{S}^2 \to \mathbb{S}^2$ such that $\mathsf{T}(\mathcal{C}) = \mathcal{C}'$.*

Our interest really is in the generalization of circle packings that we call *c-frameworks* and the circle realizations of labeled graphs, and in particular in a generalization of 3-dimensional compact polyhedra to 'polyhedral' patterns of circles in the 2-sphere. This in fact is a major goal of this paper—to introduce these new concepts that generalize inversive distance circle packings to circle realizations, to define *c-polyhedra* (with the circle packings of the Main Theorem as but special examples of *c*-polyhedra), and to study the global rigidity of these *c*-polyhedra. Finding the right notion for generalizing boundedness for polyhedra, which we term *properness* for *c*-polyhedra, from Euclidean

geometry to circle geometry will take some effort. The most important ingredients for this study of rigidity are those of the convexity and the properness of c-polyhedra, and these definitions and the pursuit of the just right notions of convexity and properness will take up a significant portion of the paper in Section 1, *Preliminaries*, before we are able to prove this more general version of the Main Theorem.

Main Theorem (General). *Any two convex and proper non-unitary c-polyhedra with Möbius-congruent faces that are based on the same oriented abstract spherical polyhedron and are consistently oriented*[2] *are Möbius-congruent.*

The Main Theorem coupled with the Ma-Schlenker example of [15] and the examples of [4] show that the uniqueness of inversive distance circle packings, and more generally, of c-polyhedra is exactly analogous to that of Euclidean polyhedra—convex and bounded polyhedra in $\mathbb{E}^3$ are prescribed uniquely by their edge lengths and face angles whereas non-convex or unbounded polyhedra are not. The proof of this for convex and bounded Euclidean polyhedra is Cauchy's celebrated rigidity theorem [9], which is reviewed in Section 2. Our proof here follows Cauchy's original argument, which splits the proof into two components—a combinatorial lemma and a geometric lemma. Cauchy's combinatorial lemma deals with a certain labeling of the edges of any graph on a sphere, and applies to our setting. The geometric lemma, known as *Cauchy's Arm Lemma*, requires that a polygon with certain properties be defined for each vertex of the polyhedron, and fails to apply here. The main work of this paper is in describing and analyzing a family of hyperbolic polygons that we call *green-black polygons* that are defined for each vertex of a c-polyhedron in a Möbius-invariant manner. We develop an analog of Cauchy's Arm Lemma for convex green-black polygons in Section 3 and use it to prove the Main Theorem in Section 4.

Though our interest in the Bowers-Stephenson question arises from our study of inversive distance packings, a second motivation for studying this topic stems from the equivalence between inversive distance circle packings, or more generally c-polyhedra, and certain generalized hyperideal polyhedra in hyperbolic 3-space $\mathbb{H}^3$. Our Main Theorem implies that the hyperideal polyhedra associated to convex c-polyhedra are globally rigid. These hyperideal polyhedra are generalizations of those that Bao and Bonahon studied in [3].

Theorem. *In Klein's projective model for the hyperbolic space in which $\mathbb{H}^3$ is identified with the unit open ball B^3 in $\mathbb{E}^3 \subset \mathbb{RP}^3$, let P be the intersection with $\mathbb{H}^3$ of the compact, convex, polyhedron P' in $\mathbb{RP}^3$, all of whose vertices lie outside the closed unit ball $\mathbb{H}^3 \cup \partial\mathbb{H}^3 = \mathrm{B}^3 \cup \mathbb{S}^2$ and each of whose faces meets the open ball $\mathbb{H}^3$. If P is*

---

[2]Without the consistent orientation assumption, the c-polyhedra are inversive-congruent.

proper and non-unitary, then P *is globally rigid, unique up to isometries of* $\mathbb{H}^3$. *(See Section 5 for the definitions.)*

To say that P is *globally rigid* is to say that P is determined uniquely by the combinatorics of P$'$ and the hyperbolic isometry classes of the faces of the polyhedron P (which are determined by the faces of the polyhedron P* dual to P$'$; see Section 5). This means that combinatorics and isometry classes of faces determine uniquely the angles of intersection of adjacent faces that meet in $\mathbb{H}^3$, or their hyperbolic distance from one another if they do not meet in $\mathbb{H}^3$, again exactly analogous to bounded, convex Euclidean polyhedra.

This generalizes rigidity results of Bao-Bonahon [3] and Rousset [18] on hyperideal polyhedra that require that, either all edges of P$'$ meet the unit ball $\mathbb{H}^3$, or all edges of P$'$ lie "beyond infinity." Of course, Bao and Bonahon are able to prove more, namely, they give an existence result that characterizes in terms of dihedral angles and combinatorics those hyperideal polyhedra where all adjacent faces meet in $\mathbb{H}^3$. We do not provide such a characterization in our context, when some adjacent faces meet in $\mathbb{H}^3$ and others do not. The question of existence of these generalized hyperideal polyhedra is equivalent to the Bowers-Stephenson question of the existence of circle patterns with preassigned inversive distances between adjacent circles. When one allows for separated circles, this problem of existence becomes quite a bit more delicate than when adjacent circles are required to meet, with the possibility of even the local assignment of inversive distances about a central circle having no realization. These issues and the preceding theorem along with several corollaries, as well as a bit of background history of the rigidity of hyperbolic polyhedra, will be discussed in the final Section 5.

1. Preliminaries

Circle packings have been studied primarily in the cases where adjacent circles intersect nontrivially. When adjacent circles are allowed to be separated, with absolute inversive distances larger than unity, new difficulties arise from several directions. Existence of packings becomes problematic, even locally. When working on the 2-sphere, the facts that there are two complementary disks that a circle bounds and that the centers and radii of these disks fail to be Möbius invariants present difficulties in even defining the right notion of a Möbius-invariant circle packing. These difficulties suggest that the more natural setting for working with circle configurations on $\mathbb{S}^2$ with prescribed inversive distances is obtained by replacing circles by oriented circles, and circle packings by oriented circle realizations. We find that a more efficient language for oriented circle realizations on the 2-sphere than that of the traditional circle packing language is one obtained by adapting some of the

terminology of rigid Euclidean frameworks or linkages to the setting of circle configurations in $\mathbb{S}^2$.[3] In this rather lengthy preliminary section, we recall the necessary properties from the inversive geometry of the 2-sphere needed for our result, we introduce oriented circle realizations and circle polyhedra, we describe appropriate notions of convexity and properness for circle polyhedra, and we review some of and develop further the elementary geometry of hyperideal hyperbolic polygons and polyhedra.

1.1. **Inversive geometry of the Riemann sphere.** We use the two usual models for the Riemann sphere, the round unit sphere $\mathbb{S}^2$ in Euclidean 3-space $\mathbb{E}^3$ and its image under stereographic projection, the extended complex plane $\widehat{\mathbb{C}} = \mathbb{C} \cup \{\infty\}$. $\mathrm{Inv}(\mathbb{S}^2)$ denotes the inversive group of the 2-sphere generated by reflections in the circles of $\mathbb{S}^2$, and $\mathrm{M\ddot{o}b}(\mathbb{S}^2)$ its index two subgroup of Möbius transformations. The absolute inversive distance between two circles in $\mathbb{S}^2$ is a Möbius invariant of the placement of the two circles in the 2-sphere. More useful for us is the general inversive distance that serves as an invariant for the placement of two relatively oriented circles in the 2-sphere. We give two equivalent definitions, the first using the algebra of the extended plane $\widehat{\mathbb{C}}$ and the second using the intrinsic spherical metric of the 2-sphere $\mathbb{S}^2$. Each has its advantages, as becomes apparent as the discussion advances.

For the definitions, note that an oriented circle determines a unique closed *companion* or *spanning disk* that the circle bounds. Indeed, assuming fixed orientations for $\mathbb{S}^2$ and $\widehat{\mathbb{C}}$ that are compatible via stereographic projection, the companion disk determined by the oriented circle C is the closed complementary disk D (of the two available) whose positively oriented boundary $\partial^+ D = C$, where of coures the orientation of D is inherited from that of $\mathbb{S}^2$ or $\widehat{\mathbb{C}}$. This is described colloquially by saying that D lies to the left of C as one traverses C along the direction of its orientation.

Definition (General inversive distance). Let C_1 and C_2 be oriented circles in the extended plane $\widehat{\mathbb{C}}$ bounding their respective companion disks D_1 and D_2, and let C be any oriented circle mutually orthogonal to C_1 and C_2. Denote the points of intersection of C with C_1 as z_1, z_2 ordered so that the oriented sub-arc of C from z_1 to z_2 lies in the disk D_1. Similarly denote the ordered points of intersection of C with D_2 as w_1, w_2. The *general inversive distance* between C_1 and C_2, denoted as $\langle C_1, C_2 \rangle$, is defined in terms of the cross ratio

$$[z_1, z_2; w_1, w_2] = \frac{(z_1 - w_1)(z_2 - w_2)}{(z_1 - z_2)(w_1 - w_2)}$$

by

$$\langle C_1, C_2 \rangle = 2[z_1, z_2; w_1, w_2] - 1.$$

---

[3]The first two authors expand upon this point of view and present a general framework for studying the rigidity of circle realizations in [6].
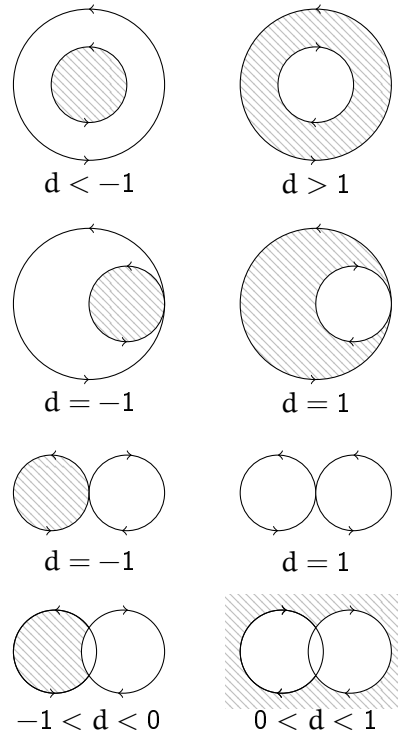
FIGURE 1. Inversive distances $d = \langle C_1, C_2 \rangle$. The shaded regions are the intersections $D_1 \cap D_2$, the points common to the spanning disks D_1 and D_2 for both circles C_1 and C_2.

Subsequently, we drop the adjective *general* and refer to the inversive distance $\langle C_1, C_2 \rangle$ with its absolute value $|\langle C_1, C_2 \rangle|$ the *absolute inversive distance*.[4]

Recall that cross ratios of ordered 4-tuples of points in $\widehat{\mathbb{C}}$ are invariant under Möbius transformations and that there is a Möbius transformation taking an ordered set of four points of $\widehat{\mathbb{C}}$ to another ordered set of four if and only if the cross ratios of the sets agree. This implies that which circle C orthogonal to both C_1 and C_2 is used in the definition is irrelevant as a Möbius transformation that set-wise fixes C_1 and C_2 can be used to move any one orthogonal circle to another. Which one of the two orientations on the orthogonal circle C is used is irrelevant as the cross ratio satisfies $[z_1, z_2; w_1, w_2] = [z_2, z_1; w_2, w_1]$.

---

[4]The second author first learned of defining inversive distance in this way from his student, Roger Vogeler. He has looked for this in the literature and, unable to find it can only surmise that it is original with Prof. Vogeler. The definition appeared in [7] in 2003.

This equation also shows that the inversive distance is preserved when the orientation of both circles is reversed so that it is only the relative orientation of the two circles that is important for the definition. In fact, the general inversive distance is a relative conformal measure of the placement of an oriented circle pair on the Riemann sphere. By this we mean that two oriented circle pairs are inversive equivalent if and only if their inversive distances agree. All of this should cause one to pause to develop some intuition about how companion disks may overlap with various values of inversive distances. See Fig. 1 for some corrections to possible misconceptions. Finally, the inversive distance is symmetric with $\langle C_1, C_2 \rangle = \langle C_2, C_1 \rangle$ since $[z_1, z_2; w_1, w_2] = [w_1, w_2; z_1, z_2]$.

The inversive distance is real since the cross ratio of points lying on a common circle is real and, in fact, every real value is realized as the inversive distance of some oriented circle pair. It is easy to give an intuitive understanding of the distance when C_1 and C_2 meet. In this case, the oriented angle α of overlap may be defined unambiguously as the angle between the unit tangent vectors to the circles at a point of overlap formed by one tangent pointing along the orientation of its parent circle and the other pointing against the orientation of its parent circle. This is precisely the *lune angle* formed by the *lune* $D_1 \cap D_2$. The inversive distance is then $\langle C_1, C_2 \rangle = \cos \alpha$. In particular, when the circles are tangent, the inversive distance is ± 1, and when orthogonal, the inversive distance is zero. When the circles do not meet, the inversive distance lies outside the interval $[-1, 1]$ and this is the range where the definitions in terms of hyperbolic, Euclidean, and spherical geometry give more geometric insight. We will give the spherical definition next, but we comment first that the general inversive distance is less than -1 if and only if one of the disks D_1 and D_2 is contained in the interior of the other.[5] We refer the reader to [7] for a more detailed treatment. Notice that if the orientation of only one member of a circle pair is reversed, the inversive distance merely changes sign. This follows from the immediate relation $[z_1, z_2; w_2, w_1] = 1 - [z_1, z_2; w_1, w_2]$. Despite its name, the inversive distance is not a metric as it fails to be non-negative and fails to satisfy the triangle inequality.[6]

The second definition is entirely in terms of the spherical metric.

Definition (Inversive distance in the spherical metric). In the 2-sphere $\mathbb{S}^2$, the inversive distance may be expressed as[7]

$$(1.1) \qquad \langle C_1, C_2 \rangle = \frac{-\cos \sphericalangle (p_1, p_2) + \cos(r_1)\cos(r_2)}{\sin(r_1)\sin(r_2)} = \frac{-p_1 \cdot p_2 + \cos(r_1)\cos(r_2)}{\sin(r_1)\sin(r_2)}.$$

---

[5] We should mention that some of the older literature use $-\langle C_1, C_2 \rangle$ for the inversive distance; for example [19].

[6] Some authors, perhaps more aptly, call the inversive *distance* the inversive *product* of C_1 and C_2.

[7] In both Luo [14] and Ma-Schlenker [15] there is a typo in the expression for the spherical formula for inversive distance. They report the negative of this formula.

Here, $\sphericalangle(p_1, p_2) = \cos^{-1}(p_1 \cdot p_2)$ denotes the spherical distance between the centers, p_1 and p_2, of the respective companion disks, $p_1 \cdot p_2$ the usual Euclidean inner product between the unit vectors p_1 and p_2, and r_1 and r_2 the respective spherical radii of the companion disks. Note that $r_i = \cos^{-1}(p_i \cdot q_i)$ for any point q_i on the circle C_i, for $i = 1, 2$.

