Analysis of Edge Deletion Processes on Faulty Random Regular Graphs

Andreas Goerdt*[1] and Mike Molloy[2]

[1] Fakultät für Informatik, TU Chemnitz, 09111 Chemnitz, Germany
goerdt@informatik.tu-chemnitz.de
[2] Department of Computer Science, University of Toronto, Toronto, Canada
molloy@cs.toronto.edu

Abstract. Random regular graphs are, at least theoretically, popular communication networks. The reason for this is that they combine low (that is constant) degree with good expansion properties crucial for efficient communication and load balancing. When any kind of communication network gets large one is faced with the question of fault tolerance of this network. Here we consider the question: Are the expansion properties of random regular graphs preserved when each edge gets faulty independently with a given fault probability? We improve previous results on this problem: Expansion properties are shown to be preserved for much higher fault probabilities and lower degrees than was known before. Our proofs are much simpler than related proofs in this area.

Introduction

A natural question in the theory of fault tolerance of communication networks reads: Is it possible to simulate the non-faulty network on the faulty one with a well determined slowdown? Here one assumes that the network proceeds in synchronous steps and in each step each processor (= node of the network) performs some local computation and some communication steps. Ideally one would like to simulate the non-faulty network in such a way that the simulation is slower only by a *constant* factor showing that the time is essentially unchanged. Whereas such efficient simulations are known for networks with unbounded degree, like the hypercube, it is still an important question whether they exist for bounded degree networks like the butterfly [CoMaSi 95]. Note that all of this paper refers to *random* faults, that is each component (normally edge or node) gets faulty independently with a given fault probability and the results only hold with high probability meaning with probability going to 1 when the network gets large.

Random regular graphs with given degree $d \geq 3$ are well known to be expander graphs (with high probability) [Bo 88]: There is a constant $C\,(< 1)$ such that each subset X of nodes has $\geq C \cdot |X|$ neighbours adjacent to X but not belonging to X (provided X contains at most half of all vertices). If we ever were

* Author's work in part performed at the University of Toronto, supported by a grant obtained through Alasdair Urquhart.

to simulate computation on a random regular graph with slowdown only by a constant factor on the faulty graph we would need a linear size expander inside the faulty graph.

The investigation of random regular graphs with edge faults starts with the paper [Sp et al. 94]. In the succeeding paper [NiSp 95] attention is drawn to the preservation of expansion properties. Some sufficient conditions are given. In work by the first author [Go 97] a threshold result for the existence of a linear size component is proved. In [Go 98] we give a sufficient condition on fault probability and degree such that we can *find* a linear size expander efficiently – a question not treated in the initial work on expansion [NiSp 95]. Crucial to our result is the notion of a $k-$core: The $k-$core of a given graph is the (unique) maximal subgraph where each node has degree at least k. In [Go 98] we first observe that the $3-$core of a faulty random regular graph is an expander (this follows simply from randomness properties of the $3-$core). Second, we present a simple edge deletion algorithm which is shown to find a $3-$core of linear size when $d \geq 42$ and each edge is non-faulty with probability at least $20/d$.

The present paper improves considerably on these results: We give a precise threshold on the fault probability for the existence of a linear size $k-$core for any $d > k \geq 3$. Thus improving the previous bounds for the existence of an expanding subgraph. For example when the degree is as low as 4 and each edge is faulty with probability $< 1/9$ we have a linear size $3-$core and thus an expanding subgraph.

Our proof uses a proof technique originally developed for [MoPi]. It is technically quite simple. This is in sharp constrast to the previous proofs of the weaker results mentioned above relying on technically advanced probability theoretic tools. This technique applies to a wide range of similar problems (see [MoWo]). The technique was inspired by the original (more involved) proof of the k-core threshold for $G_{n,p}$ given in [PiSpWo 96].

