Adversarial Deletion in a Scale Free Random Graph Process *

Abraham D. Flaxman † Alan M. Frieze † Juan Vera †

Abstract

We study a dynamically evolving random graph which adds vertices and edges using preferential attachment and is "attacked by an adversary". At time t, we add a new vertex x_t and m random edges incident with x_t, where m is constant. The neighbors of x_t are chosen with probability proportional to degree. After adding the edges, the adversary is allowed to delete vertices. The only constraint on the adversarial deletions is that the total number of vertices deleted by time n must be no larger than δn, where δ is a constant. We show that if δ is sufficiently small and m is sufficiently large then with high probability at time n the generated graph has a component of size at least $n/30$.

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 [7], Hayes [21], Watts [30], or Aiello, Chung and Lu [3]. One approach is to model these networks by random graphs. Experimental studies by Albert, Barabási, and Jeong [1], Broder *et al* [12], and Faloutsos, Faloutsos, and Faloutsos [19] have demonstrated that in the World Wide Web/Internet the proportion of vertices of a given degree follows an approximate inverse power law, which means that 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 [18] do not have power law degree sequences, so they are not suitable as models for these networks. This has driven the development of various alternative models of random graphs.

One approach is to generate graphs with a prescribed degree sequence (or prescribed expected degree sequence). This is proposed as a model for the web graph by Aiello, Chung, and Lu in [2].

An alternative approach, which we will follow in this paper, is to sample graphs via some generative procedure which yields a power law distribution. There is a long history of such models, outlined in the survey by Mitzenmacher [25]. 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. It dates back to Yule [31] and Simon [29]. It was proposed as a random graph model for the web by Barabási and Albert [4] and by Kumar, Raghavan, Rajagopolan, Sivakumar, Tomkins and Upfal [23]. Bollobás and Riordan [8] showed that at time n, with high probability (meaning with probability tending to 1 as n tends to ∞ and abbreviated **whp**), the diameter of this graph is asymptotically equal to $\frac{\log n}{\log \log n}$. Bollobás, Riordan, Spencer and Tusnády [11] showed that the degree sequence of this graph follows a power law distribution **whp**.

An evolving network such as a P2P network sometimes loses vertices. Bollobás and Riordan [9], [10] consider the effect of deleting vertices from the basic preferential attachment model of [4], [8],

*Supported by NSF Grant CCR-0200945

†Department of Mathematical Sciences, Carnegie Mellon University.

after the graph has been generated. Cooper, Frieze, and Vera [17] consider the effect of random edge and vertex deletion *while* the graph is generated. Chung and Lu [13] independently consider a similar model. In this paper we also consider the deletion of vertices while the graph is generated, but the deletions are *adversarial*, not random. In our model there is an (adaptive) adversary who decides which vertices to delete after each time step.

We will study process $\mathcal{P}$, which generates a sequence of graphs $G_t = (V_t, E_t)$, for $t = 1, 2, \ldots, n$. It is defined as follows:

Formal Definition of Process $\mathcal{P}$

Time step $t = 0$: $G_0 = (\emptyset, \emptyset)$.

Time step $t \geq 1$: We add vertex x_t to G_{t-1}.
If $E(G_{t-1})$ is empty, we add m loops incident to x_t.
Otherwise: Add m random edges $(x_t, y_i)_{i=1,2,\ldots,m}$ incident with x_t, where each y_i is chosen from V_{t-1} by *preferential attachment*, meaning for $v \in V_{t-1}$,

$$\mathbf{Pr}(y_i = v) = \frac{\deg_{t-1}(v)}{2|E_{t-1}|},$$

where $\deg_{t-1}(v)$ denotes the degree in G_{t-1}.
After the addition of x_t and the m edges, the adversary chooses a (possible empty) set of vertices and deletes all of them. The adversary does not have any knowledge of future random bits.

The only constraint the adversary has is that by time n the number of vertices he or she has deleted is at most δn, where δ is a sufficiently small constant.

We follow the convention of counting both ends of a loop when counting degree, so the degree of an isolated vertex with a single loop is 2.

2 Results.

All the asymptotic notation is with respect to n, and all other parameters are considered to be fixed.