Verifying the equivalence of the two definitions is an exercise in the use of trigonometric identities after a standard placement of C_1 and C_2 on $\mathbb{S}^2$ followed by stereographic projection. This standard placement is obtained by finding the unique great circle C orthogonal to both C_1 and C_2 and then rotating the sphere so that this great circle is the equator, which then stereographically projects to the unit circle in the complex plane. We leave the details to the reader. In Section 1.4 we will recall the standard expression for inversive distance greater than unity in terms of hyperbolic geometry.

1.2. **Circle packings, frameworks, and realizations.** We here generalize the language of circle packing and patterns of triangulations and quadrangulations of the sphere to that of circle realizations of oriented circle frameworks.[8] Let G be a graph, by which we mean a set of vertices $V = V(G)$ and simple edges $E = E(G)$. We disallow both loops and multiple edges. An oriented edge incident to the initial vertex u and terminal vertex v is denoted as uv, and $-uv$ means the oppositely oriented edge vu. We use the same notation, uv, to denote an un-oriented edge, context making the meaning clear. A *circle framework with adjacency graph* G, or *c-framework* for short, is a collection $\mathcal{C} = \{C_u : u \in V(G)\}$ of oriented circles in $\mathbb{S}^2$ indexed by the vertex set of G. We denote this by $G(\mathcal{C})$. Two *c*-frameworks $G(\mathcal{C})$ and $G(\mathcal{C}')$ are *equivalent* if $\langle C_u, C_v \rangle = \langle C_u', C_v' \rangle$ whenever uv is an edge of G. Let H be a subgroup of $\mathrm{Inv}(\mathbb{S}^2)$. Two collections $\mathcal{C}$ and $\mathcal{C}'$ of oriented circles indexed by the same set are H-*equivalent* or H-*congruent* provided there is a mapping $T \in H$ such that $T(\mathcal{C}) = \mathcal{C}'$, respecting the common indexing and the orientations of the circles. When H is not so important they are *inversive-equivalent* or *inversive-congruent*, and when T can be chosen to be a Möbius transformation, they are *Möbius-equivalent* or *Möbius-congruent*. The global rigidity theory of *c*-frameworks concerns conditions on G or $G(\mathcal{C})$ that ensure that the equivalence of the *c*-frameworks $G(\mathcal{C})$ and $G(\mathcal{C}')$ guarantees their H-equivalence. Often we restrict our attention to *c*-frameworks in a restricted collection $\mathscr{F}$ of *c*-frameworks. In this paper, $\mathscr{F}$ is the collection of non-unitary, convex and proper *c*-polyhedra, defined subsequently, and our interest is in Möbius equivalence.

We develop next some terminology to describe how much separation there is between two oriented circles. Two oriented circles C_1 and C_2 are *coupled* if one of the circles is contained in the companion disk of the other circle. This may occur in two distinct ways. It may

---

[8]The material in this section is extended and further developed in [6].

be that one of the companion disks is contained in the other. This occurs exactly when $\langle C_1, C_2 \rangle \leqslant -1$, and in this case the pair remains coupled when both their orientations are reversed. On the other hand, it may be that neither companion disk is contained in the other as in the first row of Fig. 1. In this case $\langle C_1, C_2 \rangle \geqslant 1$ and the two circles become uncoupled when both orientations are reversed. It follows then that C_1 and C_2 are *uncoupled* precisely when either $C_1 \cap C_2$ has two points, or the interiors of their companion disks are disjoint. If C_1 and C_2 are uncoupled, then $\langle C_1, C_2 \rangle > -1$. A stronger form of separation is that of segregation. The oriented circles C_1 and C_2 are *segregated* if they they are uncoupled and *overlap by at most* $\pi/2$. In particular, either the companion disks are disjoint, or when not disjoint, then the two bounding circles meet in an oriented angle of at most $\pi/2$. In this case $\langle C_1, C_2 \rangle \geqslant 0$. Even stronger, the circles are *separated* if the companion disks are disjoint, and in this case $\langle C_1, C_2 \rangle > 1$. Notice that C_1 and C_2 are uncoupled if segregated, and segregated if separated. A collection of oriented circles is respectively uncoupled, segregated, or separated if every distinct pair of the collection is uncoupled, segregated, or separated.

The oriented circles C_u and C_v of a c-framework $G(\mathcal{C})$ are *adjacent* provided uv is an edge of G. The c-framework $G(\mathcal{C})$

   (i) is *edge-uncoupled* if each pair of adjacent circles is uncoupled (implying that $\langle C_u, C_v \rangle > -1$ for all edges uv of G);

  (ii) is *edge-segregated* if each pair of adjacent circles is segregated so that the companion disks of all adjacent circles overlap by at most $\pi/2$ (implying that $\langle C_u, C_v \rangle \geqslant 0$ for all edges uv of G);

 (iii) is *edge-separated* if each pair of adjacent circles is separated so that the companion disks of all adjacent circles are disjoint (implying that $\langle C_u, C_v \rangle > 1$ for all edges uv of G);

  (iv) is *non-unitary* if $\langle C_u, C_v \rangle \neq \pm 1$ for all edges uv of G (companion disks of adjacent circles are not tangent);

   (v) has *deep overlaps* if there is at least one edge uv for which $\langle C_u, C_v \rangle < 0$ (implying that it is not edge-segregated).

Definition (Labeled graph and circle realization). An *edge-label* is a real-valued function $\beta : E(G) \to \mathbb{R}$ defined on the edge set of G, and G together with an edge-label β is denoted as G_β and called an *edge-labeled graph*. The c-framework $G(\mathcal{C})$ is a *circle realization* of the edge-labeled graph G_β provided $\langle C_u, C_v \rangle = \beta(uv)$ for every edge uv of G, and we denote it as $G_\beta(\mathcal{C})$.

Circle packings are circle realizations of edge-labeled graphs that arise as the 1-skeletons of oriented triangulations of the 2-sphere that also satisfy certain properties that ensure that the realizations of the triangular boundaries of faces respect orientation. The general definition allows for branch vertices and configurations of circles in which the open geodesic triangles cut out by connecting centers of adjacent circles overlap. There are subtleties in which we have no interest, so we adapt a restricted definition that corresponds to the circle packings that arise from spherical polyhedral metrics on triangulated surfaces. These are circle realizations of the edge-labeled 1-skeleton $G_\beta = K_\beta^{(1)}$ of an oriented triangulation K of $\mathbb{S}^2$ that produce oriented geodesic triangulations[9] of the 2-sphere when adjacent circle centers are connected by geodesic arcs. The assumption here is that the centers of no two adjacent circles are antipodal, so that there is a unique geodesic arc connecting them, and that the centers of three circles corresponding to the vertices of a face of K do not lie on a great circle. Now this causes no particular problems when all adjacent circles overlap nontrivially, the traditional playing field of circle packing, but does cause some real concern when adjacent circles may have inversive distance greater than unity. For example, a circle realization $\mathcal{C}$ may produce a geodesic triangulation of the sphere by connecting adjacent centers while its Möbius image $T(\mathcal{C})$ may not. This is traced directly to the fact that neither circle centers nor radii, nor geodesic arcs, are Möbius invariants in the inversive geometry of the sphere. This behavior does not occur for inversive distance circle packings of the Euclidean or hyperbolic planes (and surfaces), precisely because circle centers and geodesics are invariant under automorphisms and radii are invariant up to scale in Euclidean geometry and invariant in hyperbolic geometry. Our belief is that using centers and radii of circles in inversive geometry should be avoided except where these can be used to simplify computations (as in the use of the spherical definition of inversive distance). Our shift in this paper then is from inversive distance circle packings to inversive distance circle realizations. We are less concerned with possible underlying geodesic triangulations and more concerned with Möbius-invariant quantities. For example, rather than working with a geodesic face formed by connecting the centers of three mutually adjacent circles, we are more interested in the existence of an ortho-circle, a circle mutually orthogonal to the three, which is a Möbius invariant. Though our initial motivation was circle packing as reflected in the Main Theorem, our real interest has evolved to circle realizations as reflected in the general version of the Main Theorem.

1.3. **Circle polyhedra.** A precise definition of a *circle polyhedron*, or *c-polyhedron* for short, is given subsequently, but it is, essentially, a *c*-framework $G(\mathcal{C})$ where $G = P^{(1)}$, the 1-skeleton of a Euclidean 3-dimensional polyhedron P. There are conditions imposed

---

[9]By this we mean that the orientation of the geodesic triangulation determined by the packing is consistent with the orientation on K.

that correspond to the fact that the faces of a Euclidean polyhedron are planar polygons. To describe this we borrow from the incidence geometry of the space of circles. It turns out that the space of all oriented circles on the 2-sphere is topologically a twice punctured 3-sphere homeomorphic to $\mathbb{S}^2 \times \mathbb{R}$ that double covers the space of un-oriented circles, topologically a punctured projective space $\mathbb{RP}^{3*}$. What is more significant is that there is a natural incidence geometry of points, lines, and planes in each of these spaces of circles that mimics the incidence geometry of Euclidean space $\mathbb{E}^3$ in some significant ways. This is developed rather fully in [5] and put to use in studying iso-inversive embeddings. We need not develop these incidence geometries here, but we do want to borrow one of the types of planes from the incidence geometry of circle space to use as a natural notion of a planar set of circles in $\mathbb{S}^2$. These are called *hyperbolic c-planes* in [5] (as opposed to *parabolic* and *elliptic c-planes*) and, in the older literature, are examples of bundles of circles. We will call these *c-planes*. A c-plane Π_O is determined by its *generating circle* O, an un-oriented circle in $\mathbb{S}^2$, and is defined as the collection of all oriented circles in $\mathbb{S}^2$ that meet O orthogonally, or what is the same, all circles C for which $\langle C, O \rangle = 0$ when O is given either orientation. Any collection $\mathcal{C}$ of oriented circles is said to be *c-planar* provided all the circles of the collection belong to a common *c-plane*, or what is the same, all the circles of $\mathcal{C}$ have a common *ortho-circle*, a circle O orthogonal to each circle of $\mathcal{C}$. When the c-planar collection $\mathcal{C}$ has at least three distinct elements that are not coaxial, the collection has a unique ortho-circle O determining the unique *spanning plane* Π_O of $\mathcal{C}$.[10]

An *abstract polyhedral graph* G is the 1-skeleton of an *abstract spherical polyhedron* P, by which we mean an oriented cellular decomposition of the 2-sphere into a finite number of combinatorial polygons[11] for which distinct polygonal faces are either disjoint, or meet at a set of vertices and full edges.[12] We can think of P as set of topological polygons glued together along full edges to form a topological sphere. Writing $P = (V, E, F)$, where $V = V(P)$ is the finite set of vertices, $E = E(P)$ is the finite set of edges, and $F = F(P)$ is the finite set of combinatorial polygons, we have $G = (V, E)$. Each n-gon face $f \in F$ may be denoted by listing its vertices in cyclic order as $f = v_1 v_2 \ldots v_n$, respecting the orientation.

---

[10]We should point out that three distinct circles need not have a common ortho-circle, and therefore need not lie on a c-plane as defined here. Nonetheless, three distinct circles that are not coaxial always lie on a unique c-plane as defined in [5], as the three may determine a *parabolic* or *elliptic* c-plane instead of a *hyperbolic* one.

[11]We do not allow bi-gons, only n-gons for $n \geqslant 3$.

[12]In the literature a *polyhedral graph* is one isomorphic to the 1-skeleton of a convex polyhedron in $\mathbb{E}^3$. By Steinitz's Theorem, these are characterized as precisely the 3-vertex-connected planar graphs. Our abstract polyhedral graphs are more general in that we allow for 1-skeletons of non-convex polyhedra, for example, or for two faces to share two or more non-contiguous edges.

Notice that if two faces share the two adjacent vertices u and v, then one face determines the oriented edge uv, while the other, the oriented edge vu.

Definition (Circle polyhedron). Let $G = P^{(1)}$ be an abstract polyhedral graph where P is an oriented abstract spherical polyhedron. The *c-framework* $G(\mathcal{C})$ is called a *circle polyhedron based on* P, or a *c-polyhedron* for short, provided it is edge-uncoupled and, for each n-gon face $f = u_1 \ldots u_n$ of P, the *c-face* $\mathcal{C}_f = \{C_{u_i} : i = 1, \ldots, n\}$ is not coaxial but is *c-planar*. We denote the unique ortho-circle for $\mathcal{C}_f$ as O_f.

The primary interest of this paper is in the global rigidity of *c*-polyhedra. The Ma-Schlenker example of [15] and the Ma-Schlenker *c*-octahedra constructed in [4] show that in general *c*-polyhedra are not globally rigid with respect to the inversive group of the 2-sphere. We next define a notion of convexity for *c*-polyhedra. The Ma-Schlenker *c*-octahedra are not convex and our Main Theorem shows it is this lack of convexity that allows for the non-rigidity of these examples.

Definition (Convex *c*-polyhedron). Let $G(\mathcal{C})$ be a *c*-polyhedron based on the abstract spherical polyhedron P with face set $F = F(P)$. For any face $f \in F$ we say that $G(\mathcal{C})$ is *convex with respect to* f provided the ortho-circle O_f may be oriented so that every circle $C_u \in \mathcal{C}$ is segregated from the oriented ortho-circle O_f. $G(\mathcal{C})$ then is *convex* provided it is convex with respect to each of its faces so that the ortho-circles O_f for $f \in F$ may be oriented so that every circle $C_u \in \mathcal{C}$ is segregated from every oriented ortho-circle O_f for $f \in F$. The oriented ortho-circles for $f \in F$ that meet this condition are denoted as O_f^+. To avoid unnecessary pathology, we make the further assumption that the circles corresponding to three consecutive vertices on the face f are never coaxial.[13]

We have stated the Main Theorem in terms of Möbius rigidity instead of inversive rigidity. For this we need to impose a condition to make sure that two *c*-polyhedra based on the same abstract spherical polyhedron P are oriented consistently, both with the orientation of P. We can insure this with the application of the antipodal map if needed.

Proposition 1.1. *Let* $G(\mathcal{C})$ *be a convex c-polyhedron based on* P. *The either (i) for every oriented n-gon face $f = u_1 \ldots u_n$ of P, the circles $C_{u_1}, \ldots, C_{u_n}$ are met in that order as one progresses around O_f^+ in the direction of its orientation, starting at C_{u_1}; or (ii) for every oriented n-gon face $f = u_1 \ldots u_n$ of P, the circles $C_{u_1}, \ldots, C_{u_n}$ are met in that order as one progresses around O_f^+ in the direction opposite of its orientation, starting at C_{u_1}.*

---

[13]This is the analogue with convex Euclidean polyhedra of avoiding three vertices lying on a line. The middle vertex is unnecessary.

