GAFA, Geom. funct. anal.
Vol. 11 (2001) 971 – 1030
1016-443X/01/050971-60 $ 1.50+0.20/0

GAFA Geometric And Functional Analysis

PERCOLATION CLUSTERS IN HYPERBOLIC TESSELLATIONS

S.P. Lalley

Abstract

It is known that, for site percolation on the Cayley graph of a co-compact Fuchsian group of genus ≥ 2, infinitely many infinite connected clusters exist almost surely for certain values of the parameter $p = P\{\text{site is open}\}$. In such cases, the set Λ of limit points at ∞ of an infinite cluster is a perfect, nowhere dense set of Lebesgue measure 0. In this paper, a variational formula for the Hausdorff dimension $\delta_H(\Lambda)$ is proved, and used to deduce that $\delta_H(\Lambda)$ is a continuous, strictly increasing function of p that converges to 0 and 1 at the lower and upper boundaries, respectively, of the coexistence phase.

1 Introduction

This paper is a sequel to [L1], in which it was shown that Bernoulli site percolation on the Cayley graph G of a co-compact Fuchsian group of genus ≥ 2 has a *coexistence phase*. In this phase, there are infinitely many infinite connected clusters, and the set Λ of limit points at ∞ of an infinite cluster is a perfect, nowhere dense set of Lebesgue measure 0. It was conjectured in [L1] that the Hausdorff dimension $\delta_H(\Lambda)$ of Λ must be strictly less than one in the coexistence phase. In this paper we prove this conjecture, and show that $\delta_H(\Lambda)$ is a continuous function of p for $p \in [0, 1]$ that is strictly increasing in p throughout the coexistence phase. The continuity theorem shows that the phase transition for percolation differs markedly for the corresponding phase transition for branching random walks and contact processes ([LS1,2]), where the Hausdorff dimension of the limit set exhibits a jump discontinuity. We also give a new and improved proof of the existence of the coexistence phase, using recent general results of Benjamini, Lyons, Peres, and Schramm [BeLPS1,2] concerning invariant percolation processes on nonamenable, homogeneous graphs; this new proof is valid

Supported by NSF grant DMS-0071970.

for *all* co-compact Fuchsian groups, not just those of genus 2 and higher. (Benjamini and Schramm [BeS] have also proved this result, by a somewhat different argument.)

The continuity and strict monotonicity of $\delta_H(\Lambda)$ are deduced from a *variational formula* for $\delta_H(\Lambda)$, which is of interest in its own right. This formula is a simple relation among three fundamental quantities — *entropy*, *dimension*, and the *connectivity function*. It is reminiscent of the Ledrappier–Young formula [LeY] in smooth ergodic theory, which gives a similar relation among entropy, dimension, and *Lyapunov exponents*, but is actually more closely related to a generalization of Hawkes' theorem [Ha] due to the author and T. Sellke [LS3].

1.1 Fuchsian groups, hyperbolic tilings, and Cayley graphs. Recall that a *Fuchsian group* is a discrete group Γ of isometries of the hyperbolic plane $\mathbb{H}$. It is *co-compact* if the quotient space $\Gamma \setminus \mathbb{H}$ is compact, equivalently, if Γ has a compact fundamental polygon $\mathcal{R}$. The fundamental polygon $\mathcal{R}$ may be chosen so that

(a) it is closed and geodesically convex;
(b) its boundary consists of finitely many geodesic arcs;
(c) its images $\{x\mathcal{R}\}_{x\in\Gamma}$ cover $\mathbb{H}$: that is, $\mathbb{H} = \cup_{x\in\Gamma} x\mathcal{R}$; and
(d) distinct tiles $x\mathcal{R}, y\mathcal{R}$ intersect only in boundary arcs, single points, or not at all.

If the fundamental polygon $\mathcal{R}$ is so chosen, then there are only finitely many tiles $\gamma\mathcal{R}$ that intersect $\mathcal{R}$ in geodesic boundary arcs. The group elements x such that $x\mathcal{R} \cap \mathcal{R}$ is a geodesic arc are called the *side-pairing transformations*; the set Γ^* of side-pairing transformations generates Γ. Note that the choice of generators may depend on the choice of fundamental polygon.

The *Cayley graph* $G = G_{\Gamma,\Gamma^*}$ is the graph with vertex set Γ and edges connecting those pairs x, x_* of vertices such that $xx_*^{-1} \in \Gamma^*$ (equivalently, those pairs x, x_* such that the tiles $x\mathcal{R}, x_*\mathcal{R}$ overlap in a geodesic segment). It is possible to embed the Cayley graph G in the hyperbolic plane $\mathbb{H}$ in such a way that (i) it is Γ-invariant; (ii) each vertex $x \in \Gamma$ is identified with the point $x\omega \in x\mathcal{R}$, where ω is a fixed point in the interior of $\mathcal{R}$; and (iii) the edges of G are represented by geodesic segments connecting pairs $x\omega, x_*\omega$. We shall use the *disk model* of the hyperbolic plane $\mathbb{H}$, with boundary circle at infinity $\partial\mathbb{H}$ identified with the unit circle, and normalized so that ω is the Euclidean center of the disk. For ease of exposition, we shall sometimes identify the group elements x with the corresponding tiles

$x\mathcal{R}$ and sometimes with the points $x\omega$; we shall also sometimes identify connected clusters of vertices with the corresponding unions of tiles in $\mathbb{H}$.

See [B], [Kat], or [Leh] for expositions of the basic theory of Fuchsian groups, and see the appendix for certain more arcane facts concerning the geometry of $\mathbb{H}$ and hyperbolic tessellations.

1.2 Site percolation on Γ. Fix $p \in (0,1)$. Color each vertex $x \in \Gamma$ of G *blue* or *red*, blue with probability p and red with probability $q = 1-p$, with colors chosen independently for different tiles. For any vertex x, define the (blue) *cluster* K_x to be the maximal connected set of blue vertices in G containing x (if x is colored red, then $K_x = \emptyset$), and write $K = K_1$, where 1 is the group identity (identified with the tile $\mathcal{R}$). Denote by P_p the probability measure on configuration space $\{\text{Red, Blue}\}^\Gamma$ under which the coordinate random variables are i.i.d. Bernoulli-p.

Theorem 1. *Let Γ be a co-compact Fuchsian group with generating set Γ^*, and let $G = G_{\Gamma,\Gamma^*}$ be its Cayley graph. Then there exist constants $0 < p_c < p_u < 1$, depending on Γ, such that*

1. *For $p \le p_c$ there is no infinite blue cluster, with P_p-probability 1.*
2. *For $p \ge p_u$ there is a single infinite blue cluster, with P_p-probability 1.*
3. *For $p_c < p < p_u$ there are infinitely many infinite blue clusters, with P_p-probability 1.*

For Fuchsian groups of genus ≥ 2, and also for a certain series of triangle groups, Theorem 1 (minus the determination of what happens at p_c and p_u) was proved in [L1]. Theorem 1 in the generality stated above was proved by I. Benjamini and O. Schramm [BeS], by an argument that uses results and techniques from [BeLPS1,2]. We shall give a different proof of Theorem 1 in section 2 below. This proof incorporates certain elements of the Benjamini–Schramm proof, notably the use of the "density implies percolation" result of [BeLPS1] (see Proposition 2.6 below), but is based on a different strategy, stemming from the observation that the inequality $p_c < p_u$ follows from the relation

$$\lim_{p \to p_u +} \vartheta(p) > 0\,, \tag{1.1}$$

where $\vartheta(p) = P_p\{|K| = \infty\}$ is the *percolation probability* function. It is possible that this observation may be of use in extending Theorem 1 to other nonamenable groups.

If a blue cluster K_x is infinite, then its vertices (tiles) must accumulate at the boundary $\partial\mathbb{H}$ of hyperbolic space. (NOTE: Here the hyperbolic plane $\mathbb{H}$ and its boundary $\partial\mathbb{H}$ are identified with the unit disk $\mathbb{D}$ and the unit circle S^1 in the (Euclidean) plane, and "accumulation" is relative to the usual Euclidean topology on $\mathbb{D} \cup S^1$.) The limit set Λ_x of the cluster K_x is defined to be the set of all accumulation points of the cluster K_x in $\partial\mathbb{H}$ (for $x = 1$, we set $\Lambda = \Lambda_1$).

1.3 Behavior of the connectivity function along geodesics. The main results of this paper have to do with the geometry of the cluster K and its limit set Λ in the coexistence phase $p_c < p < p_u$. To state these results, we must introduce the *connectivity function* $\tau : \Gamma \to [0, 1]$. This is defined as follows:

$$\tau(x) = \tau(x; p) = P_p\{x \in K\}\,. \tag{1.2}$$

Observe that, because multiplication by x^{-1} induces an automorphism of the Cayley graph G, the connectivity function τ is invariant by inversion, i.e. for all $x \in \Gamma$,

$$\tau(x; p) = \tau(x^{-1}; p)\,. \tag{1.3}$$

LEMMA 1.1 (Log-subadditivity inequality). *For any $x, y \in \Gamma$ let xy denote the group product of x and y. Then*

$$\tau(xy) \geq \tau(x)\tau(y)\,. \tag{1.4}$$

Proof. The event $\{xy \in K\}$ is implied by the intersection of the events $\{x \in K\}$ and $\{xy \in K_x\}$. The latter two events are *increasing*, and therefore positively correlated, by the FKG inequality [G1]; hence, the probability of their intersection is no smaller than the product of their probabilities. Since the events $\{xy \in K_x\}$ and $\{y \in K\}$ have the same probability, by the Γ-invariance of the probability measure $P = P_p$, the log-subadditivity inequality (1.4) follows. □

For any $y \in \Gamma$ denote by $|y|$ the distance from vertex 1 to vertex y in the Cayley graph G. (This metric is commonly referred to as the *word metric*, since $|y|$ is the length of the shortest word in the side-pairing transformations representing y.)

LEMMA 1.2 (Harnack inequality). *For any $x, y \in \Gamma$ and any $p \in (0, 1)$,*

$$\tau(xy; p) \geq p^{|y|}\tau(x; p) \geq p^{|xy|}\,. \tag{1.5}$$

Proof. There is a path in G of length $|y|$ from 1 to y. If all of the vertices in this path are colored blue, then there is a blue path from 1 to y. Consequently, $\tau(y; p) \geq p^{|y|}$. The Harnack inequality therefore follows from the log-subadditivity inequality. □

The log-subadditivity inequality (1.4) implies that τ decays regularly along geodesics, as we now explain. Let $\{\phi_t\}_{t\in\mathbb{R}}$ be the geodesic flow in the unit tangent bundle SM over $M = \Gamma \setminus \mathbb{H}$. Elements of SM may be identified with pairs (z, θ), where $z \in \mathcal{R}$ is a point of the fundamental polygon and θ is a direction (unit vector) based at z. The geodesic flow line $\{\phi_t(z, \theta)\}_{t \in \mathbb{R}}$ in SM may be projected to M and then lifted to $\mathbb{H}$, yielding a geodesic $\{\Phi_t(z, \theta)\}$ in $\mathbb{H}$ passing through z in the direction θ; for each $t \geq 0$, $\Phi_t(z, \theta)$ is the point in $\mathbb{H}$ at distance t from z in the direction θ, while $\Phi_{-t}(z, \theta)$ is the point at distance t from z in the direction $\pi + \theta$. Observe that $\Phi_t(z, \theta)$ is a point of $\mathbb{H}$, but that $\phi_t(z, \theta)$ is an ordered pair whose first entry is a point of $\mathcal{R}$ and whose second entry is a unit vector based at this point.

We shall need to refer to the group element $x \in \Gamma$ whose associated tile contains the point $\Phi_t(z, \theta)$. If this point lies in the interior of a tile, then it lies in *only* this tile; however, if it lies on the boundary of a tile then there is an ambiguity. We resolve this ambiguity by assigning exactly *one* geodesic boundary arc of $\mathcal{R}$ from each equivalent pair (equivalent with respect to a side-pairing transformation) to the tile $\mathcal{R}$; and, for the usual orientation of $\partial\mathcal{R}$, assign the *initial* endpoint of each geodesic boundary arc to the same tile as the arc. We then extend this assignment to the other tiles of the tessellation by requiring that the assignment of boundary arcs and points to tiles be Γ-invariant. The resulting assignment is a mapping

$$\gamma : \mathbb{H} \to \Gamma$$

such that the following properties hold:

(a) If z lies in the interior of $x\mathcal{R}$, then $\gamma(z) = x$.
(b) The mapping $\gamma : \mathbb{H} \to \Gamma$ is Γ-invariant, i.e. for all $z \in \mathbb{H}$ and all $x \in \Gamma$,

$$\gamma(xz) = x\gamma(z)\,. \tag{1.6}$$

Note that the invariance (1.6) implies that

$$\Phi_{t+s}(z, \theta) = \gamma\big(\Phi_s(z, \theta)\big)\Phi_t\big(\phi_s(z, \theta)\big)\,. \tag{1.7}$$

Proposition 1.3. *Let μ be an ergodic, invariant probability measure for the geodesic flow $\{\phi_t\}$. Then for $\mu-$almost every pair (z, θ),*

$$\lim_{t \to \infty} \tfrac{1}{t} \log \tau\big(\gamma(\Phi_t(z, \theta); p)\big) = \beta(\mu) = \beta(\mu; p) \tag{1.8}$$

where

$$\begin{aligned}\beta(\mu; p) &= \lim_{t \to \infty} \tfrac{1}{t} E_\mu \log \tau(\gamma \circ \Phi_t; p) \\ &= \sup_{t \geq 1} \tfrac{1}{t} E_\mu \log \tau(\gamma \circ \Phi_t; p)\,. \end{aligned} \tag{1.9}$$

Proof. This is a consequence of Kingman's subadditive ergodic theorem [Ki]. The subadditivity hypothesis of Kingman's theorem holds for the functional $-\log\tau$, by the log-subadditivity inequality (1.4) and the invariance property (1.7). Kingman's theorem also requires that the moment hypothesis $E_\mu \log_- \tau(\gamma\circ\Phi_1) < \infty$ be satisfied. To see that this hypothesis holds, observe that for any pair (z,θ), where z is a point of the fundamental tile $\mathcal{R}$ and θ is an angle, the point $\Phi_1(z,\theta) \in \mathbb{H}$ is at distance ≤ 1 from $\mathcal{R}$. Since the tessellation $\{x\mathcal{R}\}_{x\in\Gamma}$ is locally finite, there are only finitely many tiles (say M) at distance ≤ 1 from $\mathcal{R}$; consequently, there is a path from $\mathcal{R}$ to the tile containing $\Phi_1(z,\theta)$ that passes through at most M tiles. It now follows from the Harnack inequality (1.5) that for any (z,θ),

$$\tau\big(\gamma\circ\Phi_1(z,\theta);p\big) \geq p^M .$$

Hence, the random variable $-\log\tau(\gamma\circ\Phi_1;p)$ is bounded. This proves that the moment hypothesis is satisfied. Thus, Kingman's theorem implies that the limiting relation (1.8) holds as $t\to\infty$ through the *integers*. To see that it holds also as $t\to\infty$ through $\mathbb{R}$, note that for any integer $n\geq 0$,

$$\max_{0<s<1} \big|\log\tau(\gamma\circ\Phi_{n+s};p) - \log\tau(\gamma\circ\Phi_n;p)\big| \leq -M\log p\,,$$

by virtually the same argument used to verify the moment hypothesis. □

1.4 The decay rate function $\boldsymbol{\beta(\mu;p)}$. Observe that $\beta(\mu;p)$ may be defined by (1.9) for any invariant probability measure μ, regardless of ergodicity, since the expectation is superadditive in t. Moreover, the resulting functional $\mu\to\beta(\mu;p)$ is *affine* in μ. Various properties of this function will be proved in section 6.2 below. The main results are summarized in the following theorem, which will be proved in section 6:

Theorem 2. *The decay rate function $\beta(\mu;p)$ is jointly continuous in μ and p for $\mu\in\mathcal{I}$ and $p\in(0,1]$, and for each μ is strictly increasing in p for $p\in(p_c,p_u)$. Furthermore,*

$$\beta(\mu;p) = 0 \quad \text{for all } p\geq p_u \text{ and } \mu\in\mathcal{I}\,, \tag{1.10}$$

and, uniformly for $\mu\in\mathcal{I}$,

$$\lim_{p\to 0}\beta(\mu;p) = -\infty\,. \tag{1.11}$$

The main points of Theorem 2 are the continuity and strict increase of β. The continuity of β in the Bernoulli parameter depends on the following result, which is of interest in its own right.

Theorem 3. *The connectivity function $\tau(x;p)$ is continuous in p.*

Theorem 3 will be proved in section 4. The most interesting point is the upper critical value $p=p_u$. We shall show that the continuity of $\tau(x;p)$

at $p = p_u$ is a consequence of the uniqueness of the infinite blue cluster at $p = p_u$. For Bernoulli site percolation on certain other homogeneous graphs (see [S] and [Pe]) where $p_c < p_u$, there are *infinitely many* infinite blue clusters at $p = p_u$. For such graphs, the connectivity function $\tau(x;p)$ is necessarily *discontinuous* at $p = p_u$, by Proposition 4.2 below.

1.5 Variational formula for Hausdorff dimension. Henceforth, let $\mathcal{I}$ be the set of invariant probability measures for the geodesic flow on SM, and let $\mathcal{I}_e$ denote the subset of $\mathcal{I}$ consisting of *ergodic* invariant probability measures. For $\mu \in \mathcal{I}$, let $h(\mu)$ be the Kolmogorov–Sinai entropy of the measure-preserving system (SM, Φ_1, μ).

Theorem 4. *For every $p \in (0,1]$ the Hausdorff dimension $\delta_H(\Lambda)$ of the limit set Λ is, almost surely on the event $\Lambda \neq \emptyset$, given by*

$$\delta_H(\Lambda) = \delta(p) \stackrel{\Delta}{=} \max_{\mu \in \mathcal{I}} \big(h(\mu) + \beta(\mu;p)\big)_+ . \tag{1.12}$$

Moreover, the maximum in (1.12) *is attained at an* ergodic *invariant probability measure μ, and the maximum is positive if and only if $p > p_c$.*

The proof of Theorem 4 will be carried out in sections 7–9 below. That the maximum in (1.12) is attained follows from the upper semicontinuity of the entropy functional and the continuity of the decay rate function $\beta(\mu;p)$. That it is attained at an *ergodic* measure μ follows by the affinity of the maps $\mu \to \beta(\mu;p)$ and $\mu \to h(\mu)$ (for the latter, see [W, Theorem 8.1]). The topological entropy of the geodesic flow ϕ_t is 1, and so by the Variational Principle for entropy,

$$\max_{\mu \in \mathcal{I}} h(\mu) = 1 . \tag{1.13}$$

Thus, in view of inequality (1.10), the variational formula (1.12) implies that $\delta_H(\Lambda) = \delta(p) = 1$ for all $p \geq p_u$. For $p \leq p_c$, the blue cluster K containing $\mathcal{R}$ is almost surely finite, and so $\Lambda = \emptyset$. Hence,

$$\delta(p) = 0 \quad \text{for} \quad p \leq p_c . \tag{1.14}$$

Corollary 1.4. *$\delta(p)$ is continuous in p for $p \in (0,1]$ and strictly increasing in p for $p_c < p < p_u$.*

Proof. Fix $p_c < p < p' < p_u$. By the remarks above, the maximum in (1.12) above is attained at some $\mu \in \mathcal{I}$. By Theorem 2, $\beta(\mu;p) < \beta(\mu;p')$. Hence, by the variational formula, $\delta(p) < \delta(p')$.

To prove that $\delta(p)$ is continuous in p, it is enough to show that it is right-continuous for $p \in [p_c, p_u)$ and left-continuous for $p \in (p_c, p_u]$. For each p, let $\nu_p \in \mathcal{I}$ be such that the max in (1.12) above is attained at

$\mu = \nu_p$. Let $p_n \downarrow p \in [p_c, p_u)$; by passing to a subsequence, we may arrange that, for some $\nu \in \mathcal{I}$, the measures ν_{p_n} converge to ν weakly. By Theorem 2 and the upper semicontinuity of entropy,

$$\beta(\nu; p) = \lim_{n\to\infty} \beta(\nu_{p_n}, p_n) \quad \text{and} \tag{1.15}$$

$$h(\nu) \geq \limsup_{n\to\infty} h(\nu_{p_n}) . \tag{1.16}$$

Consequently, the variational formula implies that $\delta(p) \geq \limsup \delta(p_n)$. Since $\delta(p)$ is increasing in p, we must have $\delta(p) = \lim \delta(p_n)$. This proves that $\delta(p)$ is right continuous for $p \in [p_c, p_u)$.

Since $\delta(p)$ is increasing in p, to prove that it is left-continuous it suffices to prove that $\delta(p) \leq \lim \delta(p_n)$ for any increasing sequence $p_n \uparrow p \in (p_c, p_u]$. Let $\nu \in \mathcal{I}$ be an invariant probability measure such that the max in (1.12) is attained at $\mu = \nu$. Then $\lim \beta(\nu; p_n) = \beta(\nu; p)$, by the continuity of β, and so, by the variational formula (1.12),

$$\liminf_{n\to\infty} \delta(p_n) \geq \lim_{n\to\infty} h(\nu) + \beta(\nu; p_n) = \delta(p) . \qquad \square$$

1.6 Regularly embedded subtrees. An interesting byproduct of the proof of Theorem 4 is the fact that each infinite blue cluster contains a homogeneous rooted tree of degree 3, and that the embedding of this tree respects the natural geometry of the hyperbolic plane. A *homogeneous rooted tree* $\mathcal{T}$ of degree $d+1$ is an infinite graph with no cycles (nontrivial closed paths) whose vertices are arranged in *levels* $\{\mathcal{L}_n\}_{n\geq 0}$, in such a way that

(i) $\mathcal{L}_0$ consists of a single vertex r, the *root*;
(ii) r has exactly d neighboring vertices, all in $\mathcal{L}_1$; and
(iii) each $x \in \mathcal{L}_n$, where $n \geq 1$, has exactly $d+1$ neighbors, one in $\mathcal{L}_{n-1}$ and d in $\mathcal{L}_{n+1}$.

