Memoryless Rules for Achlioptas Processes

Andrew Beveridge Tom Bohman* Alan Frieze† Oleg Pikhurko‡

Department of Mathematical Sciences
Carnegie Mellon University
Pittsburgh, PA 15213

March 4, 2007

Abstract

In an Achlioptas process two random pairs of $\{1, \ldots, n\}$ arrive in each round and the player has to choose one of them. We study the very restrictive version where player's decisions cannot depend on the previous history and only one vertex from the two random edges is revealed.

We prove that the player can create a giant component in

$$(2\sqrt{5} - 4 + o(1))n = (0.4721\ldots + o(1))n$$

rounds and this is best possible. On the other hand, if the player wants to delay the appearence of a giant, then the optimal bound is $(1/2 + o(1))n$, the same as in the Erdős-Rényi model.

1 Introduction

Let $[n] = \{1, \ldots, n\}$ be the vertex set. There are $m = m(n)$ rounds. In an *Achlioptas process* two random edges arrive in each round and the player, called *Paul*, has to accept one of them (and reject the other). It is not surprising that Paul can ensure that the obtained graph G_m differs from a typical Erdős-Rényi random graph with the same edge density. The property most frequently studied in this context is the time when a *giant component* appears **whp**, see [3, 4, 5, 6, 7, 11, 14]. Here, **whp** is an abbreviation for *with high probability*, that is, with probability $1 - o(1)$ as $n \to \infty$. A *giant component* is a component of size at least κn where $\kappa > 0$ is an absolute constant.

*Research supported in part by NSF grant DMS-0401147
†Research supported in part by NSF grant CCR-0502793
‡Research supported in part by NSF grant DMS-0457512

Unfortunately, it is still an open question what is the earliest/latest time of the birth of a giant that Paul can ensure. The asnwer is unknown even for some simple specific Paul's strategies, such as for example the Product Rule, see [3].

Here, we study the class of more restrictive strategies, which we call *memoryless rules*. Namely, Paul's choice at Round i is a function of the current two random edges, $D_i = \{x_i, y_i\}$ and $D_i' = \{x_i', y_i'\}$, and the number i. It does not depend on what happened in Rounds 1 to $i-1$, that is, Paul does not remember any previous history. Such restrictions may appear when our online algorithm has to make decisions fast and has limited computational resources, so processing or storing the current graph is infeasible or impractical.

We obtain the complete answer for the even more restrictive case when Paul can see only one vertex x_i from the first random edge. Here we assume that the first random edge $\{x_i, y_i\}$ is generated by taking $x_i \in [n]$ and then $y_i \in [n] \setminus \{x_i\}$. (An alternative non-equivalent setting, namely to reveal $\min\{x_i, y_i\}$, is not studied here.) Any memoryless algorithm of this type is described by the sequence of the *acceptance sets* $A_1, \ldots, A_m \subseteq [n]$, where Paul selects D_i if and only if the revealed vertex x_i belongs to A_i. The graph constructed is denoted by G_m.

Here is one example: suppose Paul accepts D_i if and only if $x_i \le n/2$. In Section 3 we will prove the following result about this rule.

Theorem 1 *Fix a constant $\varepsilon > 0$. Let $m = \lfloor (c+\varepsilon)n \rfloor$, where*

$$c = 2\sqrt{5} - 4 = 0.4721... \tag{1}$$

and let each A_i be equal to $\{1, \ldots, \lfloor n/2 \rfloor\}$, $i \in [m]$. Then, **whp** *the graph obtained has a* unique *giant component. Furthermore, the second largest component is of size $O((\ln n)^3)$.*

We show in Section 3 that this simple rule is asymptotically best possible.

Theorem 2 *Let c, ε be as in (1). For any sequence $A_1, \ldots, A_m$, where $m = \lfloor (c-\varepsilon)n \rfloor$,* **whp** *the maximum component of the generated graph has at most $O(\ln n)$ vertices.*

Clearly, Paul can generate a genuine Erdős-Rényi graph $\mathcal{G}(n,m)$ by always accepting the first edge. It is well-known that the (unique) giant component appears after $(1/2 + o(1))n$ rounds. Rather surprisingly, we prove in Section 4 that the constant 1/2 is best possible.

Theorem 3 *For every constant $\varepsilon > 0$ and any $A_1, \ldots, A_m \subset [n]$, $m = \lfloor (1/2+\varepsilon)n \rfloor$,* **whp** *the obtained graph has a giant component.*

Let us very briefly describe the main ideas of the proofs as well as define a few parameters that we will use later. Let m and $A_1, \ldots, A_m \subseteq [n]$ be given and G_m be the graph constructed by the corresponding memoryless rule. The expected number of rounds in which any given (unordered) pair $\{x, y\}$ is accepted is

$$p_{xy} = \frac{1}{\binom{n}{2}} \sum_{i=1}^{m} \left(\frac{\mathbf{1}_{x\in A_i} + \mathbf{1}_{y\in A_i}}{2} + \frac{n - |A_i|}{n} \right), \tag{2}$$

where the indicator function $\mathbf{1}_{x\in A_i}$ is 1 if $x \in A_i$ and 0 otherwise. For $z \in [n]$ let $\gamma_z = |\{i : \ z \in A_i\}|$. We can write as $p_{xy} = p_x + p_y$, where

$$p_x = \frac{m-\gamma}{n(n-1)} + \frac{1}{n(n-1)} \sum_{i=1}^{m} \mathbf{1}_{x\in A_i} = \frac{m-\gamma+\gamma_x}{n(n-1)}, \tag{3}$$

and $\gamma = \frac{1}{n}\sum_{i=1}^{m} |A_i| = \frac{1}{n}\sum_{z=1}^{n} \gamma_z$.

The probability that $\{x, y\}$ is chosen at least i times is $O(n^{-i})$ (assuming $m = O(n)$). It follows that the probability that $\{x, y\} \in E(G_m)$ is $p_{xy}(1 + O(n^{-1}))$. Thus one might expect that the random graph that contains $\{x, y\}$ with probability p_{xy} independently of the other edges is a good approximation to the obtained graph G_m. Such *inhomogeneous* random graph models were studied by Alon [1], Söderberg [13], Bollobás, Janson, and Riordan [8], and others. However, given some dependences between the edges of G_m, we have found it more convenient to adopt the branching process approach. We proceed in a manner similar to the proof given in Janson, Łuczak, and Rucinski's book [12, Section 5.2]. Namely, to estimate the order of the connectivity component C of G_m that contains a given vertex x, we approximate the breadth-first search from x by an appropriate ideal branching process. A new technical difficulty in comparison to [12, Section 5.2] that we face is that the vertices are not homogeneous, so we have to use multi-type branching processes. All details appear in Section 2.

The case when Paul can use more information about the two random edges than just one vertex seems to be far more complicated.