Theorem 1. *For any sufficiently small constant δ there exists a sufficiently large constant $m = m(\delta)$ and a constant $\theta = \theta(\delta, m)$ such that* **whp** *G_n has a "giant" connected component with size at least θn.*

In the theorem above, the constants are phrased to indicate the suspected relationship, although we do not attempt to optimize them. Our unoptimized calculations work for $\delta \leq 1/50$ and $m \geq \delta^{-2} \times 10^8$ and $\theta = 1/30$.

The proof of Theorem 1 is based on an idea developed by Bollobas and Riordan in [10]. There they couple the graph G_n with $G(n, p)$, the Bernoulli random graph, which has vertex set $[n]$ and each pair of vertices appears as an edge independently with probability p. We couple a carefully chosen induced subgraph of G_n with $G(n', p)$.

To describe the induced subgraph in our coupling, we now make a few definitions. We say that a vertex v of G_t is *good* if it was created after time $t/2$ and the number of its *original edges* that remain undeleted exceeds $m/6$. By original edges of v, we mean the m edges that were created when v was added. Let Γ_t denote the set of good vertices of G_t and $\gamma_t = |\Gamma_t|$. We say that a vertex of G_t is *bad* if it is not good. Notice that once a vertex becomes bad it remains bad for the rest of the process. On the other hand, a vertex that was good at time t_1 can become bad at a later time t_2, simply because it was created at a time before $t_2/2$.

Let

$$p = \frac{m}{1500n}$$

and let $\sim$ denote "has the same distribution as".

Theorem 2. *For any sufficiently small constant δ there exists a sufficiently large constant $m = m(\delta)$ such that we can couple the construction of G_n and random graph H_n, with vertex set Γ_n, such that $H_n \sim G(\gamma_n, p)$ and* **whp** $|E(H_n) \setminus E(G_n)| \leq e^{-\delta^2 m/10^6} mn$.

In Section 4 we prove Theorem 2. In Section 5 we prove a lower bound on the number of good vertices, a key ingredient for the proof of Theorem 1.

3 Proof of Theorem 1.

We will prove the following two lemmas in Section 5.

Lemma 1. *Let G obtained by deleting fewer than $n/100$ edges from a realization of $G_{n,c/n}$. If $c \geq 10$ then* **whp** *G has a component of size at least $n/3$.*

Lemma 2. **Whp**, *for all t with $n/2 < t \leq n$ we have $\gamma_t \geq \frac{t}{10}$.*

With these lemmas, the proof of Theorem 1 is only a few lines:

Let $G = G_n$ and $H = G(\gamma_n, p)$ be the graphs constructed in Theorem 2. Let $G' = G \cap H$. Then $E(H) \setminus E(G') = E(H) \setminus E(G)$ and so **whp** $|E(H) \setminus E(G')| \leq e^{-\delta^2 m/10^6} mn$. By Lemma 2, **whp** $|G'| = \gamma_n \geq n/10$. Since m is large enough, $p = m/1500n > 10/\gamma_n$ and $e^{-\delta^2 m/10^6} mn < n/1000 \leq \gamma_n/100$. Then, by Lemma 1, **whp** G' (and therefore G) has a component of size at least $|G'|/3 \geq n/30$. □

4 The Coupling: Proof of Theorem 2.

We construct $G[k] \sim G_k$ and $H[k] \sim G(\gamma_k, p)$ for $k \geq n/2$ inductively. $G[k]$ will be constructed by following the definition of the process $\mathcal{P}$ and $H[k]$ will be constructed by coupling its construction with the construction of $G[k]$.

For $k = n/2$, we only make the size of $H[k]$ correct and do not try to make the edge structure look like $G[k]$; we just take $H[n/2]$ to be an independent copy of $G(\gamma_{n/2}, p)$ with vertex set $\Gamma_{n/2}$.

For $k > n/2$, having constructed $G[k]$ and $H[k]$ with $G[k] \sim G_k$ and $H[k] \sim G(\gamma_k, p)$, we construct $G[k+1]$ and $H[k+1]$ as follows: Let $G[k] = (V_k, E_k)$, and let $\nu_k = |V_k|$, $\eta_k = |E_k|$ and recall that the number of good vertices is denoted $\gamma_k = |\Gamma_k|$.

If $\gamma_k < \frac{k}{10}$ then we call this a *failure of type 0* and generate $G[n]$ and $H[n]$ independently. (By Lemma 2 the probability of occurrence of this event is $o(1)$.)

