Discrete Mathematics and Theoretical Computer Science DMTCS vol. **15:2**, 2013, 71–86

Connectivity for line-of-sight networks in higher dimensions

Luc Devroye† and Linda Farczadi

School of Computer Science, McGill University, Montreal, Canada H3A 2K6

received 13th December 2010, *revised 3rd January 2013*, *accepted 29th April 2013*.

Let T be a d-dimensional toroidal grid of n^d points. For a given range parameter ω, and a positive integer $k \leq d$, we say that two points in T are mutually visible if they differ in at most k coordinates and are a distance at most ω apart, where distance is measured using the ℓ_p norm. We obtain a random d-dimensional line-of-sight graph G by placing a node at each point in T independently with some fixed probability p^* and connecting all pairs of mutually visible nodes. We prove an asymptotically tight connectivity result for this random graph.

Keywords: Line-of-sight networks, Connectivity, Random geometric graphs, Wireless networks

1 Introduction and Results

The study of random graphs was initiated by Erdős and Rényi (1959). Their model consists of a graph on n nodes, denoted by $G(n, p)$, where each potential edge is added independently with some probability p. A famous result of Erdős and Rényi (1960) states that if $p = (\log n + c_n)/n$, then

$$\lim_{n\to\infty} \mathbf{P}\{\text{graph is connected}\} = \begin{cases} 0 & c_n \to -\infty, \\ e^{-c} & c_n \to c \in \mathbb{R}, \\ 1 & c_n \to \infty. \end{cases}$$

A more common choice for studying wireless networks are random geometric graphs which are formed by placing n nodes uniformly at random in an underlying space and connecting all pairs of nodes that are within a distance at most r, where r is a predetermined range parameter. A detailed study of random geometric graphs can be found in a book by Penrose (2003). For the case where the underlying space is the d-dimensional torus $[0, 1]^d$ it was shown by Penrose (1999) that if $\theta r^d = (\log n + c_n)/n$, then

$$\lim_{n\to\infty} \mathbf{P}\{\text{graph is connected}\} = \begin{cases} 0 & c_n \to -\infty, \\ e^{-c} & c_n \to c \in \mathbb{R}, \\ 1 & c_n \to \infty, \end{cases}$$

†Research supported by an NSERC Discovery Grant program.

where θ is the volume of the unit ball in the norm of choice.

Line-of-Sight networks were introduced by Frieze et al. (2008) to model wireless communications in an urban setting. Their model is based on an $n \times n$ two-dimensional grid of points in the shape of a torus. The distance between two points is measured using the ℓ_1 norm. For a chosen range parameter ω, two points on the grid are said to be mutually visible if they agree in one coordinate and are a distance at most ω apart. A random graph is obtained by placing a node at each grid point independently with some probability $p^* > 0$ and connecting all mutually visible pairs of nodes. Among the results proven by Frieze et al. (2008) is the connectivity threshold for this graph. Assuming that $\omega = n^\delta$, for some $0 < \delta < 6/15$, they show that if $p^* = \left(\left(1 - \frac{1}{2}\delta\right) \ln n + \frac{1}{2} \ln \ln n + c_n\right) / 2\omega$, then

$$\lim_{n\to\infty} \mathbf{P}\left\{\text{G is connected}\right\} = \begin{cases} 0 & c_n \to -\infty, \\ e^{-\lambda} & c_n \to c \in \mathbb{R}, \\ 1 & c_n \to \infty, \end{cases}$$

where $\lambda = \frac{1}{2}\left(1 - \frac{1}{2}\delta\right) e^{-2c}$ and $c_n = o(\ln \ln n)$.

If $\omega = 1$ than the model reduces to the well studied problem of site percolation on the lattice whereas if $\omega = n$ than the model is equivalent to another well known problem consisting of the random bipartite graph $K_{n,n}$ where each edge is present with probability p^*. The motivation to take ω of the form n^δ for some $0 < \delta < 1$, is to avoid both of these extremes and focus instead on the middle region whose study requires new techniques. It was remarked in the original paper of Frieze et al. (2008) that for $\omega = o(\ln n)$ the connectivity threshold is very close to one, and therefore less interesting. This justifies the assumption that $\omega \gg \ln n$.

In this paper we extend the line-of-sight model of Frieze et al. (2008) to higher dimensions and the general ℓ_p norm. The initial motivation for higher dimensions comes from the three-dimensional case that occurs in scenarios where nodes can be placed both on the ground and also in space, for example on different floors of a building, in airplanes or in satellites Gupta and Kumar (2001) The four-dimensional case can also be of interest in a situation where communication between nodes is dependent on time which can be modeled as an additional dimension. Since the methods used for studying the connectivity of the three-dimensional model can be easily extended to any higher dimension we chose to present our results in terms of the general d-dimensional case.

1.1 Model description

Given positive integers d and k such that $d \geq 2$ and $1 \leq k < d$, we define an underlying d-dimensional grid in the shape of a torus

$$T = \left\{(x_1, \cdots, x_d) : x_i \in \{1, 2, ..., n\}, 1 \leq i \leq d\right\}.$$

We say that two points are mutually visible if they differ in at most k coordinates and are a distance at most ω apart. Distance is measured assuming that the points lie on a torus. The norm is the standard ℓ_p norm in $\mathbb{R}^d$ for $1 \leq p \leq \infty$. For reasons similar to those in Frieze et al. (2008) our range parameter ω is of the form $\omega = n^\delta$, for some constant δ satisfying $0 < \delta < \delta_0 < 1$ where δ_0 is a constant defined as

$$\delta_0 := \begin{cases} \dfrac{d}{d+k} & \text{if } d/2 < k < d, \\ \dfrac{d}{d + k\left(\lceil \frac{d}{k} \rceil - 1\right)} & \text{if } k \leq d/2. \end{cases}$$

We note that the choice of δ_0 is motivated by the technique used to prove the upper bound in Section 4.3.

We obtain a random graph G by placing a node at each grid point independently with some probability $p^* > 0$ and connecting all mutually visible pairs of nodes.

