Negatively Curved Graph and Planar Metrics with Applications to Type

PHILIP L. BOWERS

Introduction

A graph is of parabolic or hyperbolic type if the simple random walk on the vertices is, respectively, recurrent or transient. A plane triangulation graph is CP-parabolic or CP-hyperbolic if the maximal circle packing determined by the graph packs, respectively, the complex plane $\mathbb{C}$ or the Poincaré disk $\mathbb{D}$. We examine the implications that (Gromov) negative curvature carries for determining type, specifically in these settings. Our main result is encased in the following theorem.

THEOREM. *Every proper* (*Gromov*) *negatively curved metric space whose boundary contains a nontrivial continuum admits a* $(2, C)$*-quasi-isometric embedding of a uniform binary tree.*

Corollaries of this theorem include:

(1) the simple random walk on every locally finite, negatively curved graph whose boundary contains a nontrivial continuum is transient;
(2) the simple random walk on a locally finite, 1-ended negatively curved graph whose boundary contains more than one point is transient;
(3) a negatively curved plane triangulation graph is CP-hyperbolic if and only if it has a circle boundary (equivalently, CP-parabolic if and only if it has a point boundary).

The classical "type problem" is that of determining whether a given noncompact, simply connected Riemann surface is conformally equivalent to the plane $\mathbb{C}$ or the disk $\mathbb{D}$. The surface is said to be of *parabolic* type in the former case, and of *hyperbolic* type in the latter. Our concern is with two related discretizations of this classical problem, one via random walks on graphs, the other via planar circle packings. Connections between probabilistic characteristics and the type problem are deep and intimate, and have been known for a long time. For instance, a simply connected Riemann surface is hyperbolic if and only if a Brownian traveler starting at any point has a positive escape probability. This generalizes to

Received November 4, 1996. Revision received August 26, 1997.

The author thanks the Department of Pure Mathematics and Mathematical Statistics at Cambridge University for their kind hospitality during his sabbatical visit in the first half of 1996, when this work began. The author also thanks an anonymous referee, especially for simplifying the proofs of Section 3.

Michigan Math. J. 45 (1998).

higher-dimensional Riemannian manifolds. A complete Riemannian manifold is hyperbolic exactly when the Brownian motion generated by the Laplace–Beltrami operator is transient. Kanai discretized this in [26; 27] where he characterized hyperbolic Riemannian manifolds of bounded geometry as (roughly) those that are quasi-isometric to certain graphs of hyperbolic type. Another discretization of type has occurred more recently in the very geometric/combinatorial setting of circle packings, which is a purely 2-dimensional phenomenon. Here, a triangulation of the plane determines a maximal circle packing of exactly one of the plane $\mathbb{C}$ or the disk $\mathbb{D}$, and there has been interest in finding conditions on the combinatorics of the triangulation that determine type; see, for instance, [18; 23; 28].

It is of no surprise that negative curvature of a graph has implications for the type of the graph, but all work to date concerning these implications has assumed additional structure—either bounded degree or some sort of isoperimetric inequality. These, especially the latter, are very strong uniformity conditions that forbid the type of random asymptotic behavior that can easily occur in graphs and in circle packings, even in the context of negative curvature.

Our proof of the theorem occupies Section 1 and involves primarily the very fast divergence of geodesics in a negatively curved space, as well as the quasi-denseness of geodesic rays. The theorem itself in weaker forms is hinted at in Gromov's work [21] and Bowditch's exposition [8] of Gromov's work. Actually, we prove a result stronger than the theorem, namely, that every metric space satisfying the hypotheses of the theorem admits, for any positive ε, a $(1+\varepsilon, C(\varepsilon))$-quasi-isometric embedding of a uniform binary tree, where $C(\varepsilon)$ is small when measured in the length scale of the binary tree. The applications of the theorem to the determination of type are easy and appear in the subsequent sections. After recalling some standard facts about random walks in Section 2, we apply the theorem to confirm the first two corollaries listed above. These are compared with previous results of Ancona [4] and of Kaimanovich and Woess [25], both of which assume bounded degree and strong isoperimetric conditions in the context of simple random walks. Section 3 develops some of the basic topology of negatively curved planar metrics. These metrics arise naturally from negatively curved plane triangulation graphs as quasi-isometric images, and we prove that their boundaries are either singletons or topological circles. The beautiful results of He and Schramm [23] determining the CP-type of plane triangulation graphs using vertex extremal length are reviewed in Section 4, after which the theorem, as well as results from the section preceding, are applied to the determination of CP-type. In particular, the third corollary listed above is proved. A final section, Section 5, concludes with a short discussion of further applications of the theorem that are due to K. Stephenson and the author, proofs of which will appear in [9]. Since many readers whose background is in either random walks or circle packings may not be familiar with the recent developments on metric negative curvature, it seems fitting to include all the results that we will use from this theory in one section. Thus an appendix has been included that gives an overview of recent metric geometry with references (or proofs when we could find no references) for all the results stated. As there is no uniformity of terminology in this rather young field, we take this opportunity to fix the terminology and notation that we use in our proofs, and we present some results in a nonstandard way that is particularly suited for our purposes in the paper proper.

1. Proof of the Theorem

Throughout this section, ρ denotes a proper, negatively curved geodesic metric on the set X with a fixed basepoint x_0, and Z denotes a nontrivial continuum in the Gromov boundary ∂X (see the Appendix for definitions). Also, all geodesic rays are assumed to be arclength parameterized. Since ρ is negatively curved, there exists a positive constant δ such that (X, ρ) has δ-inscribed triangles.

LEMMA 1.1. *Let α be a geodesic ray based at x_0 that limits at a point of Z, and let s be any positive parameter value. If there exists a geodesic ray β based at x_0 that limits at a point of Z for which the ρ-distance between $\alpha(s)$ and $\beta(s)$ is at least 4δ, then there exists a geodesic ray σ based at x_0 that limits at a point of Z for which*

$$2\delta < \rho(\alpha(s), \sigma(s)) < 4\delta.$$

Proof. Assuming that the boundary ∂X is parameterized from the basepoint x_0 as described in the Appendix, let

$$U = \{\, \sigma(\infty) \in \partial X : \rho(\alpha(s), \sigma(s)) \leq 2\delta \,\}$$

and

$$V = \{\, \sigma(\infty) \in \partial X : \rho(\alpha(s), \sigma(s)) \geq 4\delta \,\}.$$

Observe that U contains the point $\alpha(\infty) \in Z$ and that V contains the point $\beta(\infty) \in Z$. From the definition of the topology on the boundary, as described in the Appendix, the sets U and V are closed in the boundary and, by Lemma A.4, they are disjoint. If there is no geodesic ray with the desired property, then the union of the two sets U and V covers Z, and it follows that U and V provide a separation of the connected set Z, a contradiction. □

We apply this lemma recursively to define a collection of arclength parameterized geodesic rays σ_b, indexed by binary sequences b of finite length, that are based at x_0 and limit at points of Z. First, choose such rays σ_0 and σ_1 based at x_0 that limit at respective distinct points of Z. Since $\sigma_0(\infty)$ and $\sigma_1(\infty)$ are distinct boundary points, the ρ-distance between $\sigma_0(t)$ and $\sigma_1(t)$ is unbounded as t increases, and we may choose a parameter value t_0 for which the ρ-distance between $\sigma_0(t_0)$ and $\sigma_1(t_0)$ is at least 2δ. Let $\Sigma_1 = \{\sigma_0, \sigma_1\}$.

Fix a constant $K \geq 5$ and, for the positive integer n, assume that we have constructed the set

$$\Sigma_n = \{\, \sigma_b : |b| = n \,\}$$

of geodesic rays based at x_0 that limit at points of Z such that the ρ-distance between any two distinct points from the set

$$\{\, \sigma_b(t_0 + K\delta(n-1)) : |b| = n \,\}$$

is at least 2δ. Here we are using $|b|$ to denote the length of the binary sequence b. Lemma A.4 implies that the ρ-distance between any two points from the set

$$\{\, \sigma_b(t_0 + K\delta n) : |b| = n \,\}$$

is at least $(2K+1)\delta$, which is greater than 4δ. For each σ_b in Σ_n, apply Lemma 1.1 with s of the lemma equal to $t_0 + K\delta n$ to find a geodesic ray σ_{b*} based at x_0 that limits at a point of Z and for which

$$2\delta < \rho(\sigma_b(t_0 + K\delta n), \sigma_{b*}(t_0 + K\delta n)) < 4\delta. \tag{1.1}$$

For $i = 0, 1$, define σ_{bi} as σ_b if the last digit of b equals i; otherwise, define σ_{bi} as σ_{b*}. This defines the set $\Sigma_{n+1} = \{\, \sigma_b : |b| = n+1 \}$, and our construction ensures, since $(2K+1)\delta$ is at least 11δ, that the ρ-distance between any two points from the set

$$\{\, \sigma_b(t_0 + K\delta n) : |b| = n+1 \}$$

is at least 2δ.

