

Reliable Fault Diagnosis with Few Tests

ANDRZEJ PELC¹† and ELI UPFAL²‡

¹ Département d'Informatique, Université du Québec à Hull,
Hull, Québec J8X 3X7, Canada
(e-mail: pelc@uqah.quebec.ca)

² IBM Almaden Research Center, San Jose, CA 95120, USA,
and
Department of Applied Mathematics,
The Weizmann Institute of Science, Rehovot, Israel
(e-mail: eli@wisdom.weizmann.ac.il)

Received 23 May 1996; revised 15 August 1997

We consider the problem of fault diagnosis in multiprocessor systems. Processors perform tests on one another: fault-free testers correctly identify the fault status of tested processors, while faulty testers can give arbitrary test results. Processors fail independently with constant probability $p < 1/2$ and the goal is to identify correctly the status of all processors, based on the set of test results. For $0 < q < 1$, q -diagnosis is a fault diagnosis algorithm whose probability of error does not exceed q . We show that the minimum number of tests to perform q -diagnosis for n processors is $\Theta(n \log \frac{1}{q})$ in the nonadaptive case and $n + \Theta(\log \frac{1}{q})$ in the adaptive case. We also investigate q -diagnosis algorithms that minimize the maximum number of tests performed by, and performed on, processors in the system, constructing testing schemes in which each processor is involved in very few tests. Our results demonstrate that the flexibility yielded by adaptive testing permits a significant saving in the number of tests for the same reliability of diagnosis.

1. Introduction

As the size of commercially available multiprocessor systems grows, they become increasingly vulnerable to component failures. This yields growing interest in the issue of reliability of such systems. One of the major problems in this area, known as the *fault diagnosis problem*, is to locate all faulty processors in the system. The classical approach

† Research supported in part by NSERC grant OGP 0008136.

‡ Work at the Weizmann Institute supported in part by the Norman D. Cohen Professorial Chair of Computer Science, a MINERVA grant, and a grant from the Israeli Academy of Science.

to fault diagnosis was originated by Preparata, Metze and Chien [18]. Processors perform tests on one another and diagnosis is based on the collection of test results. It is assumed that fault-free processors always give correct test results, while tests conducted by faulty processors are totally unpredictable: a faulty tester can output any test result regardless of the status of the tested processor. Faults are assumed permanent, that is, the fault status of a processor does not change during testing and diagnosis. In [18] a worst-case scenario is adopted: it is assumed that at most t processors are faulty and that they are placed in locations most detrimental for diagnosis. Also, it is assumed that all tests are determined in advance and they cannot be rescheduled during the diagnosis process.

This model and some of its variations have been thoroughly studied in the literature (see the survey [10], where an extensive bibliography can be found). It has been argued that the worst-case scenario often fails to reflect realistic diagnosis situations. As an alternative, various probabilistic models were proposed (see [6], [7], [8], [9], [15], [16], [19], [20]). Instead of imposing an upper bound on the number of faulty processors and assuming their worst-case location, an *a priori* failure probability, independent for each processor, is assumed in these models. Diagnosis is then restricted to sets of faulty processors of sufficiently high *a priori* probability [15], in which case it can be performed unambiguously [9], or is done in general and has a high probability of correctness (see [5], [6], [7], [8], [16], [19], [20]).

Nakajima [14] was the first to modify the assumption that all tests are scheduled in advance. He proposed a new approach called *adaptive diagnosis*, in which the next test can be determined after seeing the results of previous ones. (Classical diagnosis is called *nonadaptive*.) The flexibility of adaptive diagnosis increases its efficiency. In [13], [21], [2], [3] and [1], the parallel time (number of rounds) of adaptive diagnosis was investigated, assuming that tests involving disjoint pairs of processors can be conducted in the same round. It was shown that, while locating $t < \frac{n}{2}$ faults requires worst-case time t in the nonadaptive setting, adaptive diagnosis can locate fewer than $\frac{n}{2}$ faults among n in constant time. On the other hand, Blecher [4] showed that the number of tests required to identify $t < \frac{n}{2}$ faults in worst case decreases from tn for nonadaptive diagnosis to $n + t - 1$ in the adaptive setting.

