Internet Mathematics Vol. 3, No. 2: 187-205

A Geometric Preferential Attachment Model of Networks

Abraham D. Flaxman, Alan M. Frieze, and Juan Vera

Abstract. We study a random graph G_n that combines certain aspects of geometric random graphs and preferential attachment graphs. The vertices of G_n are n sequentially generated points $x_1, x_2, \ldots, x_n$ chosen uniformly at random from the unit sphere in $\mathbb{R}^3$. After generating x_t, we randomly connect that point to m points from those points in $x_1, x_2, \ldots, x_{t-1}$ that are within distance r of x_t. Neighbors are chosen with probability proportional to their current degree, and a parameter α biases the choice towards self loops. We show that if m is sufficiently large, if $r \geq \ln n / n^{1/2-\beta}$ for some constant β, and if $\alpha > 2$, then with high probabilty (whp) at time n the number of vertices of degree k follows a power law with exponent $\alpha + 1$. Unlike the preferential attachment graph, this geometric preferential attachment graph has small separators, similar to experimental observations of [Blandford et al. 03]. We further show that if $m \geq K \ln n$, for K sufficiently large, then G_n is connected and has diameter $O(\ln n/r)$ whp.

1. Introduction

Recently there has been much interest in understanding the properties of real-world, large-scale networks such as the structure of the Internet and the World Wide Web. For a general introduction to this topic, see [Bollobás and Riordan 02, Hayes 00, Watts 99, Aiello et al. 01]. One approach is to model these networks by random graphs. Experimental studies [Albert et al. 99, Broder et al. 00, Faloutsos et al. 99] have demonstrated that in the World Wide Web/Internet the proportion of vertices of a given degree follows an approximate inverse power law, i.e., the proportion of vertices of degree k is approximately $Ck^{-\alpha}$ for some constants C, α. The classical models of random graphs introduced by Erdős and Renyi

1542-7951/06 $0.50 per page

[Erdős and Rényi 59] do not have power law degree sequences, so they are not suitable for modeling these networks. This has driven the development of various alternative models for random graphs.

One approach is to generate graphs with a prescribed degree sequence (or prescribed expected degree sequence). This has been proposed as a model for the web graph [Aiello et al. 00]. Mihail and Papadimitriou also use this model [Mihail and Papadimitriou 02] in their study of large eigenvalues, as do Chung, Lu, and Vu [Chung et al. 03a].

An alternative approach, which we will follow in this paper, is to sample graphs via some generative procedure that yields a power law distribution. There is a long history of such models, outlined in the survey by Mitzenmacher [Mitzenmacher 04]. We will use an extension of the preferential attachment model to generate our random graph. The preferential attachment model has been the subject of recently revived interest, though it dates back further [Yule 25, Simon 55]. It was proposed as a random graph model for the web by Barabási and Albert [Barabasi and Albert 99], and their description was elaborated by Bollobás and Riordan who showed that at time n, with high probability (whp) the diameter of a graph constructed in this way is asymptotic to $\frac{\ln n}{\ln \ln n}$ [Bollobás and Riordan 04a]. Subsequently, Bollobás, Riordan, Spencer, and Tusnády proved that the degree sequence of such graphs does follow a power law distribution [Bollobás et al. 01].

The random graph defined in the previous paragraph has good expansion properties. For example, Mihail, Papadimitriou, and Saberi showed that whp the preferential attachment model has conductance bounded below by a constant [Mihail et al. 03]. On the other hand, Blandford, Blelloch, and Kash found that some web-related graphs have smaller separators than what would be expected in random graphs with the same average degree [Blandford et al. 03]. The aim of this paper is to describe a random graph model that has *both* a power-law degree distribution and small separators.

We study here the following process, which generates a sequence of graphs $G_t, t = 1, 2, \ldots, n$. The graph $G_t = (V_t, E_t)$ has t vertices and mt edges. Here, V_t is a subset of S, the surface of the sphere in $\mathbb{R}^3$ of radius $\frac{1}{2\sqrt{\pi}}$ (so that $Area(S) = 1$).

For $u \in S$ and $r > 0$, we let $B_r(u)$ denote the spherical cap of radius r around u in S. More precisely, $B_r(u) = \{x \in S : ||x - u|| \leq r\}$.

1.1. The Random Process

The parameters of the process are $m > 0$, the number of edges added in every step, and $\alpha \geq 0$, a measure of the bias towards self loops.

Notice that there exists a constant c_0 such that, for any $u \in S$, we have

$$A_r = Area(B_r(u)) \sim c_0 r^2.$$

- **Time step 0.** To initialize the process, we start with G_0 being the empty graph.
- **Time step $t+1$.** We choose vertex x_{t+1} uniformly at random in S and add it to G_t. Let $V_t(x_t) = V_t \cap B_r(x_{t+1})$, and let $D_t(x_t) = \sum_{v \in V_t(x_t)} \deg_t(v)$. We add m random edges (x_{t+1}, y_i), $i = 1, 2, \ldots, m$, incident with x_{t+1}. Here, each y_i is chosen independently from $V_t(x_t) \cup \{x_{t+1}\}$ (parallel edges and loops are permitted), such that for each $i = 1, \ldots, m$

$$\mathbf{Pr}(y_i = v) = \frac{\deg_t(v)}{\max(D_t(x_{t+1}), \alpha m A_r t)}$$

and

$$\mathbf{Pr}(y_i = x_{t+1}) = 1 - \frac{D_t(x_{t+1})}{\max(D_t(x_{t+1}), \alpha m A_r t)}$$

for all $v \in V_t(x_{t+1})$. (When $t = 0$, we have $\mathbf{Pr}(y_i = x_1) = 1$.)

Let $d_k(t)$ denote the number of vertices of degree k at time t, and let $\overline{d}_k(t)$ denote the expectation of $d_k(t)$. We will prove the following.

Theorem 1.1.

(a) If $0 < \beta < 1/2$ and $\alpha > 2$ are constants, $r \sim n^{\beta - 1/2} \ln n$, and m is a sufficiently large constant, then there exist constants $c, \gamma, \epsilon > 0$ such that for all $k = k(n) \geq m$

$$\overline{d}_k(n) = C_k \frac{n}{k^{1+\alpha}} + O(n^{1-\gamma}), \tag{1.1}$$

where $C_k = C_k(m, \alpha)$ tends to a constant $C_\infty(m, \alpha)$ as $k \to \infty$.

Furthermore, for n sufficiently large, the random variable $d_k(n)$ satisfies the following concentration inequality:

$$\mathbf{Pr}(|d_k(n) - \overline{d}_k(n)| \geq n^{1-\gamma}) \leq e^{-n^{\epsilon}}. \tag{1.2}$$

(b) If $\alpha \geq 0$ and $r = o(1)$, then whp V_n can be partitioned into $T, \bar{T}$ such that $|T|, |\bar{T}| \sim n/2$, and there are at most $4\sqrt{\pi} rnm$ edges between T and $\bar{T}$.

