SIAM J. DISCRETE MATH.
Vol. 23, No. 2, pp. 993–1008

MEMORYLESS RULES FOR ACHLIOPTAS PROCESSES*

ANDREW BEVERIDGE†, TOM BOHMAN‡, ALAN FRIEZE‡, AND OLEG PIKHURKO‡

Abstract. In an Achlioptas process two random pairs of $\{1,\dots,n\}$ arrive in each round and the player has to choose one of them. We study the very restrictive version where a 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 $(2\sqrt{5}-4+o(1))n = (0.4721\dots+o(1))n$ rounds and that this is the best possible. On the other hand, if the player wants to delay the appearance of a giant, then the optimal bound is $(1/2+o(1))n$, the same as in the Erdős–Rényi model.

Key words. Achlioptas process, memoryless

AMS subject classification. 05C80

DOI. 10.1137/070684148

1. Introduction. Let $[n] = \{1,\dots,n\}$ 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, 10, 13]. 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.

Unfortunately, it is still an open question as to the earliest/latest time that Paul can ensure the birth of a giant. The answer is unknown even for some of Paul's simplest strategies, such as, for example, the product rule; see [3].

Here, we study a class of more restrictive strategies, which we call *memoryless rules*. Namely, Paul's choice e_r at round r is a function of the current two random edges, $D_r = \{x_r, y_r\}$ and $D'_r = \{x'_r, y'_r\}$, and the number r. It does not depend on what happened in rounds 1 to $r-1$; that is, Paul does not remember any previous history. Such restrictions may appear when our online algorithm has to make fast decisions 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_r from the first random edge. Here we assume that the first random edge $\{x_r, y_r\}$ is generated by taking a random $x_r \in [n]$ and then a random $y_r \in [n]\setminus\{x_r\}$. (An alternative nonequivalent setting, namely to reveal $\min\{x_r, y_r\}$, is not studied here.) Any memoryless algorithm of this type is described by the

*Received by the editors March 1, 2007; accepted for publication (in revised form) March 4, 2009; published electronically May 20, 2009.
http://www.siam.org/journals/sidma/23-2/68414.html

†Department of Mathematics and Computer Science, Macalester College, 1600 Grand Ave., St. Paul, MN 55105 (abeveridge@macalester.edu).

‡Department of Mathematical Sciences, Carnegie Mellon University, Pittsburgh, PA 15213 (tbohman@moser.math.cmu.edu, alan@random.math.cmu.edu, pikhurko@andrew.cmu.edu). The second author's research was supported in part by NSF grant DMS-0401147. The third author's research was supported in part by NSF grant CCR-0502793. The fourth author's research was supported in part by NSF grant DMS-0457512.

sequence of the *acceptance sets* $A_1, \dots, A_m \subseteq [n]$, where Paul selects D_r if and only if the revealed vertex x_r belongs to A_r. 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*

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

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

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

THEOREM 2. *Let c, ε be as in* (1). *For any sequence $A_1, \dots, A_m$, 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 the best possible.

THEOREM 3. *For every constant $\varepsilon > 0$ and any $A_1, \dots, A_m \subseteq [n]$, $m = \lfloor (1/2+\varepsilon)n \rfloor$, 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 $A_1, \dots, A_m \subseteq [n]$ be given and let G_m be the graph constructed by the corresponding memoryless rule.

We will examine the state of the process after m_0, m_1 as well as m rounds. Here $m_0 = \lfloor m - 2\varepsilon n/3 \rfloor$ and $m_1 = \lfloor m - \varepsilon n/3 \rfloor$ for some small positive ε. We will show that whp there are $\Omega(n)$ vertices in components of G_{m_0} of size at least $n^{7/8}$. This is the main part of the proof. The edges $e_{m_0+1}, \dots, e_{m_1}$ are used to combine these large components into a single giant. The edges $e_{m_1+1}, \dots, e_m$ are used to absorb the components of size $\Omega(\ln n)$ into the giant.

