Ultra-small scale-free geometric networks

J. E. Yukich[1]

May 3, 2006

Abstract

We consider a family of long-range percolation models $(G_p)_{p>0}$ on $\mathbb{Z}^d$ allowing dependence between edges and having these 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 x and y is bounded by a multiple of $\log_{pd} \log_{pd} |x - y|$ 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 disconnected large subgraphs.

1 Introduction

The statistical properties of large networks have received considerable attention in the recent scientific literature [2, 16, 23, 27]. 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, 23].

Networks often have a geometric component to them where the vertices have positions in space and where geographic proximity plays a role in deciding which vertices get connected. In this context random geometric graphs are a natural alternative to the classical Erdös-Rényi random graph models. Random connection models [22] provide one way to describe networks with spatial content. In these models the event $E_{x,y}$ of a connection between points x and y has probability

American Mathematical Society 2000 subject classifications. Primary 60D05, Secondary 05C80

Key words and phrases: scale free graphs, long-range percolation, chemical distance

[1] Research supported in part by NSF Grant DMS-0203720

$p_{x,y} := P[E_{x,y}] = g(|x-y|)$, where $g : \mathbb{R}^+ \to [0,1]$ is a connection function and where $|x|$ denotes the Euclidean norm of x. The standard long-range percolation model assumes independence of $E_{x,y}$ and $E_{x,u},\ y \neq u$, 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 admits 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 x and y is bounded by a multiple of $\log_{pd} \log_{pd} |x-y|$ with probability $1 - o(1)$ where $o(1)$ denotes a quantity tending to 0 as $|x-y| \to \infty$. Ultra-small graph distances imply efficiency, are consistent with the 'small world phenomenon' [2, 16, 26, 27], and are relevant in the context of routing, 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 having nearly one half the total network nodes.

1.1 A general dependent random connection model

Let $\{U_z\}_{z \in \mathbb{Z}^d}$ be i.i.d. 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 both nodes are 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

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

Let $\delta = 1$. By independence of U_z, we have $p_{x,y} := P[E_{x,y}] = |x-y|^{-2/p}$, showing that the probability of 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 $U_z^{-p} \geq 1$ 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 have arbitrary large diameter with arbitrarily large probability. Moreover, in accordance with their Poisson Boolean model counterparts (cf. [22]), 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

$D_p(\mathbf{0})$ denotes the degree of the origin in $G_p(\mathbb{Z}^d)$, ω_d is the volume of the unit radius ball in $\mathbb{R}^d$, and $\alpha := pd - 1$. 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 $d = 1, 2, ...$ and all $p \in (1/d, \infty)$*

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

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.[14]) 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 of four in this result can be improved but have not tried for the sharpest bound.

Theorem 1.2 *($G_p(\mathbb{Z}^d)$ has small graph distance) For all $d = 1, 2, ...$ 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 zero as $|x| \to \infty$.

The network failure of $G_p(\mathbb{Z}^d)$ is easily quantified:

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

In particular, Theorem 1.3 implies 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 $p_{x,y} := P[E_{x,y}] = |x-y|^{-s+o(1)}$ as $|x-y| \to \infty$ for some constant $s \in (0, \infty)$; $E_{x,y}$ and $E_{x,u}$ are independent for all $x, y, u \in \mathbb{Z}^d$. When $s \in (0, d)$, Benjamini et al. [7] show that the graph distance $d(\mathbf{0}, x)$ behaves like the constant $\lceil \frac{s}{d-s} \rceil$ as

$|x| \to \infty$. When $s = d$, Coppersmith et al. [15] show that $d(\mathbf{0}, x)$ scales as $\log |x| / \log\log |x|$, whereas for $s \in (d, 2d)$ Biskup [10, 11] shows that $d(\mathbf{0}, x)$ scales as $(\log |x|)^{\Delta + o(1)}$, where $\Delta := \Delta(s, d) := \log 2 / \log(2d/s)$. The case $s = 2d$ is open and for $s \in (2d, \infty)$, $d(\mathbf{0}, x)$ scales at least linearly in $|x|$, as shown by Berger [8]. 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 [21] proposes a lattice model where long range contacts are added in a biased way with 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 $f : \mathbb{R}^d \to \mathbb{R}^+$ 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 [25], joins two nodes in $\mathcal{P}_f$ whenever their Euclidean distance is less than a specified cut-off. Herrman et al. [20] show that if $\int_{\mathbb{R}^d} f^r(x)dx = \infty$ for all $r > r_0$, then the degree distribution is effectively a power law (sect. II.B of [20]).

b. The *on-line nearest neighbors graph* is defined on randomly ordered point sets in $\mathbb{R}^d$ and it 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 [9, 18].

c. Franceschetti and Meester [19] develop a scale-free continuum model but do not obtain iterated log bounds on interpoint graph distances.