In this paper, which is an extended version of [17], we work in the probabilistic model previously studied in [6] and [20]. The assumptions concerning test results are the same as in the above-described model of Preparata, Metze and Chien [18] and faults are also permanent. However, unlike in [18], it is assumed that processors fail with constant probability $p < 1/2$ and all faults are independent. It should be noted that this is the only probabilistic model in which no assumption is made on the behaviour of faulty testers. Thus diagnosis algorithms working reliably under this model are very robust in that they produce correct diagnosis under any behaviour of faulty processors.

In [6] the authors show a simple nonadaptive diagnosis strategy based on majority vote whose probability of error converges to 0 as the system grows. The number of tests they use is slightly larger than linear in the number of nodes. In [11] it is shown that this majority strategy maximizes the probability of correctness for test assignments considered in [6].

The goal of the present paper is to establish a precise relation between error probability q of a diagnosis strategy and the minimum number of tests it requires, both in the

nonadaptive and in the adaptive scenario. We show that the minimum number of tests used by such a strategy for n processors is $\Theta(n \log \frac{1}{q})$ in the nonadaptive case and $n + \Theta(\log \frac{1}{q})$ in the adaptive case. In both cases we present concrete diagnosis algorithms that match these bounds. We also investigate diagnosis algorithms that minimize the maximum number of tests involving any processor in the system, presenting efficient algorithms in which each processor is involved in only very few tests.

The paper is organized as follows. In Section 2 we establish terminology, formalize the model and state some preliminary results. In Sections 3 and 4 we study nonadaptive and adaptive diagnosis, respectively. Section 5 contains conclusions.

2. Terminology and preliminaries

Let $U = \{u_1, \dots, u_n\}$ be the set of processors. A *test assignment* is represented by a directed graph $G = (U, E)$ where $(u, v) \in E$ means that processor u tests processor v . The *degree* of a test assignment G is the maximum over all $u \in U$ of $|\{v : (u, v) \in E\} \cup \{v : (v, u) \in E\}|$. For every $u \in U$ and a given test assignment G , processors $v \in U$ such that $(u, v) \in E$ or $(v, u) \in E$ are called *neighbours* of u . The *distance* between processors u and v is the smallest k such that there exists a sequence of processors $u = u_0, u_1, \dots, u_k = v$ whose consecutive terms are neighbours. The outcome of a test $(u, v) \in E$ is 1 (0) if u evaluates v as faulty (fault-free). The test assignment is given *a priori* in the case of nonadaptive diagnosis. In the adaptive scenario it is constructed dynamically, the next test depending on the results of previous tests. A complete collection of test results is called a *syndrome*. Formally, a syndrome is a function $S : E \rightarrow \{0, 1\}$. The set of all possible syndromes is denoted by Σ . The set of all faulty processors in the system is called a *fault set*. This can be any subset of U . A syndrome S is said to be *compatible* with a fault set F if, for any $(u, v) \in E$, such that $u \in U \setminus F$, $S(u, v) = 1$ if and only if $v \in F$. This corresponds to the assumption that fault-free processors always give correct test results. Since faulty testers can give arbitrary test results, any syndrome compatible with a fault set F can occur when faulty processors in the system are exactly those in F . The set of all syndromes compatible with a fault set F is denoted by $\sigma(F)$. Fault sets F_1 and F_2 are called *associated* if $\sigma(F_1) \cap \sigma(F_2) \neq \emptyset$.

We consider only deterministic diagnosis algorithms. The input of such an algorithm is a syndrome and the output is the set of processors that the algorithm diagnoses as faulty (all other processors are implicitly diagnosed as fault-free). Thus a *diagnosis* is any function $D : \Sigma \rightarrow \mathcal{P}(U)$.

We now define formally the probability of correctness of any diagnosis. Let $p < \frac{1}{2}$ be the probability that a processor is faulty. This probability is fixed and considered as a constant throughout the paper. The sample space is the set of all fault sets, that is,

$$\Omega = \{F : F \subset U\}.$$

The probability function P is defined for all subsets of Ω by the formula

$$P(X) = \sum_{F \in X} p^{|F|} (1-p)^{n-|F|},$$

for any $X \subset \Omega$. If D is a diagnosis, $\text{Cor}(D)$ is the event consisting of those fault sets F for which D returns F on any syndrome compatible with F , that is, the event that diagnosis D is correct regardless of faulty processors' behaviour. More precisely,

$$\text{Cor}(D) = \{F \subset U : \forall S \in \sigma(F) D(S) = F\}.$$

The probability of error of diagnosis D is

$$\text{Err}(D) = 1 - P(\text{Cor}(D)).$$

A diagnosis D is called a q -diagnosis if $\text{Err}(D) \leq q$.

