J. Appl. Prob. **43,** 665–677 (2006)
Printed in Israel

ULTRA-SMALL SCALE-FREE GEOMETRIC NETWORKS

J. E. YUKICH,* *Lehigh University*

Abstract

We consider a family of long-range percolation models $(G_p)_{p>0}$ on $\mathbb{Z}^d$ that allow dependence between edges and have the following connectivity properties for $p \in (1/d, \infty)$: (i) the degree distribution of vertices in G_p has a power-law distribution; (ii) the graph distance between points $\boldsymbol{x}$ and $\boldsymbol{y}$ is bounded by a multiple of $\log_{pd} \log_{pd} |\boldsymbol{x} - \boldsymbol{y}|$ with probability $1 - o(1)$; and (iii) an adversary can delete a relatively small number of nodes from $G_p(\mathbb{Z}^d \cap [0, n]^d)$, resulting in two large, disconnected subgraphs.

Keywords: Scale-free graph; long-range percolation; chemical distance

2000 Mathematics Subject Classification: Primary 60D05
Secondary 05C80

1. Introduction

The statistical properties of large networks have received considerable attention in the recent scientific literature [2], [14], [21], [25]. Of special interest are the power-law random networks in which the fraction of vertices of degree k is proportional to k^{-q} for some $q > 0$. Such networks lack an inherent scale and have been termed 'scale free'. Scale-free graphs are ubiquitous in random network theory and have been proposed as a way to model the behavior of technological, social, and biological networks [1], [21].

Networks often have a geometric component to them where the vertices have positions in space and geographic proximity plays a role in deciding which vertices get connected. In this context, random geometric graphs are a natural alternative to the classical Erdős–Rényi random graph models. Random connection models [20] provide one way to describe networks with spatial content. In these models the event, $E_{\boldsymbol{x},\boldsymbol{y}}$, of a connection between points $\boldsymbol{x}$ and $\boldsymbol{y}$ has probability $p_{\boldsymbol{x},\boldsymbol{y}} := \mathrm{P}[E_{\boldsymbol{x},\boldsymbol{y}}] = g(|\boldsymbol{x} - \boldsymbol{y}|)$, where $g\colon \mathbb{R}^+ \to [0, 1]$ is a connection function and $|\boldsymbol{x}|$ denotes the Euclidean norm of $\boldsymbol{x}$. The standard long-range percolation model assumes independence of $E_{\boldsymbol{x},\boldsymbol{y}}$ and $E_{\boldsymbol{x},\boldsymbol{u}}$, $\boldsymbol{y} \neq \boldsymbol{u}$, which may not be the case in networked systems. Moreover, the degree distribution in this connection model generally does not follow a power law.

Allowing dependency between edges will in general result in technically more complicated models. In this note we show that a natural edge dependency gives rise to a family of long-range percolation models, $(G_p)_{p>0}$, which is technically tractable and which exhibits three connectivity properties for $p \in (1/d, \infty)$. First, G_p has a power-law distribution. Second, G_p is ultra-small, in the sense that the graph distance between lattice points $\boldsymbol{x}$ and $\boldsymbol{y}$ is bounded by a multiple of $\log_{pd} \log_{pd} |\boldsymbol{x} - \boldsymbol{y}|$ with probability $1 - o(1)$, where $o(1)$ denotes a quantity tending to 0 as $|\boldsymbol{x} - \boldsymbol{y}| \to \infty$. Ultra-small graph distances imply efficiency, are consistent

Received 13 April 2005; revision received 3 May 2006.
* Postal address: Department of Mathematics, Lehigh University, Bethlehem, PA 18015, USA.
Email address: joseph.yukich@lehigh.edu
Partially supported by NSF grant DMS-0203720.

with the 'small-world phenomenon' [2], [14], [24], [25], and are relevant in the context of routeing, searching, and transport of information. Third, an adversary can delete a relatively small number of nodes from $G_p(\mathbb{Z}^d \cap [0, n]^d)$, after which there are two disconnected subgraphs, each containing nearly one-half of the total number of network nodes.

1.1. A general dependent random connection model

Let $\{U_z\}_{z \in \mathbb{Z}^d}$ be independent, identically distributed uniform[0, 1] random variables indexed by $\mathbb{Z}^d$. Let $p > 0$ and $\delta \in (0, 1]$. For each $z \in \mathbb{Z}^d$, we take δU_z^{-p} to represent a weight at node z defining the radius of the 'ball of influence' at z. Consider the graph $G_{p,\delta} := G_{p,\delta}(\mathbb{Z}^d)$ which puts an edge between nodes $x, y \in \mathbb{Z}^d$ *whenever each node is contained in the other's ball of influence*. Thus, this connection rule says that the edge (x, y) appears in $G_{p,\delta}(\mathbb{Z}^d)$ whenever

$$|x - y| \le \delta \min(U_x^{-p}, U_y^{-p}). \tag{1.1}$$

Let $\delta = 1$. By the independence of the U_z, we have $p_{x,y} := \mathrm{P}[E_{x,y}] = |x - y|^{-2/p}$, showing that the probability of (there being) long edges in $G_p := G_{p,1}$ increases with p. Edges in G_p have dependent probabilities: if $|y| < |x|$ then the probability of the edge $(\mathbf{0}, y)$ given the edge $(\mathbf{0}, x)$ is $|y|^{-1/p}$ instead of $|y|^{-2/p}$.

The family of random connection models $G_{p,\delta}$ is disconnected for general p and δ, but not for $\delta = 1$, since having $U_z^{-p} \ge 1$ for all $z \in \mathbb{Z}^d$ implies that adjacent lattice points are connected in G_p. The main results below show, for all $p \in (1/d, \infty)$, that the components of G_p are of arbitrarily large diameter with arbitrarily large probability. Moreover, in accordance with their Poisson Boolean model counterparts (cf. [20]), it is easy to check, for all $\delta \in (0, 1]$ and large p, that the expected number of nodes in the component of $G_{p,\delta}$ containing $\mathbf{0}$ is infinite, whereas, for p and δ both small, the expected number of such nodes is finite. Our purpose here is to explore the connectivity properties of G_p, $p \in (1/d, \infty)$.

1.2. Main results

Let $D_p(\mathbf{0})$ denote the degree of the origin in $G_p(\mathbb{Z}^d)$, let ω_d denote the volume of the unit-radius ball in $\mathbb{R}^d$, and let $\alpha := pd - 1$. Our first result shows that if $p \in (1/d, \infty)$ then the degree of a typical vertex follows a power law, i.e. G_p is scale free.

Theorem 1.1. ($G_p(\mathbb{Z}^d)$ has a power-law degree distribution.) *For all $d = 1, 2, \dots$ and all $p \in (1/d, \infty)$,*

$$\lim_{t \to \infty} t^{1/\alpha}\, \mathrm{P}[D_p(\mathbf{0}) > t] = (pd\omega_d/\alpha)^{1/\alpha}.$$

For all $x, y \in \mathbb{Z}^d$, $d_p(x, y)$ denotes the G_p graph distance ('chemical distance') between x and y. Our next result says that G_p is ultra-small (cf. [12]), in that $d_p(x, y)$ is bounded by $4(2 + \log\log|x - y|)$ with probability $1 - o(1)$, where throughout, for all $s > 0$, $\log s$ is short for $\log_{pd} s$. We expect that the upper bound in this result can be improved but have not tried to obtain the sharpest bound.