Consider the first m_0 rounds. The expected number of rounds in which any given (unordered) pair $\{x, y\}$ is accepted is

$$p_{xy} = \frac{1}{\binom{n}{2}} \sum_{r=1}^{m_0} \left(\frac{\mathbf{1}_{x \in A_r} + \mathbf{1}_{y \in A_r}}{2} + \frac{n - |A_r|}{n} \right), \tag{2}$$

where the indicator function $\mathbf{1}_{x \in A_r}$ is 1 if $x \in A_r$ and 0 otherwise. For $z \in [n]$, let $\gamma_z = |\{i \in [m_0] : z \in A_r\}|$. We can write $p_{xy} = p_x + p_y$, where

$$p_x = \frac{m_0 - \gamma}{n(n-1)} + \frac{1}{n(n-1)} \sum_{r=1}^{m_0} \mathbf{1}_{x \in A_r} = \frac{m_0 - \gamma + \gamma_x}{n(n-1)}, \tag{3}$$

and $\gamma = \frac{1}{n} \sum_{r=1}^{m_0} |A_r| = \frac{1}{n} \sum_{z=1}^{n} \gamma_z$.

Let

$$p_{\max} = \frac{2m_0}{n(n-1)}.$$

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

The probability that $\{x, y\}$ is chosen at least l times is $O(n^{-l})$ (assuming $m = O(n)$). It follows that the probability that $\{x, y\} \in E(G_m)$ is $p_{xy} + O(n^{-2})$. 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 of the obtained graph G_m. Such *inhomogeneous* random graph models were studied by Alon [1], Söderberg [12], Bollobás, Janson, and Riordan [8], and others. However, given some dependence 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 Ruciński's book [11, 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 (BFS) from x by an appropriate ideal branching process. A new technical difficulty that we face in comparison to [11, section 5.2] is that the vertices are not homogeneous, so we have to use multitype 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, in any inequalities we 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. For all definitions related to branching processes, we refer the reader to Athreya and Ney [2], in particular to Chapter V, which deals with multitype processes.

In the ideal multitype branching processes that 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 bounded away from zero. Therefore, in order to prove a sufficient condition for the existence of a giant component, we proceed as follows.

Let m and $A_1, \dots, A_m \subseteq [n]$ be given, all dependent 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, $n/4 \le m \le n$ because outside this range the existence of a giant is whp predetermined irrespective of what Paul does; see Bohman and Kim [6] and Bohman and Kravitz [7].

2.1. A sufficient condition for the existence of a giant. In what follows we will introduce small positive constants $c_1, c_2, \dots, c_{10}$. They will not be specified exactly, but we will indicate their relative sizes. Our assumptions will be that

$$c_7 \ll c_6 \ll c_4 \ll c_3 \ll c_2 \ll c_1 \ll c_5 \ll 1.$$

The value of c_8 depends on c_1 and can be made arbitrarily small. c_9 is defined in terms of c_8, and c_{10} can be made arbitrarily small.

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

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

and

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

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

Let

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

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

$$\frac{n_0 p_{\max}}{k} > \sum_{x \in V_0} p_x = \frac{1}{n(n-1)} \sum_{i=1}^{m_0} |V_0 \cap A_i| + \frac{n_0 m_0}{n(n-1)} - \frac{n_0}{n^2(n-1)} \sum_{i=1}^{m_0} |A_i|.$$

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

$$\frac{n_0 \times 2m_0}{kn(n-1)} > \frac{n_0^2 m_0}{n^2(n-1)} + \frac{n - n_0}{n^2(n-1)} \sum_{i=1}^{m_0} |V_0 \cap A_i| \ge \frac{n_0^2 m_0}{n^2(n-1)},$$

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

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

For $i, j \in [0, k-1]$, let

$$M_{i,j} = \max\left\{ \frac{(i+j) p_{\max}}{k} |V_j| - c_2, 0 \right\}. \tag{5}$$

If $x \in V_i$, $y \in V_j$, then $M_{i,j}$ approximates the expected number of G_{m_0}-neighbors that x has in V_j.

Let $\iota = |I|$. Let M be the $\iota \times \iota$ matrix with rows and columns indexed by I, whose entries are $M_{i,j}$.