1 Outline

We will study random regular graphs with edge faults by focussing on the configuration model (cf.[Bo 85]). It is well known that properties which hold a.s. (almost surely) for a uniformly random d-regular configuration also hold a.s. for a uniformly random d-regular simple graph. For the configuration model, we consider n disjoint $d-$element sets called *classes*; the elements of these classes are called *copies*. A configuration is a partition of the set of all copies into $2-$element sets, which are edges. Identifying classes with vertices, configurations determine multigraphs and standard graph theoretic terminology can be applied to configurations. More details can be found in [Bo 85]. We fix the degree d and the probability p for the rest of this paper and consider probability spaces $\mathrm{Con}(n, d, p)$ of random configurations where each edge is present with probability p or absent with fault probability $f = 1 - p$. We call this space the space of *faulty configurations*. An element of this space of is best considered as being generated by the following probabilistic experiment consisting of two

stages:
(1) Draw randomly a configuration $\Phi = (\mathcal{W}, E)$ where $\mathcal{W} = W_1 \dot{\cup} \ldots \dot{\cup} W_n$ and $|W_i| = d$. (2) Delete each edge of Φ, along with its end-copies, independently with fault probability f.

The probability of a fixed faulty configuration with k edges is $(n \cdot d - 2 \cdot k)!! \cdot (1-p)^{(n \cdot d/2) - k} \cdot p^k$. Given k, each set of k edges is equally likely to occur. The degree of a class W with respect to a faulty configuration Φ, $\mathrm{Deg}_\Phi(W)$, is the number of copies of W which were not deleted. Note that edges $\{x, y\}$ with $x, y \in W$ contribute with two to the degree. The degree of a copy x, $\mathrm{Deg}_\Phi(x)$, is the degree of the class to which x belongs. The k-core of a faulty configuration is the maximal subconfiguration of the faulty configuration in which each class has a degree $\geq k$. We call classes of degree less than k *light* whereas classes of degree at least k are *heavy*. By $Bin(m, \lambda)$ we denote the binomial distribution with parameters m and success probability λ.

We now give an overview of the proof of the following theorem which is the main result of this paper. For $d > k \geq 3$ we consider the real valued function $L(\lambda) = \lambda / Pr[Bin(d-1, \lambda) \geq k-1]$ which we define for $0 < \lambda \leq 1$. $L(1) = 1$ and $L(\lambda)$ goes to infinity for λ approaching 0. Moreover $L(\lambda)$ has a unique minimum for $1 \geq \lambda > 0$. Let $r(k, d) = \min\{L(\lambda) | 1 \geq lambda > 0\}$. For example we have that $r(3, 4) = 8/9$. The definition of $r(k, d)$ is, no doubt, mysterious at this point, but we will see that it has a very natural motivation.

Theorem 1. *(a) If $p > r(k,d)$ then a random $\Phi \in Con(n, d, p)$ has a k-core of linear size with high probability.*
(b) If $p < r(k,d)$ then a random $\Phi \in Con(n, d, p)$ has only the empty k-core with high probability.

Theorem 1 implies that the analogous result holds for the space of faulty random regular graphs (obtained as: first draw a graph, second delete the faulty edges). The following algorithm which can easily be executed in the faulty network itself is at the heart of our argument.

Algorithm 2. The Global Algorithm
Input: A faulty configuration Φ, output: The k-core of Φ.

while Φ has light classes do
  $\Phi :=$ the modification of Φ where all light classes are deleted.
od. Output Φ.

Specifically, when we delete a class, W, we delete (i) all copies within W, (ii) all copies of other classes which are paired with copies of W, (iii) W itself. Note that it is possible for W itself to still be undeleted but to contain no copies as they were all deleted as a result of neighbouring classes being deleted, or faulty edges. In this case, of course, W is light and so it will be deleted on the next iteration. At the end of the algorithm Φ has only classes of degree $\geq k$, which form the k-core. The following notion will be used later on: A class W of the faulty configuration Φ *survives* j $(j \geq 0)$ rounds of the global algorithm with degree t iff W has not yet been deleted and has degree t after the j'th execution of the while-loop of the algorithm with input Φ. A class simply survives if it has not yet been deleted.

