OPERATIONS RESEARCH CENTER

working paper

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

ASYMPTOTIC ANALYSIS OF NETWORK RELIABILITY MODELS I: CONNECTIVITY AND FLOW REQUIREMENT MEASURES IN A COMBINATORIAL ENVIRONMENT

by

Daniel Bienstock

OR 143-85 October 1985

The research reported in this paper was carried out at MIT and was partially supported by grant NSF-ECS-8316224.

Abstract

We consider the problem of designing reliable networks at low cost, and show that for several standard and non-standard models, 0-1 effects occur.

Asymptotic Analysis of Network Reliability Models I: Connectivity and Flow Requirement Measures in a Combinatorial Environment

Daniel Bienstock

1. Introduction

In this paper we will consider the design of large-scale, low hardware cost communication networks in the presence of probabilistic failures. Under such circumstances, the designer will ordinarily try to trade off reliability versus cost. We analyze several models of failures and performance requirements, and show that, asymptotically, 0-1 effects occur in the reliability vs. cost tradeoff (this statement will be made more precise below).

We will use the words "graph," "network" and "design" interchangeably. "Design" will also be used to imply a design rule or sequence of graphs, one for each possible number of nodes.

In order to approach the reliability vs. cost issues from a mathematical viewpoint, we must precisely indicate what we mean by a "model." In what follows, a reliability model will be a specification of

(i) A performance measure, that is, what we mean when we say that a network "works." The traditional network reliability literature has heavily favored connectivity measures. In a practical setting, however, links may have finite capacity, and certain flow levels (possibly stochastic) may have to be maintained. In this paper we will consider both connectivity and flow requirement measures.

(ii) A distribution of failures. In the standard literature network elements (e.g., arcs) fail independently. We consider models both with independent and dependent failures.

(iii) A <u>measure of cost</u> of a design. The simplest measure for the quantity of hardware is the number of arcs in a design. When considering flow requirements, the capacities of links should also be considered.

The main result concerning the asymptotic analysis of such models can be abbreviated as follows (each model requires a specific proof):

In a given model, let $\{G(n)\}$ be a sequence of graphs, where $G(n)$ has n nodes, hardware cost $H(n)$ and reliability $R(n)$. There exists a "threshold" function $T(n)$, such that if

(a) $H(n) \leq c\, T(n)$ where $c < 1$, then $R(n) \to 0$, independent of design.

On the other hand, we can produce designs with

(b) $H(n) = O(T(n))$ and $R(n) \to 1$.

Thus, the designs that attain (b) are essentially optimal. The proof of (a) yields as a byproduct tight upper bounds on reliability that depend only on cost, not on any specific design. Similarly, we can easily compute tight lower bounds on the reliability of the efficient designs that attain (6).

Kel'mans [5,6] analyzed the simplest case of a model. In his model we want to minimize the total number of arcs in a design, in order to keep all nodes connected while arcs are independently erased with probability $0 < p < 1$. Kel'mans showed that in this case $T(n) = \frac{1}{2} k(p)\, n \log n$, where $k(p) = -1/\log p$. The efficient networks are obtained by taking, roughly, n/m copies of K_m, (where $m = k(p) \log n$) and making copy i share one node with copy i + 1 (see Figure 1), for $1 \leq i < \frac{n}{m}$.

Figure 1 - Kel'mans design

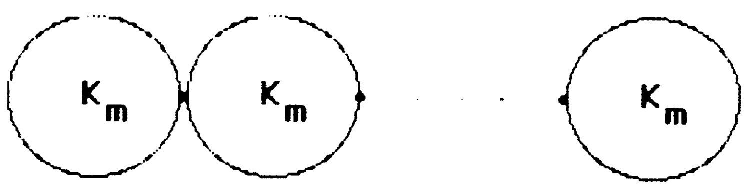

Figure 1

However, our main inspiration comes from the work of Erdös (see, for example, [4]) in the analysis of random structures in combinatorics.

We should mention that Leighton and Leiserson [7] have analyzed a set of problems which, although they do not fit in our framework, yield similar results.

2. Some Notation and Auxiliary Results

Let $\{E_n\}$ denote a sequence of events (= sets of points in a probability space). We will say that E_n holds <u>with high probability</u> if $P(E_n) \to 1$ as $n \to \infty$.