1.2 Main result

Define the constant $a_p = \binom{d}{k} \int_{B_p^k(0,1)} dx$ where $B_p^k(z,r) := \left\{x \in \mathbb{R}^k : \| x - z \|_p \leq r\right\}$. The exact expression for the integral can be found in Wang (2005) and is given by

$$\int_{B_p^k(0,1)} dx = \frac{2^k \Gamma\left(1 + \frac{1}{p}\right)^k}{\Gamma\left(1 + \frac{k}{p}\right)},$$

where Γ denotes the gamma function. The following theorem is our main result.

Theorem 1 *Let* $p^* = \frac{(d-k\delta)\ln n + \ln\ln n + c_n}{a_p \omega^k}$ *for* $\omega = n^\delta$ *and* $c_n = o(\ln\ln n)$. *Let G be the line-of-sight network defined above with placement probability p^*. Then*

$$\lim_{n\to\infty} \mathbf{P}\left\{G \textit{ is connected}\right\} = \begin{cases} 0 & c_n \to -\infty, \\ e^{-\lambda} & c_n \to c \in \mathbb{R}, \\ 1 & c_n \to \infty, \end{cases}$$

where $\lambda = \frac{(d-k\delta)e^{-c}}{a_p}$.

We observe that the connectivity threshold is the same as the threshold for the existence of isolated vertices, a result which parallels those for random graphs and random geometric graphs.

1.3 Further remarks

Even though our bounds on k exclude the case $k = d$, the corresponding threshold for this case is $p^* = ((d - d\delta)\ln n + \ln\ln n + c_n)/(a_p \omega^d)$. When $k = d$ the line-of-sight network can be thought of as a discretized version of a random geometric graph. Indeed the above threshold matches the one for random geometric graphs found in Penrose (2003). We illustrate this for the two-dimensional torus $[0,1]^2$ and the Euclidean norm.

We first set $r = (c \log N/\pi N)^{1/2}$ for some constant c and obtain a random geometric graph by placing N points uniformly at random in the torus and connecting all pairs of nodes that are within a distance r apart. Then the results of Penrose (2003) state that with high probability the graph is connected if $c > 1$ and disconnected if $c < 1$. Now suppose that on this torus we also have an n by n equally spaced grid of n^2 points. A discretized version of the random geometric graph can be obtained by taking each $1/n \times 1/n$ grid square and mapping all the nodes in this square to the grid point at the upper right corner. Even though more than one node can be mapped to the same grid point, the expected number of nodes that fall in a given square is N/n^2. Taking $n \gg \sqrt{N}$ ensures that $N/n^2 \ll 1$ and thus the probability of two nodes being mapped to the same grid point is negligible. After performing this mapping we connect a pair of mapped nodes if and only if they were connected in the original random geometric graph. Requiring $1/N \ll r$ ensures that this is roughly equivalent to connecting two nodes if and only if they are within a distance r apart.

We now set $r' = \omega/n$ and $p = N/n^2$. We obtain a line-of-sight graph by placing a node at each grid point independently with probability p and connecting all pairs of nodes that are a distance at most r' apart. Then this line-of-sight graph and the discretized random geometric graph are equivalent models. We note that the number of nodes in the two graphs is not necessarily the same. The line-of-sight graph has an expected number of n^2p nodes while the discretized random geometric graph has a total of N nodes. However since we fixed the placement probability of the line-of-sight graph as $p = N/n^2$, the total number of nodes is roughly the same. Setting $r = r'$ and solving for p we obtain $p = c\left((2-2\delta)\ln n + \ln\ln n\right)/\left(\pi n^{2\delta}\right)$ which is the connectivity threshold for the line-of-sight network with $k = d$.

2 Preliminaries and notation

All the asymptotic results in this paper are given as $n \to \infty$. We say that a certain event holds with high probability if the probability that the event holds converges to one as $n \to \infty$. For two functions $f(n)$ and $g(n)$ we write $f(n) \sim g(n)$ provided that $f(n) = g(n)(1 + o(1))$.

For an event E we let $\overline{E}$ denote its complement. We write $Bin(n, p)$ to denote a random variable with a binomial distribution with parameters n and p, and $Pois(\lambda)$ for a random variable with a Poisson distribution with parameter λ.

For a subset X of $[0, n]^d$ we write $vol(X)$ to denote the volume of X. Given two distinct points $x = (x_1, \cdots, x_d)$ and $y = (y_1, \cdots, y_d)$ we say that x is lexicographically smaller than y if there exists $j \in \{1, \cdots, d\}$ such that $x_j < y_j$ and $x_i = y_i$ for all $i < j$.

For each point $i \in T$ we define $V(i)$ to be the set of all points in T that differ in at most k coordinates from i and are within a distance ω from i. Let $S_1, \cdots, S_{\binom{d}{k}}$ be a labelling of the $\binom{d}{k}$ distinct subsets of $\{1, \cdots, k\}$ of cardinality k. For each point $i \in T$ and each S_j, we define $V_{S_j}(i)$ to be the set of all points that are within a distance ω from i and that can differ from i only in the k coordinates that are in S_j. Then $\cup_j V_{S_j}(i) = V(i)$. We refer to the sets $V_{S_j}(i)$ as *sections* of the point i. Hence the neighbourhood of each point consists of the union of $\binom{d}{k}$ sections, which are k-dimensional ℓ_p balls centered at that point. We note that the definitions of $V(i)$ and $V_{S_j}(i)$ depend on k, but this dependence is omitted from our notation.

We note that $|V(i)|$ does not depend on i and similarly $\left|V_{S_j}(i)\right|$ does not depend on i or j. Fix a point $i_1 \in T$ and define $V = |V(i_1)|$ and $V' = |V_{S_1}(i_1)|$. Then we have $|V(i)| = V$ for all $i \in T$ and $\left|V_{S_j}(i)\right| = V'$ for all $i \in T$, $1 \le j \le \binom{d}{k}$. The following lemma gives us the asymptotic values of V and V'.

Lemma 2 *Let $B_p^k(0, r)$ denote a k-dimensional ℓ_p ball of radius r centered at the origin. Then for V and V' defined as above we have*

$$V' = vol\left(B_p^k(0,\omega)\right) + O\left(\omega^{k-1}\right)$$
$$V = a_p\omega^k + O\left(\omega^{k-1}\right),$$

where a_p is the constant defined in Section 1.2.

