Failure resilience in balanced overlay networks

Ayalvadi Ganesh Laurent Massoulié
Microsoft Research, 7 J J Thomson Avenue
Cambridge CB3 0FB, United Kingdom
{ajg, lmassoul}@microsoft.com

Abstract

The growth of peer-to-peer applications on the Internet has motivated interest in general purpose overlay networks. A basic requirement of an overlay is that it be connected, so that any two peers can communicate, and that it retain connectivity in the presence of node and link failures. In this paper, we use random graphs to model of a class of overlay networks and study their connectivity in the presence of link failures.

1 Introduction

Peer-to-peer applications currently constitute one of the fastest growing uses of the Internet. A challenging research problem is to develop overlay architectures that can support such applications without overloading end systems. Features of peer-to-peer systems such as the transience of users and the absence of powerful central servers motivate fully decentralized schemes that do not require global knowledge of membership, and that are robust to node and link failures. Scalability is crucial in this context, and overlays must be built ensuring that the memory and communication load on each node grows slowly in the overlay size, while the global load is evenly balanced among the nodes. The basic property required of an overlay network is that it be connected, so that any two peers can communicate over it. Moreover, connectivity should be maintained in the presence of failures or temporary disconnections of some fraction of member nodes.

We shall model the overlay as a graph with nodes i representing members, and a directed edge (i, j) representing that i knows j. If the "who knows who" relation is symmetric, the edges are taken to be undirected. A classical random graph model, which was studied in detail by Erdös and Renyi [4], has a collection of n nodes; an edge is present between every pair of nodes with probability p (which may depend on n), independent of other edges. In the undirected case, they showed that there is a threshold for connectivity at a mean degree of $\log(n)$: precisely, if the mean degree is $\log n + x$, then the probability of connectivity goes to $e^{-e^{-x}}$ as $n \to \infty$. This was extended by Ball and Barbour [1] to random directed graphs with a wide range of degree distributions. They showed a similar threshold at a mean degree of $\log(n)$ (irrespective of the degree distribution) for the probability that there is a directed path from any specified node to all other nodes.

These results suggest that the membership protocol should be designed so as to provide each member with a set of neighbours whose size is of order $\log n$. Then, the memory requirements on each member grow slowly in the overlay size, but the overlay network nevertheless has the desired properties. In [5], a decentralized algorithm is described for constructing directed overlays in which each node has a mean out-degree of $c \log n$, where c is a specified design parameter. By the result of Ball and Barbour quoted above, these overlays can tolerate independent link failures at rates p up to $(c-1)/c$, while retaining connectivity in the sense described, with high probability. A similar result holds in the undirected case. Moreover, it turns out that at higher link failure rates, connectivity is initially lost due to the isolation of single nodes which had lower than average degree. This suggests that balancing node degrees should improve resilience to failures. In [6], a Metropolis algorithm is proposed for reshaping the overlays constructed in [5] so as to balance node degrees, without altering the mean degree, in order to achieve load balancing as well as to improve fault tolerance. The Metropolis scheme corresponds to a random walk on the space of undirected graphs on n nodes with E edges. It is shown in [6] that the equilibrium distribution of this random walk is given by:

$$\mu_n(G) = Z^{-1} \exp(-\beta \sum_{i=1}^{n} d_i^2), \tag{1}$$

where d_i denotes the degree of node i in the graph G, β is a specified parameter (inverse temperature) of the Metropolis algorithm, and Z is a normalizing constant. In the remainder of the paper, we shall study the connectivity and failure resilience properties of random graphs generated according to this distribution.

2 Degree distributions

In this section, we show for the random graph model described above that the node degrees concentrate about their mean value. Specifically, we show that the variance of node degrees remains bounded as $n \to \infty$ and, with high probability, all node degrees are within $O(\sqrt{\log n})$ of the mean. This is in contrast to the Erdös-Renyi model, where the variance grows like $\log n$, and the maximum fluctuation of node degrees is $O(\log n)$.

The probability measure μ_n on graphs induces a probability measure on degree distributions, which we denote by π_n. For $\mathbf{d} = (d_1, \ldots, d_n)$,

$$\pi_n(\mathbf{d}) = \frac{1}{Z_n} G_n(\mathbf{d}) e^{-\beta \sum_{i=1}^{n} d_i^2} \mathbf{1}_{\sum_{i=1}^{n} d_i = 2E}, \tag{2}$$

where $G_n(\mathbf{d})$ is the number of graphs having the degree sequence $\mathbf{d}$, and Z_n is a normalizing constant. We can rewrite the above as

$$\begin{aligned} \pi_n(\mathbf{d}) &= \frac{1}{Z_n(\gamma)} \left[\frac{E! 2^E}{(2E)!} G_n(\mathbf{d}) \prod_{i=1}^{n} (d_i!) \right] \prod_{i=1}^{n} \frac{1}{d_i!} e^{-\beta d_i^2 + \gamma (\log n) d_i} \mathbf{1}_{\sum_{i=1}^{n} d_i = 2E} \\ &= \frac{\tilde{G}_n(\mathbf{d})}{Z_n(\gamma)} \prod_{i=1}^{n} \frac{1}{d_i!} e^{-\beta d_i^2 + \gamma (\log n) d_i} \mathbf{1}_{\sum_{i=1}^{n} d_i = 2E}. \end{aligned} \tag{3}$$

The introduction of the tilt parameter γ doesn't change the distribution as it multiplies $\pi_n(\mathbf{d})$ by $e^{2\gamma E \log n}$. This is a constant since the total number of edges is fixed. Thus, it can be absorbed into the normalization factor $Z_n(\gamma)$ along with the term $E! 2^E / (2E)!$.