We now use the sets Σ_n to construct a mapping of a *uniform binary tree* $\mathcal{T}$ into X. The tree $\mathcal{T}$ is a nontrivial, rooted graph with intrinsic metric d (see Appendix) wherein each vertex other than the root vertex has degree three and where each edge is isometric to a Euclidean interval of some fixed length, which we take to be of length δ. The root itself, denoted as $v_\varnothing$ (where $\varnothing$ stands for the empty binary sequence), has degree 2. There are 2^n vertices in the sphere of radius $n\delta$ about the root vertex in $\mathcal{T}$, and we assume that these have been labeled as v_b, where b ranges over the binary sequences of length n; the *children* of vertex v_b are v_{b0} and v_{b1}. Define the mapping $\lambda \colon \mathcal{T} \to X$ on vertices by

$$\lambda(v_b) = \sigma_b(t_0 + K\delta(|b| - 1)) \tag{1.2}$$

if $b \neq \varnothing$, with $\lambda(v_\varnothing)$ equal to the midpoint of a geodesic segment from $\lambda(v_0) = \sigma_0(t_0)$ to $\lambda(v_1) = \sigma_1(t_0)$, and extend to edges by mapping the edge from v_b to its child v_{b*} convexly onto a geodesic segment from $\lambda(v_b)$ to $\lambda(v_{b*})$, using the segment that lies in σ_b whenever $\sigma_b = \sigma_{b*}$.

Our first task is to calculate upper and lower bounds for the ρ-distance between the λ-images of any two vertices in $\mathcal{T}$. Toward this end, first notice that the ρ-distance between $\lambda(v_b)$ and $\lambda(v_{b*})$, where v_{b*} is a child of v_b, is at least $K\delta$ if $b \neq \varnothing$; moreover, this distance is at most $(K+1)\delta$. This upper bound on the distance follows from the definition of λ on vertices, Equation (1.2), along with the fact that $\sigma_{b0}(t)$ and $\sigma_{b1}(t)$ are δ-close at parameter value $t = t_0 + K\delta(|b| - 1)$. This last fact follows from an examination of the internal points of a triangle with vertices x_0, $\lambda(v_{b0})$, and $\lambda(v_{b1})$, and the fact that $K \geq 5$ along with (1.1). Let $b0x$ and $b1y$ be binary sequences that agree in the first $|b|$ digits and observe that the upper bound of $(K+1)\delta$ for the ρ-distances between the λ-images of a vertex and its children gives the upper bound

$$\rho(\lambda(v_{b0x}), \lambda(v_{b1y})) \leq (K+1)\delta(|x| + |y| + 2) = (K+1)d(v_{b0x}, v_{b1y}). \tag{1.3}$$

A useful lower bound is a bit harder to come by. For this, let b and c be distinct binary sequences for which $|b| = |c|$. Use Lemma A.4 and the fact that the ρ-distance between $\lambda(v_b)$ and $\lambda(v_c)$ is at least 2δ, while that between the λ-images of two sibling vertices is at most 4δ, to obtain the inequality

$$2K\delta + \rho(\lambda(v_b), \lambda(v_c)) - 9\delta \leq \rho(\lambda(v_{b*}), \lambda(v_{c*})), \tag{1.4}$$

where v_{b*} and v_{c*} are respective children of v_b and v_c. If $|x| = |y|$ for the binary sequences $b0x$ and $b1y$, then starting with the fact that $\rho(\lambda(v_{b0}), \lambda(v_{b1})) \geq 2\delta$, successive applications of (1.4) imply that

$$((2K - 9)|x| + 2)\delta \leq \rho(\lambda(v_{b0x}), \lambda(v_{b1y})). \tag{1.5}$$

For the general lower bound, assume that $|x| \geq |y|$. Let z be the binary sequence of length $|x| - |y|$, all of whose digits are equal to the last digit of the binary sequence $1y$. Then $\sigma_{b1y} = \sigma_{b1yz}$ and (1.2) implies that

$$\rho(\lambda(v_{b1y}), \lambda(v_{b1yz})) = K\delta|z| = K(|x| - |y|)\delta. \tag{1.6}$$

The triangle inequality applied to the λ-images of the vertices $v_{b0x}, v_{b1y}, v_{b1yz}$, along with (1.5) and (1.6), gives the lower bound

$$((K - 9)|x| + K|y| + 2)\delta \leq \rho(\lambda(v_{b0x}), \lambda(v_{b1y})). \tag{1.7}$$

After subtracting $9\delta|y|$ from, and adding and subtracting $2(K - 9)\delta$ to, the left-hand side of (1.7), recalling also that the d-distance between the vertices v_{b0x} and v_{b1y} is $(|x| + |y| + 2)\delta$, we obtain the lower bound

$$(K - 9)d(v_{b0x}, v_{b1y}) - 2(K - 10)\delta \leq \rho(\lambda(v_{b0x}), \lambda(v_{b1y})). \tag{1.8}$$

THEOREM 1.2. *For each $\varepsilon > 0$, the space X admits a $(1 + \varepsilon, C(\varepsilon))$-quasi-isometric mapping of a uniform binary tree.*

Proof. Let $K \geq 10$ and scale the metric d by a factor of $(K - 9)$ to obtain a metric d_K on $\mathcal{T}$ in which each edge has length $(K - 9)\delta$. Then, for every pair of nonroot vertices u and v of $\mathcal{T}$, (1.3) and (1.8) imply that

$$d_K(u, v) - 2(K - 10)\delta \leq \rho(\lambda(u), \lambda(v)) \leq \frac{K + 1}{K - 9} d_K(u, v).$$

It follows that the restriction of λ to the set V of nonroot vertices of $\mathcal{T}$ is a $\left(\frac{K+1}{K-9}, 2(K - 10)\delta\right)$-quasi-isometry, when $\mathcal{T}$ is given the scaled metric d_K. An application of Lemma A.2 with $M = (K - 9)\delta$ and $N = (K + 1)\delta$ implies that λ is a (μ, C)-quasi-isometry, where $\mu = \frac{K+1}{K-9}$ and $C = 2(3K - 8)\delta$. Setting $1 + \varepsilon = \mu$, and observing that ε monotonically decreases to zero as $K = 9 + 10\varepsilon^{-1}$ monotonically increases, yields the desired result. □

COROLLARY 1.3. *The $(1 + \varepsilon, C(\varepsilon))$-quasi-isometric mapping of the previous theorem extends continuously to a homeomorphism of the Cantor set boundary $\partial\mathcal{T}$ onto a Cantor set contained in the continuum Z.*

REMARK. Let $\mathcal{T}_\varepsilon$ denote the binary tree $\mathcal{T}$ with the scaled metric d_K. Notice from the proof of Theorem 1.2 that the length of each edge of $\mathcal{T}_\varepsilon$ is $L(\varepsilon) = 10\delta/\varepsilon$. It follows that

$$C(\varepsilon) = 60\delta/\varepsilon + 38\delta = 6L(\varepsilon) + 38\delta. \tag{1.9}$$

The following corollary implies the theorem in the Introduction. It and its companion corollary say that every negatively curved space whose boundary contains a nontrivial continuum has embedded binary trees whose distortions from uniformity are small in the scale of the lengths of the edges of the embedded trees.

COROLLARY 1.4. *For each $\varepsilon > 0$, the space X admits a $(1+\varepsilon, C(\varepsilon))$-quasi-isometric embedding of a uniform binary tree $\mathcal{B}_\varepsilon$, where $C(\varepsilon)$ is $O(\varepsilon\ell(\varepsilon))$ as $\varepsilon \downarrow 0$ and $\ell(\varepsilon)$ denotes the length of each edge of $\mathcal{B}_\varepsilon$.*

Proof. First we use Theorem 1.2 and the remark following Corollary 1.3 to identify a $(1+\varepsilon, C(\varepsilon))$-quasi-isometric *mapping* with the desired properties, and then we adjust the mapping to an embedding. Fix $\varepsilon > 0$ so that $9 + 10\varepsilon^{-1} = K$ is an integer, and scale the metric d_K of $\mathcal{T}_\varepsilon$ by a factor of $(K-9)$ to obtain the uniform binary tree $\mathcal{B}_\varepsilon$ in which the length of each edge is

$$\ell(\varepsilon) = (K-9)^2\delta = (K-9)L(\varepsilon) = 10L(\varepsilon)/\varepsilon. \tag{1.10}$$

For each binary sequence b, let $\bar{b}$ denote the binary sequence obtained from b by replacing each digit i of b by a string of $(K-9)$ i's. Let $f\colon \mathcal{B}_\varepsilon \to \mathcal{T}_\varepsilon$ be the isometric embedding defined on the vertices by $f(v_b) = v_{\bar{b}}$, and recall from the proof of Theorem 1.2 that $\lambda\colon \mathcal{T}_\varepsilon \to X$ is a $(1+\varepsilon, C(\varepsilon))$-quasi-isometric mapping, where $C(\varepsilon)$ is given by (1.9). It follows that the composition $\lambda \circ f$ is a $(1+\varepsilon, C(\varepsilon))$-quasi-isometric mapping. Notice that $C(\varepsilon)$ increases without bound as $\varepsilon \downarrow 0$, but (1.9) and (1.10) imply that

$$\frac{C(\varepsilon)}{\varepsilon\ell(\varepsilon)} = \frac{3}{5} + \frac{19\varepsilon}{50},$$

so that $C(\varepsilon)$ is $O(\varepsilon\ell(\varepsilon))$ as $\varepsilon \downarrow 0$.