Proof: Let A be the collection of grid points inside $B_p^k(0, \omega)$. Then $V' = |A|$. Let C be the union of all the unit length k-dimensional cubes centered at each grid point in A. Then $|A| = vol\,(B) + O\left(\omega^{k-1}\right)$.

Furthermore $vol(C) = |A|$ and $B_p^k(0, \omega - 2) \subseteq C \subseteq B_p^k(0, \omega + 1)$, where we can assume wlog that $\omega > 2$ since $\omega \to \infty$ as $n \to \infty$. We thus obtain

$$vol\left(B_p^k(0, \omega - 2)\right) \leq |A| \leq vol\left(B_p^k(0, \omega + 1)\right).$$

Hence $|A| = vol\left(B_p^k(0, \omega)\right) + O\big(\omega^{k-1}\big)$ which proves the equation for V'. The equation for V follows from using the definition of a_p in Section 1.2 and the fact that $\cup_j V_{S_j}(i) = V(i)$ for all $i \in T$. We note that while this is not a disjoint union, any two distinct sections of point i can intersect in at most $O\big(\omega^{k-1}\big)$ points. Therefore $V = a_p \omega^k + O\big(\omega^{k-1}\big)$ as required. □

Let N denote the number of isolated nodes in G. For each grid point $i \in T$ let X_i be the random indicator variable for the event that there is an isolated node at grid point i. Then $N = \sum_{i \in T} X_i$. Recalling that $p^* = \frac{(d-k\delta)\ln n + \ln\ln n + c_n}{a_p \omega^k}$ and $c_n = o(\ln \ln n)$ we obtain

$$\mathbf{E}\{N\} \sim \frac{n^d (d - k\delta)\ln n}{a_p \omega^k} e^{-p^* V} \sim \frac{(d - k\delta)e^{-c_n}}{a_p}.$$

Therefore

$$\lim_{n\to\infty} \mathbf{E}\{N\} = \begin{cases} \infty & c_n \to -\infty, \\ \dfrac{(d-k\delta)e^{-c}}{a_p} = \lambda & c_n \to c \in \mathbb{R}, \\ 0 & c_n \to \infty. \end{cases} \tag{1}$$

This shows that the asymptotic behaviour of $\mathbf{E}\{N\}$ is determined by that of c_n. We therefore need to analyze three separate cases depending on whether $c_n \to -\infty$, $c_n \to \infty$ or $c_n \to c$ for some constant $c \in \mathbb{R}$.

3 The lower bound

In this section we consider the case where $c_n \to -\infty$. We use the second moment method to show that with high probability there are isolated nodes. Recall that for each point $i \in T$ we defined $V(i)$ to be the set of all points in T that differ in at most k coordinates from i and are within a distance ω from i. Now let

$$\mathcal{S} := \{(i, j) : i, j \in T, j \notin V(i), V(i) \cap V(j) \neq \emptyset\}.$$

We note that if $j \in V(i)$ then we cannot have isolated nodes at both i and j since the two nodes would be connected. Hence the requirement $j \notin V(i)$ excludes the case where $\mathbf{E}\{X_i X_j\} = 0$. We also observe that the indicator variables X_i and X_j are independent if and only if $V(i) \cap V(j) = \emptyset$. Hence $\mathcal{S}$ consists of all pairs of points (i, j) for which the corresponding indicator random variables X_i and X_j are dependent and $\mathbf{E}\{X_iX_j\} \neq 0$. Then

$$\begin{aligned}
\mathrm{V}\{N\} &= \sum_i \mathbf{E}\left\{(X_i - \mathbf{E}\,X_i)^2\right\} + \sum_{i\neq j} \mathbf{E}\{(X_i - \mathbf{E}\,X_i)\}\,\mathbf{E}\{(X_j - \mathbf{E}\,X_j)\} \\
&= \sum_i \mathrm{V}\{X_i\} + \sum_{(i,j)\in\mathcal{S}} (\mathbf{E}\{X_iX_j\} - \mathbf{E}\{X_i\}\,\mathbf{E}\{X_j\}) \\
&\leq \sum_i \mathbf{E}\{X_i\} + \sum_{(i,j)\in\mathcal{S}} \mathbf{E}\{X_iX_j\} \\
&= \mathbf{E}\{N\} + \sum_{(i,j)\in\mathcal{S}} \mathbf{E}\{X_iX_j\}.
\end{aligned}$$

Thus using Chebyshev's inequality

$$\mathbf{P}\{N=0\} \leq \frac{\mathrm{V}\{N\}}{\mathbf{E}^2\{N\}},$$

we obtain

$$\mathbf{P}\{N=0\} \leq \frac{1}{\mathbf{E}\{N\}} + \frac{\sum_{(i,j)\in\mathcal{S}} \mathbf{E}\{X_iX_j\}}{\mathbf{E}^2\{N\}} \to 0$$

if

$$\lim_{n\to\infty} \mathbf{E}\{N\} = \infty, \quad \text{and} \quad \lim_{n\to\infty} \frac{\sum_{(i,j)\in\mathcal{S}} \mathbf{E}\{X_iX_j\}}{\mathbf{E}^2\{N\}} = 0.$$

The first condition holds from (1). To verify the second condition we first split the sum $\sum_{(i,j)\in\mathcal{S}} \mathbf{E}\{X_iX_j\}$ according to the number of coordinates in which the points $i, j \in T$ differ. Specifically for $1 \leq t \leq d$ we define

$$\mathcal{S}_t := \{(i,j) : (i,j) \in \mathcal{S} \text{ and } i \text{ and } j \text{ differ in exactly } t \text{ coordinates}\}.$$

Then

$$\sum_{(i,j)\in\mathcal{S}} \mathbf{E}\{X_iX_j\} = \sum_{t=1}^{d} \sum_{(i,j)\in\mathcal{S}_t} \mathbf{E}\{X_iX_j\}.$$