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

3. *Power exponents* $q \in (2, 3)$. Consider a random graph on n nodes $v_1, v_2, ..., v_n$ with weight (expected degree) w_i at node v_i. Nodes v_i and v_j are connected with probability $\rho w_i w_j$, where $\rho = (\sum_{i=1}^n w_i)^{-1}$. Chung and Lu [13] provide conditions on the weights under which the degree distribution is proportional to k^{-q}, $q \in (2, 3)$ and $k \in \mathbb{Z}$, the average distance between nodes is a.s. $O(\log\log n)$, and the diameter is $O(\log n)$. In unrelated work, Cohen and Havlin [14] argue that whenever the degree distribution of a random graph on n vertices is proportional to k^{-q}, where $q \in (2, 3)$ and where k is restricted to (m, K), where m and $K := K(n)$ are well-defined 'cut-offs', 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 vertex. Thus nodes of high degree tend to acquire more new links than nodes of low degree.

Barabási and Albert [1] show that such models follow a power law, are not geometry dependent, and in general are not ultra-small [12].

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

6. *Further connectivity results.* Theorems 1.1 - 1.3 describe 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 coeffficients) on $G_p(\mathbb{Z}^d \cap [0,n]^d)$ is simplified by appealing to the stabilization properties of G_p (see especially [24]). $G_p(\mathbb{Z}^d)$ is *assortative* in that high degree nodes tend to link to high degree nodes whereas 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 scale-free ultra-small graphs and those which are not.

2 Proof of Theorem 1.1

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

$$D(u^{-p}) = \sum_{y \in L_{u^{-p}}(\mathbf{0}),\ y \neq \mathbf{0}} \mathbf{1}_{U_y \leq |y|^{-1/p}}.$$

Thus to prove Theorem 1.1 we condition on U_0 and show

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

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

Lemma 2.1 *We have for all $p \in (1/d, \infty)$*

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

where the error on the right hand side of (2.2) is for $u \to 0^+$.

Proof. Note that $\mathbb{E}\, D(u^{-p})$ is approximated by $\int_{|x| \le u^{-p}} |x|^{-1/p} dx = d\omega_d \int_0^{u^{-p}} t^{d-1-1/p} dt = \beta u^{-\alpha}$. 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 $\left|\mathbb{E}\, D(u^{-p}) - \int_{|x| \le u^{-p}} |x|^{-1/p} dx\right|$ is bounded by the sum of the following three errors:

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

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

and

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

Now

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

and thus E_1 is bounded by the product of $\mathrm{card}[(B_{u^{-p}}(\mathbf{0}) \setminus R(u)) \cap \mathbb{Z}^d]$ and $\sup_{y \in (B_{u^{-p}}(\mathbf{0}) \setminus R(u)) \cap \mathbb{Z}^d} |y|^{-1/p}$. Since the first factor is bounded by $Cu^{-p(d-1)}$ and the second by Cu, it follows that $E_1 \le Cu^{-pd+p+1}$. Similar methods show $E_3 \le Cu^{-pd+p+1}$.

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

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

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. $\square$

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

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

We note that (2.3) is plausible since Lemma 2.1 suggests that $P[D(s) > t]$ is close to one for $t << \beta s^{\alpha/p}$ and close to zero for $t >> \beta s^{\alpha/p}$, indicating that the left hand side of (2.3) behaves as

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

To put this heuristic argument on rigorous footing, we will 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 $m(v) := \sup\{s : \mathbb{E}\, D(s) \le v\}$. Lemma 2.1 implies

$$\mathbb{E}\, D(s) = \beta s^{\alpha/p} + O\left(\max(1, s^{d-1-1/p})\right) = \beta s^{\alpha/p}\left(1 + \max(O(s^{1/p-d}), O(s^{-1}))\right). \tag{2.4}$$

It follows that for v large

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

where $o(1)$ tends to zero as $v \to \infty$.

