On the Mixing Time of Geographical Threshold Graphs

Andrew Beveridge [1]

Department of Mathematics and Computer Science, Macalester College, Saint Paul, MN 55105, USA. E-mail: abeverid@macalester.edu

Milan Bradonjić [2]

Theoretical Division, and Center for Nonlinear Studies, Los Alamos National Laboratory, Los Alamos, NM 87545, USA. E-mail: milan@lanl.gov

Abstract

We study the mixing time of random graphs generated by the geographical threshold graph (GTG) model, a generalization of random geometric graphs (RGG). In a GTG, nodes are distributed in a Euclidean space, and edges are assigned according to a threshold function involving the distance between nodes as well as randomly chosen node weights. We specifically study the mixing times of random walks on d-dimensional GTGs near the connectivity threshold for $d \geq 2$. If the weight distribution function decays with $\Pr[W \geq x] = O(1/x^{d+\nu})$ for an arbitrarily small constant $\nu > 0$ then the mixing time of GTG is $O(n^{2/d}(\log n)^{(d-2)/d})$. This matches the known mixing bounds for the d-dimensional RGG.

1. Introduction

In recent years, we have witnessed the development of numerous approaches to study the structure of large real-world technological and social networks, and to optimize processes on these networks. Large networks, such as the Internet, World Wide Web, phone call graphs, infections disease contacts and financial transactions, have provided new challenges for modeling

[1] Supported in part by NSA Young Investigator Grant H98230-08-1-0064.

[2] Supported in part through the Laboratory Directed Research and Development Program, and Center for Nonlinear Studies.

and analysis [1]. As an example, Web graphs may have billions of nodes and edges, which implies that processing and extracting information on these large sets of data, is 'hard' [2]. Extensive theoretical and experimental research has been done in web-graph modeling, attempting to capture both the structure and dynamics of the web graph [3, 4, 5, 6, 7, 8, 9, 10, 11].

In general, a particularly fertile approach is to consider the network as an instance of an ensemble, arising from a suitable random generative model. Since the seminal papers on the evolution of uniform random graph model [12, 13], many other models have been proposed to better capture the structure seen in real-world networks, which are systematically covered in [14]. One straightforward example is the random geometric graph (RGG) model, where nodes are placed uniformly at random in a Euclidean space and edges are placed between any two nodes within a threshold distance. For further study of RGGs, see the monograph by Penrose [15]. The RGGs have the advantage of describing many aspects of systems such as sensor networks, while avoiding unnecessary detail. However, they fail to capture heterogeneity in the network.

Geographical threshold graphs (GTGs) are a generalization of RGGs. Heterogeneity in the network is provided via a richer stochastic model that nevertheless preserves much of the simplicity of the RGG model. GTGs assign to nodes both a location and a weight. The weight may represent a quantity such as transmission power in a wireless network or influence in a social network. Edges are placed between two nodes if a symmetric function of their weights and the distance between them exceeds a certain threshold [16].

Structural properties of GTGs, such as connectivity, clustering coefficient, degree distribution, diameter, existence and absence of the giant component, chromatic number have been recently analyzed [17, 18, 19]. These properties are not merely of theoretical importance, but also play an important role in applications. In communication networks, connectivity implies the ability to reach all parts of the network. In packet routing, diameter gives the minimal number of hops needed for transmission between two arbitrary nodes. In the case of epidemics, the existence or absence of the giant component controls whether the epidemic spreads or is contained. When treating the node colors as the different radio channels or frequencies, the chromatic number gives the minimal number of channels needed so that neighboring radios do not interfere with each other.

Herein, we consider random walks on GTGs near the connectivity threshold. Random walks (or more formally, Markov chains) on large networks have many applications. For example, random walks model the spread of disease or the dispersion of information [20]. The *mixing time* of a random walk is the expected number of random steps that are required to guarantee that the current distribution is close to the stationary distribution. Mixing times are an essential tool in both theory and practice: for example, see the recent survey of Diaconis [21] on Markov chain Monte Carlo methods. In [22], the authors derived the mixing time of exponential random graphs, a model extensively used in sociology, showing that the mixing time of the *Glauber dynamics* was $\Theta(n^2 \log n)$ (exponential) for the unimodal (multimodal) *Gibbs distribution*, respectively. For the definitions of the Glauber dynamics and Gibbs distribution see [22].

The mixing time for RGG at the connectivity threshold has been determined. For the 2-dimensional RGG, Avin and Ercal [23] showed that the mixing time is $O(n)$. More recently, Cooper and Frieze [24] proved that for $d \geq 3$, the mixing time of a d-dimensional RGG is $\tilde{O}(n^{2/d})$ (in this notation, the logarithmic factors are surpressed). In this paper, we study the mixing times of random walks on d-dimensional GTGs near the connectivity threshold, where $d \geq 2$. We prove that when the node weight distribution decays sufficiently quickly, the mixing time is $O(n^{2/d}(\log n)^{(d-2)/d}) = \tilde{O}(n^{2/d})$, which matches the mixing bounds for RGG.

Model. The GTG model is constructed from of a set of n nodes placed independently uniformly at random into the unit cube in $\mathbf{R}^d$. A non-negative weight w_i, taken randomly and independently from a continuous probability distribution function $f(w) : \mathbf{R}^+ \to \mathbf{R}^+$, is assigned to each node v_i for $i \in \{1, 2, \ldots, n\}$. For nodes i, j at distance r_{ij}, the edge $\{i, j\}$ exists if and only if the following connectivity relation is satisfied:

$$G(w_i, w_j) D(r_{ij}) \geq \theta_n \,. \tag{1}$$

Here $G(w_i, w_j)$ is the interaction strength between nodes, $D(r)$ is a decreasing function of r and θ_n is a given threshold parameter that depends on the size n of the network. The interaction strength $G(w_i, w_j)$ is usually taken to be symmetric and either multiplicatively or additively separable, i.e., in the form of $G(w_i, w_j) = g(w_i)g(w_j)$ or $G(w_i, w_j) = g(w_i) + g(w_j)$. We use $D(r) = r^{-s}$ for some positive s, which is typical for e.g., the path-loss model in wireless networks [16].

Some basic results have already been shown, e.g., the expected degree of a node with given weight w, for the case of uniformly distributed nodes over a unit space [25, 16]. In both the multiplicative and additive case of $G(w_i, w_j)$, questions of diameter, connectivity, and topology control have been addressed [16].

Here we restrict ourselves to the case of $g(w) = w$, $s = d$, and nodes distributed uniformly over a d-dimensional space. For analytical simplicity we take the space to be a unit d-torus. We concentrate on the analysis of the additive model, i.e., when the connectivity relation for nodes i, j is given by

$$\frac{w_i + w_j}{r_{ij}^d} \geq \theta_n. \tag{2}$$

Let the cumulative density function (cdf), for the distribution of node weights $f(w)$, be

$$F(x) = \Pr[w \leq x] = \int_0^x f(w)dw. \tag{3}$$

For a 2-dimensional GTG, a connectivity threshold has already been shown in Theorem 5.3 of [18]. We list the changes to the proof of this result to generalize to d-dimensions. We tile the unit space $[0,1]^d$ into $\Theta(\log n/n)$ cubes of equal volume (as opposed to 2-dimensional squares). We use the connectivity relation equation (2) for the d-dimensional space, and radii $r_c(w_t) = (\log(\alpha n)/(\alpha\pi n))^{1/d}$ and $r_t(w_t) = (\lambda_c/(\alpha n))^{1/d}$, where $\alpha \in (0,1)$. With these modifications, we obtain the proof of the connectivity threshold for a d-dimensional GTG, given by the following theorem.