In what follows, any inequalities will assume that n is sufficiently large.

2 Approximation by Branching Processes

Here we will relate the appearence of a giant component to the extinction probability of certain branching processes. We refer the Reader to Athreya and Ney [2] for all definitions related to branching processes, in particular to Chapter V in [2] which deals with multi-type processes.

It will be the case that in the ideal multi-type branching processes we will consider the size of the offspring $X_{i,j}$ of Type j that a particle of Type i produces has Poisson distribution and, for every i, the random variables $X_{i,j}$ are independent. Let us call such a process a *Poisson branching process*. Thus all the transition probabilities can be encoded by the *mean matrix* M, where $M_{i,j} = \mathbf{E}(X_{i,j})$.

Unfortunately, we cannot just take (p_{xy}) for M, since our proof will require that the square matrix M has a bounded number of rows and that all entries are strictly positive.

Therefore, in order to prove a sufficient condition for the existence of a giant component we proceed as follows.

Let m and $A_1, \ldots, A_m \subseteq [n]$ be given, all depending on n. Assume that $m = \theta_0 n$ where $\theta_0 > 0$ is a positive constant. When necessary, this dependence on n will be emphasized by a superscript, as in $A_1 = A_1^{(n)}$. Without loss of generality we can assume that, for example, $n/4 \leq m \leq n$ because outside this range the existence of a giant is **whp** predermined irrespective of what Paul does, see Bohman and Kim [6] and Bohman and Kravitz [7].

Let

$$p_{\max} = \frac{2m}{n(n-1)}.$$

We have $p_x < p_{\max}$ for every $x \in [n]$, where p_x is defined as in (3).

2.1 A Sufficient Condition For the Existence of Giant

In what follows we will introduce small positive constants $c_1, c_2, \ldots,$. They will not be specified exactly, but we will indicate there relative sizes.

Let $\nu_0 = (\ln n)^2$ and let $k = \lceil 1/c_1 \rceil$. Define

$$V_i = \{x \in [n] \mid ip_{\max} \leq kp_x < (i+1)p_{\max}\}, \quad i \in [0, k-1],$$

and

$$I = \left\{i \in [k-1] \mid |V_i| \geq n/k^2\right\}.$$

Note that, by the definition, $0 \notin I \subseteq \{1, 2, \ldots, k-1\}$.
Let

$$V = \bigcup_{i \in I} V_i.$$

We show next that $|V|$ is close to n. Indeed, let us estimate $n_0 = |V_0|$ in this case. We have

$$\frac{n_0 p_{\max}}{k} > \sum_{x \in V_0} p_x = \frac{1}{n(n-1)} \sum_{i=1}^{m} |V_0 \cap A_i| + \frac{n_0 m}{n(n-1)} - \frac{n_0}{n^2(n-1)} \sum_{i=1}^{m} |A_i|.$$

Using the estimate $|A_i| \leq |V_0 \cap A_i| + n - n_0$, we get

$$\frac{n_0 \times 2m}{kn(n-1)} > \frac{n_0^2 m}{n^2(n-1)} + \frac{n - n_0}{n^2(n-1)} \sum_{i=1}^{m} |V_0 \cap A_i| \geq \frac{n_0^2 m}{n^2(n-1)},$$

which shows that $n_0 < \frac{2n}{k}$. Since, $|\cup_{i \in [k] \setminus I} V_i| \leq k(n/k^2) = n/k$ and $\cup_{i=0}^{k} V_i = [n]$, we have

$$|V| \geq \left(1 - \frac{3}{k}\right) n. \tag{4}$$

For $i,j \in [0,k]$ let

$$M_{i,j} = \max\left\{\frac{(i+j)p_{\max}}{k}|V_j| - c_2, 0\right\}$$

where $c_2 \ll c_1$.

Let $t = |I|$. Let M be the $t \times t$-matrix with rows and columns indexed by I whose entries are $M_{i,j}$. The entries of M are strictly positive.

The definition of I implies that each entry of M is at least $\frac{2p_{\max}}{k} \times \frac{n}{k^2} - c_2 \geq 3\theta_0 c_1^3 - c_2$, that is, it is bounded from 0 by a constant independent of n. The Frobenius Theorem implies that M, as a matrix with strictly positive entries, has a maximal eigenvalue λ_1 which is positive and simple.

Lemma 4 *Suppose that there exists $\theta_1 > 0$ such that $\lambda_1^{(n)} \geq 1 + \theta_1$ for all sufficiently large n. Then* **whp** *G_m contains a unique giant component. Furthermore, the second largest component is of size $O((\ln n)^3)$.*

Proof. Fix $x \in V$. Let us expose the vertex set of the connected component containing x in the following way. Initially mark all vertices as *unsaturated.* We do a breadth first search in $G_m[V]$ starting with the vertex x. Let C_x (initially $\{x\}$) be the current set of vertices we have added to $V(C)$. If $|C_x| \geq \nu_0$, then we stop. Otherwise take an available unsaturated vertex $y \in C_x$ that is closest to x. If no such vertex y is available, then we stop — we know the vertex set of the component containing x. Expose all remaining edges of G_m that that are incident to y together with the indices of the rounds in which they were generated. Add these neighbors of y to R. Mark y as *saturated.* We let $B_x \subseteq C_x$ denote the saturated vertices of C_x. Assign to each vertex of C_x, one of the t types, depending on which of the sets V_i, $i \in I$, it comes from.

Consider the moment when we expose neighbors of some vertex y, say $y \in V_i$. Let $S \subset [m]$ consist of the indices of rounds for which we have not previously exposed the accepted edge. We have $|S| \geq m - \nu_0 - O((\ln n)^2)$. This is because **whp** every set of ν_0 vertices in Γ spans at most $\nu_0 + O(\nu_0/\ln n)$ edges and the maximum degree in Γ is $o(\ln n)$. Here Γ is the graph spanned by all $2m$ edges generated. Furthermore, it is also easy to see that

$$\mathbf{Pr}(C_x \textit{ contains a cycle}) = O(\nu_0^2/n). \tag{5}$$

(The probability that some unsaturated x chooses an unsaturated $y \in C_x$ is $\leq 4m\nu_0/n(n-1)$.)

Notice also that our bound on maximum degree implies that when we stop our process having reached $\geq \nu_0$ vertices we have

$$|C_x| \leq \nu_0 + o(\ln n).$$

Conditioning: At any stage the edges in D_i, D_i', $i \in S$ are random subject to the conditions:

- If $x_i \in A_i$ then $\{x_i, y_i\} \cap B_x = \emptyset$.
- If $x_i \notin A_i$ then $\{x_i', y_i'\} \cap B_x = \emptyset$.

Note that for any $z \in V_j \setminus C_x$ and $s \in S$, the probability that $\{y, z\}$ is selected in Round s is at least as large as the unconditional probability. This is because no outcome (D_s, D_s') that results in the acceptance of $\{y, z\}$ in Round s is ruled out by our conditioning, although some outcomes involving edges meeting C_x can be.