Since there are $O\big(n^d\omega^t\big)$ choices for selecting two points i and j such that $(i,j)\in\mathcal{S}_t$ we have

$$\begin{aligned}
\sum_{(i,j)\in\mathcal{S}_t} \mathbf{E}\{X_iX_j\} &= O\big(n^d\omega^t\big) \max_{(i,j)\in\mathcal{S}_t} \mathbf{E}\{X_iX_j\} \\
&= O\big(n^{d+t\delta}\big)\, {p^*}^2 (1-p^*)^{2V-\max_{(i,j)\in\mathcal{S}_t}|V(i)\cap V(j)|} \\
&= O\Big(n^{d+(t-2k)\delta}(\ln n)^2\Big)\,(1-p^*)^{2V-\max_{(i,j)\in\mathcal{S}_t}|V(i)\cap V(j)|}.
\end{aligned}$$

Since we are only summing over pairs of points i and j for which $j \notin V(i)$ there exists a constant $\epsilon > 0$ such that $\max_{(i,j)\in\mathcal{S}_t} |V(i) \cap V(j)| \leq (1-\epsilon)V$. We now consider two cases. First suppose that $t \leq k$. Then using the above observation we obtain

$$\begin{aligned}(1-p^*)^{2V-\max_{(i,j)\in\mathcal{S}_t}|V(i)\cap V(j)|} &\leq \exp\{-p^*(1+\epsilon)V\}\\ &= \exp\left\{-\frac{(d-k\delta)\ln n + \ln\ln n + c_n}{a_p\omega^k}(1+\epsilon)\left(a_p\omega^k + O\big(\omega^{k-1}\big)\right)\right\}\\ &\leq \exp\left\{-(1+\epsilon)(d-k\delta)\ln n - O\left(\frac{\ln n}{\omega}\right)\right\}\\ &\sim n^{-(1+\epsilon)(d-k\delta)}.\end{aligned}$$

Thus for $1 \leq t \leq k$ we have

$$\begin{aligned}\sum_{(i,j)\in\mathcal{S}_t} \mathbf{E}\{X_iX_j\} &= O\Big(n^{d+(t-2k)\delta-(1+\epsilon)(d-k\delta)}(\ln n)^2\Big)\\ &= O\big(n^{-\epsilon d+t\delta-k\delta+\epsilon k\delta}(\ln n)^2\big)\\ &= O\big(n^{-\epsilon}(\ln n)^2\big) \quad \text{since } t \leq k,\ k \leq d \text{ and } \delta < 1\\ &\to 0 \quad \text{since } \epsilon > 0.\end{aligned}$$

Now suppose $t > k$. Then for any $(i,j) \in \mathcal{S}_t$ we claim that $V(i) \cap V(j) = O\big(\omega^{k-1}\big)$. To see this recall that $V(i)$ and $V(j)$ consist of the union of $\binom{d}{k}$ sections in the form of k-dimensional ℓ_p balls centered at i and j respectively. Consider one section centered at i and one section centered at j. If they do not correspond to the same set of k coordinates then clearly they can only intersect in at most $O\big(\omega^{k-1}\big)$ points. On the other hand if they correspond to the same set of k coordinates then since i and j differ in strictly more than k coordinates there exists at least one coordinate not in this set in which i and j differ in, hence the two sections do not intersect. Therefore for $k < t \leq d$ we have

$$\begin{aligned}(1-p^*)^{2V-\max_{(i,j)\in\mathcal{S}_t}|V(i)\cap V(j)|} &\leq \exp\left\{-p^*\left(2V - O\big(\omega^{k-1}\big)\right)\right\}\\ &\sim \exp\left\{-p^*\left(2a_p\omega^k\right)\right\}\\ &= \exp\{-2((d-k\delta)\ln n + \ln\ln n + c_n)\}\\ &\leq n^{-2(d-k\delta)}.\end{aligned}$$

Thus for $t > k$ we have

$$\begin{aligned}\sum_{(i,j)\in\mathcal{S}_t} \mathbf{E}\{X_iX_j\} &= O\Big(n^{d+(t-2k)\delta-2(d-k\delta)}(\ln n)^2\Big)\\ &= O\big(n^{-d+t\delta}(\ln n)^2\big)\\ &\to 0 \quad \text{since } t \leq d \text{ and } \delta < 1.\end{aligned}$$

Therefore we have shown that for all $1 \leq t \leq d$ we have $\sum_{(i,j)\in\mathcal{S}_t} \mathbf{E}\{X_iX_j\} \to 0$ which proves that

$$\lim_{n\to\infty} \frac{\sum_{(i,j)\in\mathcal{S}} \mathbf{E}\{X_iX_j\}}{\mathbf{E}^2\{N\}} = 0$$

as required.

4 The upper bound

We now consider the case where $c_n \to \infty$. Our goal is to show that with high probability the graph is connected. In section 2 we saw that $\lim_{n\to\infty} \mathbf{E}\{N\} = 0$ when $c_n \to \infty$. Thus $\mathbf{P}\{N > 0\} \to 0$, which means that with high probability we do not have any isolated nodes. Using the same construction as in Frieze et al. (2008), we add nodes according to a two stage process. In the first stage we place a node at each point in T independently with probability

$$p_1 := p^* - \frac{1}{a_p \omega^k \ln n}.$$

In the second stage we place a node at each point in T independently with probability p_2. We pick p_2 so that the two stage process is equivalent to the original process with probability p^*. That is

$$(1-p_1)(1-p_2) = 1-p^*.$$

Thus

$$p_2 \sim \frac{1}{a_p \omega^k \ln n}.$$

We refer to the nodes placed in the first stage as *red* nodes, and nodes placed in the second stage as *blue* nodes. We let H denote the subgraph of G consisting of only red nodes.