(c) If $\alpha \geq 0$, $r \geq n^{-1/2} \ln n$, $m \geq K \ln n$, and K is sufficiently large, then whp G_n is connected.

(d) *If* $\alpha \geq 0$, $r \geq n^{-1/2} \ln n$, $m \geq K \ln n$, *and* K *is sufficiently large, then whp* G_n *has diameter* $O(\ln n/r)$.

We note that geometric models of trees with power laws have been considered [Fabrikant et al. 02, Berger et al. 03, Berger et al. 04]. We also note that Gómez-Gardeñes and Moreno have empirically analyzed a one-dimensional version of our model when $\alpha = 0$ and their experiments suggest that this yields a power-law exponent of 3 [Gómez-Gardeñes and Moreno 04].

1.2. Open Questions

In an earlier version of the paper, there was no α and we have failed to produce a proof of Theorem 1.1(a) when $\alpha \leq 2$. This remains a challenge for us at the present moment. We do not think that the $\ln n$ factors are necessary in parts (c) and (d).

1.3. Some Definitions

Given $U \subseteq S$ and $u \in S$, we define

$$V_t(U) = V_t \cap U \qquad \text{and} \qquad V_t(u) = V_t(B_r(u))$$

and

$$D_t(U) = \sum_{v \in V_t(U)} \deg_t(v) \qquad \text{and} \qquad D_t(u) = D_t(B_r(u)).$$

Given $v \in V_t$, we also define

$$\deg_t^-(v) = \deg_t(v) - m. \tag{1.3}$$

Notice that $\deg_t^-(v)$ is the number of edges of G_t that are incident to v and were added by vertices that chose v as a neighbor, including loops at v.

Given $U \subseteq S$, let $D_t^-(U) = \sum_{v \in V_t(U)} \deg_t^-(v)$. We also define $D_t^-(u) = D_t^-(B_r(u))$. Notice that $D_t(U) = m|V_t(U)| + D_t^-(U)$.

We localize some of our notation: given $U \subseteq S$ and $u \in S$, we define $d_k(t, U)$ to be the number of vertices of degree k at time t in U and $d_k(t, u) = d_k(t, B_r(u))$.

2. Outline of the Paper

In Section 3 we show that there are small separators. This is easy, since any given great circle can whp be used to define a small separator.

We prove a likely power law for the degree sequence in Section 4. We follow a standard practise and prove a recurrence for the expected number of vertices of degree k at time step t. Unfortunatley, this involves the estimation of the expectation of the reciprocal of a random variable, and to handle this we show that this random variable is concentrated. This is quite technical and is done in Section 4.3.

Section 5 proves connectivity when m grows logarithmically with n. The idea is to show that whp the subgraph $G_n(B)$ induced by a ball B of radius $r/2$ and of center $u \in S$ is connected. This is done by constructing a connected subgraph of $G_n(B)$ via a coupling argument. We then show that the union of the $G_n(B)$ for $u = x_1, x_2, \ldots, x_n$ is connected and has small diameter.

3. Small Separators

Theorem 1.1(b) is the easiest part to prove. We use the geometry of the instance to obtain a sparse cut. Consider partitioning the vertices using a great circle of S. This will divide V into sets T and $\bar{T}$ which each contain about $n/2$ vertices. More precisely, we have

$$\mathbf{Pr}\left[|T| < (1-\epsilon)n/2\right] = \mathbf{Pr}\left[|\bar{T}| < (1-\epsilon)n/2\right] \leq e^{-\epsilon^2 n/4}.$$

Edges only appear between vertices within distance r, so only vertices appearing in the strip within distance r of the great circle can appear in the cut. Since $r = o(1)$, this strip has area less than $3r\sqrt{\pi}$, and, letting U denote the vertices appearing in this strip, we have

$$\mathbf{Pr}\left[|U| \geq 4\sqrt{\pi} rn\right] \leq e^{-\sqrt{\pi} rn/9}.$$

Even if every one of the vertices chooses its m neighbors on the opposite side of the cut, this will yield at most $4\sqrt{\pi}rnm$ edges whp. So, the graph has a cut with $\frac{e(T,\bar{T})}{|T||\bar{T}|} \leq \frac{17\sqrt{\pi}rm}{n}$ with probability at least $1 - e^{-\Omega(rn)}$.

4. Proving a Power Law

4.1. Establishing a Recurrence for $\overline{d}_k(t)$

Our approach to proving Theorem 1.1(a) is to find a recurrence for $\overline{d}_k(t)$, the expected number of vertices of degree k at time t. We define $\overline{d}_{m-1}(t) = 0$ for all integers t with $t > 0$. Let $\eta_k(G_t, x_{t+1})$ denote the (conditional) probability that a parallel edge to a vertex of degree no more than k is created. Then,

$$\begin{aligned}\eta_k(G_t, x_{t+1}) &= O\left(\sum_{i=m}^{k} \frac{d_i(t, x_{t+1})\, i^2}{\max\{\alpha m A_r t, D_t(x_{t+1})\}^2}\right) \\ &= O\left(\min\left\{\frac{k^2}{\max\{\alpha m A_r t, D_t(x_{t+1})\}}, 1\right\}\right). \end{aligned} \tag{4.1}$$

Then, for $k \geq m$,

$$\begin{aligned}\mathbf{E}\left[d_k(t+1) \mid G_t, x_{t+1}\right] = d_k(t) &+ m d_{k-1}(t, x_{t+1}) \frac{k-1}{\max\{\alpha m A_r t, D_t(x_{t+1})\}} \\ &- m d_k(t, x_{t+1}) \frac{k}{\max\{\alpha m A_r t, D_t(x_{t+1})\}} \\ &+ \mathbf{Pr}\left[\deg_{t+1}(x_{t+1} = k) \mid G_t, x_{t+1}\right] \\ &+ O(m \eta_k(G_t, x_{t+1})). \end{aligned} \tag{4.2}$$

Let $\mathcal{A}_t$ be the event

$$\{|D_t(x_{t+1}) - 2mA_r t| \leq C_1 A_r m t^{\gamma} \ln n\}$$

where

$$\max\{2/\alpha, 1/2, 1-2\beta\} < \gamma < 1$$

and C_1 is some sufficiently large constant.