Given $t \ge \beta$ and $\varepsilon \in (0, 1/2)$ fixed, define the following two integration domains:

$$I_1 := \left[1,\ m(t - t^{1/2+\varepsilon})\right),$$

and

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

Rewrite the left-hand side of (2.3) as

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

provided that both limits exist.

To prove Theorem 1.1 it suffices to show $S_1 = 0$ and $S_2 = \beta^{1/\alpha}$. We first show $S_1 = 0$. Bernstein's inequality [17] for sums of independent bounded random variables yields for all $s \in I_1$

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

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

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

It follows that

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

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

Lemma 2.2 *We have for all $p \in (1/d, \infty)$*

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

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

$$d_{TV}\left(D(s), \mathrm{Po}(\mathbb{E}\, D(s)\right) \le \left(\sum_{y\in L_s(\mathbf{0}),\ y\ne\mathbf{0}} p_y\right)^{-1} \sum_{y\in L_s(\mathbf{0}),\ y\ne\mathbf{0}} p_y^2.$$

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

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

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

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

It follows by Lemma 2.1 that for $d > 2/p$

$$d_{TV}\left(D(s), \mathrm{Po}(\mathbb{E}\, D(s))\right) \le \left(\beta s^{d-1/p}(1+o(1))\right)^{-1} \beta s^{d-2/p}(1+o(1)) = O(s^{-1/p})$$

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

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

Letting

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

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 rewrite S_2 as

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

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

Recall that $m(t-t^{1/2+\varepsilon}) = \left(\frac{t-t^{1/2+\varepsilon}}{(1+o(1))\beta}\right)^{p/\alpha}$ where here and in the remainder of this section $o(1)$ tends to zero as $t \to \infty$. It follows that

$$\int_{m(t-t^{1/2+\varepsilon})}^{\infty} e(s,t) s^{-1/p-1} ds = O\left(\int_{m(t-t^{1/2+\varepsilon})}^{\infty} s^{-\xi-1/p-1} ds\right) = O(t^{-p/\alpha(\xi+1/p)})$$

and therefore

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

We thus obtain Lemma 2.2. □

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

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

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

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

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

$$S_2 = \frac{\beta^{1/\alpha}}{\alpha} \int_{1}^{\infty} z^{-1/\alpha-1} dz \ + \gamma \cdot O(1) = \beta^{1/\alpha} + \gamma \cdot O(1).$$

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

Constructing the path π would be easy if the balls of influence at $\mathbf{0}$ and x both had radius at least $|x|$, for then π would consist merely of the single edge $(\mathbf{0}, x)$. In general the balls of influence at $\mathbf{0}$ and x have much smaller radius 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 $2\log\log|x|$ with these properties: the first node $\mathbf{0}'$ is distant at most $\frac{1}{2}\log\log|x|$ from $\mathbf{0}$, the last node x' is distant at most $\frac{1}{2}\log\log|x|$ from x, and 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 $\log\log|x|$ edges and likewise with x and x', we can obtain a path π consisting of roughly $4\log\log|x|$ edges. The construction of this sequence of nodes depends critically on an intermediate node, denoted here by P_0, and having an unusually large ball of influence. Before defining $\mathbf{0}', P_0$, and x' we need some terminology.

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

3.1 Definition of $\mathbf{0}', P_0$, and x'

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

Lemma 3.1 *Let $U_1, ..., U_n$ be i.i.d. uniform 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 large, with a value to be determined later.

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

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

Clearly, $E_{\mathbf{0}}$ depends only on $U_z,\ z \in L^-_{\frac{1}{2}\log\log|x|}(\mathbf{0})$.

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

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

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 $d_p(y,x) \le 2|y-x|$ for all $x, y \in \mathbb{Z}^d$, i.e.,

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

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

$$U_{x'}^{-p} \geq 2\log\log|x|. \tag{3.3}$$

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

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

$$U_{P_0} \leq \frac{K\log(|x|)^d}{|x|^d}. \tag{3.4}$$

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

$$U_{P_0}^{-p} \geq 2|x|. \tag{3.5}$$

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

It will suffice to show that there is an event $E := E(x)$, $P[E(x)] = 1 - o(1)$, 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 a second path between P_0 and x. We first introduce some additional terminology.