The definition of I implies that each entry of M is at least $\frac{2p_{\max}}{k} \times \frac{n}{k^2} - c_2 \ge 3\theta_0 c_1^3 - c_2$, which is bounded away 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 for some constant $\theta_1 > 0$, $\lambda_1 = \lambda_1^{(n)} \ge 1 + \theta_1$ for all sufficiently large n. Then whp we have the following:*

(a) *G_m contains a giant component.*

(b) *Furthermore the second largest component is of size $O(\ln n)$.*

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 BFS in $G_{m_0}[V]$ starting with the vertex x. Let C_x (initially $\{x\}$) be the current set of vertices we have added to $V(C)$. If $|C_x| \ge \nu_0$, 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 are incident to y together with the indices of the rounds in which they were generated. Add these neighbors of y to C_x. 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 ι 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 $S \subseteq [m_0]$ consist of the indices of rounds for which we have not yet exposed the

accepted edge. We have $|S| \geq m_0 - \nu_0 - O((\ln n)^2)$. 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 $2m_0$ edges generated. Furthermore, it is also easy to see that

$$\mathbf{Pr}(C_x \text{ contains a cycle}) = O(\nu_0^2/n). \tag{6}$$

(The probability that some unsaturated x chooses an unsaturated $y \in C_x$ is at most $4m_0\nu_0^2/n(n-1)$.)

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 $D_r = \{x_r, y_r\}$, $D'_r = \{x'_r, y'_r\}$, $r \in S$ with x_r, x'_r exposed to the builder are random subject to the following conditions:

- If $x_r \in A_r$, then $\{x_r, y_r\} \cap B_x = \emptyset$.
- If $x_r \notin A_r$, then $\{x'_r, y'_r\} \cap B_x = \emptyset$.

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 B_x can be.

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

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

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

$$p_{yz} - \frac{2\nu_0}{\binom{n}{2}} - O(n^{-2}) \geq \frac{p_{\max}(i+j)}{k} - \frac{5\nu_0}{n^2}. \tag{7}$$

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 $\mathbf{Pr}(y \in D_q \cup D'_q) \leq 4/n$. Also, $\lambda = \mathbf{E}(Y_j \mid \text{previous history}) \leq m_0(4/n) = O(1)$. It follows from Durrett [9, Chapter 2, equation (6.5)] that

$$\sup_{A \subseteq \mathbb{Z}} |\mathbf{Pr}(Y_j \in A) - \mathbf{Pr}(Po(\lambda) \in A)| \leq \frac{16}{n}, \tag{8}$$

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

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

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

$$\mathbf{E}(Y_j) \leq \frac{p_{\max}(i+j+2)}{k}|V_j| \leq M_{i,j} + 5\theta_0 c_1 |V_j|/n + c_2. \tag{9}$$

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 neighbors for each $j \in I$. Then we generate

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

blue neighbors 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 neighbors. (This demon will qs[1] 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 the first ν_0 progeny of the *idealized multitype 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 toward 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 $1 + c_2$ (because it is the maximum of $\|\mathbf{x}^T M^{(\infty)}\mathbf{x}\|$ over all unit vectors $\mathbf{x} \in \mathbb{R}^\iota$). At this point it is convenient to assume that $c_2 \ll \theta_1$. From [2, Theorem V.3.2], we can assume that $\rho_i^{(\infty)} < 1 - c_3$, $c_3 = c_3(\theta_1)$ for each $i \in I$. We can decrease c_3 if necessary and so the assumption $c_3 \ll c_2$ is acceptable.

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. 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.6.1], 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 $D = \lceil (1+c_2)^{r/2} \rceil$ children with probability p and none with probability $1-p$. The expected number of progeny of an individual in this process is at least $p(1+\theta_1)^{r/2} \ge c_3(1+\theta_1)^{r/2}/2 \ge 1 + \theta_1$ for r sufficiently large. We have $1 > p_0 = 1 - p + pp_0^D$. If $p_0 = 1 - p + \delta$, then $\delta = pp_0^D$. 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 $\lim_{l\to\infty} \rho_q^{(n_l)} = 1$.

[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.