Otherwise we have $\gamma_k \geq \frac{k}{10}$. In this case, to construct $G[k+1]$ process $\mathcal{P}$ adds vertex x_{k+1} to $G[k]$ and links it to vertices $t_1, \ldots, t_m \in V_k$ chosen according to the preferential attachment rule. To construct $H[k+1]$, let $\{t_1, \ldots, t_r\} = \{t_1, \ldots, t_m\} \cap \Gamma_k$ be the subset of selected vertices that are good at time k. Let $\epsilon_0 = 1/120$. If r, the number of good vertices selected, is less than $(1-\delta)\epsilon_0 m$ then we call this a *type 1 failure* and generate $H[k+1]$ by joining x_{k+1} to each vertex in $H[k]$ independently with probability p.

Since the number of good vertices $|\Gamma_k| = \gamma_k \geq k/10$ and any $v \in \Gamma_k$ is still incident to at least $m/6$ of its original edges and $|E(G[k])| \leq mk$, we have

$$\mathbf{Pr}\left[t_i \in \Gamma_k\right] = \sum_{v \in \Gamma_k} \frac{\deg_{G[k]}(v)}{2|E(G[k])|} \geq \frac{k}{10} \frac{m}{6} \frac{1}{2mk} = \epsilon_0.$$

So, by comparing r with a Binomial random variable, we obtain an exponential upper bound on the probability of a type 1 failure:

$$\mathbf{Pr}\left[r \leq m\epsilon_0(1-\delta)\right] \leq \mathbf{Pr}\left[\mathrm{Bi}(m, \epsilon_0) \leq (1-\delta)m\epsilon_0\right] \leq e^{-\delta^2 \epsilon_0 m/2}.$$

If we do not have a type 1 failure, then for every $i = 1, \ldots, m$ and for every $v \in \Gamma_k$,

$$\mathbf{Pr}\left[t_i = v\right] = \frac{\deg_{G[k]}(v)}{2|E(G[k])|} \geq \frac{m}{12mk} = \frac{1}{12k}$$

For each $i = 1, \ldots, r$ we choose $s_i \in \Gamma_k \cup \{\perp\}$ such that for each $v \in \Gamma_k$ we have $\mathbf{Pr}\left[s_i = v\right] = \frac{1}{12k}$. We couple the selection of the s_i's with the selection of the t_i's such that if $s_i \neq \perp$ then $s_i = t_i$. Let $X = \left|\{i : s_i \neq \perp\}\right|$. Then $X \sim \mathrm{Bi}\left(r, \frac{\gamma_k}{12k}\right)$. Let $Y \sim \mathrm{Bi}\left(\gamma_k, p\right)$. Then

$$\mathbf{E}[X] = r\frac{\gamma_k}{12k} \geq (1-\delta)\epsilon_0 m \frac{\gamma_k}{12n} \geq (1+\delta)\gamma_k p \geq (1+\delta)\mathbf{E}[Y].$$

Since $\mathbf{E}[X] \geq (1+\delta)\mathbf{E}[Y]$, the probability that $Y \leq X$ is at most the probability that X or Y deviates from its mean by a factor of $\delta/2$. And, since

$$\mathbf{E}[X] \geq \mathbf{E}[Y] = \gamma_k p \geq \frac{k}{10} \frac{m}{1500n} \geq \frac{m}{30000},$$

by Chernoff's bound, $\mathbf{Pr}[X \leq (1-\delta/2)\mathbf{E}[X]]$ and $\mathbf{Pr}[Y \geq (1+\delta/2)\mathbf{E}[Y]]$ are at most $e^{-\delta^2 m/10^6}$. We say we have a *type 2 failure* if $Y > X$, so we have a type 2 failure with probability at most $2e^{-\delta^2 m/10^6}$.

Conditioning on X, the s_i's are a subset S_1 of Γ_k of size X chosen uniformly at random from all of these subsets. We choose S_2 uniformly at random between all the subsets of Γ_k of size Y, coupling the selection of S_2 to the selection of S_1 such that $S_2 \subseteq S_1$ when $Y \leq X$. Now, to generate $H[k+1]$, we join x_{k+1} to every vertex in S_2 (deterministically).