If $x \in \mathcal{L}_n$ and $x' \in \mathcal{L}_{n+1}$ are neighbors (that is, they are the endpoints of an edge) then x is called the *parent* of x' and x' is called an *offspring* of x. If the unique self-avoiding path from the root to vertex y passes through vertex x, then y is said to be a *descendant* of x, and x is an *ancestor* of y. Observe that the subgraph consisting of a vertex x and all its descendants is itself a homogeneous rooted tree of degree $d+1$, with root x.

DEFINITION 1.5. *Let $\mathcal{G} = (V, E)$ be any infinite graph, and let $\mathcal{T} = (\mathcal{V}, \mathcal{E})$ be a homogeneous rooted tree of degree $d+1$. Say that $\mathcal{T}$ is* regularly embedded *in $\mathcal{G}$, with embedding constant D, if there is an injective mapping $\iota : \mathcal{V} \to V$ such that*

(a) *for each parent/offspring pair* (x, y) *of* $\mathcal{T}$, *there is a path* $\alpha = \alpha_{x,y}$ *in* $\mathcal{G}$ *from* $\iota(x)$ *to* $\iota(y)$ *of length* $\leq D$; *and*

(b) *for distinct parent/offspring pairs* (x, y) *and* (x', y'), *the paths* $\alpha_{x,y}$ *and* $\alpha_{x',y'}$ *do not overlap, except at the vertex* $\iota x = \iota x'$ *if* $x = x'$.

Theorem 5. *For each* $d \geq 2$ *and each* $p > p_c$ *there exists* $D < \infty$ *such that, with* P_p*−probability one, every infinite blue cluster contains a regularly embedded, homogeneous rooted tree of degree* d, *with embedding constant* D. *Furthermore, for every nonempty open arc* $J \subset \partial\mathbb{H}$, *if a blue cluster* κ *has a limit point in* J *then there is a regularly embedded subtree of* κ *all of whose ends lie in* J.

Theorem 5 will be proved in section 3 below.

1.7 Comparison with other stochastic growth models. The *coexistence* phase for percolation corresponds to the *weak survival* phase for time-dependent stochastic particle systems such as *branching random walk* and the *contact process*. The existence of a weak survival phase for branching random walk on a nonelementary Fuchsian (or, more generally, a nonamenable) group follows (easily) from the well known fact [KV] that the spectral radius of a nontrivial random walk on a nonamenable group is strictly less than 1. The existence of a weak survival phase for the contact process is a much more delicate issue – as of this writing, it has been proved only for isotropic nearest neighbor contact processes on homogeneous trees (see [P], [Li], and [St]) – but it is conjectured to occur on other nonamenable groups.

Let us briefly discuss a particular stochastic growth model, *branching Brownian motion* on the hyperbolic plane $\mathbb{H}$, whose behavior is typical of all the other processes mentioned above. See [LS1] for an extended discussion, with proofs of the results stated below. Branching Brownian motion starts with a single particle, located at a specific point $\omega \in \mathbb{H}$ at time 0. This particle executes a Brownian motion, starting from ω, and at random exponentially distributed times, independent of the motion, undergoes *binary fission*. The offspring particles behave as their parents, executing Brownian motions from the places of their births and undergoing further binary fissions at exponentially distributed random times; furthermore, their behavior is completely independent of their parents' and other particles' behaviors, except for the locations of their births. Observe that the size of the population is a simple continuous-time Galton–Watson process, which grows (roughly) exponentially in time at rate λ, the inverse of the mean fission time. If $\lambda > 1/8$, this exponential growth rate is sufficient to guarantee the existence of particles in any open ball at arbitrarily large times, with probability 1 ("strong survival"). However, if $\lambda \leq 1/8$ then, although the size of the population grows exponentially in time, with probability 1 every compact subset of $\mathbb{H}$ is eventually vacated ("weak survival").

It is quite obvious that, in the strong survival phase, particle trajectories accumulate (relative to the Euclidean topology on $\mathbb{H} \cup \partial\mathbb{H}$) at every point of $\partial\mathbb{H}$. However, in the weak survival phase, the set Λ of such accumulation points is, with probability 1, a Cantor subset of $\partial\mathbb{H}$ of Lebesgue measure 0. The main result of [LS1] is that, for any $\lambda \leq 1/8$ the Hausdorff dimension of Λ is almost surely given by

$$\delta_H(\Lambda) = \tfrac{1}{2}\big(1 - \sqrt{1 - 8\lambda}\big)\,. \tag{1.17}$$

Observe that at the critical point λ=1/8 this expression takes the value 1/2; thus, the Hausdorff dimension exhibits a jump discontinuity at the critical point (for $\lambda > 1/8$, the limit set Λ is the entire circle $\partial\mathbb{H}$, which has Hausdorff dimension 1). In [LS2], it is shown that the same jump discontinuity in Hausdorff dimension occurs at the (upper) critical point of the isotropic contact process on a homogeneous tree; that the size of this jump is 1/2 the Hausdorff dimension of the entire geometric boundary; and that, at the critical point, the contact process survives weakly.

The main results of this paper, in particular, Corollary 1.4 and Proposition 4.2, suggest that the nature of the phase transition at the upper critical point p_u for percolation on a homogeneous graph with hyperbolic geometric is quite different from that of other stochastic growth models.

Notational conventions. We shall denote probability by P, and expectation by E, except in circumstances where dependence on the Bernoulli parameter p must be emphasized for clarity, in which cases we shall write P_p and E_p for probability and expectation, respectively. The letters

$$\mathbb{H}\,, \partial\mathbb{H}\,, \Gamma, \mathcal{R}, K, K_x, \Lambda, \Lambda_x, G, \tau, \vartheta, \gamma, \beta, \phi_t, \Phi_t, \delta(p), p_c, \text{ and } p_u$$

will have the same meanings throughout the paper as in this section.We shall write

$$f_n \asymp g_n \qquad \text{or} \qquad f(t) \asymp g(t)$$

to indicate that there exist constants $0 < C_1 \leq C_2 < \infty$ such that the ratio f_n/g_n (or the ratio $f(t)/g(t)$) is bounded above and below by C_2 and C_1, respectively.

Denote by d_H the hyperbolic metric on $\mathbb{H}$, and by d or d_W the word metric on Γ (thus, d_W is the usual metric on the Cayley graph G). When convenient, we write $|x|_W = d_W(1, x)$. For each $r > 0$, $z \in \mathbb{H}$, and $x \in \Gamma$

define

$$\begin{aligned}
&B(r;z) = \{z' \in \mathbb{H} : d_H(z,z') \le r\}\,;\\
&B_\Gamma(r;x) = \{x' \in \Gamma : d_H(x\mathcal{R}, x'\mathcal{R}) \le R\}\,;\\
&B_\Gamma(r) = B_\Gamma(r;1)\,;\\
&S(r;z) = \partial B(r;z)\,;\\
&S_\Gamma(r;x) = (\partial B_\Gamma(r;x)^c)^c \text{ and}\\
&S_\Gamma(r) = S_\Gamma(r;1)\,.
\end{aligned}\tag{1.18}$$

Observe that $B_\Gamma(r;x) = xB_\Gamma(r)$ and $S_\Gamma(r;x) = xS_\Gamma(r)$, because each $x \in \Gamma$ is an isometry; this is one of several reasons for defining $B_\Gamma(r;x)$ using the metric d_H instead of the graph (word) metric d_W. Since the group Γ is discontinuous, $|B_\Gamma(r)| < \infty$ for every r. Note that $B_\Gamma(r;x)$ is *not* a ball in the word metric d_W, since the metrics d_H and d_W are not the same; however, we shall refer to the sets $B(r;z)$ and $B_\Gamma(r;x)$ as "balls". (The metrics d_H and d_W are, in a certain sense, comparable – see Proposition 10.1 in the Appendix.) For $F \subset \Gamma$, the *boundary* ∂F of F is defined to be the set of all vertices $x \notin F$ such that x has a nearest neighbor $x'F$. The set $S_\Gamma(r)$ is a discrete approximation to the hyperbolic circle $S(r;\omega)$ of radius t centered at $\omega \in \mathcal{R}$: In particular, (a) the hyperbolic distance to $S(r;\omega)$ from any $x\mathcal{R}$, where $x \in S_\Gamma(r)$, is no more than diameter($\mathcal{R}$); and (b) the union of the tiles $x\mathcal{R}$ over $x \in S_\Gamma(r;1)$ is a connected subset of $\mathbb{H}$ that completely encircles $\mathcal{R}$.

2 The Three Phase Theorem

2.1 The auxiliary graph G^*. In the study of *bond* percolation on planar graphs, the use of the *dual* bond percolation on the dual graph is fundamentally important. For *site* percolation the natural dual process is again a site percolation, but on an augmented graph G^* obtained by adding edges to G according to the following rule: if $x, x' \in \Gamma$ are such that the tiles $x\mathcal{R}$ and $x'\mathcal{R}$ overlap, either in a geodesic arc *or in a point*, then G^* contains an edge with endpoints x and x'. (The vertices of G^* are again the group elements $x \in \Gamma$. In [L1], the graphs G and G^* were denoted by G_{blue} and G_{red}, respectively.) The importance of the graph G^* is this: A connected path of red vertices in G^* "blocks" blue paths in G, that is, no G-connected path of blue tiles can pass through a G^*-connected path of red tiles. Consequently, if there is a doubly infinite G^*-connected path of red tiles connecting two distinct points of $\partial\mathbb{H}$, then any infinite G-cluster of blue tiles must lie entirely on one side of this path. Similarly, if there is a doubly infinite G-connected path of blue tiles connecting distinct points of $\partial\mathbb{H}$, then any infinite G^*-cluster of red tiles must lie entirely on one side of this path. Henceforth, by a *blue path* we shall mean a G-connected path of blue tiles, and by a *red path* we shall mean a G^*-connected path of red tiles. Similarly, by a *blue cluster* we shall mean a set of vertices that are connected in the graph G, and by a *red cluster* a set of vertices that are connected in G^*.

2.2 The standard coupling. Study of the dependence of various probabilities on the Bernoulli parameter p is facilitated by the so-called *standard coupling.* Define P to be the product Lebesgue measure on $[0,1]^\Gamma$, that is, the unique probability measure under which the coordinate random variables $U_x, x \in \Gamma$ are i.i.d., each with the uniform-$[0,1]$ distribution. For each $x \in \Gamma$ and $p \in [0,1]$, define

$$\begin{aligned} \xi_x(p) &= \text{Blue} \qquad \text{if } \ U_x < p\,, \\ \xi_x(p) &= \text{Red} \qquad \text{if } \ U_x \geq p\,. \end{aligned} \tag{2.1}$$

Then for each p, the random configuration $\xi(p) = (\xi_x(p))_{x\in\Gamma}$ has distribution P_p, and the configurations $\xi(p)$ are monotone in p relative to the natural partial ordering on configuration space induced by the convention Blue>Red. Therefore, for any event F whose indicator function is nondecreasing with respect to this partial order, the probability $P_p(F)$ is nondecreasing in p. Henceforth, we shall refer to either the probability space $([0,1]^\Gamma, P)$ or the one-parameter family $\xi(p)_{0\leq p\leq 1}$ as the *standard coupling.*

The next three easy lemmas illustrate the use of the standard coupling. All will be needed in the sequel.

Lemma 2.1. *The percolation probability function $\vartheta(p)$ is* strictly *increasing in p for $p > p_c$.*

Proof. Since $\vartheta(p)$ is the P_p-probability of an increasing event, it is nondecreasing in p. Suppose that $p_c < p < p' < 1$. Then $\vartheta(p) > 0$, and so with positive P_p-probability there exist infinite blue clusters. Since $p < 1$, there is positive probability ε under P that vertex 1 has a nearest neighbor in an infinite blue cluster in the configuration $\xi(p)$ and that $p < U_1 < p'$. On this event, vertex 1 is included in an infinite blue cluster of $\xi(p')$, but not included in an infinite blue cluster of $\xi(p)$. It follows that $\vartheta(p') \geq \vartheta(p)+\varepsilon$. □

No easy argument based on the standard coupling can be given to show that the percolation probability $\vartheta(p)$ is continuous in p – indeed, one of the important unsolved problems of percolation is to show that $\vartheta(p_c) = 0$ for

Bernoulli percolation (site or bond) on the integer lattice $\mathbb{Z}^d$. However, the standard coupling does yield the following weaker result.

LEMMA 2.2. *The percolation probability function $\vartheta(p)$ is right-continuous in p.*

Proof. See [G1, section 6.3, Lemma 6.36]. The argument given there for bond percolation on the integer lattice $\mathbb{Z}^d$, due to Russo [R], is easily adapted to more general graphs. Following is a brief sketch: Let F_r be the event that there is a blue path from 1 to a vertex in $S_\Gamma(r)$, and define $\theta_r(p) = P_p(F_r)$. Since the event F_r involves only the colors of the vertices in $B_\Gamma(r)$, a finite set, for each $r < \infty$ the function $\theta_r(p)$ is continuous in p. Clearly, the functions $\theta_r(p)$ are nonincreasing in r, and

$$\vartheta(p) = \lim_{r\to\infty} \theta_r(p)\,.$$

Therefore, $\vartheta(p)$ is right-continuous in p. □

LEMMA 2.3. *The connectivity function $\tau(x;p)$ is strictly increasing and left-continuous in p.*

Proof. The proof that $\tau(x;p)$ is strictly increasing in p is very similar to the proof above that $\vartheta(p)$ is strictly increasing for $p > p_c$. For any $0 \le p < p' \le 1$ there is positive probability ε that in the standard coupling $p < U_x < p'$ and that there is a blue path in the configuration $\xi(p)$ from the vertex 1 to a vertex neighboring x. On this event there is a blue path from 1 to x in the configuration $\xi(p')$, but not in the configuration $\xi(p)$. Hence $\tau(x;p') \ge \tau(x;p) + \varepsilon$.

The proof that $\tau(x;p)$ is left-continuous in p, which we adapt from [G1, p. 118], is similar to the proof above that $\vartheta(p)$ is right-continuous. The difference is that whereas $\vartheta(p)$ is approximated from above by continuous nondecreasing functions, the function $\tau(x;p)$ is approximated from *below.* For each $n \ge 1$, define $\tau_n(x;p)$ to be the P_p-probability of the event $G_n(x)$ that there is a blue path from 1 to x that stays completely inside the ball $B_\Gamma(n)$. It is clear that the probabilities $\tau_n(x;p)$ are nondecreasing in both n and p. Since the event that there is a connecting blue path from $\mathcal{R}$ to $x\mathcal{R}$ is the monotone limit of the events G_n, the monotone convergence theorem implies that

$$\tau(x;p) =\uparrow \lim_{n\to\infty} \tau_n(x;p). \tag{2.2}$$

To complete the proof it suffices to show that each of the functions $\tau_n(x;p)$ is continuous in p. In the standard coupling, the occurrence or nonoccurrence of the event $G_n = G_n(p)$ for the configurations $\xi(p)$ is completely determined by the values of the random variables U_y, where $y \in B_\Gamma(n)$. Since the random variables U_y have nonatomic distribution, and since only finitely many U_y are involved, for any $p \in [0,1]$ there exists, with probability one, an $\varepsilon > 0$ (random) such that *none* of the random variables U_y takes a value between $p-\varepsilon$ and $p+\varepsilon$. If this is the case, then either $G_n(p')$ occurs for *all* $p' \in [p-\varepsilon, p+\varepsilon]$, or it occurs for *no* $p' \in [p-\varepsilon, p+\varepsilon]$. It follows that the event $G_n(p)$ is the P-a.s. limit, as $p' \to p$, of the events $G_n(p')$. Since the indicators of these events are bounded random variables, the dominated convergence theorem implies that $P(G_n(p'))$ converges to $P(G_n(p)$ as $p' \to p$; thus $\tau_n(x;p)$ is continuous in p. □

2.3 Invariant percolation processes. Recent work of Benjamini, Lyons, Peres, and Schramm [BeLPS1,2] has shown the usefulness of the notion of *invariant percolation.* Define an *invariant site percolation* on Γ to be any Γ-invariant Borel probability measure Q on the configuration space $\{\text{Red, Blue}\}^\Gamma$; define an *invariant bond percolation* similarly. Note that Bernoulli site percolation (product Bernoulli-p measure on configuration space) is a particular case of an invariant site percolation. Let Q be any Γ-invariant site percolation. The action of Γ on the probability space $(\{\text{Red,Blue}\}^\Gamma, Q)$ is said to be *ergodic* if all Γ-invariant events have Q-measure either zero or one. Note that if the action of Γ is ergodic, then all Γ-invariant random variables, such as the number of infinite blue clusters, are almost surely constant. Say that the action of Γ is *strongly ergodic* if, for each element $x \in \Gamma$ of infinite order, all x-invariant events have Q-measure either zero or one. Observe that if Γ contains elements of infinite order (as, for instance, is the case when Γ is a co-compact Fuchsian group), then strong ergodicity implies ergodicity. Observe that, for Bernoulli site percolation on a Fuchsian group Γ, the action of Γ is strongly ergodic, by the Kolmogorov 0-1 Law.

For any invariant site percolation Q, denote by p_Q the Q-probability that vertex x is blue (by invariance, this is independent of x).

Proposition 2.4 [BeLPS1]. *Let Γ be any nonamenable, finitely generated group. There exists a constant $0 < \varepsilon < 1$ such that for any invariant site percolation Q on Γ,*

$$p_Q > \varepsilon \Longrightarrow Q\{\exists \textit{ infinite blue clusters}\} = 1\,. \tag{2.3}$$

Here an *infinite blue cluster* is an infinite, connected set of vertices in the Cayley graph G of Γ. In fact, the result of [BeLPS1] holds, more generally, when the graph G is replaced by any locally finite graph with vertex set Γ

on which Γ acts transitively. Thus, in particular, for a co-compact Fuchsian group Γ the graph G may be replaced by the augmented graph G^*:

COROLLARY 2.5. *Let Γ be a co-compact Fuchsian group. Then there exist constants $0 < \varepsilon_0, \varepsilon_1 < 1$ such that for any Γ-invariant site percolation Q,*

(a) $p_Q > \varepsilon_1$ implies $Q\{\exists$ infinite blue clusters$\} = 1$, and
(b) $p_Q < \varepsilon_0$ implies $Q\{\exists$ infinite red clusters$\} = 1$.

We shall have occasion to use the result of Corollary 2.5 not only for co-compact Fuchsian groups, but also for invariant percolations on the *free group* with 2 generators, which occurs as a subgroup of any co-compact Fuchsian group. (See section 2.6 below.) For a finitely generated free group, the Cayley graph is a homogeneous tree. For invariant percolation on homogeneous trees, Corollary 2.5 was first proved by Häggström [H1], who introduced the powerful *mass transport technique* later exploited in [BeLPS1] and [BeLPS2].

2.4 The critical points $p_c < p_c^*$. Consider again the case of Bernoulli site percolation on a co-compact Fuchsian group. If the Bernoulli parameter p is less than $1/D$, where D is the vertex-degree of the Cayley graph G (the number of sides of the fundamental tile $\mathcal{R}$), then, with probability one, there are no infinite blue clusters (because the expected size of the blue cluster K is finite). Since the percolation probability $\vartheta(p)$ is nondecreasing in p, it therefore follows from Corollary 2.5 (a) that there is a critical value $0 < p_c < 1$ such that $\vartheta(p) = 0$ for $p < p_c$ and $\vartheta(p) > 0$ for $p > p_c$. Similarly, by Corollary 2.5 (b), there is a critical value $0 < p_c^* < 1$ such that the *red* cluster containing the vertex 1 is infinite with positive probability if $p < p_c^*$, but is almost surely finite if $p > p_c^*$. Benjamini *et al* [BeLPS1] determined what happens at the critical points p_c and p_c^*:

PROPOSITION 2.6 [BeLPS1]. *For Bernoulli site percolation on Γ, there exist values $0 < p_c, p_u < 1$ such that*

(a) $p > p_c \iff P_p\{\exists$ infinite blue clusters$\} > 0$
$\qquad\qquad\qquad\qquad \iff P_p\{\exists$ infinite blue clusters$\} = 1$;
(b) $p < p_c^* \iff P_p\{\exists$ infinite red clusters$\} > 0$
$\qquad\qquad\qquad\qquad \iff P_p\{\exists$ infinite red clusters$\} = 1$;
(c) $p = p_c \implies P_p\{\exists$ infinite blue clusters$\} = 0$; and
(d) $p = p_c^* \implies P_p\{\exists$ infinite red clusters$\} = 0$.

COROLLARY 2.7. *If $\lim_{p\to p_c^*+} \vartheta(p) > 0$ then $p_c^* > p_c$.*

Proof. If $\lim_{p\to p_c^*+} \vartheta(p) > 0$, then, by the right-continuity of $\vartheta(p)$, it must be that $\vartheta(p_c^*) > 0$. But part (c) of Proposition 2.6 implies that $\vartheta(p_c) = 0$, and since $\vartheta(p) = 0$ for all $p < p_c$, the only possibility is that $p_c^* > p_c$. ◻

Corollary 2.8. *If* $\lim_{p \to p_c^*+} \vartheta(p) = 0$ *then* $p_c > p_c^*$.