Theorem 1. *Let G be a GTG on the unit d-torus with threshold function $\theta_n = cn/\log n$ where the constant $c < \sup_{\alpha\in(0,1)} \alpha F^{-1}(1-\alpha)/4$. Then G is connected whp*[3].

This paper characterizes the mixing time for a GTG at the connectivity threshold, provided that the weight distribution decays at an adequate rate. Our main result is given by the following theorem.

[3]We will use the notation "with high probability" and denote it as *whp*, meaning with probability $1 - o(1)$ as n tends to infinity.

Theorem 2. *Let G be a connected GTG with threshold function $\theta_n = cn/\log n$ on the unit d-torus. If the weight distribution satisfies $Pr[W \geq x] = O(1/x^{d+\nu})$ for $\nu > 0$, and $c \leq \frac{1}{d+3} \min\left\{\sup_{\alpha\in(0,1/2]} \alpha F^{-1}(\alpha), \sup_{\alpha\in(1/2,1)}(1-\alpha)F^{-1}(\alpha)\right\}$, then the mixing time of G is $O(n^{2/d}(\log n)^{(d-2)/d})$.*

Note that $\Pr[W \geq x] = O(1/x^{d+\nu})$ also ensures that the expectation of the weights is finite.

The rest of the paper is organized as follows. In Section 2 we derive upper and lower bounds on the maximal weight. Consequently, we find upper and lower bounds on the node degrees, and show that $|E(G)| = \Theta(n \log n)$. In Section 3 we construct a family of canonical paths for the GTG. In Section 4, we prove Theorem 2. In Appendix Appendix A, we exemplify our results with two different weight functions. Our first example is the exponential weight distribution $f(w) = e^{-w}$ with cumulative density function $F(x) = 1 - e^{-x}$. Our second example is a power law cumulative density function $F(x) = 1 - x^{-\gamma}$ where $\gamma > d \geq 2$.

2. Node weights and node degrees in GTG

In this section, we determine the upper and lower bounds on the maximal weight $W_{\max}$ in GTG. Subsequently, we derive the upper and lower bounds on the degrees of the nodes in GTG near the connectivity threshold. Finally, we show that $|E(G)| = \Theta(n \log n)$ for these connected GTGs.

We adopt the following notation for the remainder of the paper. We have a constant $\alpha \in (0, 1)$ and we small fix constants ϵ, ν so that

$$0 < \epsilon < \nu/2d. \tag{4}$$

Furthermore, we assume that there is a weight W_0 such that if $W \geq W_0$ then $\Pr[W \geq x] = O(1/x^{d+\nu})$. For brevity, we will state this as "$\Pr[W \geq x] = O(1/x^{d+\nu})$." Finally, we use $\omega(n)$ to denote an arbitrarily slowly increasing function of n.

The maximal weight $W_{\max}$ satisfies $\Pr[W_{\max} \leq x] = F(x)^n$, since the weights are independently distributed. Consider a continuous weight distribution $f(w)$ with cdf $F(x)$. Our goal is to find two thresholds W_1, W_2, such that $\Pr[W_{\max} \leq W_1] = o(1)$ and $\Pr[W_{\max} \geq W_2] = o(1)$. We can 'invert'

$F(x)$ using the quantile function $F^{-1}(p) = \inf\{x \in \mathbf{R}^+ : p \leq F(x)\}$. Define

$$W_1 = F^{-1}\left(1 - \frac{\omega(n)}{n}\right) \quad \text{and} \quad W_2 = F^{-1}\left(1 - \frac{1}{n\,\omega(n)}\right).$$

We have $\Pr[W_{\max} \leq W_1] = (1-\omega(n)/n)^n \leq \exp(-\omega(n)) \to 0$ and $\Pr[W_{\max} \leq W_2] = (1 - 1/n\,\omega(n))^n \geq \exp(-1/\omega(n) + 1/n\,\omega(n)^2) \to 1$. In conclusion, the maximal weight $W_{\max}$ satisfies

$$\lim_{n\to+\infty} \Pr[W_{\max} \in (W_1, W_2)] = 1. \tag{5}$$

See Appendix Appendix A for concrete examples of the calculation of the bounds on the maximal weight.

Let us determine the upper and lower bounds for the node degrees. We consider the GTG around the connectivity regime, as described in Theorem 1. From [18], the degree $\deg_i$ of a node v_i with weight w_i follows the binomial distribution

$$\deg_i(\cdot|w_i) \sim \mathrm{Bin}(n-1, p(w_i)) \tag{6}$$

where $p(w_i) = \frac{\Upsilon_d}{\theta_n}(w_i + \mu)$, $\mu = \mathrm{E}[W]$ is the expected node weight and Υ_d is the volume of the unit ball in d dimensions:

$$\Upsilon_d = \frac{\pi^{d/2}}{\Gamma(d/2+1)} = \begin{cases} \pi^k/k!, & d = 2k \text{ even},\\ 2^d k!\pi^k/d!, & d = 2k+1 \text{ odd}. \end{cases}$$

Lemma 3. *Let G be a connected GTG with threshold function $\theta_n = cn/\log n$. For any constant c_0 such that $c_0 > 2c/\left(1 - \sqrt{\frac{c}{\mu\Upsilon_d}}\right)$, whp all nodes $v \in V(G)$ satisfy $\deg(v) \in I_{GTG}$, where*

$$I_{GTG} = \Big[\frac{2\Upsilon_d\mu}{c_0}\log n, \frac{\Upsilon_d}{c}F^{-1}\Big(1 - \frac{1}{n\omega(n)}\Big)(1+o(1))\log n\Big]. \tag{7}$$

Proof. We apply the Chernoff bound on the degree $\deg(v|w)$ of a node v with a given weight w:

$$\Pr\Big[\deg(v|w) \leq \mathrm{E}[\deg(v|w)](1-\delta)\Big] \leq \exp(-\mathrm{E}[\deg(v|w)]\delta^2/2).$$

