

KRZYSZTOF DIKS

ANDRZEJ PELC

**Fast diagnosis of multiprocessor systems
with random faults**

Informatique théorique et applications, tome 27, n° 5 (1993),
p. 391-401

http://www.numdam.org/item?id=ITA_1993__27_5_391_0

© AFCET, 1993, tous droits réservés.

L'accès aux archives de la revue « Informatique théorique et applications » implique l'accord avec les conditions générales d'utilisation (<http://www.numdam.org/conditions>). Toute utilisation commerciale ou impression systématique est constitutive d'une infraction pénale. Toute copie ou impression de ce fichier doit contenir la présente mention de copyright.

NUMDAM

Article numérisé dans le cadre du programme
Numérisation de documents anciens mathématiques

<http://www.numdam.org/>

FAST DIAGNOSIS OF MULTIPROCESSOR SYSTEMS WITH RANDOM FAULTS (*)

by Krzysztof DIKS ⁽¹⁾ and Andrzej PELC ⁽²⁾

Communicated by W. RYTTER

Abstract. – *Processors in a multiprocessor system fail independently with constant probability $0 < p < 1/2$. They can test one another and faulty processors have to be identified on the basis of test results. Fault-free processors diagnose other fault-free processors correctly and find faults in faulty ones with probability $q \leq 1$ in each test. Faulty testers are unreliable: they may even behave maliciously. Tests are independent and in every time unit a processor can be involved in at most one test. We propose testing schemes which are fast, use few tests and are correct with probability converging to 1 as the size of the system grows.*

Résumé. – *Les processeurs dans un système multiprocesseur tombent en panne de façon indépendante avec probabilité constante $0 < p < 1/2$. Ils peuvent effectuer des tests un sur l'autre et tous les processeurs défectueux doivent être identifiés sur la base des résultats de ces tests. Les processeurs fonctionnels diagnostiquent correctement les autres processeurs fonctionnels et trouvent des défauts dans les processeurs défectueux avec probabilité $q > 0$ pendant chaque test. Les processeurs défectueux sont non fiables comme testeurs : ils peuvent même se comporter malicieusement. Les tests sont indépendants et pendant chaque unité de temps un processeur peut être impliqué dans un seul test. Nous proposons des schémas de tests qui sont rapides, utilisent un petit nombre de tests et sont corrects avec probabilité convergente vers 1 lorsque la grandeur du système croît.*

INTRODUCTION

Reliability of large multiprocessor systems has recently gained growing attention as a result of their increasing role in computing. An important problem in this area, known as system-level fault diagnosis, is to identify all faulty processors (units) in the system. Units can test one another, the result

(*) Received Received September 26, 1991; accepted November 10, 1992.

⁽¹⁾ Instytut Informatyki, Uniwersytet Warszawski, ul. Banacha 2, 00-913 Warszawa 59, Poland.

Research supported in parts by grants KBN 2-1190-91-01 and KBN 2-2043-92-03.

⁽²⁾ Département d'Informatique, Université du Québec à Hull, C. P. 1250, succ. "B", Hull, Québec J8X 3X7, Canada.

Research supported by the Natural Science and Engineering Research Council of Canada, grant OGP 0008136.

of a test can be “faulty” or “fault-free”, and a central monitor has to identify faulty processors on the basis of all test results. Different restrictions on faults and different interpretation of test results give rise to many models used for diagnosis (*see* [8]).

Two types of assumptions about faults are usually made: either an upper bound t on the number of faults is imposed and diagnosis in worst-case fault configuration is sought (e. g. [13]) or processors are assumed to fail independently with given probability [3 to 7, 9, 10, 12, 14, 15]. Since the number of faults in a large system may be proportional to its size and some configurations of t faults in an n -unit system require $\Theta(nt)$ tests for diagnosis (*cf.* [13]) it follows that the worst case scenario leads to procedure using a quadratic number of tests. In the probabilistic setting, more efficient procedures are known to work with high probability (*cf.* [5, 6, 9, 10, 14]), whence the growing popularity of probabilistic models.

Apart from the number of tests, the time of diagnosis execution is an important measure of its efficiency. It is reasonable to assume (*cf.* [2, 16]) that every processor can be involved in at most one test in a unit of time (that is, it can test at most one other processor, or –exclusively– be tested by at most one processor). Only the number of time units used for testing is accounted for in execution time, assuming that the time used by the central monitor to process test results is negligible. The problem of minimizing testing time using this scenario was considered in [2, 16] under the assumption that at most t units can be faulty and that testing can be performed adaptively.