Abbreviate notation and put $b := pd$. Note $b > 1$ by assumption. Fix $\varepsilon \in (0, 1)$ and $x \in \mathbb{Z}^d$, $|x|$ large. For all $j = 1, 2, ...$ set

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

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

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

Consider for all $j = 1, 2, ...$ the following disjoint 'semi-annular' regions of lattice points:

$$A_j := \left[\left(L^+_{r_j}(\mathbf{0}') - L^+_{r_{j+1}}(\mathbf{0}')\right) \setminus L^+_{|x|/10}(x/2)\right]. \tag{3.6}$$

The construction of the path joining P_0 to $\mathbf{0}$ is facilitated with the following four lemmas. The first three lemmas show for all $1 \leq j \leq \lceil\log\log|x|\rceil + 1$, that 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 $P_j,\ 0 \le j \le \lceil \log\log|x|\rceil + 1$, and $\mathbf{0}'$, we construct a path joining P_0 to $\mathbf{0}'$ with $\lceil \log\log|x|\rceil + 2$ edges. Since $\mathbf{0}'$ is within $\frac{1}{2}\log\log|x|$ of $\mathbf{0}$, we need at most $\lceil \log\log|x|\rceil$ edges to join $\mathbf{0}'$ to $\mathbf{0}$ (recall (3.2)). This gives a path joining P_0 to $\mathbf{0}$ with at most $2\lceil \log\log|x|\rceil + 2$ edges. Since $2 + 2\lceil \log\log|x|\rceil \le 4 + 2\log\log|x|$ we obtain Theorem 1.2 as desired. We now turn to our four key lemmas.

Lemma 3.3 *There exists an event E_1 with $P[E_1] = 1 - O(r_1^{-dK})$, such that on E_1 there is a node $P_1 \in A_1$ which is linked to P_0, i.e., the edge (P_0, P_1) is in $G_p(\mathbb{Z}^d)$.*

Proof. The number of lattice points in A_1 is $\Theta\left(|x|^{db^{-1+\varepsilon}}\right)$. Lemma 3.1 implies that there is an event E_1 depending only on $\{U_z\}_{z\in A_1}$, with

$$P[E_1] = 1 - O\left(|x|^{-dKb^{-1+\varepsilon}}\right) \tag{3.7}$$

such that for $|x|$ large E_1 implies the existence of $P_1 \in A_1$ with

$$U_{P_1} \le \frac{K\log\left[|x|^{db^{-1+\varepsilon}}\right]}{|x|^{db^{-1+\varepsilon}}}.$$

Since $b := pd$ it follows for $|x|$ large that P_1 has weight

$$U_{P_1}^{-p} \ge \frac{|x|^{b^{\varepsilon}}}{(K\log\left[|x|^{db^{-1+\varepsilon}}\right])^p} \ge 2|x|. \tag{3.8}$$

We now show that P_1 is linked to P_0. It suffices to show

$$|P_0 - P_1| \le \min(U_{P_0}^{-p}, U_{P_1}^{-p}).$$

But $|P_0 - P_1| \le |P_0| + |P_1| \le 2|x|$ and Lemma 3.3 follows by (3.5) and (3.8). $\square$

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

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

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 P_0 to a node $P_m \in A_m$.

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

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

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

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

$$U_{P_j} \leq \frac{K\log[r_j^d]}{r_j^d} := W_j, \tag{3.9}$$

i.e., (i) holds.

Since (i) holds, it remains to show (ii), i.e., to show

$$|P_j - P_{j-1}| \leq \min(U_{P_j}^{-p}, U_{P_{j-1}}^{-p}) \tag{3.10}$$

for all $1 \leq j \leq m$. Lemma 3.3 shows (3.10) for $j = 1$. The maximal distance between points in A_j and A_{j-1} is at most twice r_{j-1}, i.e, $|P_j - P_{j-1}| \leq 2r_{j-1}$. So it suffices to show

$$2r_{j-1} \leq \min(W_j^{-p}, W_{j-1}^{-p}) = W_j^{-p} \tag{3.11}$$

since $W_{j-1}^{-p} \geq W_j^{-p}$ for all $1 \leq j \leq m$.