By having $\mathrm{E}[\deg(v|w)] = (n-1)p(w)$ and choosing $\delta = 1 - \frac{2c/c_0}{1+\frac{w}{\mu}} > 0$, it follows

$$
\begin{aligned}
\Pr\Big[\deg(v|w) \le \frac{2\Upsilon_d\mu}{c_0}\log n\Big] &\le \exp(-\mathrm{E}[\deg(v|w)]\delta^2/2) \\
&\le \exp\Big(-\frac{\mu\Upsilon_d}{c}\frac{\log n}{n}(n-1)\left(1+\frac{w}{\mu}\right)\Big(1-\frac{2c/c_0}{1+\frac{w}{\mu}}\Big)^2\Big) \\
&= n^{-\frac{\mu\Upsilon_d}{c}(1-\frac{1}{n})(1+\frac{w}{\mu})\left(1-\frac{2c/c_0}{1+\frac{w}{\mu}}\right)^2}. \qquad (8)
\end{aligned}
$$

Next, we will find conditions such that equation (8) is $o(1/n)$ for all $w \ge 0$ and sufficiently large n. For the sake of simplicity, let us denote $x = 1 + w/\mu \ge 1$, and consider the function

$$
\phi(x) = \frac{\mu\Upsilon_d}{c}x\left(1-\frac{2c}{c_0 x}\right)^2.
$$

The minimum of $\phi(x)$ is attained at $x = 2c/c_0$. Moreover, $\phi(x)$ is strictly decreasing on $(0, 2c/c_0)$ and strictly increasing on $(2c/c_0, +\infty)$. By taking $2c/c_0 \le 1$ and $\phi(1) = \frac{\Upsilon_d\mu}{c}(1-2c/c_0)^2 > 1$, or equivalently $c_0 > 2c/\left(1-\sqrt{c/\mu\Upsilon_d}\right)$, it follows that $\phi(x) > 1$ for $x \ge 1$. That is, equation (8) is $o(1/n)$, for sufficiently large n. Thus, the degree distribution satisfies

$$
\begin{aligned}
&\Pr\Big[\deg(v) \le \frac{2\Upsilon_d\mu}{c_0}\log n\Big] = \int dw f(w)\Pr\Big[\deg(v|w) \le \frac{2\Upsilon_d\mu}{c_0}\log n\Big] \\
\le &\int dw f(w) n^{-\frac{\mu\Upsilon_d}{c}(1-\frac{1}{n})(1+\frac{w}{\mu})\left(1-\frac{2c/c_0}{1+\frac{w}{\mu}}\right)^2} = o\Big(\int dw f(w) n^{-1}\Big) = o(1/n).
\end{aligned}
$$

The union bound gives the lower bound on degree of the nodes in the graph.

We now obtain the upper bound. Equation (5) ensures $F^{-1}(1-\omega(n)/n) \le W_{\max} \le F^{-1}(1-1/n\omega(n))$. Moreover, by the continuity of $F^{-1}(x)$, for any $\epsilon > 0$ and any function $\omega(n)$, there is sufficiently large $n = n(\epsilon)$, such that the upper and lower bounds on $W_{\max}$ are arbitrarily close $(W_2 - W_1) \le \epsilon$.

The degree of the node with maximal weight satisfies the binomial distribution $\mathrm{Bin}(n-1, (\Upsilon_d/\theta_n)(W_{\max}+\mu))$, which is concentrated around its mean $(\Upsilon_d/c)(1-1/n)(W_{\max}+\mu)\log n$. Finally, the union bound gives the upper bound on the degrees. □

We now partition the interval I_{GTG} of equation (7). We use this partition to calculate $|E(G)|$ and again in Section 4 to bound the mixing time. Define $h(x) = (1 - F(x))^{-1}$, or equivalently, $F(x) = 1 - 1/h(x)$. By assumption we have $h(x) = \Omega(x^{d+\nu})$. Our first interval contains the low weight nodes:

$$B = \{v \in V(G) \mid w_v \leq F^{-1}(1-\alpha)\} = \{v \in V(G) \mid F(w_v) \leq 1-\alpha\}. \quad (9)$$

Next, we partition the αn nodes with weights in $[F^{-1}(1-\alpha), W_{\max}]$. By equation (5), $W_{\max} = F^{-1}\left(1 - \frac{1}{n\omega(n)}\right) = h^{-1}(n\omega(n))$. Let

$$\begin{aligned} a_0 &= W_{\max} \\ a_k &= \max\left\{F^{-1}\left(1 - a_{k-1}^{-(1+\epsilon)}\right), F^{-1}(1-\alpha)\right\}, \text{ for } k \geq 1. \end{aligned}$$

The a_k are only defined until we reach $F^{-1}(1-\alpha)$. Call this final index M. Our partition consists of the subintervals of the form $(a_k, a_{k-1}]$ for $1 \leq k \leq M$. Note that the indexing of our endpoints is the reverse of the standard convention.

For $1 \leq k \leq M$, let

$$A_k = \{v \in V(G) \mid w_v \in (a_k, a_{k-1}]\}. \quad (10)$$

The degree of $v \in A_k$ is $O((1 + w_v)\log n) = O(a_{k-1}\log n)$. Furthermore, we have $\Pr[W > a_k] = 1/a_{k-1}^{1+\epsilon}$ and therefore $\mathrm{E}[|A_k|] = \Theta\left(n/a_{k-1}^{1+\epsilon}\right)$. By the Chebychev inequality, the actual value is tightly concentrated around its mean because $a_{k-1}^{1+\epsilon} \leq a_0^{1+\epsilon} = W_{\max}^{1+\epsilon} = (h^{-1}(n\omega(n)))^{1+\epsilon} = o(n)$. See Appendix Appendix A for concrete examples of the partition of the nodes according to weight.

Lemma 4. $\sum_{k=0}^{M} a_k^{-\epsilon} = \Theta(1)$.

Proof. We use the general d'Alembert's convergence criterion: the sum of positive terms $\sum_{k=0}^{\infty} c_k$ is convergent if there exist a positive integer N and $\eta > 0$ such that $k > N$ guarantees $c_k/c_{k-1} < 1 - \eta$.

Let $b_0 = F^{-1}(1-\alpha)$ and $b_k = h(b_{k-1})^{1/(1+\epsilon)}$ for $1 \leq k \leq M$. We have $b_k \leq a_{M-k-1}$ for $0 \leq k \leq M$. It follows that

$$\sum_{k=0}^{M} \frac{1}{a_k^{\epsilon}} \leq \frac{1}{a_M^{\epsilon}} + \sum_{k=0}^{\infty} \frac{1}{b_k^{\epsilon}} \leq 1 + \sum_{k=0}^{\infty} \frac{1}{b_k^{\epsilon}}.$$

and

$$\frac{1/b_k^{\epsilon}}{1/b_{k-1}^{\epsilon}} = \left(\frac{b_{k-1}}{b_k}\right)^{\epsilon} = \left(\frac{b_{k-1}}{h(b_{k-1})^{1/(1+\epsilon)}}\right)^{\epsilon} < 1 - \eta$$

for large k since $h(x) \gg x^{d+\nu} \gg x^{1+\epsilon}$. Therefore this sum converges to a constant (independent of n). $\square$

Lemma 5. *If G is a GTG at the connectivity threshold then $|E(G)| = \Theta(n \log n)$ whp.*