If two subsets are associated then at most one of them can belong to $\text{Cor}(D)$. Since every set F and its complement are associated (the common compatible syndrome is the one that gives result 0 for u and v both in or both outside F , and result 1 otherwise), we have the following observation (see [11]).

Proposition 2.1. *For any diagnosis D and any fault set F , F and $U \setminus F$ cannot both belong to $\text{Cor}(D)$.*

Among complementary fault sets F and $U \setminus F$, the one with smaller size has higher probability because $p < \frac{1}{2}$. Hence the highest possible value of $P(\text{Cor}(D))$ is obtained if $\text{Cor}(D)$ consists of all sets of size less than $\frac{n}{2}$ (plus half of all sets of size $\frac{n}{2}$ in the case of even n). This implies the following result. Define Q_n to be

$$\sum_{k=\frac{n+1}{2}}^n \binom{n}{k} p^k (1-p)^{n-k},$$

for odd n , and

$$\frac{1}{2} \binom{n}{n/2} p^{n/2} (1-p)^{n/2} + \sum_{k=\frac{n}{2}+1}^n \binom{n}{k} p^k (1-p)^{n-k},$$

for even n .

Proposition 2.2. *For any diagnosis D working for n processors,*

$$\text{Err}(D) \geq Q_n.$$

It was shown in [11] that the diagnosis strategy from [6], based on majority vote, achieves this bound.

We will use the following version of Chernoff's bound (see [12]).

Proposition 2.3. *Let X be the number of successes in a series of n Bernoulli trials with success probability $r > \frac{1}{2}$. Then*

$$P(X \leq \frac{n}{2}) \leq e^{-cn},$$

for some positive constant c depending on r but not on n .

Taking $r = 1 - p$ in Chernoff's bound we get the following result.

Corollary 2.1. *We have*

$$Q_n \leq e^{-c_0 n},$$

for some positive constant c_0 depending on p but not on n .

The above constant c_0 will be used often hereafter. We use $|S|$ to denote the size of a set S and $\log$ to denote the natural logarithm.

3. Nonadaptive diagnosis

In this section we establish the relation between the reliability of nonadaptive diagnosis and the minimum number of tests it requires. We also study the minimum degree of test assignments that permit q -diagnosis.

Theorem 3.1. *Let $e^{-c_0 n} \leq q < 1$. Then the minimum number of tests required by nonadaptive q -diagnosis for n processors is $\Theta(n \log \frac{1}{q})$.*

Proof. Let $e^{-c_0 n} \leq q < 1$. We first show a nonadaptive q -diagnosis using $O(n \log \frac{1}{q})$ tests. Let $t = \lceil \frac{1}{c_0} \log \frac{1}{q} \rceil$. The assumption $q \geq e^{-c_0 n}$ implies $t \leq n$. Let T be any subset of U of size t . Consider the test assignment consisting of all tests (v, u) , where $v \in T$ and $u \in U$. (This was called a *tester graph* in [6].) Consider the following diagnosis strategy, T -Majority (see [6], [11]). Given a syndrome S , T -Majority outputs the fault set F consisting of those processors u for which $|\{v \in T : S(v, u) = 1\}| > \frac{t}{2}$, that is, processors failed by the strict majority of testers are diagnosed as faulty. As shown in [6] and [11], $\text{Cor}(T\text{-Majority})$ contains all fault sets f of size less than $\frac{t}{2}$. Hence Proposition 2.3 implies

$$\text{Err}(T\text{-Majority}) \leq e^{-c_0 t} \leq e^{-c_0 \frac{1}{c_0} \log \frac{1}{q}} = q,$$

and hence T -Majority is a q -diagnosis. Clearly it uses fewer than tn tests, which is $O(n \log \frac{1}{q})$.