By replacing the c-polyhedron $G(\mathcal{C})$ by its image $G(\mathcal{C}^*)$ under the antipodal mapping of the sphere $\mathbb{S}^2$ if necessary, and then reversing all circle orientations, we may assume that property (i) of Proposition 1.1 holds. We then say that the c-polyhedron $G(\mathcal{C})$ is *oriented consistently* with the orientation of the base polyhedron P for $G(\mathcal{C})$, which will be assumed without comment in the remainder of the paper. Without the assumption of consistent orientation, we would need to work in terms of inversive rigidity, allowing for an orientation reversing map to match orientations before applying a Möbius transformation.

Before pursuing the definition of properness for c-polyhedra, we pause to develop prerequisite terminology for hyperideal polygons in the hyperbolic plane.

1.4. **Hyperideal polygons and complex angles.** We assume a fixed orientation on the hyperbolic plane $\mathbb{H}^2$, and when using an open disk D in $\mathbb{S}^2$ or $\widehat{\mathbb{C}}$ as a Poincaré-disk model of $\mathbb{H}^2$, we assume the orientation of $\mathbb{H}^2 = D$ is inherited from that of $\mathbb{S}^2$. Let ℓ_1 and ℓ_2 be oriented lines in the hyperbolic plane $\mathbb{H}^2$. We define the *complex angle* $\theta = \theta(\ell_1, \ell_2)$ between ℓ_1 and ℓ_2. When the lines meet at the point p, $\theta = \theta(\ell_1, \ell_2)$ is the angle between the unit tangent vectors to the lines at p formed by one tangent pointing along the orientation of its parent line and the other pointing against the orientation of its parent line. This is precisely the *lune angle* formed by the *lune* $h_1 \cap h_2$, where h_i is the oriented half-plane with $\partial^+ h_i = \ell_i$, the half-plane whose positively oriented boundary is ℓ_i, for $i = 1, 2$. When ℓ_1 and ℓ_2 are parallel, meeting at an ideal point $p \in \partial \mathbb{H}^2$, the angle $\theta = \theta(\ell_1, \ell_2)$ is either zero or π, depending on the relative orientations. When ℓ_1 and ℓ_2 are ultra-parallel, there is no point of intersection, ideal or finite. When the orientations are *consistent* with one another, meaning that neither of h_1 and h_2 is contained in the other, we define $\theta(\ell_1, \ell_2) = i\, d_{\mathbb{H}^2}(\ell_1, \ell_2)$, this also called the *imaginary angle* between the lines. Of course i is the imaginary unit and $d_{\mathbb{H}^2}(\ell_1, \ell_2)$ is the hyperbolic distance between the lines, the length of the unique geodesic arc orthogonal to both ℓ_1 and ℓ_2, marked as the green dotted lines of Fig. 2b. When the orientations are inconsistent, the complex angle is a *phase shift* of the imaginary angle by π, by which we mean that $\theta(\ell_1, \ell_2) = \pi + i d_{\mathbb{H}^2}(\ell_1, \ell_2)$. In a Poincaré disk model of the hyperbolic plane (which includes the upper-half-plane model), for $i = 1, 2$, ℓ_i is the intersection with the disk of an oriented circle C_i in $\mathbb{C}$ that meets the boundary of the disk orthogonally. In this case, the inversive distance satisfies $\langle C_1, C_2 \rangle = \cos \theta(\ell_1, \ell_2)$. Notice that the range of complex angles is the curve $\Theta = i\mathbb{R}_+ \cup [0, \pi] \cup (\pi + i\mathbb{R}_+)$ in the complex plane[14] on which the complex cosine function cos is a homeomorphism onto the real line $\mathbb{R}$. When we use the inverse function $\cos^{-1}$, we mean the inverse of this homeomorphism.

---

[14] $\mathbb{R}_\pm = \{x \in \mathbb{R} : \pm x > 0\}$, respectively the sets of positive and negative real numbers.

(A) The convex hyperideal polygon defined by six hyperbolic lines.

(B) The green dotted segments denote the hyperideal vertices, the unique shortest paths between consecutive non-intersecting lines.

FIGURE 2. A convex hyperideal 6-gon.

Let $H = \{h_1, \ldots, h_n\}$ be a set of half-planes in $\mathbb{H}^2$ such that the region $P = h_1 \cap \cdots \cap h_n$ is non-empty and the boundary line ℓ_i of each h_i supports a non-empty segment, ray, or line on the boundary of P. The lines ℓ_i for $i = 1, \ldots, n$ are oriented consistent with the orientation P inherits from $\mathbb{H}^2$ so that $\partial^+ h_i = \ell_i$. If P is bounded, then it is a compact convex polygon in $\mathbb{H}^2$. Whether or not P is bounded, we call P a *convex hyperideal polygon*. As is usual, we sometimes refer to the oriented polygonal boundary $\partial^+ P$ as a convex hyperideal polygon. An example is shown in Fig. 2a. Suppose that the half-planes are indexed so that $\ell_1, \ldots, \ell_n$ are in cyclic counter-clockwise order along P. Two consecutive lines ℓ_{i-1}, ℓ_i may intersect either at a point in $\mathbb{H}^2$ forming a *finite vertex* of P, or at an ideal point at infinity forming an *ideal vertex*, or not at all. In the latter case, the hyperbolic segment orthogonal to the two consecutive lines is called the *hyperideal vertex* of P determined by ℓ_{i-1} and ℓ_i. The complex angle $\theta_i = \theta(\ell_{i-1}, \ell_i)$ is the *complex interior angle* of the polygon P determined by ℓ_{i-1} and ℓ_i; see Fig. 2b . In Section 1.5 we use hyperideal polygons to construct *green-black* polygons that then are used in Section 3 to generalize the Cauchy Arm Lemma for use in the proof of the Main Theorem.

Moving up one dimension, the *complex dihedral angle* between two oriented hyperbolic planes in the hyperbolic 3-space $\mathbb{H}^3$ is defined analogously to the complex angle between oriented lines. When the oriented planes h_1 and h_2 meet, $\theta = \theta(h_1, h_2)$ is the dihedral angle between them, the angle between their normal vectors at a point of intersection, one along the orientation of the plane and the other against the orientation. When the planes meet at an ideal point, the complex dihedral angle is zero or π depending on relative orientations. When the planes are disjoint, the complex dihedral angle between them is the imaginary angle $\theta = \theta(h_1, h_2) = i\, d_{\mathbb{H}^3}(h_1, h_2)$, or this phase-shifted by π, where $d_{\mathbb{H}^3}(h_1, h_2)$ is the hyperbolic distance between the planes, the length of the unique

geodesic arc orthogonal to both h_1 and h_2. The presence or absence of the phase shift depends on relative orientations. In either the Poincaré ball or the upper-half-space model of $\mathbb{H}^3$, the oriented boundaries of the planes h_1 and h_2 are oriented circles C_1 and C_2 on the sphere at infinity, and $\langle C_1, C_2 \rangle = \cos \theta(h_1, h_2)$. This motivates the next definition.

Definition (Complex dihedral angle). Let $G(\mathcal{C})$ be a convex c-polyhedron based on the abstract spherical polyhedron P with face set $F = F(P)$. For each pair of faces $f, g \in F$ that share an edge, define the *complex dihedral angle* between the *adjacent c-faces* $\mathcal{C}_f$ and $\mathcal{C}_g$ at that edge to be $\cos^{-1} \langle O_f^+, O_g^+ \rangle$.

By the preceding discussion, the complex dihedral angle between the adjacent c-faces $\mathcal{C}_f$ and $\mathcal{C}_g$ is the complex dihedral angle between the oriented hyperbolic planes h_f and h_g in the Poincaré ball model of hyperbolic 3-space, where h_f is the oriented hyperbolic plane in the ball $B^3 = \mathbb{H}^3$ whose oriented boundary at infinity is O_f^+.

Theorem 1.2. *Let $G(\mathcal{C})$ and $G(\mathcal{C}')$ be two convex non-unitary c-polyhedra both based on the abstract spherical polyhedron P. Then $G(\mathcal{C})$ and $G(\mathcal{C}')$ are Möbius-congruent with one another if and only if the following conditions hold.*

(i) *$G(\mathcal{C})$ and $G(\mathcal{C}')$ have Möbius-congruent c-faces. This means that for each face f of P, there is an Möbius transformation $T \in M\ddot{o}b(\mathbb{S}^2)$ such that $T(\mathcal{C}_f) = \mathcal{C}'_f$, taking corresponding oriented circles to corresponding oriented circles.*

(ii) *Corresponding complex dihedral angles between adjacent c-faces agree; equivalently, for each pair of adjacent faces $f, g \in F(P)$, $\langle O_f^+, O_g^+ \rangle = \langle O'^+_f, O'^+_g \rangle$, where the prime denotes the oriented ortho-circles for $G(\mathcal{C}')$.*

Proof. The forward direction is obvious. For the reverse direction it is enough to verify the following: If $f, g \in F(P)$ are faces that share an edge uv and $T \in M\ddot{o}b(\mathbb{S}^2)$ is a Möbius transformation such that $T(\mathcal{C}_f) = \mathcal{C}'_f$, taking corresponding oriented circles to corresponding oriented circles, then $T(\mathcal{C}_g) = \mathcal{C}'_g$, taking corresponding oriented circles to corresponding oriented circles. To verify this, let f, g, uv, and T be as in the preceding sentence and observe that $T(O_f^+) = O'^+_f$. This uses property (i) of Proposition 1.1 to ensure that T preserves the orientation on the ortho-circles. The circles C'_u and C'_v belong to a unique coaxial family $\mathcal{A}$ of circles in $\mathbb{S}^2$ and the ortho-circles O'_f and O'_g belong to its orthogonal complement $\mathcal{A}^\perp$, the collection of circles orthogonal to the circles in the family $\mathcal{A}$. Now $T(O_f^+) = O'^+_f$ and the question we ask is what of the image $T(O_g^+)$: does $T(O_g^+) = O'^+_g$? Note first that $T(O_g)$, since it is orthogonal to both C'_u and C'_v, belongs to the family $\mathcal{A}^\perp$. Since the c-polyhedra are non-unitary, the coaxial family $\mathcal{A}$ is not parabolic, and hence neither is $\mathcal{A}^\perp$. This implies that one of $\mathcal{A}$ and $\mathcal{A}^\perp$ is elliptic and the other hyperbolic.

The fact we use here is that if one chooses any two distinct circles O and A in an elliptic or hyperbolic coaxial family and orients them, say to obtain the oriented circles O^+ and A^+, then there is exactly one circle B distinct from O and A in that same coaxial family that admits an orientation to give an oriented circle B^+ so that the inversive distances satisfy $\langle O^+, A^+ \rangle = \langle O^+, B^+ \rangle$. Of course then also $\langle O^+, A^- \rangle = \langle O^+, B^- \rangle$, where the minus-superscript denotes the opposite orientation. By this observation, there is precisely one oriented circle B^+ other than $A^+ = O_g'^{\,+}$ of the orthogonal complement $\mathcal{A}^\perp$ whose inversive distance to $T(O_f^+) = O_f'^{\,+}$ is $\langle O_f'^{\,+}, O_g'^{\,+} \rangle$. It follows that $T(O_g^+)$ is equal to either B^+ or to $O_g'^{\,+}$, and our claim is that the convexity of the c-polyhedra $G(\mathcal{C})$ and $G(\mathcal{C}')$ precludes B^+ as the image of O_g^+. We argue by contradiction. Suppose that $T(O_g'^{\,+}) = B^+$. Let C be an oriented circle in the c-face $\mathcal{C}_f$ that is not in the coaxial family determined by C_u and C_v so that its image $C' = T(C) \notin \mathcal{A}$. Since $G(\mathcal{C})$ is convex, C is segregated from O_g^+ and therefore its image $T(C) = C'$ is segregated from the image $T(O_g^+) = B^+$. But also since $T(\mathcal{C}_f) = \mathcal{C}_f'$, $C' = T(C)$ is an oriented circle in $G(\mathcal{C}')$ and since the c-polyhedron $G(\mathcal{C}')$ is convex, C' is segregated from $O_g'^{\,+}$. Thus the oriented circle $C' = T(C)$ is orthogonal to $O_f'^{\,+}$ and segregated from both B^+ and $A^+ = O_g'^{\,+}$. Since C is orthogonal to the ortho-circle $O_f'^{\,+}$, this violates Lemma 1.3 that is stated and proved below.

We have verified that $T(O_g^+) = O_g'^{\,+}$. Now let $T' \in \text{Möb}(\mathbb{S}^2)$ be a Möbius transformation such that $T'(\mathcal{C}_g) = \mathcal{C}_g'$, taking corresponding oriented circles to corresponding oriented circles. Let a and b be the points of intersection of C_u with the ortho-circle O_g and c and d the points of intersection of C_v with O_g. Since the maps T and T' both take C_u to C_u', C_v to C_v', and O_g^+ to $O_g'^{\,+}$, preserving all orientations, they must agree on the four points a, b, c, and d. But then $T = T'$ and we have $T(\mathcal{C}_g) = T'(\mathcal{C}_g) = \mathcal{C}_g'$. $\qquad\square$

Lemma 1.3. *Let $\mathcal{A}$ be a hyperbolic or elliptic coaxial family. Let O, A, and B be three pairwise distinct circles that belong to the orthogonal complement $\mathcal{A}^\perp$. Orient the circles to obtain O^+, A^+ and B^+ and let C be an oriented circle that does not belong to $\mathcal{A}$. Suppose that $\langle O^+, A^+ \rangle = \langle O^+, B^+ \rangle$. If O and C are orthogonal and C is segregated from A^+, then C is not segregated from B^+.*

