BILIPSCHITZ HOMOGENEOUS HYPERBOLIC NETS

CHRISTOPHER J. BISHOP

Abstract. We answer a question of Itai Benjamini by showing there is a $K < \infty$ 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 $\epsilon(K) > 0$. 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 the complex plane $\mathbb{C}$ and let ρ denote the hyperbolic metric on $\mathbb{D}$ (defined in Section 2). 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 $z \in \mathbb{D}$ is within hyperbolic distance ϵ of some point $x \in X$. 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 do not assume $\mathcal{F}$ is a group; see Remark 1 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 (we could also consider biLipschitz self-maps of X; see Remark 2 below). 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 (as K increases, the infimum is over larger sets), and the orbit of any co-compact Fuchsian group G is a $(1, \epsilon)$-net for some $\epsilon < \infty$; we can take ϵ to be the diameter of the compact quotient surface $R = \mathbb{D}/G$. An explicit bound is given by the genus two Bolza surface, whose

Date: May 27, 2020; revised March 22, 2024.

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

Key words and phrases. biLipschitz maps, hyperbolic geometry, Margulis constant, homogeneous set, quadrilateral mesh.

The author is partially supported by NSF Grant DMS 2303987.

hyperbolic diameter is $\arctan(3+2\sqrt{2}) \approx 2.45$; see [10]. 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 both inequalities are strict.

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

The upper bound follows from an explicit construction: for any $\epsilon > 0$ we build an ϵ-net that is homogeneous for K-biLipschitz maps with $K < \infty$ independent of ϵ. The lower bound is given by an indirect argument. Assuming $K_c = 1$ we construct an ϵ-dense set in $\mathbb{D}$ that "looks like" a copy of $\mathbb{Z} \times \mathbb{Z}$, and we will derive a contradiction with the exponential growth of the hyperbolic area.

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 [7], 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 intersect X. Thus $\epsilon(1) \geq \epsilon_1/2$. Alternate proofs of the Margulis lemma for Fuchsian groups are given in [8], [11], [13]; 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 it have an analog for sets of biLipschitz mappings.

Remark 1: We claim that $K_c = \infty$, if we require X to be 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 [12] says that such a group is of the form $H = hGh^{-1}$ for some quasiconformal map $h : \mathbb{D} \to \mathbb{D}$ and some Möbius group G acting on $\mathbb{D}$. By Mori's theorem [9] (or Chapter 3 of [1]), the image of a hyperbolic ϵ-disk under h or h^{-1} contains a hyperbolic disk of radius $\geq \epsilon^{K^2}/16$. 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.

Remark 2: We have assumed that X is homogeneous under biLipschitz self-maps of the disk, but we could replace this by self-maps of X. Our proof of the upper bound produces biLipschitz maps of the whole disk, and the proof of the lower bound only uses that we have self-maps of X. Thus the inequality $1 < K_c < \infty$ holds in either case, although it is not clear whether the exact value of K_c is the same in both situations (this depends partly on whether a K-biLipschitz self-map of an ϵ-net can be extended to a K-biLipschitz self-map of the disk, or whether a larger constant is needed).

2. The upper bound: $K_c < \infty$

The pseudo-hyperbolic metric on $\mathbb{D}$ is given by

$$\widetilde{\rho}(z,w) = \left| \frac{z-w}{1-\overline{w}z} \right|$$

and the hyperbolic metric by

$$\rho(z,w) = \frac{1}{2} \log \frac{1+\widetilde{\rho}(z,w)}{1-\widetilde{\rho}(z,w)}.$$

The hyperbolic metric can also be defined as $\rho(z,w) = \inf \int_\gamma ds/(1-|x|^2)$, where the infimum is over all rectifiable paths in $\mathbb{D}$ connecting z and w. This implies $\rho(z,w) > |z-w|$ whenever $z \neq w$. The (orientation preserving) isometries of the hyperbolic metric are the linear fractional transformations of the disk to itself. The geodesics for the hyperbolic metric are diameters of the circle and their images under isometries, i.e., circular arcs perpendicular to the boundary. A ball of hyperbolic radius r has hyperbolic area that grows exponentially in r. See [2] or [5] for these basic facts about the hyperbolic metric. A hyperbolic K-biLipschitz map $f : X \to Y$ between subsets of $\mathbb{D}$ is one that satisfies

$$1/K \leq \frac{\rho(f(z),f(w))}{\rho(z,w)} \leq K \text{ for all } z,w \in X.$$