Proof. Lemma 3 guarantees that all nodes have degree $\Omega(\log n)$ whp, so $|E(G)| = \Omega(n \log n)$. As for the upper bound, whp

$$\begin{aligned} 2|E(G)| &= \sum_{v \in B} \deg(v) + \sum_{k=1}^{M} \sum_{v \in A_k} \deg(v) = O\Big(\Big(|B| + \sum_{k=1}^{M} |A_k| \cdot a_{k-1}\Big) \log n\Big) \\ &= O\Big(\Big(\alpha n + \sum_{k=1}^{M} \frac{n}{a_{k-1}^{1+\epsilon}} \cdot a_{k-1}\Big) \log n\Big) = O(n \log n) \end{aligned}$$

by Lemma 4. $\square$

3. Canonical paths for GTG

We employ canonical paths (as introduced in [26]) to calculate our bound on the mixing time. In this section, we construct the canonical paths for G, a connected GTG with threshold function $\theta_n = cn/\log n$. For every ordered pair of nodes $u, v \in V(G)$ we choose a canonical path γ_{uv} between them. We define

$$\begin{aligned} \rho &= \max_{e=\{x,y\} \in E(G)} \frac{1}{\pi(x) P(x,y)} \sum_{\gamma_{uv} \ni e} \pi(u)\pi(v)|\gamma_{uv}| \\ &= \max_{e=\{x,y\} \in E(G)} \frac{1}{2|E(G)|} \sum_{\gamma_{uv} \ni e} \deg(u)\deg(v)|\gamma_{uv}| \end{aligned} \tag{11}$$

where $|\gamma_{uv}|$ is the length of the canonical path from u to v. As per [26] Proposition 12.1, the mixing time from node x satisfies

$$\tau_x(\delta) \leq \rho \left(\log \pi(x)^{-1} + \log \delta^{-1}\right) \tag{12}$$

where $\delta > 0$. We will set $\delta = 1/n$, so that $\tau_x(1/n) \leq \rho\left(\log \pi(x)^{-1} + \log n\right)$.

Our strategy for constructing the canonical paths is similar to the one used in [24] to bound the mixing time of RGG. We partition the unit cube into a toric grid of $[k]^d$ small cubes, where k is specified below. A set of canonical paths for the grid will act as the framework for our canonical paths for the GTG. Intuitively, we increment the entries $(\bmod\ k)$ in succession. So first we increase a_1 until we achieve b_1, then do the same for the second entry, and so on. The canonical path from $(a_1, a_2, \ldots, a_d)$ to $(b_1, b_2, \ldots, b_d)$ is

$$\begin{aligned}(a_1, a_2, \ldots, a_d),\ (a_1+1, a_2, \ldots, a_d),\ \ldots,\ (b_1-1, a_2, \ldots, a_d),\ (b_1, a_2, \ldots, a_d),\\ (b_1, a_2+1, \ldots, a_d),\ \ldots,\ (b_1, b_2, \ldots, b_d-1),\ (b_1, b_2, \ldots, b_d).\end{aligned} \tag{13}$$

Each path has length at most dk. Note that we always increment the index by $+1$ (even if there is a shorter path).

While there are k^{2d} canonical paths, each edge appears in no more than k^{d+1} paths. Indeed, any path that includes the edge from $(i_1, \ldots i_t, \ldots i_d)$ to $(i_1, \ldots i_t+1, \ldots i_d)$ must start at $(l_1, \ldots l_t, j_{t+1}, \ldots, j_d)$ and end at $(j_1, \ldots j_t - 1, l'_t, \ldots, l'_d)$ for some $l_1, \ldots l_t$ and $l'_t, \ldots, l'_d$. This results in t choices for $l_1, \ldots, l_t$ and $d-t+1$ choices for $l'_t, \ldots l'_d$.

Before constructing the canonical paths for our GTG, we must prove two lemmas. Let G be a GTG at the connectivity threshold $\theta_n = cn/\log n$ for $c < \sup_{\alpha \in (0,1)} \alpha F^{-1}(1-\alpha)/4$, as per Theorem 1. The proof of this result in [17] establishes the following two facts. First, a constant fraction of the nodes αn have weights greater than $F^{-1}(1-\alpha)$. We let $H(G) = \{i \in V(G) \mid w_i > F^{-1}(1-\alpha)\}$ denote this set of *high weight nodes* and let $L(G) = \{i \in V(G) \mid w_i \leq F^{-1}(1-\alpha)\}$ denote the complementary set of *low weight nodes.* Second, every high weight node is connected to every node within the *critical radius*

$$r_c = \left(\frac{\log(\alpha n)}{\alpha n \Upsilon_d}\right)^{1/d}. \tag{14}$$

Tile $[0,1]^d$ into cubes S_i with side length $1/k = (c_0 \log n/n)^{1/d} = \Theta(r_c)$, where we state the conditions on constant c_0 later. We have a $[k]^d$ grid, whose cubes each have volume $c_0 \log n/n$. Let $H(S_i)$ and $L(S_i)$ denote the high weight and low weight nodes in S_i, respectively.

Lemma 6. *There exist constants $\beta_0, \beta_1 > 0$ such that whp every cube S satisfies $\beta_0 \log n \leq |L(S)| \leq \beta_1 \log n$ and $\beta_0 \log n \leq |H(S)| \leq \beta_1 \log n$.*

Proof. This proof is similar to the proof of Theorem 3 in [17], and will use the same notation. Let $B_i = |H(S_i)|$ and $R_i = |L(S_i)|$. In expectation, there are $\mathrm{E}[B_i] = \alpha c_0 \log n$ high weight nodes within S_i. Using the lower and upper tail Chernoff bounds [27], it follows

$$\begin{aligned}
\Pr[B_i \leq (1-\delta_1)\mathrm{E}[B_i]] &\leq \exp\left(-\frac{\delta_1^2}{2}\mathrm{E}[B_i]\right), \\
\Pr[B_i \geq (1+\delta_2)\mathrm{E}[B_i]] &\leq \exp\left(-\frac{\delta_2^2}{2(1+\delta_2)}\mathrm{E}[B_i]\right).
\end{aligned}$$

Fixing $\delta_2 \in [0,1]$, we take $\delta_1 = \delta_2/\sqrt{1+\delta_2}$, and thus $\delta_1 \in [0,1]$. The number of high weighted nodes B_i within the cube S_i satisfies

$$\begin{aligned}
\Pr\left[B_i \in (1-\delta_1, 1+\delta_2)\,\mathrm{E}[B_i]\right] &\geq 1 - 2\exp\left(-\frac{\delta_2^2}{2(1+\delta_2)}\mathrm{E}[B_i]\right) \\
&= 1 - 2n^{-\alpha c_0 \delta_2^2/(1+\delta_2)}.
\end{aligned}$$

By using the spatial independence among the cubes S_i, that is, the fact that nodes are tossed over the cubes independently, it follows

$$\Pr\left[\bigcap_i \{B_i \in (1-\delta_1, 1+\delta_2)\,\mathrm{E}[B_i]\}\right] \geq \left(1 - 2n^{-\alpha c_0 \delta_2^2/(1+\delta_2)}\right)^{n/(c_0 \log n)}.$$

