

Correspondence

Optimal and Efficient Probabilistic Distributed Diagnosis Schemes

Sunggu Lee and Kang G. Shin

Abstract—The distributed self-diagnosis of a multiprocessor/multicomputer system based on interprocessor tests with imperfect fault coverage (thus also permitting intermittently faulty processors) is addressed. Focusing on probabilistic diagnosis methods, we define several different categories of probabilistic diagnosis based on the type of fault syndrome information used in the diagnosis. Rigorous probabilistic analysis is then used to derive optimal diagnosis algorithms (optimal in terms of diagnostic accuracy) for the diagnosis categories introduced. Analysis and simulations are used to evaluate the performance of the diagnosis algorithms introduced.

Index Terms—Distributed diagnosis, fault-tolerant computing, intermittent fault, multicomputer, multiprocessor, probabilistic diagnosis, system-level diagnosis.

I. INTRODUCTION

The progress made in the design of powerful single-chip computers (and even single-chip multiprocessors) have led to the construction of increasingly sophisticated multiprocessor/multicomputer systems with literally tens of thousands of processing nodes. In order to maintain a highly reliable system, faulty processing nodes must be identified and periodically removed (either physically or by reconfiguration) from the system.

This correspondence addresses the general problem of the distributed on-line self-diagnosis of processing nodes in a multiprocessor/multicomputer system. The diagnosis will be done on the basis of a *fault syndrome* consisting of a collection of binary pass-fail interprocessor test results. Diagnosis using this type of fault syndrome has its origins in the PMC model [8]. However, unlike most of the system-level diagnosis methods based on the PMC model [4], [7], we do not place an upper bound on the number of permitted faulty processors nor do we assume inter-processor tests with perfect fault coverage. Instead, we focus on probabilistic diagnosis algorithms which can achieve correct diagnosis with high probability given intermittently faulty processing nodes.

Three arguments used to support probabilistic diagnosis algorithms are: 1) using analysis to show that high *diagnostic accuracy*, defined as the percentage of diagnoses which are correct, is achieved in certain situations [5], 2) guaranteeing that the set of nodes most likely to have caused the syndrome is found [3], and 3) showing that as the number of nodes in the system grows to infinity, diagnostic accuracy approaches 100% [1], [9]. While argument 2)—guaranteeing the most probable diagnosis—is the most appealing, it has been shown

Manuscript received November 2, 1989; revised February 1, 1992 and October 11, 1992. This work was supported in part by NSF under grant MIP-9012549, NASA under Grants NAG-1-492 and NAG-1-1220, and by the Korean Research Institute of Science and Technology under Grant R92071.

S. Lee is with the Department of Electronic and Electrical Engineering, Pohang Institute of Science and Technology, Pohang, Korea.

K. G. Shin is with the Real-Time Computing Laboratory, Department of Electrical Engineering and Computer Science, University of Michigan, Ann Arbor, MI 48109.

IEEE Log Number 9207295.

that finding the most probable diagnosis given the global syndrome information is an NP-hard problem [2], [6].

There exist probabilistic diagnosis methods [1], [5] that have the property of asymptotically correct diagnosis *and* produce good diagnosis results for the example network configurations and probability parameters considered by the respective papers.¹ However, it is claimed that this is insufficient for the following reason. The property of asymptotically correct diagnosis simply states that 100% correct diagnosis is achieved as the system size grows to infinity, given certain restrictions on the interconnection network structure. However, under these same restrictions, there are many possible probabilistic diagnosis algorithms that achieve 100% correct diagnosis as the system size grows to infinity [6]. Among these diagnosis algorithms, claiming that a specific diagnosis algorithm performs well given certain network configurations and probability parameters does not provide an adequate comparison to other diagnosis algorithms. In fact, for the networks and probability parameters used in the simulations presented in Section IV of this correspondence, the diagnosis algorithms in [1] and [5] are shown to perform significantly worse (in terms of diagnostic accuracy) than some of the probabilistic diagnosis algorithms introduced here. As a general rule, interprocessor tests with higher fault coverage result in higher diagnostic accuracy. Thus, diagnosis algorithms with higher diagnostic accuracy are desirable not only because of the higher diagnostic accuracy but because, for a given level of diagnostic accuracy, interprocessor tests with lower fault coverage, and thus shorter testing times, are required.

