

APPLYING POINCARÉ'S POLYHEDRON THEOREM TO GROUPS OF HYPERBOLIC ISOMETRIES

KYLER SIEGEL

ABSTRACT. We present a computer algorithm which confirms that an approximately discrete subgroup of $PSL(2, \mathbb{C})$ is in fact discrete. The algorithm proceeds by constructing the Dirichlet domain of the subgroup in H^3 , and then checks that the hypotheses of Poincaré's Theorem for Fundamental Polyhedra are satisfied. In order to check that the hypotheses are exactly satisfied, we rely on group theoretical properties resulting from certain geometric conditions of the Dirichlet domain. We begin in Section 1 with a review of relevant background material. In Section 2 we provide a formal statement of the problem, including the restrictions we impose upon the subgroup of $PSL(2, \mathbb{C})$. In Section 3 we introduce the geometric conditions which must apply for our algorithm to hold, and we prove that these are generically satisfied. We then show in Section 4 that these conditions are in fact sufficient to verify the hypotheses of Poincaré's Theorem. In Section 5 we look at the field containing our matrix entries.

1. BACKGROUND

1.1. Models of Hyperbolic Space. Let H^n denote n -dimensional hyperbolic space. We shall represent the upper half space model of three dimensional hyperbolic space using a subset of the quaternions \mathbb{H} :

$$U^3 := \{x + yi + zj \in \mathbb{H} : z > 0\}. \quad (1.1)$$

We view $\widehat{\mathbb{C}} := \mathbb{C} \cup \{\infty\}$ as the boundary of U^3 in the obvious way by associating quaternions in U^3 of the type $x + yi$ with complex numbers, and adjoining an element ∞ in the usual manner. We sometimes refer to $\widehat{\mathbb{C}}$ as the *boundary (or sphere) at infinity*. The geodesics in this model are Euclidean half-circles and half-lines which are orthogonal to \mathbb{C} .

Remark 1.1. Since planes (lines) will most often arise as limiting cases of spheres (circles), we will refer to an n -dimensional sphere or plane as a *generalized n -sphere* or simply *G - n -sphere*. We will also sometimes add the prefixes *E* or *H* to distinguish between Euclidean and hyperbolic objects (metrics, generalized spheres, etc) when confusion may arise.

Recall that we endow U^3 with the Riemmanian metric

$$ds = \frac{dx^2 + dy^2 + dz^2}{z}. \quad (1.2)$$

Date: August 4, 2009.

*This project is based on work done as part of the 2009 Summer Undergraduate Research Program hosted at Columbia University. The project was advised by Maksim Lipyanskiy and Eric Potash and based on their previous work. I wish to thank M. Lipyanskiy for his continued advising and guidance on this project.

The distance function on U^n is given by

$$\cosh d_U(x, y) = 1 + \frac{|x - y|^2}{2x_n y_n}.$$

We shall denote by B^n the n -dimensional *conformal ball model* of hyperbolic space. The isometry relating U^n and B^n is a Möbius transformation $\eta : U^n \rightarrow B^n$ which we call the *standard transformation* from U^n to B^n .

Lastly, let D^n denote the n -dimensional *projective disk model* of hyperbolic space. The isometry $\phi : D^3 \rightarrow U^3$ relating the three dimensional projective disk and upper half space models is given by

$$\phi(r_x, r_y, r_z) = \frac{r_x + r_y i + \sqrt{1 - r_x^2 - r_y^2 - r_z^2} j}{1 - r_z} \quad (1.3)$$

with inverse

$$\phi^{-1}(x + yi + tj) = \frac{(2x, 2y, x^2 + y^2 + t^2 - 1)}{1 + x^2 + y^2 + t^2}. \quad (1.4)$$

The metric d_{D^3} induced on D^3 is given by

$$\cosh(d_{D^3}(r, r')) = \frac{1 - r \cdot r'}{\sqrt{1 - |r|^2} \sqrt{1 - |r'|^2}}. \quad (1.5)$$

More details on the models of hyperbolic space and the transformations between them can be found in many references, for example [1] or [4].

1.2. Möbius Transformations and the Isometries. Recall that a Möbius transformation of $\widehat{E}^n := E^n \cap \{\infty\}$ (here E^n denotes ordinary n -dimensional Euclidean space) is a finite composition of reflections in \widehat{E}^n about generalized $(n - 1)$ -spheres. The Möbius transformations of \widehat{E}^n form a group $Möb(\widehat{E}^n)$. The isometries (i.e. distance preserving bijections) of U^n are precisely the Möbius transformations of \widehat{E}^n (restricted to U^n) which preserve U^n . These are generated by reflections about generalized spheres which are orthogonal to the boundary at infinity of U^n (which is of course associated with \widehat{E}^{n-1}). In other words, the group of isometries of U^n is isomorphic to $Möb(\widehat{E}^{n-1})$.

In this paper we are concerned with the *orientation-preserving* transformations, i.e. those can be written as a composition of an even number of reflections about generalized spheres. Henceforth in this paper all Möbius transformations and isometries will be taken to be orientation-preserving ones unless otherwise stated.

An element $g = \begin{pmatrix} a & b \\ c & d \end{pmatrix}$ of $SL(2, \mathbb{C})$ acts on an element $q = z + tj \in H^3$ (here $z \in \mathbb{C}$ and $t > 0$) by

$$\begin{aligned} gq &:= (aq + b)(cq + d)^{-1} \\ &= \frac{(az + b)\overline{(cz + d)} + a\bar{c}t^2 + |ad - bc|tj}{|cz + d|^2 + |c|^2t^2} \end{aligned} \quad (1.6)$$

(see [1] Ch.4).

It is well-known that elements of $SL(2, \mathbb{C})$ represent Möbius transformations of $\widehat{\mathbb{C}} = \widehat{E}^2$, and in fact $Möb(\widehat{E}^2) \cong PSL(2, \mathbb{C}) := SL(2, \mathbb{C})/\{\pm I\}$. Hence, by the above discussion, the isometry group of H^3 is $PSL(2, \mathbb{C})$, with the explicit action in upper half space model given by (1.6).

Every non-identity $g \in PSL(2, \mathbb{C})$ has at least one fixed point in $\overline{H^3} := H^3 \cup \widehat{\mathbb{C}}$. If g has a fixed point in H^3 , we call g *elliptic*; if g is not elliptic, and g has exactly one fixed point in $\widehat{\mathbb{C}}$, then we call g *parabolic*; otherwise g has exactly two fixed points in $\widehat{\mathbb{C}}$ and is called *loxodromic*. Note that every non-identity power of g has the same fixed points as g and therefore the same classification.

Remark 1.2. Some authors use the term *hyperbolic* in place of our loxodromic. For us, note that any loxodromic element g can be conjugated such that it is of the form $q \mapsto kq$, with $k \in \mathbb{C}$. We shall then call g *hyperbolic* if $k \in \mathbb{R}$, and *strictly loxodromic* otherwise.

Let $tr^2(g)$ be the trace squared of (the matrix representation of) g . Note that tr^2 is a well-defined function on $PSL(2, \mathbb{C})$, although tr is not. It can be shown that g is parabolic if and only if $tr^2(g) = 4$, elliptic if and only if $tr^2(g) \in [0, 4)$, and loxodromic otherwise. We observe that every element in the same conjugacy class as g has the same classification.