So we can assume that $\rho_i^{(n)} \le 1 - c_4$ for some $c_4 \ll c_3$ and for any $i \in I$. Now for each $x \in V$ 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 $\mathbf{Pr}(z_x = 1) \ge c_4 - O(\nu_0/n)$. So if $Z = \sum_{x \in V} z_x$, then $\mathbf{E}(Z) \ge c_4 n/2$. With a view toward applying the Chebyshev inequality, we now estimate $\mathbf{Pr}(z_x = z_y = 1)$ for $x \in V_i$, $y \in V_j$, $i, j \in I$, with $x \ne y$:

$$\begin{aligned}\mathbf{Pr}(z_x = z_y = 1) &\le \mathbf{Pr}(z_x = 1 \mid z_y = 1,\, x \in C_y)\,\mathbf{Pr}(x \in C_y)\\ &\quad + \mathbf{Pr}(z_x = 1 \mid z_y = 1,\, x \notin C_y)\,\mathbf{Pr}(z_y = 1)\\ &= O(\nu_0/n) + \mathbf{Pr}(z_y = 1)(\mathbf{Pr}(z_x = 1) + O(\nu_0^2/n)). \qquad (10)\end{aligned}$$

To verify (10) we observe that

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

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 probability of $O(\nu_0/n)$.

Therefore

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

and so

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

Now let $\zeta_x = 1$ if the above BFS procedure when started with x *and allowed to continue* results in exposing a G_{m_0} component of size at least $n^{7/8}$, and let $\zeta_x = 0$ otherwise. We show that whp there exists a set X of size $O((\ln n)^3)$ such that

$$\zeta_x = z_x \text{ for all } x \notin X. \qquad (11)$$

We will show then that whp,

(12)
$x, y \in V$ and $\zeta_x = \zeta_y = 1$ implies that x and y lie in the same G_{m_1} component K_1.

Note that (11) and (12) imply that whp $|K_1| \ge c_4 n/4$ and so it is a giant component, which will prove part (a) of the lemma. We finally show that whp all components in G_{m_1} of size at least $\frac{16}{c_4 \varepsilon \eta} \ln n$ will be part of the same giant component of G_m. This will complete the proof of the lemma.

The exceptional vertices X are those for which either (i) we have to invoke the demon before exposing $n^{7/8}$ vertices during the BFS construction or (ii) C_x contains a cycle. We have $\mathbf{E}(|X|) \le 16\nu_0 + O(\nu_0^2/n)$ by (6), and so $|X| = O((\ln n)^3)$ whp.

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

$$\mathbf{Pr}(x \notin X \text{ and } z_x = 1 \text{ and } |C_x^{Blue}| \ge c_5\nu_0) \le n^{-2}. \qquad (13)$$

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 (9) that $|B_y|$ is

dominated by the number W of *proper* descendants in a single-type branching process where the number of progeny is $Po(\mu)$, $\mu = 5\theta_0 c_1 + c_2$. 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 setting $u = \ln(t/\mu)$,

$$\mathbf{Pr}(Po(\mu) \geq t) \leq e^{-tu}\mathbf{E}(e^{uPo(\mu)}) = e^{-\mu}\left(\frac{\mu e}{t}\right)^t \leq (3\mu)^t. \tag{14}$$

Now consider labeling the progeny of a branching process by sequences $i_1, i_2, \ldots, i_r$. This sequence is the i_rth child of the i_{r-1}th child of the $\ldots$ of the i_1th child of the root. The probability that such a particle exists is, by (14), 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 (13).

We now consider the red branching process at a point where $|C_x|$ reaches ν_0. The total progeny will be at least $\nu_1 = \nu_0 - c_5\nu_0$, after accounting for the blue vertices. We are now considering the case where $z_x = 1$, no demon has been invoked, and the blue branching process is small; see (13). We will now show that whp after $O(1)$ more rounds, a sufficient number of red progeny are at the bottom level. Suppose that s generations of the red branching process have been produced. Let $\mathbf{t}_{(0)} = [t_i, i \in I]^T$, 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. We do not include the initial vertex in this count. Then we have

$$b_j = \sum_{i=1}^{\iota} N_{i,j} - t_j. \tag{15}$$

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

$$\mathbf{Pr}(Po(\lambda) \geq \lambda + u) \leq \begin{cases} e^{-u/3}, & u \geq \lambda, \\ e^{-u^2/(3\lambda)}, & u \leq \lambda, \end{cases} \tag{16}$$