Let G(n) be the n^{th} graph in a sequence of designs. We will denote by **G(n)** the random variable obtained through the failure of elements of G(n). Let C(n) be any property satisfied w.h.p. by **G(n)**. Whenever possible, we will abbreviate this statement by saying: **G(n)** satisfies C(n).

Let {G(n)} be as above. Typically the reliability of G(n) will depend on a family F of parameters of the model (e.g., arc reliabilities). We will denote the reliability of G(n) by $R_n(F)$; and $R(F) = \lim_{n\to\infty} R_n(F)$. We will also say that {G(n)} is <u>reliable</u> if R(F) = 1.

Given two functions f(n) and g(n), we will say that f and g are log-equivalent if $\log f(n) \sim \log f(n)$. Unless otherwise stated, all logarithms are to the base e.

An LEF (logarithmic-exponential function) is a function f of the form

$$f(n) = (1-a(n))^{b(n)},$$

where $a(n) \to 0^+$ and $b(n) \to \infty$ as $n \to \infty$. It is a simple exercise to verify that, if

$$\begin{array}{lll} \text{(i)} \ \log b(n) + \log a(n) \to +\infty & & \text{(i)} \ f(n) \to 0 \\ & \text{then} & \\ \text{(ii)} \ \log b(n) + \log a(n) \to -\infty & & \text{(ii)} \ f(n) \to 1. \end{array} \qquad (1)$$

A useful tool in the asymptotic analysis of sums of random variables are the Chernoff bounds [2], here stated in abridged form. Let S_n be the sum of n independent variables distributed as

$$X = \begin{cases} 1, & \text{w.p. } Q \\ 0, & \text{w.p. } 1-Q, \end{cases}$$

and suppose $a < Q$. Then

$$P(S_n \leq an) \leq [r(Q,a)]^n, \tag{2}$$

where $$r(Q,a) = \left[\frac{Q}{a}\right]^a \left[\frac{1-Q}{1-a}\right]^{1-a} < 1. \tag{3}$$

3. A Connectivity Model with Independent Arc Failures

The following model is a generalization of Kel'mans'. The arcs of our network fail independently with probability $0 < p < 1$. We randomly select a subset S of f(n) nodes, where n = total number of nodes. We want to minimize the total number of arcs while guaranteeing that S remains connected. The only assumption concerning f(n) is that $f(n) \to \infty$ as $n \to \infty$.

Consider a sequence {G(n)} of graphs, let $R_n(p)$ = reliability of G(n) and $R(p) = \lim_{n\to\infty} R_n(p)$, and set a(n) = number of arcs of G(n).

Theorem 1

(a) Suppose $a(n) \sim \frac{k(p')}{2} n \log f(n)$, where $p'<p$. Then $R(p) = 0$.

Sketch of proof:

(i) Choose $x > 0$ so that $k(p') + x < k(p)$. Then $T = 1+(k(p')+x) \log p > 0$. Set $d(n) = (k(p')+x) \log f(n)$. We can find a set X, all of whose nodes have degree at most d(n), with $|X| = \Omega(n)$. Thus, we can find an independent set X' all of whose nodes have degree at most d(n), with $|X'| = \Omega(n/\log f(n))$ (this argument is similar to one used by Kel'mans).

(ii) In **G(n)**, at least

$$\Omega\left(\frac{n}{[f(n)]^{D} \log f(n)}\right)$$

nodes are isolated, where $D = 1-T > 0$. Let I be the set of isolated nodes in **G(n)**.

(iii)

$$P(I \cap S = \phi) \leq \frac{\binom{n - |I|}{f(n)}}{\binom{n}{f(n)}} \to 0$$

(strictly speaking, the above estimation is correct if $f(n) \leq n - |I|$. If $f(n) > n - |I|$, then certainly $I \cap S \neq \emptyset$). □

Before proceeding to give a tight upper bound, we point out that the above proof will still hold when there are parallel arcs.

(b) There exist designs with $a(n) = O(k(p)\, n \log f(n))$ that attain $R(p) = 1$.

These designs are obtained by taking a copy of K_m, where $m = [2k(p) \log f(n)]$, and connecting n-m additional nodes to all the nodes of K_m. We will refer to such a graph as an (n,m)-starfish. The copy of K_m will be called the hub and the set of remaining nodes, the periphery.