1.3. The Cross Ratio. We define the quaternionic cross ratio of four points in $\widehat{\mathbb{H}} := \mathbb{H} \cup \{\infty\}$ by:

$$C(q_1, q_2, q_3, q_4) := (q_1 - q_3)(q_1 - q_4)^{-1}(q_2 - q_4)(q_2 - q_3)^{-1}, \quad (1.7)$$

as in [2]. The cross ratio will play an important role in this paper. With $f, g, h \in PSL(2, \mathbb{C})$, we will be particularly interested in the function $C_{f,g,h} : U^3 \rightarrow \mathbb{H}$ given by

$$\begin{aligned} C_{f,g,h}(q) &:= C(q, fq, gq, hq) \\ &= (q - gq)(q - hq)^{-1}(fq - hq)(fq - gq)^{-1}. \end{aligned} \quad (1.8)$$

1.4. Discrete Groups and the Dirichlet Domain. Let

$$\Phi : SL(2, \mathbb{C}) \rightarrow PSL(2, \mathbb{C})$$

be the projection map. Let us give $SL(2, \mathbb{C})$ its natural metric topology as a subset of \mathbb{C}^4 and use Φ to induce a quotient topology on $PSL(2, \mathbb{C})$. We call a subset $G \subset PSL(2, \mathbb{C})$ *discrete* if G is discrete with respect to this quotient topology. Equivalently, a subgroup G of $PSL(2, \mathbb{C})$ is discrete if, for every compact subset K of H^3 , we have

$$g(K) \cap K = \emptyset, \quad (1.9)$$

except for a finite number of $g \in G$.

Now assume that G is a discrete group of isometries of H^3 and let w be some element of H^3 which is not fixed by any non-identity element of G . For each non-identity $g \in G$, the set of all points H -equidistant to w and gw is a 2- H -plane (i.e. a generalized 2- E plane in \widehat{E}^3 intersecting the boundary at infinity orthogonally),

denoted by $P(w, gw)$. Let $D(w, gw)$ denote the half space consisting of all points in H^3 which are H-closer to w than to gw . That is,

$$D(w, gw) := \{q \in H^3 : d_H(q, w) < d_H(q, gw)\},$$

where d_H is the hyperbolic distance function. We define the *Dirichlet domain* D , centered at w , to be $\bigcap D(w, gw)$, where the intersection is over all non-trivial $g \in G$. D has a natural cell decomposition in terms of the half spaces defining it. The codimension 1 cells are called *sides*, the codimension 2 cells are called *edges*, and the codimension 3 cells are called *vertices*. Note that each edge is the intersection of exactly two sides.

Remark 1.3. We have defined the Dirichlet domain in H^3 and will think of it as existing simultaneously the different models of H^3 , most often working in U^3 .

The Dirichlet domain D is a *convex fundamental polyhedron* for G , and we make the following observations:

- (1) For every non-trivial $g \in G$, $g(D) \cap D = \emptyset$.
- (2) For every $q \in H^3$, there is a $g \in G$ such that $g(q) \in \overline{D}$.
- (3) The sides of D are paired together by elements of G . That is, for each side s , there is a (distinct) partner side s' and an element $g_s \in G$ such that $g_s(s) = s'$. Also, $g_{s'} = g_s^{-1}$ and $(s')' = s$. We call the group elements g_s the *side pairing transformations*.
- (4) Any compact set meets only finitely many G -translates of D .
- (5) D is H-convex (being an intersection of half spaces).

We also observe that the edges of the Dirichlet domain D appear in *edge cycles*. That is, let $e = e_1$ be some edge belonging to a side s_1 . Let $e_2 = g_{s_1}(e_1)$, and let s_2 be the side adjacent to s_1' which contains e_2 as an edge. Similarly, let $e_3 = g_{s_2}(e_2)$, and let s_3 be the side adjacent to s_2' which contains e_3 as an edge, and so on. We find that the sequence e_1, e_2, e_3, \dots is periodic, with some smallest period k . Then e_1, e_2, \dots, e_k is called the *edge cycle* of e , and $g_{s_1}, g_{s_2}, \dots, g_{s_k}$ is the corresponding cycle of side pairing transformations. We have

$$(g_{s_k} \circ \dots \circ g_{s_1})^t = 1, \tag{1.10}$$

for some positive integer t , and that

$$\sum_{m=1}^k \alpha(e_m) = 2\pi/t, \tag{1.11}$$

where $\alpha(e)$ is the angle at e measured inside D . In fact, $t = |G_e|$, where

$$G_e := \{g \in G : g(e) = e\}. \tag{1.12}$$

1.5. Poincaré's Theorem. Whereas the Dirichlet domain is a fundamental polyhedron constructed from a discrete group of isometries, Poincaré's theorem involves the reverse process. It states that, given a polyhedron D satisfying certain properties, we can construct a discrete group of isometries for which D is a fundamental polyhedron. For the precise statement of Poincaré's theorem we will follow the formulation by Maskit (see [3]). We begin with \mathbb{X} and \mathbb{G} , where \mathbb{X} is either

H^n or E^{n-1} , and \mathbb{G} is the isometry group of \mathbb{X} . Here $n \geq 2$. We assume that D is a convex polyhedron with finitely many sides (i.e. an intersection of finitely many open half spaces in \mathbb{X}) which are paired by elements of \mathbb{G} . That is, for each side s there is a side s' and an element $g_s \in \mathbb{G}$ (called a *side pairing transformation*) such that:

- (1) $g_s(s) = s'$.
- (2) $g_{s'} = g_s^{-1}$.
- (3) $g_s(D) \cap D = \emptyset$.

As before, each edge of D belongs to an edge cycle, which is finite. Let $\{e_1, \dots, e_k\}$ be an edge cycle and let g_1, \dots, g_k be the corresponding side pairing transformations, where k is the common period. Then the *cycle transformation* $h := g_k \circ \dots \circ g_1$ keeps e_1 invariant. For a choice of an edge e_1 , we get either h or h^{-1} as the corresponding cycle transformation, depending on which initial side s_1 containing e_1 as an edge we choose. We require:

- (4) For each edge e , there is a positive integer t so that $h^t = 1$.

The relations in G of the form $h^t = 1$ are called the *cycle relations*.

For the next condition, let $\alpha(e)$ be the angle, measured from inside D , at the edge e . For each edge cycle we require:

$$(5) \sum_{m=1}^k \alpha(e_m) = 2\pi/t$$

For the last condition we are concerned with the completeness of the quotient space obtained by gluing the sides together. Note that we have not excluded the possibility that sides of D extend out to infinity. Suppose two sides are tangent at a point x_1 on the sphere at infinity. That is, the two sides do not intersect in \mathbb{X} but both get arbitrarily ϵ -close to x_1 . Note that this can only occur if $\mathbb{X} = H^n$. Call one of the sides s_1 , let g_1 be the corresponding side pairing transformation with $g_1(s_1) = s'_1$, and let $x_2 = g_1(x_1)$. If x_2 is a point of tangency of s'_1 and some other side s_2 , then let g_2 be the side pairing transformation with $g_2(s_2) = s'_2$, let $x_3 = g_2(x_2)$, and so on. If we eventually have $x_k = x_1$, for some finite k (i.e. $h := g_k \circ \dots \circ g_1$ leaves x_1 invariant), then we call x_1 an *infinite tangency point* and h the *infinite cycle transformation* at x_1 . We thus require:

- (6) Every infinite cycle transformation at every infinite tangency point is parabolic.