Proof. Choose $\varepsilon > 0$ so that $\vartheta(p) < \varepsilon_0$ for all $p \leq p_c^* + \varepsilon$, where ε_0 is as in the statement of Corollary 2.5. Let $\xi = (\xi_x)_{x \in \Gamma}$ denote a random configuration with distribution P_p. Starting from ξ, build a modified random configuration $\zeta = (\zeta_x)_{x \in \Gamma}$ as follows: If x is included in an infinite blue cluster of ξ, then set $\zeta_x =$blue; otherwise, set $\zeta_x =$red. Since the distribution P_p of the configuration ξ is invariant under the action of Γ, so is the distribution of ζ; moreover, by construction, the probability that a given vertex x is colored blue in the modified configuration ζ is just $\vartheta(p)$. If $p_c^* < p < p_c^* + \varepsilon$, then $\vartheta(p) < \varepsilon_0$, so by Proposition 2.6, with probability one the configuration ζ contains an infinite red cluster. It follows that the original configuration ξ must also contain an infinite red cluster, because every vertex x such that $\zeta_x =$red but $\xi_x =$blue is in a finite blue cluster of ξ that is completely surrounded by red. This proves that $p_c \geq p_c^* + \varepsilon$. ◻

Corollary 2.9. $p_c < p_c^*$.

Proof. By the preceding corollaries, the only other possibility is that $p_c > p_c^*$. If this were the case then for all $p \in (p_c^*, p_c)$, with P_p-probability one, in a random configuration $\xi = (\xi_x)_{x \in \Gamma}$ with distribution P_p there would be neither infinite blue clusters nor infinite red clusters. This implies that every red tile is in a finite red cluster that is completely surrounded by a finite blue cluster, and vice versa.

Fix $k \geq 1$, and define a new invariant percolation $\zeta = (\zeta_x)_{x \in \Gamma}$ as follows: For every tile $x\mathcal{R}$ that is either a member of a blue cluster of size $\leq k$, or is a member of a red cluster that is surrounded by a blue cluster of size $\leq k$, set $\zeta_x =$blue; for all other tiles, set $\zeta_x =$red. Since all clusters of ξ (red or blue) are finite and are surrounded by clusters of the opposite color, the integer k can be chosen so large that

$$P\{\zeta_x = \text{blue}\} > \varepsilon_1 \,,$$

where ε_1 is as in Corollary 2.5 above. Thus, ζ contains an infinite blue cluster, with probability one. But this implies that the original configuration ξ also contains an infinite blue cluster! This is a contradiction, and so $p_c > p_c^*$ must be impossible. ◻

2.5 Sector percolation and the coexistence phase. We will now show that, for $p_c < p < p_c^*$ there exist, with P_p-probability one, infinitely many distinct infinite blue clusters and infinitely many distinct infinite red clusters. If $p \geq p_c^*$, then by the results of the preceding section there are no infinite red clusters, and so there can be only one infinite blue cluster; similarly, if $p \leq p_c$ then there are no infinite blue clusters and only one infinite red cluster. Hence, upon proving the almost sure existence of infinitely many distinct infinite blue clusters and infinitely many distinct infinite red clusters in the regime $p_c < p < p_c^*$, we shall have completed the proof of Theorem 1, with

$$p_u = p_c^* \,. \tag{2.4}$$

Recall from [L1] the notion of red and blue *sector percolation.* This notion will also be needed for certain of the estimates in subsequent sections. Say that *i-sector percolation*, for i =red or i =blue, occurs if there exists an infinite i-path $\{\psi_n\}_{n\in\mathbb{Z}}$ (here i = red or blue) such that for some point $\zeta \in \partial\mathbb{H}$,

$$\lim_{n\to\infty} \psi_n = \zeta \,. \tag{2.5}$$

Proposition 2.10. *For any $p \in [0,1]$, if infinite i-clusters exist with P_p-probability one, then i-sector percolation occurs with P_p-probability one.*

The proof is given in section 2.6 below.

Corollary 2.11. *For all $p \in (p_c, p_c^*)$ there exist, with P_p-probability 1, infinitely many infinite blue clusters and infinitely many infinite red clusters.*

Proof. The argument is the same as that given in [L1] for Fuchsian groups of genus 2 or larger. Here is a brief resume: If $p_c < p < p_c^*$ then there are infinite blue clusters and infinite red clusters, and so by Proposition 2.10 both red and blue sector percolation occur, almost surely P_p. Fix i =red or i =blue, and let i^* denote the opposite color to i. If i-sector percolation occurs, then with P_p-probability one the set of points $\zeta \in \partial\mathbb{H}$ to which connected i-paths converge is dense in $\partial\mathbb{H}$, because the percolation process is Γ-invariant and each point $\zeta \in \partial\mathbb{H}$ has dense Γ-orbit (see, for example,[Leh, chapter III]). Consequently, for any two nonoverlapping nonempty arcs J_-, J_+ of the boundary circle $\partial\mathbb{H}$, there exist, with positive probability, a connected, doubly infinite i-paths that converge to points of J_- and J_+, respectively. By an easy argument based on the FKG inequality, it follows that there is a doubly infinite i-path that connects a point of J_- to a point of J_+. But on the event $C(J_-, J_+)$ that such a doubly infinite path exists, the number of distinct infinite i^*-clusters must be at least 2. It then follows from a well known theorem of Newman and Schulman [NS] (see also [L1]) that there are *infinitely many* distinct infinite i^*-clusters, with P_p-probability one.

To show that here are at least two distinct infinite i^*-clusters on the event $C(J_-, J_+)$, we appeal to the ergodic theorem. For $x \in \Gamma$, let F_x denote the event that x is in an infinite i-cluster. Then $P_p(F_x)$ is independent of x and strictly positive, since $p_c < p < p_c^*$. Consequently, by the ergodic theorem, for any $x \in \Gamma$ of infinite order,

$$\sum_{n=1}^{\infty} 1_{F_{x^n}} = \infty \qquad P_p\text{-almost surely}. \tag{2.6}$$

Since the group Γ is co-compact, it contains hyperbolic elements, each of which has distinct attractive and repulsive fixed points $\zeta_+, \zeta_- \in \partial\mathbb{H}$ lying in opposite arcs of $\partial\mathbb{H} \setminus (J_- \cup J_+)$. Applying (2.6) first for x and then for x^{-1}, one concludes that on the event $C(J_-, J_+)$ there must be, almost surely, infinite i^*-clusters on either side of the doubly infinite i-path connecting J_- to J_+. □

2.6 Free subgroups and sector percolation. In this section we shall prove Proposition 2.10, which states that if i-percolation (i =blue or red) occurs with positive probability then *i-sector* percolation occurs almost surely. This entails showing that some infinite i-cluster contains an infinite self-avoiding path that converges to a point ξ of the ideal boundary $\partial\mathbb{H}$. If the group Γ were free (of course, since it is co-compact, it cannot be free), so that its Cayley graph were a homogeneous tree, then the problem would be trivial: *every* infinite self-avoiding path would converge to a point of $\partial\mathbb{H}$. Thus, our strategy will be to replace Γ by a free group. That free groups occur as subgroups of nonelementary Fuchsian groups is well known. Following is a sufficient condition for a finitely generated subgroup of Γ to be free:

LEMMA 2.12. *Every co-compact Fuchsian group Γ contains elements a, b such that the subgroup $\mathcal{F} = \mathcal{F}(a, b)$ of Γ generated by a and b is free. If there exist pairwise disjoint, closed halfplanes H_1, H_2, H_3, H_4 in $\mathbb{H}$ whose bounding geodesics do not intersect or meet at ∞ such that*

$$a(H_1) = \mathit{closure}(\mathbb{H} \setminus H_2) \qquad \mathit{and} \tag{2.7}$$

$$b(H_3) = \mathit{closure}(\mathbb{H} \setminus H_4)\,, \tag{2.8}$$

then the subgroup $\mathcal{F}(a, b)$ is free. Furthermore, for any such a, b there exists a constant $\varepsilon > 0$ such that for any $x \in \mathcal{F}$,

$$d_H(1, x) \geq \varepsilon |x|_W \tag{2.9}$$

where $d_H(1,x)$ denotes the hyperbolic distance between ω and $x\omega$ and $|x|_W$ the length of the shortest word in the letters a, a^{-1}, b, b^{-1} that represents x.

Proof. Suppose first that a and b are two elements of G such that (2.7) and (2.8) hold. Then the open halfplanes $\{a^n(H_1^c)\}_{n\in\mathbb{Z}}$ are nested (for each n, the closure of $a^n(H_1^c)$ is contained in $a^{n+1}(H_1^c)$), and so a is a hyperbolic element of G with fixed points $\xi_-, \xi_+ \in \partial\mathbb{H}$, where ξ_+ is the common point in the boundary arcs of the halfplanes $a^n(H_1^c)$ and ξ_- is the common point in the boundary arcs of the halfplanes $a^n(H_2^c)$. In particular, the discrete subgroup generated by a is isomorphic to the integers $\mathbb{Z}$; similarly for the discrete subgroup generated by b. Moreover, a has an isometric circle contained in $\mathbb{H}\setminus(H_1\cup H_2)$, and b has an isometric circle contained in $\mathbb{H}\setminus(H_3\cup H_4)$. It follows that the subgroup of G generated by a and b is the free product of two copies of $\mathbb{Z}$ (see, for instance, [Leh, section IV.2]); thus, it is the free group on two generators. Moreover, by Theorem IV.2B of [Leh], a fundamental region for this subgroup is

$$\mathcal{R}(a,b) = \mathbb{H}\setminus(\cup_{i=1}^4 H_i)\,.$$

That G contains elements a, b satisfying (2.7) and (2.8) follows easily from co-compactness. Here is a sketch: Let $\xi_1, \xi_2, \ldots, \xi_K$ be the finitely many geodesics obtained by extending the sides of the fundamental polygon $\mathcal{R}$. Choose hyperbolic elements $g, h \in G$ whose fixed points on $\partial\mathbb{H}$ are distinct, and do not coincide with any of the endpoints of $\xi_1, \xi_2, \ldots, \xi_K$. Then for sufficiently large n, the elements $a = g^n$ and $b = h^n$ will satisfy the mapping requirements (2.7) and (2.8).

Finally, let x be a reduced word in the letters $a^\pm, b^\pm$ of length n. Then for any $\omega \in \mathcal{R}$, the geodesic segment from ω to $x\omega$ must cross n images of $\mathcal{R}(a,b)$. Since the minimal crossing distance of $\mathcal{R}(a,b)$ is positive, the relation (2.9) follows. □

We shall only want to consider the regime $p \to p_c^*-$ near p_c^* in the coexistence phase. The following lemma shows that this will suffice to show that both red and blue sector percolation occur for all $p \in (p_c, p_c^*)$.

Lemma 2.13. *If either red or blue sector percolation occurs with positive P_p-probability for* some $p \in (p_c, p_c^*)$ *then it occurs for* all $p \in (p_c, p_c^*)$ *with P_p-probability one.*

Proof. This is scavenged from several proofs in [L1]. Suppose, to be definite, that blue sector percolation occurs at $p_* \in (p_c, p_c^*)$. Since the occurrence of blue sector percolation is a tail event, it must happen P_{p_*}-almost surely, by the Kolmogorov 0-1 Law. It then follows, by the standard coupling, that blue sector percolation occurs P_p-almost surely for every $p \geq p_*$.

Fix $p \in [p_*, p_c^*)$. By the same argument as in the proof of Corollary 2.11, there exist, with P_p-probability one, (a) a doubly infinite blue path ψ connecting distinct points ζ_-, ζ_+ of $\partial\mathbb{H}$; and therefore also (b) infinite red clusters on either side of ψ. Any infinite self-avoiding path in such an infinite red cluster must converge to a point of $\partial\mathbb{H}$. The reason is as follows: In each of the two open arcs $J_1, J_2 \subset \partial\mathbb{H}$ with endpoints ζ_-, ζ_+, there exist countable dense sets $\{\zeta_-(n)\}$ and $\{\zeta_+(n)\}$ such that for each n there is a doubly infinite blue path connecting $\zeta_-(n)$ to $\zeta_+(n)$. (This is proved by appeal to the ergodic theorem, using almost precisely the same argument as in the proof of Corollary 2.11.) But the existence of all these doubly infinite blue paths prevents any infinite, self-avoiding red path from accumulating at more than a single point of $\partial\mathbb{H}$. This proves that red sector percolation occurs with P_p-probability one. Since $p \in [p_*, p_c^*)$, it follows that red sector percolation occurs at all values of $p \in [p_*, p_c^*)$; and by monotonicity, that red sector percolation occurs at all $p < p_c^*$.

To complete the argument, reverse the roles of the two colors to conclude that blue sector percolation occurs at all values of $p > p_c$. □

Proof of Proposition 2.10. We begin with some observations about the percolation process at the upper critical point $p = p_c^*$. By Proposition 2.6, there are no infinite red clusters at p_c^*, and so every red cluster is completely surrounded by a closed blue path. Consequently, for any n there is a closed blue path that surrounds the ball $B_\Gamma(n)$. Therefore, for any $\varepsilon > 0$ and any $n \in \mathbb{N}$ there exists $m > n$ such that, with probability at least $1-\varepsilon$, there is a closed blue path that surrounds the set B_n *and* is completely contained in $B_\Gamma(m)$.

By Corollary 2.9, $p_c < p_c^*$, so by Lemma 2.2, $\theta = \theta(p_c^*) > 0$. Thus, for any vertex x there is probability θ that x is contained in an infinite blue cluster. It follows that, for any $\varepsilon > 0$ there exists $n \in \mathbb{N}$ sufficiently large that, with probability in excess of $1 - \varepsilon$ at least one of the vertices $x \in B_n$ is in an infinite blue cluster. By the Γ-invariance of the percolation process, the same estimate holds for the set $B_n(y)$ of vertices at distance $\leq n$ from y, for any $y \in \Gamma$. But the FKG inequality implies that these events, for different y, are positively correlated. Therefore, for any $\varepsilon > 0$ there exists $n \in \mathbb{N}$ so large that for any pair $y, y' \in \Gamma$, there is probability at least $1 - \varepsilon$ that two vertices $x \in B_\Gamma(n; y)$ and $x' \in B_\Gamma(n; y')$ are connected by a blue path.

Fix $\varepsilon > 0$, and let n be sufficiently large that any two balls $B_\Gamma(n; y), B_\Gamma(n; y')$ in Γ are connected by a blue path with probability at least $1-\varepsilon$. Choose $m > n$ so large that, with probability at least $1-\varepsilon$, there is a closed blue path that surrounds the set $B_\Gamma(n; y)$ *and* is completely contained in $B_\Gamma(m; y)$. Now choose $a, b \in \Gamma$ so that the subgroup $\mathcal{F} = \mathcal{F}(a, b)$ of Γ generated by a, b is free, and so that the distance in Γ between any two of the elements $1, a, b, a^{-1}, b^{-1}$ is at least $3m$. Fix $k \geq 1$, and for any $x \in \mathcal{F}$ define $G_n(x)$ to be the event that

(a) $\exists$ blue paths connecting $B_\Gamma(n; x)$ to each of the balls $B_\Gamma(n; xz)$, where $z \in \{a, b, a^{-1}, b^{-1}\}$;
(b) each of these four paths lies completely in $B_\Gamma(k; x)$; and
(c) for each $y = xz$, where $z \in \{a, b, a^{-1}, b^{-1}\}$, there is a closed blue path that surrounds the set $B_\Gamma(n; y)$ *and* is completely contained in $B_\Gamma(m; y)$.

If k is chosen sufficiently large, then

$$P_{p_c^*}(G_n(x)) > 1 - 16\varepsilon\,. \tag{2.10}$$

Since the event $G_n(x)$ is completely determined by the colors of the vertices in a finite set, its probability varies continuously in p. Thus, there exists $p < p_c^*$ such that

$$P_p(G_n(x)) > 1 - 32\varepsilon\,. \tag{2.11}$$

Consider the set $\mathcal{S}$ of all vertices $x \in \mathcal{F}$ such that the event $G_n(x)$ occurs. This set determines an $\mathcal{F}$-invariant percolation on the free group $\mathcal{F}$. If $\varepsilon > 0$ is sufficiently small, then by Proposition 2.6, there will be infinite clusters in $\mathcal{S}$. By construction, any infinite self-avoiding infinite path in $\mathcal{S}$ will correspond to an infinite self-avoiding blue path in G that converges to a point of $\partial\mathbb{H}$. □

Having completed the proof of Theorem 1, we shall henceforth write p_u for the upper critical probability p_c^*.

2.7 Sector percolation and barriers. Proposition 2.10, which asserts that the existence of infinite clusters (blue or red) implies sector percolation, is the basis for certain technical estimates on the connectivity function that we shall need for the main results of the paper. Sector percolation – the occurrence of connected monochromatic paths converging to points of $\partial\mathbb{H}$ – creates obstructions to connected clusters of the opposite color. Any region which may completely contain such a path is a *barrier* (to clusters of the opposite color). We shall discuss two kinds of barriers: *halfplanes* and *strips.*

2.7.1 Halfplane barriers. Let Π be a hyperbolic halfplane that contains the fundamental polygon $\mathcal{R}$; define

$$F_\Pi = \{\exists \text{ infinite, self-avoiding blue path in } \Pi \text{ containing tile } \mathcal{R}\}\,, \quad (2.12)$$

$$F_\Pi^* = \{\exists \text{ infinite, self-avoiding red path in } \Pi \text{ containing tile } \mathcal{R}\}\,. \quad (2.13)$$

Note that since the event F_Π implies that there is an infinite blue cluster containing the tile $\mathcal{R}$, its probability is bounded above by the percolation probability $\vartheta(p)$. Consequently, by Lemma 2.2, $P_p(F_\Pi) \to 0$ as $p \to p_c+$. Similarly, $P_p(F_\Pi^*) \to 0$ as $p \to p_u-$.

LEMMA 2.14. *There exist monotone, nonnegative functions $\varepsilon(p), \varepsilon^*(p)$ of $p \in [0,1]$ such that*

$$\varepsilon(p) > 0 \qquad \textit{for } \; p > p_c \quad \textit{ and} \quad (2.14)$$

$$\varepsilon^*(p) > 0 \qquad \textit{for } \; p < p_u\,; \quad (2.15)$$

and such that for every hyperbolic halfplane Π that contains the fundamental polygon $\mathcal{R}$,

$$P_p(F_\Pi) \geq \varepsilon(p) \qquad \textit{and} \quad (2.16)$$

$$P_p(F_\Pi^*) \geq \varepsilon^*(p)\,. \quad (2.17)$$

Proof. Monotonicity of the bounds in p is obvious by a comparison argument based on the standard coupling. Consider the case i =blue; the case i =red is similar. For each boundary arc s of $\mathcal{R}$, denote by Π_s the halfplane containing $\mathcal{R}$ that is bounded by the geodesic obtained by extending s.

CLAIM. *To prove* (2.16)*, it suffices to prove that $P(F_\Pi) > 0$ for each of the halfplanes $\Pi = \Pi_s$.*

Proof of Claim. By Proposition 10.2 of the Appendix, there is a constant $C < \infty$ with the following property: If Π is any halfplane containing $\mathcal{R}$ in its interior, then there exists $x \in \Gamma$ such that

(i) for some boundary arc s of $\mathcal{R}$, the halfplane $x\Pi_s$ is contained in Π; and
(ii) there is a connected path of no more than C tiles, lying entirely in Π, from $\mathcal{R}$ to $x\mathcal{R}$.

By the FKG inequality, the event that the path provided by (ii) is a *blue* path is positively correlated with the event that there is an infinite blue path in $x\Pi_s$ beginning at $x\mathcal{R}$ and converging to a point of $\partial\mathbb{H}$, and by invariance, the probability of the latter event is $P(F_{\Pi_s})$. Thus,

$$P(F_\Pi) \geq p^C \min_s P(F_{\Pi_s})\,. \qquad \square$$

Fix a boundary arc s of $\mathcal{R}$; we must show that $P(F_{\Pi_s}) > 0$. We have shown in Proposition 2.10 that for any $p > p_c$, blue sector percolation occurs with positive probability. Consequently, there is a halfplane Π that contains, with positive probability, an infinite self-avoiding blue path. Now for *any* halfplane Π' there is some $x \in \Gamma$ such that $x\Pi \subset \Pi'$ (this follows from the fact that the group Γ is "nonelementary"). Hence, by the Γ-invariance of the percolation structure, the probability that Π' contains an infinite self-avoiding blue path is positive; and in particular, this is true for the halfplane $\Pi' = \Pi_s$. Since Π_s contains only countably many tiles $y\mathcal{R}$, for at least one of them, say $y^*\mathcal{R}$, the event $G(y^*)$ that $y^*\mathcal{R}$ lies in an infinite self-avoiding blue path contained in Π_s has positive probability. But the event that there is a finite blue path contained in Π_s connecting $\mathcal{R}$ to $y^*\mathcal{R}$ also has positive probability, and, by FKG, is positively correlated with the event $G(y^*)$; thus, the intersection of these two events has positive probability. This intersection contains the event F_{Π_s}. □

2.7.2 Strip barriers. A *strip* is a (nonempty) subset of $\mathbb{H}$ whose boundary consists of two nonintersecting geodesics that do not have a common endpoint on $\partial\mathbb{H}$. Observe that any strip R is bounded at ∞ by two nonintersecting closed arcs A, B of $\partial\mathbb{H}$; call these the *boundary arcs* of the strip. Call a strip with boundary arcs A, B a *barrier* (or *strip barrier*) if it (completely) contains a connected, doubly infinite path of tiles connecting the boundary arcs A and B. For any barrier R with boundary arcs A and B, define

$$F_R = \{\exists \text{ doubly infinite blue path in } R \text{ connecting } A \text{ and } B\}\,, \tag{2.18}$$

$$F_R^* = \{\exists \text{ doubly infinite red path in } R \text{ connecting } A \text{ and } B\}\,. \tag{2.19}$$

LEMMA 2.15. *There exist monotone, nonnegative functions $\varepsilon(p)$ and $\varepsilon^*(p)$ such that*

$$\varepsilon(p) > 0 \qquad \text{for } \ p > p_c \quad \text{ and} \tag{2.20}$$

$$\varepsilon^*(p) > 0 \qquad \text{for } \ p < p_u\,; \tag{2.21}$$

and such that for each stripbarrier R,

$$P_p(F_R) > \varepsilon(p)\,, \qquad \text{and} \tag{2.22}$$

$$P_p(F_R^*) > \varepsilon^*(p)\,. \tag{2.23}$$

Proof. Any strip R is, by definition, bounded by two nonintersecting geodesics α, α'. Let β, β' be the two geodesics contained in R whose four endpoints in $\partial\mathbb{H}$ are the same four points as the endpoints of α, α' (thus, the four geodesics $\alpha, \alpha', \beta, \beta'$ bound a hyperbolic quadrilateral contained in R); call β and β' the *complementary geodesics*. If the strip R is a *barrier*, then at its "thinnest" it must be wide enough so that an entire tile fits completely inside. Consequently, there is a constant $C > 0$ such that for any barrier R, the bounding geodesics α, α' are at distance C or greater. It follows that there is a constant $C' < \infty$ such that for every barrier R, the complementary geodesics β and β' are at distance less than C'.