We now partition each of the section $V_{S_j}(i)$ of a point i into equal sized subsections. Recall that each section $V_{S_j}(i)$ is a k-dimensional ℓ_p ball of radius ω centered at i. We define the positive integer

$$b_p = \begin{cases} \lfloor 2\omega \rfloor & \text{if } p = \infty, \\ \left\lfloor 2\omega / k^{\frac{1}{p}} \right\rfloor & \text{if } 1 \le p < \infty, \end{cases}$$

and let $V'_{S_j}(i)$ be the k-dimensional cube with side length b_p centered at i. Then $V'_{S_j}(i) \subseteq V_{S_j}(i)$ for each point i and set S_j. We now divide $V'_{S_j}(i)$ into 4^k equal sized subcubes of dimension k and side length $b_p/4$. These subcubes will be referred to as the *subsections* of the section $V_{S_j}(i)$. We may assume without loss of generality that $b_p/4$ is an integer since otherwise we can just decrease b_p as needed. Since $b_p \to \infty$ as $n \to \infty$ this will not affect the asymptotic analysis. We let $n' = (b_p/4)^k$ denote the number of grid points inside each subsection. Then the number of red nodes in a subsection is $Bin(n', p_1)$ with mean $n'p_1$. We define the positive constants

$$\gamma := \begin{cases} (d-k\delta)/\left(a_p 2^k\right) & \text{if } p = \infty \\ (d-k\delta)/\left(a_p 2^k k^{k/p}\right) & \text{if } 1 \le p < \infty \end{cases}$$

$$\beta := \left\lceil \frac{20d}{\gamma} \right\rceil.$$

The important thing to note is that $n'p_1 \ge \gamma \ln n$ and $(\gamma\beta)/10 > d$. Still following the proof method in Frieze et al. (2008) we say that a section $V_{S_j}(i)$ is *mighty* if each of its 4^k subsections contains at least $\ln \ln n$ red nodes.

4.1 High probability events

Lemma 3 *Let ϵ_1 denote the event for any red node v, which has a section $V_{S_j}(v)$ containing β red nodes $u_1, \cdots, u_\beta$, such that for some ℓ the sections $V_{S_\ell}(u_1), \cdots, V_{S_\ell}(u_\beta)$ are pairwise non-intersecting, at least one of the sections $V_{S_\ell}(u_1), \cdots, V_{S_\ell}(u_\beta)$ is mighty. Then ϵ_1 holds with high probability.*

Proof: Recall that the number of red nodes in a subsection is $Bin(n', p_1)$ with mean $n'p_1 \geq \gamma \ln n$. For large enough n we have $n'p_1 - \ln\ln n \geq (n'p_1)/10$. Therefore using a Chernoff bound that can be found in Chernoff (1952) we obtain

$$\mathbf{P}\left\{Bin\left(n', p_1\right) \leq \ln\ln n\right\} \leq \exp\left\{-\frac{\left(n'p_1 - \ln\ln n\right)^2}{2n'p_1}\right\} \leq \exp\left\{-\frac{-\gamma \ln n}{10}\right\}.$$

Since there are 4^k subsections in a section we have

$$\mathbf{P}\left\{\ V_{S_j}(i) \text{ not mighty }\right\} \leq 4^k n^{-\gamma/10}.$$

We now upper bound the probability of the event $\overline{\epsilon_1}$. This is the event that there exists a red node v which has a section $V_{S_j}(v)$ on which we can find β red nodes $u_1, \cdots, u_\beta$ such that the sections $V_{S_\ell}(u_1), \cdots, V_{S_\ell}(u_\beta)$ are all non-mighty and pairwise non-intersecting. There are n^d ways to select the location of the red node v, at most $\binom{d}{k}^2$ ways to choose the integers j and ℓ, and at most $\binom{V}{\beta}$ ways to select the locations of the β red nodes in $V_{S_j}(v)$. The red node v and the β red nodes in $V_{S_j}(v)$ are all present independently with probability p_1. Since the sections $V_{S_\ell}(u_1), \cdots, V_{S_\ell}(u_\beta)$ are pairwise non-intersecting each one is non-mighty independently from the others. Hence

$$\begin{aligned}
\mathbf{P}\left\{\overline{\epsilon_1}\right\} &\leq n^d \binom{d}{k}^2 \binom{V}{\beta} (p_1)^{\beta+1} \left(4^k n^{-\gamma/10}\right)^\beta \\
&= O\left(n^{d-(\gamma\beta)/10}\right) \quad \text{ since } \binom{V}{\beta} (p_1)^{\beta+1} = o(1) \\
&\to 0 \quad \text{ since } (\gamma\beta)/10 > d.
\end{aligned}$$

□

Lemma 4 *Let ϵ_2 denote the event that there is no red node v with deg$(v) < \ln\ln n$ that has a red neighbour w such that w has a non-mighty section. Then ϵ_2 holds with high probability.*

Proof: For a fixed red node v and $t \leq \ln\ln n$ we have

$$\mathbf{P}\left\{\deg(v) = t\right\} = \binom{V}{t} p_1^t (1-p_1)^{V-t} \leq \binom{V}{t} p_1^t e^{-p_1(V-t)} \sim \binom{V}{t} p_1^t e^{-p_1 V} \leq \binom{V}{t} p_1^t n^{-(d-k\delta)}.$$

Conditional on $\deg(v) = t$ there are t choices for a neighbour w of v. Recalling that each of the $\binom{d}{k}$

sections of w is non-mighty with probability at most $4^k n^{-\gamma/10}$ we obtain

$$
\begin{aligned}
\mathbf{P}\left\{\overline{\epsilon_2}\right\} &\le \sum_{t=0}^{\ln\ln n} n^d p_1 t \binom{d}{k} 4^k n^{-\gamma/10} \binom{V}{t} p_1^t n^{-(d-k\delta)} \\
&= O\Big(n^{-\gamma/10} \ln n\Big) \sum_{t=0}^{\ln\ln n} O\left(t \binom{V}{t} p_1^t\right) \\
&\to 0 \quad \text{since } \gamma > 0.
\end{aligned}
$$

□

Lemma 5 *Let ϵ_3 denote the event that every vertex has at least one red neighbour. Then both ϵ_3 holds with high probability.*

Proof: Recall that a red vertex is present with probability p_1 and a blue vertex is present with probability p_2. The probability that either a blue or red vertex does not have any red neighbours is just $(1-p_1)^V$. We note that

$$
(1-p_1)^V \le \exp\{-p_1 V\} = \exp\left\{-p^* V + \frac{V}{a_p \omega^k \ln n}\right\} \sim \exp\{-p^* V\}.
$$

Since $p_2 \le p_1 \le p^*$ we obtain

$$
\mathbf{P}\left\{\overline{\epsilon_3}\right\} \le n^d p^* (1-p_1)^V \le n^d p^* e^{-p^* V} \to 0
$$

since $c_n \to \infty$. □

4.2 The algorithm

We note here that this section differs in style from the analysis in Frieze et al. (2008). We assume that the high probability events $\epsilon_1, \epsilon_2, \epsilon_3, \epsilon_4$ all hold and that our graph G has no isolated nodes. All the nodes mentioned in this section are assumed to be red nodes.