Note that if

$$t \geq (\ln n)^{2/(1-\gamma)}$$

then

$$\mathcal{A}_t \text{ implies } D_t(x_{t+1}) \leq \alpha m A_r t.$$

Then, because

$$\mathbf{E}[d_k(t, x_{t+1})] \leq k^{-1} \mathbf{E}[m|V_t(B_{2r}(x_{t+1}))|] \leq k^{-1} m(4A_r t)$$

and

$$d_k(t, x_{t+1}) \leq k^{-1} D_t(x_{t+1}) < mt,$$

we have for $t \geq (\ln n)^{2/(1-\gamma)}$

$$\begin{aligned}
&\mathbf{E}\left[\frac{d_k(t, x_{t+1})}{\max\{\alpha m A_r t, D_t(x_{t+1})\}}\right] \\
&\quad = \mathbf{E}\left[\frac{d_k(t, x_{t+1})}{\max\{\alpha m A_r t, D_t(x_{t+1})\}} \;\middle|\; \mathcal{A}_t\right] \mathbf{Pr}\left[\mathcal{A}_t\right] \\
&\qquad + \mathbf{E}\left[\frac{d_k(t, x_{t+1})}{\max\{\alpha m A_r t, D_t(x_{t+1})\}} \;\middle|\; \neg\mathcal{A}_t\right] \mathbf{Pr}\left[\neg\mathcal{A}_t\right] \\
&\quad = \frac{\mathbf{E}\left[d_k(t, x_{t+1}) \mid \mathcal{A}_t\right]}{\alpha m A_r t} \mathbf{Pr}\left[\mathcal{A}_t\right] + \mathbf{E}\left[O\left(\frac{d_k(t, x_{t+1})}{D_t(x_{t+1})}\right) \;\middle|\; \neg\mathcal{A}_t\right] \mathbf{Pr}\left[\neg\mathcal{A}_t\right] \\
&\quad = \frac{\mathbf{E}\left[d_k(t, x_{t+1}) \mid \mathcal{A}_t\right]}{\alpha m A_r t} \mathbf{Pr}\left[\mathcal{A}_t\right] + O\left(\frac{\mathbf{Pr}\left[\neg\mathcal{A}_t\right]}{k}\right) \\
&\quad = \frac{\mathbf{E}\left[d_k(t, x_{t+1})\right]}{\alpha m A_r t} + \left(O\left(\frac{1}{k}\right) - \frac{\mathbf{E}\left[d_k(t, x_{t+1}) \mid \neg\mathcal{A}_t\right]}{\alpha m A_r t}\right) \mathbf{Pr}\left[\neg\mathcal{A}_t\right] \\
&\quad = \frac{\mathbf{E}\left[d_k(t, x_{t+1})\right]}{\alpha m A_r t} + O\left(\frac{1}{k} + \frac{1}{A_r}\right) \mathbf{Pr}\left[\neg\mathcal{A}_t\right].
\end{aligned}$$

In Lemmas 4.1 and 4.3 we prove that

$$\mathbf{E}\left[d_k(t, x_{t+1})\right] = m A_r \bar{d}_k(t)$$

and that

$$\mathbf{Pr}\left[\neg\mathcal{A}_t\right] = O\left(n^{-2}\right). \tag{4.3}$$

Thus, if $t \geq (\ln n)^{2/(1-\gamma)}$, then

$$\mathbf{E}\left[\frac{d_k(t, x_{t+1})}{\max\{\alpha m A_r t, D_t(x_{t+1})\}}\right] = \frac{\bar{d}_k(t)}{\alpha m t} + O\left(\frac{1}{n^2}\left(\frac{1}{A_r} + \frac{1}{k}\right)\right). \tag{4.4}$$

In a similar way,

$$\mathbf{E}\left[\frac{d_{k-1}(t, x_{t+1})}{\max\{\alpha m A_r t, D_t(x_{t+1})\}}\right] = \frac{\bar{d}_{k-1}(t)}{\alpha m t} + O\left(\frac{1}{n^2}\left(\frac{1}{A_r} + \frac{1}{k}\right)\right). \tag{4.5}$$

On the other hand, given G_t and x_{t+1}, if

$$p = 1 - \frac{D_t(x_{t+1})}{\max\left(D_t(x_{t+1}), \alpha m A_r t\right)},$$

then

$$\mathbf{Pr}\left[\deg_{t+1}(x_{t+1} = k) \mid G_t, x_{t+1}\right] = \mathbf{Pr}\left[\mathrm{Bi}(m, p) = k - m\right].$$

So, if $t \geq (\ln n)^{2/(1-\gamma)}$,

$$\begin{aligned}
&\mathbf{Pr}\left[x_{t+1} = k\right] \\
&\quad = \binom{m}{k-m} \mathbf{E}\left[p^{k-m}(1-p)^{2m-k} \,\middle|\, \mathcal{A}_t\right] \mathbf{Pr}\left[\mathcal{A}_t\right] + O(\mathbf{Pr}\left[\neg\mathcal{A}_t\right]) \\
&\quad = \binom{m}{k-m}\left(1-\frac{2}{\alpha}\right)^{k-m}\left(\frac{2}{\alpha}\right)^{2k-m}(1+O(t^{\gamma-1}\ln n))\mathbf{Pr}\left[\mathcal{A}_t\right] + O(n^{-2}) \\
&\quad = \binom{m}{k-m}\left(1-\frac{2}{\alpha}\right)^{k-m}\left(\frac{2}{\alpha}\right)^{2k-m} + O(t^{\gamma-1}\ln n).
\end{aligned}$$

Now note from Equations (4.1) and (4.3) that if

$$t \geq t_0 = n^{(1-2\beta)/\gamma}$$

and

$$k \leq k_0(t) = (mA_r t^{\gamma} \ln n)^{1/2},$$

then

$$\mathbf{E}(\eta_k(G_t, x_{t+1})) = O(t^{\gamma-1}\ln n). \tag{4.6}$$

Taking expectations on both sides of Equation (4.2) and using Equations (4.4), (4.5), and (4.6), we see that, if $t \geq t_0$ and $k \leq k_0(t)$, then

$$\begin{aligned}
\overline{d}_k(t+1) = \overline{d}_k(t) &+ \frac{k-1}{\alpha t}\overline{d}_{k-1}(t) - \frac{k}{\alpha t}\overline{d}_k(t) \\
&+ \binom{m}{k-m}\left(1-\frac{2}{\alpha}\right)^{k-m}\left(\frac{2}{\alpha}\right)^{2m-k} + O\left(t^{\gamma-1}\ln n\right).
\end{aligned} \tag{4.7}$$

We consider the recurrence given by $f_{m-1} = 0$, and for $k \geq m$

$$f_k = \frac{k-1}{\alpha} f_{k-1} - \frac{k}{\alpha} f_k + \binom{m}{k-m}\left(1-\frac{2}{\alpha}\right)^{k-m}\left(\frac{2}{\alpha}\right)^{2m-k},$$

which, for $k > 2m$, has solution

$$\begin{aligned}
f_k &= f_{2m} \prod_{i=m+1}^{k} \frac{i-1}{i+\alpha} \\
&= \phi_k(m,\alpha)\left(\frac{m}{k}\right)^{\alpha+1},
\end{aligned}$$

where $\phi_k(m,\alpha)$ tends to a limit $\phi_\infty(m,\alpha)$ depending only on m, α as $k \to \infty$. We can absorb the values $f_m, f_{m+1}, \ldots, f_{2m}$ into this notation.

We finish the proof of Equation (1.1) by showing that there exists a constant $M > 0$ such that

$$|\overline{d}_k(t) - f_k t| \leq M(t_0 + t^\gamma \ln n) \tag{4.8}$$

for all $0 \leq t \leq n$ and $m \leq k \leq k_0(t)$. This is trivially true for $t < t_0$. For $k > k_0(t)$, this follows from $\overline{d}_k(t) \leq 2mt/k$.