In this section, we prove the upper bound $K_c < \infty$ in Theorem 1.1 by building explicit (K,ϵ)-nets with K fixed and ϵ tending to zero. All our examples correspond to infinite quadrilateral meshes that refine a fixed tesselation of $\mathbb{D}$ by right pentagons. These meshes were constructed for different purposes in [4] (in that paper, 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 the standard tesselation of $\mathbb{D}$ by hyperbolic right pentagons. See the left side of Figure 1. Connect the center of each pentagon to the midpoint of each of its five boundary arcs. This divides the pentagon into five fundamental quadrilaterals. Each quadrilateral has three right angles and one angle of $2\pi/5$, the latter at the center of the pentagon.

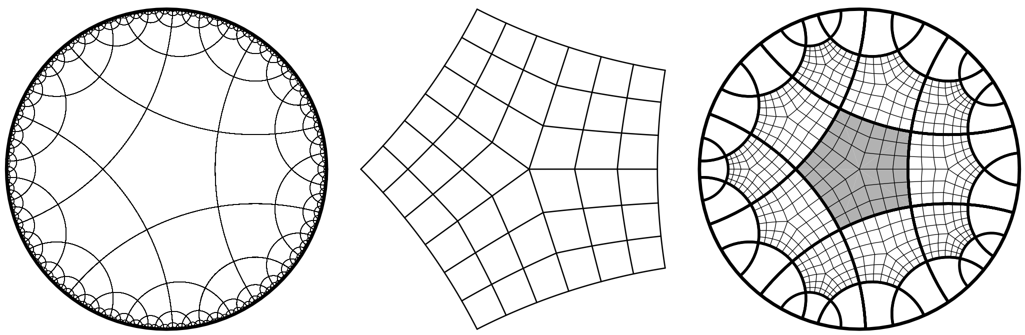

FIGURE 1. Hyperbolic right pentagons tessellate the disk. Each pentagon is divided into five quadrilaterals which are then each divided into a $N \times N$ quadrilateral mesh (here $N = 3$). The elements all have diameter and side lengths $\simeq 1/N$.

The two edges adjacent to the center of the pentagon have the same length as each other, as do the two sides opposite these. Choose a positive integer N and divide each quadrilateral into a $N \times N$ quadrilateral mesh using geodesic arcs as shown in the center of Figure 1. Each boundary arc of the fundamental quadrilateral is divided into N sub-arcs of equal length. This implies the mesh in each fundamental quadrilateral matches the mesh in all its neighbors and defines a quadrilateral mesh of the whole disk. We will call this mesh M; it is an infinite graph embedded in $\mathbb{D}$ in which every vertex has degree four or five (the latter occurs only at the centers of the hyperbolic pentagons). The set X of vertices of this mesh is our ϵ-net with $\epsilon \simeq 1/N$.

We observe for later use that the hyperbolic distance between two vertices z, w of M is comparable to their graph distance, $d_M(w, z)$, divided by N. This holds with a constant that is independent of N. To see this, note that each edge of the mesh has hyperbolic length $O(1/N)$, so $\rho(z, w) = O(d_M(z, w)/N)$. On the other hand, a geodesic segment $\gamma \subset \mathbb{D}$ connecting distinct points $z, w \in X$ can can hit at most $O(N\rho(z, w))$ faces of the mesh: each such face has hyperbolic area $\simeq 1/N^2$, and is contained in a $O(1/N)$ neighborhood of γ, so the union of faces hitting γ has area

$O(\rho(z,w)/N)$. The edges of these faces contain a path of mesh edges connecting z and w, so $d_M(z,w) = O(N\rho(z,w))$.