Taking the limit of the last expression as $n \to +\infty$, for $c_0 \geq (1+\delta_2)/\alpha\delta_2^2$, we obtain the concentration on B_i, for each cube S_i:

$$\lim_{n\to\infty} \Pr\left[\bigcap_i \{B_i \in (1-\delta_1, 1+\delta_2)\,\alpha c_0 \log n\}\right] = 1.$$

We now prove the concentration on the number of low weight nodes R_i, within each cube S_i. In expectation, there are $\mathrm{E}[R_i] = (1-\alpha)c_0 \log n$ low weight nodes, within S_i. Thus we can apply the Chernoff tail bounds on R_i, and analogously to the previous analysis we obtain

$$\lim_{n\to\infty} \Pr\left[\bigcap_i \{R_i \in (1-\delta_1, 1+\delta_2)\,(1-\alpha)c_0 \log n\}\right] = 1.$$

Finally, we can guarantee the concentration of both the high and low weight nodes by taking $c_0 \geq \max\left\{\frac{1+\delta_2}{\alpha\delta_2^2}, \frac{1+\delta_2}{(1-\alpha)\delta_2^2}\right\}$, that is,

$$c_0 \geq \begin{cases} \frac{1+\delta_2}{\alpha\delta_2^2} & \text{if } \alpha \in (0, 1/2) \\ \frac{1+\delta_2}{(1-\alpha)\delta_2^2} & \text{if } \alpha \in [1/2, 1). \end{cases} \tag{15}$$

So the lemma holds with $\beta_0 = \left(1 - \delta_2/\sqrt{1+\delta_2}\right)\min\{\alpha, 1-\alpha\}$ and $\beta_1 = (1+\delta_2)\max\{\alpha, 1-\alpha\}$. □

Two cubes are *adjacent* if they share a $(d-1)$-dimensional boundary.

Lemma 7. *Let S_i, S_j be adjacent cubes. The number of edges between $H(S_i)$ and $H(S_j)$ is $\Omega(\log^2 n)$ provided that $c \leq (d+3)^{-d/2}\min\{\sup_{x\in(1/2,1)}(1-x)F^{-1}(x), \sup_{x\in(0,1/2]} xF^{-1}(x)\}$.*

See Appendix Appendix A for explicit calculation of the constant c for two example weight distributions.

Proof. Consider any high weight node $u \in H(S_i)$ and any high weight node $v \in H(S_j)$, with the weights w_u and w_v, respectively. The distance r_{uv} between u and v is at most $\sqrt{d+3}(c_0 \log n/n)^{1/d}$. Indeed, the furthest points in two adjacent d-dimensional unit cubes are at distance $\sqrt{d+3}$ by the d-dimensional Pythagorean theorem.

Consider the connectivity relation

$$\frac{w_u + w_v}{r_{uv}^d} \geq \frac{F^{-1}(1-\alpha) + F^{-1}(1-\alpha)}{(d+3)^{d/2} c_0 \log n/n} = \frac{2F^{-1}(1-\alpha)}{(d+3)^{d/2} c_0}\frac{n}{\log n}.$$

High weight nodes u, v are connected with probability one if $(w_u + w_v)/r_{uv}^d \geq \theta_n = cn/\log n$, which is guaranteed if $2F^{-1}(1-\alpha)/((d+3)^{d/2}c_0) \geq c$. Using equation (15), we require

$$c \leq \frac{2}{(d+3)^{d/2}}\frac{\delta_2^2}{1+\delta_2}\begin{cases} \alpha F^{-1}(1-\alpha) & \text{if } \alpha \in (0, 1/2) \\ (1-\alpha)F^{-1}(1-\alpha) & \text{if } \alpha \in [1/2, 1). \end{cases}$$

Since $\delta_2 \in (0,1)$ is arbitrary, with $\sup_{\delta_2\in(0,1)} \delta_2^2/(1+\delta_2) = 1$, the conditions combine to give

$$\begin{aligned} c &\leq (d+3)^{-d/2}\min\Big\{\sup_{\alpha\in(0,1/2)}\alpha F^{-1}(1-\alpha), \sup_{\alpha\in[1/2,1)}(1-\alpha)F^{-1}(1-\alpha)\Big\} \\ &= (d+3)^{-d/2}\min\Big\{\sup_{x\in(1/2,1)}(1-x)F^{-1}(x), \sup_{x\in(0,1/2]}xF^{-1}(x)\Big\}. \end{aligned}$$

□

We now employ a randomized procedure for choosing canonical paths. This procedure guarantees that no edge appears in more than $\Theta(k^{d+1}) = \Theta(r_c^{-d-1}) = \Theta((n/\log n)^{(d+1)/d})$ canonical paths. We now denote cubes as $S(a_1,\dots a_d)$ for $1 \leq a_i \leq k$, where $(a_1,\dots,a_d)$ is the location in the d-dimensional grid. By Lemma 6, we have both $\beta_0 \log n \leq |L(S(a_1,\dots a_d))| \leq \beta_1 \log n$ and $\beta_0 \log n \leq |H(S(a_1,\dots a_d))| \leq \beta_1 \log n$ for some constants $\beta_0,\beta_1 > 0$. For each cube $S(a_1,\dots a_d)$, evenly partition $L(S(a_1,\dots a_d))$ into β_0 sets $\{L_i(S(a_1,\dots a_d))\}_{i=1}^{\beta_0}$. Each set in this partition will contain either $\lfloor\beta_1/\beta_0\rfloor$ or $\lceil\beta_1/\beta_0\rceil$ low weight nodes. Finally, choose a distinct high weight node $h_i \in H(S(a_1,\dots a_d))$ for $1 \leq i \leq k$. The high weight node h_i is the *high weight representative* for the nodes in $L_i(S(a_1,\dots a_d))$. Note that each h_i represents at most $\lceil\beta_1/\beta_0\rceil = O(1)$ low weight nodes.

Consider any ordered pair of nodes (x,y). Let $x \in S(a_1,\dots a_d)$ and $y \in S(b_1,\dots b_d)$. We choose a canonical path from x to y as follows. We use the toric grid to identify the sequence of cubes in the canonical path. Indeed, taking equation (13) as our framework, we consider cubes $S(a_1,a_2,\dots a_d)$, $S(a_1+1,a_2,\dots,a_d)$, …, $S(b_1,a_2,\dots,a_d)$, $S(b_1,a_2+1,\dots,a_d)$, …, $S(b_1,b_2,\dots,b_d)$. For brevity, call these cubes $S_0,S_1,\dots,S_t$. If x is a low weight node, then set x_0 to be the high weight representative for x. If x is a high weight node, set $x_0 = x$. For $1 \leq i \leq t$, choose x_i to be a random high weight node in S_i. Set $x_t = y$ if y is a high weight node, otherwise using the high weight representative for y. Our canonical (x,y) path is $x,x_0,x_1,\dots,x_t,y$. If $x_0 = x$ or $x_t = y$, we remove the repeated node from the path. We have

$$\max_{u,v\in V(G)}|\gamma_{uv}| \leq d\,k+2 = O\left((n/\log n)^{1/d}\right) \tag{16}$$

and furthermore, we can bound how often each edge appears in a canonical path.