In order to show that these designs are reliable, we prove

(i) As a subgraph of **G(n)**, the hub remains connected. This is an easy consequence of the Erdös results on random graphs, since with high probability $\theta(m^2)$ arcs remain operative in the hub. Also, in [5] it is proven directly that the connectivity probability of K_N, when arcs are erased with probability p, is at least

$$1 - (N-1)^2 p^{N-1}$$

if $Np^{N-3} \leq 1$.

(ii)At most f(n) nodes of S are chosen in the periphery. With high probability, there is at least one operative arc connecting each such node to the hub. □

This concludes the proof of Theorem 1. Notice that the (n,m)-starfish graphs are obtained from $K_{n-m,m}$ by adding only $O(m^2)$ arcs. In fact, with m chosen as above, one can show that $K_{n-m,m}$ is also reliable (however, the starfish graphs have some additional desirable properties described below).

We note in passing that Kel'mans' design (intended to keep all nodes connected) can only be generalized to our model when $\log f(n) = \theta(\log n)$ (otherwise the design will not be simultaneously optimal and reliable). In such cases, the starfish graphs are more desirable. For example, let us assume f(n) = n (all other cases are similar). Then the starfish design has $m = 2k(p) \log n$, and

(1) While Kel'mans' design is 1-connected, we can show that with high probability, the starfish graphs remain $\Omega(\sqrt{\log n})$-connected. We can prove this statement as follows: (i) With high probability, we can find a set S of hub nodes and a set S' of periphery nodes, with $|S| = |S'| \geq c\sqrt{\log n}$, such that every node of S remains adjacent to every node of S' (where c is a constant). (ii) With high probability, every hub node remains adjacent to (at least) $c'\sqrt{\log n}$ nodes in S (for a certain constant c'). Thus, if X is a set of hub nodes disjoint from S, and $|X| \leq c'\sqrt{\log n}$, then with high probability for every node x of X we can find a distinct node in S adjacent to x. (iii) Finally, with high probability every periphery node remains adjacent to $\Omega(\log n)$ hub nodes.

(2) The Kel'mans design has $\Omega(n/\log n)$ diameter while the starfish designs, with high probability, have diameter at most six.

One disadvantage of the starfish graphs in the case $f(n) = n$ is that they require high layout area (in the VLSI sense). However, we can produce a reliable design with $O(n \log^3 n)$ area, with diameter $O(\sqrt{n/\log n})$ and node connectivity $\Omega(\sqrt{\log n})$ (both with high probability). One such design is obtained by taking $\sim n/m$ copies of K_m (as above) and placing them in the grid points of an $\frac{n}{m}$ - node square grid. Then, the nodes in each copy of K_m are connected to all of the nodes in four neighboring copies of K_m in the grid (thus, the resulting graph has an underlying mesh structure). We will call such a graph a <u>supermesh</u>.

4. A Flow Model with Independent Arc Failures

In this section we will consider a model where flows have to be routed to the nodes of the graph. More specifically, there is a special node (called the source node) from which we have to route one unit of flow to every other node. Each arc has a certain capacity parameter, which is an upper bound to the flow on that arc. Thus, in a feasible flow assignment, the difference between the flows leaving from and arriving to a given node (different from the source) is exactly one.

Such a model arises in the analysis of power networks. Typically, the sum of all capacities will be a good measure of the cost of the network.

We will assume that all capacities are integer, and that arcs fail independently with probability $0 < p < 1$. Now let $\{G(n)\}$ be a sequence of graphs, and $c(n)$ = sum of capacities in $G(n)$. Since all capacities are integer, it is clear that we must have

$$c(n) = \Omega(k(p)\, n \log n)$$

if **G(n)** is not to fail. Surprisingly, this bound is tight. Before proving this fact, we state the following (simple) result.

Let $G = (U,V,E)$ be a bipartite graph, with $|U| = M$. Suppose $\deg u = d$ for all $u \in U$. Then to each $u \in U$ we can assign a node $v \in V$ such that $(u,v) \in E$ and no $v \in V$ is assigned more than $\lceil M/d \rceil$ times. (The proof of this fact is simple. We can construct the desired assignment by first partitioning U into $\lceil M/d \rceil$ subsets of at most d nodes each, and assigning distinct nodes to the nodes in each subset.) □