We first define some notation used in Algorithm 1. We recall the definition of b_p as

$$
b_p = \begin{cases} \lfloor 2\omega \rfloor & \text{if } p = \infty, \\ \left\lfloor 2\omega / k^{\frac{1}{p}} \right\rfloor & \text{if } 1 \le p < \infty. \end{cases}
$$

Let $s := b_p/4$ and for a point $x \in T$ define the set of torus points $B(x,s) := \{y \in T : \parallel x - y \parallel_\infty \le s\}$. For a given point $u \in T$ let u_i denote its i^{th} coordinate. Given a subsection U and a point $x \in T$ we define the function

$$
d(U,x) = \max_{u:\, u \in U} \parallel u - x \parallel_{\ell_\infty},
$$

By the *closest subsection to* x, we mean the subsection U that minimizes $d(U,x)$. If the minimizer is not unique we select one arbitrarily.

Algorithm 1 takes as input any point $x \in T$ and a node $v \in H$ and returns a node $\phi(x, v)$ that is in the same component of H as v and lies in the box $B(x, s)$. The general idea is to start at node v and move from neighbouring node to neighbouring node, until we reach a node inside the box $B(x, s)$. Initially all coordinates are declared unfixed. Each time the algorithm moves to a new node it checks the distance to the point x in each of the remaining unfixed coordinates. When this distance becomes at most s it declares the respective coordinate fixed.

There are three tracking variables used by the algorithm. The variable *currNode* keeps track of the location of the algorithm. At each step the algorithm first chooses a set of k coordinates which is tracked by the variable *currSet*, and then selects a neighbour of *currNode* which lies in the section corresponding to *currSet*. Hence at each step the algorithm moves in the direction of the k coordinates which are in *currSet*. The idea is to select the *currSet* which contains the maximum number of unfixed coordinates among all available choices and then move to a neighbour which decreases the distance to x in each of the unfixed coordinates belonging to *currSet*. This is done by selecting the subsection which minimizes the distance to x. We now present the algorithm

Algorithm 1

Require: node $v \in H$, point $x \in T$
Ensure: node $\phi(x, v) \in B(x, s)$ that is in same component of H as v
If $v \in B(x, s)$ return v and terminate
update(v)
if $deg(v) < \ln \ln n$ **then**
  currSet $\leftarrow$ set S_ℓ with max number of unfixed coordinates
  currNode $\leftarrow$ any neighbour w of v
else
  currSet $\leftarrow$ set S_ℓ chosen according to Rule 1
  currNode $\leftarrow$ neighbour w such that the section $V_{\text{currSet}}(w)$ is mighty
end if
while there are unfixed coordinates **do**
  update(currNode)
  currSubsection $\leftarrow$ subsection in V_{currSet}(currNode) that is closest to x
  currSet $\leftarrow$ set S_ℓ chosen according to Rule 2
  currNode $\leftarrow$ neighbour w such that the section $V_{\text{currSet}}(w)$ is mighty
end while
return currNode

The algorithm uses the two rules:

- **Rule 1**: select the set with the maximum number of unfixed coordinates among all sets S_ℓ for which there exists a section of *currNode* with at least β nodes having pairwise non-intersecting V_{S_ℓ} sections
- **Rule 2**: select the set with the maximum number of unfixed coordinates among all sets S_ℓ for which *currSubsection* contains at least β nodes having pairwise non-intersecting V_{S_ℓ} sections.

And the subroutine:

update(v)

for all unfixed coordinates i **do**
  if $|v_i - x_i| \leq s$ **then**
    Set coordinate i to fixed.
  end if
end for

We now prove that the algorithm runs correctly through a series of claims.

Claim 6 *If $\deg(v) < \ln \ln n$ then the algorithm is able to select a neighbour of w of v whose sections are all mighty.*

Proof: Since the event ϵ_3 holds we know that v has at least one neighbour and since the event ϵ_2 holds we know that for any neighbour of v all of its sections are mighty. □

Claim 7 *If $\deg(v) > \ln \ln n$ then the algorithm is able to select a set S_ℓ according to Rule 1 and is then able to find a neighbour w of v such that $V_{S_\ell}(w)$ is mighty.*

Proof: Since $deg(v) \to \infty$ as $n \to \infty$ for large enough n there exists a section $V_{S_j}(v)$ of v which contains at least β^k nodes. Since all of these nodes are distinct and lie in the same k-dimensional section there must exists a coordinate $i \in S_j$ such at least β of these nodes have pairwise distinct i coordinates. Therefore for any choice of S_ℓ, such that $i \notin S_\ell$, the corresponding V_{S_ℓ} sections of these β nodes will be pairwise non-intersecting. Since each set contains $k < d$ coordinates there is always a choice of such a set. Now since ϵ_1 holds we know that least one of these β nodes has a mighty V_{S_ℓ} section. □

Claim 8 *In each iteration of the while loop the algorithm is able to select a set S_ℓ according to Rule 2 and is then able to find a neighbour w of currNode such that $V_{S_\ell}(w)$ is mighty.*

Proof: We first note that at the start of each iteration the section $V_{\text{currSet}}(\text{currNode})$ is mighty. This follows from the way the algorithm always selects currNode. Every subsection of a mighty section contains at least $\ln \ln n$ nodes. Hence for large enough n we have $\ln \ln n > \beta^k$ and the same argument as in the proof of Claim 7 applies. □

Claim 9 *In each iteration of the while loop the distance to x is decreased for each unfixed coordinate belonging to currSet, and the distance to x remains at most s for each fixed coordinate belonging to currSet.*

