BILIPSCHITZ HOMOGENEOUS HYPERBOLIC NETS

CHRISTOPHER J. BISHOP

Abstract. We answer a question of Itai Benjamini by showing there is a $K > 1$ so that for any $\epsilon > 0$, there exist ϵ-dense discrete sets in the hyperbolic disk that are homogeneous with respect to K-biLipschitz maps of the disk to itself. However, this is not true for K close to 1; in that case, every K-biLipschitz homogeneous discrete set must omit a disk of hyperbolic radius $\epsilon(K) > 0$, depending only on K. For $K = 1$, this is a consequence of the Margulis lemma for discrete groups of hyperbolic isometries.

1. Introduction

Let $\mathbb{D} = \{z : |z| < 1\}$ denote the unit disk in complex plane $\mathbb{C}$. Let ρ denote the hyperbolic metric on $\mathbb{D}$. A set $X \subset \mathbb{D}$ is called discrete if it has no accumulation points in $\mathbb{D}$, and for $\epsilon > 0$ it is called ϵ-dense if every $x \in \mathbb{D}$ is within hyperbolic distance ϵ of some point of X. A set X is called homogeneous with respect to a set $\mathcal{F}$ of homeomorphisms if for any $x, y \in X$ there is a $f \in \mathcal{F}$ so that $f(X) = X$ and $f(x) = y$. In other words, $\mathcal{F}$ acts transitively on X (we emphatically do not assume $\mathcal{F}$ is a group; see below). We say that $X \subset \mathbb{D}$ is a (K, ϵ)-net if it is a discrete ϵ-net that is homogeneous with respect to the set of hyperbolic K-biLipschitz maps from $\mathbb{D}$ onto itself. Define

$$\epsilon(K) = \inf\{\epsilon : (K, \epsilon)\text{-nets exist}\}.$$

This is finite for all K since it is clearly a decreasing function of K and the orbit of any co-compact Fuchsian group is a $(1, \epsilon)$-net for some $\epsilon < \infty$. Thus

$$K_c = \inf\{K : \epsilon(K) = 0\} = \sup\{K : \epsilon(K) > 0\}$$

is well defined and $1 \leq K_c \leq \infty$. We shall prove:

Date: May 27, 2020.

1991 *Mathematics Subject Classification.* Primary: 30F45, Secondary: 30C620, 51M09 .

Key words and phrases. biLipschitz maps, hyperbolic geometry, Margulis constant, homogeneous set, quasiconformal group.

The author is partially supported by NSF Grant DMS 1906259.

Theorem 1.1. *$1 < K_c < \infty$.*

It is well known that $\epsilon(1) > 0$. For $K = 1$, the maps f are hyperbolic isometries and generate a subgroup H of the group G of all hyperbolic isometries mapping X to itself. Since X is a discrete set, G is a discrete group, i.e., a Fuchsian group, and $R = \mathbb{D}/G$ is a (possibly branched) Riemann surface and the set X projects to a single point $x \in R$. A famous result of Každan and Margulis [5], says that there is a positive constant $\epsilon_1 > 0$ (the Margulis constant) so that the injectivity radius is at least ϵ_1 at some point of R and hence R contains disk of radius at least $\epsilon_1/2$ that does not contain x. Thus $\epsilon(1) \geq \epsilon_1/2$. Alternate proofs of the Margulis lemma for Fuchsian groups are given in [6], [8], [10]; the latter gives the sharp value.

The question of whether $K_c > 1$ was raised by Itai Benjamini as a result of considering whether the Margulis lemma really requires the machinery of hyperbolic isometries, group actions and fundamental domains, or might have an analog for sets of biLipschitz mappings.

Next we observe that $K_c = \infty$ if we required that X were homogeneous with respect to some group H of K-biLipschitz maps on $\mathbb{D}$. Such a group would consist of K^2-quasiconformal maps, and a result of Tukia [9] says that $H = hGh^{-1}$ for some quasiconformal map $h : \mathbb{D} \to \mathbb{D}$ and some Möbius group G acting on $\mathbb{D}$. By Mori's theorem [7] (see also Chapter 3 of [1]) the image of a hyperbolic ϵ-disk under h or h^{-1} contains a hyperbolic disk of radius $\geq \frac{1}{16}\epsilon^{K^2}$. Since $h(X)$ is invariant under G, the previous paragraph shows it omits some disk of hyperbolic radius ϵ_1, and hence X omits some disk of radius ϵ_K depending only on K. Thus Theorem 1.1 says there is a phase transition the occurs when considering sets homogeneous with respect to all K-biLipschitz self-maps of $(\mathbb{D}, \rho)$ that does not occur when we consider set invariant under groups of such maps.