In order to prove the lower bound, assume that a nonadaptive diagnosis algorithm D uses at most nt tests. Hence at least one processor, call it u_0 , is tested by at most t processors $u_1, \dots, u_t$. For every set $F \subset U \setminus \{u_0, u_1, \dots, u_t\}$, the fault sets $F \cup \{u_0, u_1, \dots, u_t\}$ and $F \cup \{u_1, \dots, u_t\}$ are associated, hence at most one of them can be in $\text{Cor}(D)$. The latter fault set has higher probability because $p < \frac{1}{2}$. The family of all fault sets that do not include $\{u_1, \dots, u_t\}$ has probability $1 - p^t$. The family of all fault sets that include $\{u_1, \dots, u_t\}$ but do not contain u_0 has probability $p^t(1 - p)$. Consequently,

$$P(\text{Cor}(D)) \leq 1 - p^t + p^t(1 - p) = 1 - p^{t+1},$$

which implies $\text{Err}(D) \geq p^{t+1}$. Thus any nonadaptive diagnosis using at most nt tests has error probability at least p^{t+1} . It follows that every nonadaptive q -diagnosis must use $\Omega(n \log \frac{1}{q})$ tests. $\square$

The test assignment used to perform the above diagnosis, T -Majority, has the drawback that processors from the set T test all other processors and hence they require $n-1$ incident

communication links in the underlying multiprocessor system. Such interconnections are difficult and costly to implement as n grows and hence it is important to investigate the possibility of q -diagnosis for test assignments of low degree. In the rest of this section we establish the lowest degree (up to a multiplicative constant) of a test assignment that enables q -diagnosis and show a q -diagnosis for test assignments of such degree. We will use the following lemma.

Lemma 3.1. *Let $q < 1$. Then there exists a constant a depending on p such that if the degree of a test assignment for n processors is $t \leq a \log n$ then for any diagnosis D , $\text{Err}(D)$ exceeds q , for sufficiently large n .*

Proof. Take a constant $a < (3 \log \frac{1}{p})^{-1}$. Any processor in the system has at most t neighbours, and hence the set of processors at distance at most 2 of a given processor is at most $1 + t^2$. Construct inductively the following set A of $m = \lfloor \frac{n}{t^2+1} \rfloor$ processors. Choose u_1 arbitrarily. If $u_1, \dots, u_k$ are already constructed, choose as u_{k+1} any processor at distance at least 3 from all $u_1, \dots, u_k$. Thus all processors in A are at distance at least 3 and consequently they have disjoint sets of testers and do not test one another. Let $x = \lfloor m^{1/3} \rfloor$ and $y = x^2$. Let $A_1, \dots, A_x$ be disjoint subsets of A of size y . For a fixed i , the probability of the event that all processors in A_i have a fault-free tester is at most $(1 - p^{a \log n})^y$ because sets of testers of different processors in A are disjoint. Since sets A_i are disjoint, these events are independent for different indices i . Hence the probability that in every set A_i there is a processor all of whose testers are faulty is at least

$$z_n = (1 - (1 - p^{a \log n})^y)^x \geq 1 - x e^{-y p^{a \log n}} \geq 1 - x e^{-y n^{-1/3}},$$

by the choice of a .

Now suppose that in every set A_i there is a processor a_i all of whose testers are faulty. Let B_i be the set of testers of a_i . By construction, the sets $\{a_i\} \cup B_i$ are pairwise disjoint for distinct i . For any fault set F disjoint from all $\{a_i\} \cup B_i$, all fault sets $F \cup \bigcup_{i=1}^x B_i \cup S$, for $S \subset \{a_1, \dots, a_x\}$, are associated and $F \cup \bigcup_{i=1}^x B_i$ has the largest probability among them. It follows that, for any diagnosis strategy D ,

$$P(\text{Cor}(D)) \leq 1 - z_n + z_n(1 - p)^x,$$

and hence

$$P(\text{Err}(D)) \geq z_n(1 - (1 - p)^x) \geq (1 - x e^{-y n^{-1/3}})(1 - e^{-p^x}),$$

which converges to 1 as n grows. □

Theorem 3.2. *Let $e^{-c_0 n/2} \leq q < 1$. Then the minimum degree of a test assignment permitting q -diagnosis for n processors is $\Theta(\log n + \log \frac{1}{q})$.*