$$\mathbf{Pr}(Po(\lambda) \leq \lambda - u) \leq e^{-u^2/(2\lambda)}. \tag{17}$$

Now $b_j + t_j$ is the sum over $i = 1, 2, \ldots, \iota$ of t_i independent variables, each distributed as $Po(M_{i,j})$. Fix an i and imagine generating ν_0' such random variables $X_1, X_2, \ldots, X_{\nu_0'}$ distributed as $Po(M_{i,j})$. Then (16) and (17) imply that for some $c_7 \ll c_6 \ll c_1$,

$$\mathbf{Pr}(\exists t \leq \nu_0' : |X_1 + \cdots + X_t - tM_{i,j}| \geq c_6\nu_1) \leq \nu_0' e^{-c_7\nu_0}. \tag{18}$$

To see this, first note that $X_1 + \cdots + X_t$ is distributed as $Po(\lambda)$, $\lambda = tM_{i,j}$. We use (16) and (17) directly with $u = c_6\nu_1$ to get the bounds.

We put $\tau_j = t_j + b_j$ for $j \in I$.

It follows from (15) and (18) that

$$\mathbf{Pr}\left(\left|\tau_j - \sum_{i=1}^{\iota} t_i M_{i,j}\right| \geq c_6 k\nu_1\right) \leq e^{-c_7\nu_0/2}. \tag{19}$$

We continue the process for $\ell = O(1)$ more rounds to create $\mathbf{t}_{(1)}, \ldots, \mathbf{t}_{(\ell)}$. Note that $\mathbf{t}_{(1)} = (\tau_1, \ldots, \tau_\iota)$, and so

$$\|\mathbf{t}_{(1)}\|_1 \geq \nu_1. \tag{20}$$

Arguing as above with $\mathbf{t}_{(l-1)}$ taking the place of $\mathbf{t}_{(0)}$, we can then write

$$-c_6\nu_1\mathbf{e}^T \leq \mathbf{t}_{(l)}^T - \mathbf{t}_{(l-1)}^T M \leq c_6 k\nu_1\mathbf{e}^T,$$

where $\mathbf{e} = [1, 1, \ldots, 1]^T$.

Iterating, we get that qs

$$-c_6 k\nu_1\mathbf{e}^T(M^{\ell-2} + \cdots + I) \leq \mathbf{t}_{(\ell)}^T - \mathbf{t}_{(1)}^T M^{\ell-1} \leq c_6 k\nu_1\mathbf{e}^T(M^{\ell-2} + \cdots + I).$$

Now we can write

$$M^{\ell-1} = \lambda_1^{\ell-1}(\mathbf{u}\mathbf{v}^T + O(\zeta^{\ell-1})J), \tag{21}$$

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 λ_1, $0 \leq \zeta < 1$ 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, c_1, c_2.

Making ℓ large and using (20) we see that qs

$$\mathbf{t}_{(\ell)}^T \geq (1 - c_8)\lambda_1^{\ell-1}(\mathbf{t}_{(1)}^T\mathbf{u})\mathbf{v}^T, \tag{22}$$

where c_8 can be made arbitrarily small by a suitable choice of small c_6 and large ℓ.

We now return our attention to the actual BFS construction. Let $U_j^{(s)}$ denote the set of unsaturated vertices of type j at the time of construction of the $(s + \ell)$th red generation. Then, from (22), we see that given $z_x = 1$, we have that qs

$$|U_j^{(0)}| \geq c_9 v_j(\ln n)^2, \tag{23}$$

where $c_9 = ((1 - c_8)\lambda_1^\ell - (1 + c_8)\lambda_1^{\ell-1})(\mathbf{t}_{(1)}^T\mathbf{u}) - c_5 > 0$.

Now for each j, choose a subset of $U_j^{(0)}$ of size exactly $\lceil c_9 v_j(\ln n)^2\rceil$. 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

$$|U_j^{(t)}| = c_9(\lambda_1 - c_{10})^t v_j(\ln n)^2$$

for a small c_{10} and for all j.