We let $Y \subset X$ denote the vertices of the pentagonal tesselation in the left side of Figure 1; these are the points where the geodesics defining the edges of the tesselation cross each other. We call these geodesics the "bounding geodesics" and call their crossing points the "corner points". The set Y of corner points is clearly homogeneous under isometries of the hyperbolic disk. Thus to map a point $x_1 \in X$ to another point $x_2 \in X$, it suffices to map x_1 to some $y_1 \in Y$ and map some $y_2 \in Y$ to x_2, and then isometrically map y_1 to y_2. Thus it suffices to show that each x inside a fundamental quadrilateral Q can be mapped to a corner point $y \in Y$ by a K-biLipschitz map of X to itself, with K independent of N.

We do this in two steps. Given $x \in X$, let Q be the fundamental quadrilateral containing x and let $y = Q \cap Y$ be the corresponding corner point. First we will define a "discrete rotation" of X around y that maps x to a point $z \in \partial Q \cap X$ that lies on bounding geodesic γ passing though y. The second step is to define a "discrete translation" of X along γ that maps z to y. We will show both steps can be accomplished by K-biLipschitz maps, with K independent of N.

If $x \in Y$, there is nothing to do, so we assume $x \notin Y$ and choose $y = Q \cap Y$ where Q is a fundamental quadrilateral containing x. If x in on a bounding geodesic passing through y, we can continue to the second step of the construction, so for the moment, we assume this is not the case.

The corner point y is on the boundary of four hyperbolic pentagons. Let P be the union of these four pentagons. We define a series of closed cycles $\{\Gamma_k\}_1^{2N}$ in the mesh $M \cap P$. See Figure 2 for an example where $N = 5$. The first curve, Γ_1, consists of the eight points of X that are adjacent to y in the mesh M. In general, if we have already defined $\Gamma_1, \ldots \Gamma_k$, then Γ_{k+1} is the cycle consisting of points of $X \cap P$ that are adjacent to Γ_k but not in Γ_{k-1}. Note that Γ_{2N} lies on the boundary of P. Also, for $k = 1, \ldots N$, observe that Γ_{k+1} has eight more points than Γ_k. For $k \geq N$ it has sixteen more points than Γ_k (note that Γ_N is the cycle passing through the centers of the pentagons).

By assumption, $x \neq y$, but x and y are in the same fundamental quadrilateral Q, so x lies on some Γ_k with $1 \leq k \leq N$. Moreover, there is a point $w \in X \cap \Gamma_k$ that

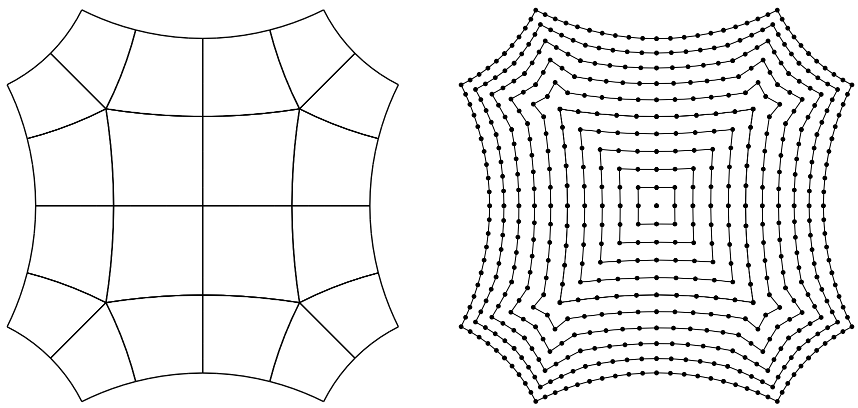

FIGURE 2. The left shows the union of four pentagons (= twelve fundamental quadrilaterals) that touch the corner point y at the center. The right picture shows the concentric cycles Γ_k surrounding y. If x is on the kth ring with $1 \leq k \leq N$, then it is at most k points away from a bounding geodesic passing through y (the vertical and horizontal lines through y). Here $N = 5$.

is on a bounding geodesic and is at most $j \leq k$ steps away from x on the cycle Γ_k. Thus we can map x to w by "rotating" Γ_k by j steps (every point of Γ_k is moved j positions in the same direction).