The aim of this paper is to give diagnosis schemes which work fast, use few tests and are correct with high probability. We work in two very general probabilistic models introduced by Blough, Sullivan and Masson (*cf.* [4, 5, 6]). Processors fail independently with fixed probability $0 < p < 1/2$. This bound cannot be relaxed: no reliable diagnosis is possible for $p \geq 1/2$. A fault-free processor always diagnoses another fault-free processor correctly. In the permanent fault model it also diagnoses every faulty processor correctly (a fault is always detected), while in the intermittent fault model it detects a fault in a faulty processor with some probability $0 < q < 1$. Results of different tests are independent. Faulty testers can report arbitrary test results, they can even act maliciously. We assume that testing is synchronous and in every time unit a processor can be involved in at most one test. Moreover, we assume that every processor can test every other processor (this assumption will be subsequently relaxed). Since our testing schemes are non-adaptive (the scheduling of all tests is done in advance), they can be modelled as labelled directed multigraphs where nodes stand for processors, arrows stand

for tests and a label of an arrow indicates at which time unit the test should be performed. All our negative results also refer to the class of non-adaptive diagnosis schemes.

We say that a diagnosis scheme is almost safe if the probability that the central monitor correctly identifies all faulty processors upon completion of all tests converges to 1, as the size of the system grows. Our schemes are almost safe and we prove that in many cases they are asymptotically optimal with respect to execution time and to the total number of tests.

For every diagnosis scheme D working for a system of n processors we use the following notation:

$T(D, n)$ – the time of execution of D (the number of time units it uses),

$S(D, n)$ – the number of tests used by D ,

$R(D, n)$ – the reliability of D (the probability that D is correct).

For every real number x , $[x]$ denotes the largest integer $\leq x$ and $\lceil x \rceil$ denotes the least integer $\geq x$. We write $\log x$ instead of $\log_2 x$. For any set S , $|S|$ denotes the size of S . For a random event E , \bar{E} denotes its complement.

2. THE PERMANENT FAULT MODEL

Denoting fault-free processors by $+$ and faulty processors by $-$, test results in this model can be summarized as follows.

$$\begin{array}{r}
 + \text{---“fault-free”---} + \\
 + \text{---“faulty”---} - \\
 - \text{---arbitrary---} + \text{ or } -
 \end{array}$$

In [6] an almost safe diagnosis for the permanent fault model was proposed, which used $nf(n)$ tests for every function $f: \mathcal{N} \rightarrow \mathcal{N}$ such that $\lim_{n \rightarrow \infty} f(n) = \infty$. It was also proved in [4] that no diagnosis using $O(n)$ tests can be almost safe. We start with a similar negative result concerning time execution.

THEOREM 1: *For every constant $c < (-1/\log p)$ and every diagnosis scheme D , if $T(D, n) \leq c \log n$ then D is not almost safe.*

Proof: Let D be a diagnosis scheme such that $T(D, n) \leq c \log n$. Every processor can be tested by at most $c \log n$ different processors and can test at most $c \log n$ different processors. Call two processors independent if they do not have common testers. Thus there exists a set S of at least $\lfloor n/(c^2 \log^2 n) \rfloor$ pairwise independent processors. Fix one processor $u \in S$. Let $E(u)$ be the event that all testers of u are faulty. Thus $\Pr(E(u)) \geq p^c \log n$. Let $E = \bigcup_{u \in S} E(u)$. Since events $E(u)$, for $u \in S$, are mutually independent, we have

$$\Pr(E) \geq 1 - (1 - p^c \log n)^{\lfloor n/(c^2 \log^2 n) \rfloor} = 1 - (1 - n^c \log p)^{\lfloor n/(c^2 \log^2 n) \rfloor},$$

hence $\lim_{n \rightarrow \infty} \Pr(E) = 1$, in view of $c < (-1/\log p)$. Clearly, if E holds, the status of the processor which has only faulty testers cannot be guessed with probability exceeding $1 - p$. Let $\text{COR}(D, n)$ denote the event that diagnosis D is correct for a system of size n . Then

$$\begin{aligned} R(D, n) &= \Pr(\text{COR}(D, n)) = \Pr(\text{CORD}(D, n) | E) \Pr(E) \\ &\quad + \Pr(\text{COR}(D, n) | \bar{E}) \Pr(\bar{E}) \leq 1 - p + \Pr(\bar{E}). \end{aligned}$$

Since $\Pr(\bar{E})$ converges to 0, D is not almost safe.