Let $\Theta_k(t) = \overline{d}_k(t) - f_k t$. Then, for $t \geq t_0$ and $m \leq k \leq k_0(t)$,

$$\Theta_k(t+1) = \frac{k-1}{\alpha t}\Theta_{k-1}(t) - \frac{k}{\alpha t}\Theta_k(t) + O(t^{\gamma-1}\ln n). \tag{4.9}$$

Let L denote the hidden constant in $O(t^{\gamma-1}\ln n)$ in Equation (4.9). Our inductive hypothesis $\mathcal{H}_t$ is that

$$|\Theta_k(t)| \leq M(t_0 + t^\gamma \ln n)$$

for every $m \leq k \leq k_0(t)$ and M sufficiently large. It is trivially true for $t \leq t_0$. So, assume that $t \geq t_0$. Then, from Equation (4.9),

$$\begin{aligned} |\Theta_k(t+1)| &\leq M(t_0 + t^\gamma \ln n) + Lt^{\gamma-1}\ln n \\ &\leq M(t_0 + (t+1)^\gamma \ln n). \end{aligned}$$

This verifies $\mathcal{H}_{t+1}$ and completes the proof by induction.

4.2. Expected Value of $d_k(t,u)$

Lemma 4.1. *Let $u \in S$, and let k and t be positive integers. Then, $\mathbf{E}[d_k(t,u)] = A_r\overline{d}_k(t)$.*

Proof. By symmetry, for any $w \in S$, $d_k(t,u)$ has the same distribution as $d_k(t,w)$. Then,

$$\begin{aligned} \mathbf{E}[d_k(t,u)] &= \int_S \mathbf{E}[d_k(t,u)]\,dw = \int_S \mathbf{E}[d_k(t,w)]\,dw \\ &= \mathbf{E}\left[\int_S d_k(t,w)dw\right] = \mathbf{E}\left[\int_S \sum_{v\in V_t} 1_{\deg v=k} 1_{v\in B_r(w)}dw\right] \\ &= \mathbf{E}\left[\sum_{v\in V_t} 1_{\deg v=k}\int_S 1_{w\in B_r(v)}dw\right] = \mathbf{E}\left[\sum_{v\in V_t} 1_{\deg v=k}A_r\right] \\ &= A_r\mathbf{E}[d_k(t)]. \end{aligned}$$

$\square$

Lemma 4.2. *Let $u \in S$ and $t > 0$. Then, $\mathbf{E}[D_t(u)] = 2A_r mt$.*

Proof.

$$\mathbf{E}[D_t(u)] = \sum_{k>0} \mathbf{E}[d_k(t,u)] = A_r \sum_{k>0} \mathbf{E}[d_k(t)] = A_r \mathbf{E}\left[\sum_{k>0} d_k(t)\right] = 2A_r mt.$$

□

4.3. Concentration of $D_t(u)$

In this section we prove the following lemma.

Lemma 4.3. *If $t > 0$ and u is chosen randomly from S, then*

$$\mathbf{Pr}\left[|D_t(u) - \mathbf{E}[D_t(u)]| \geq A_r m(t^{2/\alpha} + t^{1/2}\ln t)\ln n\right] = O\left(n^{-2}\right).$$

Proof. We think of every edge added as two directed edges. We also think of x_t, the vertex added, as being added with $(\alpha m A_r t - D_t(x_t))^+ = \max\{\alpha m A_r t - D_t(x_t), 0\}$ "phantom" edges pointing to it. Then, choosing a vertex is equivalent to choosing one of these directed edges uniformly and taking the vertex to which this edge points as the chosen vertex. So the ith step of the process is defined by a tuple of random variables $T = (X, Y_1, \ldots, Y_m) \in S \times E_i^m$ where X is the location of the new vertex, a randomly chosen point in S, and Y_j is an edge chosen uniformly at random from among the edges directed into $B_r(X)$ in G_{i-1}. The process G_t is then defined by a sequence $\langle T_1, \ldots, T_t\rangle$, where each $T_i \in S \times E_i^m$.

Let s be a sequence $s = \langle s_1, \ldots, s_t\rangle$ where $s_i = (x_i, y_{(i-1)m+1}, \ldots, y_{im})$ with $x_i \in S$ and $y_j \in E_{\lceil j/m\rceil}$. We say s is *acceptable* if, for every j, y_j is an edge entering $B_r(x_{\lceil t/j\rceil})$. Notice that non-acceptable sequences have probability zero of being realized. Fix $t > 0$. Fix an acceptable sequence $s = \langle s_1, \ldots, s_t\rangle$, and let $A_\tau(s) = \{z \in S \times E_\tau^m : \langle s_1, \ldots, s_{\tau-1}, z\rangle \text{ is acceptable}\}$. For any τ with $1 \leq \tau \leq t$ and any $z \in A_\tau(s)$, let

$$g_\tau(z) = \mathbf{E}[D_t(u) \mid T_1 = s_1, \ldots, T_{\tau-1} = s_{\tau-1}, T_\tau = z],$$

let $r_\tau(s) = \sup\{|g_\tau(z) - g_\tau(\hat{z})| : z, \hat{z} \in A_\tau(s)\}$, and let $\hat{r}^2(s) = \sum_{\tau=1}^{t}(\sup_s r_\tau(s))^2$, where the supremum is taken over all acceptable sequences.

From the Azuma-Hoeffding inequality (see, for example, [Alon and Spencer 00]) we know that, for all $\lambda > 0$,

$$\mathbf{Pr}[|D_t(u) - \mathbf{E}[D_t(u)]| \geq \lambda] < 2e^{-\lambda^2/2\hat{r}^2}. \tag{4.10}$$

Fix τ, with $1 \leq \tau \leq t$. Our goal now is to bound $r_\tau(s)$ for any acceptable sequence s. Also, fix $z, \hat{z} \in A_\tau(s)$. We define $\Omega(G_t, \hat{G}_t)$ as a coupling between $G_t = G_t(s_1, \ldots, s_{\tau-1}, z)$ and $\hat{G}_t = G_t(s_1, \ldots, s_{\tau-1}, \hat{z})$.

- **Step τ.** Start with the graphs $G_\tau(s_1, \ldots, s_{\tau-1}, z)$ and $\hat{G}_\tau(s_1, \ldots, s_{\tau-1}, \hat{z})$.
- **Step σ ($\sigma > \tau$).** Choose the same point $x_\sigma \in S$ in both processes. Let E_σ (respectively $\hat{E}_\sigma$) be the edges pointing to the vertices in $B_r(x_\sigma)$ in $G_{\sigma-1}$ (respectively $\hat{G}_{\sigma-1}$) plus the $(\alpha m A_r \sigma - D_\sigma(x_\sigma))^+$ (respectively $(\alpha m A_r \sigma - \hat{D}_\sigma(x_\sigma))^+$) phantom edges pointing to x_σ. Let $C_\sigma = E_\sigma \cap \hat{E}_\sigma$, $R_\sigma = E_\sigma \setminus \hat{E}_\sigma$, and $L_\sigma = \hat{E}_\sigma \setminus E_\sigma$.