Since the probability that a pair $\{y, z\}$ is selected in any given Round q is trivially at most

$$\mathbf{Pr}\big(D_q = \{y, z\} \vee D_q' = \{y, z\}\big) \le \frac{2}{\binom{n}{2}},$$

the probability that $\{y, z\}$ is selected in an S-round is at least

$$p_{yz} - \frac{2\nu_0}{\binom{n}{2}} \ge \frac{p_{\max}(i+j)}{k} - \frac{5\nu_0}{n^2}. \tag{6}$$

Also, the outcomes of the rounds indexed by S are still independent events. Thus, for $j \in [0, k-1]$, the number Y_j of neighbors of y in $V_j \setminus C_x$, counted with multiplicity, is the sum of $|S|$ independent Boolean random variables, each having expectation at most $\mathbf{Pr}(y \in D_q \cup D_q') \le 4/n$. Also, $\lambda = \mathbf{E}(Y_j \mid \textit{previous history}) \le m(4/n) = O(1)$. It follows from Durrett [9], Chapter 2, Equation (6.5), that

$$\sup_{A \subseteq \mathbb{Z}} |\mathbf{Pr}(Y_j \in A) - \mathbf{Pr}(Po(\lambda) \in A)| \le \frac{16}{n} \tag{7}$$

where $Po(\lambda)$ is Poisson with mean λ. Note next that if $j \in I$,

$$\mathbf{E}(Y_j \mid \textit{previous history}) \ge \left(\frac{p_{\max}(i+j)}{k} - \frac{5\nu_0}{n^2}\right)|V_j \setminus C_x| \ge \\ \left(\frac{p_{\max}(i+j)}{k} - \frac{5\nu_0}{n^2}\right)(|V_j| - \nu_0) \ge M_{i,j} + c_2/2.$$

Note also that for $i, j \in [0, k-1]$

$$\mathbf{E}(Y_j) \le \frac{p_{\max}(i+j+2)}{k}|V_j| \le M_{i,j} + 5\theta_0 c_1 |V_j|/n. \tag{8}$$

We can therefore couple the BFS tree (plus a few extra edges) that we grow from x with the following: Given the current $y \in V_i$ we generate $Po(M_{i,j})$ new Red neighbours for each $j \in I$. Then we generate

$$Po(\xi_j), \xi_j = \mathbf{E}(Y_j \mid \textit{previous history}) - M_{i,j},\ c_2/2 \le \xi_j \le 5\theta_0 c_1$$

Blue neighbours in V_j for each $j \in [0, k]$ and then with probability at most $16/n$ we invoke a *demon* who will add or delete some number of new neighbours. (This demon will **qs**[1]

[1] A sequence of events $\mathcal{E}_n$ is said to occur *quite surely* (**qs**) if $\mathbf{Pr}(\mathcal{E}_n) = 1 - O(n^{-K})$ for any constant K.

add/delete $O(\ln n)$ edges, since the maximum degree is $O(\ln n)$ with this probability). It follows that with probability at least $1-16\nu_0/n$ the Red part of the BFS tree is the same as first ν_0 progeny of the *idealised multi-type branching process*.

It will be seen that this coupling is valid as long as the number of progeny is $o(n)$.

Let $\rho_i = \rho_i^{(n)}$ be the extinction probability of $\mathcal{B}_i^{(n)}$, the ideal branching process given by $M^{(n)}$ that starts with one particle of Type i, $i \in I$. Let us show that each ρ_i is strictly bounded away from 1.

Suppose that this is not true, that is, we can choose $q \in [k]$ and a sequence $(n_l)_{l=1}^{\infty}$ such that $q \in I^{(n_l)}$ for each l and $\lim_{l\to\infty} \rho_q^{(n_l)} = 1$. We can additionally require that $I^{(n_l)}$ does not depend on l. Since the entries of M are bounded by $2kp_{\max}n = O(1)$, we can assume, by taking a subsequence, that $M^{(n_l)}$ tends to a limiting matrix $M^{(\infty)}$, that is, $\lim_{l\to\infty} M_{i,j}^{(n_l)} = M_{i,j}^{(\infty)}$ for any $i, j \in I$.

The matrix $M^{(\infty)}$ is strictly positive and its largest eigenvalue $\lambda_1^{(\infty)}$ is at least $1+\theta_1$ (because it is the maximum of $\|\mathbf{x}^T M^{(\infty)}\mathbf{x}\|$ over all unit vectors $\mathbf{x} \in \mathbb{R}^t$). Hence, [2, Theorem V.3.2], we can assume that $\rho_i^{(\infty)} < 1 - c_3$, $c_3 = c_3(\theta_1)$ for each $i \in I$.

Let $X_{i,r}^{(n)}$ (resp. $X_{i,r}^{(\infty)}$) be the event that $\mathcal{B}_i^{(n)}$ (resp. $\mathcal{B}_i^{(\infty)}$) dies within the first r rounds. We have $\rho_i^{(\infty)} = \lim_{r\to\infty} \mathbf{Pr}(X_{i,r}^{(\infty)})$, since $\cup_{r=1}^{\infty} X_{i,r}^{(\infty)}$ is the extinction event. So we can pick a sufficiently large r so that, in particular, for any $i \in I$

$$\left|\rho_i^{(\infty)} - \mathbf{Pr}(X_{i,r}^{(\infty)})\right| < c_3/4.$$

Also, by [2, Theorem V.6.2]), we can assume that the probability of $\mathcal{B}_i^{(\infty)}$ surviving r levels but having at most, say, $(1+\theta_1)^{r/2}$ particles at level r, is at most $c_3/4$.

Now choose l_0 (depending on r) so that for any $l \ge l_0$ the variation distance between the distribution on the first r levels of branching for $\mathcal{B}_i^{(\infty)}$ and $\mathcal{B}_i^{(n_l)}$ is at most $c_3/4$. For all such l, $\mathcal{B}_i^{(n_l)}$ has at least $(1+\theta_1)^{r/2}$ particles at level r with probability at least $p = 1-c_3/4-\rho_i^{(\infty)}-c_3/4 > c_3/2$. Hence, the extinction probability of $\mathcal{B}_q^{(n_l)}$ can be bounded from above by p_0, the extinction probability of the (single-type) branching process that produces $D = (1+\theta_1)^{r/2}$ children with probability p and none with probability $1-p$. The expected number of progeny of an individual in this process is $p(1+\theta_1)^{r/2} \ge c_3(1+\theta_1)^{r/2}/2 \ge 1+\theta_1$, for r sufficiently large. We have $1 > p_0 = 1-p+pp_0^D$. If $p_0 = 1-p+\varepsilon$ then $\varepsilon = pp_0^D$. As we increase r, D increases and the probability of extinction decreases. (We can see the latter by a simple coupling argument). We can therefore make r sufficiently large so that $p_0^D \le 1/2$. In which case $\rho_q^{(n_l)} \le p_0 \le 1-c_3/4$, contradicting $\lim_{l\to\infty} \rho_q^{(n_l)} = 1$.