After the adversary deletes a (possible empty) set of vertices in $G[k]$, we delete all the vertices $H[k]$ that don't belong to Γ_{k+1}, possibly including $x_{\lfloor (k+1)/2 \rfloor}$, simply because of its age.

For $k \geq n/2$ this process yields an $H[k]$ with vertex set Γ_k and identically distributed with $G(\gamma_k, p)$, so we have $H[n] \sim G\left(\gamma_n, p\right)$.

We call an edge e in $H[n]$ *misplaced* if e is not an edge of $G[n]$. We are interested in bounding the number of misplaced edges. Misplaced edges can only be created when we have a failure. The probability of having a type 1 or 2 failure at step k is at most $3e^{-\delta^2 m/10^6}$. Let M_k denote the number of misplaced edges created between good vertices when we have a failure of type 1 or 2 at step k. Then M_k is stochastically smaller than $Y \sim \mathrm{Bi}(\gamma_k, p)$, a binomial random variable with $\mathbf{E}[Y] = \gamma_k p \le (k/2)(m/1500n) \le m/3000$.

Let M denote the total number of misplaced edges at time n. Then $M \le \sum_{k=n/2}^{n} M_k$ and therefore

$$\begin{aligned}
\mathbf{E}[M] &\le \sum_{k=n/2}^{n} \mathbf{E}[M_k] \\
&\le \sum_{k=n/2}^{n} 3e^{-\delta^2 m/10^6} m/3000 \\
&= e^{-\delta^2 m/10^6} mn/2000.
\end{aligned}$$

The bounds we have estimated for M_k are independent of the history at each step in the construction so Chernoff's bound is sufficient to prove that M is concentrated around its mean. □

5 Bounding the number of good vertices: Proof of Lemma 2.

We now prove Lemma 2, which is restated here for convenience.

Lemma 2. Whp, *for all t with $n/2 < t \le n$ we have $\gamma_t \ge \frac{t}{10}$.*

Proof Let z_t denote the total number of edges created after time $t/2$ that have been deleted by the adversary, up to time t. Let ν'_t and η'_t be the number of vertices and edges respectively in G_t that were created after time $t/2$. Notice that $\eta'_t = \frac{1}{2}mt - z_t$ and $\nu'_t \le t/2$. Also, since each vertex contributes at most m edges, and bad vertices (not in Γ_t) contribute at most $m/6$ edges, we have $\eta'_t \le m\gamma_t + \frac{m}{6}(\nu'_t - \gamma_t)$. So

$$\gamma_t \ge \frac{6\eta'_t - m\nu'_t}{5m} \ge \frac{3mt - 6z_t - mt/2}{5m} = \frac{t}{2} - \frac{6z_t}{5m},$$

So if $z_t \le mt/3$ then $\gamma_t \ge t/10$. Thus, to prove the lemma, it is sufficient to show that

$$\mathbf{Pr}\left[z_t \ge \frac{mt}{3}\right] \le e^{-\delta^2 mn/10}. \tag{1}$$

To show Inequality (1) holds, we will couple our process with another process $\mathcal{P}^\star$ in which the adversary deletes no vertices until time t and then deletes the same set of vertices as in $\mathcal{P}$.

Fix $t \ge n/2$. We begin by showing that with

$$t_0 = 1000\delta n,$$

$$\mathbf{Pr}\left[z_t(\mathcal{P}) \ge z_t(\mathcal{P}^\star) + mt_0\right] = O\left(ne^{-\delta^2 mn/7}\right). \tag{2}$$

Generate G_s for $s = 1, \dots, t$ by process $\mathcal{P}$. Let $D_1, D_2, \dots$ be the sequence of vertex sets deleted by the adversary in this realization of $\mathcal{P}$. Let $D = \bigcup_{\tau=1}^{t} D_\tau$ denote the set of vertices deleted by the adversary by time t.

We define $G_s^\star$ inductively. To begin, fuuubgenerate $G_{t_0}^\star$ according to preferential attachment. Then, for every s with $t_0 \leq s < t$: For $G_s = (V_s, E_s)$ and $G_s^\star = (V_s^\star, E_s^\star)$, let $X_s = E_s^\star \setminus E_s$ be the set of edges that have been deleted by the adversary's moves.