Theorem 1.2. ($G_p(\mathbb{Z}^d)$ has small graph distance.) *For all $d = 1, 2, \dots$ and all $p \in (1/d, \infty)$,*

$$\frac{d_p(\mathbf{0}, x)}{2 + \log\log|x|} \le 4$$

with probability $1 - o(1)$, where $o(1)$ tends to 0 as $|x| \to \infty$.

The network failure of $G_p(\mathbb{Z}^d)$ is easily quantified, as follows.

Theorem 1.3. (Network failure.) *For all $d = 1, 2, \dots$ and all $p \in (1/d, \infty)$, an adversary can delete N nodes from $G_p(\mathbb{Z}^d \cap [0, n]^d)$, where $\mathrm{E}[N] = O(n^{d-1}[n^{1-1/p} \vee 1])$, resulting in two disconnected subgraphs on vertex sets of cardinality at least $n^d/2 - N$.*

Theorem 1.3 implies, in particular, that if $p \in (1/d, 1)$ then removing roughly $O(n^{d-1})$ nodes may reduce $G_p(\mathbb{Z}^d \cap [0, n]^d)$ to two large, disconnected subgraphs.

Remarks. 1. *Standard long-range percolation models.* Assume that $p_{\boldsymbol{x},\boldsymbol{y}} := \mathrm{P}[E_{\boldsymbol{x},\boldsymbol{y}}] = |\boldsymbol{x} - \boldsymbol{y}|^{-s+o(1)}$ as $|\boldsymbol{x} - \boldsymbol{y}| \to \infty$, for some constant $s \in (0, \infty)$; $E_{\boldsymbol{x},\boldsymbol{y}}$ and $E_{\boldsymbol{x},\boldsymbol{u}}$ are independent for all $\boldsymbol{x}, \boldsymbol{y}, \boldsymbol{u} \in \mathbb{Z}^d$. For $s \in (0, d)$, Benjamini *et al.* [4] showed that the graph distance $d(\boldsymbol{0}, \boldsymbol{x})$ behaves like the constant $\lceil s/(d - s) \rceil$ as $|\boldsymbol{x}| \to \infty$. Here, $\lceil x \rceil$ denotes the greatest integer less than x. For $s = d$, Coppersmith *et al.* [13] showed that $d(\boldsymbol{0}, \boldsymbol{x})$ scales as $\log |\boldsymbol{x}| / \log \log |\boldsymbol{x}|$, whereas, for $s \in (d, 2d)$, Biskup [7], [8] showed that $d(\boldsymbol{0}, \boldsymbol{x})$ scales as $(\log |\boldsymbol{x}|)^{\Delta + o(1)}$, where $\Delta := \Delta(s, d) := \log 2 / \log(2d/s)$. The case $s = 2d$ is open and, for $s \in (2d, \infty)$, $d(\boldsymbol{0}, \boldsymbol{x})$ scales at least linearly in $|\boldsymbol{x}|$, as shown by Berger [5]. The different scalings for the standard long-range percolation model suggest that G_p also has different scalings for $p \in (0, 1/d)$, but we have not determined them. Kleinberg [19] proposed a lattice model where long-range contacts are added in a biased way, there being, however, a uniform bound on the number of such contacts.

2. *Geometric networks in $\mathbb{R}^d$.* We expect that Theorems 1.1–1.3 extend to analogously defined continuum models on Poisson point sets in $\mathbb{R}^d$. This would add to the following related results.

(a) Let $f\colon \mathbb{R}^d \to \mathbb{R}^+$ and let $\mathcal{P}_f$ be a Poisson point process on $\mathbb{R}^d$ with intensity f. The *geometric graph*, described in depth by Penrose [23], joins two nodes in $\mathcal{P}_f$ whenever their Euclidean distance is less than a specified cutoff. Hermann *et al.* [18, Section II.B] showed that if $\int_{\mathbb{R}^d} f^r(\boldsymbol{x})\,\mathrm{d}\boldsymbol{x} = \infty$ for all $r > r_0$, then the degree distribution is effectively a power law.

(b) The *on-line nearest-neighbors graph* is defined on randomly ordered point sets in $\mathbb{R}^d$, and places an edge between each point and its nearest neighbor amongst the points preceding it in the ordering. Such graphs have scale-free properties over certain degree domains [6], [16].

(c) Franceschetti and Meester [17] developed a scale-free continuum model but did not obtain iterated log bounds on interpoint graph distances.

(d) The standard *Boolean connection model* puts an edge between $\boldsymbol{x}$ and $\boldsymbol{y}$ whenever the respective balls of influence overlap. In the context of (1.1), $(\boldsymbol{x}, \boldsymbol{y})$ is an edge whenever $|\boldsymbol{x} - \boldsymbol{y}| \le \delta(U_{\boldsymbol{x}}^{-p} + U_{\boldsymbol{y}}^{-p})$. These models are not in general scale free.

3. *Power exponents $q \in (2, 3)$.* Consider a random graph on n nodes $\boldsymbol{v}_1, \boldsymbol{v}_2, \dots, \boldsymbol{v}_n$ with weight (expected degree) w_i at node $\boldsymbol{v}_i$. Nodes $\boldsymbol{v}_i$ and $\boldsymbol{v}_j$ are connected with probability $\rho w_i w_j$, where $\rho = (\sum_{i=1}^n w_i)^{-1}$. Chung and Lu [10], [11] provided conditions on the weights under which the degree distribution is proportional to k^{-q}, $q \in (2, 3)$, $k \in \mathbb{Z}$, the average distance between nodes is almost surely $O(\log \log n)$, and the diameter is $O(\log n)$. In unrelated work, Cohen and Havlin [12] argued that whenever the degree distribution of a random graph on n vertices is proportional to k^{-q}, where $q \in (2, 3)$, k is restricted to (m, K), and where m and $K := K(n)$ are well-defined 'cutoffs', then the diameter behaves like $\log \log n$.

4. *Preferential attachment models.* These dynamic graphs evolve with time in such a way that a newly arriving vertex connects to an existing vertex with a probability proportional to the degree of the (latter) vertex. Thus, nodes of high degree tend to acquire more new links than do nodes of low degree. Albert and Barabási [1] showed that such models follow a power law, are not geometry dependent, and, as shown by Bollabás and Riordan [9], are not ultra-small in general.

5. *Degree dependence on p.* Theorem 1.1 tells us that $\mathrm{P}[D_p(\mathbf{0}) = k] \sim Ck^{-q}$, where $q := pd/(pd-1)$. Thus, as p increases on $(1/d, \infty)$, the exponent of the degree distribution, q, decreases to 1.

6. *Further connectivity results.* Theorems 1.1–1.3 describe the connectivity of $G_p(\mathbb{Z}^d)$. Further analysis of the connectivity of $G_p(\mathbb{Z}^d)$, such as thermodynamic and Gaussian limits for the number of three cycles (or other clustering coefficients) on $G_p(\mathbb{Z}^d \cap [0,n]^d)$, is simplified by appealing to the stabilization properties of G_p (see especially [22]). $G_p(\mathbb{Z}^d)$ is *assortative* in that high-degree nodes tend to link to high-degree nodes and low-degree nodes tend to link to low-degree nodes.

7. *The case $p \in (0, 1/d)$.* If $p \in (0, 1/d)$ then G_p has few long edges and the proofs of the scale-free and ultra-small properties break down. The scalar $1/d$ thus represents the boundary between graphs that are ultra-small scale free and those which are not.