We can now state Poincaré's theorem:

Theorem 1.4 (Poincaré). *Let D be a finite sided polyhedron with side pairing transformations satisfying conditions (1) through (6). Then the group G generated by the side pairing transformations is discrete, the cycle relations form a complete set of relations for G , and D is a fundamental polyhedron for G .*

¹The proof is also valid for other spaces which will not be considered here, as well as for polyhedra with countably many sides

2. FORMAL STATEMENT OF PROBLEM

Our problem² can now be formally stated as follows:

Let G be a finitely presented discrete subgroup of $PSL(2, \mathbb{C})$ containing no elliptic elements. Given:

- A set of generators for G with a decimal floating point approximation of the matrix entries (complex numbers) of each generator to arbitrarily high precision
- An algorithmic solution to the word problem for G
- The ability to check whether an element $g \in G$, i.e. a word in the generators, is exactly parabolic (recall that, by checking tr^2 , we can show that g is approximately parabolic)

Prove that G is discrete.

In order to guarantee the success of our algorithm we place the following additional restrictions on G :

- For any $w \in H^3$, the Dirichlet domain $D(w)$ centered at w has finitely many sides.
- For any $w \in U^3$, the set $\overline{D(w)} \cap \widehat{\mathbb{C}}$ is finite, where $\overline{D(w)}$ denotes the closure of $D(w)$ in $\overline{U^3}$.
- For any $f, g, h \in G$ loxodromic and not all hyperbolic, the four points w, fw, gw, hw do not lie on an E-circle for every choice of $w \in G$.
- For any $f, g, h, i \in G$ not all parabolic with the same fixed point and not all hyperbolic with the same fixed point, the five points w, fw, gw, hw, iw do not lie on an E-sphere for every choice of $w \in G$.

3. GENERIC CONDITIONS

3.1. Constructing the Dirichlet Domain. Throughout this chapter, let G be a discrete subgroup of $PSL(2, \mathbb{C})$ obeying the conditions of section 2, and let $D(w)$ be the Dirichlet domain centered at $w \in H^3$ which corresponds to the group G . The idea behind our solution is as follows: we construct the approximate Dirichlet domain $D(w)$ centered at some $w \in H^3$ by taking various group elements $g \in G$ and intersecting their respective half spaces $D(w, gw)$. Here our computer program uses our high accuracy decimal expansions to calculate approximate intersection points. A finite number of these half spaces determine the domain (i.e. all other half spaces contain the domain), and we stop once the domain approximately satisfies the conditions of Poincaré's theorem. It turns out that, by using a good method for deciding in what order to consider the elements of G (as words in the generators), the domain can be fully constructed by a computer in a reasonably short period of time. This gives us an approximate set of the vertices of $D(w)$ and reasonable guesses as to the side pairing relations of $D(w)$. If we can then

²First work on computer applications of Poincaré's Theorem was initiated by Robert Riley, who used his computer program to prove that certain knot complements were hyperbolic. Since he uses a Ford Domain, many of the precision issues in his program need to be handled on a case-by-case basis. See [5] for this alternative approach.

somehow check that the hypotheses of Poincaré's theorem are exactly satisfied, and verify that the side pairing relations do in fact generate the group G , then it will follow from Poincaré's theorem that G is discrete.

One initial problem that may arise is that there are vertices of $D(w)$ belonging to more than 3 sides. In this case, regardless of the accuracy we use, we can never be certain from our approximation whether the sides determine exactly one point of intersection or merely two or more very close points of intersection. This issue will be addressed in sections 3.4 and 3.3, which together show that for generic w this issue does not arise for finite vertices and can be handled easily for infinite vertices.

The calculations involved in the above construction of $D(w)$ are tedious and impractical for a human, but are a straightforward procedure for a computer to perform. Of course we cannot rigorously verify the hypotheses of Poincaré's theorem using only our approximate construction of $D(w)$, and so we will need to find some other methods. Our main task is to verify that the sides of our approximate $D(w)$ are exactly mapped to each other by the side pairing transformations. Once we have this the other conditions will follow relatively straightforwardly. In principle, we could at this point attempt to construct the relevant field containing all vertices of $D(w)$ and then verify the hypotheses exactly using algebraic numbers. The field construction is discussed in section 5. Unfortunately this method is highly impractical for a computer routine, so we instead take a different approach.

We focus our attention on the edges of $D(w)$, in particular working to show that the unique geodesics containing each edge are appropriately mapped to each other. Our most useful tool will be the following, which we dub the **cycle length three method**:

Let e and e' be edges of $D(w)$ lying on geodesics $L(e)$ and $L(e')$ respectively, and suppose we wish to prove that $g(L(e)) = L(e')$. That is, we've approximately identified $g \in G$ as one of the side pairing transformations and seen that $g(e) = e'$ holds to high accuracy (note that each edge is determined the vertices at its endpoints, so in particular this means that the two defining vertices of e map to the two defining vertices of e' to high accuracy). Now suppose that e belongs to an edge cycle of length three, say $f(e), e, g(e) = e'$. Then e must be the edge lying on adjacent sides $S(g^{-1})$ and $S(f^{-1})$, where $S(g^{-1})$ denotes the side of $D(w)$ lying in the plane $P(w, g^{-1}w)$ (and similarly for $S(f^{-1})$). Also, e' should lie on $S(g)$ and some other adjacent side, say $S(h)$. A little thought shows that f^{-1}, g, h^{-1} is part of the cycle of side pairing transformations corresponding to the edge cycle $f(e), e, g(e) = e'$, and by so assumption it is the entire cycle. Then by (1.10), since we are assuming G has no elliptic elements, we have

$$h^{-1} \circ g \circ f^{-1} = 1. \quad (3.1)$$

Now since $L(e)$ lies on the H-planes $P(w, g^{-1}w)$ and $P(w, f^{-1}w)$, $L(e)$ is H-equidistant to $w, f^{-1}w$, and $g^{-1}w$ (in the sense that every point of $L(e)$ is equidistant to these three points). Hence $L(e)$ is the *unique* geodesic equidistant to $w, f^{-1}w$, and $g^{-1}w$. This implies that $g(L(e))$ is the unique geodesic equidistant to $gw, g \circ f^{-1}w$, and $g \circ g^{-1}w = w$. Similarly, $L(e')$ lies on $P(w, gw)$ and $P(w, hw)$,

and so $L(e')$ is the unique geodesic equidistant to w, gw , and hw . But from (3.1) we have $h = g \circ f^{-1}$, and so by uniqueness it follows that $g(L(e)) = L(e')$, which is the desired result.

Remark 3.1. One should be careful about the logic of the above argument, which is used throughout this paper. We began by strongly suspecting that $D(w)$ is a Dirichlet domain of a discrete group and that $g(e) = e'$. Under this suspicion we found a condition we should expect to hold, namely that $h = g \circ f^{-1}$. We can then verify explicitly, using the solution to the word problem, that $h = g \circ f^{-1}$ indeed holds. Then by our equidistance/uniqueness argument, we have in fact proven that $g(L(e)) = L(e')$. In this manner we are able to make rigorous deductions about our approximate $D(w)$ without being certain that our suspicions that G is discrete, $D(w)$ is a Dirichlet domain, etc. are true.