We now adjust $\lambda \circ f$ to an embedding by redefining λ only near the vertices of the form $v_{\bar{b}}$. We save ourselves an unnecessary technical headache by assuming that ε is no more than 2, so that $(K-9)$ is at least 5. Let $b \neq \varnothing$ be a binary sequence and, without loss of generality, assume that the last digit of b is 0. Then the parent of $v_{\bar{b}}$ in $\mathcal{T}_\varepsilon$ is of the form v_{a0}, for the appropriate binary sequence a, since the sequence $\bar{b}$ ends in a string of $(K-9)$ zeros. Recall that λ is defined on the edge $v_{\bar{b}}v_{\bar{b}1}$ as a convex map to a geodesic segment of ρ-length between $K\delta$ and $(K+1)\delta$ from $\lambda(v_{\bar{b}})$, which lies on the ray $\sigma_{\bar{b}} = \sigma_{a0}$, to $\lambda(v_{\bar{b}1})$, which lies on the ray $\sigma_{\bar{b}1}$. Let ς be the geodesic segment contained in the image $\lambda(v_{\bar{b}}v_{\bar{b}1})$ that meets $\sigma_{\bar{b}}$ at the single point x, an endpoint of ς that lies between $\lambda(v_{a0})$ and $\lambda(v_{\bar{b}0})$, and meets the subsegment of $\sigma_{\bar{b}1}$ between $\lambda(v_{\bar{b}1})$ and $\lambda(v_{\bar{b}11})$ at the single point y, the other endpoint of ς. Define $\lambda'(v_{\bar{b}}) = x$ and $\lambda'(v_{\bar{b}1}) = y$, and extend convexly to map the edge $v_{\bar{b}}v_{\bar{b}1}$ onto ς. A further convex adjustment on the edges $v_{a0}v_{\bar{b}}$, $v_{\bar{b}}v_{\bar{b}0}$, and $v_{\bar{b}1}v_{\bar{b}11}$ produces an embedding λ' of the convex hull of the vertices v_{a0}, $v_{\bar{b}0}$, and $v_{\bar{b}11}$ that agrees with λ at these three vertices, and that is $(K+1)\delta$-close to λ. Repeating this construction at every vertex of $\mathcal{T}_\varepsilon$ of the form $v_{\bar{b}}$, as well as performing a similar construction near $v_\varnothing$, produces a mapping $\lambda' \colon \mathcal{T}_\varepsilon \to X$ that is $(K+1)\delta$-close to λ and, as the reader may check, for which $\Lambda = \lambda' \circ f$ is an embedding. An application of Lemma A.3 implies that Λ is a $(1+\varepsilon, C'(\varepsilon))$-quasi-isometric embedding, where $C'(\varepsilon) = C(\varepsilon) + 2(K+1)\delta$. Using (1.9) and (1.10),

$$\frac{C'(\varepsilon)}{\varepsilon\ell(\varepsilon)} = \frac{4}{5} + \frac{29\varepsilon}{50},$$

so that $C'(\varepsilon)$ is $O(\varepsilon\ell(\varepsilon))$ as $\varepsilon \downarrow 0$. □

The details of the constructions above imply the following corollary.

COROLLARY 1.5. *For each $\varepsilon > 0$, the $(1+\varepsilon, C(\varepsilon))$-quasi-isometric embedding Λ_ε of the uniform binary tree $\mathcal{B}_\varepsilon$ guaranteed by the previous corollary may be chosen so that the ρ-length of the image of every edge e of $\mathcal{B}_\varepsilon$ satisfies*

$$m(\varepsilon) \le \ell_\rho(\Lambda_\varepsilon(e)) \le M(\varepsilon), \tag{1.11}$$

where both bounds $m(\varepsilon)$ and $M(\varepsilon)$ are asymptotic to $\ell(\varepsilon) = 100\delta/\varepsilon^2$, the length of each edge of $\mathcal{B}_\varepsilon$, as $\varepsilon \downarrow 0$.

2. Simple Random Walks and Type

For a good exposition of important results about random walks on graphs that includes a section on the type problem, see the survey article [34]. All graphs considered in this section are connected and locally finite, but we do not assume that they have bounded degree nor that they satisfy any sort of isoperimetric inequality. The simple random walk (SRW) on the graph $X = (V(X), E(X))$ is the random walk for which the transition probability $p(x, y)$ from vertex x to vertex y is given by

$$p(x,y) = \begin{cases} 1/\deg(x) & \text{if } xy \in E(X), \\ 0 & \text{otherwise.} \end{cases}$$

Here $V(X)$ and $E(X)$ denote, respectively, the vertex and edge sets of X. The SRW on X is *recurrent* if, with probability 1, the random walk starting at some vertex x returns to x; otherwise, the SRW is *transient*. In the transient case, there is a positive probability for the event that a random walker will escape to infinity. The graph X is recurrent/transient whenever the SRW on X is recurrent/transient.

The *word metric* on the graph X is an intrinsic metric in which each edge has unit length, and the graph is said to be *negatively curved* or *word hyperbolic* if its word metric is negatively curved. We often use the adjective *combinatorial* to refer to word-metric properties of a graph; for example, combinatorial length, distance, and diameter are always with reference to the word metric.

Let $\mathcal{T}$ be a tree with word metric. A vertex of $\mathcal{T}$ of degree at least 3, as well as the root vertex if the tree is rooted, is called a *branch* vertex. Two branch vertices are *consecutive* if there is a simple edge path between them that traverses no other branch vertices. We say that $\mathcal{T}$ is *quasi-uniform* if $\mathcal{T}$ is nontrivial, there are no vertices of unit degree, and there is a finite upper bound on the combinatorial length of any simple path of degree-2 vertices. Notice that every quasi-uniform tree has uncountably many ends and there is a finite upper bound on the combinatorial distance between any two consecutive branch vertices. Also, $\mathcal{T}$ is *binary* if $\mathcal{T}$ is rooted and the degree of each nonroot vertex is at most 3 while that of the root vertex is at most 2. The following two lemmas are well known and imply, with Corollary 1.5, items (1) and (2) of the Introduction.

LEMMA 2.1. *Every quasi-uniform binary tree is transient*.

Proof. An argument as in [17, Chap. 6] implies that the electrical resistence to infinity is finite, which implies that the tree is transient. □

LEMMA 2.2. *If the locally finite graph X contains a transient subgraph, then X itself is transient*.

Proof. See, for example, [34, Sec. 2]. □

REMARK. Note that these two lemmas together imply that every locally finite, quasi-uniform tree is transient, since every such tree contains a quasi-uniform binary subtree.

THEOREM 2.3. *If X is a locally finite, negatively curved graph whose Gromov boundary contains a nontrivial continuum, then X is transient*.

Proof. Since X is locally finite, its word metric is proper. By Corollary 1.5, there is a $(2, C(1))$-quasi-isometric embedding Λ_1 of the uniform binary tree $\mathcal{B}_1$ into X. Because Λ_1 is an embedding into a graph and the degree of each nonroot vertex of $\mathcal{B}_1$ is 3, the image of each such vertex under the embedding must be a vertex of X. It follows that X contains a binary subtree $\mathcal{B} = \Lambda_1(\mathcal{B}_1)$ that is quasi-uniform. Apply the two preceding lemmas. □

COROLLARY 2.4. *If X is a locally finite, one-ended, negatively curved graph whose Gromov boundary contains more than one point, then X is transient*.

We review some of the previous work on the implications that Gromov negative curvature holds for random walks on graphs. The terms from the general theory of random walks on graphs that we use in this paragraph are defined in, for instance, [34]. These previous results have been derived in more general contexts than that of simple random walks, but have assumed hypotheses that imply stronger geometric conditions on graphs than just mere negative curvature of the word metric. We focus on two results, both derived in the context of a random walk given by a stochastic transition matrix

$$P = (p(x, y))_{x, y \in V(X)}$$

that descibes the one-step transition probabilities. The first is a result of Ancona [4], who proves that if X is negatively curved and P is uniformly irreducible, has bounded range, and has spectral radius less than unity, then the Martin boundary $\partial_M X$ coincides with the Gromov boundary ∂X, and the random walk converges to ∂X almost surely. The second is a result of Kaimanovich and Woess [25], who replace the bounded range hypothesis of Ancona with the weaker hypothesis that P satisfy a uniform first moment condition, and conclude that the random walk converges to ∂X almost surely. The uniform irreducibility condition implies that the graph X has bounded degree. In the context of a SRW, if a graph has bounded degree, then the spectral radius is less than unity if and only if the graph satisfies a strong isoperimetric inequality; [24] and [34, Thm. 3.3]. These results imply, when the random walk is simple, not only that X is transient, but the additional conclusion that the Martin boundary coincides with the Gromov boundary, to which the random walk converges almost surely. The price that is paid, though, for these stronger results are the very strong geometric limitations placed on the graph X that it both have bounded degree and satisfy a strong isoperimetric inequality, limitations absent from our results.

3. Negatively Curved Planar Metrics

Our attention now specializes to planar graphs, particularly those associated to circle packings of the Euclidean and hyperbolic planes. As a preliminary to determining CP-type in the next section, here we collect some basic topological results. The setting throughout is that of a complete, negatively curved geodesic metric ρ on the complex plane $\mathbb{C}$ that is compatible with (i.e., induces the same topology as) the Euclidean one.