So we can assume that $\rho_i^{(n)} \le 1 - c_4$ for some $c_4 \ll 1$ and for any $i \in I$. Now for each $x \in [n]$ let $z_x = 1$ if $|C_x| \ge \nu_0$. Let $z_x = 0$ otherwise. The analysis above shows that if $x \in V_i$, $i \in I$ then $\mathbf{Pr}(z_x = 1) \ge \theta_1 - O(\nu_0/n)$. So if $Z = \sum_{x\in V} z_x$, then $\mathbf{E}(Z) \ge \theta_1 n/2$. With a view to applying the Chebyshev inequality, we now estimate $\mathbf{Pr}(z_x = z_y = 1)$ for

$x \in V_i, y \in V_j,\ i,j \in I$.

$$\begin{aligned}\mathbf{Pr}(z_x = z_y = 1) &\leq \mathbf{Pr}(z_x = 1 \mid z_y = 1,\ x \in C_y)\mathbf{Pr}(x \in C_y) + \\ &\qquad\qquad \mathbf{Pr}(z_x = 1 \mid z_y = 1,\ x \notin C_y)\mathbf{Pr}(z_y = 1) \\ &= O(\nu_0/n) + \mathbf{Pr}(z_y = 1)(\mathbf{Pr}(z_x = 1) + O(\nu_0^2/n)). \end{aligned} \tag{9}$$

To verify (9) we observe that

$$\mathbf{Pr}(x \in C_y) \leq \mathbf{Pr}(\Delta(G_m) \geq \ln n) + 2(\nu_0 + \ln n)p_{\max} = O(\nu_0/n).$$

The term $2(\nu_0 + \ln n)p_{\max}$ bounds $|C_y|$ times the probability that an edge $\{v, x\}$, $v \in C_y$ is discovered in the BFS construction of C_y. Then to estimate $\mathbf{Pr}(z_x = 1 \mid z_y = 1,\ x \notin C_y)$ we use $O(\nu_0/n)$ to bound the probability that we invoke the demon when constructing C_y. Then we maximally couple the conditioned and unconditioned Red/Blue branching processes. Then there are ν_0 opportunities for these to deviate, each having a probabiliy of $O(\nu_0/n)$.

Therefore

$$\mathbf{E}(Z^2) \leq O(\nu_0^2 n) + \mathbf{E}(Z)^2$$

and so

$$\mathbf{Pr}(Z \leq \mathbf{E}(Z)/2) \leq \frac{O(\nu_0^2 n)}{\mathbf{E}(Z)^2} = o(1).$$

Now let $\zeta_x = 1$ if the above BFS procedure when started with x *and allowed to continue* results in exposing a component of size at least $n^{2/3}$. We show that **whp** there exists a set X of size $O((\ln n)^3)$ such that

$$\zeta_x = z_x \textit{ for all } x \notin X. \tag{10}$$

We will show then that **whp**, $x, y \notin X$ and

$$\zeta_x = \zeta_y = 1 \textit{ implies that } x \textit{ and } y \textit{ lie in the same component.} \tag{11}$$

It follows from (10) and (11) that **whp** there is a component with $Z \geq \theta_1 n/2 - o(n)$ vertices and that every other component is of size $O((\ln n)^3)$, completing the proof of the lemma.

The exceptional vertices X are those for which either (i) we have to invoke the demon before exposing $n^{2/3}$ vertices during the BFS construction or (ii) C_x contains a cycle. We have $\mathbf{E}(|X|) \leq 16\nu_0 + O(\nu_0^2/n)$, by (5), and so $|X| = O((\ln n)^3)$ **whp**.

We show next that if C_x^α denotes the vertices of color $\alpha = \textit{Red}, \textit{Blue}$ in C_x then for some $c_5 \ll 1$,

$$\mathbf{Pr}(x \notin X \textit{ and } z_x = 1 \textit{ and } |C_x^{Blue}| \geq c_5\nu_0) \leq n^{-2}. \tag{12}$$

For each $y \in C_x^{Red}$ let B_y be the vertices $z \in C_x^{Blue}$ that are direct descendants of y i.e. the path from y to z is Blue, except for y. It follows from (8) that $|B_y|$ is dominated by the number W of *proper* descendants in a single type branching process where the number of progeny is $Po(\mu)$, $\mu = 5\theta_0 c_1$. Let $\nu_0' = \nu_0 + \ln n$ and let $Z = W_1 + W_2 + \cdots + W_{\nu_0'}$ where the W_i are independent copies of W. The moment generating function $\mathbf{E}(e^{uPo(\mu)}) = \exp\{\mu(e^u - 1)\}$. So for any $t \geq 1$ we have, after putting $u = \ln(t/\mu)$,

$$\mathbf{Pr}(Po(\mu) \geq t) \leq e^{-tu}\mathbf{E}(e^{uPo(\mu)}) = e^{-\mu}\left(\frac{\mu e}{t}\right)^t \leq (3\mu)^t. \tag{13}$$

Now consider labelling the progeny of a branching process by sequences $i_1, i_2, \ldots, i_r$. This sequence is i_rth child of the i_{r-1}th child of the ... of the i_1th child of the root. The probability that such a particle exists is by (13), at most $(3\mu)^{i_1+\cdots+i_r}$. It follows that

$$\mathbf{Pr}(W \geq t) \leq \sum_{r=1}^{\infty}\sum_{s=t}^{\infty}\sum_{\substack{i_1+\cdots+i_r=s\\ i_1,\ldots,i_r\geq 1}}(3\mu)^s = \sum_{r=1}^{\infty}\sum_{s=t}^{\infty}\binom{s-1}{r-1}(3\mu)^s = \sum_{s=t}^{\infty}2^{s-1}(3\mu)^s \leq (6\mu)^t.$$

So,

$$\mathbf{Pr}(Z \geq t) \leq \sum_{s=t}^{\infty}\sum_{\substack{i_1+\cdots+i_{\nu_0'}=s\\ i_1,\ldots,i_{\nu_0'}\geq 0}}(6\mu)^s = \sum_{s=t}^{\infty}\binom{s+\nu_0'-1}{s-1}(6\mu)^s \leq 2^{\nu_0'}(12\mu)^t.$$

We can choose $c_1 \ll c_5 \ll 1$ so that putting $t = c_5\nu_0'$ into the above proves (12).