Fix a barrier R with boundary arcs A, A', and let Π, Π' be the non-intersecting halfplanes bounded by the complementary geodesics of R, labelled so that their boundary arcs are A, A', respectively. Since Π, Π' are at distance less than C', there exists tiles $x\mathcal{R}, x'\mathcal{R}$ contained in Π, Π, respectively, that are connected by a path of tiles in R of length $n = n(C')$. By Lemma 2.15 above, there is probability at least $\varepsilon^*(p)$ that Π (respectively Π') contains a red path starting at $x\mathcal{R}$ (respectively, $x'\mathcal{R}$) and converging to A (respectively, A'). Moreover, there is positive probability at least $(1-p)^n$ that all the tiles on the path from $x\mathcal{R}$ to $x'\mathcal{R}$ are colored red. Since all of the events mentioned above are positively correlated, by FKG, it follows that their intersection has positive probability. But on the intersection, the event F_R^* must occur. A similar argument works for the event F_R. □

3 Regularly Embedded Trees

The argument used in the proof of Proposition 2.10 shows that, when the Bernoulli parameter p is near p_u, there are infinite blue clusters that intersect certain free subgroups of Γ in infinitely many vertices. Unfortunately, this argument cannot be adapted to prove that there are infinite blue clusters that contain entire trees. In this section, we show, by a different argument, that there are regularly embedded, homogeneous rooted trees in *all* infinite blue clusters. It will be clear that the argument may be adapted to show that such trees are also contained in all infinite *red* clusters.

3.1 Expected growth of clusters I. Recall that K is the blue cluster containing vertex 1. For each $t > 0$, define

$$K(t) = K \cap S_\Gamma(t)\,;$$
$$Y(t) = Y_t = |K(t)|\,;$$
$$K^*(t) = \big\{x \in K(t)\ : \exists \text{ blue path contained in } B_\Gamma(t) \text{ from 1 to } x\big\}\,;$$
$$Y^*(t) = |K^*(t)|; \text{ and}$$
$$\delta^*(p) = \limsup_{t\to\infty} (1/t) \log EY_t\,. \tag{3.1}$$

LEMMA 3.1. *Almost surely on the event* $|K| = \infty$,

$$\lim_{t\to\infty} Y^*(t) = \infty\,. \tag{3.2}$$

Proof. If each tile not in $B_\Gamma(t)$ that borders a tile in $K^*(t)$ is colored red, then K is entirely contained in $B_\Gamma(t)$ and therefore finite. Since the random set $K^*(t)$ is completely determined by the colors of the tiles in $B_\Gamma(t)$, it follows that the conditional probability, given $K^*(t)$, that $|K| < \infty$ is at least

$$(1-p)^{Y^*(t)}.$$

Consequently, for any $\rho > 0$ and $C < \infty$, on the event that $Y^*(n\rho) \le C$ for infinitely many $n \in \mathbb{N}$, it is almost surely the case that $|K| < \infty$, by Lévy 's Borel–Cantelli lemma. This proves that on the event $|K| = \infty$ it is almost surely the case that $Y^*(n\rho) \to \infty$, for any $\rho > 0$. If $\rho > 0$ is sufficiently small then this implies that $Y^*(t) \to \infty$ as $t \to \infty$ through $\mathbb{R}$, by a simple geometric argument. □

COROLLARY 3.2. $p > p_c \Longrightarrow \lim_{t\to\infty} E_p Y^*(t) = \infty$.

Proof. If $p > p_c$ then the event $|K| = \infty$ has positive probability, so the result follows from Lemma 3.1, by the monotone convergence theorem. □

PROPOSITION 3.3. *For every* $p > p_c$,

$$\delta^*(p) > 0 \qquad \textit{and} \tag{3.3}$$

$$\delta^*(p) = \delta(p)\,, \tag{3.4}$$

where $\delta(p)$ *is defined by* (1.12).

The first assertion $\delta^*(p) > 0$ will be proved in section 3.4 below, using Corollary 3.2. The equality (3.4) will be proved in sections 8–9. For the purpose of constructing regularly embedded trees in infinite blue clusters, only the positivity of $\delta^*(p)$ is needed; the equality (3.4) will be needed for proving the variational formula (1.12).

3.2 Embedded Galton–Watson trees. To prove the existence of regularly embedded *homogeneous* trees in percolation clusters we will show that percolation clusters contain embedded *Galton–Watson* trees. The construction of embedded Galton–Watson trees will also be a key step in proving the variational formula for the Hausdorff dimension of the limit sets. A *Galton–Watson* tree is the genealogical tree $\mathcal{T}$ of a Galton–Watson process; thus, $\mathcal{T}$ is a random rooted tree whose vertices are arranged in *levels*:

$$V = \bigcup_{n=0}^{\infty} V_n\,, \tag{3.5}$$

where the vertices $v \in V_n$ represent the individuals of the nth generation, and V_0 consists of a single vertex v_*. Directed edges connect those vertices $v \in V_n$ and $v' \in V_{n+1}$ such that v' is an offspring of v. If $v \in V_n$ and $v' \in V_{n+m}$ are vertices such that the unique self-avoiding path from the root to v' passes through v, then v' is said to be a *descendant* of v.

If $\mathcal{T}$ is a Galton–Watson tree with vertex levels V_n, then for each integer $L \geq 1$ one may define a derived Galton–Watson tree $\mathcal{T}_L$ with vertex levels $V_0, V_L, V_{2L}, \dots$ and edges between those pairs of vertices $v \in V_{nL}$ and $v' \in V_{nL+L}$ for which v' is a *descendant* of v in the original tree $\mathcal{T}$. The offspring distribution of the derived Galton–Watson tree $\mathcal{T}_L$ is the L-fold convolution of the offspring distribution of $\mathcal{T}$; thus, in particular, the mean offspring number for $\mathcal{T}_L$ is the Lth power $(E|V_1|)^L$ of the mean offspring number of $\mathcal{T}$. The following theorem was proved in [LS2].

PROPOSITION 3.4. *Let $\mathcal{T}$ be a supercritical Galton–Watson tree with mean offspring number $\mu = E|V_1| > 1$. Then for each $1 < \lambda < \mu$ there exists an integer $L \geq 1$ sufficiently large that, with positive probability, the derived Galton–Watson tree $\mathcal{T}_L$ contains a homogeneous rooted subtree of degree $\lceil \lambda^L \rceil + 1$.*

3.3 Auxiliary trees $\mathbb{T}_t^\sigma$. To construct Galton–Watson trees in blue clusters, we shall first identify certain auxiliary homogeneous rooted trees in the tessellation $\{x\mathcal{R}\}_{x\in\Gamma}$. Let Σ be the (finite) set of geodesic segments that bound the fundamental tile $\mathcal{R}$. For each $\sigma \in \Sigma$, let H^σ be the hyperbolic halfplane that does *not* contain $\mathcal{R}$ whose bounding geodesic is the infinite geodesic σ^* obtained by extending σ. Observe that for each $x \in \Gamma$, the halfplane xH^σ borders the tile $x\mathcal{R}$ along the edge $x\sigma$. Define

$$H_\Gamma^\sigma = \{x \in \Gamma : x\mathcal{R} \subset H^\sigma\}\,. \tag{3.6}$$

For any $x \in \Gamma$, we shall call xH_Γ^σ the *type-σ exterior halfplane* at x. If $x, x' \in \Gamma$ are such that $x \in x'S_\Gamma(t)$, say that x is *σ-admissible* relative to $x'B_\Gamma(t)$ if

$$x'B_\Gamma(t) \cap xH_\Gamma^\sigma = \emptyset \ \text{ and} \tag{3.7}$$

$$xH_\Gamma^\sigma \subset x'H_\Gamma^\sigma\,. \tag{3.8}$$

LEMMA 3.5 (Spacing lemma). *There exists $C \in (0,\infty)$ with the following property. For all sufficiently large t and every $\sigma \in \Sigma$, there exists a C-dense subset $S_\Gamma^\sigma(t)$ of $S_\Gamma(t) \cap H_\Gamma^\sigma$ such that*

(i) every $x \in S_\Gamma^\sigma(t)$ is σ-admissible relative to $B_\Gamma(t)$; and

(ii) $\forall\ x, x' \in S_\Gamma^\sigma(t)$ distinct, the halfplanes xH^σ and $x'H^\sigma$ are properly separated.

See the Appendix for the proof. Two halfplanes H, H' are said to be *properly separated* if the distance between them is at least twice the diameter of the tile $\mathcal{R}$; thus, they intersect no common tiles of the tessellation. If $F \subset G \subset \Gamma$ then F is said to be C-dense in G if for every $x \in G$ there is an $x' \in F$ such that the distance between the tiles $x\mathcal{R}$ and $x'\mathcal{R}$ is no larger than C.

Now fix $\sigma \in \Sigma$ and $t > 0$ large (at least twice the diameter of the tile $\mathcal{R}$), and let $S_\Gamma^\sigma(t)$ be as in the conclusion of Lemma 3.5. Define sets $L_n = L_n^{\sigma,t} \subset \Gamma$ inductively as follows:

$$\begin{aligned} &L_0 = \{1\}\,; \\ &L_{n+1} = \cup_{x \in L_n} L(x)\,; \qquad \text{and} \\ &L(x) = \left\{xx' : x' \in S_\Gamma^\sigma(t)\right\}. \end{aligned} \tag{3.9}$$

For each vertex x that occurs in $\mathcal{L} = \cup_n L_n$, place edges between x and all of the vertices of $L(x)$, and let $\mathcal{E} = \mathcal{E}_t^\sigma$ be the set of all edges so defined. Observe that, by (3.8), if $x' \in L(x)$ then $x'H_\Gamma^\sigma \subset xH_\Gamma^\sigma$; and by assertion (b) of the lemma, if $x', x'' \in L(x)$ are distinct, then $x'H_\Gamma^\sigma$ and $x''H_\Gamma^\sigma$ do not intersect. Thus,

$$\mathbb{T} = \mathbb{T}_t^\sigma = (\mathcal{L}, \mathcal{E})$$

is a homogeneous tree that is regularly embedded in G. Observe that the exterior halfplanes xH^σ are nested: if x is a descendant of x' in $\mathbb{T}_t^\sigma$, then $xH^\sigma \subset x'H^\sigma$; but if neither x is a descendant of x' nor is x' a descendant of x then $xH^\sigma \cap x'H^\sigma = \emptyset$.

The next lemma shows that the trees $\mathbb{T}_t^\sigma$ are regularly embedded in G.

LEMMA 3.6. *There exists a constant $C < \infty$, independent of t, such that for all sufficiently large t and every $\sigma \in \Sigma$, the tree $\mathbb{T}_t^\sigma$ has the following property. For each vertex $x \in L_n$ and each vertex $x' \in L_{n+m}$ that is a descendant of x,*

$$mt - mC \le d_H(x\mathcal{R}, x'\mathcal{R}) \le mt\,. \tag{3.10}$$

Proof. If $xx' \in L(x)$ then $x' \in \Gamma(t)$, so the distance between the tiles $x\mathcal{R}$ and $xx'\mathcal{R}$ is no larger than t. Consequently, by the triangle inequality and induction on m, if $x' \in L_{n+m}$ is a descendant vertex of $x \in L_n$ then the distance between the tiles $x\mathcal{R}$ and $x'\mathcal{R}$ is no larger than mt.

Let $x \in L_1$; then $x \in S_\Gamma^\sigma(t)$, so, by (3.7)–(3.8), the exterior halfplane xH_Γ^σ is contained in H_Γ^σ but lies outside the ball $B_\Gamma(t)$. It follows, by an elementary geometric argument, that there is a constant $C < \infty$, independent of t, with the following property: Every geodesic ray starting at a point $\omega \in \mathcal{R}$ that enters the halfplane xH^σ passes within distance C of the tile $x\mathcal{R}$. Since the exterior halfplanes xH_Γ^σ are nested, it follows that if $x \in L_m$ then the shortest geodesic segment connecting $\mathcal{R}$ to $x\mathcal{R}$ must be of length at least $mt - 2mCt - 2mDt$, where D is the diameter of $\mathcal{R}$. The result now follows (with a new constant $C' = 2C + 2D$), because each $x \in \Gamma$ is an isometry of $\mathbb{H}$, and each $x \in \mathbb{T} = \mathbb{T}_t^\sigma$ maps $\mathbb{T}$ isometrically onto the subtree consisting of all descendant vertices of x. □

3.4 Construction of embedded Galton–Watson trees. Next, we define a *random* subtree $\mathcal{T} = \mathcal{T}_t^\sigma$ of $\mathbb{T}_t^\sigma$. Set $V_0 = L_0 = \{1\}$ if $\mathcal{R}$ is blue; if $\mathcal{R}$ is red, then set $V_0 = \emptyset$. Define sets V_n inductively by setting

$$V_{n+1} = \bigcup_{x \in V_n} V(x)\,, \tag{3.11}$$

where $V(x)$, the set of "offspring" of x, is defined to be the set of all $x' \in L(x)$ such that there is a blue path from x to x' that lies entirely in $xB_\Gamma(t) \cap xH_\Gamma^\sigma$. The edges of $\mathcal{T}$ are those inherited from $\mathbb{T}$: for each $x \in \mathcal{V} = \cup_n V_n$, there are edges connecting x to all of its offspring.

PROPOSITION 3.7. *The random tree $\mathcal{T}$ is a Galton–Watson tree, all of whose vertices are members of the same blue cluster.*

Proof. It is clear that all vertices of $\mathcal{T}$ are in the same blue cluster, because by construction there exist blue paths connecting all parent/offspring pairs. To see that $\mathcal{T}$ is a Galton–Watson tree, observe that if x, x' are distinct elements of the nth generation V_n then $xH_\Gamma^\sigma \cap x'H_\Gamma^\sigma = \emptyset$, so the assignments of colors to tiles in the regions xH_Γ^σ and $x'H_\Gamma^\sigma$ are independent. □

PROPOSITION 3.8. *For any $\varepsilon > 0$ there exists $t < \infty$ and $\sigma \in \Sigma$ such that the Galton–Watson tree $\mathcal{T} = \mathcal{T}_t^\sigma$ has mean offspring number*

$$E_p|V_1| > \exp\left\{t(\delta^*(p) - \varepsilon)\right\}. \tag{3.12}$$

Proof. Fix $\sigma \in \Sigma$, and let σ^* be the infinite geodesic that extends the edge σ. Let ζ, ζ' be the endpoints of σ^*, and let A, A' be arcs of $\partial\mathbb{H}$ that contain ζ, ζ', respectively. By Lemma 2.15, there is positive probability that there are infinite blue paths α, α' in H^σ starting at $\mathcal{R}$ that converge to points in the arcs A, A', respectively. On the event that such paths α, α' exist, it must be the case that every $x \in K(t)$ that is contained in the part of H^σ between the paths α and α' is connected to 1 by a blue path that lies entirely in $B_\Gamma(t) \cap H^\sigma$.

Now consider the event F that for all $\sigma \in \Sigma$ there are infinite blue paths $\alpha_\sigma, \alpha'_\sigma$ in H^σ that start at $\mathcal{R}$ and converge to points in A_σ, A'_σ, respectively, where A_σ and A'_σ are short arcs that contain the endpoints ζ_σ

and ζ'_σ of σ^*. By the FKG inequality, this event has positive probability. If the arcs A_σ, A'_σ are chosen so small that there are no overlaps among any of them, then for all sufficiently large t and every $x \in K^*(t)$ there will be a blue path from 1 to x lying entirely in the intersection of $B_\Gamma(t)$ with at least one of the halfplanes H^σ. Consequently, on the event F,

$$Y_t^* \leq \sum_{\sigma\in\Sigma} Y_t^\sigma \tag{3.13}$$

where Y_t^σ is defined to be the number of vertices $x \in K^*(t)$ that are connnected to 1 by a blue path that lies entirely in $B_\Gamma(t) \cap H_\Gamma^\sigma$. It now follows routinely, since $P_p(F) > 0$, that for some $\sigma \in \Sigma$,

$$\limsup_{t\to\infty} t^{-1} \log E_p Y_t^\sigma \geq \delta^*(p)\,. \tag{3.14}$$

The size $|V_1^\sigma|$ of the first generation of the Galton–Watson tree $\mathcal{T}^{\sigma,t}$ is almost, but not quite, Y_t^σ. The discrepancy is due to the fact that in V_1^σ only vertices in $S_\Gamma^\sigma(t)$ are allowed. However, by Lemma 3.5, the set $S_\Gamma^\sigma(t)$ is C-dense in $\Gamma(t) \cap H_\Gamma^\sigma$. Consequently, by the Harnack inequality and the local finiteness of the tessellation,

$$E|V_1^\sigma| \geq \eta E Y_t^\sigma \tag{3.15}$$

for a suitable constant $\eta > 0$ independent of t. The proposition now follows. □

3.5 Proof that $\boldsymbol{\delta^*(p) > 0}$. By definition, $\delta^*(p)$ is the upper exponential growth rate of the expectations EY_t. Since $Y_t \geq Y_t^*$, to prove that $\delta^*(p) > 0$ it suffices to show that the expectations EY_t^* grow exponentially as $t \to \infty$. For this, it suffices to prove that

$$\limsup_{t\to\infty} t^{-1} \log EY_t^{**} > 0\,, \tag{3.16}$$

where Y_t^{**} is defined to be cardinality of $\cup_{s\leq t} K^*(s)$, the set of all vertices connected to 1 by a blue path that lies entirely inside $B_\Gamma(t)$.

By inequality (3.13), together with Corollary 3.2, there is at least one $\sigma \in \Sigma$ for which the expectations EY_t^σ converge to ∞ along some sequence $t_n \to \infty$, where Y_t^σ is as defined in the proof of Proposition 3.8 above. Hence, by the "thinning" inequality (3.15), there exists some t such that the first generation V_1^σ satisfies $\mu = E|V_1^\sigma| > 1$. Since $|V_n^\sigma|$ is a Galton–Watson process, this implies that

$$\lim_{n\to\infty} n^{-1} \log E|V_1^\sigma| = \log \mu > 0\,.$$

But $Y_{nt}^{**} \geq |V_n^\sigma|$, by construction of the Galton–Watson trees $\mathcal{T}_t^\sigma$, and so (3.16) follows. □

3.6 Proof of Theorem 5. By Proposition 3.8 and the fact (3.3) that $\delta^*(p) > 0$, there exist $\sigma \in \Sigma$ and $t < \infty$ large such that the Galton–Watson process associated to the tree $\mathcal{T} = \mathcal{T}_t^\sigma$ is supercritical. Proposition 3.4 then implies that, with positive probability, the blue cluster K contains a regularly embedded, homogeneous, rooted tree. Since the event that every infinite blue cluster contains a regularly embedded, homogeneous, rooted tree is a tail event, it then follows that this event has probability one. This proves Theorem 5. □

4 Continuity of the Connectivity Function $\tau(x;p)$

In this section we shall prove Theorem 3, which asserts the continuity of the connectivity function $\tau(x;p)$ in p for each x. The proof breaks into two parts: (1) continuity in the region $p \geq p_u$; and (2) continuity in the region $p < p_u$. The proof of (1) uses only uniqueness of the infinite blue cluster, while the proof of (2) depends more heavily on the planarity of the graph G (and, in particular, on the results and techniques of the preceding section, specifically, the analogue for infinite *red* clusters of Theorem 5).

4.1 Continuity in the region $p \geq p_u$. As we have shown (Lemma2.3), $\tau(x;p)$ is left-continuous in p; the hard bit is to establish right-continuity. The following result shows that $\tau(x;p)$ is right-continuous for $p \geq p_u$. The proof makes no real use of the geometry of the graph G, except insofar as it is needed to imply that at $p = p_u$ there is a unique infinite blue cluster.

PROPOSITION 4.1. *Suppose that $p_* > p_c$ is such that for all $p \geq p_*$ there exists, P_p-almost surely, a unique infinite blue cluster. Then for each x the function $\tau(x;p)$ is continuous in p at $p = p_*$.*

REMARK. Van den Berg and Keane [BerK] proved that if there exists a unique infinite cluster with P_{p_*}-probability one, then the percolation probability $\theta(p)$ is also continuous at $p = p_*$.

Proof of Proposition 4.1. It suffices to establish *right*-continuity at p_*. Recall (relation (2.2)) that $\tau(x;p)$ is the increasing limit of the functions $\tau_n(x;p)$, where $\tau_n(x;p)$ is the P_p-probability of the event $G_n(x)$ that there is a blue path from $\mathcal{R}$ to $x\mathcal{R}$ that lies completely in the ball $B_\Gamma(n)$. Recall also that each of the functions $\tau_n(x;p)$ is continuous in p. Consequently, to prove that $\tau(x;p)$ is right-continuous at $p = p_*$ it suffices to show that for each $\varepsilon > 0$ there exists $n = n_\varepsilon \in \mathbb{N}$ and $\rho > 0$ such that

$$\tau(x;p) - \tau_n(x;p) < \varepsilon \qquad \forall\, p \in [p_*, p_* + \rho]\,. \tag{4.1}$$

Consider the standard coupling $(P, \{\xi(p)\})$. In each configuration $\xi(p)$, where $p \geq p_*$, there is (with P-probability one) a unique infinite blue cluster $K(p)$. These infinite blue clusters are nested: $p < p'$ implies that $K(p) \subset K(p')$; in particular, $K(p_*)$ is contained in all the others. Furthermore,

$$K(p_*) = \bigcap_{p>p_*} K(p)\,. \tag{4.2}$$

The proof is as follows: If $x \in K(p)$ for all $p > p_*$, then for every $p > p_*$ and every $n \in \mathbb{N}$ there is a blue path of length n in $K(p)$ that starts at x. Since there are only finitely many possible such paths, for any given n, it follows that for each n there is a blue path ψ_n of length n starting at x that is contained in $\cap_{p>p_*} K(p)$. But if ψ_n is contained in $\cap_{p>p_*} K(p)$, then it is a blue path in the configuration $\xi(p_*)$, because with probability 1 none of the random variables U_y takes the value p_*. This shows that, with probability one, the size of the blue cluster K_x in the configuration $\xi(p_*)$ is at least n. Since n is arbitrary, it follows that this cluster is infinite, and so must coincide with $K(p_*)$.