We extend this rotation to the rest of X as follows. For $1 \leq m < j$ we rotate Γ_{k+m} by $j - |m|$ positions. Similarly for Γ_{k-m}. On the rest of X we take the identity. Recall that the hyperbolic distance between points of X is comparable (with absolute constants) to the mesh distance in M divided by N. The map above clearly only multiplies mesh distances by at most a bounded factor, independent of N. To see this, note that if $m \neq 0$ and two points are on Γ_k and Γ_{k+m} respectively, then the mesh distance between them is at least m and it can increase by at most $O(|m|)$ (partly due to the size of the shifts varying by at most m, and partly due to the lengths of the two cycles differing by at most $16m$. If two points are on the same cycle Γ_k, then the shifts at worst multiply the mesh distance by two. Since $d_M(z, w) \simeq N\rho(z, w)$, our maps also multiply hyperbolic distances by a bounded factor, i.e. they are hyperbolically Lipschitz with a uniform constant. Moreover, the inverse map has the same form, so the inverse is also Lipschitz with a uniform bound. Thus our discrete rotation map is uniformly biLipschitz as a map $X \to X$.

We can extend the map $X \to X$ defined above to be a biLipschitz self-map of the whole disk. We define the extension as the identity outside P (the union of pentagons touching y), and within P we define it as follows. For each annular region A_k between the cycles Γ_{k-1} and Γ_k we take a biLipschitz map of A_k to a round annulus with points of X mapping to evenly spaced points on each boundary circle (this is easy). The discrete rotation maps on the cycles become Euclidean rotations on the boundary circles, and the angle of rotations differ by at most a bounded multiple of the width of the annulus. These boundary rotations can be interpolated by a biLipschitz map that just rotates each concentric circle between the boundaries, and this map is then transported back to A_k.

This completes the first step of our construction: every $x \in X$ can be mapped to a point of X lying on a bounding geodesic, by a uniformly biLipschitz map of $\mathbb{D}$ to itself. Next we need to show any mesh point on a bounding geodesic can be mapped to a corner point $y \in Y$. This is even easier that the previous step.

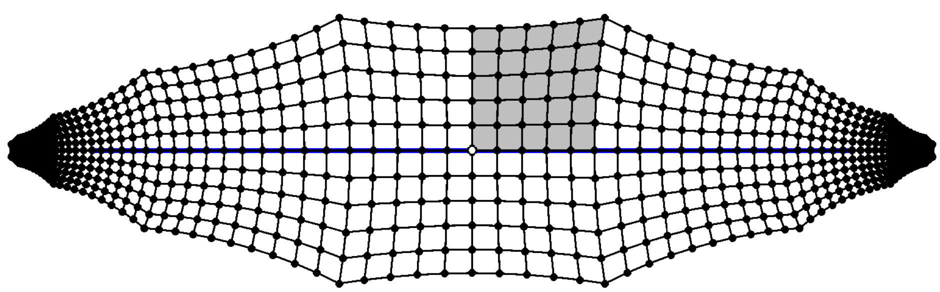

FIGURE 3. S (for strip) is the union of fundamental quadrilaterals touching a single bounding geodesic (the thickened central horizontal line passing through y, the white dot). The mesh in this region is isomorphic to the square mesh on $\mathbb{Z} \times [-N, N]$ and it is easy to define a uniformly biLipschitz map that translates the central geodesic by j mesh elements (with $|j| \le N$) and is the identity outside S. Here $N = 5$.

Suppose $y \in Y$ is on the bounding geodesic γ. Let S denote the union of all fundamental quadrilaterals Q that touch γ. In Figure 3 a single fundamental quadrilateral Q is shaded. This quadrilateral, and its three rotations around y by $\pi/2$, π and

$3\pi/2$, form a larger quadrilateral Q', and S is union of translates of Q' under powers (positive and negative) of a single hyperbolic translation along γ. See Figure 3.