THEOREM 3.1. *The Gromov boundary $\partial_\rho \mathbb{C}$ is either a singleton or a topological circle.*

The theorem is a consequence of the next two lemmas.

LEMMA 3.2. *The Gromov boundary $\partial_\rho \mathbb{C}$ is a metric continuum.*

Proof. For any positive number R, let $\bar{U}_\rho(R)$ denote the closure of the unbounded complementary domain of the ρ-ball of radius R centered at the origin. For each positive integer n, let C_n be the union $\bar{U}_\rho(n) \cup \partial_\rho \mathbb{C}$, a subspace of the compactification $\bar{\mathbb{C}} = \mathbb{C} \cup \partial_\rho \mathbb{C}$ of the plane $\mathbb{C}$. That each C_n is a metric continuum follows quickly from the fact that $\mathbb{C}$ has one end, the definition of the topology on $\bar{\mathbb{C}}$, as well as the basic facts that $\bar{\mathbb{C}}$ is compact and metrizable (see the Appendix). As the boundary $\partial_\rho \mathbb{C}$ is the decreasing intersection of the metric continua C_n, it is itself a metric continuum. □

LEMMA 3.3. *Every pair of distinct points of $\partial_\rho \mathbb{C}$ separates $\partial_\rho \mathbb{C}$.*

Proof. Let $a \neq b$ be points of $\partial_\rho \mathbb{C}$ and choose a bi-infinite geodesic γ such that $\gamma(\infty) = a$ and $\gamma(-\infty) = b$. The trace $|\gamma|$ is a closed planar set homeomorphic to a line, and hence separates the plane into exactly two components—say, X° and Y°. Note that the respective restrictions of the metric ρ to the sets $X = X^\circ \cup |\gamma|$ and $Y = Y^\circ \cup |\gamma|$ are proper, geodesic, and negatively curved. It follows as in the proof of the preceding lemma, since X and Y each have one end, that the Gromov boundaries ∂X and ∂Y are metric continua. The inclusions of X and Y in $\mathbb{C}$ induce embeddings of Gromov boundaries, so that ∂X and ∂Y are naturally (closed) subspaces of $\partial_\rho \mathbb{C}$. An easy exercise establishes that the intersection $\partial X \cap \partial Y$ is precisely the pair $\partial\gamma = \{a, b\}$ and, as ∂X and ∂Y are each connected and contain the two points a and b, the sets $U = \partial X - \partial\gamma$ and $V = \partial Y - \partial\gamma$ are nonempty. It follows that the sets U and V form a separation of $\partial_\rho \mathbb{C} - \{a, b\}$. $\square$

Proof of Theorem 3.1. With Lemmas 3.2 and 3.3, an old characterization theorem of Moore [33, Thm. 28.14] applies to show that $\partial_\rho \mathbb{C}$, if not a singleton, is a topological circle. $\square$

We next present a useful property of the metric ρ when the Gromov boundary $\partial_\rho \mathbb{C}$ is a singleton. We assume that ρ has δ-inscribed triangles for the fixed positive constant δ.

THEOREM 3.4. *If the Gromov boundary $\partial_\rho \mathbb{C}$ is a singleton, then there is a constant L such that, for any compact subset K of $\mathbb{C}$, there is a simple closed curve of ρ-length at most L that separates K from infinity.*

Proof. Let γ be a geodesic ray based at the point $\gamma(0) = z_0$. For any point p not on γ, let α_p be a shortest geodesic segment connecting p to γ. Let β_p be a shortest path connecting p to γ from the side of γ “opposite to” the side from which α_p approaches γ. This is unambiguous as long as α_p meets γ at a point other than z_0; when α_p happens to meet γ at z_0, β_p is chosen to equal α_p. The reader might notice that our description of β_p, even when α_p misses the basepoint z_0, allows for the possibility that β_p hits the point z_0. Always in this case, β_p must meet γ along an initial segment of γ. Let $\mathcal{P}$ be the set of points for which the length of the paths α_p and β_p coincide, in which case both are geodesic segments, each meeting γ in exactly one point. Continuity of the metric implies not only that $\mathcal{P}$ is not empty, but, moreover, that $\mathcal{P}$ is an unbounded set in the plane.

In the next paragraph we verify that, for points p of $\mathcal{P}$ far enough from the basepoint z_0, both companion segments α_p and β_p miss the basepoint z_0. For each such point, let α_p and β_p meet γ at the respective points $a_p \neq z_0$ and $b_p \neq z_0$, and consider the triangle (or bigon if a_p and b_p coincide) pa_pb_p with sides $|\alpha_p|$, $|\beta_p|$, and the subsegment of γ between a_p and b_p. Let x_p and y_p be points on the respective segments α_p and β_p of ρ-distance 2δ from the respective points a_p and b_p. By Lemma A.1, since α_p and β_p are shortest paths to γ from p, the ρ-distance from x_p to y_p is at most δ. Let C_p denote the piecewise geodesic path that starts at x_p, travels along α_p to a_p, continues along γ to b_p, then along β_p to y_p, and finally back to x_p along a geodesic segment. Note that the ρ-length of C_p is at most

$2\delta + 4\delta + 2\delta + \delta = 9\delta = L$ (look at the internal points of the isosceles triangle $pa_p b_p$). By the definition of α_p, the segment from x_p to y_p along C_p cannot touch γ, from which it follows that C_p meets γ only along the subsegment of γ between a_p and b_p, at which points C_p approaches γ from opposite sides. This implies that the curve C_p is essential in $\mathbb{C} - \{z_0\}$. The metric ball $B_\rho(z_0, n)$ misses C_p when $n < \rho(z_0, a_p) - 5\delta$, so C_p separates $B_\rho(z_0, n)$ from infinity. By choosing p with $\rho(z_0, a_p)$ large enough, which is possible by an argument similar to that of the next paragraph, we may separate any given compact set from infinity by one of the paths C_p, and the theorem follows.

Finally, if arbitrarily far from the basepoint z_0 there are points of $\mathcal{P}$ at least one of whose companion shortest segments to γ meets γ at z_0, then a sequence $p(i)$ of such points may be extracted for which either the sequence of segments $\alpha_{p(i)}$, or the companion sequence $\beta_{p(i)}$, converges to a ray γ' based at z_0. This limiting ray γ' cannot be asymptotic (see the Appendix) to the ray γ, for this would violate the fact that there is a shortest path from $p(i)$ to γ, for arbitrarily large i, that travels all the way to the basepoint z_0 to meet γ. It follows that the points $[\gamma]$ and $[\gamma']$ of the boundary $\partial_\rho \mathbb{C}$ are not the same, contradicting our assumption that the boundary is a singleton. □

By choosing the points x_p and y_p in the proof above to be slightly more than δ units from the respective points a_p and b_p, rather than 2δ away, one may obtain the value 5δ for the constant L.

4. CP-Type

Two discrete graphical versions of classical extremal length have appeared, the first in 1962 by Duffin [19], the *edge extremal length*, and the second more recently by Cannon [14], the *vertex extremal length*. Edge extremal length is useful for determining the type of a SRW on a locally finite graph [19]. Vertex extremal length is useful for constructing square tilings of rectangles with prescribed patterns of contact [15; 32]. In a remarkable paper [23], He and Schramm use vertex extremal length to give a complete combinatorial characterization of the CP-type of a plane triangulation graph. For bounded degree graphs, the two discrete versions of extremal length agree [23, Thm. 8.1], and one of the impressive accomplishments of [23] is the realization that vertex extremal length is a fine enough sieve with which to determine CP-type in the nonbounded degree setting, where edge extremal length fails. The results to follow concerning the CP-type of negatively curved plane triangulation graphs are verified by simple applications of the work of He and Schramm in [23], along with our results in Sections 1 and 3. We begin by recalling definitions and terminology.

For a graph $\mathcal{G}$, we use $V(\mathcal{G})$ and $E(\mathcal{G})$ to denote, respectively, the vertex and edge sets of $\mathcal{G}$. A *plane triangulation graph* is the 1-skeleton of a triangulation of the plane, and a *circle packing* for the plane triangulation graph $\mathcal{G}$ is a collection $\mathcal{C} = \{ C_v : v \in V(\mathcal{G}) \}$ of Euclidean circles in the plane $\mathbb{C}$ with pairwise disjoint interiors such that C_v is tangent to C_w whenever vw is an edge of $\mathcal{G}$. By connecting the Euclidean centers of tangent circles in the circle packing $\mathcal{C}$ by line segments, we obtain a geometric realization of the abstract graph $\mathcal{G}$ as the 1-skeleton of a geodesic triangulation of, necessarily, a simply-connected domain $D(\mathcal{C})$ in $\mathbb{C}$. This domain $D(\mathcal{C})$ is called the *carrier* of $\mathcal{C}$, though this term in the literature often refers also to its geodesic triangulation described previously. Whenever there is a circle packing for $\mathcal{G}$ with carrier D, we say that $\mathcal{G}$ *packs* the domain D. Notice that every plane triangulation graph is locally finite. The following is the basic existence–uniqueness result concerning infinite circle packings in the plane. In this form, existence is due to He and Schramm [22] and uniqueness to Schramm [31]. Previously, Beardon and Stephenson [5] had proved the result in the setting of bounded degree graphs, and subsequently, He and Schramm, in another impressive accomplishment of [23], have extended the existence result to include packings of more general simply connected domains by sets of more general shapes than circles.