Proof. Case 1: $\mathcal{A}$ is hyperbolic. Since $\mathcal{A}$ is hyperbolic, its orthogonal complement $\mathcal{A}^\perp$ is elliptic and as the circles O, A, and B belong to $\mathcal{A}^\perp$, they pass through two distinct common points, say a and b. Fig. 3(a) pictures the oriented circles O^+, A^+ and B^+ with $\langle O^+, A^+ \rangle = \langle O^+, B^+ \rangle$. Applying appropriate Möbius transformations to normalize, we may assume that $a = 0$, $b = \infty$, O^+ is the real axis of $\widehat{\mathbb{C}}$ oriented towards the right, and A^+ and B^+ are oriented lines through the origin both meeting the real axis at the same acute angle, but oriented so that $\langle O^+, A^+ \rangle = \langle O^+, B^+ \rangle$, as in Fig. 3(b). Since C is orthogonal to O and not a member of the coaxial family $\mathcal{A}$, it is either a circle in the
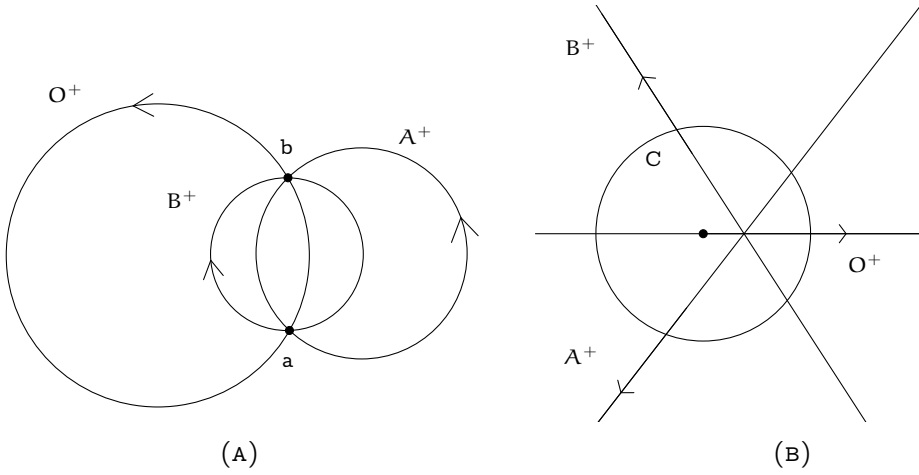
FIGURE 3. Case 1.

normalized picture centered on the real axis at a point $x \neq 0$, or a vertical line orthogonal to the real axis. Now if C is tangent to or disjoint from A, then it is tangent to or disjoint from B, and hence is contained in the companion disk of either A^+ or B^+. Therefore C is coupled with one of A^+ and B^+ and cannot be segregated from both A^+ and B^+. We may assume then that the radius r of C is greater than $|x|/\sqrt{2}$ ($r = \infty$ when a C is a line), forcing C to intersect each of A and B in exactly two distinct points (one of which may be ∞). Since $x \neq 0$, the smaller angle of intersection of C with both A and B is $\theta < \pi/2$. From the symmetry of the normalized picture of Fig. 3(b), whether C is oriented clockwise or counterclockwise, the oriented angle of C with one of A^+ and B^+ is θ, and with the other is $\pi - \theta$. Since $\theta < \pi/2$, we have $\pi - \theta > \pi/2$ so C is not segregated from one of A^+ and B^+. We conclude that if C is in fact segregated from A^+, then it cannot be segregated from B^+.

Case 2: $\mathcal{A}$ is elliptic. Since $\mathcal{A}$ is elliptic, its orthogonal complement $\mathcal{A}^\perp$ is hyperbolic and as the circles O, A, and B belong to $\mathcal{A}^\perp$, after applying an appropriate Möbius transformation to normalize, we may assume that O^+ is the real axis of $\widehat{\mathbb{C}}$ oriented toward the right and A and B are circles centered at $\pm i$ of equal radii $r < 1$, as in Fig. 4. To satisfy that $\langle O^+, A^+ \rangle = \langle O^+, B^+ \rangle$, one of A^+ and B^+ is oriented clockwise and the other counterclockwise. Since C is orthogonal to O, it is a circle in the normalized picture centered on O or a vertical line orthogonal to the real axis. If a circle, we may apply an elliptic Möbius transformation that set-wise fixes O, A, and B so that C is a vertical line. Since C is not a member of the coaxial family $\mathcal{A}$, C is a vertical line that misses the origin. The orientations on A^+ and B^+, one clockwise and the other counterclockwise, then imply that
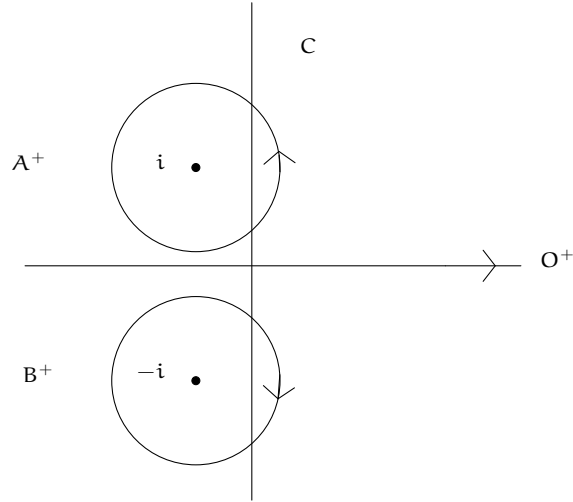
FIGURE 4. Case 2.

C meets A and B in two distinct points, since otherwise it is coupled with one of A^+ and B^+. But now, just as in the preceding paragraph, the oriented angle of C with one of A^+ and B^+ is $\theta < \pi/2$, and with the other is $\pi - \theta > \pi/2$, so C is not segregated from one of A^+ and B^+. We conclude that if C is in fact segregated from A^+, then it cannot be segregated from B^+. $\square$

The content of the Main Theorem is that when the convex c-polyhedra are non-unitary and proper, whose definition awaits, item (ii) of Theorem 1.2 is superfluous as it will be seen to follow as a consequence of item (i).

1.5. **Green-black polygons and proper c-polyhedra.** Fig. 2 might give the reader the impression that a convex hyperideal polygon may be "completed" to a compact convex hyperbolic polygon by adding the hyperideal vertices as new sides. Fig. 5 should provide a remedy to this impression. Fig. 5(a) shows three oriented lines ℓ_1, ℓ_2, and ℓ_3 that pairwise intersect at three finite vertices so there is no hyperideal vertex to exploit. With the orientations shown, the convex region P is unbounded. Of course if the orientations are reversed on the three lines, then these oriented lines cut out a bounded convex polygon, a triangle. More serious is the example pictured in Fig. 5(b) in which the hyperideal vertex between ℓ_1 and ℓ_3 fails to lie in the convex region P as it meets ℓ_2 at a point separating its meeting points with ℓ_1 and ℓ_3. This inspires the next definition.
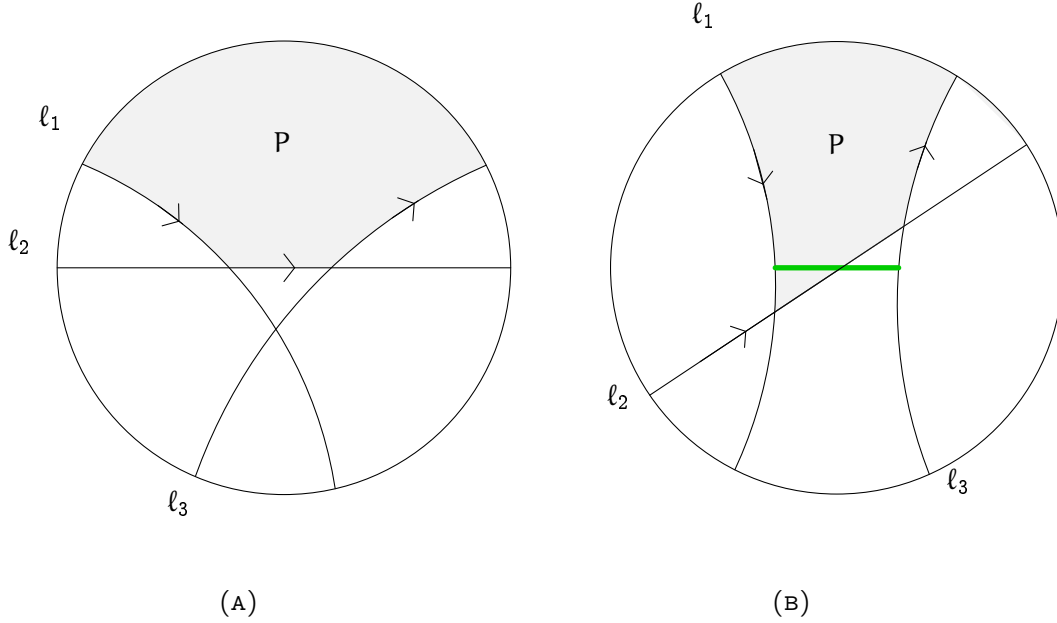
(A)   (B)

FIGURE 5. Improper green-black polygons.

Definition (Proper hyperideal polygon). The convex hyperideal polygon P determined by the cyclically ordered oriented lines $\ell_1, \ldots, \ell_n$ is said to be *proper* provided the following two conditions are met.

(1) Any hyperideal vertex, say for instance the hyperbolic line segment $s_{i,i+1}$ meeting the two consecutive lines ℓ_i and ℓ_{i+1} orthogonally, lies in the the convex hyperideal polygon P, as pictured in Fig. 6b.

(2) The oriented lines $\ell_1, \ldots, \ell_n$ along with any hyperideal vertices form the boundary of a bounded or compact convex polygon P′ contained in P, as in Fig. 6c.

We will construct a bounded convex hyperbolic polygon whose sides and vertices are colored either green or black from a proper convex hyperideal polygon.

Definition (Green-black polygon). Let P be a bounded convex hyperbolic polygon in $\mathbb{H}^2$ with vertices $p_1, \ldots, p_n$ listed in cyclic order respecting orientation. We denote this by $P = p_1 \cdots p_n$. Color each edge and each vertex of P either green or black. Then P is called a *green-black polygon* provided

(1) each green edge is adjacent to two black edges;
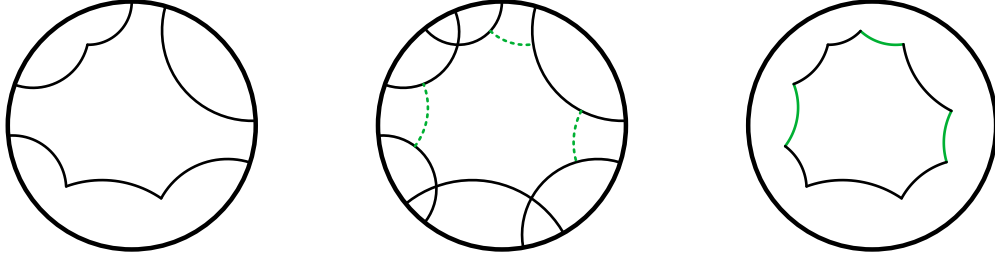
(A) A convex hyperideal 6-gon

(B) The green dotted segments denote the hyperideal vertices between consecutive non-intersecting lines.

(C) The green-black polygon defined by the convex hyperideal polygon in fig. 6a

FIGURE 6. Construction of a green-black polygon from a proper convex hyperideal polygon.

(2) each vertex incident to a green edge is colored black, and the remaining vertices are colored green;

(3) adjacent green and black edges are orthogonal.

The reason for the color coding is that the version of the Cauchy Arm Lemma we will prove in Section 3 considers green-black polygons for which the black side lengths are fixed while the green side lengths vary, and the black angles are fixed at right angles while the green ones vary—black indicates fixed measurements, green variable measurements.

We construct a green-black polygon from a proper convex hyperideal polygon P with no ideal vertices by the following procedure. First color all of the edges of P, some of which are rays or lines in $\mathbb{H}^2$, black. For each non-intersecting pair of consecutive support lines ℓ_i, ℓ_{i+1}, add the hyperideal vertex $s_{i,i+1}$ as a green colored segment along the unique shortest path between ℓ_i and ℓ_{i+1} as in Fig. 6c. Each line ℓ_i now supports a line segment, either a finite edge of P, or a truncation of a line or ray on the boundary of P to a line segment. The resulting polygon P' is compact and convex and every green edge has as neighbors two black edges, both of which it meets at right angles. Also color the vertices—if a vertex is adjacent to two black edges, color it green, otherwise black.

The analogue of boundedness for c-polyhedra turns out to be a bit more subtle than one might at first suspect. The vertices of bounded polyhedra in $\mathbb{E}^3$ have compact links while the links of vertices at infinity of unbounded polyhedra may fail to be compact. Moreover, convex bounded polyhedra have convex links. We need a condition to force this behavior

on convex c-polyhedra.[15] The subtlety arises precisely because we allow adjacent circles to be separated. When dealing with circle packings with ortho-circles where adjacent circles meet, either tangent or overlapping in an angle in the range $[0, \pi/2]$, the analogue of compact and convex links is satisfied automatically. When adjacent circles are allowed to separate, we need a condition that guarantees the behavior that the boundedness of convex Euclidean polyhedra forces upon neighborhoods of the vertices of the polyhedra. For this we introduce the next definition, which serves to force compact behavior around the vertices.

Definition (Proper c-polyhedron, and c-links). Let $G(\mathcal{C})$ be a non-unitary convex c-polyhedron based on the oriented abstract spherical polyhedron P with vertex set $V = V(P)$. For any vertex $v \in V$, give the interior of the companion disk D_v to C_v a complete hyperbolic metric of constant curvature -1 making D_v a model of the hyperbolic plane with its circle at infinity. Let $f_1, \dots, f_n$ be the faces adjacent to v ordered cyclically about v with respect to the orientation of P. Let ℓ_i be the hyperbolic line in D_v determined by the orthogonal intersection $D_v \cap O_{f_i}^+$, but oriented oppositely to that of the ortho-circle $O_{f_i}^+$. Then $G(\mathcal{C})$ is *proper at* v if the oriented lines $\ell_1 \dots, \ell_n$ are the support lines of a proper convex hyperideal polygon $P(v)$ in D_v. The c-polyhedron $G(\mathcal{C})$ then is *proper* provided it is proper at each of its vertices. The green-black polygon $L(v)$ determined by the proper hyperideal polygon $P(v)$ is the *c-link* of the vertex circle C_v.

Consistent orientation of our c-polyhedron $G(\mathcal{C})$ with its oriented abstract spherical polyhedron P implies that lines ℓ_i in the definition of properness are ordered cyclically about $P(v)$, meaning that the cyclic ordering that $P(v)$ gives to the support lines is exactly that given to the faces adjacent to v by the orientation of P. The relationship between convexity and properness is a bit subtle. We suspect that if the c-polyhedron $G(\mathcal{C})$ is edge segregated so that it avoids deep overlaps, then convexity of $G(\mathcal{C})$ implies its properness. We have yet to verify our suspicions.

1.6. **A congruence condition for c-polyhedra.** Two green-black polygons are *compatible* when they have the same number of green and black edges given in the same order. When also every pair of corresponding black edges have the same hyperbolic length, the polygons are *black-edge-congruent*. Note here that the complex interior angles—the real angles at green vertices and the imaginary unit times the lengths of the green edges—may differ between two black-edge-congruent polygons. We use the plain term *congruent* to mean the corresponding complex angles also agree.