Proof: Let u=currNode and S_ℓ= currSet. Then the algorithm will select the closest subsection to x from the 4^k subsections of the section $V_{S_\ell}(u)$ and move to a new node belonging to this subsection. Let $i \in S_\ell$ and assume without loss of generality that $u_i \geq x_i$. If i is unfixed then $u_i > x_i + s$ and the algorithm will choose the subsection corresponding to the interval $[u_i - s, u_i - 2s]$ in the i^{th} coordinate. Then for any node w whose i^{th} coordinate belongs to this interval we will have $w_i - x_i < u_i - x_i$. On the other hand if i is fixed then the algorithm will choose the subsection corresponding to the the interval $[u_i, u_i - s]$. Hence for any node w whose i^{th} coordinate belongs to this interval we will have $|w_i - x_i| \leq s$. We note that the distance to x could actually increase in this case, that is we can have $|w_i - x_i| > |u_i - x_i|$, however it will always remain at most s. □

Claim 9 implies that once the distance to x in a coordinate is at most s it will remain at most s for the remainder of the algorithm. Furthermore since the algorithm selects a new currSet in each iteration, and it always chooses the one with the maximum number of unfixed coordinates, it follows that while a coordinate remains unfixed it will keep getting selected. Hence by Claim 9 after enough iterations each unfixed coordinate will become fixed. Therefore we have proved the following lemma which is the counterpart of lemma 2.5 in Frieze et al. (2008).

Lemma 10 *Assume that the high probability events $\epsilon_1, \epsilon_2, \epsilon_3, \epsilon_4$ all hold and that G does not have any isolated nodes. Then for each point $x \in T$ and each node $v \in H$ there exists a node $\phi(x, v)$ that is in the same component of H as v and lies in the box $B(x, s)$.*

4.3 The second stage

We now show how placing the blue nodes in the second stage can guarantee that the final graph is connected with high probability, following the proof method in Frieze et al. (2008). Let L be the set of points in T whose coordinates are multiples of 3ω. For each component K of H let v_K be the lexicographically smallest node in K. Now suppose we have two distinct components K and J of H. For each $x \in L$ let $\phi(x, v_K)$ and $\phi(x, v_J)$ be the unique nodes returned by Algorithm 1. Then $\phi(x, v_K)$ and $\phi(x, v_J)$ both lie in the box $B(x, s)$. Since K and J are distinct components it follows that the nodes $\phi(x, v_K)$ and $\phi(x, v_J)$ are distinct for all $x \in L$.

For $d/2 < k < d$ let $z(J, K, x)$ be the unique point in $B(x, s)$ that agrees with $\phi(x, v_K)$ in the first $d - k$ coordinates and with $\phi(x, v_J)$ in the last k coordinates. We recall that two points are said to be mutually visible if they differ in at most k coordinates and are within a distance ω apart. Since $2s \leq \omega$ it follows that $z(J, K, x)$ is mutually visible with both $\phi(x, v_K)$ and $\phi(x, v_J)$. Therefore if we were to place a blue node at the point $z(J, K, x)$ the nodes $\phi(x, v_K)$ and $\phi(x, v_J)$, and therefore the components K and J, would become linked up.

For $k \leq d/2$ let $p(J, K, x)$ be the unique path whose points all lie in $B(x, s)$ that is picked as follows: we let the first point of the path agree with $\phi(x, v_K)$ in the first $d - k$ coordinates and with $\phi(x, v_J)$ in the last k coordinates. Hence the first point of the path is mutually visible with $\phi(x, v_K)$. We let the second point agree with $\phi(x, v_K)$ in the first $d - 2k$ coordinates and with $\phi(x, v_J)$ in the last $2k$ coordinates. This ensures that the second point of the path is mutually visible with the first point of the path. We continue in this way until the last point is mutually visible with $\phi(x, v_J)$. Then this path has at most $\lceil d/k \rceil - 1$ points. If we were to place a blue node at each point of the path then $\phi(x, v_K)$ and $\phi(x, v_J)$, and therefore the components K and J, would become linked up.

Lemma 11 *For distinct points x and y in L the points $z(J, K, x)$ and $z(J, K, y)$ are distinct and the paths $p(J, K, x)$ and $p(J, K, y)$ do not have any points in common.*

Proof: This follows since the point $z(J, K, x)$ and the path $p(J, K, x)$ lie in $B(x, s)$, while the point $z(J, K, y)$ and the path $p(J, K, y)$ lie in $B(y, s)$. These boxes are disjoint since they each have side length $2s \leq \omega$ and x and y are at least 3ω apart. □

□

Hence for a fixed pair of components J and K there are $(n/3\omega)^d$ such points (or paths) that could link them up. The probability of not placing a blue node at a fixed point is $1-p_2$, while the probability of not placing a blue node at every point in a fixed path of length ℓ is $1-p_2^\ell$. The total number of components of H is upper bounded by $(2s)^d \le \omega^d$ since for any fixed point $x \in L$ each component has a point in the box $B(x,s)$ and this box has a total of $(2s)^d$ points.

Therefore for $d/2 < k < d$ we have

$$\begin{aligned}
\mathbf{P}\{H \text{ is not connected}\} &\le \mathbf{P}\{\text{there exist components } J \text{ and } K \text{ that are not linked up}\}\\
&= \omega^d(1-p_2)^{n^d/(3\omega)^d}\\
&\le \omega^d \exp\left\{-\frac{n^{d-(d+k)\delta}}{a_p 3^d \ln n}\right\}\\
&\to 0 \quad \text{if } \delta < \frac{d}{d+k}.
\end{aligned}$$

And similarly for $k \le d/2$

$$\begin{aligned}
\mathbf{P}\{H \text{ is not connected}\} &\le \mathbf{P}\{\text{there exist components } J \text{ and } K \text{ that are not linked up}\}\\
&\le \omega^d(1-p_2^{\lceil d/k\rceil -1})^{n^d/(3\omega)^d}\\
&\le \omega^d \exp\left\{-\frac{n^{d-\left(d+k\left(\lceil \frac{d}{k}\rceil -1\right)\right)\delta}}{3^d(a_p \ln n)^{\lceil d/k\rceil -1}}\right\}\\
&\to 0 \quad \text{if } \delta < \frac{d}{d+k\left(\lceil \frac{d}{k}\rceil -1\right)}.
\end{aligned}$$

Thus assuming the high probability events $\epsilon_1, \epsilon_2, \epsilon_3, \epsilon_4$ hold and G has no isolated nodes, H is connected with high probability. From ϵ_3 we know that every blue node has at least one red neighbour, and therefore it follows that G must be connected with high probability. This completes the proof of the upper bound.