2. Proof of Theorem 1.1

Throughout, we adopt the following notation: $B_r(\boldsymbol{x})$ denotes the Euclidean ball of radius r centered at $\boldsymbol{x} \in \mathbb{R}^d$, $L_r(\boldsymbol{x}) := B_r(\boldsymbol{x}) \cap \mathbb{Z}^d \setminus \{\boldsymbol{x}\}$ denotes the lattice points a distant at most r from $\boldsymbol{x}$, and C denotes a generic positive constant whose value may change from line to line. The underlying probability space is $\Omega := [0,1]^{\mathbb{Z}^d}$ and is equipped with the product probability measure $\mathrm{P} := \mu^{\mathbb{Z}^d}$, where μ is the uniform probability measure on $[0,1]$. Conditional on $U_{\mathbf{0}} = u$, $D_p(\mathbf{0})$ is the number of points $\boldsymbol{y}$ in $L_{u^{-p}}(\mathbf{0})$ with weight, $U_{\boldsymbol{y}}^{-p}$, exceeding $|\boldsymbol{y}|$; hence, $U_{\boldsymbol{y}} \in [0, |\boldsymbol{y}|^{-1/p}]$. Writing $D(u^{-p})$ for the value of $D_p(\mathbf{0})$ conditioned on $\mathbf{0}$ having weight u^{-p}, we have

$$D(u^{-p}) = \sum_{\boldsymbol{y} \in L_{u^{-p}}(\mathbf{0})} 1_{\{U_{\boldsymbol{y}} \le |\boldsymbol{y}|^{-1/p}\}}.$$

Thus, to prove Theorem 1.1 we condition on $U_{\mathbf{0}}$ and show that

$$\lim_{t\to\infty} t^{1/\alpha} \int_0^1 \mathrm{P}[D(u^{-p}) > t]\,\mathrm{d}u = \left(\frac{pd\omega_d}{\alpha}\right)^{1/\alpha}, \tag{2.1}$$

where, recall, $\alpha := pd - 1$. The next lemma will be useful in establishing (2.1). Let $\beta := pd\omega_d/\alpha$.

Lemma 2.1. *For all $p \in (1/d, \infty)$, we have*

$$\mathrm{E}[D(u^{-p})] = \beta u^{-\alpha} + O(\max(1, u^{-pd+p+1})), \tag{2.2}$$

where the error on the right-hand side of (2.2) holds as $u \to 0^+$.

Proof. Note that $\mathrm{E}[D(u^{-p})]$ is approximated by

$$\int_{|\boldsymbol{x}|\le u^{-p}} |\boldsymbol{x}|^{-1/p}\,\mathrm{d}\boldsymbol{x} = d\omega_d \int_0^{u^{-p}} t^{d-1-1/p}\,\mathrm{d}t = \beta u^{-\alpha}.$$

https://doi.org/10.1239/jap/1158784937 Published online by Cambridge University Press

Let $R := R(u)$ be the maximal collection of grid cubes (cubes centered at points in $\mathbb{Z}^d$ with edge length 1) contained within $B_{u^{-p}}(\mathbf{0})$. The approximation error

$$\left|\mathrm{E}[D(u^{-p})] - \int_{|\boldsymbol{x}|\le u^{-p}} |\boldsymbol{x}|^{-1/p}\,\mathrm{d}\boldsymbol{x}\right|$$

is bounded by the sum of the following three errors:

$$E_1 := \left|\mathrm{E}[D(u^{-p})] - \sum_{\boldsymbol{y}\in R(u)\cap\mathbb{Z}^d,\ \boldsymbol{y}\neq\mathbf{0}} |\boldsymbol{y}|^{-1/p}\right|,$$

$$E_2 := \left|\sum_{\boldsymbol{y}\in R(u)\cap\mathbb{Z}^d,\ \boldsymbol{y}\neq\mathbf{0}} |\boldsymbol{y}|^{-1/p} - \int_{R(u)} |\boldsymbol{x}|^{-1/p}\,\mathrm{d}\boldsymbol{x}\right|,$$

$$E_3 := \left|\int_{R(u)} |\boldsymbol{x}|^{-1/p}\,\mathrm{d}\boldsymbol{x} - \int_{|\boldsymbol{x}|\le u^{-p}} |\boldsymbol{x}|^{-1/p}\,\mathrm{d}\boldsymbol{x}\right|.$$

Now,

$$E_1 = \sum_{\boldsymbol{y}\in(B_{u^{-p}(\mathbf{0})}\setminus R(u))\cap\mathbb{Z}^d,\ \boldsymbol{y}\neq\mathbf{0}} |\boldsymbol{y}|^{-1/p}$$

and, so, is bounded by the product of

$$\mathrm{card}\{(B_{u^{-p}}(\mathbf{0})\setminus R(u))\cap\mathbb{Z}^d\} \quad \text{and} \quad \sup\{y\in(B_{u^{-p}}(\mathbf{0})\setminus R(u))\cap\mathbb{Z}^d : |y|^{-1/p}\}.$$

Since the first factor is bounded by $Cu^{-p(d-1)}$ and the second by Cu, it follows that $E_1 \le Cu^{-pd+p+1}$. A similar method shows that $E_3 \le Cu^{-pd+p+1}$.

We estimate E_2 as follows. For all $\boldsymbol{y}\in\mathbb{Z}^d$, let $Q_{\boldsymbol{y}}$ denote the grid cube with center $\boldsymbol{y}$. For all $s = 1, 2, \dots$, let $M(s) := \mathrm{card}\{y\in\mathbb{Z}^d : |\boldsymbol{y}|\in[s, s+1)\}$. Since there is a constant $C > 0$ such that, for all $\boldsymbol{x}\in Q_{\boldsymbol{y}}$ and all $\boldsymbol{y}\in\mathbb{Z}^d$,

$$||\boldsymbol{y}|^{-1/p} - |\boldsymbol{x}|^{-1/p}| \le C|\boldsymbol{y}|^{-1/p-1},$$

it follows that

$$E_2 \le C\sum_{s=1}^{u^{-p}} s^{-1/p-1}M(s) \le C\sum_{s=1}^{u^{-p}} s^{-1/p+d-2} \le C\max(1, u^{-pd+p+1}),$$

since $M(s) \le Cs^{d-1}$. Combining the bounds for E_1, E_2, and E_3 yields Lemma 2.1.

Letting $s := u^{-p}$ in (2.1), note that, to prove Theorem 1.1, it suffices to show that

$$\lim_{t\to\infty} t^{1/\alpha}\int_1^\infty \mathrm{P}[D(s) > t]\frac{1}{p}s^{-1/p-1}\,\mathrm{d}s = \beta^{1/\alpha}. \tag{2.3}$$

We observe that (2.3) is plausible because Lemma 2.1 suggests that $\mathrm{P}[D(s) > t]$ is close to 1 for $t \ll \beta s^{\alpha/p}$ and close to 0 for $t \gg \beta s^{\alpha/p}$, indicating that the left-hand side of (2.3) behaves like

$$\lim_{t\to\infty} t^{1/\alpha}\int_{(t/\beta)^{p/\alpha}}^\infty \frac{1}{p}s^{-1/p-1}\,\mathrm{d}s = \beta^{1/\alpha}.$$

To put this heuristic argument on a rigorous footing, we rewrite the integral in (2.3) as a sum of two integrals. The first integral is estimated via Bernstein's inequality and the second is handled using Poisson approximation arguments. We do this as follows.