Proof. Let $e^{-c_0 n/2} \leq q < 1$. We first show a test assignment of degree $O(\log n + \log \frac{1}{q})$ that permits nonadaptive q -diagnosis. Let $t = \lceil \frac{1}{c_0}(\log n + \log \frac{1}{q}) \rceil$. Thus $t < n$ and $n e^{-c_0 t} \leq q$. Partition all processors into $\lfloor \frac{n}{t} \rfloor$ disjoint sets, of sizes between t and $2t$. Consider the test assignment consisting of all tests (u, v) , where u and v belong to the same set of the

partition. The degree of this test assignment is $O(\log n + \log \frac{1}{q})$. Let Group-Maj be the diagnosis that considers a processor faulty if and only if the majority of its testers consider it faulty. If fault-free processors are a majority in every set of the partition, diagnosis Group-Maj is correct. Hence

$$P(\text{Err}(\text{Group-Maj})) \leq ne^{-c_0 t} \leq q.$$

On the other hand, the lower bound argument from the proof of Theorem 3.1 shows that the degree of a test assignment permitting q -diagnosis must be $\Omega(\log \frac{1}{q})$. Lemma 3.1 shows that this degree must be $\Omega(\log n)$. $\square$

Our next result shows that restricting the degree of test assignments to the least possible order $O(\log n + \log \frac{1}{q})$ increases the number of tests required by nonadaptive q -diagnosis. For a fixed q this minimum number of tests is linear in n if arbitrary test assignments are possible (cf. Theorem 3.1), while it turns out to be $\Theta(n \log n)$ for test assignments of logarithmic degree.

Theorem 3.3. *Let $e^{-c_0 n/2} \leq q < 1$. Then the minimum number of tests in a test assignment of degree $O(\log n + \log \frac{1}{q})$ permitting nonadaptive q -diagnosis is $\Theta(n(\log n + \log \frac{1}{q}))$.*

Proof. The diagnosis Group-Maj from the proof of Theorem 3.2 is a nonadaptive q -diagnosis for a test assignment of degree $O(\log n + \log \frac{1}{q})$ (and thus using $O(n(\log n + \log \frac{1}{q}))$ tests). By Theorem 3.1 the number of tests for any q -diagnosis must be $\Omega(n \log \frac{1}{q})$. It remains to show that the number of tests is $\Omega(n \log n)$.

Fix $q < 1$ and consider a test assignment of degree $t \leq c \log n$, for any constant c , with at most $dn \log n$ tests, for $d < (6 \log \frac{1}{p})^{-1}$. There is a subset B of at least $\frac{n}{2}$ processors that have at most $2d \log n$ testers. Construct the set A as in the proof of Lemma 3.1 choosing processors only from B . Since B has size at least $\frac{n}{2}$, it is possible to construct such a set $A \subset B$ of size $m = \lfloor \frac{n}{2(t^2+1)} \rfloor$. As before, processors in A have disjoint sets of testers and do not test one another. Repeat the rest of the argument from the proof of Lemma 3.1 with $a = 2d$ (d has been chosen to satisfy the condition imposed on a). As before, for any diagnosis strategy D , $P(\text{Err}(D))$ converges to 1 as n grows and hence exceeds q for sufficiently large n . This concludes the proof that the number of tests is $\Omega(n \log n)$. $\square$

4. Adaptive diagnosis

In this section we consider the minimum number of tests required by adaptive q -diagnosis. In the case of adaptive diagnosis this number of tests is not fixed and hence we consider the worst case.

Theorem 4.1. *Let $e^{-c_0 n} \leq q < 1$. Then the minimum number of tests required in the worst case by adaptive q -diagnosis for n processors is $n + \Theta(\log \frac{1}{q})$.*

Proof. Let $e^{-c_0 n} \leq q < 1$. We first show an adaptive q -diagnosis using $n + O(\log \frac{1}{q})$ tests. Let $t = \lceil \frac{1}{c_0} \log \frac{1}{q} \rceil$. The assumption $q \geq e^{-c_0 n}$ implies $t \leq n$. Let T be any subset

of U of size t . Blecher [4] showed an adaptive diagnosis working correctly in a set of t processors under the assumption that fewer than $\frac{t}{2}$ processors are faulty. This diagnosis was proved to use fewer than $\frac{3}{2}t$ tests in the worst case. Call this diagnosis, applied to the set T , T -Adaptive. Our diagnosis algorithm, Ada, works as follows. First run diagnosis T -Adaptive on the set T . If all processors in T are diagnosed as faulty, diagnose all other processors as faulty and stop. Otherwise pick any processor $v \in T$ that was diagnosed as fault-free by T -Adaptive and perform all tests (v, u) , for $u \in U \setminus T$. Diagnose as faulty (fault-free) processors that v considers faulty (fault-free).