Thus every almost safe diagnosis scheme must work in at least logarithmic time and use a superlinear number of tests. The main goal of this section is to establish trade-offs between the number of tests and execution time of every almost safe diagnosis scheme.

We start with the description of a general testing procedure. Let $A = \{0, \dots, t-1\}$ be a set of processors and $B \subset A$, $B = \{0, \dots, s-1\}$, be a set of testers. The following procedure tests all processors from A by all testers from B in time $2s + t - 3$.

```

Procedure Pipeline-Test(A, B)
  for time ← 1 to t - 1 + 2(s - 1) do
    for all 0 ≤ j ≤ s - 1 in parallel do
      if 2j + 1 ≤ time ≤ 2j + t - 1 then
        j tests (time - j) mod t
      endif
    endfor
  endfor
end Pipeline-Test

```

Next we describe a class of diagnosis schemes. Let $t, s: \mathcal{N} \rightarrow \mathcal{N}$ be functions such that $1 \leq s(n) \leq t(n) \leq n$, for every natural n . We denote by $GD(s, t)$ (GD stands for group diagnosis) the following scheme.

Partition all n processors into $\lfloor n/t(n) \rfloor$ disjoint subsets $G_1, \dots, G_{\lfloor n/t(n) \rfloor}$, each containing $t(n)$ or $t(n) + 1$ processors. Choose a subset S_i of size $s(n)$ in every group G_i . Testing proceeds for all $i = 1, \dots, \lfloor n/t(n) \rfloor$ in parallel, using Pipeline-Test(G_i, S_i). For any distinct u, v , such that $u \in S_i$ and $v \in G_i$, let $result(u, v)$ denote the result ("faulty" or "fault-free") of the test of v by u . Additionally put $result(u, u) =$ "fault-free" for any u . Then the status of each processor is decided as follows. If $v \in G_i$ then v is diagnosed as faulty if $|\{u \in S_i : result(u, v) = \text{"faulty"}\}| > s(n)/2$; otherwise v is diagnosed as fault-free. Clearly, the diagnosis scheme $GD(s, t)$ uses less than $s(n)n$ tests and works in time at most $3t(n)$.

It should be noted that for $t(n) = n$ our diagnosis scheme coincides with that from [6].

LEMMA 1: Let $r = 1 - p$. If $\lim_{n \rightarrow \infty} (1 - \exp(-((r - 1 + 1/(4r))s(n))/2))^{n/t(n)} = 1$ then $GD(s, t)$ is almost safe.

Proof: Call a set S_i of testers doubtful if at most $s(n)/2$ of its elements are fault-free. By Chernoff bound (cf. [1, 11]) the probability that S_i is doubtful does not exceed $\exp(-\epsilon^2 s(n)r/2)$, where $\epsilon = 1 - 1/(2(1 - p)) = (2r - 1)/(2r)$. If no set S_i is doubtful then $GD(s, t)$ is correct. Hence

$$R(GD(s, t), n) \geq \left(1 - \exp\left(\frac{-\epsilon^2 s(n)r}{2}\right)\right)^{\lfloor n/t(n) \rfloor} \geq \left(1 - \exp\left(\frac{-(r - 1 + 1/(4r))s(n)}{2}\right)\right)^{n/t(n)}$$

The next lemma gives combinations of the number of tests and execution time which prohibit a diagnosis scheme from being almost safe.

LEMMA 2: For any diagnosis scheme D , if $S(D, n) \leq s(n)n$, $T(D, n) \leq t(n)$ and

$$\lim_{n \rightarrow \infty} (1 - p^{2s(n)n/(4s(n)t(n))}) = 0$$

then D is not almost safe.

Proof: Consider a scheme D satisfying the above assumptions. At least $n/2$ processors are tested by at most $2s(n)$ testers. Since $2s(n)$ testers can test

at most $2s(n)t(n)$ processors in time $t(n)$, there are at least $n/(4s(n)t(n))$ processors which have pairwise disjoint groups of testers of size at most $2s(n)$. Similarly as in the proof of Theorem 1 we can show that with probability converging to 1, all testers are faulty in one of those groups. As before this implies that D is not almost safe.

The main result of this section shows how restricting testing time forces to increase the number of tests in order to make almost safe diagnosis possible. While linear time permits any superlinear number of tests (and going down to $O(n)$ tests prohibits almost safe diagnosis by the result from [4]), time $O(n^\alpha)$ for any $0 < \alpha < 1$ forces $\Theta(n \log n)$ tests. However, for a suitable constant k , performing $kn \log n$ tests permits to go down to logarithmic testing time and this in turn cannot be improved in view of Theorem 1.