Suppose that, with positive P-probability, there exist p_n decreasing to p_* such that, for each n, there exists in the configuration $\xi(p_n)$ a blue path from 1 to x that exits the circle $S_\Gamma(n)$. Then by an argument like that of the preceding section, for each m there exist blue paths ψ_m, ψ'_m of length m starting at 1 and x, respectively, that are contained in *all* the configurations $\xi(p)$, for $p > p_*$. This implies that they are contained in the configuration $\xi(p_*)$. By yet another "selection" argument, it follows that there are *infinite* blue paths in $\xi(p_*)$ starting at 1 and x. Since there is, with probability one, only one infinite blue cluster in $\xi(p_*)$, it must be the case that x and 1 are connected by a blue path in $\xi(p_*)$.

Suppose now that (4.1) is false. Then for some $\varepsilon > 0$ there exists a sequence p_n decreasing to p_* such that $\tau(x; p_n) - \tau_n(x; p_n) \geq \varepsilon$ for each $n \in \mathbb{N}$. Let F_n be the event that in the configuration $\xi(p_n)$ there exists a blue path from $\mathcal{R}$ to $x\mathcal{R}$, but that any such path must exit the circle $S_\Gamma(n)$. Then for each $n \geq 1$,

$$P(F_n) = \tau(x; p_n) - \tau_n(x; p_n) \geq \varepsilon\,.$$

Consequently, $P(F_n \text{ i.o.}) > 0$. Now on the event F_n, there is no blue path from $\mathcal{R}$ to $x\mathcal{R}$ in the configuration $\xi(p_n)$ that does not exit $S_\Gamma(n)$; since the configurations $\xi(p)$ are nested, it follows that there is no blue path from $\mathcal{R}$ to $x\mathcal{R}$ in the configuration $\xi(p_*)$ that does not exit $S_\Gamma(n)$. Hence, on the event F_n i.o., there is no blue path from $\mathcal{R}$ to $x\mathcal{R}$ in the configuration $\xi(p_*)$, since any such path would have to lie entirely inside at least one of the circles $S_\Gamma(n)$. On the other hand, the argument of the preceding section shows that on the event F_n i.o. there *is* a blue path from $\mathcal{R}$ to $x\mathcal{R}$ in the configuration $\xi(p_*)$. This is contradiction, so we may conclude that (4.1) is true. □

Proposition 4.1 has an interesting converse:

PROPOSITION 4.2. *Suppose that $p_* > p_c$ is such that for all $p > p_*$ there exists, P_p-almost surely, a unique infinite blue cluster. Suppose further that for each x the function $\tau(x;p)$ is continuous in p at $p = p_*$. Then*

$$P_{p_*}\{\exists \text{ unique infinite blue cluster}\} = 1\,. \tag{4.3}$$

The proof will use only the vertex-transitivity and local finiteness of the graph G, and therefore the result is valid for those graphs which satisfy only these weaker hypotheses. In particular, the result and proof are valid for the graphs considered by Schonmann [S] and Peres [Pe]. Since (4.3) does not hold at $p_u = p_*$ for these graphs, it follows that the connectivity function $\tau(x;p)$ is *not* continuous at $p = p_u$ for all x. On the other hand, for Bernoulli site percolation on Cayley graphs G of co-compact Fuchsian groups, there exists, almost surely, only one infinite blue cluster at $p = p_u$, as we have seen. Propositions 4.1 and 4.2 show that the nature of the phase transition at $p = p_u$ is tied up with the continuity properties of the connectivity function $\tau(x;p)$.

Proof of Proposition 4.2. This very simple argument was suggested by Olle Häggström [H2], improving on an earlier, more complicated proof by the author. Suppose that the conclusion of the proposition is not true. Then at $p = p_*$ there is positive probability that distinct infinite blue clusters exist, and so for some $x \in \Gamma$ there is positive probability that the blue clusters K and K_x are infinite and distinct. Consider the the standard coupling $(P, \{\xi(p)\})$, and let $K(p), K_x(p)$ be the blue clusters containing vertices 1 and x in the configuration $\xi(p)$. Then the event

$$F = \big\{|K(p_*)| = |K_x(p_*)| = \infty \text{ and } K(p_*) \neq K_x(p_*)\big\}$$

has positive P-probability. Since there is only one infinite blue cluster in $\xi(p)$ for any $p > p_*$, and since connected blue clusters are nondecreasing in p, it follows that on the event F the vertices x and 1 are in the *same* infinite blue cluster $K(p)$ for every $p > p_*$. Consequently, the connectivity function $\tau(x;p)$ has a discontinuity of size at least $P(F) > 0$ at $p = p_*$. □

4.2 Continuity in the region $p < p_u$. The proof of right-continuity of the connectivity function in this case, unlike that for the case $p \geq p_u$, relies heavily on the planarity of the graphs G, G^* and geometric features of the hyperbolic plane. We shall also use the following analogue of Theorem 5 for infinite *red* clusters.

PROPOSITION 4.3. *With probability one, every infinite red cluster contains a regularly embedded rooted red tree of degree $d \geq 2$. If J is a nonempty open arc of $\partial\mathbb{H}$, then every red cluster with a limit point in J contains a regularly embedded rooted red tree of degree $d \geq 2$ all of whose ends are in J.*

Proof. Essentially the same argument as used in section 3 for blue clusters applies. □

PROPOSITION 4.4. *For each x, the connectivity function $\tau(x;p)$ is continuous in p for $p < p_c$.*

Proof. As in the proof of Proposition 4.1, it suffices to prove that for each $\varepsilon > 0$ and $p < p_u$ there exist $n = n_\varepsilon \in \mathbb{N}$ and $\rho > 0$ such that

$$\tau(x;p') - \tau_n(x;p') < \varepsilon \qquad \forall\, p' \in [p, p+\rho]\,, \tag{4.4}$$

where $\tau_n(x;p')$ is the $P_{p'}$-probability that there is a blue path from 1 to x that does not exit the ball $B_\Gamma(n)$ of radius n. Suppose that (4.4) is false. Then for some $\varepsilon > 0$ there exists a sequence p_n decreasing to p such that $\tau(x;p_n) - \tau_n(x;p_n) \geq \varepsilon$ for each $n \in \mathbb{N}$. Consider the standard coupling $(P, \{\xi(p')\})$: Let F_n be the event that in the configuration $\xi(p_n)$ there exists a blue path from $\mathcal{R}$ to $x\mathcal{R}$, but that every such path must exit the circle S_n. Then for each $n \geq 1$,

$$P(F_n) = \tau(x;p_n) - \tau_n(x;p_n) \geq \varepsilon\,.$$

Consequently, $P(F_n \text{ i.o.}) > 0$. Now on the event F_n, there is no blue path from $\mathcal{R}$ to $x\mathcal{R}$ in the configuration $\xi(p_n)$ that does not exit the circle $S_\Gamma(n)$, but there *is* a blue connecting path that exits $S_\Gamma(n)$. Since the configurations $\xi(p')$ are nested, it follows that in the configuration $\xi(p)$, the vertices 1 and x are both members of infinite blue clusters, but these are distinct. We will show, using Proposition 4.3, that this impossible, contradicting the hypothesis that (4.4) is false.

Since for $p < p_u$ red sector percolation must occur with P_p-probability one, the set of $\xi \in \partial\mathbb{H}$ to which red paths converge is dense in $\partial\mathbb{H}$ (by strong ergodicity). Hence, if the infinite clusters K and K_x are distinct, then their limit sets Λ and Λ_x must be contained in nonintersecting closed arcs J, J_x of $\partial\mathbb{H}$, which are separated by nonempty open intervals I_1, I_2. Moreover, there must be an infinite red cluster C in the configuration $\xi(p)$

with limit points in both $\bar{I}_1$ and $\bar{I}_2$ (for instance, the red cluster containing the connected component of ∂K lying in the same component of $\mathbb{H} \setminus K$ as K_x). By Proposition 4.3, there are regularly embedded rooted trees $\mathcal{T}_1$ and $\mathcal{T}_2$ of C such that, for each $i = 1, 2$, all ends of $\mathcal{T}_i$ are in the interval $\bar{I}_i$.

On the event $F = (F_n \text{ i.o.})$, vertices 1 and x are in the same (infinite) blue cluster of the configuration $\xi(p_n)$. Consequently, since the configurations $\xi(p)$ and $\xi(p_n)$ coincide in any finite subset of Γ for all sufficiently large n, in at least one of the two trees $\mathcal{T}_i$ the root is "cut off" from the ends in each configuration $\xi(p_n)$. Because the trees $\mathcal{T}_i$ are regularly embedded, each parent/offspring pair are connected by a red path in C of length at most D, where D is the embedding constant; hence, the probability that a particular parent/offspring pair is cut off in configuration $\xi(p_n)$ is at most $D(p_n - p)$, and these "cutoff" events are independent for different parent/offspring pairs. Since $D(p_n - p) \to 0$ as $n \to \infty$, the event that the root of $\mathcal{T}_i$ is cut off from the ends in configuration $\xi(p_n)$ has vanishingly small probability as $n \to \infty$, by basic results in the theory of Galton–Watson processes. But this contradicts the fact that at least one of the two trees $\mathcal{T}_i$ has its root cut off from its ends in every configuration $\xi(p_n)$. □

5 Approximate Additivity of the Log-connectivity Function

The main result of this section (and the only result that will be needed for the arguments of sections 6–9) is the following proposition, which, together with the log-subadditivity lemma of section 1, shows that the log of the connectivity function $\tau(x; p)$ is nearly additive along geodesics.

PROPOSITION 5.1. *For each $\varepsilon > 0$ there exist constants $\eta = \eta(T) \to 0$ as $T \to \infty$, such that for every $p < p_u - \varepsilon$ and every integer $m \geq 1$,*

$$\tau(\gamma \circ \Phi_{mT}; p) \leq (1 + \eta)^{mT} \prod_{j=0}^{m-1} \tau(\gamma \circ \Phi_T \circ \phi_{jT}; p) \tag{5.1}$$

Observe that the inequality

$$\tau(\gamma \circ \Phi_{mT}; p) \geq \prod_{j=0}^{m-1} \tau(\gamma \circ \Phi_T \circ \phi_{jT}; p) \tag{5.2}$$

follows directly from the log-subadditivity lemma.

We shall deduce Proposition 5.1 from the following inequality.

PROPOSITION 5.2. *For all sufficiently small $\varepsilon > 0$, there exist constants $C_1, C_2 < \infty$ such that the following statement holds. For any $p \in (p_c, p_u - \varepsilon)$*

and all $x, y \in \Gamma$ such that there is a geodesic ray that passes through the tiles $\mathcal{R}, x\mathcal{R}$, and $xy\mathcal{R}$ (in this order),

$$\tau(xy;p) \leq C_1|xy|^{C_2}\tau(x;p)\tau(y;p)\,. \tag{5.3}$$

The proof of Proposition 5.2 will depend crucially on the estimates provided by Lemma 2.14, and thus on the planarity of the Cayley graph G. Before proceeding to this proof, we show how Proposition 5.1 follows from Proposition 5.2.

Proof of Proposition 5.1. Let C_1, C_2 be as in the statement of Proposition 5.2. By Proposition 10.1 of the Appendix, there exists $A < \infty$ such that $|\gamma \circ \Phi_t| \leq At$. Consider first the case where $m = 2^k$ is a power of two. By Proposition 5.2 and equation (1.7),

$$\tau(\gamma \circ \Phi_{2^kT}) \leq C_1(A2^kT)^{C_2}\tau(\gamma \circ \Phi_{2^{k-1}T})\tau(\gamma \circ \Phi_{2^{k-1}T} \circ \phi_{2^{k-1}T})\,. \tag{5.4}$$

Here we have omitted the functional dependence of τ on the parameter p, as $p < p_u - \varepsilon$ will be fixed for the remainder of the argument. Iteration of inequality (5.4) now gives

$$\begin{aligned}
&\tau(\gamma \circ \Phi_{2^kT}) \\
&\leq (C_1(AT)^{C_2})^{\sum_{j=0}^{k-1}2^j}(2^{C_2})^{k+2(k-1)+4(k-2)+\cdots+2^{k-1}}\prod_{j=0}^{2^k-1}\tau(\gamma \circ \Phi_T \circ \phi_{jT}) \\
&\qquad\qquad \leq (1+\rho(T))^{2^kT}\prod_{j=0}^{2^k-1}\tau(\gamma \circ \Phi_T \circ \phi_{jT})
\end{aligned}$$

where $\rho(T)$ is defined by

$$\log(1+\rho(T)) = \left(2\log(C_1(AT)^{C_2}) + C_2(\log 2)\sum_{j=1}^{\infty} j2^{-j}\right)\Big/T\,.$$

Observe that $\rho(T) \to 0$ as $T \to \infty$. This proves that (5.1) holds for powers of two, with $\eta(T) = \rho(T)$. To complete the proof, we must show that the gaps between successive powers of two can be filled. Choose $\eta = \eta(T)$ so that $\eta(T) \to 0$ as $T \to \infty$, and so that

$$C_1(2^{k+1}AT)^{C_2}\big(1+\rho(T)\big)^{2^kT} \leq \big(1+\eta(T)\big)^{2^kT} \qquad \forall k \geq 0\,.$$

Let $m \geq 1$ be any integer, and let k be the largest integer such that $2^k \leq m$, so that $m = 2^k + n$ for some $n \leq m/2$. Then by Propositions 5.2 and 10.1 and the result of the previous section,

$$\tau(\gamma \circ \Phi_{mT}) \leq C_1(AmT)^{C_2}\tau(\gamma \circ \Phi_{2^kT})\tau(\gamma \circ \Phi_{nT} \circ \phi_{2^kT})$$

$$\leq \Big(C_1(AmT)^{C_2}(1+\rho)^{2^kT}\prod_{j=0}^{2^k-1}\tau(\gamma\circ\Phi_T\circ\phi_{jT})\Big)\tau(\gamma\circ\Phi_{nT}\circ\phi_{2^kT})$$

$$\leq \Big((1+\eta)^{2^kT}\prod_{j=0}^{2^k-1}\tau(\gamma\circ\Phi_T\circ\phi_{jT})\Big)\tau(\gamma\circ\Phi_{nT}\circ\phi_{2^kT})\,.$$

This inequality may now be applied iteratively, each time reducing the index n by at least a factor of 2. The powers of $(1+\eta)$ that are generated in this iteration are the terms of the binary expansion of m. Thus, the last inequality of the iteration is precisely (5.1). □

Proof of Proposition 5.2. The inequality (5.3) is a manifestation of the general tendency of blue paths to follow geodesics in $\mathbb{H}$. Since, by hypothesis, the tiles $\mathcal{R}, x\mathcal{R}$, and $xy\mathcal{R}$ lie on a common geodesic, the existence of a blue path from 1 to xy necessitates (with high probability) the existence of blue paths from 1 to x' and from x' to xy, for some vertex x' near x. Thus $\tau(xy)\approx\tau(x)\tau(y)$. (Note the similarity with the behavior of the Green's function for random walk on the group Γ – see [A].) The following argument makes this line of reasoning precise.

First note that it suffices to prove the inequality (5.3) for x, y satisfying the additional hypothesis

$$\min(d_H(\omega, x\omega), d_H(x\omega, xy\omega)) > C_2\log|xy| + 2\text{diameter}(\mathcal{R})\,, \tag{5.5}$$

where ω is a distinguished point of the tile $\mathcal{R}$. For if (5.5) does *not* hold, then (say) $|x| < C'\log|xy|$ for a suitable C', by Proposition 10.1, and so

$$\tau(x)\geq p^{|x|} = |xy|^{C'\log p}.$$

But the Harnack inequality implies that

$$\begin{aligned}
&\tau(y)\geq\tau(xy)\tau(x^{-1}) && \Longrightarrow\\
&\tau(xy)\leq\tau(y)/\tau(x) && \Longrightarrow\\
&\tau(xy)\leq|xy|^{C'\log p}\tau(x)\tau(y)\,,
\end{aligned}$$

and so inequality (5.3) holds for *all* x, y, with a possibly larger value of C_2.

Assume then that (5.5) holds. Then the hyperbolic circle $S(r; x\omega)$ of radius r centered at $x\omega$ intersects the geodesic segment γ from ω to $xy\omega$ twice. Hence, the circle $S(r; x\omega)$ may be subdivided into four nonoverlapping arcs A_1, A_2, A_3, A_4 of equal lengths, in such a way that γ intersects each arc A_1 and A_3 at its midpoint. Let A_1 be the arc closer to ω and A_3 the arc closer to $(xy)\omega$. Let $\gamma^\perp$ be the geodesic through $x\omega$ that intersects the geodesic arc γ orthogonally. If there is a blue path connecting the vertices 1 and xy, say that it is a *good* path if it first intersects the arc A_1 *before* it intersects the geodesic $\gamma^\perp$, and does not intersect $\gamma^\perp$ *after* last intersecting the arc A_3. Say that a path is *bad* if it is not good. Note that the partition of paths into "good" and "bad" depends on the choice of radius R. Fix R, and define events

$$\begin{aligned} F_{xy} &= \{\exists \text{ blue path } 1 \to xy\}\,, \\ F_{xy}(x;R) &= \{\exists \text{ good blue path } 1 \to xy\}\,. \end{aligned} \tag{5.6}$$

LEMMA 5.3. *If $C > 0$ is sufficiently large then there exists $\varepsilon = \varepsilon_C < 1$ such that*

$$P\big(F_{xy} - F_{xy}(x; 2C\log|xy|)\big) \le \varepsilon^{|xy|} P(F_{xy})\,. \tag{5.7}$$

Proof. Let $R = C\log|xy|$. Consider the event that there is a *bad* blue path connecting the tile $\mathcal{R}$ to the tile $xy\mathcal{R}$. Such a path ψ must wind around outside the ball $B_\Gamma(R;x)$ along either the top side A_1^t or the bottom side A_1^b of the arc A_1, or along the top side A_3^t or the bottom side A_3^b of A_3. Since each half A_j^i has length $\asymp e^R$, there exist $\asymp e^R$ *pairwise disjoint* halfplanes Π whose boundary geodesics intersect A_j^i twice, and whose intersections $\Pi \cap B_\Gamma(R;x)$ with the ball $B_\Gamma(R;x)$ all contain tiles $z\mathcal{R}$ of the tessellation. (see Proposition 10.5 of the Appendix). Any red path in a halfplane Π that connects the tile $z\mathcal{R}$ to $\partial\mathbb{H}$ will prevent any blue path from winding around A_j^i. Hence, by Lemma 2.14, the probability that there is a bad path connecting $\mathcal{R}$ to $xy\mathcal{R}$ is no larger than

$$2\varepsilon^*(p)^{C'e^R} = 2|xy|^{CC'\log\varepsilon^*(p)}$$

for a suitable choice of C'. By the Harnack inequality, $\tau(xy;p) \ge p^{|xy|}$. Thus, if C is chosen sufficiently large, then the conditional probability that there is a bad blue path, given that a blue path connects $\mathcal{R}$ and $xy\mathcal{R}$, is exponentially small in $|xy|$. □

By the preceding lemma, the conditional probability that there is a *good* open blue path connecting the tiles $\mathcal{R}$ and $xy\mathcal{R}$, given that there is *some* blue path connecting them, is close to one. If there is a good blue path, then

(a) there is an open blue path from $\mathcal{R}$ to some tile $z_1\mathcal{R}$ touching the arc A_1 that does not enter the halfplane Π_+ containing $xy\omega$ bounded by the geodesic $\gamma^\perp$; and

(b) there is an open blue path to $xy\mathcal{R}$ from some tile $z_2\mathcal{R}$ touching the arc A_3 that does not enter the halfplane Π_- containing ω bounded by the geodesic $\gamma^\perp$.

Fix z_1, z_2 such that the tiles $z_1\mathcal{R}, z_2\mathcal{R}$ intersect the arcs A_1, A_3, respectively. Denote by $G(z_1)$ the event that there is an open blue path from $\mathcal{R}$ to

$z_1\mathcal{R}$ that does not enter the halfplane Π_+, and denote by $H(z_2)$ the event that there is an open blue path to $xy\mathcal{R}$ from $z_2\mathcal{R}$ that does not enter the halfplane Π_-. Then the events $G(z_1)$ and $H(z_2)$ are independent, since they involve nonoverlapping sets of tiles. Consequently,

$$P\big(G(z_1) \cap H(_2)\big) \le \tau(z_1;p)\tau(z_2^{-1}xy;p)\,.$$

Since the distance from $z_i\omega$ to $x\omega$ is no more than $R = 2C\log|xy|$ (see equation (5.7) above), the Harnack inequality implies that for a suitable constant $C < \infty$ (possibly different from the constant C in equation (5.7)),

$$\tau(z_1;p) \le |xy|^C\tau(x;p) \quad \text{and} \quad \tau(z_2^{-1}xy;p) \le |xy|^C\tau(y;p)\,.$$

Hence, for a suitable constant $C < \infty$

$$P\big(G(z_1) \cap H(_2)\big) \le |xy|^C\tau(x;p)\tau(y;p)\,. \tag{5.8}$$

The number of tiles $z_j\mathcal{R}$ that touch A_j is (roughly) proportional to the arc length of A_j, which is $\asymp |xy|^C$, for some constant C (different from the constant C in Lemma 5.3). Consequently, the number of possible pairs z_1, z_2 is $\asymp |xy|^C$, for yet another constant $C < \infty$. Summing over all possible choices of z_1, z_2, we conclude that the probability that there is a *good* open blue path from 1 to xy is no larger than $|xy|^C\tau(x;p)\tau(y;p)$. □

6 Properties of the Decay Rate Function

In this section we shall prove Theorem 2, which states that the decay rate function $\beta(\mu;p)$ is jointly continuous in μ and p and strictly increasing in p for $p_c < p < p_u$.