Random graphs with a specified degree sequence can be constructed using the configuration model [2], which we describe briefly. Associated with each vertex i are d_i labelled configuration points, called half-edges. A configuration is a matching of configuration points. If a configuration point of i is matched with a configuration point of j, this is interpreted as an edge between i and j. If two configuration points of i are matched, this corresponds to a loop. Thus, each configuration corresponds to a multigraph. Every simple graph corresponds to exactly $\prod_{i=1}^n d_i!$ distinct configurations since the d_i edges incident on node i can be assigned to its d_i configuration points in $d_i!$ ways.

The configuration model implies the upper bound $\tilde{G}_n(\mathbf{d}) \leq 1$ for any degree sequence $\mathbf{d}$. Moreover, if the maximum degree $\Delta = o(E^{1/4})$, then McKay and Wormald [8] establish the equivalent

$$\tilde{G}_n(\mathbf{d}) \sim e^{-\lambda-\lambda^2}, \text{ where } \lambda = \frac{1}{4E}\sum_{i=1}^n d_i(d_i-1). \tag{4}$$

Given a degree sequence $\mathbf{d}$, we define the mean degree $\overline{d} = \sum_{i=1}^n d_i/n$ and the variance $\mathrm{Var}(\mathbf{d}) = \frac{1}{n}\sum_{i=1}^n (d_i - \overline{d})^2$. We are interested in a regime where $\overline{d} = c\log n$ for some specified constant c, so that $E = cn\log n/2$. For fixed constants α_1, α_2, c_1 and c_2, we define the following sets of degree sequences:

$$\begin{aligned} A &= \{\mathbf{d}:\ \overline{d} = c\log n\}, \\ A_1(\alpha_1,\alpha_2) &= \{\mathbf{d}:\ -\sqrt{\alpha_1\log n} \leq d_i - \overline{d} \leq \sqrt{\alpha_2 \log n}\ \forall\, i\}, \\ A_2 &= \{\mathbf{d}:\ d_i \leq n^{1/4}\ \forall\, i\}. \end{aligned}$$

Define $\hat{A}_1(\alpha_1,\alpha_2) = A \cap A_1(\alpha_1,\alpha_2)$, $\hat{A}_2 = A \cap A_2$.

We show below that $\pi_n(A_2^c)$ is negligible, and so we need only consider graphs with degree sequences $\mathbf{d} \in A_2$. For such graphs, the estimate in (4) is applicable. We can use this to show that $\pi_n(\hat{A}_1(\alpha_1,\alpha_2)) \to 1$ as $n \to \infty$, for suitable α_1 and α_2.

Observe from (4) that $4E\lambda = n(\mathrm{Var}(\mathbf{d}) + \overline{d}^2 - \overline{d})$. But if $\mathbf{d} \in \hat{A}_1(\alpha_1,\alpha_2)$, then $\mathrm{Var}(\mathbf{d}) \leq \max\{\alpha_1,\alpha_2\}\log n$, and so $\lambda \leq \frac{1}{2}(c\log n - 1 + \frac{1}{c}\max\{\alpha_1,\alpha_2\})$. Hence,

$$\mathbf{d} \in \hat{A}_1(\alpha_1,\alpha_2) \ \Rightarrow\ \frac{1}{\tilde{G}_n(\mathbf{d})} \sim e^{\lambda+\lambda^2} \leq e^{\frac{c^2\log^2 n}{2}},$$

for all n sufficiently large. Recall that $\tilde{G}_n(\mathbf{d}) \leq 1$ for all $\mathbf{d}$ and, in particular, for $\mathbf{d} \in A_2^c$, the complement of A_2. Thus, we have from (3) that, for n sufficiently large,

$$\frac{\pi_n(A_2^c)}{\pi_n(\hat{A}_1(\alpha_1,\alpha_2))} \leq e^{\frac{c^2\log^2 n}{2}} \frac{\sum_{\mathbf{d}\in A_2^c}\prod_{i=1}^n \frac{1}{d_i!}e^{-\beta d_i^2+\gamma(\log n)d_i}}{\sum_{\mathbf{d}\in \hat{A}_1(\alpha_1,\alpha_2)}\prod_{i=1}^n \frac{1}{d_i!}e^{-\beta d_i^2+\gamma(\log n)d_i}}. \tag{5}$$

Let $D_1,\dots,D_n$ be independent and identically distributed (iid) random variables, with

$$P(D_1 = k) = \frac{1}{F(\gamma)}\frac{1}{k!}e^{-\beta k^2+\gamma(\log n)k}, \quad ,k \in \mathbb{N}, \tag{6}$$

where $F(\gamma)$ is a normalization constant. The dependence of the D_i on n and γ hasn't been made explicit in the notation. We choose γ so that $\mathbf{E}D_1 = c\log n$, for a specified constant, c; this is possible because $\mathbf{E}D_1$ is a continuous function of γ.