Selecting a vertex by preferential attachment is equivalent to choosing an edge uniformly at random and then randomly selecting one of the end points of the edge. So we can view the transition from G_s to G_{s+1} as adding x_{s+1} to G_s, and then choosing m edges $e_1, \ldots, e_m$. Then for each i select a random endpoint y_i of e_i, and add an edge between x_{s+1} and y_i.

To construct $G_{s+1}^\star$, we first add x_{s+1} to $G_{s+1}^\star$, and then to choose $y_1^\star, \ldots, y_m^\star$ we apply the following procedure, for each i:

- With probability $1 - |X_s|/(ms)$ we set $e_i^\star = e_i$ and $y_i^\star = y_i$
- With probability $|X_s|/(ms)$, we choose $e_i^\star$ uniformly at random from X_s. Notice that $e_i^\star$ has already been deleted from G_s by the adversary and therefore it is incident to at least one deleted vertex, $v_i \in D$. Now, we randomly choose $y_i^\star$ from the two end points of $e_i^\star$. If the total degree T_s of the vertices $V_s \cap D$ that $\mathcal{P}$ will delete in the future is at most $ms/2$ then $\mathbf{Pr}\left[y_i \in D\right] \leq 1/2$ and we couple the (random) decisions in such way that if y_i is going to be deleted by time t then $y_i^\star = v_i$. Otherwise we say we have a *failure* and choose $y_i^\star$ independently of y_i.

In the coupling, after time t_0 and before the first failure, an edge incident with x_{s+1} and destined for deletion in $\mathcal{P}$ is matched with an edge incident with x_{s+1} and destined for deletion in $\mathcal{P}^\star$. So, until the first failure, T_s is bounded by $T_s^\star$, the corresponding total degree of $V_s \cap D$ in $G_s^\star$. In Lemma 3 below, we prove that $\mathbf{Pr}\left[T_s^\star > sm/2\right] = O\left(e^{-\delta^2 mn/6}\right)$ and therefore the probability of having a failure is $O\left(ne^{-\delta^2 mn/6}\right) = O\left(e^{-\delta^2 mn/7}\right)$.

To repeat, if there is no failure and if e_i is deleted in $\mathcal{P}$ before time t we have two possibilities: x_{s+1} is deleted or y_i is deleted. In either case, x_{s+1} or $y_i^\star$ will be deleted by time t in $\mathcal{P}^\star$ and therefore $e_i^\star$ will be deleted, and Equation (2) follows.

We will show that

$$\mathbf{Pr}\left[z_t(\mathcal{P}^\star) \geq \frac{mt}{4}\right] \leq O(e^{-\delta^2 mn}), \tag{3}$$

and then Inequality (1) follows from Equation (2).

To prove Inequality (3) let s be such that $t/2 \leq s \leq t$ and $x_s \notin D$. We want to upper bound the probability in the process $\mathcal{P}^\star$ that an edge created at time s chooses its end point in D. For $i = 1, \ldots, m$,

$$\mathbf{Pr}\left[y_i^\star \in D \,\middle|\, T_s^\star\right] = \frac{T_s^\star}{2ms}.$$

By Lemma 3 (below), we have $\mathbf{Pr}\left[T_s^\star \geq mt/2\right] \leq O(e^{-\delta^2 mn})$ so

$$\mathbf{Pr}\left[y_i^* \in D\right] \leq \frac{1}{4} + o(1).$$

Therefore $z_t(\mathcal{P}^\star)$ is stochastically dominated by Bi $\left(\frac{mt}{2}, \frac{1}{4} + o(1)\right)$. Inequality (3) now follows from Chernoff's bound. This completes the proof of Lemma 2. $\square$

Lemma 3. *Let $A \subset \{x_1, \ldots, x_t\}$, with $|A| \leq \delta n$. Let $t \geq 1000\delta n$ and let G_t be a graph generated by preferential attachment (i.e. the process $\mathcal{P}$, but without an adversary). Let T_A denote the total degree of the vertices in A. Then*

$$\mathbf{Pr}\left[\exists A:\ T_A \geq mt/2\right] = O\left(e^{-\delta^2 mn}\right).$$

Proof Let $A' = \{x_1, \ldots, x_{\delta n}\}$ be the set of the oldest δn vertices. We can couple the construction of G_t with G'_t, another graph generated by preferential attachment, such that $T_{A'} \geq T_A$. Therefore $\mathbf{Pr}\left[T_A \geq mt\right] \leq \mathbf{Pr}\left[T_{A'} \geq mt\right]$, and we can assume $A = A'$.