Lemma 8. *Every edge in G appears in at most $O((n/\log n)^{(d+1)/d})$ canonical* *paths whp.*

Proof. Let Z_{xy} denote the number of times the edge xy is chosen. If both x and y are low weight nodes then $Z_{xy} = 0$ (we always move to a high weight node from a low weight node). When x is low weight and y is high weight, then the edge xy can only be used when x and y are in the same cube. In this case, the edge xy can only be chosen when the canonical path has x as one of its endpoints. Therefore $Z_{xy} \leq 2(n-1)$.

Now suppose that both x and y are high weight nodes. The edge xy is only used if x, y are in the same cube or in adjacent cubes. First, we consider high weight x, y in the same cube S. The edge xy will be used only by paths between $L_i(S) \cup \{x\}$ and $L_j(S) \cup \{y\}$. If one or both is not a high weight representative, then this edge will be used even fewer times. Therefore $Z_{xy} \leq (1 + \lceil \beta_1/\beta_0 \rceil)^2 = O(1)$ whp.

Consider $x \in S$ and $y \in T$ in adjacent cubes S, T. Let $\mathcal{P}_{xy} = \{\gamma_{uv} \mid xy \in \gamma_{uv}\}$. We consider four cases, according to the locations of u and v. If $u \in S$ and $v \in T$ then just as in the previous case, whp there are at most $(1 + \lceil \beta_1/\beta_0 \rceil)^2 = O(1)$ such paths.

The total number of paths γ_{uv} with $u \in S$ and $v \notin T$ is $O(n)$. Indeed, x must be the high weight representative of u, which gives $1 + \lceil \beta_1/\beta_0 \rceil = O(1)$ choices for u, and there are $O(n)$ end nodes for paths that start in S_1. Similarly, the total number of paths with $u \notin S$ and $v \in T$ is $O(n)$.

The remaining case is when $u \notin S$ and $v \notin S$. Let Z'_{xy} denote the number of times the edge xy is chosen as the random edge from S to T for some canonical path. Let $\mathcal{P}(S,T) = \{\gamma_{uv} \mid \gamma_{uv}$ traverses from S to T and $u \notin S$ and $v \notin T\}$. Then whp

$$\mathrm{E}[Z'_{xy}] = \sum_{\gamma_{uv} \in \mathcal{P}(S,T)} \Pr(xy \in \gamma_{uv}) \leq r_c^{-d-1} (\beta_1 \log n)^2 \frac{1}{(\beta_0 \log n)^2} = \left(\frac{\beta_1}{\beta_0}\right)^2 r_c^{-d-1}.$$

Indeed, the number of canonical paths for the toric grid which pass from S to T is $O(r_c^{-d-1})$. Each edge between cubes corresponds to at most $(\beta_1 \log n)^2$ toric paths in the GTG. Since both x and y are internal nodes of these paths, they were each chosen uniformly and independently with probability at most $1/(\beta_0 \log n)$.

Using the Chernoff bound for this binomial distribution,

$$\begin{aligned}
&\Pr\left[|Z'_{xy} - \mathrm{E}[Z'_{xy}]| \geq \eta \mathrm{E}[Z'_{xy}]\right] \;\leq\; 2\exp\Big(-\frac{\eta^2}{3}\mathrm{E}[Z'_{xy}]\Big) \\
&\leq\; 2\exp\Big(-\frac{\eta^2}{3}\big(\frac{\beta_1}{\beta_0}\big)^2 r_c^{-d-1}\Big) \\
&=\; 2\exp\Big(-\frac{\eta^2}{3}\big(\frac{\beta_1}{\beta_0}\big)^2\big(\frac{n}{\alpha\Upsilon_d \log n}\big)^{(d+1)/d}\Big) \;=\; o(e^{-n}).
\end{aligned}$$

The union bound now gives

$$\begin{aligned}
&\Pr\left[\wedge_{x,y}\left(|Z'_{xy} - \mathrm{E}[Z'_{xy}]| \geq \eta \mathrm{E}[Z'_{xy}]\right)\right] \\
\leq\; &\sum_{x,y} \Pr\left[|Z'_{xy} - \mathrm{E}[Z'_{xy}]| \geq \eta \mathrm{E}[Z'_{xy}]\right] \;\leq\; n^2 o(e^{-n}) \to 0.
\end{aligned}$$

Therefore whp, every edge between high weight nodes in adjacent cubes is used by $(1 \pm \epsilon) r_c^{-d-1} = \Theta((n/\log n)^{(d+1)/d})$ canonical paths. In this case, $\mathrm{E}[Z_{xy}] = \Theta((n/\log n)^{(d+1)/d})$. □

4. The mixing time for GTG

In this section, we prove that the mixing time for a d-dimensional GTG near the threshold for connectivity is $O(n^{2/d}(\log n)^{(d-2)/d})$.

Lemma 9. *Our choice of canonical paths gives $\rho = O\left((\log n)^{2/d}\right)$ when $Pr[W \geq x] = O(1/x^{d+\nu})$ for any constant $\nu > 0$.*

Proof. We have $|E(G)| = \Theta(n \log n)$ by Lemma 5 and $\max_{u,v \in V(G)} |\gamma_{uv}| = O((n/\log n)^{1/d})$ by equation (16). Substituting these values into equation (11) yields

$$\rho \leq O\left(\frac{1}{n^{(d-1)/d}(\log n)^{(d+1)/d}}\right)\left(\max_{e=\{x,y\}\in E(G)} \sum_{\gamma_{uv}\ni e} \deg(u)\deg(v)\right). \quad (17)$$

Fix an edge $e = \{x, y\}$ between high weight nodes in adjacent cubes. Suppose that these high weight cubes differ in the kth coordinate. Specifically, the cubes containing x and y are indexed by $(j_1, \ldots, j_{k-1}, j_k, j_{k+1}, \ldots, j_d)$

and $(j_1,\ldots,j_{k-1},j_k+1,j_{k+1},\ldots,j_d)$, respectively. If $e\in\gamma_{uv}$ then u must be in an initial cube indexed by $(i_1,\ldots,i_{k-1},i_k,j_{k+1},\ldots,j_d)$ and v must be in a target cube indexed by $(j_1,\ldots,j_{k-1},\ell_k,\ell_{k+1},\ldots,\ell_d)$. The volume corresponding to the set of possible initial cubes is $(n/\log n)^{k/d}/(n/\log n) = (\log n/n)^{(d-k)/d}$. Similarly, the volume corresponding to the set of all target cubes is $(n/\log n)^{(d-k+1)/d}/(n/\log n) = (\log n/n)^{(k-1)/d}$.