Let $\mathbf{D}$ denote the random vector $(D_1,\dots,D_n)$. We can now rewrite (5) as

$$\frac{\pi_n(A_2^c)}{\pi_n(\hat{A}_1(\alpha_1,\alpha_2))} \leq e^{\frac{c^2\log^2 n}{2}} \frac{\mathbf{P}(\mathbf{D}\in A_2^c)}{\mathbf{P}(\mathbf{D}\in \hat{A}_1(\alpha_1,\alpha_2))}. \tag{7}$$

The advantage of this is that we have reformulated statements about graphs as statements about iid random vectors. If we can show that the RHS above is small, then necessarily $\pi_n(A_2^c)$ is small, as we set out to show.

We now state some properties of the random variable D_1 that can be derived by straightforward but tedious calculation. We omit the proofs due to lack of space.

Lemma 1 *Let $k_\gamma - 1$ denote the integer part of $x_\gamma = \frac{1}{2\beta}\left(\gamma \log n + \log\log n + \frac{\gamma}{2\beta}\right)$, and let $\eta = 2\beta(x_\gamma - k_\gamma + \frac{1}{2})$. Then, $\mathbf{E}D_1 - k_\gamma$ and $Var(D_1)$ remain bounded as n to ∞, and the moment generating function of D_1 satisfies*

$$\mathbf{E}[e^{\theta D_1}] \sim e^{\theta k_\gamma} \frac{\psi(\theta+\eta)}{\psi(\eta)}$$

where $\psi(\theta) = \sum_{j=-\infty}^{\infty} e^{\theta j - \beta j^2} / \sum_{j=-\infty}^{\infty} e^{-\beta j^2}$.

Intuitively, the lemma says that D_1 has mean of order $\log n$ and constant variance; moreover, the moment generating function of $D_1 - k_\gamma$ remains bounded as $n \to \infty$. Finally, given a constant c, the parameter γ can be tuned so that $\mathbf{E}D_1 = c \log n$. It can be shown using the above that

$$\mathbf{P}(\mathbf{D} \in A_2^c) \le \sum_{i=1}^{n} \mathbf{P}(D_i > n^{1/4}) \le K n e^{-\beta\sqrt{n}}, \tag{8}$$

for some constant K.

Define the centered random variables, $X_{nj} = D_j - \mathbf{E}D_j$. It is a straightforward consequence of the above lemma that:

Lemma 2 *The random variables, $\{X_{nj}, j = 1, \ldots, n, n \in \mathbb{N}\}$, satisfy the following conditions:*
(i) $\limsup_{n\to\infty} \mathbf{E}[e^{\theta |X_{n1}|}] < \infty$ *for some* $\theta > 0$.
(ii) $\liminf_{n\to\infty} Var(X_{n1}) > 0$.
(iii) $\liminf_{n\to\infty} \sum_{j=-\infty}^{\infty} \min\{\mathbf{P}(X_{n1} = j), \mathbf{P}(X_{n1} = j+1)\} > 0$.

Using the above estimates and Chernoff's bound, it can easily be shown that, given any $K > 0$, we can choose α_1 and α_2 such that, for all n sufficiently large,

$$\mathbf{P}(\mathbf{D} \in A_1(\alpha_1, \alpha_2)^c) < e^{-K \log n}. \tag{9}$$

Let $(\tilde{D}_1, \ldots, \tilde{D}_n)$ have the joint distribution of $(D_1, \ldots, D_n)$ conditional on $\mathbf{D} \in A_1(\alpha_1, \alpha_2)$. Equivalently, $\tilde{D}_1, \ldots, \tilde{D}_n$ are iid, with $\tilde{D}_j$ having the distribution of D_j conditional on $-\alpha_1\sqrt{\log n} \le D_j - \mathbf{E}D_j \le \alpha_2\sqrt{\log n}$. Now

$$\begin{aligned}\mathbf{P}(\mathbf{D} \in \hat{A}_1(\alpha_1, \alpha_2)) &= \mathbf{P}(\mathbf{D} \in A_1(\alpha_1, \alpha_2))\mathbf{P}(\sum_{j=1}^{n} D_j = cn\log n | \mathbf{D} \in A_1(\alpha_1, \alpha_2)) \\ &= \mathbf{P}(\mathbf{D} \in A_1(\alpha_1, \alpha_2))\mathbf{P}(\sum_{j=1}^{n} \tilde{D}_j = cn\log n). \end{aligned} \tag{10}$$

Suppose $\alpha_1, \alpha_2 > 0$ are chosen so that $E\tilde{D}_1 = ED_1 = c\log n$. It can be verified that Lemma 2 holds for $\tilde{D}_j$ as well as for D_j. Properties (i)-(iii) established by this lemma are precisely what we need to apply a local limit theorem of McDonald [7, Theorem 1]; using this theorem, we obtain that,

$$\mathbf{P}(\sum_{j=1}^{n} \tilde{D}_j = cn\log n) = \frac{1}{\sqrt{2\pi n\tilde{\sigma}}}\Big(1 + O\Big(\frac{1}{n}\Big)\Big), \tag{11}$$

where $\tilde{\sigma} = \mathrm{Var}(\tilde{D}_1)$ remains bounded as $n \to \infty$. Combining this with (10) yields

$$\mathbf{P}(\mathbf{D} \in \hat{A}_1(\alpha_1, \alpha_2)) = \frac{1}{\sqrt{2\pi n\tilde{\sigma}}}\Big(1 + O\Big(\frac{1}{n}\Big)\Big).$$