2. The construction

In this section we prove $K_c < \infty$ by building explicit (K, ϵ)-nets with K fixed and ϵ tending to zero. All our examples are infinite quadrilateral meshes that refine the same pentagonal tesselation. These meshes were constructed for different purposes in [3] (they are part of the proof that any simple planar n-gon can be quad-meshed in time $O(n)$ using elements with all new angles between 60° and 120°). We start with a standard tesselation of $\mathbb{D}$ by hyperbolic right pentagons. See Figure 1. Connect the center of each pentagon to the midpoint of each edge; this divides the pentagon into five quadrilaterals. Choose a positive integer N and divide each quadrilateral into a $N \times N$ quadrilateral mesh using geodesic arcs as shown in Figure 1. Each boundary arc of the larger quadrilateral is divided into N sub-arcs of equal length. This implies the mesh in each quadrilateral matches the mesh in its neighbors and defines a quadrilateral mesh of the whole disk. Every element of the mesh has hyperbolic diameter $\simeq 1/N$ and belongs to a compact family of quadrilaterals. The vertices of this mesh will be our (K, ϵ)-net X_N with $\epsilon \simeq 1/N$. Below we refer to X_N simply as X.

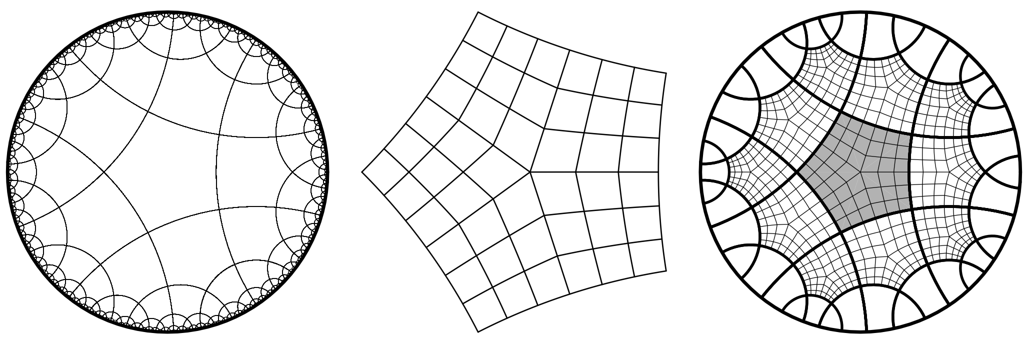

FIGURE 1. Hyperbolic right pentagons tessellate the disk. Each pentagon is divided into five quadrilaterals which are then each divided into a $N \times N$ quadrilateral mesh (here $N = 3$). The elements all have diameter $\simeq 1/N$.

We will show that there is a fixed $1 < K < \infty$, independent of N so that any element $x \in X$ can be mapped to any other $y \in X$ by a hyperbolic K-biLipschitz map f that also maps X to itself. This will prove Theorem 1.1

Because the pentagonal tesselation is invariant under a group of hyperbolic isometries, it is enough to prove this when x and y belong to the same pentagon. Indeed, by composing two maps, it suffices to assume y is the center of the pentagon. We will show that any vertex in a pentagon can be mapped to the center vertex by composing two K-biLipschitz maps called "discrete rotations" that we now describe.

Fix an element of A_0 of our quadrilateral mesh. Let A_1 be the collection of mesh elements (other than A_0) that hit A_0; here we consider all elements to be closed sets, so a mesh element Q is in A_1 if is shares an edge or vertex with A_0. Unless one of the vertices of A_0 was a center point of a pentagon, A_1 will be the union of eight quadrilaterals. Four of these are "side quadrilaterals" meaning they share one side with A_0 and four are "corner quadrilaterals", meaning they share a vertex, but no side, with A_1. Let $B_1 = A_0 \cup A_1$.

In general, we let A_n be the union of mesh elements that are not in $B_{n-1} = \cup_{k<n} A_k$, but that share an edge or vertex with a quadrilateral in A_{n-1}. As above, side elements share a side with some element of A_{n-1} and corner elements do not. The number of corner elements will remain four until the first time that A_n contains a vertex that is the center of a pentagon in our tesselation. Such points are vertices of five quadrilaterals, and if such a point x occurs on the boundary of B_n, then A_{n+1} will contain either three or four quadrilaterals containing this point: three if x is a side vertex of A_n and four if x is a corner vertex.