Now $|U_j^{(t+1)}|$ is the sum of $\sum_i |U_i^{(t)}|(|V_j| - O(n^{7/8}))$ Bernoulli random variables and

$$\mathbf{E}(|U_j^{(t+1)}|) \geq \sum_i |U_i^{(t)}| M_{i,j} = c_9(\ln n)^2(\lambda_1 - c_{10})^t \sum_i v_i M_{i,j} = c_9(\lambda_1 - c_{11})^t \lambda_1 (\ln n)^2 v_j.$$

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

$$|U_j^{(t+1)}| \geq c_9(\lambda_1 - c_{10})^{t+1}(\ln n)^2 v_j.$$

Then taking subsets we can assume that

$$|U_j^{(t+1)}| = c_9(\lambda_1 - c_{10})^{t+1}(\ln n)^2 v_j.$$

It follows, after iterating $O(\ln n)$ times, that we reach τ such that $|U_j^{(\tau)}| = \Theta(n^{7/8})$ for each j. This verifies (11).

Now suppose that $\zeta_x = \zeta_y = 1$ and $C_x \cap C_y = \emptyset$. The probability that none of the rounds $r = m_0 + 1, \ldots, m_1$ connects C_x and C_y is at most

$$\prod_{i,j \in I} \left(1 - \left(\frac{n^{7/4}}{n^2}\right)^2\right)^{m_1 - m_0} = e^{-\Omega(n^{1/2})}.$$

The term $\left(\frac{n^{7/4}}{n^2}\right)^2$ is a lower bound on the probability that *both* edges D_r, D_r' connect C_x and C_y.

This proves (12) and shows that G_{m_1} has a giant component K_1, where $|K_1| \geq c_4 n/4$.

We now consider two cases. Fix a component $C \neq K_1$ of G_{m_1} that is of size at least $\kappa \ln n$.

Case 1. $|\{i \in [m_1 + 1, m] : |A_i \cap K_1| \geq |K_1|/2\}| \geq (m - m_1)/2$.

Fix $r \in [m_1 + 1, m]$. If $|A_i \cap K_1| \geq |K_1|/2$, then the probability that Paul accepts D_r in this round and it joins K_1 and C is at least $\frac{|K_1|}{2n} \cdot \frac{|C|}{n} \geq \frac{c_4 \kappa \ln n}{8n}$. So

$$\mathbf{Pr}(K_1, C \text{ are in different } G_m \text{ components}) \leq \left(1 - \frac{c_4 \kappa \ln n}{8n}\right)^{(m - m_1)/2} = o(n^{-1})$$

if $\kappa \geq \frac{50}{c_4 \varepsilon}$.

Case 2. $|\{i \in [m_1 + 1, m] : |A_i \cap K_1| < |K_1|/2\}| \geq (m - m_1)/2$.

Fix $r \in [m_1 + 1, m]$. If $|A_i \cap K_1| < |K_1|/2$, then the probability that Paul accepts D_r' in this round and it joins K_1 and C is at least $\left(1 - \frac{|K_1|}{2n}\right) \cdot \frac{|K_1|}{2n} \cdot \frac{|C|}{n} \geq \frac{c_4 \kappa \ln n}{9n}$. So

$$\mathbf{Pr}(K_1, C \text{ are in different } G_m \text{ components}) \leq \left(1 - \frac{c_4 \kappa \ln n}{9n}\right)^{(m - m_1)/2} = o(n^{-1})$$

if $\kappa \geq \frac{60}{c_4 \varepsilon}$.

This completes the proof of the lemma. ∎

2.2. A sufficient criterion for the nonexistence of a giant. Let p_x, p_{xy}, $p_{\max}, c_1, c_2, k, V_i$ be as above. We redefine M to be the $k \times k$ matrix with entries

$$M_{i,j} = \frac{((i+1)+(j+1))p_{\max}}{k}|V_j| + c_1^2, \quad 0 \le i, j \le k-1. \tag{24}$$

LEMMA 5. *Suppose that there are $c_2 > 0$ and n_0 such that for all $n \ge n_0$ 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 BFS is *dominated* by the Poisson multitype 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 $\mathcal{B}_i$, for any $0 \le i \le k-1$, reaches at least $C \ln n$ vertices is at most n^{-2}. 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. *For all $\varepsilon > 0$ and integer t 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, $s \ge s_0$, 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 $\mathbf{n} = (n_1, \dots, n_t)^T$, where n_i is the total number of particles of type i generated.