However, by Lemma 3.2

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

Thus for all $1 \leq j \leq m$

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

since r_j are decreasing. By definition of r_m and since $b^\varepsilon - 1 > 0$, the last ratio clearly exceeds 2 for $|x|$ large, showing (3.11) and completing Lemma 3.4. □

The next lemma shows that we may link P_m and $\mathbf{0}'$ via a node $P_{m+1} \in A_{m+1}$. Combined with Lemmas 3.2 and 3.3, this builds a path between P_0 and $\mathbf{0}'$ of $m + 2$ edges.

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

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

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

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

$$U_{P_{m+1}} \leq \frac{K \log[r_{m+1}^d]}{r_{m+1}^d} \leq \frac{K \log(\log\log|x|)^d}{(\log\log|x|)^{\beta d}} \leq \frac{K \log(\log\log|x|)^d}{(\log\log|x|)^{(pd)^\varepsilon \frac{1}{p}}}$$

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

$$U_{P_{m+1}}^{-p} \geq 2 \log\log|x|. \tag{3.12}$$

Following the arguments of Lemma 3.4 (with j equal to $m+1$ there), we find that on $E_m \cap E_{m+1}$, (P_m, P_{m+1}) is an edge in $G_p(\mathbb{Z}^d)$. Furthermore, on $E_{\mathbf{0}} \cap E_m \cap E_{m+1}$, the edge $(P_{m+1}, \mathbf{0}')$ belongs to $G_p(\mathbb{Z}^d)$ iff

$$|\mathbf{0}' - P_{m+1}| \leq \min(U_{\mathbf{0}'}^{-p}, U_{P_{m+1}}^{-p}). \tag{3.13}$$

However,

$$|\mathbf{0}' - P_{m+1}| \leq |\mathbf{0}' - \mathbf{0}| + |\mathbf{0} - P_{m+1}| \leq \log\log|x| + r_{m+1} \leq 2\log\log|x|,$$

showing that (3.13) follows by (3.12) and (3.1). □

The last lemma completes the proof of Theorem 1.2.

Lemma 3.6 *For all $x \in \mathbb{Z}^d$ there is an event $E(x), P[E(x)] = 1 - o(1)$, such that on $E(x)$ there exists a path joining P_0 to $\mathbf{0}$ with $4 + 2\log\log|x|$ edges.*

Proof. Put $E(x) := E_{\mathbf{0}} \cap E_{x/2} \cap \left(\cap_{j=1}^{m+1} E_j\right)$. On $E(x)$ we have shown that there is a path π joining P_0 to $\mathbf{0}$ via the successive nodes $P_1, P_2, ..., P_m, P_{m+1}, \mathbf{0}', \mathbf{0}$. The number of edges in π is bounded by $m + 2 + \lceil \log\log|x| \rceil$, where $\lceil \log\log|x| \rceil$ denotes an upper bound on the number of edges between $\mathbf{0}'$ and $\mathbf{0}$. Since ε is arbitrary in the definition of t it follows that $m \leq \lceil \log\log|x| \rceil$. Thus $\mathrm{card}\pi \leq 4 + 2\log\log|x|$.

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

By independence

$$P[E(x)] = P[\cap_{j=1}^{m+1} E_j] \cdot P[E_{\mathbf{0}}] \cdot P[E_x] \cdot P[E_{x/2}] = (1 - o(1))^3 \Pi_{j=1}^{m+1} P[E_j].$$

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

$$\Pi_{j=1}^{m+1} P[E_j] = \Pi_{j=1}^{m+1}(1 - O(r_j^{-dK})) \ge \exp\left(-C\sum_{j=1}^{m+1} r_j^{-dK}\right)$$

$$\ge 1 - C\sum_{j=1}^{m+1} r_j^{-dK} \ge 1 - C\sum_{j=1}^{m+1} r_{m+1}^{-dK}.$$

This yields $P[E(x)] = 1 - o(1)$ 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 $k = 0, 1, 2, ..., n/2$, write $Q_{1,k} := \{n/2 - k\} \times [0,n]^{d-1} \cap \mathbb{Z}^d$ and note that $Q_1 = \cup_{k=0}^{\frac{n}{2}} Q_{1,k}$.

The number of nodes in Q_1 whose balls of influence have non-empty intersection with Q_2 is

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

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

$$\mathbb{E}\, N = \sum_{k=0}^{n} n^{d-1} P[U_{\mathbf{0}}^{-p} \ge k+1] = n^{d-1} \sum_{k=0}^{n} (k+1)^{-1/p} \le C n^{d-1}[n^{1-1/p} \vee 1],$$

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] R. Albert and A.-L. Barabási (2002), Statistical mechanics of complex networks, *Reviews of Modern Physics*, **74**, 47-97.