The number of pentagon centers included in B_n is bounded by the number of distinct pentagons of the tesselation hit by B_n, which is bounded in terms of its hyperbolic diameter, which is $\simeq n/N$. In particular, for $n = 2N$ (the only case we will consider) it is uniformly bounded.

Note that vertices on ∂B_k form a cycle; a shift by j on ∂B_k is a map from this cycle to itself that moves every element by j positions clockwise (and thus preserves the adjacencies among vertices and preserves the orientation of the cycle). A "discrete rotation" will be a biLipschitz map from B_n that is the identity on ∂B_n (i.e., a shift by zero) and for $0 \leq k < n$ is a shift on ∂B_k by $j(k)$ where $|j(k) - j(k+1)| \leq C$. It is clear this map on vertices can be extended to a map in B_n that maps each A_k to itself and is biLipschitz with constant depending only on C. The biLipschitz constant is determined by mapping each quadrilateral in A_k (which are all drawn from some compact family of quadrilaterals) to a quadrilateral inside A_k whose outer edge has been shifted by $j(k)$ and the inner edge has been shifted by $j(k-1)$; this is for a side quadrilateral that has one edge on each of ∂B_{k-1} and ∂B_k (we call these the inner and outer edges, respectively), but the idea is the same for corner quadrilaterals that have two edges on ∂B_k and a single vertex on ∂B_{k-1}. See Figure 2 for an example on a square grid; combinatorially this is the same as a B_7 on our grid which does not contain any pentagon centers.

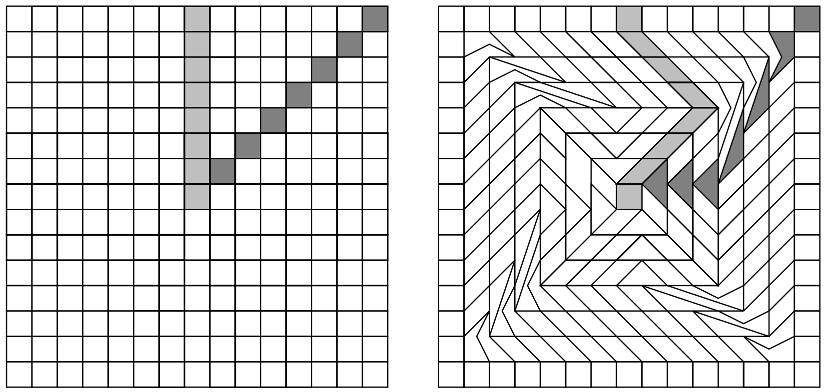

FIGURE 2. A discrete rotation on a square grid. This map is the identity at the center and on the outer boundary. in the center levels is acts a clockwise shift by varying amounts. We can map any vertex on $\partial B_{n/2}$ to any other vertex on the same curve by such a discrete rotation that is the identity on ∂B_n and K-biLipschitz on the interior, with K independent of n.

If there are pentagon centers in B_n then instead of ∂B_{k+1} having four more vertices than ∂B_k it might have as many $4 + 4c$ vertices, where c is the number pentagon centers occurring on ∂B_k. However, this remains bounded as long as $n = O(N)$ so the biLipschitz constant of interpolating between a shift by $j(k)$ on ∂B_k and a shift by $j(k+1)$ on ∂B_{k-1} is bounded by some constant depending on C and the hyperbolic diameter of B_k (these determine how far the outer edge of a quadrilateral in A_n is shifted compared to its inner edge (or two outer edges and inner vertex in the case of a corner quadrilateral).

An alternate approach to constructing the discrete rotations might be to use a biLipschitz map to map A_k to a round Euclidean annulus $\{k \leq |z| \leq k+1\}$, do the interpolation there, then map back to A_k.

Discrete rotations can similarly be constructed in discrete balls that are centered, not on a quadrilateral in the mesh, but on a mesh vertex x. Then A_1 contains four quadrilaterals if x is not the center of a pentagon, or five quadrilaterals if it is. In both cases, all elements of A_1 are corners. Examples of discrete balls around centers of pentagons are illustrated in the center (radius 3) and right (radius 9) of Figure 1. All previous estimates hold for these balls and discrete rotations as well.