Let $\sigma(B,B)$ denote the contribution of paths between low weight nodes to $\sum_{\gamma_{uv}\ni e}\deg(u)\deg(v)$. We have

$$\begin{aligned}
&\mathrm{E}[\sigma(B,B)]\\
&= O\left(\left(n\left(\frac{\log n}{n}\right)^{(d-k)/d}\log n\right)\left(n\left(\frac{\log n}{n}\right)^{(k-1)/d}\log n\right)\left(\frac{1}{\log n}\right)^2\right)\\
&= O\left(n^{(d+1)/d}(\log n)^{(d-1)/d}\right).
\end{aligned}$$

Indeed, we have $\Theta(n)$ low weight nodes, so the set of initial cubes contains $n(\log n/n)^{(d-k)/d}$ of the low weight nodes whp. Similarly, the set of target cubes contains $n(\log n/n)^{(k-1)/d}$ low weight nodes whp. Each of these vertices has degree $\Theta(\log n)$. Finally, the $\Theta(1/\log^2 n)$ term is the probability that we choose e to realize path for a given pair of low weight nodes. We are guaranteed that $\sigma(B,B)$ is concentrated around its mean by Chebychev's inequality.

We consider the contribution of paths between high weight nodes by using the partition $A_1,A_2,\ldots,A_M$ specified in equation (10). Arguing similarly to the calculation above, the contribution to ρ from paths between nodes in A_r and A_s where $0\le r,s\le M$ is whp

$$\begin{aligned}
\sigma(A_r,A_s) &= O\Bigg(\left(\frac{n}{a_r^{1+\epsilon}}\left(\frac{\log n}{n}\right)^{(d-k)/d}a_r\log n\right)\\
&\qquad\cdot\left(\frac{n}{a_s^{1+\epsilon}}\left(\frac{\log n}{n}\right)^{(k-1)/d}a_s\log n\right)\left(\frac{1}{\log n}\right)^2\Bigg)\\
&= O\left(\frac{n^{(d+1)/d}(\log n)^{(d-1)/d}}{a_r^{\epsilon}a_s^{\epsilon}}\right).
\end{aligned}$$

Here the number of nodes from A_r (respectively, A_s) in the initial set (respectively, final set) of cubes is concentrated around its mean. Indeed, the expected number of such nodes is

$$
\begin{aligned}
\mathrm{E}[\sigma(A_r, A_s)] &= \left(\frac{\log n}{n}\right)^{(d-k)/d} \frac{n}{a_r^{1+\epsilon}} \geq \frac{n^{1/d}(\log n)^{1/d}}{W_{\max}^{1+\epsilon}} \\
&\geq \frac{n^{1/d}(\log n)^{1/d}}{(n\omega(n))^{(1+\epsilon)/(d+\nu)}} \geq n^{1/d-(1+\epsilon)/(d+\nu)} \\
&= \Omega\left(n^{\nu/(2d^2(1+\nu/d))}\right)
\end{aligned}
$$

since $\epsilon < \nu/2d$ by equation (4) and $\omega(n)$ grows arbitrarily slowly. Chebyschev's inequality ensures that this value is concentrated around its mean.

Next we consider the paths between low weight and high weight nodes: whp

$$
\begin{aligned}
\sigma(B, A_s) &= O\left(\left(n\left(\frac{\log n}{n}\right)^{(d-k)/d} \log n\right)\right. \\
&\qquad \left. \cdot \left(\frac{n}{a_s^{1+\epsilon}}\left(\frac{\log n}{n}\right)^{(k-1)/d} a_s \log n\right)\left(\frac{1}{\log n}\right)^2\right) \\
&= O\left(\frac{n^{(d+1)/d}(\log n)^{(d-1)/d}}{a_s^{\epsilon}}\right),
\end{aligned}
$$

and similarly, $\sigma(A_r, B) = O\left(n^{(d+1)/d}(\log n)^{(d-1)/d}/a_r^{\epsilon}\right)$. Finally, by equation (17) and Lemma 4, we have whp

$$
\begin{aligned}
\rho &= O\left(\left(\frac{n}{\log n}\right)^{2/d}\left(1 + \sum_{i=1}^{M}\frac{1}{a_i^{\epsilon}} + \sum_{j=1}^{M}\sum_{k=1}^{M}\frac{1}{a_j^{\epsilon}a_k^{\epsilon}}\right)\right) \\
&= O\left(\left(\frac{n}{\log n}\right)^{2/d}\left(1 + \sum_{i=1}^{M}\frac{1}{a_i^{\epsilon}} + \left(\sum_{j=1}^{M}\frac{1}{a_j^{\epsilon}}\right)^2\right)\right) = O\left(\left(\frac{n}{\log n}\right)^{2/d}\right).
\end{aligned}
$$

□

Proof of Theorem 2 Equation (12) gives $\tau_x(\delta) \leq \rho\left(\log \pi(x)^{-1} + \log \delta^{-1}\right)$. The previous lemma ensures that $\rho = O((n/\log n)^{2/d})$. Meanwhile, we have $\pi(x) = \deg(x)/2|E(G)| = \Omega(\log n/(n \log n)) = \Omega(n)$ by Lemma 3 and Lemma 5. In summary, $\tau_x(\delta) = O(n^{2/d}(\log n)^{(d-2)/d})$ for $\delta = 1/n$. □

References

[1] A. Bonato, A survey of models of the web graph, in: Combinatorial and Algorithm Aspects of Networking, Vol. 3405 of Lecture Notes in Computer Sciences, Springer Berlin, Heidelberg, 2005, pp. 159–172.

[2] J. Abello, P. M. Pardalos, M. G. C. Resende (Eds.), Handbook of massive data sets, Kluwer Academic Publishers, Norwell, MA, USA, 2002.

[3] A. Bonato, N. Hadi, P. Horn, P. Pralat, C. Wang, A dynamic model for on-line social networks, in: WAW, 2009, pp. 127–142.

[4] W. Aiello, A. Bonato, C. Cooper, J. C. M. Janssen, P. Pralat, A spatial web graph model with local influence regions, in: WAW, 2007, pp. 96–107.

[5] A. D. Flaxman, A. M. Frieze, J. Vera, A geometric preferential attachment model of networks ii, in: WAW, 2007, pp. 41–55.

[6] A. D. Flaxman, A. M. Frieze, J. Vera, A geometric preferential attachment model of networks, Internet Mathematics 3 (2).

[7] R. Kumar, P. Raghavan, S. Rajagopalan, D. Sivakumar, A. Tomkins, E. Upfal, Stochastic models for the web graph, in: FOCS '00: Proceedings of the 41st Annual Symposium on Foundations of Computer Science, IEEE Computer Society, Washington, DC, USA, 2000, p. 57.

[8] A.-L. Barabási, R. Albert, Emergence of scaling in random networks, Science 286 (5439) (1999) 509 – 512.

[9] W. Aiello, F. Chung, L. Lu, A random graph model for massive graphs, in: STOC '00: Proceedings of the thirty-second annual ACM symposium on Theory of computing, ACM Press, New York, NY, USA, 2000, pp. 171–180. doi:http://doi.acm.org/10.1145/335305.335326.