[2] A.-L. Barabási (2002), Linked: The New Science of Networks, Perseus, Cambridge, Mass.

[3] A. D. Barbour, L. Holst, and S. Janson (1992), Poisson Approximation, Oxford University Press.

[4] A. D. Barbour and G. Reinert (2001), Small worlds, *Random Structures and Algorithms*, **19**, 54-74.

[5] M. Barthélemy (2003), Crossover from scale-free to spatial networks, arXiv:cond-mat/0212086 v3 (22 Sept, 2003).

[6] I. Benjamini and N. Berger (2001), The diameter of long-range percolation clusters on finite cycles, *Random Structures and Algorithms*, **19**, 102-111.

[7] I. Benjamini, H. Kesten, Y. Peres, and O. Schramm (2005), The geometry of the uniform spanning forests: transitions dimensions $4, 8, 12, ...$, *Ann. Math.*, to appear.

[8] N. Berger (2004), A lower bound for the chemical distance in sparse long-range percolation models, preprint.

[9] N. Berger, B. Bollabás, C. Borgs, J. Chayes, and O. Riordan (2003), Degree distribution of the FKP model, *Automata, Languages and Programming: 30th Int. Colloquium, ICALP 2003, Lecture Notes in Computer Science* **2719**, Springer, Heidelberg.

[10] M. Biskup (2004), On scaling of the chemical distance in long-range percolation models, *Ann. Prob.*, **32**, 4, 2938-2977.

[11] M. Biskup (2004), Graph diameter in long-range percolation, preprint.

[12] B. Bollabás and O. M. Riordan (2003), The diameter of a scale-free graph, to appear in Combinatorica.

[13] F. Chung and L. Lu (2002, 2003), The average distance in a random graph with given expected degree.*Proc. National Academy of Sciences*, 99, no. 25 (December, 2002), 15879-15882; (long version), Internet Mathematics, 1 (2003), 91-114.

[14] R. Cohen and S. Havlin (2003), Scale-free networks are ultrasmall, *Phys. Rev. Letters*, **90**, 5. (90.058701)

[15] D. Coppersmith, D. Gamarnik and M. Sviridenko (2002), The diameter of a long-range percolation graph, *Random Structures and Algorithms*, **21**, 1-13.

[16] S. N. Dorogovtsev and J. F. F. Mendes (2003), Evolution of Networks: From Biological Nets to the Internet and World Wide Web, Oxford University Press, Oxford.

[17] R. M. Dudley (1999), Uniform Central Limit Theorems, Cambridge University Press.

[18] A. Fabrikant, E. Koutsoupias, and C. M. Papadimitriou (2002), Heuristically optimized trade-offs: a new paradigm for power laws in the Internet, *ICALP 2002*, Lecture Notes in Computer Science, **2380**, Springer, 110-122.

[19] M. Franceschetti and R. Meester (2004), Navigation in small world networks, a scale-free continuum model, preprint.

[20] C. Hermann, M. Barthélemy, and P. Provero (2003), Connectivity distribution of spatial networks, *Phys. Rev. E* 68 (2003), see also 026128 arXiv:cond-mat/0302544 v1.

[21] J. M. Kleinberg (2000), Navigation in the small world, *Nature*, **406**, 845.

[22] R. Meester and R. Roy (1996), *Continuum Percolation*, Cambridge Tracts in Mathematics.

[23] M. E. J. Newman (2003), The structure and function of complex networks, *Siam Review* **45**, 2, 167-256.

[24] M. D. Penrose (2001), A spatial central limit theorem with applications to percolation, epidemics and Boolean models, *Ann. Prob.* **29**, 4, 1515-1546.

[25] M.D. Penrose (2003), Random Geometric Graphs, Clarendon Press, Oxford.

[26] D. J. Watts (1999), Small Worlds, Princeton University Press, Princeton.

[27] D. J. Watts and S. H. Strogatz (1998), Collective dynamics of 'small world' networks, *Nature*, **393**, 440-442.

J. E. Yukich, Department of Mathematics, Lehigh University, Bethlehem PA 18015:
joseph.yukich@lehigh.edu