For all $v > 0$, let $m(v) := \sup\{s\colon \mathrm{E}[D(s)] \le v\}$. Recalling that $\alpha := pd - 1$, from Lemma 2.1 we obtain

$$\mathrm{E}[D(s)] = \beta s^{\alpha/p} + O(\max(1, s^{d-1-1/p})) = \beta s^{\alpha/p}(1 + \max(O(s^{1/p-d}), O(s^{-1}))). \quad (2.4)$$

It follows, for large v and $p \in (1/d, \infty)$, that

$$m(v) = \left(\frac{v}{(1+o(1))\beta}\right)^{p/\alpha},$$

where $o(1)$ tends to 0 as $v \to \infty$. Given fixed $t \ge \beta$ and $\varepsilon \in (0, \frac{1}{2})$, define the following two integration domains:

$$I_1 := [1,\ m(t - t^{1/2+\varepsilon})), \qquad I_2 := [m(t - t^{1/2+\varepsilon}),\ \infty).$$

Rewrite the left-hand side of (2.3) as

$$\lim_{t\to\infty} t^{1/\alpha} \int_{I_1} \mathrm{P}[D(s) > t]\frac{1}{p}s^{-1/p-1}\,\mathrm{d}s + \lim_{t\to\infty} t^{1/\alpha} \int_{I_2} \mathrm{P}[D(s) > t]\frac{1}{p}s^{-1/p-1}\,\mathrm{d}s =: S_1 + S_2,$$

provided that both limits exist.

To prove Theorem 1.1 it suffices to show that $S_1 = 0$ and $S_2 = \beta^{1/\alpha}$. We first show that $S_1 = 0$. Bernstein's inequality [15, p. 12] for sums of independent, bounded random variables yields, for all $s \in I_1$,

$$\mathrm{P}[D(s) > t] \le \exp\left(\frac{-(t - \mathrm{E}[D(s)])^2}{2\ \mathrm{E}[D(s)] + 4t/3}\right).$$

Using the bounds $\inf_{s\in I_1}(t - \mathrm{E}[D(s)]) \ge t^{1/2+\varepsilon}$ and $\sup_{s\in I_1} \mathrm{E}[D(s)] \le t - t^{1/2+\varepsilon} < t$, for all $s \in I_1$ we thus obtain

$$\mathrm{P}[D(s) > t] \le \exp\left(\frac{-(t^{1/2+\varepsilon})^2}{10t/3}\right) = \exp\left(-\frac{3t^{2\varepsilon}}{10}\right).$$

It follows that

$$S_1 \le \limsup_{t\to\infty} t^{1/\alpha} \exp\left(-\frac{3t^{2\varepsilon}}{10}\right) \int_1^\infty \frac{1}{p}s^{-1/p-1}\,\mathrm{d}s = 0.$$

We next show that $S_2 = \beta^{1/\alpha}$. By approximating $D(s)$ with a Poisson random variable we establish the following simplified expression for S_2. Here and elsewhere, $\mathrm{Po}(\lambda)$ denotes a Poisson random variable with mean λ.

Lemma 2.2. *For all $p \in (1/d, \infty)$, we have*

$$S_2 = \lim_{t\to\infty} t^{1/\alpha} \int_{m(t-t^{1/2+\varepsilon})}^\infty \mathrm{P}[\mathrm{Po}(\mathrm{E}[D(s)]) > t]\frac{1}{p}s^{-1/p-1}\,\mathrm{d}s.$$

Proof. For all $\mathbf{y} \in \mathbb{Z}^d$, let $p_{\mathbf{y}} := \mathrm{E}[1_{\{U_{\mathbf{y}} \le |\mathbf{y}|^{-1/p}\}}] = |\mathbf{y}|^{-1/p}$. Letting d_{TV} be the total variation distance, it follows from well-known Poisson approximation bounds (e.g. Equation (1.23) of [3]) that

$$d_{\mathrm{TV}}(D(s), \mathrm{Po}(\mathrm{E}[D(s)])) \le \biggl(\sum_{y\in L_s(\mathbf{0})} p_y\biggr)^{-1} \sum_{y\in L_s(\mathbf{0})} p_y^2.$$

By an analysis similar to that in the proof of Lemma 2.1 and (2.4), for $d > 2/p$ we obtain

$$\sum_{y\in L_s(\mathbf{0})} p_y^2 = \frac{pd\omega_d}{pd-2} s^{d-2/p}(1+o(1)),$$

whereas, for $1/p < d \le 2/p$, we have

$$\sum_{y\in L_s(\mathbf{0})} p_y^2 = O(1).$$

It follows from Lemma 2.1 that, for $d > 2/p$, we obtain

$$\begin{aligned} d_{\mathrm{TV}}(D(s), \mathrm{Po}(\mathrm{E}[D(s)])) &\le (\beta s^{d-1/p}(1+o(1)))^{-1}\beta\biggl(\frac{pd-1}{pd-2}\biggr)s^{d-2/p}(1+o(1)) \\ &= O(s^{-1/p}), \end{aligned}$$

whereas, for $1/p < d \le 2/p$, we have

$$d_{\mathrm{TV}}(D(s), \mathrm{Po}(\mathrm{E}[D(s)])) = O(s^{-d+1/p}).$$

Letting

$$e(s,t) := \mathrm{P}[D(s) > t] - \mathrm{P}[\mathrm{Po}(\mathrm{E}[D(s)]) > t],$$

it follows that, uniformly in $t \in (0,\infty)$, we have $|e(s,t)| = O(s^{-\xi})$, where $\xi = 1/p$ for $d > 2/p$ and $\xi = d - 1/p$ for $1/p < d \le 2/p$. We now rewrite S_2 as

$$S_2 = \lim_{t\to\infty} t^{1/\alpha} \int_{m(t-t^{1/2+\varepsilon})}^{\infty} (\mathrm{P}[\mathrm{Po}(\mathrm{E}[D(s)]) > t] + e(s,t))\frac{1}{p}s^{-1/p-1}\,\mathrm{d}s$$

and show that the term containing $e(s,t)$ is negligible.

Recall that

$$m(t - t^{1/2+\varepsilon}) = \biggl(\frac{t - t^{1/2+\varepsilon}}{(1+o(1))\beta}\biggr)^{p/\alpha},$$

where, here and in the remainder of this section, $o(1)$ tends to 0 as $t\to\infty$. It follows that

$$\int_{m(t-t^{1/2+\varepsilon})}^{\infty} e(s,t)s^{-1/p-1}\,\mathrm{d}s = O\biggl(\int_{m(t-t^{1/2+\varepsilon})}^{\infty} s^{-\xi-1/p-1}\,\mathrm{d}s\biggr) = O(t^{-p/\alpha(\xi+1/p)})$$

and, therefore, that

$$\lim_{t\to\infty} t^{1/\alpha}\int_{m(t-t^{1/2+\varepsilon})}^{\infty} e(s,t)s^{-1/p-1}\,\mathrm{d}s = 0.$$

We thus obtain Lemma 2.2.