In section 2 we analyze this algorithm when run for $j-1$ executions of the loop where we set $j = j(n) = \sqrt{\log_d n}$ throughout. We prove that the number of classes surviving $j-1$ rounds with degree $t \geq k$ is linear in n with high probability when $p > r(k,d)$ whereas the number of light classes is $o(n)$. (Initially this number is linear in n.) An extra argument presented in section 4 will show how to get rid of these few light classes leaving us with a linear size $k-$core provided $p > r(k,d)$. On the other hand, if $p < r(k,d)$ then we show that the expected number of classes surviving $j-1$ rounds with any degree is $o(n)$ and that we have no longer enough classes to form a $k-$core. This is shown in section 3.

2 Reduction of the number of light classes

For $d \geq t \geq 0$, and for a particular integer j, we let

$$X_t : \mathrm{Con}(n, d, p) \to \mathcal{N} \tag{1}$$

be the number of classes surviving $j-1$ rounds of the global algorithm with degree equal to t. As usual we can represent $X = X_t$ as a sum of indicator random variables

$$X = X_{W_1} + \cdots + X_{W_n}, \tag{2}$$

where X_W assumes the value 1 when the class W survives $j-1$ rounds with degree equal to t and 0 when this is not the case. Then $EX = n \cdot E[X_W] =$ $n \cdot Pr[W \text{ survives with degree } t]$ for W arbitrary. We determine $Pr[W \text{ survives with degree } t]$ approximately, that is an interval of width $o(1)$ which includes the probability. The probability of the event: W survives $j-1$ rounds with degree t, turns out to depend only on the $j-$environment of W defined as: For a class W the $j-$environment of W, $j - \mathrm{Env}_\Phi(W)$, is that subconfiguration of Φ which has as classes the classes whose distance from W is at most j. Here distance means the number of edges in a shortest path. The edges of $j - \mathrm{Env}_\Phi(W)$ are those induced from Φ.

The proof of the following lemma follows with standard conditioning techniques observing that the $j-$environment of a class W in a random configuration can be generated by a natural probabilistic breadth first generation process (cf. [Go 97] for details on this.) Here it is important that j only slowly goes to infinity.

Lemma 3. *Let W be a fixed class then $Pr\{j - Env_\Phi(W)$ is a tree$\} \geq 1 - o(1)$.*

Note that the lemma does not mean: Almost always the j-environment of *all* classes is a tree. The definition of j-environment extends to faulty configurations in the obvious manner. Focussing on a j-environment which is a tree is very convenient since in a faulty configuration, it can be thought of as a branching process whereby the number of children of the root is distributed as $Bin(d,p)$, and the number of children of each non-root as $Bin(d-1,p)$.

The following algorithm approximates the effect the global algorithm has on a fixed class W, provided the $j-$environment of W is a tree.

Algorithm 4. The Local Algorithm.
Input: A (sub-)configuration Γ, which is a $j-$environment of a class W in a faulty configuration. Γ is a tree with root W.
$\Phi := \Gamma$
for $i = j-1$ downto 0 do
  Modify Φ as follows: Delete all light classes in depth i of the tree Φ.
od.
The output is "W survives with degree t" if W is not deleted and has final degree t. If W is deleted then the output is "W does not survive".

Note that it is not possible for W to survive with degree less than k. By round l of the algorithm we mean an execution of the loop with $i = j-l$ where $1 \leq l \leq j$. A class in depth i where $0 \leq i \leq j$ survives with degree t iff it is not deleted and has degree t after round $j - i$ of the algorithm. Note that classes in depth j are never deleted and so they are considered to always survive. The next lemma states in which respect the local algorithm approximates the global one. The straightforward formal proof is omitted in this abridged version.

Lemma 5. *Let $j \geq 1$. For each class W and each faulty configuration Φ where $j - Env_\Phi(W)$ is a tree we have: After $j-1$ rounds of the global algorithm with Φ the class W survives with degree $t \geq k$. $\Leftrightarrow$ After running the local algorithm with $j - Env_\Phi(W)$ the class W survives with degree $t \geq k$.*

Note that W either survives $j-1$ rounds of the global algorithm and the whole local algorithm with the same degree $t \geq k$ or does not survive the local algorithm in which case it does or does not survive $j-1$ global rounds, but does certainly not survive j global rounds.