The final step is to observe that two such discrete rotations can map any vertex x inside a pentagon to the center of that pentagon. First rotate around the center of the pentagon in a discrete ball of radius $2N$ to move the point to the segment connecting the center to a midpoint of a boundary segment of the pentagon. This point is then on the boundary of one of our $N \times N$ grids and we can discrete rotate around the center of this grid (in a discrete ball of radius $2N$) to bring the point to the center. This completes the proof of Theorem 1.1.

Although we have not attempted this, it should be possible to give a bound for K by constructing the discrete rotation maps explicitly.

3. The phase transition

Next we prove $K_c > 1$, by assuming $K_c = 1$ and constructing a quasiconformal map between $\mathbb{C}$ and $\mathbb{D}$, contradicting Liouville's theorem. Thus there is a type of phase transition below which Margulis type separation holds for biLipschitz homogeneous discrete nets and above which it does not.

Proof. Suppose that there were a sequence $\{X_n\}$ of (K_n, ϵ_n) nets with $K_n \to 1$ and $\epsilon_n \to 0$. By making ϵ_n smaller, if necessary, we may assume that X_n is a (K_n, ϵ_n)-net and also omits some disk of radius $\epsilon_n/2$ around a point w. Since w is within distance ϵ_n of some point of X, and X is biLipschitz homogeneous with constant K_n close to 1, every point of X is within $2\epsilon_n$ of the center of an omitted disk of radius $\epsilon_n/4$. By applying an isometry we may also assume $0 \in X_n$.

Restrict each X_n to $D(0, \sqrt{\epsilon_n})$ and expand it by a factor of $1/\epsilon_n$. This gives a sequence of sets that are Euclidean 1-nets in $D(0, 1/\sqrt{\epsilon_n})$, and by passing to a subsequence we may assume that $\{X_n\}$ converges locally in the Hausdorff metric to a closed set $X \subset \mathbb{C}$ so that (1) X is a 1-net, (2) X is homogeneous with respect to Euclidean isometries, (3) every point of X is within distance 2 of a point that is distance $1/4$ from X.

The set of isometries that map X into itself is a closed subgroup G of the Euclidean isometry group. Since G acts transitively on the 1-net X, it must be infinite. Thus G is a closed, infinite Lie subgroup of the isometry group of the plane and must be either a discrete group (and X is a Euclidean lattice) or it is $\mathbb{R} \times \mathbb{Z}$ (and X is a union of evenly spaced parallel lines).

In the first case, assume the minimum separation between elements of X is δ. We can choose n large enough so that elements of $X_n/\epsilon_n \cap D(0, 100)$ each lie within a $\delta/100$ neighborhood of X, and every such neighborhood of $X \cap D(0, 100)$ contains a point of X_n/ϵ_n. The chosen points define quadrilaterals that correspond to the fundamental parallelograms of X and we can quasiconformally map each parallelogram to the corresponding quadrilateral quasiconformally (with uniformly bounded dilatation). By starting at one point of X_n and working our way outwards, this implies $\mathbb{D}$ can be written as a quadrilateral mesh isomorphic to $\mathbb{Z} \times \mathbb{Z}$ where each quadrilateral comes from a compact family. This implies there is a quasiconformal map from $\mathbb{C}$ into $\mathbb{D}$, which is impossible (together with the measurable Riemann mapping theorem, it violates Liouville's theorem that $\mathbb{D}$ and $\mathbb{C}$ are conformally distinct).

In the other case, X consists of parallel lines that are between distance 1/2 and 2 apart (since X is within 1 of every point and its complement contains disks of radius 1/4.) Thus for n large enough, X_n/ϵ_n locally approximates parallel lines separated by $\simeq \epsilon_n$ and for any $x \in X$ we can choose a subset of X_n that approximates a square 10×10 grid with x as the center point. By using the homogeneous property of X_n we can map the center of the grid to one of its boundary points and extend the approximation. Continuing in this way we can find an infinite subset of X_n that approximates a square grid of fixed size in a neighborhood of any of its points. We can use this subset as the vertices of a quadrilateral mesh whose elements are approximate squares, and use this to construct a quasiconformal map between the disk and the plane. As above, this is impossible. Since the existence of the sequence $\{X_n\}$ leads to a contradiction, no such sequence exists and we must have $K_c > 1$. □

4. Questions and remarks

Is the function $\epsilon(K)$ strictly deceasing on $[1, K_c]$? Is $\epsilon(K_c) = 0$? Does $\epsilon(K)$ tend to the Margulis constant as $K \searrow 1$? Is $\epsilon(K)$ continuous? It seems possible that the nets that minimize ϵ for a given K could have some special combinatorial structure, and that when this is changed, the optimal ϵ is different. This it seems possible that jumps in $\epsilon(K)$ could occur.