Now we will produce some provably good designs. These graphs are (n,m) - starfish graphs, with $m = 2k(p) \log n$, and the source located in the periphery. The capacities of the arcs are as follows: all hub arcs have capacity $\sim 2n^{\alpha}$, where $0 \le \alpha < 1$ will be specified below. The source to hub arcs have capacity $\sim \frac{2}{q^2} \frac{n}{m}$ where $q = 1-p$. Finally, the remaining arcs joining the periphery and the hub have capacity 1. We will refer to such a design as a capacitated (n,m)-starfish, or capacitated starfish for short. The total capacity is

$$O(\frac{n}{m}\, m) + O(n^{\alpha} m^2) + O(nm) =$$
$$= O(nm),$$

which is optimal provided we show that the designs are reliable. Let $0 < x < 1 - \frac{1}{\sqrt{2}}$, and $r = r(q, (1-x)q)$ be the Chernoff bound constant. Set

$$\alpha = \max\{0,\ 1 + 2(1-x)q\ k(p)\ \log r\} < 1$$

since $\log r < 0$.

Theorem 2

With α defined as above, the capacitated starfish designs are reliable.

Proof: (a) With high probability the source node remains adjacent to a set L of hub nodes, with $|L| \geq q(1-x)m$.

(b) With high probability, every hub node not in L remains adjacent to $q^2(1-x)^2m$ nodes of L.

(c) The probability that a periphery node remains adjacent to at least $q^2(1-x)^2m$ nodes of L is at least

$$1 - r^{q(1-x)m}.$$

Let us denote by S the set of periphery nodes satisfying this condition. Clearly, in **G(n)** there are at most $2nr^{q(1-x)m} \leq 2n^{\alpha}$ periphery nodes not in S.

(d) With high probability, <u>every</u> periphery node remains adjacent to at least one hub node.

Thus, we can route flow as follows: (1) the flow to nodes in L is sent directly. (2) In order to send flow to the remaining nodes in the hub, and the nodes in S, we send the flow from the source to L, and then directly to the respective node. Clearly this can be performed so that no more than

$$\left\lceil \frac{n}{q^2(1-x)^2 m} \right\rceil$$

units of flow are switched by any node of L (and thus the source to hub capacities we assigned are sufficient, since $\frac{1}{(1-x)^2} < 2$). (3) Finally, the flow to any remaining periphery node v is sent to the hub, then arbitrarily routed within the hub to any node that remains adjacent to v, and finally to v. This concludes the proof of Theorem 2. □

There are other designs that are nearly optimal but superior to the starfish graphs. For example, the supermesh graphs described in the previous section can be adapted to the flow model, with arcs of bounded capacity and total capacity

$$O\left(\frac{k(p)}{q} \, n \log n\right),$$

i.e., off by a factor of $\frac{1}{q}$ from the optimum.

4. A model with dependent arc failures

In the previous two models failures occurred independently. In this section we will consider a case of dependencies in node failures. This is intended to model (for instance) the spreading of damage in a communication network.

The failures will occur in two rounds. In the first round, nodes fail independently, with probability $0 < p < 1$. In the second round, for every node v that failed in the first round, a random neighbor w of v also fails (unless it already failed in the first round). We will then say: v kills w. All the kills are carried out independently (and as a result, damage overlap is allowed).

We want to keep all operative nodes connected. At the same time, we want to guarantee that there are "many" operative nodes (for instance, $\theta(n)$ operative nodes, where n = total number of nodes).

It turns out that the minimum number of arcs that are then required is $\theta(k(p)n \log n)$. Before proving the lower bound, we will sketch the proof of a simple graph-theoretic fact.

Lemma 1

Let G be a connected graph with $n \gg 1$ nodes and at most $cn\log n$ arcs, where c is a constant. Let $d > 2c$ be another constant. Then we can find an independent set I in G with the following properties:

(i) $\log|I| \sim \log n$, and every node of I has degree at most $d\log n$.

(ii) The neighbors of every node of I have degree at least two.

(iii) If a node of I has degree larger than loglogn, then the neighbors of that node also have degree larger than loglogn.

(iv) For every node $z \in G$ that is adjacent to more than one node of I, there is a set A(z) of nodes adjacent to z, such that

(1) $I \cap A(z) = \emptyset$, and no node of A(z) is adjacent to a node of I.

(2) $|A(z)| = \Omega(\log^3 n)$

(3) Every node of A(z) has degree at most dlogn.

(4) $A(z) \cap A(z') = \emptyset$ for any $z' \neq z$.