We now consider the Red branching process at a point where s generations have been produced and the total progeny is at least $\nu_1 = \nu_0 - c_5\nu_0$. We will now show that **whp** after $O(1)$ more rounds, a sufficient number of Red progeny are at the bottom level. Let $\mathbf{t}_{(0)} = [t_1, t_2, \ldots, t_k]^T$ where t_i is the number of progeny of Type i that are in the first $s-1$ levels. Let b_i be the number of progeny of Type i at the bottom level and let $N_{i,j}$ be the number of progeny of Type j that are children of progeny of Type i. Then we have

$$b_j = \sum_{i=1}^{k} N_{i,j} - t_j. \tag{14}$$

We will use the following concentration inequalities (see Appendix):

$$\mathbf{Pr}(Po(\lambda) \geq \lambda + u) \leq \begin{cases} e^{-u/3} & u \geq \lambda \\ e^{-u^2/(3\lambda)} & u \leq \lambda \end{cases} \tag{15}$$

$$\mathbf{Pr}(Po(\lambda) \leq \lambda - u) \leq e^{-u^2/(2\lambda)} \tag{16}$$

Now b_j is the sum of t_i variables each distributed as $Po(M_{i,j})$. We can imagine generating ν_0' such random variables $X_1, X_2, \ldots, X_{\nu_0'}$. Then (15) and (16) imply that for some c_6, c_7,

$$\mathbf{Pr}(\exists t \leq \nu_0' :\ |X_1 + \cdots + X_t - tM_{i,j}| \geq c_6\nu_1) \leq \nu_0' e^{-c_7\nu_0'}. \tag{17}$$

To see this, first note that $X_1 + \cdots + X_t$ is distributed as $Po(\lambda)$, $\lambda = tM_{i,j}$. We use (15) and (16) directly with and $u = c_6\nu_1$ to get the bounds.

It follows from (14) and (17) that

$$\mathbf{Pr}\left(\left|b_j - \sum_{i=1}^{k} t_i M_{i,j} + t_j\right| \geq c_6\nu_1\right) \leq e^{-c_7\nu_1/2}. \tag{18}$$

If we continue the process for $\ell = O(1)$ more rounds and let $\mathbf{t}_{(\ell)}$ take the place of $\mathbf{t}_{(0)}$ then we can write

$$\mathbf{t}_{(l)}^T = \mathbf{t}_{(l-1)}^T M \pm c_8 \nu_1 \mathbf{e}^T$$

where $\mathbf{e} = [1, 1, \ldots, 1]^T$ and c_8 can be made arbitrarily small.

Iterating we get that **qs**,

$$\mathbf{t}_{(\ell)}^T = \mathbf{t}_{(0)}^T M^\ell \pm c_8 \nu_1 \mathbf{e}^T (M^{\ell-1} + \cdots + I).$$

Now we can write

$$M^\ell = \rho^\ell \mathbf{u}\mathbf{v}^T + O((\rho\zeta)^\ell J)$$

where, using the Frobenius Theorem for positive matrices, $\mathbf{u}, \mathbf{v} > 0$ are the left and right (column) eigenvectors of norm 1 corresponding to ρ, $0 \le \zeta < 1$ is a constant and J is the all 1's matrix. Note that the coordinates of $\mathbf{u}, \mathbf{v} > 0$ are bounded below by some positive constant that depends only on θ_0.

Making ℓ large we have **qs**,

$$\mathbf{t}_{(\ell)}^T = (1 \pm c_9)\rho^\ell (\mathbf{t}_{(0)}^T \mathbf{u})\mathbf{v}^T.$$

We now return our attention to the actual BFS construction. Let $U_j^{(0)}$ denote the set of unsaturated vertices of Type j at the time of construction of the $(s+\ell)$th Red generation. Then, given $z_x = 1$, we have that **qs**

$$|U_j^{(0)}| \ge c_{10} v_j (\ln n)^2 \tag{19}$$

where $c_{10} = ((1 - c_9)\rho^\ell - (1 + c_9)\rho^{\ell-1})(\mathbf{t}_{(0)}^T \mathbf{u}) - c_7 > 0$.

Now for each j, choose a subset of $U_j^{(0)}$ of size exactly $c_{10} v_j (\ln n)^2$. Now let $U^{(0)} = \bigcup_j U_j^{(0)}$ and then define the sequence $U^{(0)}, U^{(1)}, \ldots$, where $U^{(t+1)}$ is a subset of the neighbors of $U^{(t)}$ in $V \setminus C_x$. Then let $U_j^{(t)} = V_j \cap U^{(t)}$ for $j \in [k]$.

Assume that

$$|U_j^{(t)}| = c_{10}(\rho - c_{11})^t v_j (\ln n)^2$$

for a small c_{11} and for all j.

Now $|U_j^{(t+1)}|$ is the sum of $\sum_i |U_i^{(t)}|(|V_j| - O(n^{2/3}))$ Bernouilli random variables and

$$\mathbf{E}(|U_j^{(t+1)}|) \ge \sum_i |U_i^{(t)}| M_{i,j} = c_{10}(\ln n)^2 (\rho - c_{11})^t \sum_i v_i M_{i,j} = c_{10}(\rho - c_{11})^t \rho (\ln n)^2 v_j.$$

Furthermore, it dominates a sum of independent variables where for each i, j we have probability of being 1 is equal to $M_{i,j}$. So, applying Hoeffding's inequality we see that **qs**

$$|U_j^{(t+1)}| \ge c_{10}(\rho - c_{11})^{t+1} (\ln n)^2 v_j.$$

Then taking subsets we can assume that

$$|U_j^{(t+1)}| = c_{10}(\rho - c_{11})^{t+1} (\ln n)^2 v_j.$$

It follows, after iterating $O(\ln n)$ times that we reach τ such that $|U_j^{(\tau)}| = \Theta(n^{2/3})$ for each j. This verifies (10).

Now suppose that $\zeta_x = \zeta_y = 1$ and $C_x \cap C_y = \emptyset$. We use the notation $U_j^{(x)}$ to denote the unsaturated vertices of C_x of Type j at the time its size reaches $\Theta(n^{2/3})$ and a similar notation for y. We let $U^{(x)} = \bigcup_i U_i^{(x)}$.