For the mesh constructed in Section 2 what is the optimal K? Is there a different mesh that allows for all $\epsilon > 0$ and smaller value of K than give by this mesh?

Is the set of (K, ϵ) in $Q = [1, \infty) \times (0, \infty)$ for which a (K, ϵ)-net exists a closed set in Q; if we take a sequence of such nets, we can extract sequence converging to a K-biLipschitz homogeneous ϵ-dense set X, but is X discrete? What if in the definition of homogeneous we only require $f(X) \subset X$ instead of $f(X) = X$? What if we consider K-biLipschitz maps of X instead of $\mathbb{D}$?

What can happen if X is a K-biLipschitz ϵ-net, but we don't require X be discrete? Then we could have $X = \mathbb{D}$; what else is possible? In general, a K-biLipschitz homogeneous compact set in $\mathbb{R}^2$ can be a Cantor set, even with K close to 1 (think of a thin Cantor set constructed using very thick annuli; the outer and inner boundary boundaries can be rotated all the way around by a biLipschitz map with small constant). What if X has non-trivial connected components? Hoehn and Oversteegen [4] proved any compact planar set that is homogeneous under self-homeomorphisms is necessarily either a finite set, a Cantor set, a Jordan curve, a pseudo-arc, a circle of pseudo-arcs or the product of one of the first two with one of the latter three. It is still unknown whether a biLipschitz homogeneous continuum must be a Jordan curve. However, it is known that a biLipschitz homogeneous Jordan curve must be a quasicircle [2].

There is nothing special about negative curvature in Theorem 1.1. Analogous arguments hold for the sphere. The dodecahedron divides the sphere into 12 congruent spherical pentagons and the same division of each pentagon into $5N^2$ quadrilaterals gives an ϵ-net with $\epsilon \simeq 1/N$ that is K-biLipschitz homogeneous for a fixed $K < \infty$, independent of N. Conversely, the blowing argument shows that if the sphere had (K, ϵ)-nets with K arbitrarily close to 1 and ϵ arbitrarily close to 0, then we could construct a covering map from the plane to the sphere, which is impossible since they are both simply connected but not homeomorphic. Thus the sphere also has a critical exponent strictly between 1 and ∞ for the existence of arbitrarily fine, discrete K-biLipschitz nets. In fact, it seems reasonable that under some type of smoothness assumption the Euclidean plane is the only 2-manifold that has (K, ϵ)-nets with $(K - 1) + \epsilon$ arbitrarily close to zero. What are appropriate assumptions? Note that "snowflaking the plane" (i.e., replacing the metric $|x - y|$ by $|x - y|^\alpha$ for some $0 < \alpha < 1$) gives a metric space that is distinct from the plane and that does have arbitrarily fine bilipschitz homogeneous nets (even with $K = 1$).

There should be an alternate approach to proving $K_c > 1$ using discrete curvature of meshes. The mesh we constructed in Section 2 looks very Euclidean at most points, indeed, it is a union of $N \times N$ grids. But there are occasional vertices of degree 5. These special vertices introduce the negative curvature into the mesh structure. We can make this precise by defining the curvature at a vertex x as

$$C(x) = 1 - \frac{1}{2}\deg(x) + \sum_{f:x\in f} \frac{1}{\#(f)},$$

where $\deg(x)$ is the degree of the vertex x, the sum is over all faces with x as a vertex, and $\#(f)$ is the number of sides of the face f. If we sum this over all the vertices of a finite planar mesh we get $\sum_x C(x) = V - E + F$ where V is the number of vertices, E the number of edges and F the number of faces. By Euler's formula this is 2.