CIRCLE PACKING THEOREM. *Every plane triangulation graph packs exactly one of the plane $\mathbb{C}$ or the unit disk $\mathbb{D}$. A circle packing for a plane triangulation graph with carrier either $\mathbb{C}$ or $\mathbb{D}$ is unique up to Möbius transformations that fix the carrier.*

DEFINITION. The plane triangulation graph $\mathcal{G}$ is *CP-parabolic* if it packs the plane $\mathbb{C}$ and *CP-hyperbolic* if it packs the disk $\mathbb{D}$. The circle packing theorem implies that every plane triangulation graph is either CP-parabolic or CP-hyperbolic, but never both. A *maximal* circle packing for $\mathcal{G}$ is one with carrier either $\mathbb{C}$ or $\mathbb{D}$.

DEFINITION. The locally finite, connected graph $\mathcal{G}$ is *RW*-parabolic if the SRW on $\mathcal{G}$ is recurrent and *RW*-hyperbolic if it is transient.

For bounded degree plane triangulation graphs, these two notions of type coincide [23], but—though every CP-hyperbolic graph is RW-hyperbolic—there are plane triangulation graphs of unbounded degree that are CP-parabolic and, at the same time, RW-hyperbolic [23, Thm. 8.2]. Duffin's notion of edge extremal length captures the RW-type of a graph while Cannon's notion of vertex extremal length, as shown in [23], captures the CP-type. We refer the reader to Section 2 of [23] for the general definitions of combinatorial extremal length, and are content with quoting the results from [23] that meet our purposes.

DEFINITION. An infinite graph $\mathcal{G}$ is *VEL*-parabolic if, for some (hence, any) vertex v, the vertex extremal length of the family $\Gamma(v,\infty)$ of infinite, unbounded paths of vertices based at v is infinite; otherwise, $\mathcal{G}$ is *VEL*-hyperbolic. Similarly, $\mathcal{G}$ is *EEL*-parabolic if, for some (hence, any) vertex v, the edge extremal length of the family $\Gamma(v,\infty)$ is infinite; otherwise, $\mathcal{G}$ is *EEL*-hyperbolic.

We now quote [23, Thm. 7.2], He and Schramm's combinatorial characterization of CP-type in terms of vertex extremal length.

CHARACTERIZATION THEOREM. *A plane triangulation graph is CP-parabolic if and only if it is VEL-parabolic; equivalently, CP-hyperbolic if and only if VEL-hyperbolic.*

Recall that a graph is negatively curved if its word metric is so. The main purpose of this section is to verify the following theorem. Its proof consists of making several observations that allow us to apply Corollary 1.5, Theorem 3.1, Theorem 3.4, and the characterization theorem.

THEOREM 4.1. *A negatively curved plane triangulation graph is CP-parabolic if and only if its Gromov boundary is a singleton. A negatively curved plane triangulation graph is CP-hyperbolic if and only if its Gromov boundary is a topological circle.*

Let K be a triangulation of the plane $\mathbb{C}$ whose 1-skeleton is the graph $\mathcal{G}$, and define $|K|_{\text{eq}}$ to be the metric realization of K obtained by identifying each face of K with a Euclidean unit equilateral triangle, and using the induced intrinsic metric. Our claim is that the resulting geodesic metric space $|K|_{\text{eq}}$ is quasi-isometric with the graph $\mathcal{G}$ with word metric. This follows quickly from the next observation, whose easy verification is left to the reader.

LEMMA 4.2. *Let (X, ρ) be a metric space and Y a quasi-dense subset of X. Suppose that there is a constant $\mu \geq 1$ such that each pair of points y_0 and y_1 of Y are joined by a ρ-rectifiable path in Y of ρ-length at most $\mu\rho(y_0, y_1)$. Then (Y, d) is quasi-isometric to (X, ρ), where d is the intrinsic metric determined by the restriction of ρ to Y; indeed, the inclusion of Y in X is a $(\mu, 0)$-quasi-isometry with quasi-dense image.*

EXERCISE. Obviously, the 1-skeleton $\mathcal{G}$ of K is $1/\sqrt{3}$-dense in $|K|_{\text{eq}}$. Moreover, if x and y are points on two sides of a Euclidean unit equilateral triangle $\triangle$ that share the vertex v, then the path from x to y through v in $\partial\triangle$ has length at most $4/\sqrt{3}$ times the straight line distance in $\triangle$ from x to y. Since the metric of $|K|_{\text{eq}}$ is geodesic and equilateral on faces, Lemma 4.2 applies with $\mu = 4/\sqrt{3}$ to show that $\mathcal{G}$ with word metric is quasi-isometric to $|K|_{\text{eq}}$.

LEMMA 4.3. *The Gromov boundary $\partial\mathcal{G}$ of the negatively curved plane triangulation graph $\mathcal{G}$ is either a singleton or a topological circle.*

Proof. As $\mathcal{G}$ is the 1-skeleton of the triangulation K, the previous exercise implies that $|K|_{\text{eq}}$ is quasi-isometric to $\mathcal{G}$ with word metric. Because negative curvature is a quasi-isometry invariant, the metric of $|K|_{\text{eq}}$ is negatively curved and, since $\mathcal{G}$ is locally finite, the metric of $|K|_{\text{eq}}$ is proper and hence complete. Since K triangulates $\mathbb{C}$, Theorem 3.1 implies that the Gromov boundary $\partial|K|_{\text{eq}}$ is either a singleton or a topological circle. By the last paragraph of the Appendix and Lemma 4.2, the inclusion of $\mathcal{G}$ into $|K|_{\text{eq}}$ induces a homeomorphism of Gromov boundaries. $\square$

Proof of Theorem 4.1. The forward implications of the two statements of the theorem follow from Lemma 4.3, the circle packing theorem, and the reverse implications, which are proved next.

According to the monotonicity property [23, 2.1], if $\mathcal{G}$ contains a VEL-hyperbolic subgraph then $\mathcal{G}$ itself is VEL-hyperbolic. Assume that the Gromov boundary $\partial\mathcal{G}$ is a topological circle. By Corollary 1.5, as in the proof of Theorem 2.3, $\mathcal{G}$ contains a quasi-uniform binary tree $\mathcal{B}$ as a subgraph. It is a nice exercise in the calculation of combinatorial extremal length to show that every quasi-uniform binary tree is VEL-hyperbolic. Alternately, as $\mathcal{B}$ has bounded degree, the VEL-type of $\mathcal{B}$ coincides with the EEL-type [23, Thm. 8.1], and since EEL-type coincides with RW-type for locally finite graphs [23, Thm. 2.6], Lemma 2.1 implies that $\mathcal{B}$ is VEL-hyperbolic. Therefore, $\mathcal{G}$ is VEL-hyperbolic and the characterization theorem applies to show that $\mathcal{G}$ is CP-hyperbolic.

Assume now that the Gromov boundary $\partial\mathcal{G}$ is a singleton and that the word metric on $\mathcal{G}$ has δ-inscribed triangles for a fixed positive constant δ. Notice that since $\mathcal{G}$ is not a tree, δ is at least unity. Fix a basepoint v_0, a vertex of $\mathcal{G}$. By Theorem 3.4, there is a positive constant L such that every compact subset of $|K|_{\mathrm{eq}}$ is separated from infinity by a path of length at most L. By the exercise after Lemma 4.2, each such separating path may be replaced by a separating path in the 1-skeleton $\mathcal{G}$ of $|K|_{\mathrm{eq}}$ of length at most $(4/\sqrt{3})L$. This allows us to construct inductively a sequence of pairwise disjoint cycles c_n, each an edge path in the graph $\mathcal{G}$ of combinatorial length at most $(4/\sqrt{3})L$, and each separating v_0 from infinity. It follows that every element of $\Gamma(v_0,\infty)$ meets every $|c_n|$. Define a vertex label (or v-metric in the terminology of [23]) $m\colon V(\mathcal{G}) \to [0,\infty)$ by

$$m(v) = \begin{cases} 1/n & \text{if } v \text{ is a vertex of the cycle } c_n, \\ 0 & \text{otherwise,} \end{cases}$$

which is well-defined since the cycles c_n are pairwise disjoint. Then the m-length of each path γ in $\Gamma(v_0,\infty)$,

$$L_m(\gamma) = \sum_{i=1}^{\infty} m(\gamma(i)),$$

is infinite while the m-area $\sum_{v\in V(\mathcal{G})} m(v)^2$ is finite. It follows that the vertex extremal length of the family $\Gamma(v_0,\infty)$ is infinite (see [23, Sec. 2]) and, therefore, $\mathcal{G}$ is VEL-parabolic. The characterization theorem applies to show that $\mathcal{G}$ is CP-parabolic. $\square$

The calculation of vertex extremal length in case the boundary is a singleton in the foregoing proof works in more general settings, and is easily modified to prove the following corollary.

COROLLARY 4.4. *Let v_0 be a vertex in the graph $\mathcal{G}$ and let $\{V_n\}$ be a sequence of pairwise disjoint sets of vertices, each of which separates v_0 from ∞. Suppose there exist positive constants C and ε such that, for each n,*

$$\mathrm{Card}(V_n) \le C n^{1-\varepsilon}.$$

Then the graph $\mathcal{G}$ is VEL-parabolic.