Hence, by (7) and (8),

$$\pi_n(A_2^c) \leq \frac{\pi_n(A_2^c)}{\pi_n(\hat{A}_1(\alpha_1,\alpha_2))} \leq e^{\frac{c^2\log^2 n}{2}} \sqrt{2\pi}\tilde{\sigma} K n^{3/2} e^{-\beta\sqrt{n}}\Big(1+O\Big(\frac{1}{n}\Big)\Big), \tag{12}$$

and so $\pi_n(A_2^c) \to 0$ as $n \to \infty$. Thus, we need only consider graphs with degree sequences in A_2, for which we can use the estimate in (4).

Theorem 1 *There exist constants α_1, α_2 such that $\pi_n(\hat{A}_1(\alpha_1,\alpha_2))$ goes to 1 as n goes to infinity.*

Proof: Observe that

$$\begin{aligned}\pi_n(\hat{A}_1(\alpha_1,\alpha_2)) &= \pi_n(A) - \pi_n(A \setminus A_1(\alpha_1,\alpha_2)) \\ &\geq \pi_n(A) - \pi_n((A \setminus A_1(\alpha_1,\alpha_2)) \cap A_2) - \pi_n(A_2^c).\end{aligned}$$

But $\pi_n(A) = 1$ by definition, and we have shown above that $\pi_n(A_2^c) \to 0$ as $n \to \infty$. Hence, it suffices to show that

$$\pi_n((A \setminus A_1(\alpha_1,\alpha_2)) \cap A_2) \to 0 \text{ as } n \to \infty. \tag{13}$$

Recall from (4) that, if $\mathbf{d} \in A_2$, then $\tilde{G}_n(\mathbf{d}) \sim e^{-\lambda(\mathbf{d})-\lambda(\mathbf{d})^2}$. Now,

$$\lambda(\mathbf{d}) = \frac{\mathrm{Var}(\mathbf{d}) + \overline{d}^2 - \overline{d}}{2\overline{d}} \geq \frac{c\log n - 1}{2} \quad \forall\, \mathbf{d} \in A,$$

since the mean degree, $\overline{d}$, is $c\log n$. In particular, this lower bound on $\lambda(\mathbf{d})$ holds for all $\mathbf{d}$ in $(A \setminus A_1(\alpha_1,\alpha_2)^c) \cap A_2$, since this is a subset of A. In addition, we saw earlier that, if $\mathbf{d} \in \hat{A}_1(\alpha_1,\alpha_2)$, then

$$\lambda(\mathbf{d}) \leq \frac{1}{2}\Big(c\log n - 1 + \frac{1}{c}\max\{\alpha_1,\alpha_2\}\Big).$$

Denote $\max\{\alpha_1,\alpha_2\}$ by α. Now, by (3),

$$\begin{aligned}&\frac{\pi_n((A \setminus A_1(\alpha_1,\alpha_2)) \cap A_2)}{\pi_n(\hat{A}_1(\alpha_1,\alpha_2))} \\ &= \frac{\sum_{\mathbf{d}\in(A\setminus A_1(\alpha_1,\alpha_2))\cap A_2} e^{-\lambda(\mathbf{d})-\lambda(\mathbf{d})^2}\prod_{i=1}^n \frac{1}{d_i!}e^{-\beta d_i^2+\gamma(\log n)d_i}}{\sum_{\mathbf{d}\in\hat{A}_1(\alpha_1,\alpha_2)} e^{-\lambda(\mathbf{d})-\lambda(\mathbf{d})^2}\prod_{i=1}^n \frac{1}{d_i!}e^{-\beta d_i^2+\gamma(\log n)d_i}} \\ &\leq e^{\frac{\alpha}{2c}\left(c\log n+\frac{\alpha}{2c}\right)}\frac{\sum_{\mathbf{d}\in(A\setminus A_1(\alpha_1,\alpha_2))\cap A_2}\prod_{i=1}^n \frac{1}{d_i!}e^{-\beta d_i^2+\gamma(\log n)d_i}}{\sum_{\mathbf{d}\in\hat{A}_1(\alpha_1,\alpha_2)}\prod_{i=1}^n \frac{1}{d_i!}e^{-\beta d_i^2+\gamma(\log n)d_i}}.\end{aligned}$$

In other words, there are constants κ_1 and κ_2 such that

$$\begin{aligned}\frac{\pi_n((A \setminus A_1(\alpha_1,\alpha_2)) \cap A_2)}{\pi_n(\hat{A}_1(\alpha_1,\alpha_2))} &\leq \kappa_1 e^{\kappa_2\log n}\frac{\mathbf{P}(\mathbf{D} \in (A \setminus A_1(\alpha_1,\alpha_2)) \cap A_2)}{\mathbf{P}(\mathbf{D} \in \hat{A}_1(\alpha_1,\alpha_2))} \\ &\leq \kappa_1 e^{\kappa_2\log n}\frac{\mathbf{P}(\mathbf{D} \in A \setminus A_1(\alpha_1,\alpha_2))}{\mathbf{P}(\mathbf{D} \in \hat{A}_1(\alpha_1,\alpha_2))}.\end{aligned} \tag{14}$$