  Notice that $|E_\sigma|, |\hat{E}_\sigma| \geq \alpha m A_r \sigma$. Notice also that if $D_\sigma(x_\sigma), D'_\sigma(x_\sigma) \leq \alpha m A_r \sigma$, then $|E_\sigma| = |\hat{E}_\sigma|$ and $|R_\sigma| = |L_\sigma|$. Without loss of generality, assume that $|E_\sigma| \leq |\hat{E}_\sigma|$.

  Now, define $p = 1/|E_\sigma|$ and $\hat{p} = 1/|\hat{E}_\sigma|$. Construct G_σ by choosing m edges uniformly at random $e_1^\sigma, \ldots, e_m^\sigma$ in E_σ and then joining x_σ to their endpoints, $y_1^\sigma, \ldots, y_m^\sigma$. For each of the m edges $e_i = e_i^\sigma$, we define $\hat{e}_i = \hat{e}_i^\sigma$ as

  - if $e_i \in C_\sigma$ then, with probability $\hat{p}/p$, $\hat{e}_i = e_i$. With probability $1 - \hat{p}/p$, $\hat{e}_i$ is chosen from L_σ uniformly at random.
  - if $e_i \in R_\sigma$, $\hat{e}_i \in L_\sigma$ is chosen uniformly at random.

  Notice that, for every $i = 1, \ldots, m$ and every $e \in \hat{E}_\sigma$, $\mathbf{Pr}[\hat{e}_i = e] = \hat{p}$. To finish, in $\hat{G}_\sigma$ join x_σ to the m vertices pointed to by the edges $\hat{e}_i$.

$\square$

Now let

$$\Delta_\sigma = \sum_{\rho=\tau}^{\sigma} \sum_{i=1}^{m} 1_{y_i^\sigma \neq \hat{y}_i^\sigma},$$

and for $u \in S$ let

$$\Delta_\sigma(u) = \sum_{\rho=\tau}^{\sigma} \sum_{i=1}^{m} 1_{|\{y_i^\sigma, \hat{y}_i^\sigma\} \cap B_r(u)| = 1}.$$

Lemma 4.4.

$$|g_\tau(z) - g_\tau(\hat{z})| \leq \mathbf{E}[\Delta_t(u)].$$

Proof.

$$\begin{aligned} |g_\tau(z) - g_\tau(\hat{z})| &= |\mathbf{E}_{G_t}[D_t(u)] - \mathbf{E}_{\hat{G}_t}[D_t(u)]| \\ &= |\mathbf{E}_{\Omega(G_t, G'_t)}[D_t(u) - D'_t(u)]| \\ &\leq \mathbf{E}_{\Omega(G_t, G'_t)}[\Delta_t(u)], \end{aligned}$$

since only when $|\{y_i^\sigma, \hat{y}_i^\sigma\} \cap B_r(u)| = 1$ do we add ± 1 to the *difference* $D_\rho(u) - D'_\rho(u)$. $\square$

Recall that $A_r = Area(B_r(u)) \sim c_0 n^{2\beta-1}(\ln n)^2$ and that we have fixed τ to be an integer with $1 \le \tau \le t$.

Lemma 4.5. *Let $t \ge 1$ and $u \in S$. Then, for some constant $C > 0$,*

$$\mathbf{E}\left[\Delta_t(u)\right] \le CmA_r\left(\frac{t}{\tau}\right)^{2/\alpha}.$$

Proof. Let $\tau < \sigma \le t$. We start with

$$\Delta_\sigma = \Delta_{\sigma-1} + \sum_{i=1}^{m} 1_{y_i^\sigma \neq \hat{y}_i^\sigma}. \tag{4.11}$$

Now fix $G_{\sigma-1}, \hat{G}_{\sigma-1}$, x_σ, and i. Then, taking expectations with respect to our coupling,

$$\begin{aligned}\mathbf{E}\left[1_{y_i^\sigma \neq \hat{y}_i^\sigma}\right] &= \mathbf{Pr}(y_i^\sigma \neq \hat{y}_i^\sigma) = \mathbf{Pr}(e_i^\sigma \neq \hat{e}_i^\sigma)\\ &= 1 - \frac{|C_\sigma|}{|E_\sigma|}\frac{\hat{p}}{p} = 1 - \frac{|C_\sigma|}{|\hat{E}_\sigma|} = \frac{|L_\sigma|}{|\hat{E}_\sigma|} = \frac{\max\{|L_\sigma|, |R_\sigma|\}}{\max\{|E_\sigma|, |\hat{E}_\sigma|\}} \le \frac{|L_\sigma| + |R_\sigma|}{\alpha m A_r \sigma}.\end{aligned} \tag{4.12}$$

Therefore,

$$\mathbf{E}\left[\Delta_\sigma \;\middle|\; G_{\sigma-1}, \hat{G}_{\sigma-1}, x_\sigma\right] \le \Delta_{\sigma-1} + m\frac{|L_\sigma| + |R_\sigma|}{\alpha m A_r \sigma}. \tag{4.13}$$

For each $e \in E(\hat{G}_{\sigma-1}) \setminus E(G_{\sigma-1})$, $e \in L_\sigma$ implies that x_σ is in the ball of radius r centered at the end point e, similarly for $e \in R_\sigma$. Therefore,

$$\mathbf{E}\left[|L_\sigma| + |R_\sigma| \mid G_{\sigma-1}, \hat{G}_{\sigma-1}\right] \le 2A_r\Delta_{\sigma-1}. \tag{4.14}$$

Then,

$$\begin{aligned}\mathbf{E}\left[\Delta_\sigma\right] \le \mathbf{E}\left[\Delta_{\sigma-1}\right] + m\frac{\mathbf{E}\left[|L_\sigma| + |R_\sigma|\right]}{\alpha m A_r \sigma} &\le \mathbf{E}\left[\Delta_{\sigma-1}\right] + \frac{2\mathbf{E}\left[\Delta_{\sigma-1}\right]}{\alpha\sigma}\\ &= \mathbf{E}\left[\Delta_{\sigma-1}\right]\left(1 + \frac{2}{\alpha\sigma}\right),\end{aligned}$$

so $\mathbf{E}\left[\Delta_t\right] \le e^{10/\alpha^2}\left(\frac{t}{\tau}\right)^{2/\alpha}\mathbf{E}\left[\Delta_\tau\right]$. Now, $\Delta_\tau \le m$, because the graphs G_τ and $\hat{G}_\tau$ differ at most in the last m edges. Therefore, $\mathbf{E}\left[\Delta_t\right] \le me^{10/\alpha^2}\left(\frac{t}{\tau}\right)^{2/\alpha}$.