Finally, we mention a result of Northshield [29; 30], who constructed a boundary for bounded degree, planar graphs that satisfy a strong isoperimetric inequality and for which every circuit surrounds only finitely many vertices. He then proved that his boundary, say $\partial_N \mathcal{G}$ for the graph $\mathcal{G}$, is either a topological circle or a singleton, and that the SRW on $\mathcal{G}$ converges almost surely to a $\partial_N \mathcal{G}$-valued random variable. If $\mathcal{G}$ is in addition negatively curved, the Northshield boundary can be identified with the Poisson boundary, and if $\mathcal{G}$ is also a plane triangulation graph then the Northshield boundary is the Martin boundary. See [34, Sec. 7.E] for definitions and further references.

5. Further Applications

We do not need that the embedded binary tree $\mathcal{B}$ used in the proof of Theorem 4.1 be quasi-uniform to conclude that $\mathcal{G}$ is CP-hyperbolic, only the weaker condition that it be transient. That $\mathcal{B}$ is in fact quasi-uniform gives more information about the maximal circle packing for $\mathcal{G}$ than merely that its carrier is the disk $\mathbb{D}$. For example, the quasi-uniform condition can be used to derive lower bounds on the hyperbolic radii of certain circles in the maximal packing, as well as various quasi-denseness results about circles that have uniformly large hyperbolic radii. We quote two illustrative results that will appear later [9]. In both, we assume that the unit disk $\mathbb{D}$ carries its Poincaré metric, making it a model for hyperbolic geometry. Recall that there is a finite upper bound on the combinatorial length of any simple path of degree-2 vertices in a quasi-uniform binary tree $\mathcal{B}$. The smallest such upper bound is called the *fundamental length* for $\mathcal{B}$.

THEOREM 5.1. *If $\mathcal{G}$ is a bounded-degree plane triangulation graph with maximal packing $\mathcal{C}$ that contains a quasi-uniform, binary subtree $\mathcal{B}$, then $\mathcal{G}$ is CP-hyperbolic and there is a uniform positive lower bound λ on the hyperbolic radii of any circles C_v of $\mathcal{C}$ that correspond to vertices v of the tree $\mathcal{B}$. The bound λ depends only on the maximum degree of $\mathcal{G}$ and the fundamental length of $\mathcal{B}$.*

THEOREM 5.2. *If $\mathcal{G}$ is a CP-hyperbolic, negatively curved plane triangulation graph of bounded degree with maximal packing $\mathcal{C}$, then there is a positive constant λ for which the union of the circles of $\mathcal{C}$ of hyperbolic radii at least λ forms a quasi-dense subset of the disk in the Poincaré metric. The constant λ depends only on the maximum degree of $\mathcal{G}$ and the thinness constant δ for the triangles of $\mathcal{G}$.*

Appendix: Negatively Curved Metrics—an Overview

References for recent metric geometry of negatively curved spaces are Gromov [21], Cannon [13], and Alonso et al. [3], and, for the related topic of nonpositively curved spaces, Bridson and Haefliger [10]. The older references Blumenthal [6], Blumenthal and Menger [7], Busemann [11; 12], Aleksandrov et al. [1], and Aleksandrov and Zalgaller [2] are invaluble both for comprehensive treatments of metric geometry and for expositions of the initial developments on metric curvature by a previous generation of mathematicians.

The metric ρ on the space X is *geodesic* if, for each pair of points x and y of X, there is a ρ*-segment* with endpoints x and y. This means that there is an isometric map σ into X defined on the interval $[0, \rho(x, y)]$ with $\sigma(0) = x$ and $\sigma(\rho(x, y)) = y$. The metric ρ is *proper* if closed metric balls $\bar{B}_\rho(x, R)$, for x in X and $R > 0$, are compact.

EXERCISE. The geodesic metric ρ is proper if and only if it is complete and the space X is locally compact.

The metric ρ is *rectifiable* if each pair of points of X are endpoints of a ρ-rectifiable path, a path γ of finite length

$$\ell_\rho(\gamma) = \sup_{\mathcal{P}} \left\{ \sum \rho(\gamma(t_i), \gamma(t_{i+1})) \right\},$$

where the supremum is taken over all partitions $\mathcal{P}$ of the domain interval of γ. The *intrinsic metric* on X determined by a rectifiable metric ρ is denoted as ρ_* and defined by

$$\rho_*(x, y) = \inf\{\, \ell_\rho(\gamma) : \gamma \text{ is a path containing } x \text{ and } y \,\}.$$

The rectifiable metric ρ is *intrinsic* if $\rho = \rho_*$. All geodesic metrics are intrinsic.

EXERCISE. If the metric ρ on the space X is intrinsic, then X is connected and locally connected.

INSCRIBED TRIANGLES AND METRIC NEGATIVE CURVATURE. There are various equivalent ways of formulating an asymptotic version of negative curvature for geodesic metrics that captures the behavior that one expects from experience with simply connected, negatively curved Riemannian manifolds. We prefer to work with the geometrically appealing notions of thin and inscribed triangles, rather than Gromov's original approach of using hyperbolic inner products that, though at times offering cleaner proofs and constructions, is less intuitive to the uninitiated.

Let (X, ρ) be a geodesic metric space and δ a nonnegative constant. Although there may fail to be unique geodesics between points of X, it should cause no confusion to use the notation xy to denote some geodesic segment with endpoints x and y. With this notation, xyz denotes a set consisting of three geodesic segments xy, yz, and xz. The triangle xyz is δ*-thin* provided the ρ-distance from any point on any side of xyz to the union of the other two sides is at most δ. The *internal points* of xyz are the points $\mu(x)$ on segment yz, $\mu(y)$ on segment xz, and $\mu(z)$ on segment xy for which

$$\rho(x, \mu(y)) = \rho(x, \mu(z)), \quad \rho(y, \mu(x)) = \rho(y, \mu(z)),$$
$$\text{and} \quad \rho(z, \mu(x)) = \rho(z, \mu(y)).$$

If xyz is a Euclidean triangle in $\mathbb{C}$, the internal points are the points of tangency of the circle inscribed in xyz. The triangle xyz is δ*-inscribed* if the ρ-diameter of the set $\{\mu(x), \mu(y), \mu(z)\}$ of internal points is at most δ.

REMARK. The reader should be cautious as the topic is young enough that terminology has not solidified. In [3], for example, thin triangles are referred to as "slim triangles", whereas [13] conforms to our usage. A seemingly stronger notion of inscribed triangles than ours is referred to in [3] as "thin triangles", while having δ-inscribed triangles is rendered as "the insize is bounded by δ."

DEFINITION. We say that the geodesic metric ρ on X has *δ-thin (-inscribed) triangles* if every geodesic triangle in X is δ-thin (-inscribed), and has *thin (-inscribed) triangles* if it has δ-thin (-inscribed) triangles for some nonnegative constant δ. In either case, the constant δ is called a *thinness constant* for ρ or X.

In this paper, we primarily use the property of δ-inscribed triangles showcased in the following lemma. It is for this reason that we have introduced the lesser-used term "inscribed triangles" rather than only the more common "thin triangles".

LEMMA A.1. *If (X, ρ) has δ-inscribed triangles, then every geodesic triangle xyz is δ-*uniform, *meaning that if the points a on segment $x\mu(z)$ and b on segment $x\mu(y)$ satisfy $\rho(x, a) = \rho(x, b)$, then the ρ-distance between a and b is at most δ, and similarly with x, y, and z permuted.*

Proof. Since ρ has δ-inscribed triangles, it suffices to find points y' on segment xy and z' on segment xz such that the points a and b are two of the internal points of the triangle $xy'z'$. Let $\beta, \gamma : [0, 1] \to X$ be unit-time parameterizations (i.e., proportional to arclength) of the respective geodesic segments xy and xz with $\beta(0) = x = \gamma(0)$ and $\beta(1) = y$ and $\gamma(1) = z$. Define $s : [0, 1] \to \mathbb{R}$ by

$$s(t) = \tfrac{1}{2}[\rho(x, \beta(t)) + \rho(x, \gamma(t)) - \rho(\beta(t), \gamma(t))],$$

so that the internal points of triangle $x\beta(t)\gamma(t)$ opposite $\beta(t)$ and $\gamma(t)$, respectively, are

$$\mu(\beta(t)) = \gamma\left(\frac{s(t)}{\rho(x, z)}\right) \quad \text{and} \quad \mu(\gamma(t)) = \beta\left(\frac{s(t)}{\rho(x, y)}\right).$$

Because s is continuous with $s(0) = 0$ and

$$s(1) = r = \tfrac{1}{2}(\rho(x, y) + \rho(x, z) - \rho(y, z)) = \rho(x, \mu(y)) = \rho(x, \mu(z)),$$

s takes on every value between 0 and r. Let $r_0 = \rho(x, a)$ and observe that r_0 is in the interval $[0, r]$. It follows that there is a number t_0 in $[0, 1]$ such that $s(t_0) = r_0$, and a and b are internal points opposite the respective vertices $z' = \gamma(t_0)$ and $y' = \beta(t_0)$ in the triangle $xy'z'$. $\square$

EXERCISE. Prove that every δ-thin triangle is 4δ-inscribed.