The restriction of the mesh M to S is isomorphic to the graph $\mathbb{Z} \times [-N, N]$. Fixing j with $1 \leq j < N$, we can define a discrete translation by j by sending $(n, m) \to (n + j - |m|, m)$ for $|m| < j$ and taking the identity map on the rest of X. This is clearly uniformly biLipschitz in both the graph and hyperbolic distances on X and can move a point on a bounding geodesic up to N positions. This is enough to move any point onto a corner, as desired. The extension to a biLipschitz self-map of the disk is similar but even simpler than in the previous case. We take the identity map outside S. We map the strip S to the Euclidean strip $\mathbb{R} \times [-N, N]$, extend the translations on the top and bottom edges to the interior via a shear map preserving each horizontal line, and then map this back to S.

This completes the proof of upper bound in Theorem 1.1.

3. The lower bound: $K_c > 1$

Next we prove $K_c > 1$ by contradiction. We will assume that $K_c = 1$ and construct an ϵ-dense mesh that only has $O(r/\epsilon^2)$ points within hyperbolic distance r of the origin. However, this contradicts the well known fact that a hyperbolic ball of radius r has area that grows exponentially with r. In this section, we need only assume that X is homogeneous with respect to K-biLipschitz homeomorphisms of X to itself.

If $K_c = 1$, then for any $K > 1$ and any $\epsilon > 0$ we can find a (K, ϵ) net. Fix one such net X. Let δ be the supremum of numbers $r > 0$ so that $\mathbb{D} \setminus X$ contains a disk of hyperbolic radius r. Then $\delta > 0$ since X discrete and hence its complement is open and non-empty. Also $\delta \leq \epsilon$, since every point is within ϵ of some point of X. So if we make ϵ smaller by setting $\epsilon = \delta$, the set X is still a (K, ϵ)-net, since every point of $\mathbb{D}$ is within hyperbolic distance ϵ of X, but there exist disks of any radius strictly less than ϵ that are not intersected by X. In particular, for any $\lambda \in (0, 1)$, there is a point $x_\lambda \in X$ and a hyperbolic disk D_λ of radius $\lambda\epsilon$ so that $D_\lambda \cap X = \emptyset$, and whose center is within distance ϵ of x_λ.

Lemma 3.1. *If $K > 1$ is close enough to 1 and $\epsilon > 0$ is close enough to 0, then every disk of radius ϵ around any point $w \in X$ contains the center of a disk of radius $\epsilon/4$ that is not hit by X.*

Proof. Suppose not. Note that the regular 12-gon inscribed in a Euclidean disk of radius 1 has side lengths $2\sin(\pi/12) \approx .5176$, so that in hyperbolic space 12 equally spaced points on a circle of hyperbolic radius ϵ will more than $2\sin(\pi/12)\epsilon > \epsilon/2$ apart. Placing twelve disjoint disks of hyperbolic radius $\epsilon/4$ on the hyperbolic circle of radius ϵ around w, we see that, under our assumption, each such disk must contain a point of X. See Figure 4.

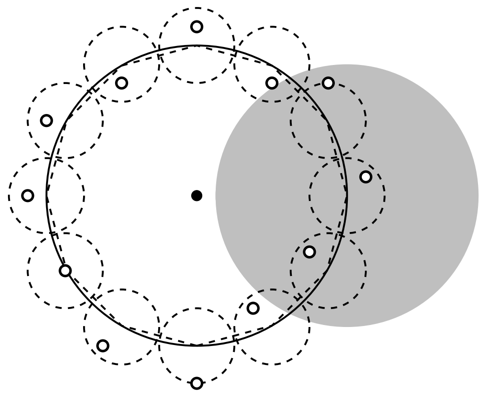

FIGURE 4. In hyperbolic space we can place twelve disks of radius $\epsilon/4$ on a circle of radius ϵ around $x \in X$ (black dot). Under our assumptions, each smaller disk contains a point of X (white dots) and, up to an isometry, each point's position is determined by its distances to x and the other chosen points. Hence approximately the same configuration occurs when we map w to x_λ by a K-biLipschitz map with $K \approx 1$. Thus one of the twelve image points must land inside D_λ (shaded) if λ is close to 1.

Map w to x_λ by a K-biLipschitz map of X and, recalling that we can take K as close to 1 as we wish. Thus the distances of the twelve image points from x_λ are almost the same as the original distances to w, and the distances between the image points is almost the same as between the original points. Thus there is a hyperbolic isometry of the disk taking w to x_λ so that the images of the twelve points under the isometry approximate the images under the K-biLipschitz map as closely as we wish.