It is now straightforward to show that $S_2 = \beta^{1/\alpha}$. Letting $z := \beta s^{d-1/p}/t$, whence $s = (tz/\beta)^{p/\alpha}$ and $\mathrm{E}[D(s)] = tz(1+O((tz)^{-\rho}))$ with $\rho := \rho(p,d) > 0$, we obtain, via Lemma 2.2,

$$S_2 = \lim_{t\to\infty}\frac{\beta^{1/\alpha}}{\alpha}\int_{1+o(1)}^{\infty} \mathrm{P}[\mathrm{Po}(tz(1+O((tz)^{-\rho}))) > t]z^{-1/\alpha-1}\,\mathrm{d}z.$$

The integrability of the integrand on $[1 + o(1), \infty)$ gives, for all $\gamma > 0$,

$$S_2 = \lim_{t\to\infty} \frac{\beta^{1/\alpha}}{\alpha} \int_{1+\gamma}^{\infty} \mathrm{P}[\mathrm{Po}(tz(1 + O((tz)^{-\rho}))) > t] z^{-1/\alpha-1}\, \mathrm{d}z + \gamma \cdot O(1).$$

For all $z \in [1+\gamma, \infty)$, we have $\mathrm{P}[\mathrm{Po}(tz(1+O((tz)^{-\rho}))) > t] \to 1$ as $t \to \infty$. The dominated convergence theorem yields

$$S_2 = \frac{\beta^{1/\alpha}}{\alpha} \int_1^{\infty} z^{-1/\alpha-1}\, \mathrm{d}z + \gamma O(1) = \beta^{1/\alpha} + \gamma \cdot O(1).$$

Now let $\gamma \to 0$ to obtain $S_2 = \beta^{1/\alpha}$, as desired.

3. Proof of Theorem 1.2

We prove Theorem 1.2 by showing, for all $\boldsymbol{x} \in \mathbb{Z}^d$, the existence of an event $E := E(\boldsymbol{x}) \subset \Omega$, with $\mathrm{P}[E] = 1 - o(1)$, such that on E there is a path π consisting of N edges in $G_p(\mathbb{Z}^d)$ joining $\mathbf{0}$ to $\boldsymbol{x}$, where $N \le 4(2 + \log\log|\boldsymbol{x}|)$. Here and in the sequel, $o(1)$ denotes a quantity tending to 0 as $|\boldsymbol{x}| \to \infty$.

Constructing the path π would be easy if the balls of influence at $\mathbf{0}$ and $\boldsymbol{x}$ both had radius at least $|\boldsymbol{x}|$, for then π would consist merely of the single edge $(\mathbf{0}, \boldsymbol{x})$. In general, the balls of influence at $\mathbf{0}$ and $\boldsymbol{x}$ have much smaller radii and the path π thus needs to join a sequence of balls such that consecutive balls contain each other's centers.

The heart of the proof will consist of constructing a sequence of nodes of cardinality roughly $2\log\log|\boldsymbol{x}|$ with these properties: the first node, $\mathbf{0}'$, is at distance at most $\frac{1}{2}\log\log|\boldsymbol{x}|$ from $\mathbf{0}$; the last node, $\boldsymbol{x}'$, is at distance at most $\frac{1}{2}\log\log|\boldsymbol{x}|$ from $\boldsymbol{x}$; and the edges defined by consecutive nodes are in G_p, i.e. the balls of influence at consecutive nodes contain each other's centers. Since $\mathbf{0}$ and $\mathbf{0}'$ can be joined with a path of at most $\log\log|\boldsymbol{x}|$ edges, and likewise for $\boldsymbol{x}$ and $\boldsymbol{x}'$, we can obtain a path π consisting of roughly $4\log\log|\boldsymbol{x}|$ edges. The construction of this sequence of nodes depends critically on an intermediate node, denoted here by $\boldsymbol{P}_0$, that has an unusually large ball of influence. Before defining $\mathbf{0}'$, $\boldsymbol{P}_0$, and $\boldsymbol{x}'$, we need some terminology.

For all $\boldsymbol{x} \in \mathbb{R}^d$ and $r > 0$, let $L_r^+(\boldsymbol{x})$ and $L_r^-(\boldsymbol{x})$ denote the lattice points in the upper and lower hemispheres of radius r centered at $\boldsymbol{x}$. That is, $L_r^+(\boldsymbol{x}) := B_r(\boldsymbol{x}) \cap (\mathbb{Z}^{d-1} \times \mathbb{Z}^+)$ and, similarly, $L_r^-(\boldsymbol{x}) := B_r(\boldsymbol{x}) \cap (\mathbb{Z}^{d-1} \times \mathbb{Z}^-)$. Here $\mathbb{Z}^+ := \{1, 2, \dots\}$ and $\mathbb{Z}^- := \{-1, -2, \dots\}$.

3.1. Definition of $\mathbf{0}'$, $\boldsymbol{P}_0$, and $\boldsymbol{x}'$

Throughout, we appeal to the following elementary fact. Recall that $\log s$ is short for $\log_{pd} s$.

Lemma 3.1. *Let $U_1, \dots, U_n$ be independent and identically uniformly distributed on* $[0, 1]$. *Then, for all $n > pd$, we have*

$$\min_{i\le n} U_i \le \frac{K \log n}{n}$$

with probability at least $1 - n^{-K}$.

In the sequel, we fix K to be large, with a value to be determined later.

3.1.1. *Definition of* $\mathbf{0}'$. Let $E_{\mathbf{0}} := E_{\mathbf{0}}(\boldsymbol{x})$ be the event that there is a node $z \in L^-_{(1/2)\log\log|\boldsymbol{x}|}(\mathbf{0})$ such that

$$U_z \le \frac{K \log(\log\log|\boldsymbol{x}|)^d}{(\log\log|\boldsymbol{x}|)^d}.$$

https://doi.org/10.1239/jap/1158784937 Published online by Cambridge University Press

Clearly, $E_{\mathbf{0}}$ depends only on U_z, $z \in L^-_{(1/2)\log\log|x|}(\mathbf{0})$. In case more than one node in $L^-_{(1/2)\log\log|x|}(\mathbf{0})$ satisfies the last bound, we choose z to be that node with smallest lexicographical order.

By Lemma 3.1, $\mathrm{P}[E_{\mathbf{0}}] \geq 1 - C(\log\log|x|)^{-dK}$. Given $E_{\mathbf{0}}$ we let $\mathbf{0}' := z$. Note that $\mathbf{0}'$ is random and that, since $pd > 1$, for all large $|x|$ we have

$$U_{\mathbf{0}'}^{-p} \geq 2\log\log|x|. \tag{3.1}$$

Inequality (3.1) will be important in the sequel. For now note that, since $G_p(\mathbb{Z}^d)$ connects adjacent lattice points, it follows that $d_p(y, x) \leq 2|y - x|$ for all $x, y \in \mathbb{Z}^d$, i.e. that

$$d_p(\mathbf{0}, \mathbf{0}') \leq \log\log|x|. \tag{3.2}$$

3.1.2. *Definition of* x'. Similarly, given x, with probability at least $1 - C(\log\log|x|)^{-dK}$ there is an event E_x such that there is a node $x' \in L^-_{(1/2)\log\log|x|}(x)$ on E_x with weight

$$U_{x'}^{-p} \geq 2\log\log|x|.$$

Clearly $d_p(x, x') \leq \log\log|x|$ and E_x depends only on U_z, $z \in L^-_{(1/2)\log\log|x|}(x)$.

3.1.3. *Definition of* P_0. Assume without loss of generality that the components of x have even parity, meaning that $x/2 \in \mathbb{Z}^d$. Consider the event, $E_{x/2}$, that there is a node $P_0 \in L_{|x|/10}(x/2)$ with

$$U_{P_0} \leq \frac{K\log(|x|)^d}{|x|^d}.$$

Lemma 3.1 implies that $\mathrm{P}[E_{x/2}] \geq 1 - C(|x|^{-dK})$. We note that, for large $|x|$,

$$U_{P_0}^{-p} \geq 2|x| \tag{3.3}$$

since $pd > 1$.