---

[15]One can in fact describe convex circle configurations that mimic unbounded polyhedra in $\mathbb{E}^3$ and fail to be globally rigid.
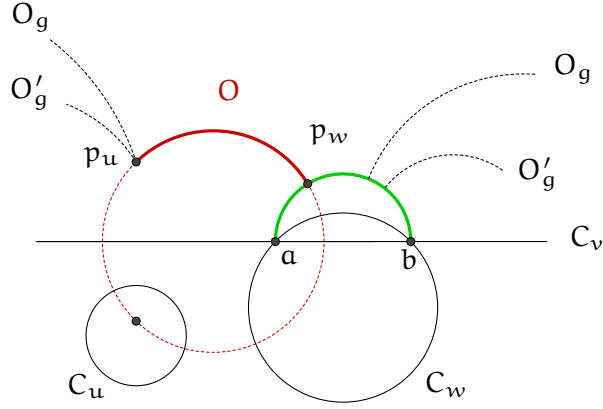
FIGURE 7. The construction of the points p_u and p_w.

We end this rather lengthy preliminary section with an important observation. If the non-unitary c-polyhedron $G(\mathcal{C})$ is both convex and proper, then the c-links are Möbius invariants. This means that if $T \in \text{Möb}(\mathbb{S}^2)$ and v is a vertex of G, then the c-link of C_v in $G(\mathcal{C})$ and that of $T(C_v)$ in $T(G(\mathcal{C}))$ are congruent. This is obvious from the fact that the Möbius transformation T restricts to an isometry of D_v° with $T(D_v^\circ)$ when both open disks are given the complete hyperbolic metric as in the definition of c-link. Here D° means the open interior of the closed disk D. This goes the other direction as well and allows us to replace condition (ii) in Theorem 1.2 by an equivalent one on c-links, as in Corollary 1.6 below.

Lemma 1.4. *Let* $G(\mathcal{C})$ *and* $G(\mathcal{C}')$ *be two proper, convex, non-unitary c-polyhedra both based on the abstract spherical polyhedron* P. *If* $G(\mathcal{C})$ *and* $G(\mathcal{C}')$ *have Möbius-congruent c-faces, then corresponding c-links are black-edge-congruent.*

Proof. We show that the corresponding black edges of the links $L(v)$ in $G(\mathcal{C})$ and $L(v)'$ in $G(\mathcal{C}')$ have the same length. Let u and w be vertices of adjacent edges uv and vw of P that lie on the face f. First by applying an appropriate Möbius transformation, assume that the c-polyhedra $G(\mathcal{C})$ and $G(\mathcal{C}')$ have been placed in a *normalized* position where $C_u = C_u'$, $C_v = C_v'$, and $C_w = C_w'$. Let O be the unique circle orthogonal to C_u, C_v, and C_w and note that the normalization implies that $O_f = O = O_f'$. Our claim is that the black edges of $L(v)$ and $L(v)'$ that lie along O are completely determined by the circles $C_u = C_u'$, $C_v = C_v'$, and $C_w = C_w'$. To see this, we define a point p_u in the open disk $D^\circ = D_v^\circ = D_v'^\circ$ as follows. The circles C_u and C_v determine a unique coaxial family $\mathcal{A}_u$ and the ortho-circle O is a member of the complementary family $\mathcal{A}_u^\perp$. There are two

cases to consider depending on whether or not C_u and C_v meet. When C_u is disjoint from C_v, the family $\mathcal{A}_u$ is a hyperbolic coaxial family with two foci, one of which is in D°. This one we call p_u. On the other hand, when C_u and C_v meet at two points a and b on C_v, the coaxial family $\mathcal{A}_u$ is elliptic. In this case, let λ be the hyperbolic line in D° with ideal vertices a and b. Note that the circle containing λ is the unique member of $\mathcal{A}_u$ that is orthogonal to C_v. Let p_u be the point of intersection of O with λ in the disk D°. The observation we make is that in both c-links $L(v)$ and $L(v)'$, the black edge determined by $O_f = O_f'$ is the edge $e = p_u p_w$, the hyperbolic segment connecting p_u and p_w. This follows from the observation that both ortho-circles O_g and O_g' for the respective c-polyhedra $G(\mathcal{C})$ and $G(\mathcal{C}')$ that are determined by the face g adjacent to f along the edge uv are members of the complementary coaxial family $\mathcal{A}_u^\perp$. In particular, in the case that $\mathcal{A}$ is hyperbolic, both pass through the same point p_u, and in the case that $\mathcal{A}$ is elliptic, the unique green edge from O to both O_g and O_g' lies along the hyperbolic line λ and so meets O at p_u. See Fig. 7 for these cases. $\qquad\square$

Remark 1.5. Note that the complex angle at a green vertex or a green edge of a c-link $L(v)$ is the complex dihedral angle between the two adjacent c-faces $\mathcal{C}_f$ and $\mathcal{C}_g$ incident to C_v.

This observation along with the preceding lemma gives the following corollary.

Corollary 1.6. *Let $G(\mathcal{C})$ and $G(\mathcal{C}')$ be two proper, convex, non-unitary c-polyhedra both based on the abstract spherical polyhedron P. Then $G(\mathcal{C})$ and $G(\mathcal{C}')$ are Möbius-congruent with one another if and only if the following conditions hold.*

  (i) *$G(\mathcal{C})$ and $G(\mathcal{C}')$ have Möbius-congruent c-faces.*

  (ii) *The complex angles at the corresponding green (finite and hyperideal) vertices of corresponding c-links agree; equivalently, by the preceding lemma, corresponding c-links are congruent.*

2. Cauchy's Rigidity Theorem

In this section we review Cauchy's celebrated rigidity theorem on the uniqueness of convex, bounded polyhedra in $\mathbb{E}^3$. This serves as the model for our proof of the Main Theorem. In fact, the bulk of the remainder of this paper is dedicated to an appropriate generalization of Cauchy's Arm Lemma from the context of Euclidean and spherical geometry to that of Möbius geometry.

2.1. **Cauchy's argument recalled.** Suppose we have two convex polyhedra in Euclidean 3-space $\mathbb{E}^3$ with the same combinatorics and with corresponding faces congruent. Cauchy's theorem states that the two polyhedra must be congruent, meaning that there is a Euclidean isometry mapping one to the other.

Cauchy's proof is split into two components—the one geometric and the other combinatorial. The geometric component is the Discrete Four Vertex Lemma (Lemma 2.2), which follows from an application of Cauchy's Arm Lemma (Lemma 2.1), presented after a bit of notation. We denote a planar or spherical polygon P merely by listing its vertices in cyclic order, say as $P = p_1 \ldots p_n$. The Euclidean or spherical length of the side $p_i p_{i+1}$ is denoted as $|p_i p_{i+1}|$ and the interior angle at p_i is denoted as $\angle p_i$, though on occasion, for either emphasis or to avoid confusion, we will use $\angle p_{i-1} p_i p_{i+1}$.

Lemma 2.1 (Cauchy's Arm Lemma). *Let $P = p_1 \ldots p_n$ and $P' = p'_1 \ldots p'_n$ be two convex (planar or spherical) polygons such that, for $1 \leqslant i < n$, $|p_i p_{i+1}| = |p'_i p'_{i+1}|$, and for $1 \leqslant i < n-1$, $\angle p_{i+1} \leqslant \angle p'_{i+1}$. Then $|p_n p_1| \leqslant |p'_n p'_1|$ with equality if and only if $\angle p_{i+1} = \angle p'_{i+1}$ for all $1 \leqslant i < n-1$.*

Cauchy's original proof of the lemma had a gap that subsequently was filled. A straightforward inductive proof, such as the one in [11], relies on the law of cosines and the triangle inequality. Imagine now that we have convex planar or spherical polygons P and P' with the same number of sides with corresponding sides of equal length. Label each vertex of P with a plus sign $+$ or a minus sign $-$ by comparing its angle with the corresponding angle in P': if the angle at p_i is larger than that at p'_i, label it with a $+$, if smaller, a $-$, and if equal, no label at all. Using the Cauchy Arm Lemma, the proof of the following lemma is straightforward.

Lemma 2.2 (Discrete Four Vertex Lemma). *Let P and P' be as in the preceding paragraph and label the vertices of P as described. Then either P and P' are congruent (and thus no vertex is labeled with $+$ or $-$), or a walk around P encounters at least four sign changes (from $-$ to $+$ or from $+$ to $-$).*

Proof. First note that because a polygon is a cycle the number of sign changes must be even. Suppose there are no sign changes, i.e., either no vertex is labeled or all of the labeled vertices have the same label. If no vertex is labeled, then the two polygons are congruent. Assume then that some of the vertices are labeled, but all with the same label. Then Cauchy's Arm Lemma implies that there exists a pair of corresponding edges in P and P' with different lengths, a contradiction.

Assume now that there are exactly two sign changes of the labels of P. Select two edges $p_i p_{i+1}$ and $p_j p_{j+1}$ (oriented counter-clockwise) of P such that all of the $+$ signs are along

the subchain from p_{i+1} to p_j and all of the $-$ signs are along the subchain from p_{j+1} back to p_i. Subdivide both edges in two by adding a vertex at the respective midpoints X and Y of $p_i p_{i+1}$ and $p_j p_{j+1}$. Similarly, subdivide the corresponding edges $p_i' p_{i+1}'$ and $p_j' p_{j+1}'$ in P′ at midpoints X′ and Y′. Denote the subchain of P from X to Y by P_+ and the subchain from Y back to X by P_-. Similarly for P_+' and P_-' in P′. Applying the arm lemma to P_+ and P_+' implies that $|XY| > |X'Y'|$, and, similarly, an application to P_- and P_-' implies that $|XY| < |X'Y'|$, a contradiction. □

This brings us to the combinatorial component of Cauchy's proof. A nice proof of the following lemma appears in [11] and follows from an argument based on the Euler characteristic of a sphere.

Lemma 2.3 (Cauchy's Combinatorial Lemma). *Let P be an abstract spherical polyhedron. Then for any labeling of any non-empty subset of the edges of P with $+$ and $-$ signs, there exists a vertex v that is incident to an edge labeled with a $+$ or a $-$ sign for which one encounters at most two sign changes in labels on the edges adjacent to v as one walks around the vertex.*

The proof of Cauchy's Rigidity Theorem is now straightforward. Assume we have convex polyhedra P and P′ that have the same combinatorics and congruent corresponding faces. For each edge of P label its dihedral angle with a $+$ or a $-$ depending on whether it is larger or smaller than the corresponding dihedral angle in P′. If P and P′ are not congruent, Cauchy's Combinatorial Lemma provides a vertex v that is incident to an edge labeled with a $+$ or a $-$ sign, and around which there are at most two sign changes. Intersect P with a small sphere centered at v (one that contains no other vertex of P on its interior) to obtain a convex spherical polygon, and intersect P′ with a sphere centered at the corresponding vertex v′ and of the same radius. By construction both spherical polygons have the same edge lengths, and the angles between edges are given by the dihedral angles between faces at v and v′. An application of the Four Vertex Lemma implies that there are at least four sign changes, contradicting that there are at most two. We conclude that P and P′ are congruent.

2.2. Observations towards a proof for c-polyhedra. The remainder of this paper is concerned with developing a Cauchy-style proof of the global rigidity of convex, proper, non-unitary c-polyhedra. To that end we make the following observations. First, Cauchy's Combinatorial Lemma is valid for abstract polyhedral graphs, which means that it holds for c-polyhedra. Second, though it is straightforward to extend the Four Vertex Lemma to convex hyperbolic polygons, doing so would only get us halfway to our desired result. Instead we generalize the Four Vertex Lemma to the class of green-black polygons.
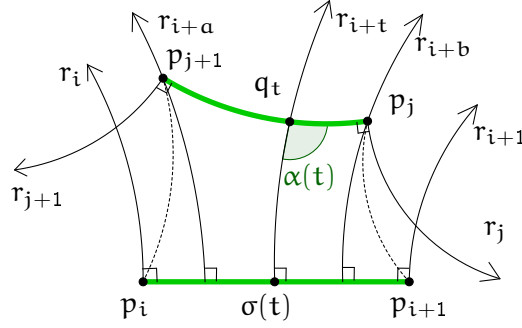
FIGURE 8. The construction for Lemma 3.1.

Finally, the crucial point of connection between the combinatorial and geometric components of Cauchy's proof is made by intersecting small spheres centered at the vertices of the polyhedra to obtain spherical polygons, the spherical links of the vertices. In the case of c-polyhedra it is not at all obvious how to go about obtaining a similar construction as there is no obvious "sphere" with which to intersect a c-polyhedron to obtain the analogue of a link of a vertex. But this is precisely where the c-links of circles in a c-polyhedron come into play, which play the role of links of vertices of Euclidean polyhedra. Given our construction of c-links, and our version of the Four Vertex Lemma for green-black polygons, the remainder of the proof is essentially Cauchy's original argument.

3. THE GREEN-BLACK ARM AND FOUR VERTEX LEMMAS

3.1. **An important property.** An important property of green-black polygons is stated in the next lemma.

Lemma 3.1. *Let P be a green-black polygon with ordered vertices $p_1, \ldots, p_n$, some of which are green with the remaining ones black. Let $p_i p_{i+1}$ be a green edge. Then the shortest path from $p_i p_{i+1}$ to a vertex p_j meets $p_i p_{i+1}$ at a right angle. Similarly, if $p_j p_{j+1}$ is another green edge, then the shortest path between $p_i p_{i+1}$ and $p_j p_{j+1}$ meets both green edges at right angles.*

Proof. Since the edge $p_i p_{i+1}$ is green, the respective edges $p_{i-1} p_i$ and $p_{i+1} p_{i+2}$ are black and meet the support line l_i of the edge $p_i p_{i+1}$ orthogonally at the respective points p_i and p_{i+1}. Let σ be a unit time parameterization of $p_i p_{i+1}$ with $\sigma(0) = p_i$ and $\sigma(1) = p_{i+1}$. Extend rays r_i and r_{i+1} outwards from p_i and p_{i+1} along the edges $p_{i-1} p_i$ and $p_{i+1} p_{i+2}$

and observe that the convex region R_i bounded by r_i, r_{i+1}, and $p_i p_{i+1}$ is the disjoint union of rays r_{i+t} for $0 \leqslant t \leqslant 1$, where r_{i+t} is the ray in the region R_i orthogonal to the segment $p_i p_{i+1}$ at the point $\sigma(t)$. This implies that the nearest point of the edge $p_i p_{i+1}$ to any point q of R_i is that point $\sigma(t)$ for which $q \in r_{i+t}$, and in this case the shortest path from q to $p_i p_{i+1}$ follows the hyperbolic segment from $\sigma(t)$ to q, a subsegment of the ray r_{i+t}, which therefore meets $p_i p_{i+1}$ orthogonally.