The diagnosis Ada uses fewer than $\frac{3}{2}t + n - t$ tests, which is $n + O(\log \frac{1}{q})$. It remains to show that it is indeed a q -diagnosis. If the majority of processors in T are fault-free, T -Adaptive works correctly and finds a fault-free processor (in fact more than $\frac{t}{2}$ such processors). Thus all processors in $U \setminus T$ are also diagnosed correctly and the entire diagnosis Ada is correct. It follows that $\text{Err}(\text{Ada})$ does not exceed the probability that at most $\frac{t}{2}$ processors in T are fault-free. In view of Proposition 2.3 the latter event has probability at most

$$e^{-c_0 t} \leq e^{-c_0 \frac{1}{c_0} \log \frac{1}{q}} = q,$$

hence Ada is indeed a q -diagnosis.

In order to prove the lower bound assume that an adaptive diagnosis algorithm D uses at most $n + t - 2$ tests. Suppose that the first t tests give result 1. Consider the remaining $n - 2$ tests. There exists a processor u_0 not tested by any of them. Hence u_0 was tested by at most t processors $u_1, \dots, u_t$ and all tests were in the first series. For every set $F \subset U \setminus \{u_0, u_1, \dots, u_t\}$, the fault sets $F_1 = F \cup \{u_0, u_1, \dots, u_t\}$ and $F_2 = F \cup \{u_1, \dots, u_t\}$ have a common compatible syndrome S , which gives value 1 on the first t tests. On the remaining tests (u, v) this syndrome gives value 1 if $v \in F$ and value 0 otherwise. Thus at most one of the fault sets F_1, F_2 can belong to $\text{Cor}(D)$. Repeating the argument from the proof of Theorem 3.1 we get $P(\text{Err}(D)) \geq p^{t+1}$. Hence any adaptive diagnosis using at most $n + t - 2$ tests has error probability at least p^{t+1} . It follows that every adaptive q -diagnosis must use $n + \Omega(\log \frac{1}{q})$ tests. $\square$

In a similar way to the nonadaptive case, it is important to construct adaptive q -diagnosis algorithms whose test assignments have low degree and that use as few tests as possible. We present here two algorithms: a simple algorithm that uses the optimal number of tests and whose test assignment has degree $O(\log \frac{1}{q})$, and a more elaborate algorithm whose test assignment has degree $O(1)$ and that uses fewer than $3n$ tests.

Theorem 4.2. *Let $2e^{-c_0 n} \leq q < 1$. Then there exists an adaptive q -diagnosis for n processors with test assignment of degree $O(\log \frac{1}{q})$, and using $n + O(\log \frac{1}{q})$ tests in the worst case.*

Proof. Let $t = \text{MAX}[\frac{1}{c_0} \log \frac{2}{q}, 4] = O(\log \frac{1}{q})$. The assumption $q \geq 2e^{-c_0 n}$ implies $t \leq n$. Partition the n processors into $\ell = O(\log n)$ sets $A_1, \dots, A_\ell$, where $|A_1| = t$, $|A_2| = \frac{1}{2}t^2$, and $|A_i| = (\frac{1}{2})^{i-1}t^i$ ($|A_\ell| \leq (\frac{1}{2})^{\ell-1}t^\ell$).

The probability that the set A_i has a majority of correct processors is (by Proposition 2.3) at least

$$1 - e^{-c_0 \frac{1}{c_0} \log \frac{2}{q} (t/2)^{i-1}} \geq 1 - \left(\frac{q}{2}\right)^i.$$

Thus with probability $1 - q$ all sets have a majority of correct processors.

The algorithm starts by running the Blecher algorithm [4] on the set A_1 . If a majority of processors in the set A_1 are correct, each of the first $t/2$ correct processors in A_1 tests t processors in the set A_2 , otherwise the algorithm fails. If a majority of processors in the set A_i , $2 \leq i \leq \ell - 1$ are found to be correct, each of the first $t^2/2$ correct processors in A_i tests t processors in A_{i+1} , otherwise the algorithm fails.