Proof: In Section 2 we essentially proved that (i) can be attained. Let I be the obtained set, and replace every node of I not satisfying (ii) by a neighbor of degree 1. We will then have a new independent set of the same cardinality as previously which satisfies (ii). Without loss of generality we will also call the new set I. Next, we show how to modify I in order to achieve (iii), while still satisfying (i) and (ii). The proof of (iv) is similar in spirit to that of (iii).

So let us assume that I satisfies (i) and (ii). Algorithm A given below will add nodes to and drop nodes from I until we obtain a new set I all of whose nodes satisfy (iii).

Algorithm A

(1) If every node of I satisfies (iii) STOP.

(2) Otherwise, let v with $\deg v \leq$ loglogn be adjacent to a node of I with degree larger than loglogn. Then we

(2a) <u>Drop</u> from I all neighbors of v.

(2b) <u>Drop</u> from I every node w with $\deg w >$ loglogn, and such that w and v have a common neighbor v' with $\deg v' \leq$ loglogn.

(2c) <u>Add</u> v to I.

(3) Go to (1).

Notice that we only add to I nodes that satisfy (iii) and that we add one such node at every iteration. Step (2b) guarantees that once we add a node, that node will not be dropped in a later iteration. Thus Algorithm A terminates after a finite number T of iterations. Step (2a) guarantees that the new set I is independent, and clearly it still satisfies (ii). Now let N denote the cardinality of the original set I (i.e., before running A), and M the cardinality of the obtained set I.

At every iteration of A we drop at most $(\log\log n)^2$ nodes. Thus, if

$$T \geq n/(2 \log^2 \log n)$$

since we added T nodes, we will have

$$M \geq T, \text{ and thus } \log M \sim \log N \sim \log n.$$

On the other hand, if $T < N/(2\log^2 \log n)$, then

$$M \geq N - T\,(\log\log n)^2 > \frac{N}{2}, \text{ and again}$$

$$\log M \sim \log n.$$

Hence the new set I still satisfies (i). □

Now let $\{G(n)\}$ be a sequence of graphs, and suppose $G(n)$ has at most $\sim \frac{1}{2}k(p')n \log n$ arcs. Let $R_n(\bullet)$ and $R(\bullet)$ have the usual meanings.

Theorem 2

$R(p) = 0$ for $p > p'$.

Proof: Apply Lemma 1 to $G(n)$, with $c = \frac{1}{2} k(p')$ and $k(p') < d < k(p)$. Let C be the set of common neighbors of nodes of I. If $C \neq \emptyset$ then with high probability every node $z \in C$ will be killed by a node in $A(z)$. The probability of this event is at least

$$\prod_{C} \left(1 - \left(1 - \frac{1}{d\log n}\right)^{|A(z)|}\right) \quad \geq$$

$$\left(1 - \left(1 - \frac{1}{d\log n}\right)^{\Omega(\log^3 n)}\right)^n \to 1.$$

We will next show that in **G(n)** at least one node v of I remains operative but all its neighbors are failed, after the two rounds. Given $v \in I$, let $Y(v)$ denote this event.

In order for $Y(v)$ to occur, v must survive the first round, which occurs with probability $q = 1 - p$. If $w \notin C$ is adjacent to v the probability that w fails in the first round but does not kill v is $p(1 - \frac{1}{\deg w})$. Finally, if $z \in C$ is adjacent to v, the probability that z does not kill v is at least

$$q + p\left(1 - \frac{1}{1 + |A(z)|}\right) \geq 1 - O\left(\frac{1}{\log^3 n}\right).$$

Consequently, if $\deg v \leq \log\log n$

$$P(Y(v)) \geq q\left[\left(1 - O\left(\frac{1}{\log^3 n}\right)\right)\left(\frac{p}{2}\right)\right]^{\log\log n} \equiv \alpha(n).$$

If degv > loglogn then

$$P(Y(v)) \geq q\left[(1 - O(\frac{1}{\log^3 n}))p(1 - \frac{1}{\log\log n})\right]^{d\log n} \equiv \beta(n).$$

since in that case every neighbor of v has degree at least loglogn. Clearly $\beta(n) = o(\alpha(n))$, and thus for any $v \in I$, $P(Y(v)) \geq \beta(n)$.