Now let $p_j p_{j+1}$ be another green edge and extend rays r_j and r_{j+1} outwards along the black edges incident to p_j and p_{j+1}. By convexity of the polygon P, p_j and p_{j+1} are contained in the convex region R_i, and p_i and p_{i+1} are contained in the convex region R_j bounded by r_j, r_{j+1}, and $p_j p_{j+1}$. This implies that all four interior angles of the hyperbolic quadrilateral $p_i p_{i+1} p_j p_{j+1}$ are at most $\pi/2$; see Fig. 8. Let $0 \leqslant a < b \leqslant 1$ be the t-values for which $p_{j+1} \in r_{i+a}$ and $p_j \in r_{i+b}$, and for $a \leqslant t \leqslant b$, let q_t be the point on the edge $p_j p_{j+1}$ that lies on the ray r_{i+t}. Finally, let $\alpha(t)$ be the measure of the angle $\angle \sigma(t) q_t p_j$; again see Fig 8. By our construction and the observation that the angles of the quadrilateral $p_i p_{i+1} p_j p_{j+1}$ are acute or right, $\alpha(a) \leqslant \pi/2$ and $\alpha(b) \geqslant \pi/2$. By continuity of the function α, there is a parameter value $a \leqslant t \leqslant b$ such that $\alpha(t) = \pi/2$. The hyperbolic segment from $\sigma(t)$ to q_t, which lies on the ray r_{i+t}, meets both edges $p_i p_{i+1}$ and $p_j p_{j+1}$ orthogonally. This implies that this segment is the shortest path connecting edges $p_i p_{i+1}$ and $p_j p_{j+1}$. $\qquad\square$

3.2. Relaxed green-black polygons.
We generalize the concept of green-black polygons slightly. Instead of requiring that green edges meet black edges at right angles, we relax this requirement and allow the measures of these angles to be less than or equal to $\pi/2$. Call these polygons *relaxed green-black polygons*. The proof of Lemma 3.1 may be modified to verify the following.

Corollary 3.2. *Let P be a relaxed green-black polygon with vertices $p_1, \ldots, p_n$. Let $p_i p_{i+1}$ be a green edge. Then the shortest path from $p_i p_{i+1}$ to any vertex p_j meets $p_i p_{i+1}$ at a right angle. Similarly, if $p_j p_{j+1}$ is a green edge, then the shortest path between $p_i p_{i+1}$ and $p_j p_{j+1}$ meets both at right angles.*

3.3. A four vertex lemma for green-black polygons.
When X and Y are points in the hyperbolic plane $\mathbb{H}^2$, we use $|XY|$ to denote the hyperbolic distance between the points. We now prove an analog of Cauchy's Discrete Four Vertex Lemma for green-black polygons. Let P and P' be two black-edge-congruent green-black polygons. We label a green vertex (or a green edge) of P with a plus sign $+$ if the angle (or edge-length) is larger than the corresponding angle (or edge-length) in P', a minus sign $-$ if it is smaller, or no label if equal. The statement of the lemma is now the same as Lemma 2.2.
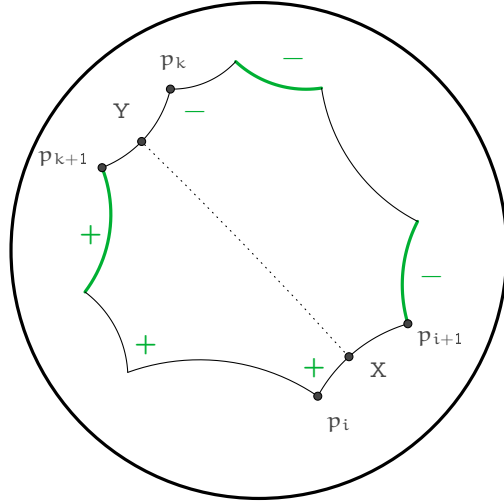
FIGURE 9. The construction for the contradiction in the proof of Lemma 3.3.

Lemma 3.3 (Green-Black Polygon Four Vertex Lemma). *Let P and P' be two black-edge-congruent green-black polygons labeled as in the preceding paragraph. Then either P and P' are congruent (in the usual sense), or a walk around P encounters at least four sign changes (from $-$ to $+$ or from $+$ to $-$).*

Proof. The proof follows the same sketch as in Cauchy's original. New are the details of extending the arm lemma to green-black polygons, which is accomplished by Lemma 3.4 below. As before the number of sign changes must be even, so we assume there are either two or zero sign changes. Denote the vertices of P in (counter-clockwise) order by $p_1, \ldots, p_n$ and the corresponding vertices of P' by $p'_1, \ldots, p'_n$.

Suppose there are exactly two sign changes. Pick a pair of black edges $p_i p_{i+1}$ and $p_k p_{k+1}$ of P such that all of the $-$ signs occur on the counter-clockwise walk from p_i to p_{k+1}, and all the $+$ signs occur on the counter-clockwise walk from p_k to p_{i+1}. Let X, Y, X', and Y' be the respective midpoints of the edges $p_i p_{i+1}, p_k p_{k+1}, p'_i p'_{i+1}$, and $p'_k p'_{k+1}$. Fig. 9 shows the construction in P. By applying the green-black arm lemma, Lemma 3.4 below, to the chains $X, p_{i+1}, p_{i+2}, \ldots, p_k, Y$ and $X', p'_{i+1}, p'_{i+2}, \ldots, p'_k, Y'$, we have $|XY| < |X'Y'|$. Similarly, by applying the green-black arm lemma to $Y, p_{k+1}, p_{k+2}, \ldots, p_i, X$ and $Y, p_{k+1}, p_{k+2}, \ldots, p_i, X$, we have that $|XY| > |X'Y'|$, a contradiction. Hence there cannot be two sign changes on P.

Suppose then there are no sign changes but some green angles or edges of P are marked, say with a $-$ sign. Choose any black edge $p_i p_{i+1}$ and choose two distinct points X and Y on the open edge from p_i to p_{i+1} ordered as p_i, X, Y, p_{i+1}. Since the green-black polygons P and P$'$ are black-edge congruent, $|p_i p_{i+1}| = |p_i' p_{i+1}'|$ and we may choose corresponding points X$'$ and Y$'$ on the open edge from p_i' to p_{i+1}', so that $|p_i X| = |p_i' X'|$, $|XY| = |X'Y'|$, and $|Yp_{i+1}| = |Y'p_{i+1}'|$. By applying the green-black arm lemma to the chains Y, $p_{i+1}, p_{i+2}, \ldots, p_i$, X and Y$', p_{i+1}', p_{i+2}', \ldots, p_i'$, X, we have that $|XY| < |X'Y'|$, a contradiction. Hence either P and P$'$ are congruent, or P contains at least four sign changes. $\qquad\square$

3.4. **The green-black arm lemma.** In this section we state and prove the Green-Black Arm Lemma, beginning with a definition.

Definition (Green-black arm chain). Any subchain of consecutive edges of a relaxed green-black polygon that begins and ends with black edges is called a *green-black arm chain*. The two vertices of unit valence are called the *free vertices*. We define two green-black arm chains to be *compatible* and *black-edge congruent* in the same way as for green-black polygons. We say that the green-black arm chain is *determined* by its vertices and we use the list of its vertices in order to name the chain.

Notice that Corollary 3.2 applies to green-black arm chains, and that connecting the two free vertices of a green-black arm chain with a black segment produces a relaxed green-black polygon, which always is convex. This latter observation, that a green-black arm chain produces a convex polygon when the free vertices are connected by a segment, is important in the proof.

Lemma 3.4 (Green-Black Arm Lemma). *Let* $p_1, p_2, \ldots, p_n$ *and* $p_1', p_2', \ldots, p_n'$ *determine the two respective compatible, black-edge congruent, green-black arm chains* P *and* P$'$ *such that:*

(1) *if edge* $p_i p_{i+1}$ *is green, then* $|p_i p_{i+1}| \leqslant |p_i' p_{i+1}'|$,

(2) *if* $\angle p_i$ *is green, then* $\angle p_i \leqslant \angle p_i'$, *and otherwise,* $\angle p_i = \angle p_i'$,

(3) *if* $p_i p_{i+1}$ *is green, then* $\angle p_i$ *equals* $\pi/2$.

Then $|p_1 p_n| \leqslant |p_1' p_n'|$, *with equality if and only if all of the corresponding angles and edges are congruent.*

3.4.1. *How not to prove the Green-Black Arm Lemma.* To prove the lemma, we will induct on N, the number of green edges of the green-black arm chain. Before the proof, though, we think it instructive to discuss how this induction argument may lead to a subtle mistake if one is not careful. If $N = 0$, The Green-Black Arm Lemma is just the classical
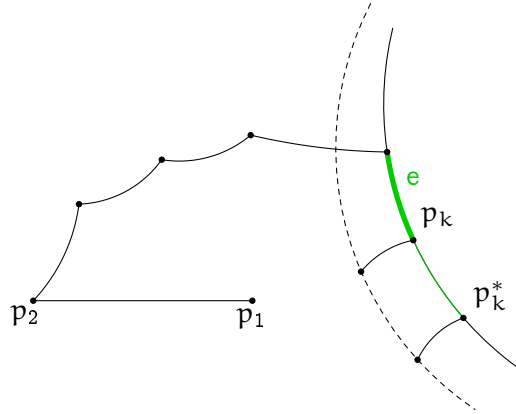
FIGURE 10. Convexity is lost in flowing p_k to p_k^*.

Cauchy Arm Lemma in the hyperbolic plane. Assume then that the lemma holds for $N-1$ green edges and consider a green-black arm chain with N green edges. Choose a green edge e and lengthen it continuously keeping all other edge-lengths and angles constant. Then the free vertices p_1 and p_n of the arm chain increase their distance. Now change the color of e from green to black and apply the inductive hypothesis. QED?

Fig. 10 illustrates the problem. When increasing the length of the one green edge e, there is no guarantee that the final result is a green-black arm chain to which the inductive hypothesis applies. The problem is that of *convexity*. The final resulting arm after lengthening e might fail to be convex when the free vertices are connected, so the inductive hypothesis fails to apply. This is the green-black version of the mistake Cauchy made in his original argument of 1813 that subsequently was noticed and repaired by Ernst Steinitz over 100 years later. This loss of convexity is the problem that necessitates a more careful analysis, to which we now turn.

3.4.2. *The proof of the Green-Black Arm Lemma.* Despite the caveat of the preceding paragraph, the proof does proceed by induction on N, the number of green edges. The basis of the induction is when $N = 0$, and as stated already, is just the hyperbolic version of the classical Cauchy Arm Lemma. The proof in this case is exactly that of I. Schoenberg's reported on pp. 340-341 of [11], but in which the distances are interpreted as hyperbolic distances. Assume then that for some $N \geqslant 1$, the lemma is true whenever there are no more than $N - 1$ green edges. By reversing the order if necessary, we may assume that the number of vertices between the initial vertex p_1 and the nearest green edge is at least

as large as the number between the terminal vertex p_n and its nearest green edge. Let $e = p_i p_{i+1}$ denote the last green edge one meets as one traverses the chain P from its first to last vertex and note that $i \leqslant n - 2$.

In the paragraph following, we are going to produce a new green-black arm chain that will be obtained by fixing the vertices $p_1, \ldots, p_i$ of P and elongating the edge e to obtain the edge $e^* = p_i p_{i+1}^*$ by translating p_{i+1} to p_{i+1}^* along the hyperbolic line ℓ supporting e. The remaining vertices p_j for $j > i$ will rigidly translate via the hyperbolic flow T_ℓ with axis ℓ to the respective vertices p_j^*. This new green-black arm chain determined by $p_1 \ldots, p_i, p_{i+1}^*, \ldots, p_n^*$ is denoted as P^* and will have the same corresponding interior angles as P as well as the same green edge lengths, except for the single edge e for which $|e| \leqslant |e^*| \leqslant |e'|$, where $e' = p_i' p_{i+1}'$. In particular, P^* and P' will be compatible, black-edge congruent, green-black arm chains that satisfy conditions (1)–(3) of the lemma.

To obtain the intermediate green-black arm chain P^*, continuously increase the length of e by flowing the points p_j for $j > i$ via T_ℓ until one of two possibilities occurs. Either (I) the hyperbolic length of $e^* = p_i p_{i+1}^*$ is equal to that of $e' = p_i' p_{i+1}'$ and the polygon formed by connecting the free vertices of P^* is convex, or (II) we reach the situation where p_1, p_2, and p_n^* become collinear, all lying on the hyperbolic line supporting the edge $p_1 p_2$, before e^* is equal in length to e'. In this latter case, to continue the flow until the lengths of e^* and e' agree would produce a non-convex polygon. Our first task in either of the cases is to show that (†) $|p_1 p_n| \leqslant |p_1 p_n^*|$ and that the inequality is strict if and only if $e \neq e^*$, i.e., the length of e^* is strictly larger than that of e. After this is verified, we consider the two cases, (I) and (II), in turn.

Verifying (†): $|p_1 p_n| \leqslant |p_1 p_n^*|$ *with strict inequality if and only if* $e \neq e^*$. We use the following fact of elementary hyperbolic geometry.

Lemma 3.5. *Fix a point* p *and a hypercycle* h[16] *in the hyperbolic plane* $\mathbb{H}^2$. *Let* m *be the unique hyperbolic line passing through* p *and perpendicular to* h *and let* q *be the unique point of intersection of* m *with* h. *Let* r *and* s *be points on* h *in one of the half planes determined by* m *with* r *strictly between* q *and* s *in the ordering on the hypercycle* h. *Then* $|pr| < |ps|$.

Let σ be the geodesic arc of shortest length that connects the point p_1 to the edge e. By Lemma 3.1, σ meets e orthogonally, say at the point u. Since the polygon formed by adding the edge $p_1 p_n$ to the arm chain P is convex, the line ℓ_1 supporting σ meets P only at the vertex p_1 and the point u of e, unless $\sigma = p_1 p_2$ with $u = p_2$, in which case ℓ_1

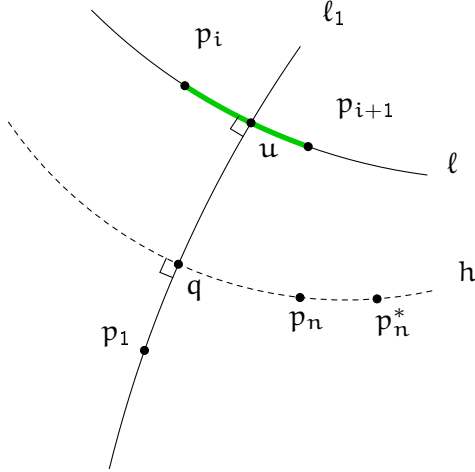[16]This includes the case where h is a hyperbolic line.

FIGURE 11. Flow along a hypercycle increases distance.

meets P along its first edge $p_1 p_2$. In either case, as the chain of edges from p_{i+1} to p_n along P is connected and does not meet ℓ_1, that chain lies in one half-space determined by ℓ_1. Let h be the hypercycle determined by p_n and ℓ, the component containing p_n of the set of points whose distance to ℓ is equal to the distance from p_n to ℓ. This hypercycle h is exactly the flow line of p_n under the hyperbolic flow T_ℓ. Since ℓ_1 is orthogonal to e and therefore to ℓ, ℓ_1 is orthogonal to all the hypercycles determined by ℓ, or what is the same, to all the flow lines of T_ℓ. In particular, ℓ_1 is orthogonal to h. Let q be the point of intersection of ℓ_1 with h. Then since e lengthens to e^* when $e \neq e^*$ and p_{i+1} and p_n are in the same half-plane determined by ℓ_1, p_n is strictly between q and p_n^* in the ordering of h; see Fig. 11. Apply the preceding lemma with $p = p_1$, $r = p_n$ and $s = p_n^*$ to conclude that $|p_1 p_n| < |p_1 p_n^*|$.