As to the usefulness of the cycle length three method, we are led to the obvious follow-up question:

Can we expect the edges to lie in cycles of length three?

This question is addressed in section 3.2 and 3.3, which together provide an affirmative answer for the generic case.

3.2. Generic Edge Cycles. We shall call an edge e of $D(w)$ *loxodromically (parabolically) paired* if there is a loxodromic (parabolic) isometry $f \in G$ such that $f(e)$ is also an edge of $D(w)$. Here f is called a *pairing* of e . We call e *exclusively loxodromically paired (exclusively parabolically paired)* if e is not parabolically paired (loxodromically paired).

We will show that, for generic $w \in H^3$, every edge e of $D(w)$ is either exclusively loxodromically paired or exclusively parabolically paired. In the former case e belongs to a cycle of length three, while in the latter case e belongs to either a cycle of length three (the “typical” case) or a cycle of length four. For edge cycles of length three we freely implement the cycle length three method. For edge cycles of length four, it will turn out that each pairing of e is parabolic with the same fixed point, and we can thus appeal to the simple properties of Euclidean Dirichlet domains to verify edge pairings. We begin with some lemmas.

Lemma 3.2. *Four pair-wise distinct points $q_1, q_2, q_3, q_4 \in U^3$ lie on a same (one-dimensional) generalized E -circle if and only if $C(q_1, q_2, q_3, q_4)$ is real.*

Proof. See Theorem 4.9 of [2]. □

Lemma 3.3. *Any 2 – H -sphere in U^3 is an E -sphere.*

Proof. Let S be a 2- H -sphere in U^3 . Using the standard isometry from the upper half space model to the ball model, we map S to a 2- H -sphere in B^3 . Now apply an isometry of B^3 which takes S to a 2- H -sphere S' centered at the origin. Now since rotations about the origin are hyperbolic isometries of B^3 and therefore fix S' , it is clear that S' is also a Euclidean sphere. But then S must be a Euclidean 2-sphere in U^3 , since the standard isometry from U^3 to B^3 is a Mobius transformation, as are the isometries of B^3 , and hence they map Euclidean spheres to Euclidean spheres. The result follows. □

Lemma 3.4. *Let $R : \widehat{\mathbb{C}} \rightarrow \widehat{\mathbb{C}}$ be any non-constant rational function. Then the set $\{z : R(z) \in \mathbb{R}\}$ has measure zero.*

Proof. We note that R is a local diffeomorphism at every point p of \mathbb{C} apart from the finite set of points $\{z_1, \dots, z_n\}$ at which the denominator of R vanishes. Then each p contains a neighborhood U such that $U \cap \{z : R(z) \in \mathbb{R}\}$ has measure zero. Since countably many such neighborhoods U cover $\mathbb{C} \setminus \{z_1, \dots, z_n\}$, the desired result follows. \square

Lemma 3.5. *Suppose that $f, g \in G$, with f loxodromic and g parabolic. Then f and G do not share a common fixed point in $\widehat{\mathbb{C}}$.*

Proof. Assume f and g share a fixed point. By applying an appropriate isometry, we may assume that f and g both fix ∞ . Then we can represent f and g as $f = \begin{pmatrix} a & b \\ 0 & 1 \end{pmatrix}$ and $g = \begin{pmatrix} 1 & c \\ 0 & 1 \end{pmatrix}$, where $a, b, c \in \mathbb{C}$ with $a, c \neq 0$ and $|a| \neq 1$. Now observe that G contains the element $f^{-n}gf^n = \begin{pmatrix} 1 & c/a^n \\ 0 & 1 \end{pmatrix}$. Since n can be chosen to be arbitrarily large and hence $f^{-n}gf^n$ can be made arbitrarily close to the identity isometry, this contradicts the discreteness of G . \square

Lemma 3.6. *For any $f, g \in G$, the fixed point sets of f and g are either disjoint or coincide.*

Proof. We see from Lemma 3.5 that f and g are either both loxodromic or both parabolic, and if they are both parabolic we are done. Now suppose that f and g are both loxodromic and share a fixed point. By applying an appropriate isometry, we may assume that f and g both fix ∞ . Then we can represent f and g by $f = \begin{pmatrix} a & b \\ 0 & 1 \end{pmatrix}$ and $g = \begin{pmatrix} c & d \\ 0 & 1 \end{pmatrix}$, where $a, b, c, d \in \mathbb{C}$, with $a, c \neq 0$ and $|a|, |c| \neq 1$. Now observe that G contains the element $fgf^{-1}g^{-1} = \begin{pmatrix} 1 & -d - bc + ad + b \\ 0 & 1 \end{pmatrix}$. But then we must have $fgf^{-1}g^{-1} = 1$, otherwise $fgf^{-1}g^{-1}$ is a parabolic element sharing the same fixed point as hyperbolic elements f and g , contradicting Lemma 3.5. Therefore we must have

$$d - bc + ad + b = 0 \implies \frac{b}{1-a} = \frac{d}{1-c}.$$

But this is precisely the statement that the finite fixed points of f and g are the same. \square

Lemma 3.7. *Let $f, g, h \in G$ be pairwise distinct isometries. If $C_{f,g,h} : U^3 \rightarrow \mathbb{H}$ is non-constant, then the set $E = \{q \in U^3 \mid C_{f,g,h}(q) \in \mathbb{R}\}$ has (Lebesgue) measure zero. Otherwise, if $C_{f,g,h}$ is constant, then f, g, h all share the same fixed point set.*

Proof. Recall that $C_{f,g,h}$ is given by

$$C_{f,g,h}(q) = (q - gq)(q - hq)^{-1}(fq - hq)(fq - gq)^{-1}.$$

From (1.6), it is clear that f, g, h can each be thought of as real analytic functions on $U^3 \subset \mathbb{R}^3$, as can the differences $(q - gq), (q - hq), (fq - hq), (fq - gq)$. Moreover, it is not hard to see that inverting non-vanishing quaternions is an analytic function from $\mathbb{H} \setminus \{0\}$ to itself. We conclude that $C_{f,g,h}$ is in fact a real analytic function from U^3 to \mathbb{R}^4 . Then the set E on which $C_{f,g,h}$ is real is just the intersection of the vanishing sets of the latter three components (i.e. the i, j, and k components) of $C_{f,g,h}$, each of which is real analytic on U^3 . Since the vanishing set of a non-constant real analytic function has measure zero, we see that E has measure zero unless the latter three components vanish identically, i.e. $C_{f,g,h}(q)$ is real for all $q \in U^3$.

Suppose that this is the case. Note that we can define $C_{f,g,h}(z)$ even when $z \in \widehat{\mathbb{C}}$, and $C_{f,g,h}(q)$ is actually continuous from $U^3 \cup \widehat{\mathbb{C}}$ to $\widehat{\mathbb{H}}$ since it is constructed from continuous functions. Since $C_{f,g,h}$ is real for all $q \in U^3$, by continuity we see that $C_{f,g,h}(z)$ is real for all $z \in \widehat{\mathbb{C}}$ except at the finitely many points where $C_{f,g,h}(z)$ is infinite. On $\widehat{\mathbb{C}}$, $C_{f,g,h}$ is given by

$$C_{f,g,h}(z) = \frac{(z - gz)(fz - hz)}{(z - hz)(fz - gz)},$$

which is a rational function of z since fz, gz , and hz are.

