BILIPSCHITZ NETS IN THE HYPERBOLIC DISK

CHRISTOPHER J. BISHOP

Abstract. We answer a question of Itai Benjamini by showing that there can be ϵ-dense discrete sets in the hyperbolic disk that are homogeneous with respect to a set of K-biLipschitz maps for the hyperbolic metric, where $K > 1$ is fixed, but $\epsilon > 0$ may be as small as we wish.

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 $\mathbb{D}$ to itself. More concisely, X is a K-biLipschitz homogeneous discrete ϵ-net. (Note that our maps are biLipschitz on $\mathbb{D}$, not just X; one could consider just biLipschitz maps $f : X \to X$, but this is a less restrictive concept since not every biLipschitz map on X need extend to a biLipschitz map on $\mathbb{D}$). Define

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

This is clearly a decreasing function of K and we shall prove it is eventually zero:

Theorem 1.1. $K_c = \sup\{K : \epsilon(K) > 0\} < \infty$.

Date: May 27, 2020.

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

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

The author is partially supported by NSF Grant DMS 1906259.

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 [3], 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 Fuschsian groups are given in [4], [6], [8]; the latter gives the sharp value.

The question of whether Theorem 1.1 holds 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 Theorem 1.1 cannot hold if X is 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 [7] 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 [5] (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 significant difference between a discrete ϵ-net being homogeneous with respect to all K-biLipschitz self-maps of $(\mathbb{D}, \rho)$ and being homogeneous with respect to a group of such maps.

2. The construction

Our (K, ϵ)-net X will consist of the vertices of a infinite quadrilateral mesh of $\mathbb{D}$ that was constructed for different purposes in [2] (it is 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. Connect the center of each polygon 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 as shown in Figure 1 so that 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 and every element of the mesh has hyperbolic diameter $\simeq 1/N$. The vertices of this mesh will be our (K, ϵ)-net X_N with $\epsilon \simeq 1/N$. Below we will drop the N and refer to X_N simply as X.

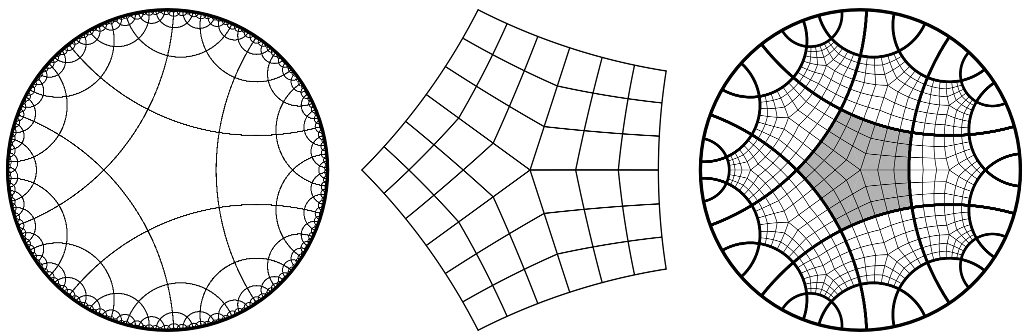

FIGURE 1. Hyperbolic right pentagons tessellate the disk. Each pentagon is divided into five quadrilaterals which are then each divided into a $N \times N$ quadrilateral mesh (here $N = 3$). The elements all have diameter $\simeq 1/N$.

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 polygon. 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 polygon 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 $n = O(N)$ (which is 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 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.