Case (I). In this case, the new arm chain P^* obtained from lengthening e is convex when the free vertices are joined by a hyperbolic segment, so that it is a green-black arm chain. Moreover, conditions (1)–(3) hold for the two compatible, black-edge congruent, green-black arm chains P^* and P' with the additional fact that the lengths of the green edges e^* and e' agree. Recolor these green edges, e^* and e', black and note that P^* and P' with these recolored edges are two compatible, black-edge congruent green-black arm chains that satisfy conditions (1)–(3) and have $N - 1$ green edges. Apply the inductive hypothesis for N to conclude that $|p_1 p_n^*| \leqslant |p_1' p_n'|$ with equality if and only if corresponding angles and edge lengths agree. Applying (†), we have $|p_1 p_n| \leqslant |p_1 p_n^*| \leqslant |p_1' p_n'|$, again with equality if and only if corresponding angles and edge lengths agree.

Case (II). By increasing the size of the edge e, we have arrived at a point where p_1, p_2 and p_n^* are collinear with p_1 between p_2 and p_n^*, so that $|p_1 p_n^*| = |p_2 p_n^*| - |p_1 p_2|$. Were e to increase its length further, convexity would be violated. Using (†) exactly as in the preceding paragraph, we need only verify that (‡) $|p_1 p_n^*| \leqslant |p_1' p_n'|$. We verify (‡) by inducting on M, the number of black edges between p_1 and the first green edge encountered in a walk from p_1 to p_n in P. The basis of the induction when $M = 1$ is a bit more difficult to prove than the inductive step, so we will delay its verification until after we verify the inductive step.

(‡) *The inductive step,* $M \geqslant 2$: Let $M \geqslant 2$ and assume that the result holds for $M - 1$. Since M is at least 2, the initial two edges, $p_1 p_2$ and $p_2 p_3$, are black. Consider the chains $p_2, \cdots, p_n^*$ and $p_2', \cdots, p_n'$ and apply the inductive hypothesis for M to conclude that $|p_2 p_n^*| < |p_2' p_n'|$. We have

$$|p_1 p_n^*| = |p_2 p_n^*| - |p_1 p_2| \leqslant |p_2' p_n'| - |p_1' p_2'| \leqslant |p_1' p_n'|,$$

where the first inequality follows since $|p_1 p_2| = |p_1' p_2'|$, and the last inequality is the triangle inequality.

(‡) *The basis of the induction,* $M = 1$: When $M = 1$, by our initial requirement that the number of vertices between the initial vertex p_1 and the nearest green edge is at least as large as the number between the terminal vertex p_n and its nearest green edge, both edges $p_2 p_3$ and $e = p_{n-2} p_{n-1}$ are green. By Corollary 3.2 the shortest path between $p_2 p_3$ and $p_{n-2} p_{n-1}^*$ meets both orthogonally. Let X and Y be the endpoints of this path. We obtain Fig. 12 by applying a Möbius transformation taking X and Y to the respective points $-A$ and A on the x-axis, with p_1, p_2, p_{n-1}^* and p_n^* below the x-axis and the remaining vertices of P above the x-axis.

Next, along the lines k and ℓ supporting the respective green edges $p_2 p_3$ and $p_{n-2} p_{n-1}^*$, stretch both green edges so that their lengths equal the corresponding lengths in P'. To do this, flow $p_1 p_2$ using the hyperbolic flow T_k with axis k to edge $p_1^* p_2^*$ so that $|p_2^* p_3| = |p_2' p_3'|$. By hypothesis, the edge $p_1 p_2$ meets k orthogonally, and so the edge $p_1^* p_2^*$ must also meet k orthogonally. Similarly flow $p_{n-1}^* p_n^*$ to $p_{n-1}^{**} p_n^{**}$ using T_ℓ so that $|p_{n-2} p_{n-1}^{**}| = |p_{n-2}' p_{n-1}'|$.

If the resulting polygon $p_1^* p_2^* p_3 \cdots p_{n-2} p_{n-1}^{**} p_n^{**}$ fails to be convex, the line through one of the edges, either $p_1^* p_2^*$ or $p_{n-1}^{**} p_n^{**}$, meets the polygon at a point not on that edge. In this case, we flow the other edge backwards until convexity is restored. The argument is symmetric in the two cases, so to continue we will assume, without loss of generality, that the support line m of $p_{n-1}^{**} p_n^{**}$ meets the polygon at a point not on that edge. Now apply the flow T_k in reverse to move the edge $p_1^* p_2^*$ back until its image $p_1^{**} p_2^{**}$ under
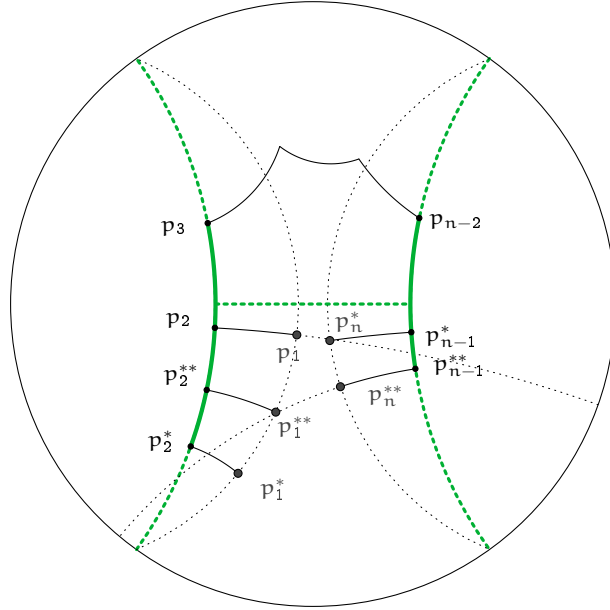
FIGURE 12. The construction for the proof of the basis of the induction
(‡). The dotted green line on the x-axis is the shortest path between $p_2 p_3$
and $p_{n-2} p_{n-1}^*$.

the flow meets m at the point p_1^{**}; see Fig. 12. The points $p_1^{**}, p_2^{**}, p_3, \cdots p_{n-2}, p_{n-1}^{**} p_n^{**}$
determine a green-black arm chain P^{**} that is compatible and black-edge congruent with
P'. Moreover, by the construction of P^{**}, it satisfies conditions (1)–(3) of the lemma and
the green edge $e^{**} = p_{n-2} p_{n-1}^{**}$ has the same length as the green edge $e' = p_{n-2}' p_{n-1}'$.
As in Case (I), recolor these green edges, e^{**} and e', black and note that P^{**} and P' with
these recolored edges are two compatible, black-edge congruent green-black arm chains
that satisfy conditions (1)–(3) and have $N-1$ green edges. Apply the inductive hypothesis
for N to conclude that $|p_1^{**} p_n^{**}| \leqslant |p_1' p_n'|$ with equality if and only if corresponding angles
and edge lengths agree.

The three inequalities

$$|p_1 p_n| \leqslant |p_1 p_n^*|, \quad |p_1 p_n^*| \leqslant |p_1^{**} p_n^{**}|, \quad |p_1^{**} p_n^{**}| \leqslant |p_1' p_n'|$$

imply that $|p_1 p_n| \leqslant |p_1' p_n'|$, confirming the basis step of induction on M. The first in-
equality is from (†), the third is confirmed in the preceding paragraph, and the second is
a consequence of the next lemma. That equality holds if and only if corresponding angles
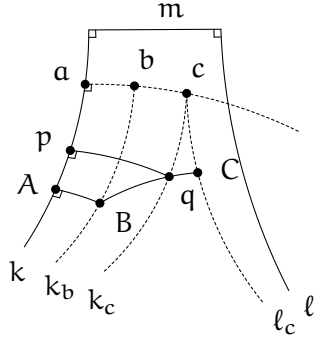and edge lengths agree is straightforward.

FIGURE 13. The proof of Lemma 3.6.

This completes the proof of the Green-Black Arm Lemma modulo the verification of the lemma following. In applying the lemma, we set $p_1 = b$, $p_n^* = c$, $p_1^{**} = B$ and $p_n^{**} = C$.

Lemma 3.6. *Let* R *be the open convex region in* $\mathbb{H}^2$ *bounded by the hyperbolic rays* k *and* ℓ *and the hyperbolic segment* m, *with* m *orthogonal to both* k *and* ℓ. *Let* $c \in R$ *and let* a *be that unique point on* k *for which the segment* ac *is orthogonal to* k. *Choose any point* b *in* ac *other than one of the endpoints. Let* B *and* C *be points in* R *such that the hyperbolic distance from* B *to* k *is equal to that from* b *to* k, *and the hyperbolic distance from* C *to* ℓ *is equal to that from* c *to* ℓ. *Assume further that* B *and* C *lie in the half-plane bordered by the line through* b *and* c *that does not meet* m. *Then* $|bc| \leqslant |BC|$, *with equality if and only if* $b = B$ *and* $c = C$. *See Fig. 13.*

Proof. The proof uses the labeling of Fig. 13. The dotted lines labeled k_b and k_c are the hypercycles determined by the line k and the respective points b and c; similarly, ℓ_c is the hypercycle determined by ℓ and c. The point B lies on k_b and the point C lies on ℓ_c. Let A be the unique point on k for which AB is orthogonal to k and let q be the point of intersection of the hypercycle k_c with the segment BC. Finally, let p be the unique point on k for which pq is orthogonal to k. Then $|ab| + |bc| = |ac| = |pq| \leqslant |AB| + |Bq|$. Since $|ab| = |AB|$, we conclude that $|bc| \leqslant |Bq| \leqslant |Bq| + |qC| = |BC|$.

The verification that $|bc| = |BC|$ implies $b = B$ and $c = C$ is left as an easy exercise. $\square$

4. The Proof of the Main Theorem

The remainder of the argument is Cauchy's. Let $G(\mathcal{C})$ and $G(\mathcal{C}')$ be two proper, convex, non-unitary c-polyhedra, both based on the same abstract spherical polyhedron P with 1-skeleton $G = P^{(1)}$, that have Möbius-congruent c-faces. Assume that $G(\mathcal{C})$ and $G(\mathcal{C}')$ are not Möbius-congruent to one another. Label the edge e of P that is adjacent to the faces f and g as follows. When the oriented ortho-circles O_f^+ and O_g^+ meet at an angle larger than the angle of intersection of $O_f'^+$ and $O_g'^+$, label with a plus sign, a minus sign if smaller, and no sign if equal; or if $O_f'^+$ and $O_g'^+$ fail to meet, then a plus sign if $\langle O_f{}^+, O_g{}^+ \rangle > \langle O_f'{}^+, O_g'{}^+ \rangle$, a minus sign if $\langle O_f{}^+, O_g{}^+ \rangle < \langle O_f'{}^+, O_g'{}^+ \rangle$, and no sign otherwise.[17] By Theorem 1.2, there is at least one edge of P labeled with a plus or a minus sign. Apply Cauchy's Combinatorial Lemma 2.3 to conclude that there is a vertex v that is incident to at least one edge labeled with a plus or minus sign for which one encounters at most two sign changes in labels on the edges adjacent to v as one walks around the vertex.

Consider now the green-black polygon $L(v)$, the c-link of circle C_v in $G(\mathcal{C})$. Label a green vertex or a green edge of $L(v)$ with a plus sign if the interior angle or edge-length is larger than the corresponding one in $L(v)'$, a minus sign if it is smaller, and no label if equal. By Remark 1.5 the green vertex or edge of $L(v)$ determined by the edge e of P has the same sign that the preceding paragraph assigns to the edge e, and so there are at most two sign changes on a walk around the c-link $L(v)$. But this contradicts the Green-Black Polygon Four Vertex Lemma 3.3, which guarantees at least four sign changes. Therefore, the c-polyhedra $G(\mathcal{C})$ and $G(\mathcal{C}')$ are Möbius-congruent, verifying the Main Theorem.

5. Hyperideal Polyhedra in $\mathbb{H}^3$

Convex hyperbolic polyhedra have been a topic of interest since Poincaré's 1881 study [16] and Dehn's 1905 paper [10]. Andre'ev in two 1970 papers classified all compact convex hyperbolic polyhedra with acute dihedral angles [1] and studied non-compact ones with ideal vertices [2]. Thurston [20] in the Princeton notes for his 1978-79 course on hyperbolic 3-manifolds recovered and popularized Andre'ev's results in the context of circle packings on the Riemann sphere. In 1993, Hodgsen and Rivin characterized all compact ones [13], even with dihedral angles greater than $\pi/2$, and in 1996 Rivin [17] characterized the ones with ideal vertices. Bao and Bonahon [3] in 2002 characterized convex hyperideal hyperbolic polyhedra in terms of their combinatorial type and their dihedral angles. The results

---

[17]Notice that in terms of the inversive distance, when O_f^+ and O_g^+ meet, e gets a plus sign when $\langle O_f{}^+, O_g{}^+ \rangle < \langle O_f'{}^+, O_g'{}^+ \rangle$.

of Andre'ev, Thurston, Hodsgen-Rivin, Rivin, and Bao-Bonahon give conditions that guarantee existence of convex polyhedra in $\mathbb{H}^3$ as well as their global rigidity—uniqueness up to hyperbolic isometries.