Since $C_{f,g,h}(z)$ is a rational function on $\widehat{\mathbb{C}}$ which is real on a set of positive measure, it follows from Lemma 3.4 that $C_{f,g,h}(z)$ is constant, say $C_{f,g,h}(z) \equiv \lambda$ for some $\lambda \in \mathbb{R}$. By evaluating $C_{f,g,h}(z)$ at a point z different from any of the fixed points of $g, h, h^{-1}f$, or $g^{-1}f$, we find that $\lambda \neq 0, \infty$. Now let z approach one of the fixed points v of g . Then $z - gz$ approaches 0 and so we see that either $z - hz$ or $fz - gz$ must approach 0 as well. Hence either $hv = v$ or $fv = v$, and we can assume without loss of generality that $fv = v$. It follows from Lemma 3.6 that f and g have the same fixed point sets. Finally, we note that the permutation $C_{f,h,g}(z)$ must also be constant, and a similar argument (with the roles of g and h interchanged) shows that h has the same fixed point as f and g . \square

Theorem 3.8. *For all $w \in H^3$ outside of a locally closed set of measure zero, each edge e of $D(w)$ is either exclusively loxodromically paired with cycle length three or exclusively parabolically paired, and in either case all pairings of e have the same fixed point sets.*

Proof. We begin by defining a set

$$E := \bigcup \{q \in U^3 : C_{f,g,h}(q) \in \mathbb{R}\}, \quad (3.2)$$

where the union is over all triples of distinct $f, g, h \in G$ such that the $C_{f,g,h}$ is nonconstant. We proceed to show that the result of the theorem fails only when $w \in E$.

To begin, we make the simple but important observation that by (1.11), since there are no elliptic elements in G , we must have

$$\sum_{i=1}^k \alpha(e_i) = 2\pi$$

for each edge cycle e_1, \dots, e_k of $D(w)$. By convexity of the Dirichlet domain, each edge subtends an angle strictly less than π , and therefore there must be at least 3 edges in each cycle.

Now suppose that an edge e of $D(w)$ belongs to a cycle of length at least 4. In particular, suppose that $f^{-1}(e), g^{-1}(e), h^{-1}(e)$ are edges of $D(w)$, with $f^{-1}, g^{-1}, h^{-1} \in G$. Then by the construction of the Dirichlet domain, w, fw, gw , and hw are equidistant to any point p lying on $L(e)$ (the geodesic containing e), and this is clearly only possible if w, fw, gw , and hw lie on an H -circle (hence an E -circle by Lemma 3.3). It follows from Lemma 3.2 that $C_{f,g,h}(w)$ is real. Then $w \in E$ unless $C_{f,g,h}$ is constant, in which case we know from Lemma 3.7 that f, g, h all have the same fixed point set. For any other edge of $D(w)$ of the form r^{-1} ($r^{-1} \in G$), an examination of $C_{r,f,g}$ as in the proof of Lemma 3.7 reveals that r also has the same fixed point set as f, g, h (assuming $w \notin E$). It follows that every pairing of e has the same fixed point set and therefore they are either all loxodromic or all parabolic. Observe that f, g, h cannot all be hyperbolic, since then w, fw, gw , and hw never lie on an H -circle (consider the case when f, g , and h fix the j -axis). Since w, fw, gw, hw lie identically on an E -circle, it follows from our assumption in Section 2 that f, g, h cannot all be loxodromic.

Finally, we show that the exceptional set E' of points in U^3 for which the conditions of the theorem fail is locally closed with measure zero. Then E has measure zero, since by Lemma 3.7 E is a union of at most countably many sets of measure zero. As we have shown, $E' \subset E$ and therefore E' of course has measure zero as well. Now for a given $w \in U^3$, observe that, by discreteness of G , there is some small neighborhood V of w such that only finitely many elements of G are involved in the construction of $D(w)$ for $w \in V$. That is, if we denote by H the set of all $g \in G$ such that $P(w, gw)$ contains a side of $D(w)$ for some $w \in V$, then $|H| < \infty$. Then we see that the set $E' \cap V$ of points in V for which the conditions of the theorem fail is actually the intersection with V of a *finite* union of closed sets of measure zero, and hence forms a closed set. That is, locally the union in (3.2) only needs to be over triples of distinct f, g, h in a finite subset H . Then we have shown that every point in E' has a neighborhood V such that $E' \cap V$ is closed, and it follows that E' is locally closed. \square

Remark 3.9. The advantage in the above theorem of showing that the exceptional set is closed is that we can rule out pathological examples like the exceptional set being the rational points in U^3 , in which case a computer will always randomly pick an exceptional point even though they have measure zero.

3.3. Generic Infinite Vertices. We have thus given a very useful partial answer to the question posed above. Namely, in the generic case, the only edge cycles which are *not* of length three are those consisting of exclusively parabolically paired edges, and for which all pairings have the same fixed point. We now proceed to consider such edge cycles.

Suppose that e is an exclusively parabolically paired edge of $D(w)$ and all pairings of e fix some point v . After applying an isometry taking v to ∞ , the pairings of e become simple translations $q \mapsto q + c$, $c \in \mathbb{C}$. Hence we are essentially in the case of Euclidean Dirichlet domains in E^2 . It is clear that the subgroup F

of G of parabolic isometries fixing v must itself be discrete, and can be identified with a discrete group of translations in E^2 . Thus we begin by looking at Dirichlet domains of discrete groups of translations in E^2 . We shall find that they always form either a hexagon or a rectangle. We can define edge cycles on such a Dirichlet domain just as before (we can think of our work in E^2 as applying to a horizontal cross section of U^3), and we shall find that in the case of a hexagon the edge cycles have length three as desired.

Moreover, the hexagonal case is the “typical” one (not the same as our earlier concept of generic $w \in U^3$), in the sense that the rectangular case occurs only when there are two perpendicular translations in the group (by perpendicular translations we mean that they are perpendicular after applying an isometry sending their fixed point to ∞). The usefulness of this result is that, in terms of our three dimensional hyperbolic Dirichlet domain, the typical case corresponds to e belonging to an edge cycle of length three. In the “atypical” case of a rectangular domain, we will need to check separately (or simply be given the fact) that the two generating parabolics are perpendicular in order to prove that the corresponding sides are exactly mapped.

Lemma 3.10. *A group Γ of translations in n -dimensional Euclidean space E^n is discrete if and only if Γ is generated by a set of linearly independent translations.*

Proof. See Theorem 5.3.2 of [4]. □

Theorem 3.11. *Let f, g be two linearly independent translations in E^2 . The Dirichlet domain $D(w)$ of $\langle f, g \rangle$ constructed at any $w \in E^2$ is either a rectangle or a hexagon with paired sides parallel to each other.*

Proof. As the Dirichlet domain is translationally invariant, choose w to be the origin of E^2 . Then f and g generate a lattice consisting of all points of the type $f^k g^l(0)$. Pick a to be a closest such point to the origin. By applying a rotation about the origin we may assume that a lies along the positive x-axis. Now let b be a closest point to the origin which is along a direction independent from a . We may pick b such that the angle $a - 0 - b$ is non-obtuse. Observe that the lattice $\langle a, b \rangle$ can be viewed as a tiling of the plane by triangles congruent to triangle $a - 0 - b$ (some of them flipped).