The number of edges between $U^{(x)}$ and $U^{(y)}$ is the sum, over unexposed rounds, of Bernouilli random variables and the expected number is at least (see (6) and replace ν_0 by $O(n^{2/3})$)

$$\sum_{i,j \in I} |U_i^{(x)}|\,|U_j^{(y)}| \left(\frac{i+j}{k} p_{\max} - \frac{O(n^{2/3})}{n^2} \right) \geq c_{12} n^{1/3}.$$

Furthermore, it dominates such a sum of independent variables and we see that **qs** there is at least one edge joining $U^{(x)}$ and $U^{(y)}$. This proves (11) and completes the proof of the lemma. ▮

2.2 A Sufficient Criterion for the Non-Existence of Giant

Let $p_x, p_{xy}, p_{\max}, c_1, c_2, k, V_i$ be as above. We re-define M to be the $k \times k$-matrix with entries

$$M_{i,j} = \frac{((i+1)+(j+1))p_{\max}}{k} |V_j| + c_1^2, \quad 0 \leq i, j \leq k-1.$$

Lemma 5 *Suppose that there are $c_2 > 0$ and n_0 such that for all $n \geq n_0$ we have $\lambda_1^{(n)} < 1 - c_2$. Then* **whp** *each component of G_m has order $O(\ln n)$.*

Proof. As before, one can argue that the breadth first search *is dominated* by the Poisson multi-type branching process given by M as long as we have exposed at most $o(n)$ edges. Lemma 6 below shows that for some C, the probability that the branching process $\mathcal{B}_i$, any $0 \leq i \leq k-1$ reaches at least $C \ln n$ vertices is at most n^{-2} say. Hence, the expected number of vertices of G_m in components of size at least $C \ln n$ is $o(1)$ and by Markov's inequality **whp** there is no such component. ▮

Lemma 6 *There is a positive constant $\delta = \delta(\varepsilon, t)$ such that the following holds. Let $\mathcal{B}$ be the Poisson branching process with mean $t \times t$-matrix M that starts with one particle (of any type). If the largest eigenvalue of M is $\lambda_1 < 1 - \varepsilon$ then, for every s, the probability that $\mathcal{B}$ reaches at least s vertices is at most $(1+\delta)^{1-s}$.*

Proof. Let $\varepsilon_0 > 0$ be sufficiently small to satisfy

$$1 - (1-\varepsilon)(1+\varepsilon_0) > 2t^2 \varepsilon_0.$$

It is enough to prove the claim for all sufficiently large s, $s \geq s_0$, where $s_0 = s_0(\varepsilon, t, \varepsilon_0)$. Let us run the process $\mathcal{B}$ level by level until the process dies out or we reach at least s

particles in total, after some level has been added. We do not expose the whole process but only the following information: the vector $\mathbf{n} = (n_1, \ldots, n_t)^T$, where n_i is the total number of particles of Type i generated.

Suppose that we have reached $s \geq s_0$ particles, $s = n_1 + \cdots + n_k$. We claim that for some i, j we have

$$n_{i,j} \geq (1+\varepsilon_0) n_i M_{i,j} + \varepsilon_0 s, \tag{20}$$

where $n_{i,j}$ is the total output of Type j particles that are produced (looking forward and including the next level) by the n_i particles of Type i that were born by now. Note that we do not expose $n_{i,j}$'s but only state that whatever feasible values these variables have, they must satisfy (20) for some i, j. Suppose on the contrary that (20) is false for all i, j. Let $\mathbf{1} = (1, \ldots, 1)^T$ be all-1s column vector. For every $j \in [t]$ we have $n_j \leq \sum_{i=1}^t n_{i,j} + 1$, where the last term accounts for the initial vertex which is not born but given. This implies that we have the following coordinawise domination:

$$\mathbf{n} \leq (1+\varepsilon_0) M^T \mathbf{n} + \varepsilon_0 t s \mathbf{1} + \mathbf{1}.$$

Taking the l_2-norm, using the triangle inequality and the fact that $\|M^T \mathbf{n}\| \leq (1-\varepsilon)\|\mathbf{n}\|$, we conclude

$$\|\mathbf{n}\| - (1-\varepsilon)(1+\varepsilon_0)\|\mathbf{n}\| < \varepsilon_0 t^{3/2} s' + \sqrt{t}.$$

But $\sum_{i=1}^t n_i = s$ implies that $\|\mathbf{n}\| \geq s/\sqrt{t}$, which contradicts the choice of ε_0.

We can generate the branching process by first generating t^2 infinite sequences

$$(X_{i,j,1}, X_{i,j,2}, \ldots), \quad i, j \in [t],$$

of independent Poisson outcomes with means $\mathbf{E}(X_{i,j,l}) = M_{i,j}$. Now, whenever we have to determine the offspring of Type j of a particle of Type i, we take the first unused $X_{i,j,l}$. Under this coupling the probability that (20) holds is at most the probability that there exist $n_1, \ldots, n_t$ with $s = n_1 + \cdots + n_t \geq s_0$ and $i, j \in [t]$ such that

$$X_{i,j,1} + \cdots + X_{i,j,n_i} \geq (1+\varepsilon_0) n_i M_{i,j} + \varepsilon_0 s. \tag{21}$$

Note that the sum $X_{i,j,1} + \cdots + X_{i,j,n_i}$ is distributed as the Poisson variable $Po(\mu_{i,j})$ with mean $\mu_{i,j} = n_i M_{i,j}$. Thus, by the union bound, the probability of (20) is at most

$$\sum_{s \geq s_0} \sum_{i,j} s^t \mathbf{Pr}(Po(\mu_{i,j}) \geq (1+\varepsilon_0)\mu_{i,j} + \varepsilon_0 s) \leq \sum_{s \geq s_0} \sum_{i,j} s^t e^{-Ls}$$

for some $L = L(\varepsilon_0)$.

To define L we observe that if $\mu_{i,j} \leq \varepsilon_0 s/2$ then we can take $L \geq \varepsilon_0/3$ whereas if $\mu_{i,j} > \varepsilon_0 s/2$ then we can take $L \geq \varepsilon_0^3/6$, after using the bounds in (15). ∎

3 Creating a Giant

Here we prove Theorems 1 and 2.

3.1 Proof of Theorem 1

Let us first investigate the following strategy for Paul. We have m rounds and Paul uses the same acceptance set A in each round, that is, $A_1 = \cdots = A_m = A$. Let $|A| = a$. Assume that the limits

$$\alpha = \lim_{n\to\infty} a/n, \qquad \mu = \lim_{n\to\infty} m/n$$

exist and $0 < \alpha < 1$. Let $B = [n] \setminus A$ be the complement of A. The function p_x assumes two possible values:

$$p_x = \begin{cases} \frac{2mn-ma}{n^2(n-1)}, & x \in A, \\ \frac{mn-ma}{n^2(n-1)}, & x \in B. \end{cases}$$

When we construct the upper bound matrix M as in Section 2.1, then when c_1, c_2 approach 0, we eventually obtain 2×2 matrix which for large n is approximately

$$\mu \times \left(\begin{array}{cc} (4-2\alpha)\alpha & (3-2\alpha)(1-\alpha) \\ (3-2\alpha)\alpha & (2-2\alpha)(1-\alpha) \end{array} \right). \tag{22}$$

The largest eigenvalue of this limiting matrix is

$$\lambda_1 = \mu(1+\sqrt{1+\alpha-\alpha^2}) \tag{23}$$

If $\alpha = 1/2$ (the case of Theorem 1), then taking μ to be strictly larger than $2\sqrt{5}-4$, we ensure that λ_1 is strictly larger than 1, say $\lambda_1 \geq 1 + 2c_2$ for some positive constant c_2. Since the limit of $\lambda_1^{(n)}$ is λ_1, we have $\lambda_1^{(n)} \geq 1 + c_2$ for all sufficiently large n. Now Theorem 1 follows from Lemma 4.