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

$$n_{i,j} \ge (1+\varepsilon_0) n_i M_{i,j} + \varepsilon_0 s, \tag{25}$$

where $n_{i,j}$ is the total output of type j particles that are produced (looking forward and including the last 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 state only that whatever feasible values these variables have, they must satisfy (25) for some i, j. Suppose on the contrary that (25) is false for all i, j. Let $\mathbf{1} = (1, \dots, 1)^T$ be the all 1's column vector. For every $j \in [t]$ we have $n_j \le \sum_{i=1}^t n_{i,j} + 1$, where the last term accounts for the initial vertex which is not born but given. This implies that we have the following coordinatewise domination:

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

Taking the l_2-norm and using the triangle inequality and the fact that $\|M^T \mathbf{n}\| \le (1-\varepsilon)\|\mathbf{n}\|$, 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 $\|\mathbf{n}\| \ge s/\sqrt{t}$, which contradicts the choice of ε_0.

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

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

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 (25) holds is at most the probability that there exist $n_1, \dots, n_t$ with $s = n_1 + \cdots + n_t \geq s_0$ and $i, j \in [t]$ such that

$$X_{i,j,1} + \cdots + X_{i,j,n_i} \geq (1+\varepsilon_0) n_i M_{i,j} + \varepsilon_0 s. \tag{26}$$

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 (25) 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 (16). This implies the existence of the required $\delta > 0$. ∎

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 a 2×2 matrix which for large n is approximately

$$\mu \times \begin{pmatrix} (4-2\alpha)\alpha & (3-2\alpha)(1-\alpha) \\ (3-2\alpha)\alpha & (2-2\alpha)(1-\alpha) \end{pmatrix}. \tag{27}$$

The largest eigenvalue of this limiting matrix is

$$\lambda_1 = \mu(1 + \sqrt{1 + \alpha - \alpha^2}). \tag{28}$$

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 $\lambda_1 \geq 1 + 2\theta_1$ for some positive constant θ_1. Since the limit of $\lambda_1^{(n)}$ is λ_1, we have $\lambda_1^{(n)} \geq 1 + \theta_1$ 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 (24). 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 $\nu_i = |V_i|$. Further, $\max_{i,j} |\widehat{M}(i,j)| \le c_1^2$ and $\mathbf{1}^T D\mathbf{1} = \sum_{i=0}^{k-1} \nu_1 = n$.

Now for $|\mathbf{x}| = 1$, we have $|\mathbf{x}^T \widehat{M}\mathbf{x}| \le kc_1^2 \le 2c_1$, 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

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

appears in 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 is not).

Write

$$\mathbf{v} = \mu\mathbf{p} + \nu\mathbf{1}, \tag{29}$$

with $\mu, \nu \in \mathbb{R}$. Note that $\mu \neq 0$. Substituting (29) 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

$$\lambda_1 \frac{\lambda_1 - \mathbf{1}^T D\mathbf{p}}{1^T D\mathbf{1}} = \mathbf{p}^T D\mathbf{p} + \frac{\lambda_1 - \mathbf{1}^T D\mathbf{p}}{1^T D\mathbf{1}} (\mathbf{p}^T D\mathbf{1}). \tag{30}$$

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 (30) is

$$\lambda_1 = \mathbf{1}^T D\mathbf{p} + \sqrt{n\mathbf{p}^T D\mathbf{p}}. \tag{31}$$

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

$$\mathbf{1}^T D\mathbf{p} \le \frac{m}{n-1} + \frac{np_{\max}}{k} \le \frac{m}{n-1} + \frac{1}{k}. \tag{32}$$

Suppose we fix γ and we want to maximize the right-hand side of (31). 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