Moreover, triangle $a - 0 - b$ is non-obtuse. To see this, first recall that angle $a - 0 - b$ was chosen to be non-obtuse. Angle $0 - b - a$ is acute since b is not shorter than a . Angle $0 - a - b$ is acute since otherwise $b - a$ is shorter than b and linearly independent from a , contradicting our choice of b .

Observe that there are six triangles sharing the origin as a vertex. We claim that the Dirichlet domain is then the polygon P whose vertices are the circumcenters of these six triangles. If a and b are perpendicular, then triangle $a - 0 - b$ is a right triangle. Consequently, the circumcenters occur along the hypotenuses and therefore P is a rectangle. Otherwise, P is a hexagon with opposite sides paired by translations and therefore parallel.

To check that P is truly the Dirichlet domain for $\langle f, g \rangle$, we check that the hypotheses of Poincaré’s theorem are satisfied (see section 1.5):

- Each side of P is paired with the opposite parallel side by a translation, and these clearly satisfy conditions (1)-(3).
- As there are no elliptic elements, one easily checks that each edge cycle is length 3 or 4, and that the corresponding cycle transformations are all the identity, so that condition (4) is satisfied.
- In the case of the rectangle, the four angles add up to 2π , so condition (5) holds. For the hexagon, the edges come in cycles of length 3. By condition (4), the angles add up to a multiple of 2π . But each of the three angles is less than π , and hence they add up to exactly 2π .
- Condition (6) does not apply to Euclidean space.

Finally, to see that $\langle a, b \rangle = \langle f, g \rangle$, suppose by contradiction that there is some point p such that $p \in \langle f, g \rangle$ but $p \notin \langle a, b \rangle$. Then since any a and b translates of p also lie in $p \in \langle f, g \rangle$, we can assume that p lies in either triangle $0, a, b$ or triangle $0, -a, -b$. If p lies in triangle $0, a, b$ then p must lie on either line segment $0, a$ or line segment $0, b$, since otherwise p is closer to the origin than b and is linearly independent from a , contradicting minimality of b . But then p must be either a or b by minimality of a and b . If p lies in triangle $0, -a, -b$, then similarly p must lie on either line segment $0, -a$ or line segment $0, -b$, since otherwise the pair $-a$ and p contradicts our choice of a and b . Then we conclude that p is either $-a$ or $-b$.

□

From the above it follows that every infinite vertex of a generic $D(w)$ has six sides in the typical case or four sides in the atypical case, and we can expect the cycle length three method to apply in most applications for an arbitrary choice of w . In the case of a domain containing an infinite vertex with four sides (an issue that is not resolved by varying w), it is necessary to know that the translations in question are perpendicular, and a little thought shows that this is enough to verify the presence of an exact pairing.

We are also now in a position to give a partial resolution to the issue raised at the end of the second paragraph of section (3.1). The question is, if we have an ideal vertex (or at least what appears to be one to high precision), how do we verify that the constituent sides all intersect at exactly the same point, and that this point lies on the sphere at infinity? Suppose the sides in question lie on planes $P(w, g_1w), \dots, P(w, g_kw)$, for $g_1, \dots, g_k \in G$ (of course we expect that $k = 6$ or $k = 4$). We have seen that (in the generic case) g_1, \dots, g_k are all parabolic with the same fixed point v , and as usual we apply an isometry sending v to ∞ . In this view $P(w, g_1w), \dots, P(w, g_kw)$ are just vertical E-planes, and it is clear that they all intersect the sphere at infinity at their fixed point (in this case ∞). We see that in general it suffices to know that g_1, \dots, g_n are parabolic with the same fixed point. It is for this reason that we included in section (2) the ability to check whether elements of G are parabolic. As to checking that g_1, \dots, g_n have the same fixed point, we can already do this with the solution to the word problem in G , thanks to the following lemma.

Lemma 3.12. *Let $f, g \in PSL(2, \mathbb{C})$ be two parabolic transformations. Then f and g have the same fixed point in $\widehat{\mathbb{C}}$ if and only if $fgf^{-1}g^{-1} = 1$.*

Proof. Suppose that $fg = gf$ and that $gv = v$ for some $v \in \widehat{\mathbb{C}}$. Then we have

$$g(fv) = fgv = fv.$$

But since g has a unique fixed point, and g fixes both fv and v , we must have $fv = v$.

Conversely, suppose that $fv = gv = v$ for some $v \in \widehat{\mathbb{C}}$. Let h be an isometry with $h(v) = \infty$. Then hfh^{-1} and hgh^{-1} both fix ∞ . But any parabolic element fixing ∞ is of the form $q \mapsto q + c$, $c \in \mathbb{C}$. Clearly any of two maps of this form commute, and consequently

$$(hfh^{-1})(hgh^{-1}) = (hgh^{-1})(hfh^{-1}) \implies fg = gf.$$

□

3.4. Generic Finite Vertices. We now wish to address the remainder of the issue raised at end of the second paragraph of section (3.1). That is, what happens if we have a finite vertex with more than three constituent sides. Recall that the issue here is that we have no method of ruling out the case in which the sides intersect at two or more extremely close points. Theorem 3.15 which follows shows that in the generic case we need not worry about this, since every vertex will be formed by exactly three sides.

Lemma 3.13. *Five points $P_i = (P_{ix}, P_{iy}, P_{iz}) \in \mathbb{R}^3$, $i = 1, 2, 3, 4, 5$ lie on a two-dimensional generalized E-sphere in \mathbb{R}^3 if and only if we have*

$$\begin{vmatrix} |P_1|^2 & P_{1x} & P_{1y} & P_{1z} & 1 \\ |P_2|^2 & P_{2x} & P_{2y} & P_{2z} & 1 \\ |P_3|^2 & P_{3x} & P_{3y} & P_{3z} & 1 \\ |P_4|^2 & P_{4x} & P_{4y} & P_{4z} & 1 \\ |P_5|^2 & P_{5x} & P_{5y} & P_{5z} & 1 \end{vmatrix} = 0.$$

Proof. Suppose that P_1, P_2, P_3, P_4, P_5 lie on a 2-E-sphere with center $C = (C_x, C_y, C_z) \in \mathbb{R}^3$ and radius $R > 0$. Then for $i = 1, 2, 3, 4, 5$ we have

$$|P_i - C|^2 - R^2 = 0 \implies |P_i|^2 - 2P_{ix}C_x - 2P_{iy}C_y - 2P_{iz}C_z + |C|^2 - R^2 = 0.$$

This can be written in a more linear fashion as

$$\begin{pmatrix} |P_1|^2 & -2P_{1x} & -2P_{1y} & -2P_{1z} & 1 \\ |P_2|^2 & -2P_{2x} & -2P_{2y} & -2P_{2z} & 1 \\ |P_3|^2 & -2P_{3x} & -2P_{3y} & -2P_{3z} & 1 \\ |P_4|^2 & -2P_{4x} & -2P_{4y} & -2P_{4z} & 1 \\ |P_5|^2 & -2P_{5x} & -2P_{5y} & -2P_{5z} & 1 \end{pmatrix} \begin{pmatrix} 1 \\ C_x \\ C_y \\ C_z \\ |C|^2 - R^2 \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \\ 0 \\ 0 \\ 0 \end{pmatrix}$$