3.2 Proof of Theorem 2

Let us turn to Theorem 2. Suppose that $m \leq 2\sqrt{5} - 4 - \varepsilon$ for some constant $\varepsilon > 0$. We define the $k \times k$-matrix M as in Section 2.2. In order to prove Theorem 2 it is enough, by Lemma 5 to show that there is $\delta > 0$, depending only on ε, such that for all sufficiently large k (= small c_1) and all large n, we have $\lambda_1 < 1 - \delta$.

Our matrix M has the form:

$$M = (\mathbf{1}\mathbf{p}^T + \mathbf{p}\mathbf{1}^T)D + \widehat{M}.$$

Here $\mathbf{p}$ is the column vector $(p_{\max}/k, 2p_{\max}/k, \ldots, p_{\max})$ and D is the diagonal matrix with $\nu_0, \ldots, \nu_{k-1}$ on diagonal, where we denote $\nu_i = |V_i|$. Furthermore, $\max_{i,j} |\widehat{M}(i,j)| \leq c_1^2$ and $\mathbf{1}^T D \mathbf{1} = \sum_{i=0}^{k-1} \nu_1 = n$.

Now for $|\mathbf{x}| = 1$, we have $|\mathbf{x}^T \widehat{M} \mathbf{x}| \leq kc_1^2 \leq 2c_1$ and so it suffices to prove that the largest eigenvalue of $M - \widehat{M}$ is bounded away from 1. With this in mind, M is really $M - \widehat{M}$ in the analysis below.

Suppose that $\mathbf{v}$ is a column eigenvector corresponding to eigenvalue λ_1 of M. Then

$$\lambda_1 \mathbf{v} = M\mathbf{v} = \mathbf{1}(\mathbf{p}^T D\mathbf{v}) + \mathbf{p}(\mathbf{1}^T D\mathbf{v}),$$

belongs to the span of $\mathbf{p}$ and $\mathbf{1}$. Note that $\mathbf{v}$ is not a multiple of $\mathbf{1}$ (otherwise $\mathbf{p}$ would also be such a multiple, and it clearly isn't).

Write

$$\mathbf{v} = \mu\mathbf{p} + \nu\mathbf{1}, \tag{24}$$

with $\mu, \nu \in \mathbb{R}$. Note that $\mu \neq 0$. Substituting (24) into $\lambda_1 \mathbf{v} = M\mathbf{v}$ and equating the coefficients, when expanded in basis $\{\mathbf{p}, \mathbf{1}\}$, we obtain

$$\begin{cases} \lambda_1 \mu &= \mu \mathbf{1}^T D\mathbf{p} + \nu \mathbf{1}^T D\mathbf{1}, \\ \lambda_1 \nu &= \mu \mathbf{p}^T D\mathbf{p} + \nu \mathbf{p}^T D\mathbf{1}. \end{cases}$$

From the first equation, we get $\nu = \mu(\lambda_1 - \mathbf{1}^T D\mathbf{p})/1^T D\mathbf{1}$. Substituting this into the second equation and cancelling μ, we obtain

$$\lambda_1 \frac{\lambda_1 - \mathbf{1}^T D\mathbf{p}}{1^T D\mathbf{1}} = \mathbf{p}^T D\mathbf{p} + \frac{\lambda_1 - \mathbf{1}^T D\mathbf{p}}{1^T D\mathbf{1}} (\mathbf{p}^T D\mathbf{1}). \tag{25}$$

Since D is a diagonal matrix, we have $\mathbf{1}^T D\mathbf{p} = \mathbf{p}^T D\mathbf{1}$. Substituting $1^T D\mathbf{1} = n$ we see that the larger root of the quadratic equation (25) is

$$\lambda_1 = \mathbf{1}^T D\mathbf{p} + \sqrt{n\mathbf{p}^T D\mathbf{p}}. \tag{26}$$

Now,

$$\mathbf{1}^T D\mathbf{p} = \sum_{i=0}^{k} \nu_i \frac{(i+1)p_{\max}}{k} \in \left[\sum_{x\in[n]} p_x, \sum_{x\in[n]} \left(p_x + \frac{p_{\max}}{k} \right) \right].$$

But

$$\sum_{x\in[n]} p_x = \sum_{x\in[n]} \frac{m - \gamma + \gamma_x}{n(n-1)} = \frac{m}{n-1} - \frac{\gamma}{n-1} + \frac{\sum_x \gamma_x}{n(n-1)} = \frac{m}{n-1}$$

and so

$$\mathbf{1}^T D\mathbf{p} \leq \frac{m}{n-1} + \frac{np_{\max}}{k} \leq \frac{m}{n-1} + \frac{1}{k}. \tag{27}$$

Suppose we fix γ and we want to maximize the right-hand side of (26). Then up to $O(1/k)$ we have to maximize

$$\mathbf{p}^T D\mathbf{p} = \sum_{i=0}^{k} \nu_i \left(\frac{(i+1)p_{\max}}{k} \right)^2 \in \left[\sum_{x\in[n]} p_x^2, \sum_{x\in[n]} \left(p_x^2 + \frac{2p_{\max}}{k} \sum_{x\in[n]} p_x + \frac{np_{\max}^2}{k^2} \right) \right].$$

But,

$$\begin{aligned} \sum_{x\in[n]} p_x^2 &= \sum_{x\in[n]} \left(\frac{m - \gamma + \gamma_x}{n(n-1)} \right)^2 \\ &= \frac{m^2 - \gamma^2}{n(n-1)^2} + \sum_{x\in[n]} \frac{\gamma_x^2}{n^2(n-1)^2}. \end{aligned}$$

and so

$$\begin{aligned} n\mathbf{p}^T D\mathbf{p} &\leq \frac{m^2-\gamma^2}{(n-1)^2} + \sum_{x\in[n]} \frac{\gamma_x^2}{n(n-1)^2} + \frac{2n^2p_{\max}^2}{k} + \frac{n^2p_{\max}^2}{k^2} \\ &\leq \frac{m^2-\gamma^2}{(n-1)^2} + \sum_{x\in[n]} \frac{\gamma_x^2}{n(n-1)^2} + \frac{2}{k}. \end{aligned}$$

Now, if we want to maximize this expression over reals, given $\gamma = \frac{1}{n}\sum_x \gamma_x$ and the constraints $0 \leq \gamma_x \leq m$, then we should take each γ_x either 0 or m (except at most one value) to have equality in $\gamma = \frac{1}{n}\sum_x \gamma_x$. So $\sum_{x\in[n]} \gamma_x^2 \leq mn\gamma + m^2$ and

$$n\mathbf{p}^T D\mathbf{p} \leq \frac{(m^2-\gamma^2)n + mn\gamma}{n(n-1)^2} + \frac{3}{k}. \tag{28}$$

Going back to (26), we see from (27) and (28) that

$$\lambda_1 \leq \frac{m}{n} + \left(\frac{(m^2-\gamma^2)n + mn\gamma}{n(n-1)^2}\right)^{1/2} + \frac{4}{k}.$$

This expression is maximised when $\gamma = m/2$ and so if $m = cn$ then

$$\lambda_1 \leq c\left(1 + \frac{\sqrt{5}}{2}\right) + \frac{4}{k} + o(1) \leq 1 - \frac{2\sqrt{5}+4}{4}\varepsilon + \frac{5}{k}$$

which is bounded away from 1 from sufficiently large k. This completes the proof of Theorem 2.