DEFINITION. The previous lemma and exercise show that the geodesic metric ρ has thin triangles if and only if it has inscribed triangles. When ρ has thin or inscribed triangles, we say that ρ is *asymptotically negatively curved* or *hyperbolic.* We shall usually delete the descriptively correct adjective "asymptotically" and refer simply to a *negatively curved metric.*

QUASI-ISOMETRIES. A not necessarily continuous function $\lambda\colon T \to X$ between the metric spaces (T, d) and (X, ρ) is a (μ, C)-*quasi-isometry*, where $\mu \geq 1$ and $C \geq 0$ are constants, if for all points u and v of T,

$$\frac{1}{\mu} d(u, v) - C \leq \rho(\lambda(u), \lambda(v)) \leq \mu d(u, v) + C.$$

A *quasi-isometric mapping* is a continuous quasi-isometry, and a *quasi-isometric embedding* is one that is also a topological embedding. A subset V of T is M*-dense* if every point of T is M-close to some point of V, and is *quasi-dense* in T if it is M-dense for some $M \geq 0$. Both (T, d) and (X, ρ) are *quasi-isometric* if there is a quasi-isometry $T \to X$ whose image is quasi-dense in X. In this case, there is a quasi-isometry $X \to T$ whose image is quasi-dense in T. In fact, quasi-isometry is an equivalence relation on the class of metric spaces.

LEMMA A.2. *For positive constants M and N, if V is an M-dense subspace of T and if $\lambda\colon T \to X$ is a function for which the restriction $\lambda|_V$ is a (μ, C)-quasi-isometry such that, for every v in V, the λ-image of the closed metric ball $\bar{B}_d(v, M)$ is contained in the closed metric ball $\bar{B}_\rho(\lambda(v), N)$, then λ is a $(\mu, C+2\mu M+2N)$-quasi-isometry.*

Proof. The proof is an exercise in the use of the triangle inequality. □

LEMMA A.3. *Let $\lambda, \lambda'\colon T \to X$ be functions such that λ is a (μ, C)-quasi-isometry and λ' is N-close to λ. Then λ' is a $(\mu, C + 2N)$-quasi-isometry.*

Proof. The proof is an even easier exercise in the use of the triangle inequality than that of the previous lemma. □

EXERCISE. Any geodesic metric quasi-isometric to a negatively curved metric is itself negatively curved.

DIVERGENCE OF GEODESIC RAYS. Throughout the remainder of this appendix, (X, ρ) is a negatively curved geodesic metric space with, say, δ-inscribed triangles for some fixed positive δ, and x_0 denotes a fixed basepoint in X. The next lemma presents the key divergence property of geodesic rays in a negatively curved metric space that is used several times in Section 1 to prove the theorem of the Introduction. Since we use it so often in this paper, and since we could find no proof of exactly the inequality presented in the lemma, we include a proof.

LEMMA A.4. *Let σ_0 and σ_1 be arclength parameterized geodesic rays (or segments) in X based at x_0 and suppose that, for some parameter value t_0, the ρ-distance between $\sigma_0(t_0)$ and $\sigma_1(t_0)$ is at least 2δ. Then*

$$\rho(\sigma_0(t), \sigma_1(t)) \geq 2(t - t_0) + \rho(\sigma_0(t_0), \sigma_1(t_0)) - \delta$$

for all $t \geq t_0$ (for which $\sigma_0(t)$ and $\sigma_1(t)$ are defined).

Proof. Let $2r$ denote the ρ-distance between $\sigma_0(t_0)$ and $\sigma_1(t_0)$. Since $r \leq t_0$,

$$\delta/2 \leq t_0 - r + \delta/2.$$

For any τ in the half-open interval $(t_0 - r + \delta/2, t_0]$, the triangle inequality implies that

$$2r \leq 2(t_0 - \tau) + \rho(\sigma_0(\tau), \sigma_1(\tau)) < 2r - \delta + \rho(\sigma_0(\tau), \sigma_1(\tau)),$$

implying that

$$\delta < \rho(\sigma_0(\tau), \sigma_1(\tau)). \tag{A.1}$$

Suppose there exists a parameter value t greater than or equal to t_0 for which

$$\rho(\sigma_0(t), \sigma_1(t)) \leq \delta. \tag{A.2}$$

Let $t_{\min}$ be the least such value of t. By continuity of the metric,

$$\rho(\sigma_0(t_{\min}), \sigma_1(t_{\min})) = \delta,$$

and the triangle inequality implies, since the ρ-distance between $\sigma_0(t_0)$ and $\sigma_1(t_0)$ is at least 2δ, that

$$t_0 \leq t_{\min} - \delta/2. \tag{A.3}$$

Because the points $\sigma_0(t_{\min} - \delta/2)$ and $\sigma_1(t_{\min} - \delta/2)$ are internal points of a triangle with vertices x_0, $\sigma_0(t_{\min})$, and $\sigma_1(t_{\min})$, the ρ-distance between them is at most δ. By (A.3), $t_{\min}$ is not the least parameter value greater than or equal to t_0 that satisfies (A.2), a contradiction. This with the previous calculation shows that (A.1) holds for every τ larger than $t_0 - r + \delta/2$.

Now let t be any parameter value greater than t_0 and set $2s$ equal to the ρ-distance between $\sigma_0(t)$ and $\sigma_1(t)$. Then $\sigma_0(t-s)$ and $\sigma_1(t-s)$ are internal points of a triangle with vertices x_0, $\sigma_0(t)$, and $\sigma_1(t)$, and hence the ρ-distance between them is at most δ. From the previous paragraph it follows that

$$t - s \leq t_0 - r + \delta/2.$$

Multiplication by 2 and rearrangement of this inequality yield the inequality that we seek. □

Corollary A.5. *Under the hypotheses of the previous lemma,*

$$\rho(\sigma_0(s), \sigma_1(t)) \geq (t - t_0) + (s - t_0) + \rho(\sigma_0(t_0), \sigma_1(t_0)) - \delta$$

for all $s, t \geq t_0$ (*for which* $\sigma_0(s)$ *and* $\sigma_1(t)$ *are defined*).

Proof. Use the triangle inequality and the previous lemma. □

The next theorem presents a characteristic property of geodesic rays in a negatively curved space. Its proof is essentially that of [3, Thm. 2.19], except that the proof there assumes (without so stating) that the constant δ is greater than $\frac{1}{2}$.

THEOREM A.6 (Exponential Divergence of Geodesics). *Let σ_0 and σ_1 be arc-length parameterized geodesic rays in X based at x_0 and suppose that the ρ-distance between $\sigma_0(t_0)$ and $\sigma_1(t_0)$ is greater than δ for some parameter value t_0. Then, for each nonnegative integer n and each path γ from $\sigma_0(t_0+n\delta)$ to $\sigma_1(t_0+n\delta)$ in the complement of the open ball $B_\rho(x_0, t_0+n\delta)$, the ρ-length of γ satisfies*

$$\ell_\rho(\gamma) \geq 2^{n-1}\delta.$$

THE GROMOV BOUNDARY. There are two equivalent approaches to the Gromov boundary and its topology in the literature—in terms of Gromov's hyperbolic inner product and equivalence classes of sequences that are convergent at infinity (see [21]) and, alternately, in terms of equivalence classes of fellow-travelling geodesic rays and Cannon's combinatorial half-spaces (see [13]). The former has the advantage of oftentimes providing very clean proofs of convergence results whereas the latter has the advantage of providing good, accurate geometric intuition. We present an approach here that has the advantages of both Gromov's and Cannon's and, at the same time, is very concise in its description and confirmation that it defines a topology. This approach uses Cannon's preference in [13] for describing the boundary in terms of equivalence classes of fellow-travelling geodesic rays, but describes the topology by prescribing precisely when a sequence of such classes converges. Any readers familiar with the two standard descriptions of the boundary will see immediately that our description defines the same topology as the former ones. We shall restrict our attention to the setting in which the metric, in addition to being negatively curved, is also proper, in which case the boundary provides a compactification $\bar{X} = X \cup \partial X$.

A *geodesic ray* is an isometric embedding of the interval $[0,\infty)$ into X, and two geodesic rays σ_0 and σ_1 are *asymptotic,* denoted as $\sigma_0 \sim \sigma_1$, if the Hausdorff distance between their images is finite, meaning that each is contained in the n-neighborhood of the other, for some positive constant n. In this case, it is easy to see that in fact they $(\varrho + 2n)$-*fellow travel,* where ϱ is the ρ-distance between $\sigma_0(0)$ and $\sigma_1(0)$; this means that the ρ-distance between $\sigma_0(t)$ and $\sigma_1(t)$ is at most $(\varrho + 2n)$ for every parameter value t. We say that the ray σ is *based* at $\sigma(0)$ and that its *trace* is the image $|\sigma| = \sigma([0,\infty))$. In our context in which (X,ρ) has δ-inscribed triangles, Lemma A.4 implies that when two asymptotic rays σ_0 and σ_1 are both based at the common point x_0, they 2δ-fellow travel. We use the notation $[\sigma]$ to denote the equivalence class of the ray σ under the equivalence relation of being asymptotic, and ∂X to denote the set of such equivalence classes. For each geodesic ray σ, extend its domain to the extended interval $[0,\infty]$ by defining $\sigma(\infty)$ to be the point $[\sigma]$ of the boundary ∂X. We then say that σ is a geodesic *segment* from $\sigma(0)$ to $\sigma(\infty) = [\sigma]$.