This shows that the leftmost matrix above has non-trivial kernel and consequently has vanishing determinant:

$$\begin{aligned} & \begin{vmatrix} |P_1|^2 & -2P_{1x} & -2P_{1y} & -2P_{1z} & 1 \\ |P_2|^2 & -2P_{2x} & -2P_{2y} & -2P_{2z} & 1 \\ |P_3|^2 & -2P_{3x} & -2P_{3y} & -2P_{3z} & 1 \\ |P_4|^2 & -2P_{4x} & -2P_{4y} & -2P_{4z} & 1 \\ |P_5|^2 & -2P_{5x} & -2P_{5y} & -2P_{5z} & 1 \end{vmatrix} = 0 \\ \implies & \begin{vmatrix} |P_1|^2 & P_{1x} & P_{1y} & P_{1z} & 1 \\ |P_2|^2 & P_{2x} & P_{2y} & P_{2z} & 1 \\ |P_3|^2 & P_{3x} & P_{3y} & P_{3z} & 1 \\ |P_4|^2 & P_{4x} & P_{4y} & P_{4z} & 1 \\ |P_5|^2 & P_{5x} & P_{5y} & P_{5z} & 1 \end{vmatrix} = 0. \end{aligned}$$

Now suppose that P_1, P_2, P_3, P_4, P_5 lie on a 2-E-plane. Then for some $a, b, c, d \in \mathbb{R}$ we have $P_1 \cdot (a, b, c) = d \iff P_{ix}a + P_{iy}b + P_{iz}c - d = 0$ for $i = 1, 2, 3, 4, 5$. We rewrite this in matrix form as

$$\begin{pmatrix} |P_1|^2 & P_{1x} & P_{1y} & P_{1z} & 1 \\ |P_2|^2 & P_{2x} & P_{2y} & P_{2z} & 1 \\ |P_3|^2 & P_{3x} & P_{3y} & P_{3z} & 1 \\ |P_4|^2 & P_{4x} & P_{4y} & P_{4z} & 1 \\ |P_5|^2 & P_{5x} & P_{5y} & P_{5z} & 1 \end{pmatrix} \begin{pmatrix} 0 \\ a \\ b \\ c \\ -d \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \\ 0 \\ 0 \\ 0 \end{pmatrix}.$$

As before, this implies

$$\begin{vmatrix} |P_1|^2 & P_{1x} & P_{1y} & P_{1z} & 1 \\ |P_2|^2 & P_{2x} & P_{2y} & P_{2z} & 1 \\ |P_3|^2 & P_{3x} & P_{3y} & P_{3z} & 1 \\ |P_4|^2 & P_{4x} & P_{4y} & P_{4z} & 1 \\ |P_5|^2 & P_{5x} & P_{5y} & P_{5z} & 1 \end{vmatrix} = 0.$$

Conversely, suppose we have

$$\begin{vmatrix} |P_1|^2 & P_{1x} & P_{1y} & P_{1z} & 1 \\ |P_2|^2 & P_{2x} & P_{2y} & P_{2z} & 1 \\ |P_3|^2 & P_{3x} & P_{3y} & P_{3z} & 1 \\ |P_4|^2 & P_{4x} & P_{4y} & P_{4z} & 1 \\ |P_5|^2 & P_{5x} & P_{5y} & P_{5z} & 1 \end{vmatrix} = 0.$$

Then the matrix

$$\begin{pmatrix} |P_1|^2 & -2P_{1x} & -2P_{1y} & -2P_{1z} & 1 \\ |P_2|^2 & -2P_{2x} & -2P_{2y} & -2P_{2z} & 1 \\ |P_3|^2 & -2P_{3x} & -2P_{3y} & -2P_{3z} & 1 \\ |P_4|^2 & -2P_{4x} & -2P_{4y} & -2P_{4z} & 1 \\ |P_5|^2 & -2P_{5x} & -2P_{5y} & -2P_{5z} & 1 \end{pmatrix}$$

has nontrivial kernel, so we have

$$\begin{pmatrix} |P_1|^2 & -2P_{1x} & -2P_{1y} & -2P_{1z} & 1 \\ |P_2|^2 & -2P_{2x} & -2P_{2y} & -2P_{2z} & 1 \\ |P_3|^2 & -2P_{3x} & -2P_{3y} & -2P_{3z} & 1 \\ |P_4|^2 & -2P_{4x} & -2P_{4y} & -2P_{4z} & 1 \\ |P_5|^2 & -2P_{5x} & -2P_{5y} & -2P_{5z} & 1 \end{pmatrix} \begin{pmatrix} M_1 \\ M_2 \\ M_3 \\ M_4 \\ M_5 \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \\ 0 \\ 0 \\ 0 \end{pmatrix}$$

for some $M_1, M_2, M_3, M_4, M_5 \in \mathbb{R}$. If $M_1 = 0$, then for $i = 1, 2, 3, 4, 5$ we have

$$-2P_{ix}M_2 - 2P_{iy}M_3 - 2P_{iz}M_4 + M_5 = 0 \implies (P_{ix}, P_{iy}, P_{iz}) \cdot (M_2, M_3, M_4) = -M_5,$$

so we see that P_1, P_2, P_3, P_4, P_5 lie on a 2-E-plane.

If $M_1 \neq 0$, then we can assume without loss of generality that $M_1 = 1$, since

any multiple of $\begin{pmatrix} M_1 \\ M_2 \\ M_3 \\ M_4 \\ M_5 \end{pmatrix}$ also lies in the kernel. We can also rename M_2 as C_x , M_3

as C_y , M_4 as C_z , and M_5 as $C_x^2 + C_y^2 + C_z^2 - Q = |C|^2 - Q$. Then for $i = 1, 2, 3, 4, 5$ we have

$$\begin{aligned} |P_i|^2 - 2P_{ix}C_x - 2P_{iy}C_y - 2P_{iz}C_z + |C|^2 - Q &= 0 \\ \implies |P_i - C|^2 &= Q. \end{aligned}$$

Then $Q \geq 0$, and in fact $Q > 0$ since the P_i 's are distinct. It then follows that P_1, P_2, P_3, P_4, P_5 lie on a 2-E-sphere (centered at C and with radius \sqrt{Q}). \square

Lemma 3.14. *Let $f, g, h, i \in G$ be distinct isometries which are not all parabolic with the same fixed point. With $P(w, fw)$ as before denoting the set of points in U^3 which are H -equidistant to w and fw (and similarly for the other isometries), the set*

$$S := \{w \in U^3 \mid P(w, fw) \cap P(w, gw) \cap P(w, hw) \cap P(w, iw) \neq \emptyset\}$$

is locally closed with measure zero.

Proof. Note that $P(w, fw) \cap P(w, gw) \cap P(w, hw) \cap P(w, iw) \neq \emptyset$ is equivalent to w, fw, gw, hw, iw lying on an H -sphere. Observe that if f, g, h, i are all hyperbolic with the same fixed point set, then w, fw, gw, hw, iw never lie on a 2- H -sphere (as usual this is most easily seen by considering the case when f, g, h, i fix 0 and ∞), so in this case S is the empty set and we are done.