6.1 Continuity in the region $p \ge p_u$. The continuity of β in p for $p < p_u$ follows, as we will show, from that of the connectivity function τ. Continuity in the region $p \ge p_u$, on the other hand, is a simple consequence of the uniqueness of the infinite cluster at $p = p_u$, as the following arguments show.

PROPOSITION 6.1. $p \ge p_u \Longrightarrow \beta(\mu;p) = 0$.

This is an immediate consequence of the following lemma.

LEMMA 6.2. *For all $p \ge p_u$ and all $x \in \Gamma$,*

$$\tau(x;p) \ge \vartheta(p)^2 \ge \vartheta(p_u)^2 > 0\,. \tag{6.1}$$

Proof. Since $\tau(x;p)$ is nondecreasing in p, it is enough to consider the case $p = p_u$. For any vertex $x \in \Gamma$, the event $\{|K_x| = \infty\}$ has probability $\vartheta(p)$. For any two vertices x and y, the events $\{|K_x| = \infty\}$ and $\{|K_y| = \infty\}$ are positively correlated, by the FKG inequality. But for $p = p_u$, there is, P_p-almost surely, only one infinite blue cluster; consequently, for every $x \in \Gamma$, the probability that $K_x = K$ is at least $\vartheta(p_u)^2$. □

COROLLARY 6.3. *The decay rate function $\beta(\mu; p)$ is jointly continuous in μ and p at every μ and $p \geq p_u$.*

Proof. Clearly, it is enough to prove that $\lim \beta(\nu; p) = \beta(\nu; p_u) = 0$ uniformly in ν as $p \to p_u-$. By Proposition 4.1, the connectivity function $\tau(x; p)$ is continuous in p for $p \geq p_u$ and each $x \in \Gamma$; and by the preceding lemma, $\tau(x; p_u)$ is bounded below by $\vartheta(p_u)^2$. Hence, for each $t < \infty$ there exists p_t so large that if $p \geq p_t$, then for every $x \in \Gamma$ such that the tile $x\mathcal{R}$ intersects the circle $S_t(\omega)$ centered at a point $\omega \in \mathcal{R}$,

$$\tau(x; p) \geq \theta(p_u)^2/2 \stackrel{\Delta}{=} C > 0 .$$

In view of relation (1.9), this implies that for every $\nu \in \mathcal{I}$, and $p \geq p_t$,

$$\inf_{p \geq p_t} \beta(\nu; p) \geq (1/t) \log C .$$

Since t is arbitrary, it follows that $\beta(\nu; p) \to 0$ uniformly in ν as $p \to p_u-$. □

6.2 Continuity in the region $p < p_u$. Continuity in the region $p < p_u$ is considerably more difficult to prove. The arguments below are based on Proposition 5.1 and the log-subadditivity inequality. By Proposition 1.3,

$$\beta(\mu; p) = \lim_{t \to \infty} t^{-1} \beta_t(\mu; p) \qquad \text{where} \tag{6.2}$$

$$\beta_t(\mu; p) = E_\mu \log \tau(\gamma \circ \Phi_t; p) . \tag{6.3}$$

Unfortunately, the functions $\beta_t(\mu; p)$ may not be jointly continuous in μ, p (they may be discontinuous at certain degenerate invariant measures μ). To prove the continuity of β, we shall need the following estimate on the size of possible discontinuities of β_t.

LEMMA 6.4. *Let $\kappa = |\Gamma_*|$ be the maximum number of tiles of the tessellation that meet at a point. Suppose that μ_n is a sequence of invariant probability measures for the geodesic flow such that $\mu_n \to \mu$ weakly, and let $p_n \to p \in (0, 1)$. Then for each $t > 0$,*

$$\begin{aligned}\beta_t(\mu; p) + \kappa \log p &\leq \liminf_{n \to \infty} \beta_t(\mu_n; p_n) \\ &\leq \limsup_{n \to \infty} \beta_t(\mu_n; p_n) \\ &\leq \beta_t(\mu; p) - \kappa \log p .\end{aligned} \tag{6.4}$$

Proof. As in the proof of Lemma 2.3, we may build Bernoulli-p percolation processes simultaneously for all values of $p \in [0,1]$ on the same probability space, using i.i.d. uniform-$[0,1]$ random variables $\{U_x\}_{x\in\Gamma}$ to determine the blue tiles for each value of p. Now let μ_n be a sequence of invariant probability measures for the geodesic flow ϕ_t such that $\mu_n \to \mu$ weakly, for some probability measure μ, which must also be invariant for the geodesic flow. By a well known theorem of Skorohod, on some probability space are defined random vectors (z_n, θ_n) and (z,θ) with distributions μ_n and μ, respectively, such that

$$z_n \longrightarrow z \quad \text{and} \quad \theta_n \longrightarrow \theta \qquad \text{a.s.}$$

We may take this underlying probability space to be the same as that on which the uniform random variables U_x are defined, and we may assume that these random variables are jointly independent of the random vectors (z_n, θ_n). By the continuity of the exponential map, it follows that for any finite $t > 0$,

$$\Phi_t(z_n, \theta_n) \longrightarrow \Phi_t(z, \theta) \quad \text{a.s.}$$

Recall that $\Phi_t(z,\theta)$ is defined to be the endpoint of the unique geodesic segment of length t with initial point z and initial direction θ, and $\gamma : \mathbb{H} \to \Gamma$ is the mapping that assigns points $z \in \mathbb{H}$ to the group element x corresponding to the tile $x\mathcal{R}$ containing z. However, it is *not* necessarily the case that $\gamma(\Phi_t(z_n,\theta_n)) \to \gamma(\Phi_t(z,\theta))$, because it is possible that the point $\Phi_t(z,\theta)$ lies on the boundary of a tile $x\mathcal{R}$. This is the reason for the possible discontinuity of the expectation in μ. Nevertheless, in such a case, for all sufficiently large n the point $\Phi_t(z_n,\theta_n)$ must lie in one of the tiles intersecting (possibly only in one point) the tile $x\mathcal{R}$. Since the G-distance between any two vertices x, y such that the tiles $x\mathcal{R}$ and $y\mathcal{R}$ overlap in a point or geodesic arc is at most κ, it follows by the Harnack inequality that for all sufficiently large n,

$$p_n^{\kappa}\tau\big(\gamma(\Phi_t(z_n,\theta_n)); p_n\big) \le \tau\big(\gamma(\Phi_t(z,\theta)); p_n\big) \le p_n^{-\kappa}\tau\big(\gamma(\Phi_t(z_n,\theta_n)); p_n\big) .$$

The result now follows by taking logs, then expectations, using the continuity in p of τ, and applying the bounded convergence theorem. □

PROPOSITION 6.5. *The function $\beta(\mu;p)$ is jointly continuous in μ and p.*

Proof. We have already proved that $\beta(\mu;p)$ is continuous at all (μ,p) such that $p \ge p_u$. Thus, it suffices to prove continuity at (μ,p) for $p < p_u$.

By the log-subadditivity lemma, for all $t, s \ge 1$,

$$\begin{aligned} \beta_{t+s}(\mu;p) &\ge \beta_t(\mu;p) + \beta_s(\mu;p) \qquad &\Longrightarrow \\ \beta_{2^k}(\mu;p) &\ge 2^{k-m}\beta_{2^m}(\mu;p) \qquad &\Longrightarrow \end{aligned}$$

$$\beta(\mu;p) \geq 2^{-m}\beta_{2^m}(\mu;p)\,, \tag{6.5}$$

the last by relation (6.2). On the other hand, Proposition 5.1 implies that for every $\eta > 0$ there exist constants $\varepsilon = \varepsilon_m$, not depending on μ, satisfying $\varepsilon_m \to 0 \to$ as $m \to \infty$, such that for all $p < p_u - \eta$ and all μ,

$$\begin{aligned} \beta_{2^k}(\mu;p) &\leq 2^{k-m}\beta_{2^m}(\mu;p) + 2^k \log(1+\varepsilon_m) \qquad \Longrightarrow \\ \beta(\mu;p) &\leq 2^{-m}\beta_{2^m}(\mu;p) + \log(1+\varepsilon_m)\,. \end{aligned} \tag{6.6}$$

Let μ_n be invariant probability measures for the geodesic flow such that $\mu_n \to \mu$ weakly, and let $p_n \to p$ for some $p \in (0, p_u - \eta)$. Fix $\delta > 0$, and choose m so large that

$$\varepsilon_m < \delta \text{ and } -\kappa(\log p) < \delta 2^m\,, \tag{6.7}$$

where κ is the number of side-pairing transformations (see Lemma 6.4). By inequalities (6.5)–(6.6),

$$\begin{aligned} 2^{-m}\beta_{2^m}(\mu;p) \leq \beta(\mu;p) \leq 2^{-m}\beta_{2^m}(\mu;p) + \delta \qquad &\text{and} \\ 2^{-m}\beta_{2^m}(\mu_n;p_n) \leq \beta(\mu_n;p_n) \leq 2^{-m}\beta_{2^m}(\mu_n;p_n) + \delta \qquad &\forall\, n \geq 1\,. \end{aligned}$$

By Lemma 6.4 and inequalities (6.7), for all sufficiently large n,

$$2^{-m}\beta_{2^m}(\mu;p) - \delta \leq 2^{-m}\beta_{2^m}(\mu_n;p_n) \leq 2^{-m}\beta_{2^m}(\mu;p) + \delta\,.$$

It now follows that for all sufficiently large n,

$$\left|\beta(\mu_n;p_n) - \beta(\mu;p)\right| < 2\delta$$

Since $\delta > 0$ is arbitrary, this proves that β is jointly continuous at every pair (μ, p) such that $p < p_u$. □

6.3 Strict increase of the decay rate function. Strict increase of the decay rate function $\beta(\mu;p)$ in p for $p \in (p_c, p_u)$ will be established by a strategy similar to that used in [L2] to prove the strict increase in the infection rate λ of the Hausdorff dimension of the limit set of the contact process on a homogeneous tree. The key to the proof is that the number of *pivotal sites* for the event $x \in K$, where K is the blue cluster containing 1, grows linearly with the distance $|x|$. A site $z \in \Gamma$ is called *pivotal* for the event $\{y \in K_x\}$ if the event occurs and every open (blue) path from x to y passes through the site z. Define $N_{x,y}$ to be the number of pivotal sites for the event $\{y \in K_x\}$. Set $N_x = N_{1,x}$.

LEMMA 6.6. *For every $p \in (p_c, p_u)$ there exist constants $C > 0$ and $\varepsilon = \varepsilon(p) > 0$ such that for every $x \in \Gamma$,*

$$P_p\big(N_x < C|x| \mid x \in K\big) \leq (1-\varepsilon)^{|x|}\,. \tag{6.8}$$

Proof. The proof is a fairly standard application of the BK inequality ([G1], ch. 3) and Menger's theorem ([Bo, ch. 3]), together with Lemma 2.15

above and Proposition 10.4 of the Appendix. Denote by A_x the event $\{x \in K\}$ that there is an open blue path from 1 to x. Each pivotal site other than 1 or x has the property that its deletion would disconnect x from 1. Consequently, the pivotal sites x_i for the event A_x occur in the same order of occurrence in all blue paths connecting 1 to x:

$$1 = x_0 \to x_1 \to x_2 \to \cdots \to x_{N_x} = x\,.$$

By Menger's theorem, for each successive pair x_i, x_{i+1} of pivotal sites there exist *vertex-disjoint* blue paths connecting x_i to x_{i+1}.

Let γ denote the geodesic segment with endpoints $\omega \in \mathcal{R}$ and $x\omega \in x\mathcal{R}$. By Proposition 10.4 of the Appendix, for a suitable constant $C' > 0$ independent of x, there are at least $m = C'|x|$ pairwise disjoint *barriers* that cross the geodesic segment γ. These may be listed $R_1, R_2, \dots, R_m$ according to the order of their intersections with γ. For each i, define L_i to be the (minimum) number of barriers R_j that must be (completely) crossed by any path from $x_i\mathcal{R}$ to $x_{i+1}\mathcal{R}$. By the BK inequality (see [G1, secs. 2.3 and 5.2], and [G2, sec. 6.2], for a similar argument)

$$P(L_{j+1} > n \mid L_0, L_1, \dots, L_j, A_x) \\ \le P\{n \text{ successive barriers crossed by a blue path}\}\,. \tag{6.9}$$

By Lemma 2.15, this last probability is no greater than $(1-\varepsilon')^n$, for suitable $\varepsilon' = \varepsilon(p) > 0$. Hence, conditional on the event A_x, the random variables $L_1, L_2, \dots$ are stochastically dominated by i.i.d. random variables with a geometric distribution. Therefore, by Chernoff's large deviation inequality for sums of i.i.d. random variables, if $C > 0$ is sufficiently small then for some $\varepsilon > 0$,

$$P\big(L_1 + l_2 + \cdots + L_{C|x|} \ge (C' - 2C)|x| \bigm| A_x\big) \le (1-\varepsilon)^{|x|} \tag{6.10}$$

Since

$$\sum_{i=0}^{N_x} (L_i + 2) \ge C'|x|$$

on the event A_x, the lemma follows. □

PROPOSITION 6.7. *For each probability measure μ invariant under the geodesic flow, the function $\beta(\mu; p)$ is strictly increasing in p for $p \in (p_c, p_u)$.*

Proof. Fix $p_c < p' < p < p_u$, and set $r = p'/p$. Let $\{Y_x, Z_x\}_{x\in\Gamma}$ be mutually independent Bernoulli random variables indexed by the set Γ of sites, with

$$P\{Y_x = 1\} = p \quad\text{ and }\quad P\{Z_x = 1\} = r\,,$$

and set

$$Y'_x = Y_x Z_x\,.$$

Thus, the random variables Y'_x are Bernoulli-p'. Declare a site x to be p-open (respectively, p'-open) if $Y_x = 1$ (respectively, $Y'_x = 1$). For each $x \in \Gamma$, denote by K_x and K'_x the connected p-cluster and the connected p'-cluster, respectively, containing the site x, and denote by K, K' the clusters K_1, K'_1 containing the group identity 1. For any $x \in \Gamma$, let N_x be the number of pivotal sites for the event $\{x \in K\}$. Since $K' \subset K$, the event $\{x \in K'\}$ can occur only if $\{x \in K\}$ occurs *and* $Z_y = 1$ for every pivotal site y. Since the random variables Z_y are independent of the random variables Y_y, it follows from Lemma 6.6 that

$$P\{x \in K'\} \leq Er^{N_x} 1\{x \in K\} \leq r^{C|x|} + (1 - \varepsilon(p))^{|x|} P\{x \in K\}\,.$$

Consequently, for a suitable constant $0 < \rho < 1$,

$$\tau(x; p') \leq \rho^{|x|} \tau(x; p)\,. \tag{6.11}$$

In view of Proposition 10.1, this implies that for a suitable constant $C > 0$,

$$\beta(\mu; p') \leq \beta(\mu; p) + C \log \rho < \beta(\mu; p)\,. \tag{6.12}$$

□

7 Variational Formula for Hausdorff Dimension I. The Case $p \geq p_u$

As we have already observed, Lemma 6.2 implies that $\beta(\mu; p) = 0$ whenever $p \geq p_u$. Thus, $\delta(p) = 1$ for $p \geq p_u$, by equation (1.13), and so to prove the variational formula (1.12) for $p \geq p_u$ it suffices to prove that, on the event $\{|K| = \infty\}$, the limit set Λ is the entire boundary circle $\partial\mathbb{H}$.

For any probability measure μ invariant for the geodesic flow, define a measure $\hat{\mu}$ on $\partial\mathbb{H}$ as follows: Choose a point $z \in \mathcal{R}$ and a direction (unit vector) θ at random in such a way that the pair (z, θ) is distributed according to μ. Let Φ_t be the geodesic ray in $\mathbb{H}$ with initial point z and initial direction θ. The ray Φ_t converges, as $t \to \infty$, to a point $\Xi \in \partial\mathbb{H}$; this point Ξ is random, since z, θ are random. Define $\hat{\mu}$ to be the distribution of Ξ.

Now let μ be any *ergodic* invariant probability measure for the geodesic flow such that the derived measure $\hat{\mu}$ gives positive mass to every interval of $\partial\mathbb{H}$ of positive length (for instance, choose μ to be the *Liouville measure*). Choose z, θ at random according to the distribution μ, and let Φ_t be the geodesic ray in $\mathbb{H}$ with initial point z and initial direction θ. As $t \to \infty$, the ray Φ_t crosses infinitely many tiles $x\mathcal{R}$ of the tessellation. For each tile, there is probability $\vartheta(p) > 0$ that the tile is part of an infinite blue cluster. Thus, by the ergodic theorem, it is almost surely the case that Φ_t crosses *infinitely many* tiles that are part of infinite clusters. But if $p \geq p_u$ there is only one infinite cluster! Consequently, with probability one, the endpoint Ξ of the ray Φ_t is an accumulation point of the infinite cluster. Since Ξ is independent of the tile-colorings, and since by choice of μ the distribution of Ξ gives positive mass to every arc of $\partial\mathbb{H}$, it follows that, with probability one, almost every point of $\partial\mathbb{H}$ is an accumulation point of the infinite cluster. As the set Λ of all such accumulation points is closed, it must be that it is the entire circle $\partial\mathbb{H}$. □

8 Variational Formula for Hausdorff Dimension II. Upper Bound

Assume throughout this section and the next that the Bernoulli parameter p satisfies $0 < p < p_u$, and set

$$\delta = \delta(p) = \max_{\mu \in \mathcal{I}} \big(h(\mu) + \beta(\mu; p) \big) . \tag{8.1}$$

To establish (8.1) as an upper bound for the Hausdorff dimension of the limit set Λ we must exhibit efficient coverings of Λ by intervals of small (Euclidean) diameters. Recall that the hyperbolic plane is represented by the unit disk, endowed with the Poincaré metric d_H, and that the ideal boundary $\partial\mathbb{H}$ in this model is the unit circle, which is given the Euclidean metric. In this representation, the (hyperbolic) circle S_m with hyperbolic radius m (and hyperbolic perimeter $\asymp \exp\{m\}$) centered at ω is a Euclidean circle of perimeter $\to 2\pi$ as $m \to \infty$, and arcs of S_m with hyperbolic lengths $\asymp \exp\{\rho m\}$ are Euclidean arcs with lengths $\asymp \exp\{(\rho - 1)m\}$ as $m \to \infty$. Throughout this section, we denote by ω the Euclidean center of $\mathbb{H}$, and we assume that $\omega \in \mathcal{R}$.

8.1 Expected growth of clusters II. The key to constructing efficient coverings is the following proposition relating the expected growth rate of connected blue clusters to the function $\delta(p)$. Recall (section 3) that $Y_t = Y(t) = |K(t)|$, where $K(t) = K \cap \Gamma(t)$, and $\Gamma(t)$ is the set of all $x \in \Gamma$ such that $d_H(\omega', x\omega'') = t$ for some $\omega', \omega'' \in \mathcal{R}$.