THEOREM 2: *Consider the permanent fault model.*

1. *If $t(n) \in \Theta(n)$ and $s: \mathcal{N} \rightarrow \mathcal{N}$ is such that $\lim_{n \rightarrow \infty} s(n) = \infty$ then there exists an almost safe diagnosis scheme D for which $T(D, n) = t(n)$ and $S(D, n) = s(n)n$.*

2. *If $t(n) \in \Theta(n/\log n)$ then there exist two constants $c_1 > c_2 > 0$ such that:*

(i) *there exists an almost safe diagnosis scheme D for which $T(D, n) = t(n)$ and $S(D, n) \leq c_1 n \log \log n$;*

(ii) *no almost safe diagnosis scheme D can have $T(D, n) = t(n)$ and $S(D, n) \leq c_2 n \log \log n$.*

3. *If $t(n) \in \Theta(n^\alpha)$, $0 < \alpha < 1$, then there exist constants $d_1 > d_2 > 0$ such that:*

(i) *there exists an almost safe diagnosis scheme D for which $T(D, n) = t(n)$ and $S(D, n) \leq d_1 n \log n$;*

(ii) *no almost safe diagnosis scheme D can have $T(D, n) = t(n)$ and $S(D, n) \leq d_2 n \log n$.*

4. *There exist constants $k_1, k_2 > 0$ and an almost safe diagnosis scheme D such that $T(D, n) \leq k_1 \log n$ and $S(D, n) \leq k_2 n \log n$.*

Proof: 1. If $t(n) \geq cn$ for some constant $c > 0$ and $\lim_{n \rightarrow \infty} s(n) = \infty$ then

$$\lim_{n \rightarrow \infty} \left(1 - \exp \left(\frac{-(r-1 + 1/(4r))s(n)}{2} \right) \right)^{n/t(n)} = 1$$

and consequently $GD(s, t)$ is almost safe in view of Lemma 1.

2. (i) If $t(n) \geq cn/\log n$ then for any constant d and $s(n) \geq d \log \log n$ we have (putting $x = (r - 1 + 1/(4r))/2$)

$$\begin{aligned} \lim_{n \rightarrow \infty} \left(1 - \exp\left(\frac{-(r-1+1/(4r))s(n)}{2}\right) \right)^{n/t(n)} &\geq \lim_{n \rightarrow \infty} (1 - \exp(-xd \log \log n))^{\log n/c} \\ &= \lim_{n \rightarrow \infty} (1 - (\log n)^{-xd \log e})^{\log n/c}. \end{aligned}$$

The latter limit is 1 for $d > 1/(x \log e)$. Hence it suffices to take $c_1 > d$ and for some $s(n) \leq c_1 \log \log n$ $GD(s, t)$ will be almost safe, in view of Lemma 1.

(ii) If $t(n) \leq cn/\log n$ then for any constant c_2 and $s(n) \leq c_2 \log \log n$ we have

$$\begin{aligned} \lim_{n \rightarrow \infty} (1 - p^{2s(n)n/(4s(n)t(n))}) &\leq \lim_{n \rightarrow \infty} (1 - p^{2c_2 \log \log n \log n / (4cc_2 \log \log n)}) \\ &= \lim_{n \rightarrow \infty} (1 - (\log n)^{2c_2 \log p})^{\log n / (4cc_2 \log \log n)}. \end{aligned}$$

The latter limit is 0 for $c_2 \leq -1/(2 \log p)$. Hence, in view of Lemma 2 no diagnosis scheme D with $T(D, n) \leq cn/\log n$ and $S(D, n) \leq c_2 n \log \log n$ (for such c_2) can be almost safe.

The proofs of 3 and 4 are similar to 2.

3. THE INTERMITTENT FAULT MODEL

Test results in this model can be summarized as follows.

$$\begin{aligned} &+ \frac{\text{“fault-free”}}{\quad} + \\ &+ \frac{\text{“faulty” with probability } 0 < q < 1}{\quad} - \\ &- \frac{\text{arbitrary}}{\quad} + \text{or} - \end{aligned}$$

In [5] an almost safe diagnosis scheme using $\omega(n)n \log n$ tests for any function $\omega(n) \rightarrow \infty$, was described and it was proved that no almost safe diagnosis can use $o(n \log n)$ tests. Later Berman and Pelc [3] gave an almost safe diagnosis using $O(n \log n)$ tests and proved that for some positive constant c , no diagnosis scheme using less than $cn \log n$ tests can be almost safe. The aim of this section is to show an almost safe diagnosis scheme

working in time $O(\log n)$ and using $O(n \log n)$ tests. In view of Theorem 1 (this negative result clearly applies to the intermittent fault model as well) and of [3, 5], both these characteristics are asymptotically optimal for this model.