Now we consider the process $\mathcal{P}$ in δ^{-1} rounds, Each round consisting of δn steps. Let T_i be the total degree of A at the end of the ith round. Notice that $T_1 = 2\delta mn$ and $T_2 \leq 3\delta mn$. For $i \geq 2$, fix s with $i\delta n < s \leq (i+1)\delta n$. Then the probability that x_s chooses a vertex in A is at most $\frac{T_i + \delta mn}{2i\delta mn}$. So given T_i, the difference $T_{i+1} - T_i$ is stochastically dominated by $Y_i \sim \mathrm{Bi}\left(\delta mn, \frac{T_i+\delta mn}{2i\delta mn}\right)$.

Therefore, for $i \geq 2$,

$$\begin{aligned}\mathbf{Pr}\left[T_{i+1} > 3i^{2/3}\delta mn\right] &\leq \mathbf{Pr}\left[T_{i+1} \geq 3i^{2/3}\delta mn \,\middle|\, T_i \leq 3(i-1)^{2/3}\delta mn\right] \\ &\qquad + \mathbf{Pr}\left[T_i > 3(i-1)^{2/3}\delta mn\right] \\ &\leq \mathbf{Pr}\left[T_{i+1} \geq 3i^{2/3}\delta mn \,\middle|\, T_i = 3(i-1)^{2/3}\delta mn\right] \\ &\qquad + \mathbf{Pr}\left[T_i > 3(i-1)^{2/3}\delta mn\right].\end{aligned}$$

For $i \geq 2$,

$$\mathbf{E}[Y_{i+1} | T_i = 3(i-1)^{2/3}\delta mn] = \left(\frac{3(i-1)^{2/3}+1}{2i}\right)\delta mn \leq \frac{4}{3} i^{-1/3}\delta mn$$

and, since $3(i^{2/3} - (i-1)^{2/3}) \leq 2i^{-1/3}$ and $i \leq \delta^{-1}$, by Chernoff's bound we have

$$\mathbf{Pr}\left[T_{i+1} \geq 3i^{2/3}\delta mn \,\middle|\, T_i = 3(i-1)^{2/3}\delta mn\right] \leq e^{-\delta^{4/3}mn/9}.$$

Hence, for any $k \leq \delta^{-1}$,

$$\mathbf{Pr}\left[T_k > 3(k-1)^{2/3}\delta mn\right] \leq \sum_{i=2}^{k-2} e^{-\delta^{4/3}mn/9} \leq e^{-2\delta^2 mn}.$$

Now, if $t \geq t_0$ then $k = \lfloor \frac{t}{\delta n} \rfloor \geq 10^3$ and so

$$3(k-1)^{2/3}\delta mn \leq tm/2.$$

Thus

$$\mathbf{Pr}\left[T_t \geq tm/2\right] \leq e^{-2\delta^2 mn}.$$

We inflate the above by $\binom{n}{\delta n}$ to get the bound in the lemma. □

Proof of Lemma 1 If after deleting $n/100$ edges the maximum component size is at most $n/3$ then $G_{n,c/n}$ contains a set S of size $n/3 \leq s \leq n/2$ such that there are at most $n/100$ edges joining

S to $V \setminus S$. The expected number of edges across this cut is $s(n-s)c/n$ so when $1-\epsilon = \frac{9}{200c}$ we have $n/100 \leq (1-\epsilon)s(n-s)c/n$ and by applying the union bound and Chernoff's bound we have

$$\begin{aligned}\mathbf{Pr}\left[\exists S\right] &\leq \sum_{s=n/3}^{n/2} \binom{n}{s} e^{-\epsilon^2 s(n-s)c/(2n)} \\ &\leq \sum_{s=n/3}^{n/2} \left(\frac{ne}{s} e^{-\epsilon^2 (n-s)c/(2n)}\right)^s \\ &= o(1).\end{aligned}$$

□

References

[1] R. Albert and A. Barabási and H. Jeong, Diameter of the world wide web, *Nature* 401 (1999) 103-131,

[2] W. Aiello and F. R. K. Chung and L. Lu, A random graph model for massive graphs, *Proc. of the 32nd Annual ACM Symposium on the Theory of Computing*, (2000) 171–180.