PROPOSITION 8.1.

$$\limsup_{n \to \infty} \frac{\log EY_t}{t} \leq \delta(p) . \tag{8.2}$$

Proof of Proposition 8.1: Outline. We shall prove Proposition 8.1 by appealing to the *Variational Principle* for the *topological pressure* functional

([W, sec. 9.3, Theorem 9.10]) of the geodesic flow. We begin by relating the expectation EY_t to the topological pressure of an appropriate function. First, observe that

$$EY_t = \sum_{x \in \Gamma(t)} \tau(x;p)\,, \tag{8.3}$$

so the growth of the expectation $EY(t)$ with t is controlled by the asymptotic behavior of the connectivity function τ. By Proposition 5.1, there exist constants $\varepsilon = \varepsilon(T) \to 0$ as $T \to \infty$ such that

$$\begin{aligned}\tau(\gamma \circ \Phi_{mT};p) &\leq (1+\varepsilon)^m \prod_{j=0}^{m-1} \tau(\gamma \circ \Phi_T \circ \phi_{jT};p) \\ &= (1+\varepsilon)^m \exp\left\{\sum_{j=0}^{m-1} \psi_T(\phi_{jT})\right\}\end{aligned} \tag{8.4}$$

where ψ_T is defined to be the logarithm of $\tau \circ \gamma \circ \Phi_T$. For each $x \in \Gamma(t)$, there is a unique unit vector $\theta(x)$ based at ω such that the geodesic ray with initial point ω and initial direction $\theta(x)$ passes through $x\omega$. The segment of this geodesic ray from ω to $x\omega$ has length $t \pm O(1)$, by definition of $\Gamma(t)$. Hence, by equations (8.3)–(8.4) and the Harnack inequality,

$$EY_{mT} \leq C \sum_{x \in \Gamma(mT)} (1+\varepsilon)^m \exp\left\{\sum_{j=0}^{m-1} \psi_T\big(\phi_{jT}(\omega, \theta(x))\big)\right\} \tag{8.5}$$

for a constant $C < \infty$ independent of m and T. The last sum is of the same type as the sums that occur in the definition of the topological pressure of the function ψ_T – see [W, Definition 9.5]. Unfortunately, the function ψ_T is not continuous, and so the Variational Principle for topological pressure does not apply directly; however, we will show (see Lemma 8.2 below) that ψ_T may be approximated above by a continuous function in such a way that the error C in the approximation remains bounded as $T \to \infty$. We will then prove (see Lemmas 8.3–8.4) that, for some constant $L < \infty$ independent of m, the set $\{(\omega, \theta(x))\}_{x \in \Gamma(mT)}$ may be partitioned into L subsets, each of which is $(m,\rho)-$*separated* (see [W, sec. 9.2], for the definition) for a suitable $\rho > 0$. This will imply that EY_{mT} is bounded above by the sum of L terms that are asymptotically bounded above by the topological pressure of a continuous function close to ψ_T. The Variational Principle will then imply that

$$\limsup_{m\to\infty} m^{-1} \log \sum_{x\in\Gamma(mT)} \exp\left\{ \sum_{j=0}^{m-1} \psi_T\big(\phi_{jT}(\omega,\theta(x))\big)\right\} \le \sup_{\mu\in\mathcal{I}} \big(Th(\mu) + E_\mu\psi_T + C\big) . \quad (8.6)$$

(Note: (A) The term C enters because of the approximation of ψ_T by a continuous function – see Lemma 8.2 below. (B) If μ is an invariant measure for the geodesic flow $\{\phi_t\}_{t\ge 0}$ with entropy $h(\mu)$, then for each $T>0$ the entropy of the measure-preserving system (SM,μ,ϕ_T) is $Th(\mu)$ – see [W, Theorem 4.13].) Dividing equation (8.6) by T and letting $T\to\infty$ will then yield Proposition 8.1, because by the log-subadditivity inequality, for every $T<\infty$,

$$T^{-1}E_\mu\psi_T = T^{-1}E_\mu \log\tau(\gamma\circ\Phi_T;p) \le \beta(\mu;p) . \quad (8.7)$$

□

LEMMA 8.2. *There exist continuous functions $\tilde\psi_T$ and a constant $C<\infty$, independent of T, such that*

$$\tilde\psi_T - C \le \psi_T \le \tilde\psi_T . \quad (8.8)$$

Proof. This is proved by a routine smoothing argument. For any point $z\in\mathbb{H}$, define $g(z)$ to be the maximum value of $\tau(\gamma(z'))$ over all points $z'\in\mathbb{H}$ such that $d_H(z,z')\le 1$. Clearly, $\tau(\gamma(z))\le g(z)$. On the other hand, the Harnack inequality implies that there exists a constant $C<\infty$ such that $g(z')\le e^C\tau(\gamma(z))$ for all $z,z'\in\mathbb{H}$ such that $d_H(z,z')\le 1$. Define

$$\tilde\psi_T = \log\int g(z')dU_z(z') , \quad (8.9)$$

where U_z denotes the uniform distribution on the (hyperbolic) disk of radius 1 centered at z. Then $\tilde\psi_T$ is certainly continuous, and (8.8) holds by construction. □

LEMMA 8.3. *There exist constants $A<\infty$ and $\rho>0$ with the following property: If F is any subset of $\Gamma(mT)$ such that for any pair x,x' of distinct elements,*

$$d_H(x\mathcal{R},x'\mathcal{R}) \ge A \quad (8.10)$$

then the set $\{(\omega,\theta(x))\}_{x\in F}$ is (m,ρ)-separated relative to the transformation ϕ_T.

Proof. By definition of $\Gamma(t)$ and the fact that the fundamental tile $\mathcal{R}$ has finite diameter, there is a finite constant $\Delta\ge$ diameter$(\mathcal{R})$ such that for each $x\in\Gamma(mT)$,

$$d_H(x\mathcal{R},S_{mT}) < \Delta .$$

Hence, the hyperbolic distance from $x\omega$ to S_{mT} is no greater than 2Δ. The point on S_{mT} nearest to $x\omega$ is $\Phi_{mT}(\omega, \theta(x))$; consequently, the distance from $x\omega$ to $\Phi_{mT}(\omega, \theta(x))$ is no larger than 2Δ. Thus, hypothesis (8.10) implies that for any two distinct elements $x, x' \in F_m$,

$$d_H\big(\Phi_{mT}(\omega, \theta(x)), \Phi_{mT}(\omega, \theta(x'))\big) \geq A - 4\Delta \,. \tag{8.11}$$

The geodesic flow ϕ_t on SM is jointly continuous in t and the initial point, and therefore the projection/lifting Φ_t has the same property. Consequently, since SM is compact, there exist constants $0 < \rho < \rho'$ such that

$$d_{SM}((z,\theta),(z',\theta')) < \rho' \quad \Longrightarrow \quad \max_{0\leq t\leq T} d_H(\Phi_t(z,\theta), \Phi_t(z',\theta')) \leq \Delta \quad \text{and}$$
$$d_{SM}((z,\theta),(z',\theta')) < \rho \quad \Longrightarrow \quad d_{SM}(\phi_T(z,\theta), \phi_T(z',\theta')) < \rho' \,.$$

Here d_{SM} denotes distance in SM, computed in the metric induced by the Riemannian metric on M. It follows that if, for every integer $k \in [0, m)$,

$$d_{SM}\big(\phi_{kT}(z,\theta), \phi_{kT}(z',\theta')\big) < \rho \,, \quad \text{then}$$
$$d_H\big(\Phi_{mT}(z,\theta), \Phi_{mT}(z',\theta')\big) \leq \Delta \,.$$

Therefore, by (8.11), if $A > 5\Delta$ then the set $\{(\omega, \theta(x))\}_{x\in F}$ is (m, ρ)-separated. □

LEMMA 8.4. *For any $A < \infty$ there exists $L < \infty$ such that, for any m, T the set $\Gamma(mT)$ may be partitioned into L subsets F_i^m, each having the property that for any two distinct elements x, x' the inequality* (8.10) *holds.*

Proof. Denote by S_t the circle of radius t centered at ω. Divide S_{mT} into arcs J_i of equal length in such a way that the endpoints of any arc J_i are at hyperbolic distance at least $4A$ and no greater than $5A$. Assume that $A > 5\Delta$, where Δ is as in the preceding proof. Since $\Delta \geq$ diameter$(\mathcal{R})$, no tile can intersect more than two arcs J_i. Moreover, if tiles $x\mathcal{R}$ and $x'\mathcal{R}$ intersect *nonadjacent* arcs J_i and $J_{i'}$, respectively, then

$$d_H(x\mathcal{R}, x'\mathcal{R}) > A \,.$$

Therefore, to prove the lemma it suffices to show that there is a constant $LK < \infty$, independent of m and T, such that the set of tiles $x\mathcal{R}$ that intersect any arc J_i has cardinality no larger than $L/3$. But for any J_i, any tile that intersects J_i is contained in a circle of radius $3A + \Delta$ centered at one of the endpoints of J_i. By the local finiteness of the tessellation, there is an upper bound $L/3$ to the nuber of tiles that intersect any disk of fixed radius. □

8.2 Construction of coverings. Fix $\rho > 0$ (small), and, for each $x \in K(m)$, define J_x^m to be the arc of S_m consisting of all points on S_m at hyperbolic distance less than $\exp\{\rho m\}$ from the tile $x\mathcal{R}$. For each arc $J = J_x^m$, denote by $U(J)$ the *shadow* of J on the boundary $\partial\mathbb{H}$, that is, the set of all points $\zeta \in \partial\mathbb{H}$ such that the geodesic ray beginning at ω and converging to ζ passes through the arc J. For each m, define

$$\mathcal{U}_m = \left\{U(J_x^n) : x \in K(n)\right\},$$
$$\mathcal{U}_m^* = \bigcup_{n=m}^{\infty} \mathcal{U}_m .$$

PROPOSITION 8.5. *With probability one, for every $m \geq 1$ the collection $\mathcal{U}_m^*$ is a covering of Λ.*

Before proceeding to the proof, we show that Propositions 8.1 and 8.5 imply the following corollary.

COROLLARY 8.6. *With probability one, the Hausdorff dimension $\delta_H(\Lambda)$ satisfies*

$$\delta_H(\Lambda) \leq \delta(p) . \tag{8.12}$$

Proof. First, observe that for each $J = J_x^m$ the set $U(J)$ is itself an arc of the boundary circle $\partial\mathbb{H}$ whose Euclidean length is approximately $\exp\{m(\rho - 1)\}$. By Proposition 8.5, the intervals $J = J_x^n$, where $n \geq m$ and $x \in K(m)$, cover Λ with probability one, no matter how small ρ is chosen. Consequently, to prove the inequality (8.12) it suffices to prove that, for any $\delta_* > \delta(p)/(1-\rho)$,

$$\begin{aligned} &\sum_{n=m}^{\infty} \sum_{x\in K(n)} |J_x^m|^{\delta_*} < \infty \qquad \text{a.s.} \\ &\qquad \Longleftrightarrow \sum_{n=m}^{\infty} Y_m \exp\{m(\rho-1)\delta_*\} < \infty \qquad \text{a.s.} \\ &\qquad \Longleftarrow \sum_{n=m}^{\infty} EY_m \exp\{m(\rho-1)\delta_*\} < \infty \end{aligned}$$

almost surely. But the last inequality follows from Proposition 8.1. □

Proof of Proposition 8.5. Fix m large. Let G_m be the union of the arcs $\{J_x^m\}_{x\in K(m)}$, and let $F_m = S_m - G_m$ be the complement of G_m in the circle S_m. The set F_m is itself the union of nonoverlapping arcs I, each bordered by two arcs J_I, J_I' in the collection $\{J_x^m\}_{x\in K(m)}$; moreover, the number of arcs I is no larger than the number of arcs J, which is by definition Y_m. Observe that the composition of the set F_m depends only on the Bernoulli random variables used to determine the colors of the tiles in the "ball" $B(m)$ (see (3.1) for the definition of $B_\Gamma(t)$), and therefore is independent of the assignment of colors to tiles *outside* $B_\Gamma(m)$.

By construction, the minimal distance *along the circle* S_m from any arc I to the nearest tile $x\mathcal{R}$ such that $x \in K$ is at least $\exp\{m\rho\}$. Along this distance there are $\asymp \exp\{m\rho\}$ pairwise disjoint halfplanes exterior to the circle S_m, by Lemma 3.5. By Lemma 2.14, the probability that *none* of these halfplanes contains a red path connecting S_m to $\partial\mathbb{H}$ is less than $\varepsilon^{\exp\{m\rho\}}$ for a suitable $\varepsilon < 1$ independent of m. Since the number Y_m of arcs I is no larger than $\exp\{t\}$ (by a simple counting argument) it follows, by the Borel-Cantelli Lemma, that with probability one, for suffices large m *every* arc I is bordered by connected red clusters connecting S_m to $\partial\mathbb{H}$. These red clusters separate the shadows $U(I)$ of the arcs I from the blue cluster K. Consequently, with probability one, for suffices large m,

$$\Lambda \subset \bigcup_{x\in(m)} U(J_x^m)\,. \tag{8.13}$$

□

9 Variational Formula for Hausdorff Dimension III. Lower Bound

The proof that $\delta(p)$ is a *lower* bound for the Hausdorff dimension of Λ when $0 < p < p_u$ will use the existence of regularly embedded, homogeneous, rooted trees, which was established in section 3. To relate the Hausdorff dimensions of these embedded trees to the function $\delta(p)$, we shall prove that the expected growth rate of connected blue clusters is at least $\delta(p)$.

9.1 Expected growth of clusters III.

PROPOSITION 9.1.

$$\liminf_{t\to\infty} \frac{\log EY(t)}{t} \geq \delta(p)\,. \tag{9.1}$$

The proof of Proposition 9.1 uses a characterization of the Kolmogorov–Sinai entropy of an invariant probability measure due to Katok [Ka]. Fix $t > 0$ and $\rho > 0$, and let μ be an invariant probability measure for the geodesic flow. Say that a finite subset F of $\Gamma(t)$ is a (t,ρ)-*cover* for the measure μ if

$$\mu\big\{(z,\theta) : (\Phi_s(z,\theta))_{s\geq 0} \text{ intersects } x\mathcal{R} \text{ for some } x \in F\big\} \geq \rho\,. \tag{9.2}$$

Lemma 9.2. *For any invariant probability measure μ, let $N(\rho, t, \mu)$ be the minimum cardinality of a (t, ρ)-cover of μ. Then for every $1 > \rho > 0$,*

$$\liminf_{t\to\infty} \frac{\log N(\rho, t, \mu)}{t} = h(\mu)\,. \tag{9.3}$$

Lemma 9.2 is nothing more than a translation of Katok's result from the geodesic flow on SM to the geodesic flow in the universal covering space $\mathbb{H}$. Next, we show that this lemma implies Proposition 9.1.

Proof of Proposition 9.1. Recall that $\delta(p) = \max_{\mu\in\mathcal{I}}(h(\mu) + \beta(\mu; p))$ where the maximum is taken over all *ergodic* invariant probability measures μ. Also, $\beta(\mu; p) = \lim_{t\to\infty} t^{-1}\beta_t(\mu; p)$ where, by equation (6.3), $\beta_t(\mu; p) = E_\mu \log\tau(\gamma \circ \Phi_t; p)$. Consequently, for any $\varepsilon > 0$ there exist $T \geq 1$ and an ergodic $\mu \in \mathcal{I}$ such that

$$h(\mu) + T^{-1}\beta_T(\mu; p) \geq \delta(p) - \varepsilon\,. \tag{9.4}$$

Since μ is ergodic, the log-subadditivity inequality and Birkhoff's Ergodic theorem, together with the definition (6.3), imply that, for all sufficiently large m,

$$\mu\big\{\tau(\gamma \circ \Phi_{mT}; p) \leq \exp\{m\beta_T(\mu; p) - m\varepsilon\}\big\} \leq \varepsilon/2\,. \tag{9.5}$$

Now let F_m be a (mT, ε)-cover for the measure μ, and let F_m^* be the set of all $x \in F_m$ such that $\tau(x; p) \geq \exp\{m\beta_T(\mu; p) - m\varepsilon\}$. By inequality (9.5), the set F_m^* must be a $(mT, \varepsilon/2)$-cover, provided that m is sufficiently large, and so by Lemma 9.2 (with $\rho = \varepsilon/2$) its cardinality must satisfy

$$|F_m^*| \geq \exp\big\{m(Th(\mu) - \varepsilon)\big\}.$$

Hence, by (9.4), for all large m,

$$\begin{aligned}
EY_{mT} &= \sum_{x\in\Gamma(mT)} \tau(x; p) \\
&\geq \sum_{x\in F_m^*} \tau(x; p) \\
&\geq \sum_{x\in F_m^*} \exp\big\{m\beta_T(\mu; p) - m\varepsilon\big\} \\
&\geq \exp\big\{m(Th(\mu) + \beta_T(\mu; p) - 2\varepsilon)\big\} \\
&\geq \exp\big\{mT(\delta(p) - \varepsilon - 2\varepsilon/T)\big\}\,.
\end{aligned}$$

It follows from the Harnack inequality, by a routine argument, that

$$\liminf_{t\to\infty}(EY_t)^{1/t} \geq \exp\big\{\delta(p) - \varepsilon - 2\varepsilon/T\big\}\,.$$

Since $\varepsilon > 0$ was arbitrary, the proof of the proposition is complete. □

9.2 Expected growth of clusters IV. Recall that $Y^*(t) = |K^*(t)|$, where $K^*(t)$ is the set of all $x \in K(t)$ connected to 1 by a blue path that lies entirely in $B_\Gamma(t)$. Thus, $Y^*(t) \le Y(t)$. The next proposition shows that $EY^*(t)$ and $EY(t)$ have the same asymptotic rate of growth.

PROPOSITION 9.3.

$$\liminf_{t\to\infty} \frac{\log EY^*(t)}{t} \ge \delta(p)\,. \tag{9.6}$$

Together with Proposition 8.1, Proposition 9.3 implies that

$$\lim_{t\to\infty} t^{-1}\log EY^*(t) = \delta^*(p) = \delta(p)\,, \tag{9.7}$$

proving assertion (3.4) of Proposition 3.3. This is what will be needed for the proof of the lower bound $\delta_H(\Lambda) \ge \delta(p)$ in section 9.3 below.

The proof of Proposition 9.3 is based on an idea that, in a simpler context, originates in [LS2, section 3]. The main step is to show that $\tau(x) = \tau(x;p)$ is of the same exponential order of magnitude as $\tau^*(x) = \tau^*(x;p)$, where

$$\tau^*(x;p) \overset{\Delta}{=} P_p\big\{\exists \text{ blue path } 1 \to x \text{ in } B_\Gamma(R)\big\} \quad \forall\, x \in \Gamma(R)\,. \tag{9.8}$$

LEMMA 9.4. *For any $\varepsilon > 0$ there exists $R = R_\varepsilon < \infty$ sufficiently large that for all $x \in \Gamma$* not *in the ball $B_\Gamma(R)$ of radius R,*

$$\tau^*(x) \ge (1-\varepsilon)^{|x|}\tau(x)\,. \tag{9.9}$$

Proof. By Proposition 5.1, the connectivity function τ satisfies the system of inequalities (5.1). By the monotone convergence theorem, for each T there exists $A < \infty$ such that $\tau(\gamma\circ\Phi(T)) \le (1+\eta)\tau^A(\gamma\circ\Phi(T))$, where $\tau^A(x)$ denotes the probability of the event $H^A(y,yx)$ that there is a blue path from y to yx that lies entirely in the ball $yB_\Gamma(T+A)$, and T is such that $x \in \Gamma(T)$. Hence,

$$\tau(\gamma\circ\Phi_{mT}) \le (1+\eta)^{2m}\prod_{j=0}^{m-1}\tau^A(\gamma\circ\Phi_T\circ\phi_{jT})\,. \tag{9.10}$$

Now by the FKG inequality, the events $H^A(y;yx)$ are positively correlated; thus, (9.10) implies that $\tau(\gamma\circ\Phi_{mT})/(1+\eta)^{2m}$ is bounded above by the probability that there is a blue path from 1 to $\gamma\circ\Phi_{mT}$ that stays within a tube of radius $A + \text{diameter}(\mathcal{R})$ surrounding the geodesic segment from ω to $(\gamma\circ\Phi_{mT})\omega$. Now if there is such a blue path, and if all the tiles along the geodesic segment from Φ_{mT} to Φ_{mT+A} are blue, then there is a blue path from 1to $\gamma\circ\Phi_{mT+A}$ that lies entirely in the ball $B_\Gamma(mT+A)$. Consequently, by the Harnack inequality, for a suitable constant C,

$$(1+\eta)^{2m}\tau^*(\gamma\circ\Phi_{mT+A}) \ge p^{CA}\tau(\gamma\circ\Phi_{mT+A}) \qquad \forall\, m \ge 1\,. \tag{9.11}$$

Inequality (9.9) follows easily from this, by another application of the Harnack inequality. □

Proof of Proposition 9.3. Fix $\varepsilon > 0$; then by Lemma 9.4, for all R sufficiently large,

$$\begin{aligned} EY^*(R) &= \sum_{x\in\Gamma(R)} \tau^*(x) \\ &\geq \sum_{x\in\Gamma(R)} (1+\varepsilon)^{-R}\tau(x) \\ &= (1+\varepsilon)^{-R} EY(R)\,. \end{aligned}$$

This implies that

$$\liminf_{R\to\infty}(EY^*(R))^{1/R} \geq (1+\varepsilon)^{-1}\lim_{R\to\infty}(EY^*(R))^{1/R} = \delta(p)\,,$$

by Propositions 8.1 and 9.1. Since $\varepsilon > 0$ is arbitrary, the result follows. □

9.3 Lower bound for the Hausdorff dimension of Λ.

Proposition 9.5. *With probability one, the Hausdorff dimension $\delta_H(\Lambda)$ satisfies*

$$\delta_H(\Lambda) \geq \delta(p)\,. \tag{9.12}$$

Together with Corollary 8.6, Proposition 9.5 implies Theorem 4. The proof of Proposition 9.5 uses the existence of the regularly embedded homogeneous rooted trees constructed in section 3. Recall that these trees are contained in the Galton–Watson trees $\mathcal{T}_t^\sigma$ of section 3.4, which themselves are subtrees of the homogeneous trees $\mathbb{T}_t^\sigma$ of section 3.3. Since $\mathbb{T}_t^\sigma$ is regularly embedded in G, its space of ends $\partial\mathbb{T}_t^\sigma$ is embedded homeomorphically as a compact subset of $\partial\mathbb{H}$. The next lemma allows us to relate the induced Euclidean metric d_E on $\partial\mathbb{T}_t^\sigma$ to one of the natural metrics d_θ on $\partial\mathbb{T}$ defined by

$$d_\theta(\xi,\xi') = \theta^{N(\xi,\xi')}\,, \tag{9.13}$$

where $N(\xi,\xi')$ denotes the depth of the tree where the ends ξ and ξ' separate.

Lemma 9.6. *There is a constant $C > 0$ with the following property. Let $x, x' \in L_n = L_n^{\sigma,t}$ be distinct vertices of $\mathbb{T}_t^\sigma$ at depth n, and let $\zeta,\zeta' \in \partial\mathbb{H}$ be points of the ideal boundary contained in the arcs of $\partial\mathbb{H}$ that bound the exterior halfplanes xH^σ and $x'H^\sigma$, respectively. Then*

$$d_E(\zeta,\zeta') \geq C\exp\{-nt\}\,. \tag{9.14}$$

Proof. Since exterior halfplanes yH_Γ^σ, where $y \in \mathcal{L}$, are nested (as noted in the construction of $\mathbb{T}$), we may assume that both x and x' are descendants of a common vertex $y \in L_{n-1}$. By Lemma 3.6, the tile $y\mathcal{R}$ is at hyperbolic distance $\leq nt$ from the fundamental tile $\mathcal{R}$, and tiles $x\mathcal{R}$ and $x'\mathcal{R}$ are at distances $\leq t$ from $y\mathcal{R}$. By construction of $\mathbb{T}_t^\sigma$, the halfplanes xH^σ and $x'H^\sigma$ are properly separated. Since these halfplanes border tiles $x\mathcal{R}$ and $x'\mathcal{R}$ that are at distances no more than nt from $\mathcal{R}$, which contains the Euclidean center of the Poincaré disk, it follows by Proposition 10.3 that the Euclidean distance between their boundary arcs is at least $C\exp\{-nt\}$, for an appropriate constant $C > 0$ independent of x, x', and n. □

COROLLARY 9.7. *The metrics d_E and d_θ, with $\theta = e^{-t}$, satisfy*

$$d_E \geq C d_\theta\,. \tag{9.15}$$

This follows trivially from the preceding lemma. The upshot is that, to establish a lower bound for the Hausdorff dimension of the space of ends of a subtree of $\mathbb{T}_t^\Sigma$, it suffices to compute in the metric d_θ, with $\theta = e^{-t}$.