3.2. Construction of the path π via $\mathbf{0}'$, P_0, and x'

It will suffice to show that there is an event $E := E(x)$, with $\mathrm{P}[E(x)] = 1 - o(1)$, such that on E there are two paths, each having at most $2 + 2\lceil\log\log|x|\rceil$ edges, with one path joining P_0 to $\mathbf{0}$ and the other joining P_0 to x. It will be enough to show the existence of a path between P_0 and $\mathbf{0}$, for the method can be repeated verbatim to yield the path between P_0 and x. We first introduce some additional terminology.

We abbreviate our notation by letting $b := pd$. Note that $b > 1$ by assumption. Fix $\varepsilon \in (0, 1)$ and $x \in \mathbb{Z}^d$ with $|x|$ large. For all $j = 1, 2, \ldots$, let

$$r_j := r_j(x, \varepsilon) := |x|^{b^{-j(1-\varepsilon)}}$$

and note that $r_j \downarrow 1$ and $1 < r_j < |x|$ for all $j = 1, 2, \ldots$. We record an elementary fact.

Lemma 3.2. $r_{j+1} = r_j^{\beta(p,d,\varepsilon)}$, *where* $\beta(p, d, \varepsilon) := b^{-1+\varepsilon}$.

For all $j = 1, 2, \ldots$, consider the following disjoint 'semi-annular' regions of lattice points:

$$A_j := [(L^+_{r_j}(\mathbf{0}') - L^+_{r_{j+1}}(\mathbf{0}')) \setminus L^+_{|x|/10}(x/2)].$$

The construction of the path joining P_0 to $\mathbf{0}$ is facilitated by the following four lemmas. The first three lemmas show that, for all j, $1 \leq j \leq \lceil\log\log|x|\rceil + 1$, there are points $P_j \in A_j$ such that (P_j, P_{j-1}) and $(P_{\lceil\log\log|x|\rceil+1}, \mathbf{0}')$ belong to $G_p(\mathbb{Z}^d)$. The fourth lemma shows that this happens on an event with probability $1 - o(1)$. By consecutively linking P_j, $0 \leq j \leq \lceil\log\log|x|\rceil + 1$, and $\mathbf{0}'$, we construct a path joining P_0 to $\mathbf{0}'$ with $\lceil\log\log|x|\rceil + 2$

https://doi.org/10.1239/jap/1158784937 Published online by Cambridge University Press

edges. Since $\mathbf{0}'$ is within $\frac{1}{2}\log\log|\boldsymbol{x}|$ of $\mathbf{0}$, we need at most $\lceil\log\log|\boldsymbol{x}|\rceil$ edges to join $\mathbf{0}'$ to $\mathbf{0}$ (recall (3.2)). This gives a path joining $\boldsymbol{P}_0$ to $\mathbf{0}$ with at most $2\lceil\log\log|\boldsymbol{x}|\rceil+2$ edges. Since $2+2\lceil\log\log|\boldsymbol{x}|\rceil \le 4+2\log\log|\boldsymbol{x}|$, we obtain Theorem 1.2, as desired. We now turn to our four lemmas.

Lemma 3.3. *There exists an event* E_1, *with* $\mathrm{P}[E_1]=1-O(r_1^{-dK})$, *such that on* E_1 *there is a node* $\boldsymbol{P}_1\in A_1$ *which is linked to* $\boldsymbol{P}_0$, *i.e. the edge* $(\boldsymbol{P}_0,\boldsymbol{P}_1)$ *is in* $G_p(\mathbb{Z}^d)$.

Proof. The number of lattice points in A_1 is $\Theta(|\boldsymbol{x}|^{db^{-1+\varepsilon}})$, i.e. there is a constant $C>0$ such that the number of lattice points is bounded from above by $C(|\boldsymbol{x}|^{db^{-1+\varepsilon}})$ and bounded from below by $C^{-1}(|\boldsymbol{x}|^{db^{-1+\varepsilon}})$. Lemma 3.1 implies that there is an event E_1, depending only on $\{U_z\}_{z\in A_1}$ and with

$$\mathrm{P}[E_1]=1-O(|\boldsymbol{x}|^{-dKb^{-1+\varepsilon}}),$$

such that, for large $|\boldsymbol{x}|$, E_1 implies the existence of $\boldsymbol{P}_1\in A_1$ with

$$U_{\boldsymbol{P}_1}\le\frac{K\log(|\boldsymbol{x}|^{db^{-1+\varepsilon}})}{|\boldsymbol{x}|^{db^{-1+\varepsilon}}}.$$

Again, if there is more than one node in A_1 satisfying this inequality, we choose the one with smallest lexicographical order. Since $b:=pd$ it follows for large $|\boldsymbol{x}|$ that $\boldsymbol{P}_1$ has weight

$$U_{\boldsymbol{P}_1}^{-p}\ge\frac{|\boldsymbol{x}|^{b^{\varepsilon}}}{(K\log(|\boldsymbol{x}|^{db^{-1+\varepsilon}}))^p}\ge 2|\boldsymbol{x}|. \tag{3.4}$$

We now show that $\boldsymbol{P}_1$ is linked to $\boldsymbol{P}_0$. It suffices to show that

$$|\boldsymbol{P}_0-\boldsymbol{P}_1|\le\min(U_{\boldsymbol{P}_0}^{-p},U_{\boldsymbol{P}_1}^{-p}).$$

However, $|\boldsymbol{P}_0-\boldsymbol{P}_1|\le|\boldsymbol{P}_0|+|\boldsymbol{P}_1|\le 2|\boldsymbol{x}|$, so Lemma 3.3 follows from (3.3) and (3.4).

Given $\boldsymbol{x}$, let $m:=m(\boldsymbol{x})$ denote the largest integer such that $r_m\ge\log\log|\boldsymbol{x}|$; m is well defined since $r_j\downarrow 1$. If $t:=[1/(1-\varepsilon)]\log\log|\boldsymbol{x}|$ then

$$|\boldsymbol{x}|^{b^{-t(1-\varepsilon)}}=|\boldsymbol{x}|^{1/\log|\boldsymbol{x}|}=b,$$

showing that m is bounded by t. The next lemma extends the arguments of Lemma 3.3 and builds a path of m edges from $\boldsymbol{P}_0$ to a node $\boldsymbol{P}_m\in A_m$.

Lemma 3.4. *For all* j, $1\le j\le m$, *there is an event* E_j, *depending only on* $\{U_z\}_{z\in A_j}$, *such that*

(i) $\mathrm{P}[E_j]=1-O(r_j^{-dK})$, *and*

(ii) *on each* E_j *there is a node* $\boldsymbol{P}_j\in A_j$ *such that, on* $E_{j-1}\cap E_j$, *the edge* $(\boldsymbol{P}_{j-1},\boldsymbol{P}_j)$ *is in* G_p.