We condition the following considerations on the almost sure event that for $j = j(n)$ the $j-$environment of the class W in the underlying fault free configuration is a tree (cf. Lemma 3). We denote this environment in a random faulty configuration by Γ. We turn our attention to the calculation of the survival probability with the *local* algorithm.

For i with $0 \leq i \leq j-1$ let ϕ_i be the probability that a class in level (=depth) $j-i$ of Γ survives the local algorithm applied to Γ. As the j-enviroment in the underlying fault-free configuration is a tree, the survival events of the children of given class are independent. Therefore:

$$\phi_0 = 1 \text{ and } \phi_i = Pr[Bin(d-1, p \cdot \phi_{i-1}) \geq k-1]. \tag{3}$$

And furthermore, considering now the root W of the j-environment, we get for $t \geq k$ by analogous considerations:
$Pr[W \text{ survives the local algorithm with degree } t.] = Pr[Bin(d, p \cdot \phi_{j-1}) = t]$.

We have that the sequence of the ϕ_i's is monotonically decreasing and in the interval $[0, 1]$. Hence $\phi = \phi(p) = \lim_{i\to\infty} \phi_i$ is well defined and as all functions involved are continuous we get: $\phi = Pr[Bin(d-1, p\cdot\phi) \geq k-1]$. (Note that this is no definition of ϕ, the equation is always satisfied by $\phi = 0$.)

Two further notations for subsequent usage: $\lambda_{t,i} = \lambda_{t,i}(p) = Pr[Bin(d, p\cdot \phi_{i-1}) = t]$ for $i \geq 1$. Again we have that the $\lambda_{t,i}$'s are monotonically decreasing and between 0 and 1 and $\lambda_t = \lambda_t(p) = \lim_{i\to\infty}\lambda_{t,i}$. exists. Hence for our fixed class W, considering $j \to \infty$, we get:

$$Pr[W \text{ survives the local algorithm with degree } t.] = \lambda_{t,j} = \lambda_t + o(1). \quad (4)$$

Here is where our formula for $r(k, d)$ comes from:

Lemma 6. *$\phi > 0$ iff $p > r(k, d)$.*

Proof. First let $\phi > 0$. As stated above we have $\phi = Pr[Bin(d-1, p\phi) \geq k-1]$. Therefore $Pr[Bin(d-1, p\phi) \geq k-1] > 0$ and setting $\lambda = p\cdot\phi$, we get $\lambda/p = Pr[Bin(d-1, \lambda) \geq k-1]$ and so $p = \lambda/Pr(Bin(d-1, \lambda) \geq k-1) = L(\lambda)$ and the result follows.

Now let $p > r(k, d)$. Let λ_0 be such that $r(k, d) = L(\lambda_0)$. We show by induction on i that $p\cdot\phi_i \geq \lambda_0$. For the induction base we get: $p\cdot\phi_0 = p > r(k, d) \geq \lambda_0$ where the last estimate holds because the denominator in the definition of $L(\lambda_0)$ always is ≤ 1. For the induction step we get:
$p\cdot\phi_{i+1} = p\cdot Pr[Bin(d-1, p\cdot\phi_i) \geq k-1] \geq p\cdot Pr[Bin(d-1, \lambda_0) \geq k-1] > \lambda_0$ where the last but one estimate uses the induction hypothesis and the last one follows from the assumption. □

We now return to the analysis of the *global* algorithm. The next corollary follows directly with Lemma 3, Lemma 5, and (4).

Corollary 7. *Let W be a fixed class, $t \geq k$ and let $j = j(n) = \sqrt{\log_d n}$. In the space of faulty configurations we have (cf.(2)):*
$Pr[X_W = 1] = Pr\{W$ survives $j(n)-1$ global rounds with degree $t\}$
$= \lambda_t + o(1)$.