Finally, note that if v is a random point in S then $\mathbf{E}\left[\Delta_t(v)\right] = A_r\mathbf{E}\left[\Delta_t\right]$. For this, fix u and let ϕ denote a random rotation of S. Let $v = \phi(u)$, and then run the first process with $\phi(G_\tau), \phi(\hat{G}_\tau)$ and $x_\sigma, \sigma > \tau$. Then consider the second process starting with $G_\tau, \hat{G}_\tau$ and $\phi^{-1}(x_\sigma), \sigma > \tau$. The mapping ϕ^{-1} does not disturb the distribution of $x_\sigma, \sigma > \tau$. Therefore $\Delta_t(u)$ in the second process is equal to $\Delta_t(v)$ in the first process. □

By applying Lemma 4.5, we have that for any acceptable sequence

$$R^2(s) = \sum_{\tau=1}^{t} r_\tau(s)^2 \leq (CmA_r)^2 t^{4/\alpha} \sum_{\tau=1}^{t} \tau^{-4/\alpha} = O\left(A_r^2 m^2 (t \ln t + t^{4/\alpha})\right).$$

Therefore, by using Equation (4.10), we have that there is C_1 such that

$$\mathbf{Pr}\left[|D_t(u) - \mathbf{E}\left[D_t(u)\right]| \geq C_1 A_r m (t^{2/\alpha} + t^{1/2} \ln t)(\ln n)^{1/2}\right] \leq e^{-2\ln n} = n^{-2}.$$

4.4. Concentration of $d_k(t)$

We follow the proof of Lemma 4.3, replacing $D_t(u)$ by $d_k(t)$ and using the same coupling. When we reach Lemma 4.4, we find that $|g_\tau(z) - g_\tau(\hat{z})| \leq 2\mathbf{E}[\widehat{D}_t]$ (i.e., each edge discrepancy can affect two vertices); the rest is the same.

This proves Equation (1.1) and completes the proof of Theorem 1.1(a).

5. Connectivity

Here we are going to prove that for $r \geq n^{-1/2} \ln n$, $m > K \ln n$, and K sufficiently large, whp G_n is connected and has diameter $O(\ln n/r)$. Notice that G_n is a subgraph of the graph $G(n,r)$—the intersection graph of the caps $B_r(x_t)$, $t = 1, 2, \ldots, n$—and therefore it is disconnected for $r = o((n^{-1} \ln n)^{1/2})$ [Penrose 03]. We denote the diameter of G by $\mathrm{diam}(G)$ and follow the convention of defining $\mathrm{diam}(G) = \infty$ when G is disconnected. In particular, when we say that a graph has finite diameter, this implies it is connected.

Let

$$T = \frac{K_1 \ln n}{A_r} = O(n/\ln n),$$

where K_1 is sufficiently large and $K_1 \ll K$.

Lemma 5.1. *Let $u \in S$, and let $B = B_{r/2}(u)$. Then,*

$$\mathbf{Pr}\left[\mathrm{diam}(G_n(B)) \geq 2(K_1 + 1) \ln n\right] = O(n^{-3}),$$

where $G_n(B)$ is the induced subgraph of G_n in B.

Proof. Given τ_0 and N, we consider the following process, which generates a sequence of graphs $H_s = (W_s, F_s)$, $s = 1, 2, \dots, N$. (The meanings of N and τ_0 will become apparent soon).

- **Time step 1.** To initialize the process, we start with H_1 consisting of τ_0 isolated vertices $y_1, \dots, y_{\tau_0}$.
- **Time step $s \geq 1$.** We add vertex $y_{s+\tau_0}$. We then add $\frac{m}{8000(\alpha+1)^2}$ random edges incident with $y_{s+\tau_0}$ of the form $(y_{s+\tau_0}, w_i)$ for $i = 1, 2, \dots, \frac{m}{8000(\alpha+1)^2}$. Here each w_i is chosen uniformly from W_s.

The idea is to couple the construction of G_n with the construction of H_N for $N \sim \mathrm{Bi}(n-T, A_r/4)$ and $\tau_0 = \mathrm{Bi}(T, A_r/4)$ such that whp H_N is a subgraph of G_n with vertex set $V_n(B)$. We are then going to show that whp $\mathrm{diam}(H_N) \leq 2(K_1+1)\ln n$, and therefore $\mathrm{diam}(G_n(B)) \leq 2(K_1+1)\ln n$.

To do the coupling we use two counters: t for the steps in G_n and s for the steps in H_N.

- Given G_{τ_0}, set $s = 0$. Let $W_0 = V_T(B)$. Notice that $\tau_0 = |W_0| \sim \mathrm{Bi}(T, A_r/4)$ and that $\tau_0 \leq K_1 \ln n$ whp.
- For every $t > T$,
  - if $x_t \notin B$, do nothing in H_s.
  - if $x_t \in B$, set $s := s+1$. Set $y_{s+\tau_0} = x_t$. Since we want H_N to be a subgraph of G_n, we must choose the neighbors of $y_{s+\tau_0}$ among the neighbors of x_t in G_n. Let A be the set of vertices chosen by x_t in $V_t(B)$. Notice that $|A|$ stochastically dominates
  $$a_t \sim \mathrm{Bi}\left(m, \frac{D_t(B)}{\max\{\alpha m A_r t, D_t(x_t)\}}\right).$$
  If
  $$\frac{D_t(B)}{\max\{\alpha m A_r t, D_t(x_t)\}} \geq \frac{1}{50(\alpha+1)},$$
  then a_t stochastically dominates $b_t \sim \mathrm{Bi}(m, \frac{1}{50\alpha})$ and so whp is at least $\frac{m}{100(\alpha+1)}$. If
  $$\frac{D_t(B)}{\max\{\alpha m A_r t, D_t(x_t)\}} < \frac{1}{50(\alpha+1)}$$
  we declare failure (but as we see below this is unlikely to happen).

For any $R > 0$,

$$\begin{aligned} m|V_t(B_R(w))| \leq D_t(B_R(w)) &= m|V_t(B_R(w))| + D_t^-(B_R(w)) \\ &\leq 2m|V_t(B_{R+r}(w))|, \end{aligned} \tag{5.1}$$

where $D_t^-(B_R(w))$ is the sum over vertices $x \in B_R(w)$ of the in-degree $\deg_t(x) - m$ of x.

Now $|V_t(B_R(w)| \sim \mathrm{Bi}(t, (R/r)^2 A_r)$, and so

$$\begin{aligned} &\mathbf{Pr}(D_t(x_t) \geq 8mA_r t \text{ OR } D_t(B) \notin [mA_r t/5, 3mA_r t] \\ &\qquad\qquad \text{OR } |V_t(B)| < A_r t/5) \leq n^{-K_1/100}. \end{aligned} \tag{5.2}$$

So, we assume that G_t is such that the event described in Equation (5.2) does not happen. Thus, each vertex of B has probability at least

$$\frac{m}{8(\alpha+1)mA_r t} \geq \frac{1}{40(\alpha+1)|V_t(B)|}$$

of being chosen under preferential attachment. Thus, as insightfully observed by Bollobás and Riordan [Bollobás and Riordan 04b], we can legitimately start the addition of x_t in G_t by choosing $\frac{m}{8000(\alpha+1)^2}$ random neighbours uniformly in B.