Clearly the degree of the algorithm is $O(t) = O(\log \frac{1}{q})$, the algorithm fails with probability bounded by q , and the total number of tests is $n + O(\log \frac{1}{q})$. $\square$

Theorem 4.3. *There exists an adaptive q -diagnosis for n processors that uses fewer than $3n$ tests in the worst case, has test assignment of degree $O(1)$, and failure probability $q = e^{-\Omega(n)}$.*

Proof. Since $p < \frac{1}{2}$ is a constant, there is a constant p' such that $p < p' < \frac{1}{2}$ and by the Chernoff bound [12] the probability that the system has more than $p'n$ faulty processors is bounded by

$$q = e^{-\frac{1}{3}(\frac{p'}{p}-1)^2 np} = e^{-\Omega(n)}.$$

Thus, with probability $1 - e^{-\Omega(n)}$, the difference between the number of fault-free processors and faulty processors in the system is at least $2(\frac{1}{2} - p')n$.

The algorithm is based on the following observation.

Proposition 4.1. *Assume that two processors v and u test each other.*

- (1) *If both tests return a non-faulty outcome, then either both processors are fault-free or both are faulty.*
- (2) *If at least one test returns a faulty outcome then at least one of the processors is faulty.* $\square$

The algorithm works in iterations. In each iteration we have a number of *active* sets of processors, and several *left-over* sets. We start the first iteration with n *active* sets, each with one processor; each processor is in exactly one set.

If n is odd we switch the status of one set from *active* to a *left-over* set.

We partition the remaining active sets into $\lfloor \frac{n}{2} \rfloor$ pairs. Assume that the sets $\{v_1\}$ and $\{v_2\}$ are paired; then v_1 tests v_2 , and v_2 tests v_1 . If both tests return a fault-free outcome, the two sets are combined into one active set $\{v_1, v_2\}$. If at least one of the tests has a faulty outcome, the two sets are eliminated. Note that, because of property (1) above, the two processors in a new active set are either both fault-free or both faulty. Property (2) guarantees that the number of fault-free processors that were eliminated is no larger than the number of faulty processors eliminated at that iteration.

Iteration i starts with k_i active sets, each with 2^{i-1} processors. In each active set either all processors are faulty or all processors are fault-free. Furthermore, each active set has two processors that so far have in- and out-degree 1.

If k_i is odd, one active set is marked left-over. The remaining active sets are partitioned into pairs. Given a pair of active sets A_j, A_ℓ , let $v_j \in A_j$ and $v_\ell \in A_\ell$ be two processors with in- and out-degree 1. The two processors test each other. If the outcome of both tests is fault-free, the two sets are combined into one active set of size 2^i . Note that the processors in the new set are either all fault-free or all faulty, and there are two processors in the new set with in- and out-degree 1. If at least one of the tests gives a faulty outcome, the two sets are eliminated. Note that in this case the number of fault-free processors eliminated is bounded by the number of faulty processors eliminated.

The process terminates when there is only one active set left. Since no more than one set was marked left-over in each iteration, and since the size of the active sets doubles in each iteration, the size of the remaining active set A is larger than the sum of the sizes of all the left-over sets. Since we have eliminated as many faulty processors as fault-free processors, and since the majority of processors are fault-free (with probability at least $1 - q$), all the processors in A must be fault-free, and, with probability at least $1 - q$, the set A has more than $(\frac{1}{2} - p')n = cn$ processors for some constant $c > 0$. In the final phase of the algorithm each processor in A tests $\frac{1}{c} = O(1)$ processors outside this set to complete the diagnosis of all the n processors.

The degree of this test assignment is bounded by $2 + \frac{1}{c} = O(1)$, the total number of tests executed by the algorithm is fewer than $3n$ (fewer than $2n$ during the iteration and fewer than n in the final phase), and its failure probability is bounded by $e^{-\Omega(n)}$. $\square$

5. Conclusion

We have established the minimum number of tests (up to a multiplicative constant) needed to diagnose an n -processor system with a given bound on error probability, both for the nonadaptive and adaptive testing scenarios. In both cases we have given concrete algorithms that match these bounds. We have also established the minimum order of magnitude of the degree of test assignments permitting diagnosis with a given bound on error probability. We have investigated test assignments permitting such diagnosis in spite of low degree and few tests. Our results demonstrate that the flexibility yielded by adaptive testing permits a significant saving in the number of tests for the same reliability of diagnosis.