For a quadrilatal mesh, $C(x) = 1 - \deg(x)/4$. Hence for the mesh constructed in Section 2, we have $C(x) = 0$ everywhere except at the pentagon centers, where $C(x) = -1/4$. If we took a large but finite piece of the mesh, e.g., a discrete ball B_n for $n \gg N$, then the sum of $C(X)$ over the interior vertices is a large negative value, so the sum over the boundary vertices must be a large positive number (recall the total sum is 2). Suppose ∂B_n has m edges; this is $\#(f)$ for the unbounded face of the finite mesh. Then at boundary vertices of degree 3, $C(X) = 1 - \frac{3}{2} + (\frac{1}{4} + \frac{1}{4} + \frac{1}{m}) = \frac{1}{m}$ and at corner vertices $C(x) = 1 - \frac{2}{2} + (\frac{1}{4} + \frac{1}{m}) = \frac{1}{4} + \frac{1}{m}$. The $1/m$ terms sum to 1 over the boundary, hence there are about as many corners on the boundary as there are degree five vertices in the interior. Each corner gives rise to an "extra" vertex in ∂B_{n+1}, so the number of boundary elements of B_n grows exponentially with its hyperbolic diameter. This is expected for a reasonable mesh in hyperbolic space.

However, if the mesh was biLipschitz homogeneous with constant close to 1, we might hope that $C(x) = 0$ at "most vertices" might imply $C(x) = 0$ at all vertices, which would imply the number of corners of ∂B_n is small compared to the number of interior vertices (perhaps even bounded). This should mean that the mesh is "not reasonable", e.g., a sub-exponential number of mesh elements covers the exponentially long boundary of B_n. Thus there should be mesh elements of arbitrarily large size, not what we would expect from a biLipschitz homogeneity.

Since we are given a set of points X, not a mesh, so we would have to introduce our own mesh associated to X, e.g., the Voronoi diagram or the dual mesh. It is easy to see the Voronoi cells of an ϵ-net have diameters bounded by $O(\epsilon)$, so perhaps the idea in the previous paragraph could be applied.

One way to explain why the Margulis lemma holds is that hyperbolic space is not scale invariant, i.e., on small scales it looks Euclidean (i.e., tiny triangles have angle sum close to π) but it looks negatively curved at larger scales. So if a mesh in the hyperbolic disk is very fine, it looks like a Euclidean mesh on small scales, and if it homogeneous enough, this means it will look Euclidean everywhere, and thus have the wrong properties (area growth, discrete curvature,..) to be a mesh in a negative curved space. Either a mesh in hyperbolic space can be very fine but not very homogeneous, or homogeneous and coarse enough that the pieces exhibit "enough" negative curvature. It would be interesting to try to make this vague idea into an actual proof.

References

[1] Lars V. Ahlfors. *Lectures on quasiconformal mappings*, volume 38 of *University Lecture Series*. American Mathematical Society, Providence, RI, second edition, 2006. With supplemental chapters by C. J. Earle, I. Kra, M. Shishikura and J. H. Hubbard.

[2] Christopher J. Bishop. Bi-Lipschitz homogeneous curves in $\mathbb{R}^2$ are quasicircles. *Trans. Amer. Math. Soc.*, 353(7):2655–2663, 2001.

[3] Christopher J. Bishop. Optimal angle bounds for quadrilateral meshes. *Discrete Comput. Geom.*, 44(2):308–329, 2010.

[4] Logan C. Hoehn and Lex G. Oversteegen. A complete classification of homogeneous plane continua. *Acta Math.*, 216(2):177–216, 2016.

[5] D. A. Každan and G. A. Margulis. A proof of Selberg's hypothesis. *Mat. Sb. (N.S.)*, 75 (117):163–168, 1968.

[6] Albert Marden. Universal properties of Fuchsian groups in the Poincaré metric. In *Discontinuous groups and Riemann surfaces (Proc. Conf., Univ. Maryland, College Park, Md., 1973)*, pages 315–339. Ann. of Math. Studies, No. 79, 1974.

[7] Akira Mori. On an absolute constant in the theory of quasi-conformal mappings. *J. Math. Soc. Japan*, 8:156–166, 1956.

[8] Jacob Sturm and Meir Shinnar. The maximal inscribed ball of a Fuchsian group. In *Discontinuous groups and Riemann surfaces (Proc. Conf., Univ. Maryland, College Park, Md., 1973)*, pages 439–443. Ann. of Math. Studies, No. 79, 1974.

[9] Pekka Tukia. On two-dimensional quasiconformal groups. *Ann. Acad. Sci. Fenn. Ser. A I Math.*, 5(1):73–78, 1980.

[10] Akira Yamada. On Marden's universal constant of Fuchsian groups. *Kodai Math. J.*, 4(2):266–277, 1981.

C.J. Bishop, Mathematics Department, Stony Brook University, Stony Brook, NY 11794-3651

Email address: `bishop@math.stonybrook.edu`