Notice that N, the number of times s is increased, is the number of steps for which $x_t \in B$, and so $N \sim \mathrm{Bi}(n - T, A_r/4)$. Now we are ready to show that H_N is connected whp.

By Chernoff's bound we have that

$$\mathbf{Pr}\left[\left|\tau_0 - \frac{K_1}{4}\ln n\right| \geq \frac{K_1}{8}\ln n\right] \leq 2n^{-K_1/48}$$

and

$$\mathbf{Pr}\left[N \leq \frac{1}{3}(\ln n)^2\right] \leq e^{-c(\ln n)^2}$$

for some $c > 0$. Therefore, we can assume that $\ln n \leq \tau_0 \leq K_1 \ln n$ and $N \geq \frac{1}{3}(\ln n)^2$.

Let X_s be the number of connected components of H_s. Then,

$$X_{s+1} = X_s - Y_s \text{ and } X_0 = \tau_0,$$

where $Y_s \geq 0$ is the number of components (minus one) collapsed into one by $y_{s+\tau_0}$. So,

$$\mathbf{Pr}\left[Y_s = 0 \mid H_s\right] \leq \sum_{i=1}^{X_s} \left(\frac{c_i}{s+\tau_0}\right)^{m/8000(\alpha+1)^2},$$

where the c_i are the component sizes of H_s. If $s < 2K_1 \ln n$ then, because $m \geq K \ln n$, we have

$$\begin{aligned}\mathbf{Pr}\left[Y_s = 0 \mid X_s \geq 2\right] &\leq 2\left(1 - \frac{1}{s+\tau_0}\right)^{m/8000(\alpha+1)^2}\\ &\leq 2e^{-m/(8000(\alpha+1)^2(s+\tau_0))}\\ &\leq \frac{1}{10}.\end{aligned}$$

Thus, X_s is stochastically dominated by the random variable $\max\{1, \tau_0 - Z_s\}$ where $Z_s \sim \mathrm{Bi}(s, 9/10)$. We then have

$$\mathbf{Pr}\left[X_{2K_1 \ln n} > 1\right] \leq \mathbf{Pr}\left[Z_{2K_1 \ln n} < \tau_0\right] \leq \mathbf{Pr}\left[Z_{2K_1 \ln n} < K_1 \ln n\right] \leq n^{-3}.$$

Therefore,

$$\mathbf{Pr}\left[H_{2K_1 \ln n} \text{ is not connected}\right] \leq n^{-3}.$$

Now, to obtain an upper bound on the diameter, we run the process of construction of H_N by rounds. The first round consists of $2K_1 \ln n$ steps, and in each new round we double the size of the graph: i.e., it consists of as many steps as the total number of steps of all the previous rounds. Notice that we have less than $\ln n$ rounds in total. Let $\mathcal{A}$ be the event "for all $i > 0$ every vertex created in the $(i+1)$th round is adjacent to a vertex in $H_{2^{i-1}K_1 \ln n}$, the graph at the end of the ith round."

On the event $\mathcal{A}$, every vertex in H_N is at distance at most $\ln n$ of $H_{2K_1 \ln n}$, whose diameter is not greater than $2K_1 \ln n$. Thus, the diameter of H_N is smaller than $2(K_1 + 1) \ln n$.

Now, we have that if v is created in the $(i+1)$th round,

$$\mathbf{Pr}\left[v \text{ is not adjacent to } H_{2^{i-1}K_1 \ln n}\right] \leq \left(\frac{1}{2}\right)^m.$$

Therefore,

$$\mathbf{Pr}\left[\neg\mathcal{A}\right] \leq \left(\frac{1}{2}\right)^m n(\ln n) \leq \frac{\ln n}{n^{K \ln 2 - 1}}. \qquad \square$$

To finish the proof of connectivity and the diameter, let u and v be two vertices of G_n. Let $C_1, C_2, \ldots, C_M$, $M = O(1/r)$, be a sequence of spherical caps of radius $r/4$ such that u is the center of C_1, v is the center of C_M, and the centers of C_i and C_{i+1} are distance $\leq r/2$ apart. The intersections of C_i and C_{i+1} have area at least $A_r/40$, and so whp each intersection contains a vertex. Using Lemma 5.1 we deduce that whp there is a path from u to v in G_n of size at most $O(\ln n/r)$.

Acknowledgments. We thank Olivier Riordan for detailed comments that pointed to a major error in our proof in an earlier version of this paper. We also thank Zeng Jianyang for his comments. Work by the second author was supported in part by NSF grant CCR-0200945.

References

[Aiello et al. 00] W. Aiello, F. R. K. Chung, and L. Lu. "A Random Graph Model for Massive Graphs." In *Proceedings of the Thirty-Second Annual ACM Symposium on the Theory of Computing*, pp. 171–180. New York: ACM Press, 2000.

[Aiello et al. 01] W. Aiello, F. R. K. Chung, and L. Lu. "Random Evolution in Massive Graphs." In *Proceedings of the 42nd Annual IEEE Symposium on Foundations of Computer Science*, pp. 510–519. Los Alamitos, CA: IEEE Press, 2001.

[Albert et al. 99] R. Albert, A. Barabási, and H. Jeong. "Diameter of the World Wide Web." *Nature* 401 (1999), 103–131.

[Alon and Spencer 00] N. Alon and J. Spencer. *The Probabilistic Method*, Second edition. New York: Wiley-Interscience, 2000.

[Barabasi and Albert 99] A. Barabasi and R. Albert. "Emergence of Scaling in Random Networks." *Science* 286 (1999), 509–512.

[Berger et al. 03] N. Berger, B. Bollobas, C. Borgs, J. Chayes, and O. Riordan. "Degree Distribution of the FKP Network Model." In *Automata, Languages and Programming: 30th International Colloquium, ICALP 2003, Eindhoven, The Netherlands, June 30–July 4, 2003, Proceedings*, pp. 725–738, Lecture Notes in Computer Science 2719. Berlin: Springer, 2003.

[Berger et al. 04] N. Berger, C. Borgs, J. Chayes, R. D'Souza, and R. D. Kleinberg. "Competition-Induced Preferential Attachment." In *Automata, Languages and Programming: 31st Internaional Colloquium, ICALP 2004, Turku, Finland, July 12–16, 2004, Proceedings*, pp. 208–221, Lecture Notes in Computer Science 3142. Berlin: Springer, 2004.

[Blandford et al. 03] D. Blandford, G. E. Blelloch, and I. Kash. Compact Representations of Separable Graphs. In *Proceedings of the Fourteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, pp. 679–688. Philadelphia: SIAM, 2003.