4.4 The middle case

We finally consider the case $c_n \to c$. The proof in this section also follows the general structure of the one in Frieze et al. (2008). In Section 2 we have shown that if $c_n \to c$ then $\lim_{n\to\infty} \mathbf{E}\{N\} = \lambda$. For a given node v of the graph G let v' denote its location in the torus. Then define N'' to be the number of pairs of isolated nodes (u,v) such that $(u',v') \in \mathcal{S}$, where the set $\mathcal{S}$ was defined in Section 3. Let N' denote the number of isolated nodes u such that for all other isolated nodes v we have $(u',v') \in \mathcal{S}$. Then

$$N' \le N \le N' + N''.$$

In Section 3 we have shown that

$$\mathbf{E}\{N''\} = \sum_{i\sim j} \mathbf{E}\{X_i X_j\} \to 0,$$

Therefore $N = N'$ with high probability. Let t be a positive integer. We obtain the following upper and lower bounds on $\mathbf{E}\{(N')_t\}$, where $(N')_t = t!\binom{N'}{t}$,

$$\begin{aligned}\mathbf{E}\{(N')_t\} &\leq t!\binom{n^d}{t}\left(p^*(1-p^*)^V\right)^t = \left(n^d p^*(1-p^*)^V\right)^t,\\ \mathbf{E}\{(N')_t\} &\geq \left(\left(n^d - t\,(V)^2\right)p^*(1-p^*)^V\right)^t.\end{aligned}$$

Therefore

$$\lim_{n\to\infty}\mathbf{E}\{(N')_t\} = \lambda^t,$$

and thus N' is asymptotically Poisson with mean λ, Alon and Spencer (2000). This implies that

$$\lim_{n\to\infty}\mathbf{P}\{G\text{ has an isolated vertex}\} \quad = \quad 1 - \mathbf{P}\{Pois(\lambda)\text{=0}\} = 1 - e^{-\lambda}.$$

Now consider the two stage process from Section 4 and recall the definition of the events ϵ_1, ϵ_2 and ϵ_3. The only part of the analysis in Section 4 that used the fact that $c_n \to \infty$ was showing that the event ϵ_3 holds with high probability. We now have

$$\begin{aligned}\mathbf{P}\{\overline{\epsilon_3}\} &= \mathbf{P}\{\ G\text{ has a vertex with no red neighbour}\}\\ &\sim \mathbf{P}\{G\text{ has an isolated vertex}\}\\ &\to 1 - e^{-\lambda}.\end{aligned}$$

Therefore if we condition on the event ϵ_3 the rest of the analysis in Section 4 holds for the case $c_n \to c$ as well and it shows that G is connected with high probability. Therefore

$$\lim_{n\to\infty}\mathbf{P}\{G\text{ is connected}\} \quad = \quad e^{-\lambda}.$$

This completes the proof of the middle case.

Acknowledgements

The authors would like to extend special thanks to the referees whose useful comments were indispensable to improving the contents of this paper

References

N. Alon and J. Spencer. *The Probabilistic Method*. Wiley-Interscience, 2000.

M. Appel and R. Russo. The Maximum Vertex Degree of a Graph on Uniform Points in $[0,1]^d$. *Advances in Applied Probability*, 29:567–581, 1997a.

M. Appel and R. Russo. The minimum vertex degree of a graph on uniform points in $[0,1]^d$. *Advances in Applied Probability*, 29(3):582–594, 1997b.

M. Appel and R. Russo. The connectivity of a graph on uniform points on $[0,1]^d$. *Statistics & Probability Letters*, 60(4):351–357, 2002.

B. Bollobás. *Random Graphs*. Cambridge University Press, 2001.

H. Chernoff. A measure of asymptotic efficiency for tests of a hypothesis based on the sum of observations. *The Annals of Mathematical Statistics*, 23(4):493–507, 1952.

A. Czumaj and X. Wang. Communication problems in random line-of-sight ad-hoc radio networks. *Stochastic Algorithms: Foundations and Applications*, pages 70–81, 2007.

P. Erdős and A. Rényi. On random graphs. *Publ. Math. Debrecen*, 6(290-297):156, 1959.

P. Erdős and A. Rényi. On the evolution of random graphs. *Publ. Math. Inst. Hung. Acad. Sci*, 5(1):17–60, 1960.

A. Frieze, J. Kleinberg, R. Ravi, and W. Debany. Line-of-sight networks. *Combinatorics, Probability and Computing*, 18(1-2):145–163, 2008.

A. Goel, S. Rai, and B. Krishnamachari. Sharp thresholds for monotone properties in random geometric graphs. In *Proceedings of the Thirty-Sixth Annual ACM Symposium on Theory of Computing*, page 586. ACM, 2004.

P. Gupta and P. Kumar. Critical power for asymptotic connectivity in wireless networks. *Stochastic Analysis, Control, Optimization and Applications: A Volume in Honor of WH Fleming*, 3(20):547–566, 1998.

P. Gupta and P. Kumar. Internets in the sky: The capacity of three dimensional wireless networks. *Communications in Information and Systems*, 1(1):33–49, 2001.

M. Penrose. On k-connectivity for a geometric random graph. *Random Structures and Algorithms*, 15(2): 145–164, 1999.

M. Penrose. *Random Geometric Graphs*. Oxford University Press, USA, 2003.

R. van der Hofstad. Random graphs and complex networks, 2009. URL `http://www.win.tue.nl/~rhofstad/NotesRGCN.pdf`. in preparation.

X. Wang. Volumes of generalized unit balls. *Mathematics Magazine*, 78(5):390–395, 2005.