Now, by (9), for any given $K > 0$, we can choose α_1 and α_2 such that $\mathbf{P}(\mathbf{D} \in A_1(\alpha_1,\alpha_2)^c) \leq e^{-K\log n}$. Thus,

$$\mathbf{P}(\mathbf{D} \in A \setminus A_1(\alpha_1,\alpha_2)) \leq \mathbf{P}(\mathbf{D} \in A_1(\alpha_1,\alpha_2)^c) \leq e^{-K\log n}. \tag{15}$$

Moreover, analogous to (11), we have

$$\mathbf{P}(\mathbf{D} \in A) = P(\sum_{j=1}^{n} D_j = cn\log n) = \frac{1}{\sqrt{2\pi n}\sigma}\left(1 + O\left(\frac{1}{n}\right)\right),$$

where $\sigma = \mathrm{Var}(D_1)$ remains bounded as $n \to \infty$. Therefore,

$$\begin{aligned}\mathbf{P}(\mathbf{D} \in \hat{A}_1(\alpha_1, \alpha_2)) &= \mathbf{P}(\mathbf{D} \in A) - \mathbf{P}(\mathbf{D} \in A \cap A_1(\alpha_1, \alpha_2)^c) \\ &\geq \mathbf{P}(\mathbf{D} \in A) - \mathbf{P}(\mathbf{D} \in A_1(\alpha_1, \alpha_2)^c) \\ &= \frac{1}{\sqrt{2\pi n}\sigma}\left(1 + O\left(\frac{1}{n}\right)\right). \qquad (16)\end{aligned}$$

Substituting (15) and (16) in (14), we have

$$\begin{aligned}\pi_n((A \setminus A_1(\alpha_1, \alpha_2)) \cap A_2) &\leq \frac{\pi_n(A \setminus A_1(\alpha_1, \alpha_2))}{\pi_n(\hat{A}_1(\alpha_1, \alpha_2))} \\ &\leq \kappa_1 \sigma \sqrt{2\pi n}\, e^{(\kappa_2 - K)\log n}\left(1 + O\left(\frac{1}{n}\right)\right).\end{aligned}$$

Since K can be chosen arbitrarily large, the above quantity goes to zero as $n \to \infty$, which establishes (13) and the claim of the theorem. □

3 Failure resilience

In the following, we work with graphs whose degree sequence belongs to the set $A_1(\alpha_1, \alpha_2)$ for some specified α_1 and α_2. We are interested in the probability that the graph remains connected when links fail independently with probability p. It is straightforward to compute the probability that a given node i becomes isolated due to link failures; it is simply p^{d_i}. Thus, by the union bound, the probability that some node becomes isolated is at most

$$\sum_{i=1}^{n} p^{d_i} \leq np^{c\log n - \sqrt{\alpha_1 \log n}} = \exp[(1 + c\log p)\log n - c\log p\sqrt{\alpha_1 \log n}].$$

Hence, if $c\log p < -1$ or, equivalently, $p < \exp(-1/c)$, then the probability that some node becomes isolated goes to zero as n increases to infinity.

By way of comparison, consider the classical random graph model of Erdös and Renyi [4] with the same mean degree. Here, an edge is present between each pair of nodes with probability $c\log n/n$, independent of all other edges. After taking failures into account, the edge probability becomes $(1-p)c\log n/n$, and the presence of edges continues to be mutually independent. It is well known for this model that, if $(1-p)c < 1$, then the graph is disconnected with high probability. Moreover, in a sense that can be made precise, the main reason for disconnection when $(1-p)c$ is "close to" 1 is the isolation of individual nodes. Intuitively, these arguments suggest that random graphs drawn from the distribution (1) can tolerate link failure rates up to $e^{-1/c}$ while retaining connectivity, whereas classical random graphs with the same mean degree can only tolerate failure rates up to $(c-1)/c$. We now establish this claim rigorously.

Given a graph G and a subset U of its vertex set, let $e_U(G)$ denote the number of edges incident within U (i.e., having both their vertices with U); let $e_{U,U^c}(G)$ denote the number of edges having one vertex in U and the other in its complement, U^c. Let

$d(G) = \{d_1, d_2, \dots, d_n\}$ denote the degree sequence of G. For positive constants δ, β and C, for $n \in \mathbb{N}$ and a degree sequence $\mathbf{d}$, we define the following subsets of graphs on a vertex set V of cardinality n:

$$\mathcal{E}_1(n,\delta,\beta,\mathbf{d}) = \{G : d(G) = \mathbf{d} \ \&\ \exists U : |U| \le \beta n, e_{U,U^c}(G) < (1-\delta)|U| c \log n\}, \quad (17)$$

$$\mathcal{E}_2(n,\beta,C,\mathbf{d}) = \{G : d(G) = \mathbf{d} \ \&\ \exists U : \beta n < |U| \le n/2, e_{U,U^c}(G) < Cn\}. \quad (18)$$

We also define

$$\mathcal{E}_1(n,\delta,\beta) = \bigcup_{\mathbf{d}} \mathcal{E}_1(n,\delta,\beta,\mathbf{d}), \quad \mathcal{E}_2(n,\beta,C) = \bigcup_{\mathbf{d}} \mathcal{E}_2(n,\beta,C,\mathbf{d}). \quad (19)$$

We shall derive bounds on the probabilities of these sets using the configuration model [2] and adapting the techniques of [3]. To this end, we define the analogous sets of configurations $\hat{\mathcal{E}}_1(n,\delta,\beta,\mathbf{d})$, $\hat{\mathcal{E}}_2(n,\beta,C,\mathbf{d})$, $\hat{\mathcal{E}}_1(n,\delta,\beta)$ and $\hat{\mathcal{E}}_2(n,\beta,C)$. Recall that configurations correspond to multigraphs, i.e, there may be loops or multiple edges. A multiple edge is counted the corresponding number of times in the above definitions.