Moreover, as ϵ tends to zero, the length of the gaps on circle $\{z : |z - w| = \epsilon\}$ between these evenly spaced disks tends to $\epsilon(2\pi - 24\sin(\pi/12))/12 \approx (.006)\epsilon$. Thus by choosing λ close enough to 1, we can ensure the center of D_λ is as close to $\{|z - x_\lambda| = \epsilon\}$ as we wish, and that the center of D_λ is within distance $3\epsilon/4$ of one of the image points. This means that D_λ contains a point of X. But D_λ was chosen to not intersect X, and this contradiction proves the lemma. □

Now restrict each X to the Euclidean disk $D(0, \sqrt{\epsilon}) \subset \mathbb{D}$ and expand it by the Euclidean dilation $z \to z/\epsilon$. Taking sequences $K_n \searrow 1$ and $\epsilon_n \searrow 0$ gives a sequence of sets $Z_n \subset \mathbb{C}$ that are Euclidean 1-nets in $D(0, 1/\sqrt{\epsilon_n})$, and by passing to a subsequence we may assume that $\{Z_n\}$ converges locally in the Hausdorff metric to a closed set $Z \subset \mathbb{C}$ so that

(1) Z is a (Euclidean) 1-net,
(2) Z is homogeneous with respect to Euclidean isometries,
(3) Any 2-ball centered in Z contains a $\frac{1}{4}$-ball disjoint from Z.

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

In either case, Z contains a lattice Z' whose fundamental parallelogram is close to a square. In the case $Z = \mathbb{R} \times \mathbb{Z}$, we can take an actual square sub-lattice, and otherwise we can choose elements of Z that are within distance 1 of the points 10 and $10i$; these give a fundamental parallelogram that is approximately a square.

This means if ϵ and K are close enough to 0 and 1 respectively, then taking $\delta = 10\epsilon$ and any point $x \in X$ we can find eight other points in X that approximate a Euclidean square 3×3 grid with side length δ centered at x. Applying the same argument to each of the eight boundary points of this grid, we can expand it to a 5×5 grid. Continuing, we can build a $(2n+1) \times (2n+1)$ grid centered at x that is a union of approximate δ sized quadrilaterals that approximate squares uniformly. More precisely, we obtain a subset of X that is a 2δ-net in $\mathbb{D}$ whose points are $\delta/2$ separated, and has the structure of a Euclidean square mesh. But then $O(n^2)$ disks of radius 2δ centered on this grid cover a ball of radius $\simeq \delta n$ around x. In other words, $O(n^2)$ disks, each of hyperbolic area $O(\delta^2)$, cover a ball of hyperbolic radius $\simeq n\delta$, and this ball must have hyperbolic area at least $\exp(c\delta n)$ for some $c > 0$. This is impossible for large n, and the contradiction implies $K_c > 1$.

4. Questions and remarks

The constructions in this paper are quite explicit. What are the explicit bounds for K_c given by this construction?

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. Thus it seems possible that jumps in $\epsilon(K)$ could occur.

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 [6] 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 (at least to the author) whether a biLipschitz homogeneous continuum in $\mathbb{R}^2$ must be a Jordan curve. However, it is known that a biLipschitz homogeneous Jordan curve in the plane must be a quasicircle [3].

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-up 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 $(K - 1)$ and ϵ both arbitrarily close to zero. If so, what are appropriate assumptions? Note that "snow-flaking 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 $K=1$).

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] Alan F. Beardon. *The geometry of discrete groups*, volume 91 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 1983.

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

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

[5] John B. Garnett and Donald E. Marshall. *Harmonic measure*, volume 2 of *New Mathematical Monographs*. Cambridge University Press, Cambridge, 2008. Reprint of the 2005 original.

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

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

[8] 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.

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

[10] H. Stepanyants, A.F. Beardon, J. Paton, and D. Krioukov. *Diameter of compact Riemann surfaces*. 2024. preprint, arXiv:2301.10844v3.

[11] 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.

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

[13] 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