[3] W. Aiello and F. R. K. Chung and L. Lu, Random Evolution in Massive Graphs, *Proc. of IEEE Symposium on Foundations of Computer Science* (2001) 510-519.

[4] A. Barabasi and R. Albert, Emergence of scaling in random networks, *Science* 286 (1999) 509-512.

[5] N. Berger and B. Bollobas and C. Borgs and J. Chayes and O. Riordan, Degree distribution of the FKP network model, *Proc. of the 30th International Colloquium of Automata, Languages and Programming*, (2003,) 725-738.

[6] N. Berger and C. Borgs and J. Chayes and R. D'Souza and R. D. Kleinberg, Competition-induced preferential attachment.

[7] B. Bollobás and O. Riordan, Mathematical Results on Scale-free Random Graphs, in *Handbook of Graphs and Networks*, Wiley-VCH, Berlin, 2002.

[8] B. Bollobás and O. Riordan, The diameter of a scale-free random graph, *Combinatorica* 24 (2004) 5-34.

[9] B. Bollobás and O. Riordan, Robustness and vulnerability of scale-free random graphs, *Internet Mathematics* 1 (2003) 1-35.

[10] B. Bollobás and O. Riordan, Coupling scale free and classical random graphs.

[11] B. Bollobás and O. Riordan and J. Spencer and G. Tusanády, The degree sequence of a scale-free random graph process, *Random Structures and Algorithms* 18 (2001) 279-290.

[12] A. Broder, R. Kumar, F. Maghoul, P. Raghavan, S. Rajagopalan, R. Stata, A. Tomkins and J. Wiener, Graph structure in the web, *Proc. of the 9th Intl. World Wide Web Conference* (2002) 309-320.

[13] F.R.K. Chung and L. Lu, Coupling on-line and off-line analyses for random power law graphs.

[14] F.R.K. Chung and L. Lu and V. Vu, Eigenvalues of random power law graphs, *Annals of Combinatorics* 7 (2003) 21-33.

[15] F.R.K. Chung and L. Lu and V. Vu, The spectra of random graphs with expected degrees, *Proc. of National Academy of Sciences* 100 (2003) 6313-6318.

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

[17] C. Cooper, A. Frieze and J. Vera, Random Deletion in a Scale Free Random Graph Process, *Internet Mathematics*, to appear.

[18] P. Erdős and A. Rényi, On random graphs I, *Publicationes Mathematicae Debrecen* 6 (1959) 290-297.

[19] M. Faloutsos and P. Faloutsos and C. Faloutsos, On Power-law Relationships of the Internet Topology, *SIGCOMM* (1999) 251-262 .

[20] Alex Fabrikant and Elias Koutsoupias and Christos H. Papadimitriou, Heuristically Optimized Trade-Offs: A New Paradigm for Power Laws in the Internet, *Proc. of 29th International Colloquium of Automata, Languages and Programming* (2002).

[21] B. Hayes, Graph theory in practice: Part II, American Scientist 88 (2000) 104-109.

[22] J. M. Kleinberg and R. Kumar and P. Raghavan and S. Rajagopalan and A. S. Tomkins, The Web as a Graph: Measurements, Models and Methods, *Lecture Notes in Computer Science* (1999).

[23] R. Kumar and P. Raghavan and S. Rajagopalan and D. Sivakumar and A. Tomkins and E. Upfal, Stochastic Models for the Web Graph, *Proc. IEEE Symposium on Foundations of Computer Science* (2000).

[24] C. J. H. McDiarmid, Concentration, in Probabilistic methods in algorithmic discrete mathematics, (1998) 195-248.

[25] M. Mitzenmacher, A brief history of generative models for power law and lognormal distributions.

[26] M. Mihail and C. H. Papadimitriou, On the Eigenvalue Power Law, *Proc. of 6th International Workshop on Randomization and Approximation Techniques* (2002) 254-262.

[27] M. Mihail, C. H. Papadimitriou and A. Saberi, On Certain Connectivity Properties of the Internet Topology, *Proc. IEEE Symposium on Foundations of Computer Science* (2003).

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

[29] H. A. Simon, On a class of skew distribution functions, *Biometrika* 42 (1955) 425-440.

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

[31] 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.