Lemma 3 *Given $\delta > 0$, we can choose $\beta > 0$ such that*

$$\lim_{n\to\infty} \mu_n(\mathcal{E}_1(n,\delta,\beta)) = 0,$$

where the distribution μ_n was defined in (1).

Proof: Denote $|U|$ by u. Suppose first that $u \le 2\epsilon c \log n$, for a given $\epsilon > 0$. The number of edges incident within U can be at most $\binom{u}{2}$, so $e_U(G) \le \epsilon u c \log n$ for all U. But

$$2e_U(G) + e_{U,U^c}(G) = D_U := \sum_{i\in U} d_i.$$

Since $D_U \sim cu \log n$ for degree sequences $\mathbf{d} \in A_1(\alpha_1,\alpha_2)$ and any $U \subseteq V$, it follows from the above that, if $\delta > 2\epsilon$ and n is sufficiently large, then $e_{U,U^c}(G) \ge (1-\delta)uc\log n$ whenever $u \le 2\epsilon c \log n$.

Next, fix a degree sequence $\mathbf{d} \in A_1(\alpha_1,\alpha_2)$. By (1), all graphs with the same degree sequence are equally likely under the distribution μ_n, so we can use the configuration model to generate a random graph with this distribution, conditional on the degree sequence. Recall that each graph with a given degree sequence corresponds to the same number of configurations, namely $\prod_{i=1}^n d_i!$, but a configuration may not yield a simple graph (it could have loops and multiple edges). Assuming that every "bad" configuration (namely, a configuration H with $e_{U,U^c}(H) < (1-\delta)uc\log n$ for some $U \subseteq V$) corresponds to a simple graph yields an upper bound on the fraction of graphs which are bad. Using the enumeration formula of McKay and Wormald [8], this bound says that,

$$\mu_n(\mathcal{E}_1(n,\delta,\beta,\mathbf{d})|\mathbf{d}) \le e^{\lambda+\lambda^2} \mathbf{P}(H \in \hat{\mathcal{E}}_1(n,\delta,\beta,\mathbf{d})|\mathbf{d}), \quad (20)$$

where $\mathbf{P}(\cdot|\mathbf{d})$ denotes the probability with respect to the uniform distribution on configurations with degree sequence $\mathbf{d}$, and λ was defined in (4).

Fix a subset U and let $D_U = \sum_{i\in U} d_i$. The number of edges incident within U in a random configuration is bounded above by a binomial random variable X with parameters D_U and $D_U/(2E - D_U)$. Using Chernoff's bound, we can show, for any $\hat{\delta} \in (0,1)$ and any degree sequence $\mathbf{d} \in A_1(\alpha_1,\alpha_2)$, that

$$\log \mathbf{P}(X > \hat{\delta} D_U) \le -uc\log n\Big[\hat{\delta}\log\frac{\hat{\delta}(n-u)}{u} - \hat{\delta}\Big]\Big(1 + O\Big(\frac{1}{\sqrt{\log n}}\Big)\Big). \quad (21)$$

We omit the details for brevity. We can use this bound, and the fact that X stochastically dominates $e_U(H)$ (conditional on $\mathbf{d}$), to show the following, for n sufficiently large:

$$\mathbf{P}(\exists U \subseteq V : e_U(H) > \hat{\delta} D_U \text{ and } 2\epsilon c \log n \leq u \leq \sqrt{n}) \leq \exp\Big(-\frac{\epsilon \hat{\delta} c^2}{2} \log^3 n\Big), \tag{22}$$

and, for small enough β,

$$\mathbf{P}(\exists U \subseteq V : e_U(H) > \hat{\delta} D_U \text{ and } \sqrt{n} \leq u \leq \beta n) \leq e^{-\sqrt{n}}. \tag{23}$$

Now, for degree sequences $\mathbf{d} \in A_1(\alpha_1, \alpha_2)$, $D_U \sim uc \log n$ for all $U \subseteq V$, and so $e_U(H) > u\delta c \log n$ implies that $e_U(H) > \delta' D_U$ for any $\delta' < \delta$, if n is sufficiently large. Hence, we have from (20), (22) and (23) that

$$\mu_n(\mathcal{E}_1(n, \delta, \beta, \mathbf{d}) | \mathbf{d}) \leq e^{\lambda + \lambda^2} (e^{-k \log^3 n} + e^{-\sqrt{n}}),$$