Proof of Proposition 9.5. By Propositions 3.7 and 3.8, for each $\varepsilon > 0$ the tree $\mathbb{T} = \mathbb{T}_t^\sigma$ contains a Galton–Watson subtree $\mathcal{T}$ whose mean offspring number satisfies

$$E|V_1| > \exp\left\{t(\delta^*(p) - \varepsilon)\right\}.$$

By Hawkes' theorem (see [Ha], also [Ly]), on the event of nonextinction, the space of ends of this Galton–Watson tree has Hausdorff dimension $\delta^*(p)-\varepsilon$, in the metric d_θ. It follows that, with positive probability, the Hausdorff dimension of the limit set Λ is at least $\delta^*(p) - \varepsilon$. But it is easily seen, by use of the Kolmogorov zero-one law, that $\delta_H(\Lambda)$ is almost surely constant on the event $\Lambda \neq \emptyset$; thus, $\delta_H(\Lambda) \geq \delta^*(p)$ on the event $\Lambda \neq \emptyset$. By equation (9.7), $\delta^*(p) = \delta(p)$, so Proposition 9.5 follows. □

10 Appendix: Hyperbolic Geometry

In this section we collect several relatively simple but slightly arcane facts about the geometry of the tessellation $\{x\mathcal{R}\}_{x\in\Gamma}$ associated with a co-compact Fuchsian group Γ.

10.1 Comparability of metrics. Recall that, for any $x \in \Gamma$, we denote by $|x|_G = d_G(1,x)$ and $|x|_{G^*} = d_{G^*}(1,x)$ the distances from 1 to x in the graphs G and G^*, respectively. The distance $|\cdot|_G$ is often called the *word length metric*, because $|x|_G$ is the length of the shortest product (word) in the generators Γ_* representing the group element x. The next proposition relates the word length metric and the hyperbolic metric on $\mathbb{H}$.

Fix $\omega \in \mathcal{R}$; for any $x \in \Gamma$, define $|x|_H = |x|_{H,\omega}$ to be the length of the hyperbolic geodesic segment with endpoints ω and $x\omega$.

Proposition 10.1. *For any $\omega \in \mathcal{R}$,*

$$|x|_G \asymp |x|_{G^*} \asymp d_H(\omega, x\omega)\,,$$

and the implied constants may be chosen independent of the reference point ω.

Proof. We will only show that $|x|_G \asymp d_H(\omega, x\omega)$, as the proof that $|x|_{G^*} \asymp d_H(\omega, x\omega)$ is similar. Let C be the diameter of the fundamental region $\mathcal{R}$ – by hypothesis, Γ is co-compact, so $C < \infty$. Fix $x \in \Gamma$. Suppose that x has the representation $x = y_1 y_2 \dots y_n$ as a word in the generators Γ_*; then there is a continuous path in $\mathbb{H}$ from ω to $x\omega$ of length $\leq Cn$, gotten by concatenating the geodesic segments from $y_1 y_2 \dots y_k \omega$ to $y_1 y_2 \dots y_{k+1}\omega$. Consequently, the length of the *geodesic* path from ω to $x\omega$ is no greater than Cn. This proves that

$$d_H(\omega, x\omega) \leq C|x|\,.$$

To prove an inequality in the opposite direction, we appeal to the fact that the tessellation $\{x\mathcal{R}\}_{x\in\Gamma}$ is *locally finite.* In particular, (a) for each radius $r < \infty$ there exists a constant $D = D(r) < \infty$ such that no ball of hyperbolic radius r intersects more than $D(r)$ distinct tiles; and (b) the infimal hyperbolic distance ρ between distinct points $x\omega, x'\omega$ in the Γ-orbit of ω is strictly positive. Fix $x \in \Gamma$, and consider the geodesic segment ψ from ω to $x\omega$. This segment visits a sequence of adjacent tiles and therefore determines a word $y_1 y_2 \dots y_n$ in the side-pairing transformations. (Unless it passes through a "corner" point of some tile. In this case we alter ψ by inserting small circular loops at all such corner points, so that the modified path avoids the corners. This can be done without multiplying the length of ψ by more than a constant factor, say 2.) The length of this word is no larger than $D(\rho)|\psi|_H/\rho)$, where $|\psi|_H$ is the hyperbolic length of ψ. This proves that

$$|x| \leq 2D(\rho) d_H(\omega, x\omega)/\rho$$

for all $x \neq 1$. □

10.2 Halfplanes. Let H and H' be hyperbolic halfplanes with bounding geodesics ψ and ψ', respectively. Say that H and H' are *properly separated* if they do not intersect and are at distance at least twice the diameter of $\mathcal{R}$. Similarly, say that H is *properly contained* in H' if $H \subset H'$ and the halfplanes $H$and $(H')^c$ are properly separated. Note that if H is properly contained in H' then $H' \setminus H$ is a *strip* (see section 2.7.2).

Proposition 10.2. *For any hyperbolic halfplane H there exists a finite subset F of Γ with the following property: Every halfplane H' of $\mathbb{H}$ that contains the fundamental tile $\mathcal{R}$ properly contains a halfplane xH, where $x \in F$.*

Proof. Use the disk model $\mathbb{D}$ for $\mathbb{H}$, normalized so that the reference point $\omega \in \mathcal{R}$ is the Euclidean center of $\mathbb{D}$. For any halfplane H', let $A_{H'}$ be the (open) arc of the ideal boundary $\partial\mathbb{H}$ that bounds H' (that is, the set of all points of $\partial\mathbb{H}$ to which geodesic rays in H' converge). Because the fundamental tile $\mathcal{R}$ contains the Euclidean center of $\mathbb{D}$, if H' contains $\mathcal{R}$ then the arc $A_{H'}$ must have Euclidean arclength at least π.

Because the group Γ is co-compact, for any two nonoverlapping open arcs I_-, I_+ of $\partial\mathbb{H}$ there exists a hyperbolic element $x \in \Gamma$ whose repulsive and attractive fixed points ζ_-, ζ_+ lie in I_-, I_+, respectively. Let ψ be the axis of x (that is, ψ is the geodesic with endpoints ζ_-, ζ_+ at ∞). If the halfplane H does not intersect ψ then, for all sufficiently large n the boundary arc A_{x^nH} of the halfplane x^nH will be contained in the arc I_+. Thus, for any $\varepsilon > 0$ there is a finite subset $F \subset \Gamma$ such that

(a) for every $x \in F$ the arclength of the boundary arc A_{xH} is less than ε; and

(b) the set $\{A_{xH} : x \in F\}$ covers the boundary circle $\partial\mathbb{H}$.

Recall that if H' is any halfplane containing the tile $\mathcal{R}$, then the arc $A_{H'}$ has Euclidean length at least π. Consequently, the midpoint of $A_{H'}$ is in some arc xA_{xH} of length $< \varepsilon < \pi/3$, where $x \in F$. It follows that the arc $x(A_H)$ is completely contained in an arc of length 3ε centered at the midpoint of $A_{H'}$. If ε is sufficiently small, then this implies that the halfplane xH is properly contained in H'. □

If two halfplanes are properly separated, then their boundary arcs (on $\partial\mathbb{H}$) cannot be too close. The next proposition makes this precise. Let ω be the Euclidean center of $\mathbb{H}$ in the Poincaré disk model, and let d_E be the Euclidean metric on the boundary circle $\partial\mathbb{H}$.

Proposition 10.3. *There exists a constant $C > 0$ with the following property. If H and H' are properly separated halfplanes, each within hyperbolic distance r of ω, with boundary arcs A and A', respectively, then*

$$d_E(A, A') \geq Ce^{-r}. \tag{10.1}$$

Proof. Suppose that, in the upper halfplane representation of $\mathbb{H}$, the hyperbolic halfplanes H and H' are properly separated and situated so that their bounding geodesics σ, σ' are both tangent to the horizontal line

$\{x+i : x \in \mathbb{R}\}$, at the points $x+i, x'+i$, respectively. Since H and H' are properly separated, their boundary arcs are separated by a Euclidean distance $D > 0$ independent of x, x'.

Now consider any two hyperbolic halfplanes H and H' at distance $\leq r$ from ω. Then, by appropriate choice of isometry, it may be arranged that in the upper halfplane representation, $\omega = i$ and the bounding geodesics of H and H' are tangent to a horizontal line $\{x+iy : x \in \mathbb{R}\}$ such that $y \geq r$. Since homotheties of the upper halfplane are hyperbolic isometries, it follows by the preceding section that the boundary arcs of H and H' are separated by a Euclidean distance at least De^{-r}. Since there is a linear fractional transformation that maps the upper halfplane isometrically onto the unit disk, sending i to 0, and since this linear fractional transformation is bi-Lipshitz on the ideal boundary (except at $i\infty$), the result follows. $\square$

The next two propositions are elementary exercises in hyperbolic geometry.

Proposition 10.4. *There exists a constant $C > 0$ such that every geodesic segment α of finite length intersects at least $C|\alpha|_H$ pairwise disjoint barriers. Here $|\alpha|_H$ denotes the hyperbolic length of α.*

Proposition 10.5. *There exists a constant $C > 0$ such that the following is true. For any arc α of a hyperbolic circle bounding a ball B, there are at least $C|\alpha|_H$* nonintersecting *geodesics exterior to B, each tangent to α.*

10.3 Horocycles and horocyclic regions. A *horocycle* asymptotic to a point $\xi \in \partial\mathbb{H}$ is defined to be a curve that intersects *orthogonally* every geodesic with ξ as a limit point. In the Poincaré disk model of the hyperbolic plane, the horocycles asymptotic to ξ are the Euclidean circles tangent to $\partial\mathbb{H}$ at ξ (and contained in the disk). For each $\xi \in \partial\mathbb{H}$, the horocycles asymptotic to ξ may be parametrized by $t \in \mathbb{R}$ as follows: Fix a geodesic $\gamma(t)$ that converges to ξ as $t \to \infty$, and let ψ_t be the (unique) horocycle asymptotic to ξ that contains the point $\gamma(-t)$. It is easily proved that this parametrization is unique (up to an isometry of $\mathbb{R}$), so we may refer to it as the *natural parametrization.*

Define a *horocyclic region* to be a geodesically convex region bounded by a horocyle. One may easily check that, for any horocyle ψ, there is one and only one horocyclic region bounded by ψ (in the disk model, it is the interior of the Euclidean circle ψ). If $\{\psi_t\}_{t\in\mathbb{R}}$ is a natural parametrization of the family of horocycles asymptotic to $\xi \in \partial\mathbb{H}$, then we shall denote by Ψ_t the horocyclic region bounded by ψ_t. For any horocyclic region Ψ with bounding horocycle ψ, and any point $z \in \psi$, define the *exterior halfplane* to Ψ at z to be the halfplane whose bounding geodesic is tangent to ψ at z that does not intersect Ψ.

Let H, H' be nonintersecting halfplanes. Say that H and H' are *properly separated* if the hyperbolic distance between them is at least twice the diameter of the fundamental tile $\mathcal{R}$. If H and H' are properly separated then no tile of the tessellation can intersect both $H$and H'. Moreover, if H and H' are properly separated then $\mathbb{H} \setminus H$ is properly contained in H'.

PROPOSITION 10.6. *There exists a constant $C < \infty$ such that the following is true: If z and z' are any two points on a horocycle ψ at distance at least C, then the exterior halfplanes H, H' to the horocyclic region Ψ bounded by ψ at the points z and z', respectively, are properly separated.*

Proof. Without loss of generality, we may use the upper halfplane representation of the hyperbolic plane $\mathbb{H}$, and take the horocycle $\psi = \psi_0$ to be the horizontal line $\{x + i : x \in \mathbb{R}\}$. For any $z = x + i$ on this horocycle, the exterior halfplane H_z is bounded by the Euclidean semi-circle of radius 1 centered at $x \in \mathbb{R}$. It is now clear that if $z = x + i$ and $z' = x' + i$ are such that $|x - x'| > 2$ then the bounding geodesics of H_z and $H_{z'}$ will not intersect and they will have distinct endpoints on $\partial\mathbb{H}$; and hence that if $|x - x'| > C$, where C is sufficiently large, then the hyperbolic distance between H_z and $H_{z'}$ will be at least $2\times$diameter$(\mathcal{R})$. □

It is possible to base the construction of embedded Galton–Watson trees (see section 3.2) solely on this simple geometric result; however, the construction is somewhat less arduous if the more sophisticated *spacing lemma* (Lemma 3.5) is used. To prove this lemma, we shall first prove an analogous result for horocycles, and then deduce the spacing lemma by using the fact that long arcs of hyperbolic circles are well approximated by horocycle segments. Recall that Σ is the (finite) set of geodesic segments that bound the fundamental tile $\mathcal{R}$; for each $\sigma \in \Sigma$, H^σ is the hyperbolic halfplane containing $\mathcal{R}$ whose bounding geodesic is the infinite geodesic σ^* obtained by extending σ; and $x_\sigma \in \Gamma$ is the side-pairing transformation that maps $\mathcal{R}$ onto the tile $x_\sigma\mathcal{R}$ on the other side of the segment σ. If $xx_\sigma\mathcal{R}$ is a tile of the tessellation that intersects a horocycle ψ, then xx_σ (or $xx_\sigma\mathcal{R}$) is said to be of type σ (relative to ψ) if the halfplane xH^σ does not intersect the horocyclic region Ψ bounded by ψ.

PROPOSITION 10.7. *There exists $C < \infty$ such that the following is true: For each $\sigma \in \Sigma$ and each horocyclic region Ψ with bounding horocycle ψ, there is a sequence $x_n x_\sigma \mathcal{R}$ of tiles of type σ relative to ψ such that*

(a) If $n \neq m$ then the halfplanes $x_n H^\sigma$ and $x_m H^\sigma$ are properly separated;
(b) For each tile $x\mathcal{R}$ that intersects the horocycle ψ there is a tile $x_n\mathcal{R}$ within distance C of $x\mathcal{R}$.

Proof. This proposition follows from Proposition 10.6 and the unique ergodicity of the horocycle flow (see [F1]). Unique ergodicity and the fact that the unique invariant measure has full support in the unit tangent bundle implies that the horocycle flow is *uniformly recurrent* (see [F2, section 1.4 and Proposition 3.6]).

Call a point z on the horocycle ψ a *good* point if it lies in a tile of type σ relative to ψ. By uniform recurrence, for any $C < \infty$ there exists a sequence $\{z_n\}_{n\in\mathbb{Z}}$ of good points such that no two are within distance C, but such that for some $C' > C$, every $z \in \psi$ is within distance C' of some z_n. □

The next proposition quantifies the extent to which circles of large radii are approximated by horocycles.

PROPOSITION 10.8. *For any $\varepsilon > 0$ there exist $\eta = \eta(\varepsilon) > 0$ and $\rho = \rho(\varepsilon) > 0$ with the following property. For any parametrized arc $\{\alpha(t)\}_{t\in J}$ of hyperbolic length ηe^r on a circle of radius $r \geq \rho$, there is a parametrized horocycle segment $\{\psi(t)\}_{t\in J}$ such that*

$$\max_{t\in J} d_H\big(\alpha(t), \psi(t)\big) \leq \varepsilon\,. \tag{10.2}$$

Proof. Without loss of generality, we may assume that, in the upper halfplane representation, the circle is centered at i and the arc $\alpha(t)$ passes through the imaginary axis at the point ie^r. The arc $\alpha(t)$ must then meet the horocycle $\psi = \{x + ie^{-r} : x \in \mathbb{R}\}$ tangentially at ie^{-r}, and lie above it elsewhere. Take the approximating horocycle segment to be the parametrized segment of ψ obtained by projecting the parametrized arc $\alpha(t)$ vertically downward. It is clear that if the arc $\alpha(t)$ is sufficiently short, then (10.2) will hold. That $\eta > 0$ can be chosen so that (10.2) holds uniformly, as advertised, for all r sufficiently large, follows from the fact that, in the upper halfplane representation, the mapping $z \to az$ is an isometry for any $a > 0$. □

Proof of Lemma 3.5. Proposition 10.8 implies that any arc of circle of large radius may be arbitrarily well approximated by long horocycle segments. If the approximation of a circular arc α by a horocycle segment ψ is sufficiently good, then any tile $x\mathcal{R}$ that is of type σ relative to ψ will be of type σ relative to $B_\Gamma(R; z)$, where R is the radius of α and $z\mathcal{R}$ is the tile containing the center of the circle. The lemma now follows routinely. □

References

[A] A. Ancona, Positive harmonic functions and hyperbolicity, in "Potential theory – surveys and problems (Prague, 1987)", Springer Lecture Notes in Math. 1344 (1988), 1–23.

[AtN] K. Athreya, P. Ney, Branching Processes, Springer–Verlag, NY, 1972.

[B] A. Beardon, The Geometry of Discrete Groups, Springer–Verlag, NY, 1983.

[BeS] I. Benjamini, O. Schramm, Percolation in the hyperbolic plane, preprint (2000).

[BeLPS1] I. Benjamini, R. Lyons, Y. Peres, O. Schramm, Critical percolation on any nonamenable group has no infinite clusters, Annals of Probability 27 (1999), 1347–1356.

[BeLPS2] I. Benjamini, R. Lyons, Y. Peres, O. Schramm, Group invariant percolation, Geometric And Functional Analysis 9 (1999), 29–66.

[BerK] J. van den Berg, M. Keane, On the continuity of the percolation probability function, in "Conference in modern analysis and probability (New Haven, Conn., 1982)", Amer. Math. Soc. (1984), 61–65.

[Bo] B. Bollobas, Graph Theory, Springer-Verlag, NY, 1979.

[BuK] R. Burton, M. Keane, Density and uniqueness in percolation, Comm. Math. Phys. 121 (1989), 501–505.

[F1] H. Furstenberg, Unique ergodicity of the horocycle flow, Recent advances in topological dynamics, Springer Lecture Notes in Math. 318 (1973), 95–115.

[F2] H. Furstenberg, Recurrence in ergodic theory and combinatorial number theory, Princeton University Press, Princeton, NJ, 1981.

[G1] G. Grimmett, Percolation, Springer-Verlag, NY, 1989.

[G2] G. Grimmett, Lecture Notes on Disordered Systems, Springer Lecture Notes in Math. 1665 (1997), 153–300.

[H1] O. Häggström, Infinite clusters in dependent automorphism invariant percolation on trees, Annals of Probability 25 (1996), 1423–1436.

[H2] O. Häggström, Personal communication, 1999.

[Ha] J. Hawkes, Trees generated by a a simple branching process, J. London Math. Soc. (2) 24 (1981), 373–384.

[KV] V. Kaimanovich, A. Vershik, Random walks on discrete groups: boundary and entropy, Annals of Probability 11 (1983), 457–490

[Ka] A. Katok, Lyapunov exponents, entropy and periodic orbits of diffeomorphisms, Publ. Math. IHES 50 (1980), 137–173.

[Kat] S. Katok, Fuchsian Groups. University of Chicago Press, 1992.

[Ki] J.F.C. Kingman, The ergodic theory of subadditive stochastic processes, J. Royal Statistical Soc. B30 (1968), 499–510.

[L1] S. Lalley, Percolation on Fuchsian groups. Ann. Inst. H. Poincaré

Probab. Statist. 34 (1998), 151–177.

[L2] S. Lalley, Growth profile and invariant measures for the contact process on a homogeneous tree, Annals of Probability 27 (1999), 206–225.

[LS1] S. Lalley, T. Sellke, Hyperbolic branching Brownian motion, Probability Theory & Related Fields 108 (1997), 171–192.

[LS2] S. Lalley, T. Sellke, Limit set of a weakly supercritical contact process on a homogeneous tree, Annals of Probability 26 (1998), 644–657

[LS3] S. Lalley, T. Sellke, An Extension of Hawkes' Theorem on the Hausdorff Dimension of a Galton–Watson Tree, Probability Theory & Related Fields 116 (2000) 41–56.

[LeY] F. Ledrappier, L.-S. Young, The metric entropy of diffeomorphisms, II Relations between entropy, exponents and dimension, Annals of Mathematics (2) 122 (1985) 540–574.

[Leh] J. Lehner, Discontinuous Groups and Automorphic Functions, American Math. Soc., Providence RI, 1964.

[Li] T. Liggett, Multiple transition points for the contact process on the binary tree, Annals of Probability 24 (1996), 1675–1710.

[Ly] R. Lyons, Random walks and percolation on trees, Annals of Probability 18 (1990), 931–958.

[NS] C.M. Newman, L. Schulman, Infinite clusters in percolation models, J. Statistical Ph. 26 (1981), 613–628.

[P] R. Pemantle, The contact process on trees, Annals of Probability 20 (1992), 2089–2169.

[Pe] Y. Peres, Percolation on nonamenable products at the uniqueness threshold, Annals Inst. H. Poincaré, Prob. et Statist., to appear.

[R] L. Russo, A note on percolation, Z. Wahrscheinlichkeitstheorie und Verw. Gebiete 43 (1978), 39–48.

[S] R. Schonmann, Percolation in $\infty+1$ dimensions at the uniqueness threshold, in "Volume in Honor to Harry Kesten" (M. Bramson, R. Durrett), to appear.

[St] A. Stacey, The existence of an intermediate phase for the contact process on trees, Annals of Probability 24:4 (1996) 1711–1726.

[W] P. Walters, An Introduction to Ergodic Theory, Springer-Verlag, NY, 1982.

Steven P. Lalley, Department of Statistics, University of Chicago, Chicago, IL 60637-1514, USA lalley@galton.uchicago.edu

Submitted: July 2000