Next the announced concentration result:

Theorem 8. *Let $t \geq k$, $X = X_t$ be the random variable defined as in (1), and let $\lambda = \lambda_t$; then we have:*
(1) $EX = \lambda\cdot n + o(n)$. *(2) Almost surely $|X - \lambda\cdot n| \leq o(n)$.*

Proof. (1) The claim follows from the representation of X as a sum of indicator random variables (cf. (2)) and with Corollary 7.
(2) We show that $VX = o(n^2)$. This implies the claim with an application of Tschebycheff's inequality. We have $X = X_{W_1} + X_{W_2} + \ldots + X_{W_n}$ (cf. (2)). This and (1) of the present theorem implies $VX = E[X^2] - (EX)^2 =$

$E[X^2] - (\lambda^2 \cdot n^2 + o(n) \cdot n)$. Moreover, $E[X^2] = EX + n \cdot (n-1) \cdot E[X_U \cdot X_W] = \lambda \cdot n + o(n) + n \cdot (n-1) \cdot E[X_U \cdot X_W]$, where U and W are two arbitrary distinct classes. We need to show that $E[X^2] = \lambda^2 \cdot n^2 + o(n^2)$. This follows from $E[X_U \cdot X_W] = \lambda^2 + o(1)$ showing that the events $X_U = 1$ and $X_W = 1$ are asymptotically independent. This follows by conditioning on the event that the $j-$environments of U and W are disjoint trees and analyzing the breadth first generation procedure for the $j-$ environment of a given class. Again we need that j goes only slowly to infinity. □

3 When there is no k-core

The proof of Theorem 1(b) is now quite simple. First we need the following fact:

Lemma 9. *A.s. a random member of Con(n, d, p) has no k-core on $o(n)$ vertices.*

Proof. The lemma follows from the fact that a random member of Con(n, d) a.s. has no subconfiguration with average degree at least 3 on at most ϵn vertices, where $\epsilon = \epsilon(d)$ is a small positive constant. Consider any $s \leq \epsilon n$. The number of choices for s classes, $1.5s$ edges from amongst those classes, and copies for the endpoint of each edge, is at most:

$$\binom{n}{s}\binom{\binom{s}{2}}{1.5s} d^{3s}.$$

Setting $M(t) = t!/(2^{t/2}(t/2)!)$ to be the number of ways of pairing t copies, we have that for any such collection, the probability that those pairs lie in our random member of Con(n, d) is

$$M((d-3s)n)/M(dn) < (\frac{e}{n})^{1.5s}.$$

Therefore, the expected number of such subconfigurations is at most:

$$\begin{aligned}\binom{n}{s}\binom{\binom{s}{2}}{1.5s} d^{3s}(\frac{e}{n})^{1.5s} &< (\frac{en}{s})^s(\frac{e(s^2/2)}{1.5s})^{1.5s}(\frac{e}{n})^{1.5s} \\ &\leq (\frac{20d^6 s}{n})^{.5s} = f(s).\end{aligned}$$

Therefore, if $\epsilon = 1/40d^6$ then the expected number of such subconfigurations is less than $\sum_{s=2}^{n} f(s)$ which is easily verified to be $o(1)$. □

Now, by Lemma 6, we have for $p < r(k, d)$ that $\phi = 0$. Therefore, as j goes to infinity the expected number of classes surviving j rounds with degree at least k is $o(n)$ and so almost surely is $o(n)$. With the last lemma we get Theorem 1 (b).

4 When there is a $k-$core

In this section, we prove Theorem 1(a). So we assume that $p > r(k, d)$. We start by showing that almost surely very few light clauses survive the first $j(n) - 1$ iterations:

Lemma 10. *In Con(n, d, p) almost surely: The number of light classes after $j(n) - 1 = \sqrt{\log_d n} - 1$ rounds of the global algorithm is reduced to $o(n)$.*

Proof. The proof follows with Theorem 8 applied to $j - 2$ and $j - 1$ (which both go to infinity). □

In order to eliminate the light classes still present after $j(n) - 1$ global rounds, we need to know something about the distribution of the configurations after $j(n) - 1$ rounds. As usual in similar situations the uniform distribution needs to be preserved. For $\bar{n} = (n_0, n_1, n_2, \ldots, n_d)$ where the sum of the n_i is at most n we let Con$(\bar{n})$ be the space of all configurations with n_i classes consisting of i copies. Each configuration is equally likely. The following lemma is proved in [Go 98].