Proof. Since $\mathrm{card}\{A_j\}=\Theta(r_j^d)$, Lemma 3.1 implies that for large $|\boldsymbol{x}|$ there is an event E_j, with $\mathrm{P}[E_j]=1-O(r_j^{-dK})$, which depends only on $\{U_z\}_{z\in A_j}$ and which implies the existence of a $\boldsymbol{P}_j\in A_j$ satisfying

$$U_{\boldsymbol{P}_j}\le\frac{K\log(r_j^d)}{r_j^d}=:W_j.$$

It remains to show that

$$|\boldsymbol{P}_j - \boldsymbol{P}_{j-1}| \le \min(U_{\boldsymbol{P}_j}^{-p}, U_{\boldsymbol{P}_{j-1}}^{-p}) \tag{3.5}$$

for all j, $1 \le j \le m$. Lemma 3.3 shows (3.5) for $j = 1$. The maximal distance between points in A_j and A_{j-1} is at most twice r_{j-1}, i.e. $|\boldsymbol{P}_j - \boldsymbol{P}_{j-1}| \le 2r_{j-1}$. It thus suffices to show that

$$2r_{j-1} \le \min(W_j^{-p}, W_{j-1}^{-p}) = W_j^{-p}, \tag{3.6}$$

which holds since $W_{j-1}^{-p} \ge W_j^{-p}$ for all j, $1 \le j \le m$.

However, by Lemma 3.2,

$$W_j^{-p} = \frac{r_j^{pd}}{(Kd \log r_j)^p} = \frac{((r_{j-1})^{b^{-1+\varepsilon}})^{pd}}{(Kdb^{-1+\varepsilon} \log(r_{j-1}))^p}.$$

Thus, for all j, $1 \le j \le m$,

$$\frac{W_j^{-p}}{r_{j-1}} = \frac{(r_{j-1})^{b^\varepsilon - 1}}{(Kdb^{-1+\varepsilon} \log(r_{j-1}))^p} \ge \frac{(r_m)^{b^\varepsilon - 1}}{(Kdb^{-1+\varepsilon} \log(r_m))^p},$$

since the r_j are decreasing. By definition of r_m and since $b^\varepsilon - 1 > 0$, the last ratio clearly exceeds 2 for large $|\boldsymbol{x}|$, showing (3.6) and completing the proof of Lemma 3.4.

The next lemma shows that we may link $\boldsymbol{P}_m$ and $\boldsymbol{0}'$ via a node $\boldsymbol{P}_{m+1} \in A_{m+1}$. Combined with Lemmas 3.2 and 3.3, this builds a path between $\boldsymbol{P}_0$ and $\boldsymbol{0}'$ which contains $m + 2$ edges.

Lemma 3.5. *There is an event E_{m+1}, depending only on $\{U_z\}_{z \in A_{m+1}}$, such that* $\mathrm{P}[E_{m+1}] = 1 - O(r_{m+1}^{-dK})$, *and on $E_{\boldsymbol{0}} \cap E_m \cap E_{m+1}$ there is a point $\boldsymbol{P}_{m+1} \in A_{m+1}$ such that the edges $(\boldsymbol{P}_m, \boldsymbol{P}_{m+1})$ and $(\boldsymbol{P}_{m+1}, \boldsymbol{0}')$ both belong to $G_p(\mathbb{Z}^d)$.*

Proof. First, by definition of m and by Lemma 3.2 we have

$$(\log\log|\boldsymbol{x}|)^\beta \le r_m^\beta = r_{m+1} < \log\log|\boldsymbol{x}|.$$

By Lemma 3.1, for large $|\boldsymbol{x}|$ there is an event E_{m+1}, with $\mathrm{P}[E_{m+1}] = 1 - O(r_{m+1}^{-dK})$, which depends only on $\{U_z\}_{z \in A_{m+1}}$ and which implies the existence of a point $\boldsymbol{P}_{m+1} \in A_{m+1}$ with

$$U_{\boldsymbol{P}_{m+1}} \le \frac{K \log(r_{m+1}^d)}{r_{m+1}^d} \le \frac{K \log(\log\log|\boldsymbol{x}|)^d}{(\log\log|\boldsymbol{x}|)^{\beta d}} \le \frac{K \log(\log\log|\boldsymbol{x}|)^d}{(\log\log|\boldsymbol{x}|)^{(pd)^\varepsilon/p}}$$

since $\beta d = (pd)^\varepsilon/p$. Since $(pd)^\varepsilon > 1$, it follows that, for large $|\boldsymbol{x}|$, on E_{m+1} we have

$$U_{P_{m+1}}^{-p} \ge 2\log\log|\boldsymbol{x}|. \tag{3.7}$$

Following the arguments in the proof of Lemma 3.4 (with j equal to $m + 1$ there), we find that, on $E_m \cap E_{m+1}$, $(\boldsymbol{P}_m, \boldsymbol{P}_{m+1})$ is an edge in $G_p(\mathbb{Z}^d)$. Furthermore, on $E_{\boldsymbol{0}} \cap E_m \cap E_{m+1}$, the edge $(\boldsymbol{P}_{m+1}, \boldsymbol{0}')$ belongs to $G_p(\mathbb{Z}^d)$ if and only if

$$|\boldsymbol{0}' - \boldsymbol{P}_{m+1}| \le \min(U_{\boldsymbol{0}'}^{-p}, U_{\boldsymbol{P}_{m+1}}^{-p}). \tag{3.8}$$

However,

$$|\boldsymbol{0}' - \boldsymbol{P}_{m+1}| \le |\boldsymbol{0}' - \boldsymbol{0}| + |\boldsymbol{0} - \boldsymbol{P}_{m+1}| \le \log\log|\boldsymbol{x}| + r_{m+1} \le 2\log\log|\boldsymbol{x}|,$$

showing that (3.8) follows, using (3.7) and (3.1).

The last lemma completes the proof of Theorem 1.2.

Lemma 3.6. *For all* $\boldsymbol{x} \in \mathbb{Z}^d$*, there is an event* $E(\boldsymbol{x})$*, with* $\mathrm{P}[E(\boldsymbol{x})] = 1 - o(1)$*, such that on* $E(\boldsymbol{x})$ *there exists a path joining* $\boldsymbol{P}_0$ *to* $\boldsymbol{0}$ *with* $4 + 2\log\log|\boldsymbol{x}|$ *edges.*

Proof. Let $E(\boldsymbol{x}) := E_{\boldsymbol{0}} \cap E_{\boldsymbol{x}/2} \cap (\bigcap_{j=1}^{m+1} E_j)$. On $E(\boldsymbol{x})$ we have shown that there is a path, π, joining $\boldsymbol{P}_0$ to $\boldsymbol{0}$ via the successive nodes $\boldsymbol{P}_1, \boldsymbol{P}_2, \dots, \boldsymbol{P}_m, \boldsymbol{P}_{m+1}, \boldsymbol{0}', \boldsymbol{0}$. The number of edges in π is bounded by $m + 2 + \lceil \log\log|\boldsymbol{x}| \rceil$, where $\lceil \log\log|\boldsymbol{x}| \rceil$ denotes an upper bound on the number of edges between $\boldsymbol{0}'$ and $\boldsymbol{0}$. Since ε is arbitrary in the definition of t, it follows that $m \le \lceil \log\log|\boldsymbol{x}| \rceil$. Thus, $\mathrm{card}\{\pi\} \le 4 + 2\log\log|\boldsymbol{x}|$.