References

- [1] Beigel, R., Hurwood, W. and Kahale, N. (1995) Fault diagnosis in a flash. In *Proc. 36th Symp. on Found. of Comput. Sci.*, pp. 571–580.
- [2] Beigel, R., Kosaraju, S. R. and Sullivan, G. F. (1989) Locating faults in a constant number of testing rounds. In *Proc. 1st Ann. ACM Symp. Par. Alg. and Arch.*, pp. 189–198.
- [3] Beigel, R., Margulis, G. and Spielman, D. A. (1993) Fault diagnosis in a constant number of parallel testing rounds. In *Proc. 5th Ann. ACM Symp. Par. Alg. and Arch.*, pp. 21–29.
- [4] Blecher, P. M. (1983) On a logical problem. *Discrete Math.* **43** 107–110.

- [5] Blough, D. M. and Pelc, A. (1992) Complexity of fault diagnosis in comparison models. *IEEE Trans. Comput.* **41** 318–324.
- [6] Blough, D. M., Sullivan, G. F. and Masson, G. M. (1992) Efficient diagnosis of multiprocessor systems under probabilistic models. *IEEE Trans. Comput.* **41** 1126–1136.
- [7] Blough, D. M., Sullivan, G. F. and Masson, G. M. (1992) Intermittent fault diagnosis in multiprocessor systems. *IEEE Trans. Comput.* **41** 1430–1441.
- [8] Blount, M. (1977) Probabilistic treatment of diagnosis in digital systems. In *Dig. 7th Int. Symp. Fault-Tolerant Computing*, IEEE Computer Society Press, pp. 72–77.
- [9] Dahbura, A. T. (1986) An efficient algorithm for identifying the most likely fault set in a probabilistically diagnosable system. *IEEE Trans. Comput.* **35** 354–356.
- [10] Dahbura, A. T. (1988) System-level diagnosis: A perspective for the third decade. In *Concurrent Computation: Algorithms, Architectures, Technologies*, Plenum Press, New York.
- [11] Diks, K. and Pelc, A. (1997) Globally optimal diagnosis in systems with random faults. *IEEE Trans. Comput.* **46** 200–204.
- [12] Hagerup, T. and Rub, C. (1989/90) A guided tour of Chernoff bounds. *Inform. Proc. Letters* **33** 305–308.
- [13] Hakimi, S. L. and Nakajima, K. (1984) On adaptive system diagnosis. *IEEE Trans. Comput.* **33** 234–240.
- [14] Nakajima, K. (1981) A new approach to system diagnosis. In *Proc. 19th Allerton Conf. Commun. Contr. and Computing*, pp. 697–706.
- [15] Maheshwari, S. N. and Hakimi, S. L. (1976) On models for diagnosable systems and probabilistic fault diagnosis. *IEEE Trans. Comput.* **25** 228–236.
- [16] Pelc, A. (1991) Undirected graph models for system-level fault diagnosis. *IEEE Trans. Comput.* **40** 1271–1276.
- [17] Pelc, A. (1996) Efficient fault location with small risk. In *Proc. 3rd Coll. on Structural Information and Communication Complexity, SIROCCO'96*, pp. 292–300.
- [18] Preparata, F., Metze, G. and Chien, R. (1967) On the connection assignment problem of diagnosable systems. *IEEE Trans. Electron. Comput.* **16** 848–854.
- [19] Rangarajan, S. and Fussell, D. (1988) A probabilistic method for fault diagnosis of multiprocessor systems. In *Dig. 18th Int. Symp. Fault-Tolerant Computing*, IEEE Computer Society Press, pp. 278–283.
- [20] Scheinerman, E. (1987) Almost sure fault-tolerance in random graphs. *SIAM J. Comput.* **16** 1124–1134.
- [21] Schmeichel, E., Hakimi, S. L., Otsuka, M. and Sullivan, G. (1988) On minimizing testing rounds for fault identification. In *Dig. 18th Int. Symp. Fault-Tolerant Computing*, IEEE Computer Society Press, pp. 266–271.