4 Delaying a Giant

As we have already observed in the introduction, by always selecting the first edge Paul can ensure that a giant component appears **whp** after about $n/2$ rounds. Here we show that, essentially, he cannot delay the birth of giant any longer, that is, prove Theorem 3.

Let $\varepsilon > 0$ be given. Choose sufficiently small $c_1 \gg c_2 \gg c_3 > 0$. Let M be the matrix as defined in Section 2.1. As we have already observed, we have $|V| \geq (1-3/k)n$. Let $l = \lfloor 1/c_2 \rfloor \gg k$. We are going to define a new square matrix L which has approximately l rows and also describes a multi-type branching process. Replace row i by $\lfloor |V_i|/(n/l) \rfloor$ new rows and each column j by $\lfloor |V_j|/(n/l) \rfloor$ new columns. The entry of L in the intersection of a row and a column that are derived from row i and column j is

$$\frac{(i+j)p_{\max}}{k} \times \frac{n}{l} - c_3.$$

The square matrix L has ρ rows where $l(1-4/k) \leq \rho \leq l$. To understand its construction, here is the interpretation in terms of graphs. To this end, we partition each V_i, into as many disjoint sets U_j as possible with each U_i having precisely $\lfloor n/l \rfloor$ elements. All the remaining vertices (at most $k \times (n/l) \ll n/k$) we "discard". The entry $L_{i,j}$ is approximately the expected number of offspring in the set U_j of any vertex in U_i. Thus L also well describes a Red branching process as defined in Lemma 4. In order to finish the proof it is enough to argue that the largest eigenvalue of L is at least $1+\varepsilon$.

But the number $\mathbf{1}^T L \mathbf{1}$ is approximately the sum over U_i of the expected number of edges from a vertex of U_i into $\cup_j U_j$. More precisely, $(n/l) \times \mathbf{1}^T L \mathbf{1} \geq (2 - O(1/k))e(G_m)$, so

$$\mathbf{1}^T L \mathbf{1} > (1 + 3\varepsilon/2)l.$$

Also, $\mathbf{1}^T \mathbf{1} = l(1 - O(1/k))$, so

$$\frac{\mathbf{1}^T L \mathbf{1}}{\mathbf{1}^T \mathbf{1}} \geq 1 + \varepsilon,$$

which implies that the largest eigenvalue of L is at least $1+\varepsilon$, as required.

References

[1] N. Alon. A note on network reliability. In *Discrete probability and algorithms (Minneapolis, MN, 1993)*, volume 72 of *IMA Vol. Math. Appl.*, pages 11–14. Springer Verlag, 1995.

[2] K. B. Athreya and P. E. Ney. *Branching Processes*. Springer Verlag, 1972.

[3] A. Beveridge, T. Bohman, A. Frieze, and O. Pikhurko. Product rule wins a competitive game. To appear in Trans. Amer. Math. Soc., 2006.

[4] T. Bohman and A. Frieze. Avoiding a giant component. *Random Struct. Algorithms*, 19:75–85, 2001.

[5] T. Bohman, A. Frieze, and N. Wormald. Avoidance of a giant component in half the edge set of a random graph. *Random Struct. Algorithms*, 25:432–449, 2004.

[6] T. Bohman and J. H. Kim. A phase transition for avoiding a giant component. To appear in Random Struct. Algorithms, 2006.

[7] T. Bohman and D. Kravitz. Creating a giant component. To appear in Combin. Prob. Computing, 2006.

[8] B. Bollobás, S. Janson, and O. Riordan. The phase transition in inhomogeneous random graphs. Submitted, 2006.

[9] R. Durrett, Probability: Theory and Examples, 2nd Edition, Brooks-Cole, Belmont, CA, 2004.

[10] W. Feller. *An Introduction to Probability Theory and Its Applications*. Wiley-Intersci. Publ., 3d edition, 1968.

[11] A. Flaxman, D. Gamarnik, and G. Sorkin. Embracing the giant component. In M. Farach-Coltin, editor, *Latin 2004: Theoretical Informatics*, pages 69–79. Springer Verlag, 2004.

[12] S. Janson, T. Łuczak, and A. Ruciński. *Random Graphs*. Wiley-Intersci. Publ., 2000.

[13] B. Söderberg. General formalism for inhomogeneous random graphs. *Phys. Rev. E*, 66:066121, 2002.

[14] J. Spencer and N. Wormald. Birth control for giants. To appear in Combinatorica, 2006.

A Proof of (15) and (16)

Let $e^\theta = 1 + u/\lambda$.

$$\begin{aligned} \mathbf{Pr}(Po(\lambda) \geq \lambda + u) &\leq e^{-\theta(\lambda+u)}\mathbf{E}(e^{\theta Po(\lambda)}) \\ &= \exp\{\lambda(e^\theta - 1) - \theta(\lambda + u)\} \\ &= \exp\{u(1 - (1 + \lambda/u)\ln(1 + u/\lambda)\} \end{aligned} \tag{29}$$

If $u \leq \lambda$ then we use

$$(1 + 1/x)\ln(1 + x) \geq 1 + x/3$$

for $0 \leq x \leq 1$.

Plugging $x = u/\lambda$ into (29) gives the required inequality (15) in the case $u \leq \lambda$. If $\lambda \leq u$ then we use

$$(1 + 1/x)\ln(1 + x) \geq 2\ln 2 = 1.386924...$$

for $x \geq 1$.

For (16) we have, with $e^{-\theta} = 1 - u/\lambda$,

$$\begin{aligned} \mathbf{Pr}(Po(\lambda) \leq \lambda - u) &\leq e^{\theta(\lambda-u)}\mathbf{E}(e^{-\theta Po(\lambda)}) \\ &= \exp\{\lambda(e^{-\theta} - 1) + \theta(\lambda - u)\} \\ &= \exp\{u((1 - \lambda/u)\ln(1 - u/\lambda) - 1)\} \\ &\leq e^{-u^2/(2\lambda)}. \end{aligned}$$