For each point c in the boundary, there is a geodesic ray σ based at x_0 such that $c = [\sigma]$. To see this, let τ be a ray for which $c = [\tau]$ and, for each positive integer n, let σ_n be a geodesic segment from x_0 to $\tau(n)$. Use the fact that ρ is proper in conjunction with a diagonal argument to show that a subsequence of the segments σ_n converges to a geodesic ray σ. Use thin triangles to show that σ and τ

are asymptotic. We use the phrase ∂X *is parameterized from* x_0 to indicate a context in which all geodesic rays representing any point of the boundary are based at x_0. Notice that we may write

$$\partial X = \{\, \sigma(\infty) : \sigma \text{ is based at } x_0 \,\}.$$

Let $\bar{X}$ denote the union $X \cup \partial X$, and define a sequence of points c_n in $\bar{X}$ to *converge* to the point c of $\bar{X}$ in the following way, depending on whether or not c is a boundary point. When c is in X, $c_n \to c$ means that there exists a positive integer N such that the sequence $\{\, c_n : n \geq N \,\}$ is contained in X and converges to c with respect to the metric ρ. When c is a boundary point, $c_n \to c$ means that for every $R \geq 0$ there exists a positive integer N such that

(i) the ρ-ball $B_\rho(x_0, R)$ does not contain c_n whenever $n \geq N$, and
(ii) for each $n \geq N$, there exist segments σ_n from x_0 to c_n and σ from x_0 to c that 6δ-fellow travel on the interval $[0, R]$.

Easily, limits of convergent sequences are unique and subsequences of a convergent sequence converge to the limit of the sequence.

Exercise. Use Lemma A.4 to show that the convergent sequences and their limits remain unchanged if item (ii) is replaced by the following statement: For each $n \geq N$, *all* segments σ_n from x_0 to c_n and σ from x_0 to c 6δ-fellow travel on the interval $[0, R]$.

Exercise. Convergence of points of $\bar{X}$ to a boundary point is defined with reference to a basepoint x_0. Use thin triangles and a diagonal argument to verify that the definition does not depend on the basepoint.

Define a subset of $\bar{X}$ to be *closed* if it contains all its limit points, where the point c of $\bar{X}$ is a *limit point* of a set C if there is a sequence of points c_n in C that converges to c. It is an easy exercise to verify that this defines a topology on $\bar{X}$; that the union of two closed sets is closed uses the observation that limits are preserved by subsequences. Throughout the paper proper, $\bar{X}$ carries this topology and ∂X carries the subspace topology it inherits from $\bar{X}$. Obviously, the subspace topology that X inherits from $\bar{X}$ is exactly the ρ-metric topology, and it is easy to see that X is open in $\bar{X}$. Moreover, the diagonal argument (alluded to previously for constructing rays from sequences of segments using the properness of the metric) shows quickly that $\bar{X}$ is sequentially compact. That $\bar{X}$ is metrizable is proved in [16; 20] and, coupled with the sequential compactness, shows that both $\bar{X}$ and its closed subspace ∂X are compact. Thus, $\bar{X}$ is a compactification of X and, for each geodesic ray σ in X, the extended ray gives an embedding of the extended interval $[0, \infty]$ into $\bar{X}$—where, of course, $[0, \infty]$ has the usual topology, making it an arc.

One of the most important properties of the Gromov boundary is that quasi-isometries between proper, negatively curved metric spaces extend continuously to boundaries. This in turn follows from the fact that *quasi-geodesic rays*, that is, quasi-isometries of the half-line $[0, \infty)$ into X, fellow-travel actual geodesic rays. This may be proved with the aid of [3, Prop. 3.3], and a consequence is that quasi-isometric, proper, negatively curved metric spaces have homeomorphic boundaries.

References

[1] A. D. Aleksandrov, V. N. Berestovskii, and I. G. Nikolaev, *Generalized Riemannian spaces,* Russian Math. Surveys 41 (1986), 1–54.

[2] A. D. Aleksandrov, V. A. Zalgaller, and J. M. Danskin, *Intrinsic geometry of surfaces,* Transl. Math. Monographs, 15, Amer. Math. Soc., Providence, RI, 1967.

[3] J. Alonso, T. Brady, D. Cooper, T. Delzant, M. Lustig, M. Mihalik, M. Shapiro, H. Short, and A. N. Other, *Notes on word hyperbolic groups,* Group theory from a geometrical viewpoint (E. Ghys, A. Haefliger, A. Verjovsky, eds.), pp. 3–63, World Scientific, River Edge, NJ, 1991.

[4] A. Ancona, *Positive harmonic functions and hyperbolicity,* Lecture Notes in Math., 1344, pp. 1–23, Springer, Berlin, 1988.

[5] A. F. Beardon and K. Stephenson, *The uniformization theorem for circle packings,* Indiana Univ. Math. J. 39 (1990), 1383–1425.

[6] L. M. Blumenthal, *Theory and application of distance geometry,* 2nd ed., Chelsea, New York, 1970 (reprint of 1953 edition).

[7] L. M. Blumenthal and K. Menger, *Studies in geometry,* Freeman, San Francisco, 1970.

[8] B. H. Bowditch, *Notes on Gromov's hyperbolicity criterion for path-metric spaces,* Group theory from a geometrical viewpoint (E. Ghys, A. Haefliger, A. Verjovsky, eds.), pp. 64–167, World Scientific, River Edge, NJ, 1991.

[9] P. L. Bowers and K. Stephenson, in preparation.

[10] M. Bridson and A. Haefliger, book in preparation.

[11] H. Busemann, *Spaces with nonpositive curvature,* Acta Math. 48 (1948), 259–310.

[12] ———, *The geometry of geodesics,* Academic Press, New York, 1955.

[13] J. W. Cannon, *The theory of negatively curved spaces and groups,* Ergodic theory, symbolic dynamics, and hyperbolic spaces (T. Bedford, M. Keane, C. Series, eds.), pp. 315–369, Oxford Univ. Press, 1991.

[14] ———, *The combinatorial Riemann mapping theorem*, Acta Math. 173 (1994), 155–234.

[15] J. W. Cannon, W. J. Floyd, and W. R. Parry, *Squaring rectangles: The finite Riemann mapping theorem*, Contemp. Math., 169, pp. 133–212.

[16] M. Coornaert, T. Delzant, and A. Papadopoulus, *Notes sur les groupes hyperboliques de Gromov,* Lecture Notes in Math., 1441, Springer, Berlin, 1990.

[17] P. G. Doyle and J. L. Snell, *Random walks and electric networks,* Carus Math. Monographs, 22, Math. Assoc. Amer., Washington, DC, 1984.

[18] T. Dubejko, *Recurrent random walks, Liouville's theorem, and circle packings,* Math. Proc. Cambridge Philos. Soc. 121 (1997), 531–546.

[19] R. J. Duffin, *The extremal length of a network,* J. Math. Anal. Appl. 5 (1962), 200–215.

[20] E. Ghys and P. de la Harpe (eds.), *Sur les groupes hyperboliques d'après Mikhael Gromov,* Progr. in Math., 83, Birkhäuser, Boston, 1990.

[21] M. Gromov, *Hyperbolic groups,* Essays in group theory (S. M. Gersten, ed.), pp. 75–263, Springer, Berlin, 1987.

[22] Z-X. He and O. Schramm, *Fixed points, Koebe uniformization and circle packings,* Ann. of Math. (2) 137 (1993), 369–406.

[23] ———, *Hyperbolic and parabolic packings,* Discrete Comput. Geom. 14 (1995), 123–149.

[24] V. A. Kaimanovich, *Dirichlet norms, capacities, and generalized isoperimetric inequalities for Markov operators,* Potential Anal. 1 (1992), 61–82.

[25] V. A. Kaimanovich and W. Woess, *The Dirichlet problem at infinity for random walks on graphs with a strong isoperimetric inequality,* Probab. Theory Related Fields 91 (1992), 445–466.

[26] M. Kanai, *Rough isometries and combinatorial approximations of geometries of non-compact Riemannian manifolds,* J. Math. Soc. Japan 37 (1985), 391–419.

[27] ———, *Rough isometries and the parabolicity of Riemannian manifolds,* J. Math. Soc. Japan 38 (1986), 227–238.

[28] G. J. McCaughan, *Some results on circle packings,* Ph.D. thesis, University of Cambridge, 1996.

[29] S. Northshield, *Geodesics and bounded harmonic functions on infinite planar graphs,* Proc. Amer. Math. Soc. 113 (1991), 229–233.

[30] ———, *Circle boundaries of planar graphs,* Potential Anal. 2 (1993), 299–314.

[31] O. Schramm, *Rigidity of infinite (circle) packings,* J. Amer. Math. Soc. 4 (1991), 127–149.

[32] ———, *Square tilings with prescribed combinatorics,* Israel J. Math. 84 (1993), 97–118.

[33] S. Willard, *General topology,* Addison-Wesley, Reading, MA, 1970.

[34] W. Woess, *Random walks on infinite graphs and groups—a survey of selected topics,* Bull. London Math. Soc. 26 (1994), 1–60.

Department of Mathematics
Florida State University
Tallahassee, FL 32306-3027

bowers@math.fsu.edu