Let k be the smallest positive integer for which $x = p + (1 - q)^k < 1/2$. Let c be a positive constant to be determined later. The diagnosis scheme $MGT(c)$ (MGT stands for multiple group testing) is described as follows. Partition all n processors into $m = \lceil n / \lceil c \log n \rceil \rceil$ disjoint subsets (groups) G_1, \dots, G_m , each containing $\lceil c \log n \rceil$ or $\lceil c \log n + 1 \rceil$ processors. Testing proceeds for all $i = 1, \dots, m$ in parallel, by k consecutive calls of *Pipeline-Test*(G_i, G_i). This way, each processor is tested k times by all other processors of its group. For any $i \leq m$ and every distinct $u, v \in G_i$, define $result(u, v)$ to be "fault-free" if all k tests of v by u give result "fault-free" and to be "faulty" otherwise. Additionally put $result(u, u) =$ "fault-free" for any u . Then the status of each processor is described as follows. If $v \in G_i$ then v is diagnosed as faulty if

$$|\{u \in G_i : result(u, v) = \text{"faulty"}\}| > \frac{|G_i|}{2};$$

otherwise v is diagnosed as fault-free. Clearly $MGT(c)$ uses less than $k \lceil c \log n \rceil n$ tests and works in time at most $3k \lceil c \log n \rceil$.

THEOREM 3: *Consider the intermittent fault model.*

There exists an almost safe diagnosis scheme working in time $O(\log n)$ and using $O(n \log n)$ tests.

Proof: It suffices to show that for some positive constant $c > 0$, the diagnosis scheme $MGT(c)$ is almost safe. Denote $r = 1 - p$.

Define two random events: E_1^n is the event that in every group G_i , $i \leq m$, more than $|G_i|/2$ processors are fault-free and E_2^n if the event that for every $i \leq m$ and every faulty processor $v \in G_i$,

$$|\{u \in G_i : result(u, v) = \text{"faulty"}\}| > \frac{|G_i|}{2}.$$

Clearly, if $\lim_{n \rightarrow \infty} \Pr(E_1^n) = 1$ and $\lim_{n \rightarrow \infty} \Pr(E_2^n) = 1$ then $MGT(c)$ is almost safe.

By Chernoff bound, the probability that in a fixed group G_i at most $|G_i|/2$ processors are fault-free is at most $\exp(-\varepsilon^2 |G_i|/2)$, where $\varepsilon = 1 - 1/(2r) > 0$

because $r > 1/2$. Thus

$$\Pr(E_1^n) \geq \left(1 - \exp\left(\frac{-\epsilon^2 \lceil c \log n \rceil r}{2}\right)\right)^m$$

which converges to 1 for $c > 8r / ((2r - 1)^2 \log e)$.

Next, consider a faulty processor $v \in G_i$. For any other processor $u \in G_i$ the probability that either u is faulty or is fault-free but v passed all k tests performed by u , is at most $x = p + (1 - q)^k < 1/2$. Thus with probability $1 - x > 1/2$, $result(u, v) = \text{"faulty"}$. By Chernoff bound, the probability that

$$|\{u \in G_i : result(u, v) = \text{"faulty"}\}| \leq \frac{|G_i|}{2}$$

is at most $\exp((-\lambda^2 |G_i| (1 - x))/2)$, where $\lambda = 1 - 1/(2(1 - x)) > 0$ because $1 - x > 1/2$. Thus

$$\Pr(E_2^n) \geq \left(1 - \exp\left(\frac{-\lambda^2 \lceil c \log n \rceil (1 - x)}{2}\right)\right)^n$$

which converges to 1 for $c > 8(1 - x) / ((1 - 2x)^2 \log e)$.

It follows that for

$$c > \max\left(\frac{8(1 - p)}{(1 - 2p)^2 \log e}, \frac{8(1 - x)}{(1 - 2x)^2 \log e}\right)$$

the diagnosis scheme $MGT(c)$ is almost safe.