Our Main Theorem implies nothing about existence, but does imply the global rigidity of certain generalized convex hyperideal polyhedra in $\mathbb{H}^3$. We first recall the Bao-Bonahon definition of hyperideal polyhedron. We use the Klein projective model of $\mathbb{H}^3$ where the hyperbolic 3-space is identified with the open unit ball B^3 in $\mathbb{E}^3 \subset \mathbb{RP}^3$. In this model, the hyperbolic lines and totally geodesic hyperbolic planes are the intersections of Euclidean lines and planes of $\mathbb{E}^3$ with $\mathbb{H}^3 = B^3$. A *hyperideal polyhedron* P is the intersection of $\mathbb{H}^3$ with a compact convex projective polyhedron P$'$ of $\mathbb{RP}^3$ with two properties: (1) no vertex of P$'$ is in $\mathbb{H}^3$, and (2) every edge of P$'$ meets $\mathbb{H}^3$ nontrivially. Our interest is in a generalization of hyperideal polyhedra to regions P obtained by intersecting $\mathbb{H}^3$ with a compact convex projective polyhedron P$'$ of $\mathbb{RP}^3$ with the two properties: (1$'$) no vertex of P$'$ is in the closed unit ball $\mathbb{H}^3 \cup \partial\mathbb{H}^3 = B^3 \cup \mathbb{S}^2$ and (2$'$) every face of P$'$ meets $\mathbb{H}^3$ nontrivially. In particular, we allow that some or all of the edges of P$'$ lie outside the closed unit ball. The terminology *hyperideal polyhedron* suggests a polyhedron whose vertices either lie at infinity (=*ideal*), or lie beyond (=*hyper*) infinity, per condition (1). The edges and faces do not lie, at least not entirely, beyond infinity per condition (2). We will use the term *strictly-hyperideal polyhedron*[18] to mean those regions P obtained by intersecting $\mathbb{H}^3$ with compact convex projective polyhedra P$'$ that satisfy properties (1$'$) and (2$'$). All vertices lie beyond infinity per condition (1$'$), and the edges may or may not lie entirely beyond infinity, but the faces do not lie, at least not entirely, beyond infinity per condition (2$'$). If in addition to conditions (1$'$) and (2$'$), P$'$ satisfies condition (3$'$) no edge of P$'$ is tangent to $\partial\mathbb{H}^3 = \mathbb{S}^2$, we say that the strictly-hyperideal polyhedron P is *non-unitary*.

Let P be a strictly-hyperideal polyhedron obtained from P$'$ and let G be the 1-skeleton of the polyhedron P* dual to P$'$. We describe a c-framework $G(\mathcal{C})$ as follows. Let u be a vertex of G. Then u is a face of P$'$ and as such has a support plane Π_u. Define the oriented circle C_u by $C_u = \Pi_u \cap \mathbb{S}^2$, oriented so that its companion disk D_u is the disk on $\mathbb{S}^2$ it bounds that does not meet the interior of P$'$.[19] Let $\mathcal{C} = \{C_u : u \in V(G)\}$.

---

[18]Bao and Bonahon [3] use this term to indicate that conditions (1$'$) and (2) adhere. We loosen the term by allowing full edges to lie beyond infinity while still requiring condition (2$'$).

[19]Since all the vertices of P$'$ lie exterior to the closed unit ball exactly one of the two complementary domains of C_u in $\mathbb{S}^2$ meets the interior of P$'$. Another way to describe the orientation on C_u is using the unit normal vector n_u in $\mathbb{E}^3$ to the plane $\Pi_u \cap \mathbb{E}^3$ that points exterior to P$'$. The circle C_u is oriented so that n_u points toward D_u, or what is the same, so that C_u is oriented counter-clockwiase looking from the tip of n_u.

Lemma 5.1. *The c-framework* $G(\mathcal{C})$ *is a convex c-polyhedron, and is non-unitary if* P *is non-unitary.*

Proof. Let f be a face of P* with vertices $u_1, \ldots, u_n$. Then $u_1, \ldots, u_n$ are the faces of P' that are incident at the vertex f of P' with respective support planes Π_{u_i}, for $i = 1, \ldots, n$. Since P' is strictly-hyperideal, the vertex f of P' lies outside the closed ball $B^3 \cup \mathbb{S}^2$. Let O be the circle of intersection of $\mathbb{S}^2$ with the cone with vertex f that circumscribes $\mathbb{S}^2$. Since each support plane Π_{u_i} passes through f, the circle $C_{u_i} = \Pi_{u_i} \cap \mathbb{S}^2$ meets O orthogonally. It follows that the *c*-face $\mathcal{C}_f = \{C_{u_i} : i = 1, \ldots, n\}$ is *c*-planar with ortho-circle O. That $\mathcal{C}_f$ is not coaxial is a consequence of the fact that the vertices $u_1, \ldots, u_n$ are not collinear and the fact that the *c*-framework $G(\mathcal{C})$ is edge-uncoupled is a consequence of the convexity of P'. It follows that $G(\mathcal{C})$ is a *c*-polyhedron, and it remains to prove that it is convex.

To see that $G(\mathcal{C})$ is convex, recall that as P' is a convex projective polyhedron, it is the closure of one of the components of the complement of the support planes of P' that does not contain a projective line. Normalize P' by applying, if needed, a projective transformation set-wise fixing the 2-sphere $\mathbb{S}^2$ and moving the ortho-circle O to the equator of $\mathbb{S}^2$, which moves the vertex f of P' to infinity, by which we mean that f lies in the complement of $\mathbb{E}^3$ in $\mathbb{RP}^3$. Of course such a projective transformation acts as a Möbius transformation of $\mathbb{S}^2$. Let P'' be the closure of the component of the complement of the n support planes $\Pi_{u_1}, \ldots, \Pi_{u_n}$ that contains P', and note that our normalization implies that there is a projective line containing the vertex f of P' that runs parallel to the z axis and is contained in P''. We need to show that O may be oriented to O^+ so that all the oriented circles of the *c*-framework $G(\mathcal{C})$ are segregated from O^+. If this is not possible, then there are faces v and w of P' whose corresponding circles C_v and C_w are not both segregated from O no matter which orientation is assigned to O. This means, without loss of generality, that the companion disk D_v of C_v overlaps the northern hemisphere of $\mathbb{S}^2$ more than the southern, and that the companion disk D_w of C_w overlaps the southern more than the northern. From the definition of the circles C_v and C_w, since these disks D_v and D_w do not meet the interior of P', it follows that P' is contained in a region Q, bounded in $\mathbb{E}^3$, cut out by the planes $\Pi_{u_1}, \ldots, \Pi_{u_n}, \Pi_v, \Pi_w$. But then f, which lies at infinity, cannot lie in P' since Q is bounded in $\mathbb{E}^3$, a contradiction. $\square$

Not all strictly-hyperideal polyhedra are rigid. The work of the present paper guarantees that one ingredient that ensures rigidity is that of properness. The strictly-hyperideal polyhedron $P = P' \cap \mathbb{H}^3$ is *proper* provided its corresponding *c*-polyhedron $G(\mathcal{C})$ is proper. Recall from the introduction that the *rigidity* of P means that it is determined up to isometries of $\mathbb{H}^3$ by the combinatorics of P' and the hyperbolic isometry classes of the faces of P, which are determined by the faces of the dual polyhedron P*. This means
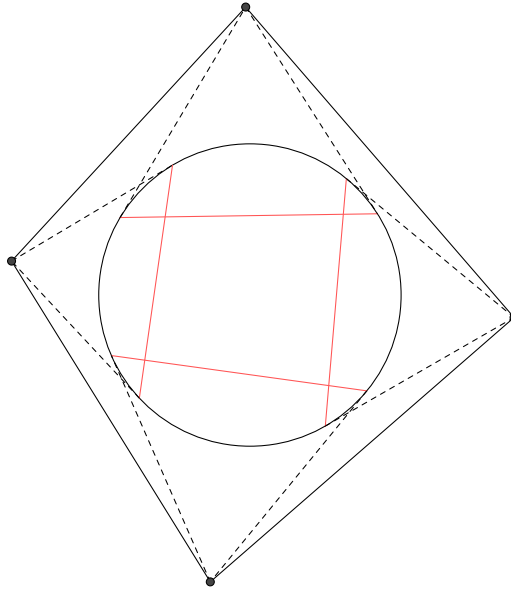
FIGURE 14. In Corollary 5.3, the c-link is a compact hyperbolic polygon. This is the projection of the upper hemisphere of $\mathbb{S}^2$ to the equatorial plane. The red lines are the ortho-circles and cut out a compact convex hyperbolic polygon.

that combinatorics and hyperbolic equivalence of faces determines the complex angles in which adjacent faces meet, which in turn determines how P' meets $\mathbb{H}^3$ up to projective equivalence that fixes $\mathbb{S}^2$.

Theorem 5.2. *Proper, non-unitary, strictly-hyperideal polyhedra in $\mathbb{H}^3$ are globally rigid, unique up to hyperbolic isometries.*

Proof. Apply The Main Theorem of the introduction to the corresponding c-polyhedra $\mathrm{G}(\mathcal{C})$. $\qquad\square$

Theorem 5.2 offers an alternate proof of two corollaries that follow from Bao and Bonahon's work. The proofs are obtained from the observation that in the setting of either of the corollaries, the obviously non-unitary strictly-hyperideal polyhedron P automatically is proper. The details of the proofs are left to the reader, but to close out this paper, we do provide a nice "proof by picture" in each case.

Corollary 5.3 (Bao-Bonahon [3]). *Let P be a strictly-hyperideal polyhedron in $\mathbb{H}^3$ for which no edge of P' meets the boundary $\partial\mathbb{H}^3$. Then P is globally rigid.*
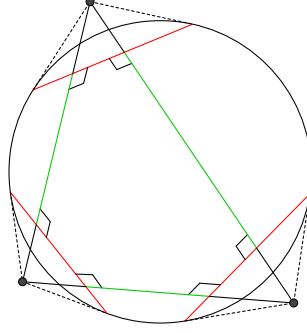
FIGURE 15. In Corollary 5.4, the c-link is a right-angled green-black polygon with alternating green and black sides. This is the projection of the upper hemisphere of $\mathbb{S}^2$ to the equatorial plane. The black sides are colored red for emphasis, and are the ortho-circles, and the green ones are the orthogonal segments between successive ortho-circles.

Proof. Obviously P is non-unitary. We show that P is proper and apply Theorem 5.2. Let u be a vertex of P^* with corresponding circle $C_u = \Pi_u \cap \mathbb{S}^2$ in the c-framework $G(\mathcal{C})$. By applying a projective transformation that fixes $\mathbb{S}^2$ if necessary, assume that the vertex u lies at infinity in the direction of the north pole of $\mathbb{S}^2$ so that C_u is the equator of $\mathbb{S}^2$ with Π_u the equatorial plane and the companion disk D_u the upper hemisphere. The vertex u of P^* now considered as a face of P' is then a bounded convex Euclidean polygon in the equatorial plane Π_u that separates the equator from infinity, none of whose edges meets C_u. Fig. 14 is a view of D_u from above, projected to the plane Π_u. The ortho-circles that meet D_u are precisely the intersections of the cones that inscribe $\mathbb{S}^2$ whose vertices are the vertices of the face u of P'. This implies that the ortho-circles of the c-polyhedron $G(\mathcal{C})$ that meet the companion disk D_u cut out a compact convex hyperbolic polygon in D_u with its hyperbolic metric, and so the c-link of the vertex u is a compact convex hyperbolic polygon and $G(\mathcal{C})$ is proper at u. In this case the c-link has no green edges. $\qquad\square$

Corollary 5.4 (Bao-Bonahon [3]). *Let P be a strictly-hyperideal polyhedron in $\mathbb{H}^3$ for which every edge of P' meets $\mathbb{H}^3$. Then P is globally rigid.*

Proof. The proof proceeds exactly as in the preceding lemma, except that all the edges of the face u of P' meet the the open unit disk in Π_u. This implies, as in Fig. 15, that the ortho-circles are pairwise disjoint. In the hyperbolic plane D_u, the common orthogonal segment to two adjacent ortho-circles with cone points a and b, adjacent vertices of the face u, project orthogonally to the edge between a and b. These give the green segments

that alternate with the segments that lie along the ortho-circles to cut out a compact convex green-black polygon with alternating edge colors. Again, the c-link of the vertex u is a compact convex hyperbolic polygon and G($\mathcal{C}$) is proper at u. $\qquad\square$

References

[1] E.M. Andre'ev. On convex polyhedra in lobachevski spaces. *Mat. Sbornik*, 81(123):445–478, 1970.

[2] E.M. Andre'ev. On convex polyhedra of finite volume in lobachevski spaces. *Mat. Sbornik*, 83(125):256–260, 1970.

[3] Xiliang Bao and Francis Bonahon. Hyperideal polyhedra in hyperbolic 3-space. *Bull. Soc. Math. France*, 130(3):457–491, 2002.

[4] John C. Bowers and Philip L. Bowers. Ma-Schlenker c-octahedra in the 2-sphere. 2016.

[5] John C. Bowers and Philip L. Bowers. Iso-inversive embeddings of circles into the 2-sphere. 2017.

[6] John C. Bowers and Philip L. Bowers. Rigidity of circle configurations in the planes and the sphere. 2017.

[7] Philip L. Bowers and Monica K. Hurdal. Planar conformal mappings of piecewise flat surfaces. *Visualization and Mathematics III*, (Chapter 1):3–34, 2003.

[8] Philip L. Bowers and Kenneth Stephenson. Uniformizing dessins and Belyĭ maps via circle packing. *Memoirs of the AMS*, 170(805):1–97, 2004.

[9] A. Cauchy. Sur les polygones et les polyhedres. *J. Ecole Polytechnique XVIe Cahier*, pages 87–98, 1813.

[10] Max Dehn. Die eulersche formel im zusammenhang mit dem inhalt in der nicht-euklidischen geometrie. *Math. Ann.*, 61:561–586, 1905.

[11] Dmitry Fuchs and Serge Tabachnikov. *Mathematical Omnibus: Thirty Lectures on Classic Mathematics*. American Mathematical Society, 2007.

[12] Ren Guo. Local rigidity of inversive distance circle packing. *Transactions of the American Mathematical Society*, 363(9):4757–4776, September 2011.

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

[14] Feng Luo. Rigidity of polyhedral surfaces, III. *Geometry & Topology*, 15(4):2299–2319, December 2011.

[15] Jiming Ma and Jean-Marc Schlenker. Non-rigidity of spherical inversive distance circle packings. *Discrete & Computational Geometry*, 47(3):610–617, February 2012.

[16] Henri Poncaré. Sur les groupes kleinéens. *C. R. Acad. Sci. Paris*, 93:44–46, 1881.

[17] Igor Rivin. A characterization of ideal polyhedra in hyperbolic 3-space. *Ann. of Math.*, 143:51–70, 1996.

[18] Mathais Rousset. Sur la rigidité de polyèdres hyperboliques en dimension 3 : cas de volume fini, cas hyperidéal, cas fuchsien. *Bulletin de la Société Mathématique de France*, 132(2):233–261, 2004.

[19] Hans Schwerdtfeger. *Geometry of Complex Numbers*. Number 13 in Mathematical Expositions. Univiversity of Toronto Press, 1962.

[20] William P. Thurston. The geometry and topology of 3-manifolds. Lecture Notes: Princeton University, 1980.

DEPARTMENT OF COMPUTER SCIENCE, JAMES MADISON UNIVERSITY, HARRISONBURG VA 22807

E-mail address: bowersjc@jmu.edu

DEPARTMENT OF MATHEMATICS, THE FLORIDA STATE UNIVERSITY, TALLAHASSEE FL 32306

E-mail address: bowers@math.fsu.edu

DEPARTMENT OF MATHEMATICS, UNIVERSITY OF CONNECTICUT, STORRS CT 06269

E-mail address: kevin.pratt@uconn.edu