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

Now, if we want to maximize this expression over reals, given $\gamma = \frac{1}{n}\sum_x \gamma_x$ and the constraints $0 \le \gamma_x \le m$, then we should take for 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 $\sum_{x\in[n]} \gamma_x^2 \le mn\gamma + m^2$ and

$$n\mathbf{p}^T D\mathbf{p} \le \frac{(m^2-\gamma^2)n + mn\gamma}{n(n-1)^2} + \frac{3}{k}. \tag{33}$$

Going back to (31), we see from (32) and (33) that

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

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

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

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 the giant any longer; that is, we 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 $|V| \ge (1-3/k)n$. 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 multitype 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 at the intersection of a row and a column that is 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 $l(1-4/k) \le \rho \le l$. 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, $(n/l) \times \mathbf{1}^T L \mathbf{1} \ge (2 - O(1/k))e(G_m)$, so

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

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

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

which implies that the largest eigenvalue of L is at least $1+\varepsilon$. Now apply Lemma 4 and the proof is complete. ∎

Appendix. Proofs of (16) and (17). Let $e^\theta = 1 + u/\lambda$. Then

$$\begin{aligned} \mathbf{Pr}(Po(\lambda) \ge \lambda + u) &\le e^{-\theta(\lambda+u)} \mathbf{E}(e^{\theta Po(\lambda)}) \\ &= \exp\{\lambda(e^\theta - 1) - \theta(\lambda + u)\} \\ &= \exp\{u(1 - (1+\lambda/u)\ln(1+u/\lambda)\}. \end{aligned} \tag{34}$$

If $u \le \lambda$, then we use

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

for $0 \le x \le 1$.

Plugging $x = u/\lambda$ into (34) gives the required inequality (16) in the case $u \le \lambda$. If $\lambda \le u$, then we use

$$(1 + 1/x)\ln(1+x) \ge 2\ln 2 = 1.386924\ldots$$

for $x \ge 1$.

For (17) we have, with $e^{-\theta} = 1 - u/\lambda$ and using $(1 - 1/x)\ln(1-x) - 1 \le -x/2$ for $0 \le x \le 1$,

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

REFERENCES

[1] N. Alon, *A note on network reliability*, in Discrete Probability and Algorithms, IMA Vol. Math. Appl. 72, Springer-Verlag, New York, 1995, pp. 11–14.
[2] K. B. Athreya and P. E. Ney, *Branching Processes*, Springer-Verlag, New York, 1972.
[3] A. Beveridge, T. Bohman, A. Frieze, and O. Pikhurko, *Product rule wins a competitive game*, Proc. Amer. Math. Soc., 135 (2007), pp. 3061–3071.
[4] T. Bohman and A. Frieze, *Avoiding a giant component*, Random Structures Algorithms, 19 (2001), pp. 75–85.
[5] T. Bohman, A. Frieze, and N. Wormald, *Avoidance of a giant component in half the edge set of a random graph*, Random Structures Algorithms, 25 (2004), pp. 432–449.
[6] T. Bohman and J. H. Kim, *A phase transition for avoiding a giant component*, Random Structures Algorithms, 28 (2006), pp. 195–214.
[7] T. Bohman and D. Kravitz, *Creating a giant component*, Combin. Probab. Comput., 15 (2006), pp. 489–511.
[8] B. Bollobás, S. Janson, and O. Riordan, *The phase transition in inhomogeneous random graphs*, Random Structures Algorithms, 31 (2007), pp. 3–122.
[9] R. Durrett, *Probability: Theory and Examples*, 2nd ed., Duxbury Press, Belmont, CA, 1996.
[10] A. Flaxman, D. Gamarnik, and G. B. Sorkin, *Embracing the giant component*, in LATIN 2004: Theoretical Informatics, M. Farach-Coltin, ed., Springer-Verlag, Berlin, 2004, pp. 69–79.
[11] S. Janson, T. Łuczak, and A. Ruciński, *Random Graphs*, Wiley-Intersci. Ser. Discrete Math. Optim., Wiley-Interscience, New York, 2000.
[12] B. Söderberg, *General formalism for inhomogeneous random graphs*, Phys. Rev. E (3), 66 (2002), 066121.
[13] J. Spencer and N. Wormald, *Birth control for giants*, Combinatorica, 27 (2007), pp. 587–628.