Lemma 11. *Conditioning the space Con$(n\,, d\,, p)$ on those configuration which give a configuration in Con$(\bar{n})$ after i global rounds, each configuration from Con$(\bar{n})$ has the same probability to occur after i global rounds.*

After running the global algorithm for $j(n) - 1$ rounds we get by Lemma 10 a configuration uniformly distributed in Con$(\bar{n})$ where $n_1 + n_2 + ... + n_{k-1} = o(n)$ and $|n_t - \lambda_t \cdot n| \le o(n)$ for $t \ge k$ with high probability. A probabilistic analysis of the following algorithm eliminating the light classes one by one shows that we obtain a linear size $k-$core with high probability.

Algorithm 12.
Input: A faulty configuration Φ.
Output: The $k-$core of Φ.
while There exist light classes in Φ do
  Choose uniformly at random a light class W from all light classes
  and delete W and the edges incident with W.
od. The classes of degree $\ge k$ are the $k-$core of Φ.

In order to perform a probabilistic analysis of this algorithm it is again important that the uniform distribution is preserved. A similar result is Proposition 1 in [PiSpWo 96] (for the case of graphs instead of configurations).

Lemma 13. *If we apply the algorithm above to a uniformly random $\Phi \in$ Con$(\bar{n})$, ($\bar{n}$ fixed) for a given number of iterations we get: Conditional on the event (in Con$(\bar{n})$) that the configuration obtained, Ψ, is in Con$(n'_0, n'_1\, n'_2,\ n'_3, \ldots, n'_d)$ the configuration Ψ is a uniformly random configuration from this space.*

Lemma 14. *We consider probability spaces $Con(\bar{n})$ where the number of heavy vertices is $\geq \delta \cdot n$. In one round of Algorithm 12 one light class disappears and we get $\leq k-1$ new light classes. Let $Y : Con(\bar{n}) \to \mathcal{N}$ be the number of new light classes after one round of Algorithm 12. Let $\nu = \sum_i i \cdot n_i$ and $\pi = (k \cdot n_k)/\nu$. Thus π is the probability to pick a copy of degree k when picking uniformly at random from all copies belonging to edges. Then:*

(a) $Pr[Y = l] = Pr[Bin(deg(W), \pi) = l] + o(1)$.

(b) $EY \leq (k-1) \cdot \pi + o(1)$.

The straightforward proof of this lemma is omitted due to lack of space. Our next step is to bound π.

Lemma 15. *$\pi \leq (1-\epsilon)/(k-1)$ for some $\epsilon > 0$.*

Proof. We will prove that when $p = r(k,d)$ then $\pi = 1/(k-1)$. Since π is easily shown to be decreasing in p, this proves our lemma. Recall that $r(k,d)$ is defined to be the minimum of the function $L(\lambda)$. Therefore, at $L(\lambda) = r(k,d)$, we have $L^{'}(\lambda) = 0$. Differentiating L, we get:

$$\sum_{i=k-1}^{d-1} \binom{d-1}{i} \lambda^i (1-\lambda)^{d-1-i} = \sum_{i=k-1}^{d-1} \binom{d-1}{i} \lambda^i (1-\lambda)^{d-2-i} (i-(d-1)\lambda). \quad (5)$$

A simple inductive proof shows that the RHS of (5) is equal to

$$(k-1)\binom{d-1}{k-1} \lambda^{k-1}(1-\lambda)^{d-k}. \quad (6)$$

Indeed, it is trivially true for $k = d$, and if it is true for $k = r+1$ then for $k = r$ the RHS is equal to

$$\binom{d-1}{r-1}\lambda^{r-1}(1-\lambda)^{d-1-r}(r-1-(d-1)\lambda) + r\binom{d-1}{r}\lambda^r(1-\lambda)^{d-1-r}$$
$$= (r-1)\binom{d-1}{r-1}\lambda^{r-1}(1-\lambda)^{d-r}$$