for every $\mathbf{d} \in A_1(\alpha_1, \alpha_2)$. But $\lambda = O(\log n)$, so $\mu_n(\mathcal{E}_1(n, \delta, \beta, \mathbf{d}) | \mathbf{d} \in A_1(\alpha_1, \alpha_2))$ goes to 0 as $n \to \infty$. By Theorem 1, $\mu_n(\mathbf{d} \notin A_1(\alpha_1, \alpha_2))$ goes to 0 as well. Since

$$\mu_n(\mathcal{E}_1(n, \delta, \beta)) \leq \mu_n(\mathcal{E}_1(n, \delta, \beta, \mathbf{d}) | \mathbf{d} \in A_1(\alpha_1, \alpha_2)) + \mu_n(\mathbf{d} \notin A_1(\alpha_1, \alpha_2)),$$

the claim of the lemma is established. □

Lemma 4 *For arbitrarily small $\beta > 0$ and arbitrarily large $C > 0$*

$$\lim_{n \to \infty} \mu_n(\mathcal{E}_2(n, \beta, C)) = 0.$$

Proof: As in the proof of Lemma 3, we fix a degree sequence $\mathbf{d} \in A_1(\alpha_1, \alpha_2)$ and a subset U, and bound the probability that $e_{U,U^c}(G) < Cn$ in terms of the probability that $e_{U,U^c}(H) < Cn$, where H is drawn uniformly at random from configurations with degree sequence $\mathbf{d}$. Analogous to (20), we have

$$\mu_n(\mathcal{E}_2(n, \beta, C, \mathbf{d}) | \mathbf{d}) \leq e^{\lambda + \lambda^2} \mathbf{P}(H \in \hat{\mathcal{E}}_2(n, \beta, C, \mathbf{d}) | \mathbf{d}). \tag{24}$$

Fix positive constants β and C. Let U be a subset of the vertex set with $\beta n < u \leq n/2$, and let $j < Cn$. Recall that the number of configurations with degree sequence $\mathbf{d}$ is

$$H_n(\mathbf{d}) = \frac{(2E)!}{E! 2^E} \prod_{i=1}^{n} d_i!, \tag{25}$$

where $E = \sum_{i=1}^n d_i/2$ is the total number of edges. The number of these configurations with exactly j edges crossing the cut between U and U^c is

$$\begin{aligned} \mathcal{H}_{U,U^c}(j) \leq & \binom{D_U}{j} \binom{2E - D_U}{j} j! \times \\ & \frac{(D_U - j)!}{(\frac{D_U - j}{2})! \, 2^{\frac{D_U - j}{2}}} \frac{(2E - D_U - j)!}{(E - \frac{D_U - j}{2})! \, 2^{\frac{2E - D_U - j}{2}}} \prod_{i=1}^{n} d_i!, \end{aligned} \tag{26}$$

where $D_U = \sum_{i \in U} d_i$. The dependence of $\mathcal{H}$ on $\mathbf{d}$ has been suppressed for notational convenience. The first two terms on the right above count the number of ways we can choose j configurations points each from U and U^c to match up. The term $j!$ counts the number of ways of matching them. The remaining configuration points have to be matched within the sets U and U^c as there are only j edges crossing the cut. The number of ways of doing this is the number of configurations on U with $D_U - j$ points, times the number of configurations on U^c with $2E - D_U - j$ points, and with a degree sequence strictly bounded by $\mathbf{d}$ (since j points each in U and U^c have been used up). This yields the remaining terms in the bound above. We obtain from (25) and (26) after some simplification that

$$\mathbf{P}(e_{U,U^c}(H) = j) = \frac{\mathcal{H}_{U,U^c}(j)}{H_n(\mathbf{d})} \le \frac{\binom{E}{D_U/2}\binom{D_U/2}{j/2}\binom{E-(D_U/2)}{j/2}}{\binom{2E}{D_U}\binom{j}{j/2}} 2^j.$$

Using Stirling's formula, it can be shown that, for n sufficiently large,

$$\log \mathbf{P}(e_{U,U^c}(H) = j) \le -\frac{cn \log n}{4} h\Big(\frac{\beta}{2}\Big), \tag{27}$$

where, for $x \in [0,1]$, $h(x) = -x\log x - (1-x)\log(1-x)$ denotes the binary entropy of x. This bound applies to arbitrary $U \subseteq V$ with $\beta n < u < n/2$, and it doesn't depend on u. The number of subsets U with cardinality between βn and $n/2$ is smaller than the total number of subsets, which is 2^n. Hence, by the union bound,

$$\mathbf{P}(H(\mathbf{d}) : \exists U \text{ with } \beta n < u < n/2 \text{ and } e_{U,U^c}(H) = j) \le 2^n e^{-\frac{cn\log n}{4} h(\frac{\beta}{2})}.$$

The above holds for each $j < Cn$. Applying the union bound once more,

$$\mathbf{P}(H \in \hat{\mathcal{E}}_2(n,\beta,C,\mathbf{d})|\mathbf{d}) \le Cn2^n e^{-\frac{cn\log n}{4} h(\frac{\beta}{2})}, \text{ for all } \mathbf{d} \in A_1(\alpha_1,\alpha_2).$$

Substituting this in (24) and noting that $\lambda = O(\log n)$, we see that

$$\mu_n(\mathcal{E}_2(n,\beta,C,\mathbf{d})|\mathbf{d}) \to 0 \text{ as } n \to \infty, \text{ for all } \mathbf{d} \in A_1(\alpha_1,\alpha_2).$$