Notice that if Y(v) does not occur, for a certain v, this will imply that at most dlogn nodes of I are killed by neighbors of v. Consequently,

$$P(\text{for at least one } v,\ Y(v)) \geq 1-(1 - \beta(n))^{\Omega(|I|/\log n)} \to 1,$$

since $d < k(p)$, and $-\log\beta(n) \sim -d\log p\log n$. □

Notice that Theorem 1 shows, in fact, that $\Omega(n^x)$ nodes remain operative but isolated in **G(n)**, where $x > 0$.

We will next produce optimal designs. We partition the n nodes into $\sim \frac{n}{m}$ layers of (roughly) m nodes each, where $m = O(k(p)\log n)$, with the precise constant given below. Number the layers 1, 2, ..., n/m. Then, for $1 \leq i < n/m$ we connect all the nodes of layer i to all the nodes of layer (i+1) (i.e., we form a chain of complete bipartite graphs). The total number of arcs is $\sim \frac{n}{m} m^2 = O(k(p)n\log n)$.

Consider first the case $p \geq \frac{3}{4}$, i.e., $q \leq \frac{1}{4}$. We then choose $m = \frac{b}{-\log r(q/2,\ q/2 + q/4)} \log n$, where r is the Chernoff bound constant and b is a large enough constant independent of p. Without loss of generality, all logarithms are to the base e. One can show that

$$-\log r(Q,\ Q + Q/2) = \theta(-\log(1-Q))$$

whenever Q is bounded away from 1. Also, notice that in this case $-\log(1-Q) = \theta(Q)$. Consequently, we will have that

$$m = \theta(\frac{1}{q/2} \log n) = \theta(k(p)\log n).$$

Now we will sketch the proof of reliability of the layered designs. A complete proof is contained in [1]. First, with high probability at least $(q/2)m$ nodes survive the first round, in each layer. Consider layer i. Then, with high probability, the second round effects on layer i from layers (i-1) and (i+1) will overlap with each other (and with the first round failures in layer i) so that we can still count $\sim [(q/32) - \theta(q^2)]m$ in layer i that survive both rounds. Thus, with high probability, $\Omega(qn)$ nodes survive both rounds and remain connected.

The case $p < \frac{1}{4}$ is similar. □

* * * *

There is a variation of the model described in this section which is also interesting. Suppose that in the second round, <u>all</u> neighbors of every node that failed in the first round are killed. Not surprisingly, the main result concerning this model is that there are <u>no</u> survivable designs, meaning:

With high probability, either

(i) Most nodes $(\sim n)$ do not survive both rounds, or

(ii) Many nodes remain operative but isolated, or

(iii) Both (i) and (ii) occur.

The proof of this fact (contained in [1]) is somewhat lengthy, but similar in spirit to the lower bound proof in this section.

5. Some extensions

The problems considered in the previous sections can be extended to spatial settings. These extensions are interesting because they allow geometry to play a role in costs and failure dependencies, and they are meaningful becuase such problems do arise in practice.

In [1] we have considered two closely related spatial settings: in one setting the nodes occur randomly within a square, and in the other the nodes occupy grid points within a square (this particular setting is similar to that used in [7]). There are two possible ways of approaching the hardware vs. reliability tradeoffs in such spatial settings:

(1) The edges have negligible width. In this case, appropriate cost measures include total arc lengths, length of longest arc, and product of arc capacity times arc length.

(2) The edges have a nonzero (say, constant), width. This assumption gives rise to area/reliability trade-off problems.

These topics will be discussed in a future paper.

* * *

References

[1] D. Bienstock, Large-Scale Network Reliability, Ph.D. Thesis, MIT Operations Research Center (1985).

[2] H. Chernoff, A Measure of Asymptotic Efficiency for Tests of a Hypothesis Based on a Sum of Observations, Ann. Math. Stat. 23 (1952) 311-327.

[3] P. Erdös and A. Rényi, On Random Graphs I, Publicationes Matematicae 6 (1959) 290-297.

[4] P. Erdös and J. Spencer, Probabilistic Methods in Combinatorics, Academic Press, New York (1960).

[5] A. K. Kel'mans, Some Problems of Network Reliability Analysis, Automation and Remote Control 3 (1965) 564-573.

[6] A. K. Kel'mans, Connectivity of a Probabilistic Network, Automation and Remote Control 3 (1967) 444-460.

[7] F. T. Leighton and C. E. Leiserson, Water-Scale Integration of Systolic Arrays, Report MIT/LCS/TM-236 (1983), MIT Laboratory for Computer Science.