Setting $j = i+1$, and multiplying by λd, the LHS of (5) comes to:

$$\sum_{j=k}^{d} d\binom{d-1}{j-1}\lambda^j(1-\lambda)^{d-j} = \sum_{j=k}^{d} j\binom{d}{j}\lambda^j(1-\lambda)^{d-j},$$

and (6) comes to

$$d(k-1)\binom{d-1}{k-1}\lambda^k(1-\lambda)^{d-k} = k(k-1)\binom{d}{k}\lambda^{k-1}(1-\lambda)^{d-k}.$$

Now, since

$$\pi = \frac{k\binom{d}{k}\lambda^{k-1}(1-\lambda)^{d-k}}{\sum_{j=k}^{d} j\binom{d}{j}\lambda^j(1-\lambda)^{d-j}} + o(1),$$

this establishes our lemma. □

Lemma 16. *Algorithm 12 stops after $o(n)$ rounds of the while loop with a linear size $k-$core with high probability (with respect to $Con(\bar{n})$).*

Proof. We define Y_i to be the number of light classes remaining after i steps of Algorithm 12. By assumption, $Y_0 = o(n)$. Furthermore, by Lemmas 14 and 15, we have $EY_1 \leq Y_0 - 1 + (k-1)\pi < Y_0 - \epsilon$. Furthermore, it is not hard to verify that, since there are $\Theta(n)$ classes of degree k, then so long as $i = o(n)$ we have

$$EY_{i+1} \leq Y_i - \frac{1}{2}\epsilon,$$

and in particular, the probability that at least ℓ new light vertices are formed during step i is less than the probability that the binomial variable $Bin(k-1, \pi)$ is at least ℓ.

Therefore, for any $t = o(n)$, $Y_0, Y_1, ..., Y_t$ is statistically dominated by a random walk defined as:

$$Z_0 = Y_0; Z_{i+1} = Z_i - 1 + Bin(k-1, \frac{1 - \frac{1}{2}\epsilon}{k-1}).$$

Since Z_i has a drift of $-\frac{1}{2}\epsilon$, it is easy to verify that with high probability, $Z_t = 0$ for some $t = o(n)$, and thus with high probability $Y_t = 0$ as well.

If $Y_t = 0$ then we are left with a k-core of linear size. □

Clearly Lemma 16 implies Theorem 1(a).

References

[Bo 85] Bela Bollobas. Random Graphs. Academic Press. 1985.

[Bo 88] –.The isoperimetric number of random regular graphs. European Journal of Combinatorics. 1988, 9, 241-244.

[CoMaSi 95] Richard Cole, Bruce Maggs, Ramesh Sitaraman. Routing on Butterfly networks with random faults. In Proceedings FoCS 1995. IEEE. 558-570.

[Go 97] Andreas Goerdt. The giant component threshold for random regular graphs with edge faults. In Proceedings MFCS 1997. LNCS 1295. 279-288.

[Go 98] –. Random regular graphs with edge faults: expansion through cores. In Proceedings ISAAC 1998. LNCS 1533, 219-228.

[PiSpWo 96] Boris Pittel, Joel Spencer, Nicholas Wormald. Sudden emergence of a giant k-core in a random graph. Journal of Combinatorial Theory B 67,1996,111-151.

[MoPi] Mike Molloy and Boris Pittel. Subgraphs with average degree 3 in a random graph. In preparation.

[MoWo] Mike Molloy and Nick Wormald.In preparation.

[Sp et al. 94] S. Nikoletseas, K. Palem, P. Spirakis, M. Yung. Vertex disjoint paths and multconnectivity in random graphs: Secure network computing. In Proceedings ICALP 1994. LNCS 820. 508-519.

[NiSp 95] Paul Spirakis and S. Nikoletseas. Expansion properties of random regular graphs with edge faults. In Proceedings STACS 1995. LNCS 900. 421-432.