[Bollobás and Riordan 02] B. Bollobás and O. Riordan. "Mathematical Results on Scale-Free Random Graphs." In *Handbook of Graphs and Networks*, edited by S. Bornholdt and H. G. Schuster, pp. 1–34. Berlin: Wiley-VCH, 2002.

[Bollobás and Riordan 04a] B. Bollobás and O. Riordan. "The Diameter of a Scale-Free Random Graph." *Combinatorica* 4 (2004), 5–34.

[Bollobás and Riordan 04b] B. Bollobás and O. Riordan. "Coupling Scale Free and Classical Random Graphs." *Internet Mathematics* 1:2 (2004), 215–225.

[Bollobás et al. 01] B. Bollobás, O. Riordan, J. Spencer and G. Tusanády. "The Degree Sequence of a Scale-Free Random Graph Process." *Random Structures and Algorithms* 18 (2001), 279–290.

[Broder et al. 00] A. Broder, R. Kumar, F. Maghoul, P. Raghavan, S. Rajagopalan, R. Stata, A. Tomkins, and J. Wiener. "Graph Structure in the Web." In "Proceedings of the 9th International World Wide Web Conference," special issue, *Computer Networks* 33:1–6 (2000), 309–320.

[Buckley and Osthus 04] G. Buckley and D. Osthus. "Popularity Based Random Graph Models Leading to a Scale-Free Degree Distribution." *Discrete Mathematics* 282 (2004), 53–68.

[Chung et al. 03a] F. R. K. Chung, L. Lu, and V. Vu. "Eigenvalues of Random Power Law Graphs." *Annals of Combinatorics* 7 (2003), 21–33.

[Chung et al. 03b] F. R. K. Chung, L. Lu, and V. Vu. "The Spectra of Random Graphs with Expected Degrees." *Proceedings of National Academy of Sciences* 100 (2003), 6313–6318.

[Cooper and Frieze 03] C. Cooper and A. M. Frieze. "A General Model of Undirected Web Graphs." *Random Structures and Algorithms* 22 (2003), 311–335.

[Drinea et al. 01] E. Drinea, M. Enachescu, and M. Mitzenmacher. "Variations on Random Graph Models for the Web." Harvard Technical Report TR-06-01, 2001.

[Erdős and Rényi 59] P. Erdős and A. Rényi. "On Random Graphs I." *Publicationes Mathematicae Debrecen* 6 (1959), 290–297.

[Fabrikant et al. 02] A. Fabrikant, E. Koutsoupias, and C. H. Papadimitriou. "Heuristically Optimized Trade-Offs: A New Paradigm for Power Laws in the Internet." In *Automata, Languages and Programming: 29th Internaional Colloquium, ICALP 2002, Malaga, Spain, July 8–13, 2002, Proceedings*, pp. 110–122, Lecture Notes in Computer Science 2380. Berlin: Springer, 2002.

[Faloutsos et al. 99] M. Faloutsos, P. Faloutsos, and C. Faloutsos. "On Power-Law Relationships of the Internet Topology." *ACM SIGCOMM Computer Communication Review* 29 (1999), 251–262.

[Gómez-Gardeñes and Moreno 04] J. Gómez-Gardeñes and Y. Moreno. "Local Versus Global Knowledge in the Barabási-Albert Scale-Free Network Model." *Physical Review E* 69 (2004), 037103.

[Hayes 00] B. Hayes. "Graph Theory in Practice: Part II." *American Scientist* 88 (2000), 104–109.

[Kleinberg et al. 99] J. M. Kleinberg, R. Kumar, P. Raghavan, S. Rajagopalan, and A. S. Tomkins. "The Web as a Graph: Measurements, Models and Methods." In *Computing and Combinatorics: 5th Annual International Conference, COCOON '99, Tokyo, Japan, July 26–28, 1999, Proceedings*, pp. 1–17, Lecture Notes in Computer Science 1627. Berlin: Springer, 1999.

[Kumar et al. 99] R. Kumar, P. Raghavan, S. Rajagopalan, and A. Tomkins. "Trawling the Web for Emerging Cyber-Communities." *Computer Networks* 31 (1999), 1481–1493.

[Kumar et al. 00a] R. Kumar, P. Raghavan, S. Rajagopalan, D. Sivakumar, A. Tomkins, and E. Upfal. "Stochastic Models for the Web Graph." In *Proceedings of the 41st Annual Symposium on Foundations of Computer Science*, pp. 57–65. Los Alamitos, CA: IEEE Press, 2000.

[Kumar et al. 00b] R. Kumar, P. Raghavan, S. Rajagopalan, D. Sivakumar, A. Tomkins, and E. Upfal. "The Web as a Graph." In *Proceedings of the Nineteenth ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems*, pp. 1–10. New York: ACM Press, 2000.

[McDiarmid 98] C. J. H. McDiarmid. "Concentration." In *Probabilistic Methods in Algorithmic Discrete Mathematics*, edited by M. Habib et al., pp. 195–248. New York: Springer, 1998.

[Mihail and Papadimitriou 02] M. Mihail and C. H. Papadimitriou. "On the Eigenvalue Power Law." In *Randomization and Approximation Techniques in Computer Science: 6th International Workshop, RANDOM 2002, Cambridge, MA, USA, September 13–15, 2002, Proceedings*, pp. 254–262, Lecture Notes in Computer Science 2483. New York: Springer, 2002.

[Mihail et al. 03] M. Mihail, C. H. Papadimitriou, and A. Saberi. "On Certain Connectivity Properties of the Internet Topology." In *Proceedings of the 44th Annual IEEE Symposium on Foundations of Computer Science*, pp. 28–35. Los Alamitos, CA: IEEE Press, 2003.

[Mitzenmacher 04] M. Mitzenmacher. "A Brief History of Generative Models for Power Law and Lognormal Distributions." *Internet Mathematics* 1:2 (2004), 226–251.

[Penrose 03] M. D. Penrose. *Random Geometric Graphs*. Oxford, UK: Oxford University Press, 2003.

[Simon 55] H. A. Simon. "On a Class of Skew Distribution Functions." *Biometrika* 42 (1955), 425–440.

[Watts 99] D. J. Watts. *Small Worlds: The Dynamics of Networks between Order and Randomness*. Princeton, NJ: Princeton University Press, 1999.

[Yule 25] G. Yule. "A Mathematical Theory of Evolution Based on the Conclusions of Dr. J. C. Willis." *Philosophical Transactions of the Royal Society of London (Series B)* 213 (1925), 21–87.

Abraham D. Flaxman, Department of Mathematical Sciences, Carnegie Mellon University, Pittsburgh, PA 15213 (abie@cmu.edu)

Alan M. Frieze, Department of Mathematical Sciences, Carnegie Mellon University, Pittsburgh, PA 15213 (alan@random.math.cmu.edu)

Juan Vera, Department of Mathematical Sciences, Carnegie Mellon University, Pittsburgh, PA 15213 (jvera@andrew.cmu.edu)

Received March 27, 2006; accepted September 19, 2006.