[10] B. Bollobás, O. Riordan, J. Spencer, G. Tusnády, The degree sequence of a scale-free random graph process, Random Struct. Algorithms 18 (3) (2001) 279–290. doi:http://dx.doi.org/10.1002/rsa.1009.

[11] C. Cooper, A. M. Frieze, A general model of undirected Web graphs, in: ESA '01: Proceedings of the 9th Annual European Symposium on Algorithms, Springer-Verlag, London, UK, 2001, pp. 500–511.

[12] P. Erdős, A. Rényi, On random graphs, Publ. Math. Inst. Hungar. Acad. Sci.

[13] P. Erdős, A. Rényi, On the evolution of random graphs, Publ. Math. Inst. Hungar. Acad. Sci.

[14] R. Durrett, Random Graph Dynamics, 1st Edition, Cambridge University Press, 2006.

[15] M. D. Penrose, Random Geometric Graphs, Oxford University Press, 2003.

[16] M. Bradonjić, J. Kong, Wireless ad hoc networks with tunable topology, in: Proceedings of the 45th Annual Allerton Conference on Communication, Control and Computing, 2007.

[17] M. Bradonjić, A. Hagberg, A. G. Percus, The structure of geographical threshold graphs, Internet Mathematics 5 (1-2) (2009) 113–139.

[18] M. Bradonjić, A. A. Hagberg, A. G. Percus, Giant component and connectivity in geographical threshold graphs, in: Proceedings of the 5th Workshop on Algorithms and Models for the Web-Graph (WAW2007), 2007, pp. 209–216.

[19] M. Bradonjić, T. Müller, A. G. Percus, Coloring geographical threshold graphs, in: Proceedings of the Fifth Workshop on Analytic Algorithmics and Combinatorics (ANALCO 09), 2009, pp. 11–16.

[20] S. Boyd, A. Ghosh, B. Prabhakar, D. Shah, Randomized gossip algorithms, IEEE Transactions on Information Theory, Special issue of IEEE Transactions on Information Theory and IEEE ACM Transactions on Networking 52 (6) (June 2006) 2508–2530.

[21] P. Diaconis, The markov chain monte carlo revolution, Bull. Amer. Math. Soc. 46 (2009) 127–205.

[22] S. Bhamidi, G. Bresler, A. Sly, Mixing time of exponential random graphs, in: 49th Annual IEEE Symposium on Foundations of Computer Science, 2008, pp. 803–812.

[23] C. Avin, G. Ercal, On the cover time and mixing time of random geometric graphs, Theor. Comput. Sci. 380 (1-2) (2007) 2–22. doi:http://dx.doi.org/10.1016/j.tcs.2007.02.065.

[24] C. Cooper, A. Frieze, The cover time of random geometric graphs, Proceedings of SODA (2009) 48–57.

[25] N. Masuda, H. Miwa, N. Konno, Geographical threshold graphs with small-world and scale-free properties, Physical Review E 71 (3) (2005) 036108. doi:10.1103/PhysRevE.71.036108.
URL `http://link.aps.org/abstract/PRE/v71/e036108`

[26] M. Jerrum, A. Sinclair, The markov chain monte carlo method: an approach to approximate counting and integration, in: D. S. Hochbaum (Ed.), Approximation Algorithms for NP-hard Problems, PWS Publishing, 1996.

[27] H. Chernoff, A note on an inequality involving the normal distribution, Ann. Probab. 9 (3) (1981) 533–535.

Appendix A. Characteristics of GTG for Example Weight Distributions

We describe the relevant characteristics of GTGs for two different weight distributions: exponential decay and polynomial decay.

Appendix A.1. Exponential Weight Distribution

Our first example is the exponential weight distribution $f(w) = e^{-w}$ with cumulative density function $F(x) = 1 - e^{-x}$. Inverting the cdf gives $F^{-1}(x) = -\log(1-x)$.

We first discuss the weights and degrees of the nodes in the GTG, as described in Section 2. As per equation (5), the maximum weight satisfies

$$W_{\max} \in [\log n - \omega(n), \log n + \omega(n)]$$

whp. Lemma 3 guarantees that whp all the node degrees are in the interval

$$I_{GTG} = \left[\Theta(\log n), \Theta(\log^2 n)\right].$$

Next, we partition the interval I_{GTG}. By equation (9), the cutoff for low weight nodes is $F^{-1}(1-\alpha) = \log \alpha^{-1}$. We partition the high weight nodes into disjoint subsets as specified in equation (10). Note that $F^{-1}\left(1 - x^{-(1+\epsilon)}\right) = \log(x^{1+\epsilon}) = \Theta(\log x)$. Therefore, ignoring leading constants, the sequence of endpoints (in descenting order) is

$$\big(\log n, \log\log n, \log\log\log n, \ldots, \log^{(M)} n\big)$$

where $M \leq \log^* n$, also known as the iterative logarithm function. This is the number of iterations of the log function required to obtain a result is less than 1.

Finally, we calculate upper bound on the constant c required by Lemma 7 in Section 3. For the exponential distribution, taking $c \leq ((d+3)^{d/2}e)^{-1}$ is sufficient.

Appendix A.2. Power Law Cumulative Density Function

We now give a parallel characterization of our second example: a power law cumulative density function $F(x) = 1 - x^{-\gamma}$ where $\gamma > d \geq 2$. Inverting this cdf gives $F^{-1}(x) = (1-x)^{-1/\gamma}$.

We consider our results concerning node weights and node degrees. We have

$$W_{\max} \in \left[\left(\frac{n}{\omega(n)}\right)^{1/\gamma}, (n\omega(n))^{1/\gamma}\right]$$

whp by equation (5). By Lemma 3, whp all the node degrees are in the interval

$$I_{GTG} = [\Theta(\log n), \Theta((n\,\omega(n))^{1/\gamma}\log n)].$$

We separate the low weight nodes from the high weight nodes using the weight cutoff $F^{-1}(1-\alpha) = \alpha^{-1/\gamma}$. Next, we partition the high weight nodes as per equation (10). Note that $F^{-1}\left(1 - x^{-(1+\epsilon)}\right) = x^{(1+\epsilon)/\gamma}$. It follows that our sequence of endpoints is

$$((n\,\omega(n))^{1/\gamma}, (n\,\omega(n))^{\beta/\gamma}, (n\,\omega(n))^{\beta^2/\gamma}, \ldots, (n\,\omega(n))^{\beta^M/\gamma})$$

where $\beta = (1+\epsilon)/\gamma < 1$. Here M is the smallest integer such that $(n\omega(n))^{\beta^M/\gamma} \leq F^{-1}(1-\alpha) = \alpha^{-1/\gamma}$. The latter requirement is equivalent to the condition $M \geq (\log\log(n\omega(n)) - \log\log(\alpha^{-1}))/\log(\beta^{-1})$.

Finally, we calculate an upper bound on the constant c required by Lemma 7 in Section 3. Since $F^{-1}(x) = (1-x)^{-1/\gamma}$, by Lemma 7 it follows $c \leq ((d+3)^{d/2} \cdot 2^{(\gamma-1)/\gamma})^{-1}$.