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

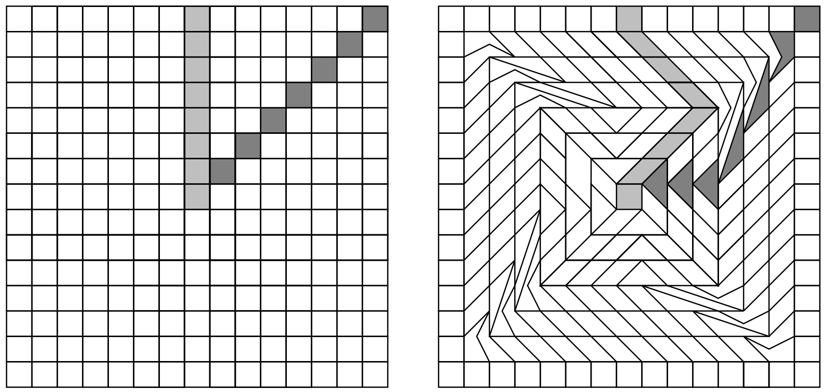

FIGURE 2. A discrete rotation on a square grid. This map is the identity at the center and on the outer boundary. in the center levels is acts a clockwise shift by varying amounts. We can map any vertex on $\partial B_{n/2}$ to any other vertex on the same curve by such a discrete rotation that is the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

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 would to biLipschitz 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 polygon, 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 the estimates described above hold for these balls and the corresponding 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. A phase transition

Can K be taken close to 1 in Theorem 1.1? We can't do this in the above construction because the discrete rotations sometimes must map of the inner and outer edges of a side quadrilateral into a quadrilateral whose inner and outer sides are separated by a corner vertex and hence are almost perpendicular to each other; this requires a fixed amount of distortion to accomplish and keeps K bounded away from 1.

I currently expect the following holds:

Theorem 3.1. $K_c = \inf\{K : \epsilon(K) = 0\} > 1.$

Here is a sketch of how we might prove this, but I have not checked it carefully and some details are missing, so it may be incorrect.

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 translating we may also assume $0 \in X_n$.

Restrict X_n to $D(0, \sqrt{\epsilon_n})$ neighborhood of and and expand it by a factor of $1/\epsilon_n$. We get a sequence of sets that are Euclidean 1-nets in $D(0, 1/\sqrt{\epsilon})$, 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 Euclidean isometries that map X into itself is a closed subgroup of the Euclidean isometry group that acts transitively on the 1-net X. Thus it is a closed, infinite Lie subgroup of the isometry group of the plane and hence is 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 $X \cap D(0,1000)$ contains a point of X_n/ϵ_n. The chosen points define quadrilaterals that correspond to the fundamental parallelograms of X and we can 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 we can choose a subset of X_n that approximates a square 10×10 grid on a neighborhood of some 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 a 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 to construct a quasiconformal map between the disk and the plane. As above, this is impossibility. Thus the existence of the sequence $\{X_n\}$ leads to a contradiction, so no such sequence exists. $\square$

4. Questions and remarks

Is the function $\epsilon(K)$ strictly deceasing on $[1, K_c]$? Is $\epsilon(K_c) = 0$? Is $\epsilon(K)$ continuous? Smooth? Does $\epsilon(K)$ tend to the Margulis constant as $K \searrow 1$? 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-net X, but is X discrete? 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? What can happen 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 $\mathbb{D}$?

It seems like there should be an alternate approach to proving Theorem 3.1 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 where 4 of the Euclidean grids meet. These special vertices introduce the negative curvature into the mesh structure. We can make this more precise by defining the curvature at a vertex x of a mesh 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, so we deduce 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 this means the number of boundary elements grows exponentially with n. This is what we expect 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 cover up the exponentially long boundary of B_n. Thus it seems that there are mesh elements of arbitrarily large size, not what we would expect from a biLipschitz homogeneity.

Since 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. The proof we gave of Theorem 3.1 is another variation on this; we showed that a (K, ϵ) net with K near 1 and ϵ near zero would have to look like the vertices of a Euclidean mesh that has curvature zero at every vertex.

Thus 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. Optimal angle bounds for quadrilateral meshes. *Discrete Comput. Geom.*, 44(2):308–329, 2010.

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

[4] 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.

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

[6] 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.

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

[8] 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`