4. CONCLUSIONS

We presented fast diagnosis schemes using few tests and correctly locating all faulty processors with probability converging to 1, as the number of processors grows. In the permanent fault model we showed trade-offs between time and the total number of tests for almost safe diagnosis schemes. For example, while for linear time, $ns(n)$ tests, for any function $s(n) \rightarrow \infty$, are sufficient, time $\Theta(n^\alpha)$, $0 < \alpha < 1$, requires $\Theta(n \log n)$ tests. In the intermittent fault model, we showed an almost safe diagnosis scheme working in time $O(\log n)$ and using $O(n \log n)$ tests. None of these asymptotic bounds can be improved.

Although we assumed that every processor can test every other processor (which requires full interconnection of the system) this hypothesis can in fact be relaxed. All our schemes use at most $O(n \log n)$ tests and consequently, only this many links in the system are required.

When proving our negative results we argued that some combinations of time and number of tests prohibit almost safe diagnosis by showing that no diagnosis satisfying these constraints can exceed some constant probability of correctness, for large n . Using slightly more complicated arguments these results can be sharpened: it can often be proved that the probability of correctness of every diagnosis satisfying given time and number of tests constraints must converge to 0, as the number of processors grows.

REFERENCES

1. D. ANGLUIN and L. G. VALIANT, Fast Probabilistic Algorithms for Hamiltonian Circuits and Matchings, *J. Comput. System Sci.*, 1979, 18, pp. 155-193.
2. R. BEIGEL, S. R. KOSARAJU and G. F. SULLIVAN, Locating Faults in a Constant Number of Parallel testing Rounds, in *Proceedings of the 1989 ACM Symposium on Parallel Algorithms and Architectures*, pp. 189-198.
3. P. BERMAN and A. PELC, Distributed Probabilistic Fault Diagnosis for Multiprocessor Systems, *Digest of Papers, FTCS-20, 1990*, pp. 340-346.
4. D. M. BLOUGH, Fault Detection and Diagnosis in Multiprocessor Systems, Ph. D. Thesis, The John Hopkins University, 1988.
5. D. M. BLOUGH, G. F. SULLIVAN and G. M. MASSON, Almost Certain Diagnosis for Intermittently Faulty Systems, *Digest of Papers, FTCS-18, 1988*, pp. 260-271.
6. D. M. BLOUGH, G. F. SULLIVAN and G. M. MASSON, Fault Diagnosis for Sparsely Interconnected Multiprocessor Systems, *Digest of Papers, FTCS-19, 1989*, pp. 62-69.
7. M. L. BLOUNT, Probabilistic Treatment of Diagnosis in Digital Systems, *Digest of Papers, FTCS-7, 1977*, pp. 72-77.
8. A. T. DAHBURA, System-Level Diagnosis: a Perspective for the Third Decade, in *Concurrent Computation: Algorithms, Architectures, Technologies, Plenum Publ. Corp.*, 1988.
9. A. T. DAHBURA, K. K. SABNANI and L. L. KING, The Comparison Approach to Multiprocessor Fault Diagnosis, *IEEE Trans. Comput.*, March 1987, 36, pp. 373-378.
10. D. FUSSELL and S. RANGARAJAN, Probabilistic Diagnosis of Multiprocessor Systems with Arbitrary Connectivity, *Digest of Papers, FTCS-19, 1989*, pp. 560-565.
11. T. HAGERUP and Ch. RÜB, A Guided Tour of Chernoff Bounds, *Inf. Proc. Lett.*, 1989/1990, 33, pp. 305-308.
12. S. N. MAHESHWARI and S. L. HAKIMI, On Models for Diagnosable Systems and Probabilistic Fault Diagnosis, *IEEE Trans. Comput.*, March 1976, 25, pp. 228-236.

13. F. P. PREPARATA, G. METZE and R. T. CHIEN, On the Connection Assignment Problem of Diagnosable Systems, *IEEE Trans. Electr. Comput.*, December 1967, 16, pp. 848-854.
14. S. RANGARAJAN and D. FUSSEL, A Probabilistic Method for Fault Diagnosis of Multiprocessor Systems, *Digest of Papers, FTCS-18*, 1988, pp. 278-283.
15. E. R. SCHEINERMAN, Almost Sure Fault Tolerance in Random Graphs, *SIAM J. Comput.*, 1987, 16, pp. 1124-1134.
16. E. SCHMEICHEL, S. L. HAKIMI, M. OTSUKA and G. SULLIVAN, On Minimizing Testing Rounds for Fault Identification, *Digest of Papers, FTCS-18*, 1988, pp. 266-271.