Finally, we show that $\mathrm{P}[E(\boldsymbol{x})] = 1 - o(1)$. For all j, $1 \le j \le m+1$, E_j depends only on $\{U_z\}_{z \in A_j}$ and, since the A_j are disjoint, the $\{E_j\}_{1 \le j \le m+1}$ are independent. Clearly, since $E_{\boldsymbol{0}}$ depends on $\{U_z\}_{z \in \mathbb{Z}^{d-1} \times \mathbb{Z}^-}$, we have independence of $E_{\boldsymbol{0}}, E_1, E_2, \dots, E_{m+1}$. Similarly, $E_{\boldsymbol{x}/2}, E_{\boldsymbol{0}}, E_1, E_2, \dots, E_{m+1}$ are independent.

By independence, we have

$$\mathrm{P}[E(\boldsymbol{x})] = \mathrm{P}\Bigg[\bigcap_{j=1}^{m+1} E_j\Bigg] \mathrm{P}[E_{\boldsymbol{0}}]\,\mathrm{P}[E_{\boldsymbol{x}}]\,\mathrm{P}[E_{\boldsymbol{x}/2}] = (1 - o(1))^3 \prod_{j=1}^{m+1} \mathrm{P}[E_j].$$

Now, m is bounded by $C \log\log|\boldsymbol{x}|$ and the definition of r_m shows, for large K, that $m r_{m+1}^{-dK} \to 0$ as $|\boldsymbol{x}| \to \infty$. Since $1 - 2s \le \exp(-s) \le 1 - s/2$ for small, positive s, it follows that

$$\begin{aligned}
\prod_{j=1}^{m+1} \mathrm{P}[E_j] &= \prod_{j=1}^{m+1} (1 - O(r_j^{-dK})) \\
&\ge \exp\Bigg(-C \sum_{j=1}^{m+1} r_j^{-dK}\Bigg) \\
&\ge 1 - C \sum_{j=1}^{m+1} r_j^{-dK} \\
&\ge 1 - C \sum_{j=1}^{m+1} r_{m+1}^{-dK}.
\end{aligned}$$

This yields $\mathrm{P}[E(\boldsymbol{x})] = 1 - o(1)$, as desired, completing the proof of Lemma 3.6.

4. Proof of Theorem 1.3

Assume without loss of generality that n has even parity. Partition $[0, n]^d \cap \mathbb{Z}^d$ into $Q_1 := [0, \frac{1}{2}n] \times [0, n]^{d-1} \cap \mathbb{Z}^d$ and $Q_2 := (\frac{1}{2}n, n] \times [0, n]^{d-1} \cap \mathbb{Z}^d$. For all $k = 0, 1, 2, \dots, n/2$, write $Q_{1,k} := \{n/2 - k\} \times [0, n]^{d-1} \cap \mathbb{Z}^d$ and note that $Q_1 = \bigcup_{k=0}^{n/2} Q_{1,k}$.

The number of nodes in Q_1 whose balls of influence have nonempty intersection with Q_2 is

$$N := \sum_{k=0}^{n} \sum_{i \in Q_{1,k}} 1_{\{U_i^{-p} \ge k+1\}}.$$

If we remove these N nodes from Q_1 then $G_p(Q_1)$ and $G_p(Q_2)$ are disconnected, i.e. the graphs have no edges joining them. Moreover, as the number of nodes in $Q_{1,k}$ equals n^{d-1}, we obtain

$$\mathrm{E}[N] = \sum_{k=0}^{n} n^{d-1}\,\mathrm{P}[U_{\mathbf{0}}^{-p} \geq k+1] = n^{d-1}\sum_{k=0}^{n}(k+1)^{-1/p} \leq Cn^{d-1}[n^{1-1/p} \vee 1],$$

which is exactly the desired upper bound.

Acknowledgements

I thank an anonymous referee for helpful comments and for pointing out an error in the original proof of Theorem 1.1. I also thank Mathew Penrose for helpful conversations on power-law graphs.

References

[1] Albert, R. and Barabási, A.-L. (2002). Statistical mechanics of complex networks. *Rev. Modern Physics* **74,** 47–97.
[2] Barabási, A.-L. (2002). *Linked: The New Science of Networks*. Perseus, Cambridge, MA.
[3] Barbour, A. D., Holst, L. and Janson, S. (1992). *Poisson Approximation*. Oxford University Press.
[4] Benjamini, I., Kesten, H., Peres, Y. and Schramm, O. (2004). The geometry of the uniform spanning forests: transitions in dimensions 4, 8, 12, ... *Ann. Math.* **160,** 465–491.
[5] Berger, N. (2004). A lower bound for the chemical distance in sparse long-range percolation models. Preprint 0409021, Department of Mathematics, University of California, Davis. Available at http://front.math.ucdavis.edu.
[6] Berger, N. *et al.* (2003). Degree distribution of the FKP network model. In *Automata, Languages and Programming* (Lecture Notes Comput. Sci. **2719**), Springer, Heidelberg, pp. 725–738.
[7] Biskup, M. (2004). Graph diameter in long-range percolation. Submitted.
[8] Biskup, M. (2004). On scaling of the chemical distance in long-range percolation models. *Ann. Prob.* **32,** 2938–2977.
[9] Bollabás, B. and Riordan, O. M. (2003). The diameter of a scale-free random graph. *Combinatorica* **24,** 5–34.
[10] Chung, F. and Lu, L. (2002). The average distance in a random graph with given expected degrees. *Proc. Nat. Acad. Sci. USA* **99,** 15879–15882.
[11] Chung, F. and Lu, L. (2003). The average distance in a random graph with given expected degrees. *Internet Math.* **1,** 91–113. (This is an expanded version of [10].)
[12] Cohen, R. and Havlin, S. (2003). Scale-free networks are ultrasmall. *Phys. Rev. Lett.* **90,** 058701.
[13] Coppersmith, D., Gamarnik, D. and Sviridenko, M. (2002). The diameter of a long-range percolation graph. *Random Structures Algorithms* **21,** 1–13.
[14] Dorogovtsev, S. N. and Mendes, J. F. F. (2003). *Evolution of Networks: From Biological Nets to the Internet and World Wide Web*. Oxford University Press.
[15] Dudley, R. M. (1999). *Uniform Central Limit Theorems*. Cambridge University Press.
[16] Fabrikant, A., Koutsoupias, E. and Papadimitriou, C. H. (2002). Heuristically optimized trade-offs: a new paradigm for power laws in the internet. In *Automata, Languages and Programming* (Lecture Notes Comput. Sci. **2380**), Springer, Berlin, pp. 110–122.
[17] Franceschetti, M. and Meester, R. (2004). Navigation in small world networks, a scale-free continuum model. Tech. Rep. UCB/ERL M03/33, EECS Department, University of California, Berkeley.
[18] Hermann, C., Barthélemy, M. and Provero, P. (2003). Connectivity distribution of spatial networks. *Phys. Rev. E* **68,** 026128.
[19] Kleinberg, J. M. (2000). Navigation in the small world. *Nature* **406,** 845.
[20] Meester, R. and Roy, R. (1996). *Continuum Percolation* (Camb. Tracts Math. **119**). Cambridge University Press.
[21] Newman, M. E. J. (2003). The structure and function of complex networks. *SIAM Rev.* **45,** 167–256.
[22] Penrose, M. D. (2001). A spatial central limit theorem with applications to percolation, epidemics and Boolean models. *Ann. Prob.* **29,** 1515–1546.
[23] Penrose, M. D. (2003). *Random Geometric Graphs*. Clarendon Press, Oxford.
[24] Watts, D. J. (1999). *Small Worlds*. Princeton University Press.
[25] Watts, D. J. and Strogatz, S. H. (1998). Collective dynamics of 'small world' networks. *Nature* **393,** 440–442.