Otherwise, we proceed to construct a real analytic function of $w \in U^3$ which vanishes whenever w, fw, gw, hw, iw lie on a 2- H -sphere, and the desired result will follow as long as our function does not vanish identically on U^3 .

We define our function $F : U^3 \rightarrow \mathbb{R}$ à la the previous lemma to be

$$F(w) := \begin{vmatrix} |w|^2 & w_x & w_y & w_z & 1 \\ |fw|^2 & (fw)_x & (fw)_y & (fw)_z & 1 \\ |gw|^2 & (gw)_x & (gw)_y & (gw)_z & 1 \\ |hw|^2 & (hw)_x & (hw)_y & (hw)_z & 1 \\ |iw|^2 & (iw)_x & (iw)_y & (iw)_z & 1 \end{vmatrix},$$

which as we have seen vanishes precisely when w, fw, gw, hw , and iw lie on a generalized 2-E-sphere. Then a fortiori F vanishes whenever w, fw, gw, hw , and iw lie on a 2-H-sphere. Since the determinant of a matrix is just a polynomial in the matrix entries, this function F is easily seen to be real analytic. Moreover, by our assumption in Section 2 w, fw, gw, hw , and iw cannot lie on a generalized 2-E-sphere for every $w \in U^3$. \square

Theorem 3.15. *For all $w \in U^3$ outside a closed set of measure zero, every finite vertex of $D(w)$ is formed by exactly three sides.*

Proof. We begin by defining a set

$$E := \bigcup \{w \in U^3 \mid P(w, fw) \cap P(w, gw) \cap P(w, hw) \cap P(w, iw) \neq \emptyset\},$$

where the union is over all quadruples of distinct $f, g, h, i \in G$. We proceed to show that $D(w)$ has a finite vertex formed by more than three sides only when $w \in E$.

Suppose that $D(w)$ has a finite vertex formed by more than three sides. In particular, suppose that sides s_1, s_2, s_3, s_4 meet at v , with corresponding planes $s_1 \subset P(w, fw), s_2 \subset P(w, gw), s_3 \subset P(w, hw), s_4 \subset P(w, iw)$, with $f, g, h, i \in G$. It is clear that f, g, h, i are not all parabolic with the same fixed point, since then as we have seen v would be an infinite vertex. Then w, fw, gw, hw , and iw are all H-equidistant to v , so $w \in E$. It follows that E is a union of at most countably sets, each of which has measure zero by Lemma 3.14. Therefore E has measure zero.

Now let E' be the set of all $w \in U^3$ such that $D(w)$ has a finite vertex formed by more than three sides. Clearly $E' \subset E$, so E' has measure zero. An argument very similar to the one given at the end of the proof of Theorem 3.8 shows that E' is also locally closed. \square

4. CHECKING THE CONDITIONS OF POINCARÉ'S THEOREM

Using our computer program we construct an approximate Dirichlet domain. We expect all vertices to be unambiguously distinct from each other, and in general there may be some (approximate) infinite vertices. Our program checks that each finite vertex has exactly three sides and that each infinite vertex has six sides (or four, in which case we must verify that the generating isometries are perpendicular). We check that each side has a partner side corresponding to the inverse isometry, and that each edge cycle has length three. We also verify that the sides corresponding to each infinite vertex are parabolic with the same fixed point. It then follows that all the edge geodesics are exactly mapped. Since the vertices are by definition just the intersections of these edge geodesics, it follows immediately that the vertices and hence the sides are exactly mapped as well.

In fact, this establishes conditions (1) - (5), of Poincaré's Theorem. Condition (6) is trivial, since each infinite cycle transformation is a composition of parabolics with the same fixed point, hence parabolic.

It remains to show that the side pairing elements generate the original group G . To show that they generate one of the original generators g_i of G , we draw a line segment from w to fw and observe which sides it intersects in the tiling of H^3 by $D(w)$ -translates. This gives us a sequence of side pairing elements whose composition should be g_i , and we verify this using the solution to the word.

5. THE DIRICHLET DOMAIN IN THE PROJECTIVE DISK MODEL

In this section we consider the field extension of \mathbb{Q} necessary to construct $D(w)$ centered at 0 (the origin) in the projective disk model. Let K denote the field generated by the entries of G and let K' be a field containing the coordinates of $g(j)$ in the upper half space model for each $g \in G$.

In D^3 the bisecting plane defined by g is given by

$$P(0, g(0)) = \{r \in D^3 \mid d_{D^3}(r, 0) = d_{D^3}(r, g(0))\},$$

and from (1.5) we also have

$$P(0, g(0)) = \{r \in D^3 \mid \sqrt{1 - |g(0)|^2} = 1 - r \cdot g(0)\}$$

A glance at (1.4) shows that the isometry $\phi^{-1} : U^3 \rightarrow D^3$ is such that $\phi^{-1}(K') \subset K'$. Now suppose that $g(0) = (r_x, r_y, r_z)$ and $g(j) = x + yi + tj$. Then from (1.3) we have

$$t = \frac{\sqrt{1 - |g(0)|^2}}{1 - r_z}.$$

This shows that $\sqrt{1 - |g(0)|^2} \in K'$, which means that $P(0, g(0))$ is of the form $\{r \in D^3 \mid r \cdot n = t\}$ for $t \in K', n \in K' \times K' \times K'$. Now suppose that for g_1, g_2, g_3 , the vertex defined by the intersection of their bisecting planes is given in the projective disk by $n_1 \cdot r = t_1, n_2 \cdot r = t_2, n_3 \cdot r = t_3$. This can be written in matrix form as

$$\begin{pmatrix} n_{11} & n_{12} & n_{13} \\ n_{21} & n_{22} & n_{23} \\ n_{31} & n_{32} & n_{33} \end{pmatrix} \begin{pmatrix} r_x \\ r_y \\ r_z \end{pmatrix} = \begin{pmatrix} t_1 \\ t_2 \\ t_3 \end{pmatrix}$$

or $NR = T \implies R = N^{-1}T$. It follows that the coordinates of r_x, r_y, r_z of the vertex are contained in K' . Then we have established the following theorem:

Theorem 5.1. *Let G be a discrete group of isometries of the projective disk model D^3 such that the origin is not fixed by any element of G . Let K' be the field containing the coordinates of the images of the origin under the group action. Then the coordinates of all the vertices of the Dirichlet domain based at the origin are contained in K' .*

REFERENCES

- [1] Alan F. Bearden. *The Geometry of Discrete Groups*, Springer, 1983.
- [2] Cinzia Bisi and Graziano Gentili. *Möbius Transformations and the Poincaré Distance in the Quaternionic Setting*, 2008.
- [3] B. Maskit, *Kleinian Groups*, Spring-Verlag, Berlin, c1988
- [4] John G. Ratcliffe. *Foundations of Hyperbolic Manifolds*, volume 149 of *Graduate Texts in Mathematics*. Springer, New York, second edition, 2006.
- [5] Robert Riley. *Applications of a Computer Implementation of Poincaré's Theorem on Fundamental Polyhedra*, Mathematics of Computation, Vol 40, No. 162 (Apr. 1983), pp. 607 - 632.

DEPARTMENT OF MATHEMATICS, COLUMBIA UNIVERSITY, NEW YORK, NY 10027