We also know from Theorem 1 that $\mu_n(\mathbf{d} \notin A_1(\alpha_1,\alpha_2))$ goes to zero. But,

$$\mu_n(\mathcal{E}_2(n,\beta,C)) \le \mu_n(\mathcal{E}_2(n,\beta,C,\mathbf{d})|\mathbf{d} \in A_1(\alpha_1,\alpha_2)) + \mu_n(\mathbf{d} \notin A_1(\alpha_1,\alpha_2)),$$

and so $\mu_n(\mathcal{E}_2(n,\beta,C)) \to 0$ as $n \to \infty$, as claimed. □

Finally, we shall use the results above to show that random graphs drawn from the distribution μ_n can tolerate link failure rates up to $e^{-1/c}$ without losing connectivity. Fix $p < e^{-1/c}$ and assume that links fail independently with probability p each. For a subset U of the vertex set, let $\hat{e}_{U,U^c}$ denote the number of edges between U and U^c that have not failed. We shall show that, with high probability, $\hat{e}_{U,U^c} > 0$ for all subsets U, i.e., the graph is connected. Now,

$$\mathbf{P}(\hat{e}_{U,U^c}(G) = 0|e_{U,U^c}(G)) = p^{e_{U,U^c}(G)}.$$

Suppose first that $G \notin \mathcal{E}_1(n,\delta,\beta)$. Then, $e_{U,U^c}(G) \ge (1-\delta)uc\log n$ for all $U \subseteq V$ with $|U| < \beta n$. Hence,

$$\mathbf{P}(\exists U : |U| < \beta n,\ \hat{e}_{U,U^c}(G) = 0) \le \sum_{u=1}^{\beta n} \binom{n}{u} p^{(1-\delta)uc\log n}.$$

Since $p < e^{-1/c}$, given we can choose $\delta > 0$ so that $p^{(1-\delta)c} < e^{-(1+\epsilon)}$, for some $\epsilon > 0$. Therefore, using the inequality $\binom{n}{u} \leq n^u/u!$, we get

$$\mathbf{P}(\exists U : |U| < \beta n,\, e_{U,U^c}(G) = 0) \leq \sum_{u=1}^{\beta n} \frac{1}{u!} e^{-\epsilon u \log n} \leq \exp(e^{-\epsilon \log n}) - 1, \qquad (28)$$

which goes to zero as $n \to \infty$.

Suppose next that $G \notin \mathcal{E}_2(n, \beta, C)$. Then, $e_{U,U^c}(G) \geq Cn$ for all $U \subseteq V$ with $\beta n < |U| \leq n/2$. Hence,

$$\mathbf{P}(\exists U : \beta n < |U| \leq n/2,\, \hat{e}_{U,U^c}(G) = 0) \leq \sum_{|U|:\beta n < u \leq n/2} p^{Cn} \leq 2^n p^{Cn}. \qquad (29)$$

Since C can be chosen arbitrarily large, the above quantity goes to zero as $n \to \infty$.

We see from (28) and (29) that, for suitably chosen δ and β,

$$\mathbf{P}(\exists U : \hat{e}_{U,U^c}(G) = 0 | G \notin \{\mathcal{E}_1(n, \delta, \beta) \cup \mathcal{E}_2(n, \beta, C)\}) \to 0 \text{ as } n \to \infty.$$

Also, by Lemmas 3 and 4,

$$\mathbf{P}(G \in \{\mathcal{E}_1(n, \delta, \beta) \cup \mathcal{E}_2(n, \beta, C)\}) \to 0 \text{ as } n \to \infty,$$

when G is chosen according to the distribution μ_n. Hence, for any $p < e^{-1/c}$, a graph G chosen at random from the distribution μ_n, and subjected to independent link failures with probability p, remains connected, with probability going to 1 as $n \to \infty$.

References

[1] F. Ball and A. Barbour, "Poisson approximation for some epidemic models", *J. Appl. Prob.*, vol. 27, pp. 479–490, 1990.

[2] Béla Bollobás, *Random graphs*, Cambridge University Press, 2001.

[3] Colin Cooper, Alan Frieze and Bruce Reed, "Random-regular graphs of non-constant degree: connectivity and Hamiltonicity", *Combin., Probab. and Comput.*, vol. 11, pp. 249–261, 2002.

[4] P. Erdös and A. Renyi, "On the evolution of random graphs", *Mat Kutato Int. Közl*, vol. 5, pp. 17–60, 1960.

[5] A. Ganesh, A.-M. Kermarrec, and L. Massoulié, "Peer-to-peer membership management for gossip-based protocols", *IEEE Trans. Comp.*, vol. 52, pp. 139–149, 2003.

[6] L. Massoulié, A.-M. Kermarrec and A. J. Ganesh, "Network awareness and failure resilience in self-organising overlay networks", in *Proc. Symp. Reliable Distributed Systems (SRDS)*, 2003.

[7] David McDonald, "A local limit theorem for large deviations of sums of independent, nonidentically distributed random variables", *Ann. Probab.*, vol. 7, pp. 526–531, 1979.

[8] B.D. McKay and N.C. Wormald, "Asymptotic enumeration by degree sequence of graphs of high degree", *Europ. J. Combinatorics*, vol. 11, pp. 565–580, 1990.