In terms of diagnostic accuracy, the most desirable probabilistic diagnosis algorithm is the *optimal* diagnosis algorithm, which has been shown to be the algorithm that solves the NP-hard problem of producing the most probable diagnosis given the global syndrome information [2]. However, it is noted that in many previous probabilistic diagnosis algorithms with quadratic or lower computational complexity [1], [5], [9], the diagnosis of each processing node is based on only limited partial syndrome information. Given that only partial syndrome information is used in the diagnosis of each node, it is possible to design a diagnosis algorithm that has polynomial computational complexity and is optimal among all diagnosis algorithms that use the same type of partial syndrome information. In this correspondence, diagnosis algorithms are categorized based on the type of syndrome information used in the diagnosis and probability analysis is used to derive an optimal diagnosis algorithm for each category.

II. PRELIMINARIES

A. Notation and Testing Model

A system S is composed of N nodes, denoted by the set $V = \{u_0, \dots, u_{N-1}\}$, where each node $u_i \in V$ is assigned a particular subset of the nodes in V to test. The set of testing assignments in S is represented by a directed graph $G = (V, E)$, called the *testing graph*, where vertex $u_i \in V$ represents a processing node and edge $e_{ij} \in E$ represents the fact that u_i tests u_j . The testing graph is assumed to be a subgraph of the graph representing the interconnection structure of the system. Test outcomes are represented by binary variables a_{ij} such that $a_{ij} = 1$ if u_j fails u_i 's test and $a_{ij} = 0$ if u_j passes u_i 's

¹ Although [5] does not claim that their algorithm produces asymptotically correct diagnosis, it has been proven in [6] that this is the case.

TABLE I
PROBABILITY PARAMETERS

fs_i	fs_j	m	$P(a_{ij} = m fs_i, fs_j)$
$\bar{\delta}_i$	$\bar{\delta}_j$	0	1
$\bar{\delta}_i$	δ_j	1	0
$\bar{\delta}_i$	δ_j	1	p_{ij}
$\bar{\delta}_i$	δ_j	0	$1 - p_{ij}$
δ_i	$\bar{\delta}_j$	0	r_{ij}
δ_i	$\bar{\delta}_j$	1	$1 - r_{ij}$
δ_i	δ_j	1	s_{ij}
δ_i	δ_j	0	$1 - s_{ij}$

test. a_{ij} is undefined if u_i does not test u_j . A (fault) syndrome SD is a function from E to $\{0, 1\}$. The function SD is defined such that for all $e_{ij} \in E$, $SD(e_{ij}) = a_{ij}$. Given a syndrome SD , a diagnosis is said to be *correct* if the set of nodes diagnosed to be faulty is the same as the actual fault set.

The set of nodes that a given node u_i tests will be denoted by $\Gamma(u_i)$. Likewise, the set of nodes that test u_i will be denoted by $\Gamma^{-1}(u_i)$ and $\Gamma^{-1}(u_i) = \Gamma_1^{-1}(u_i) \cup \Gamma_0^{-1}(u_i)$, where $\Gamma_k^{-1}(u_i) = \{u_j \in \Gamma^{-1}(u_i) : a_{ji} = k\}$, $k = 0, 1$. The *one-condensation* of G , denoted by $G_1 = (V_1, E_1)$, is the subgraph of G with $E_1 = \{e_{ij} \in E : a_{ij} = 1\}$ and $V_1 = \{u_i \in V : u_i \text{ is the endpoint of an edge in } E_1\}$. Given a node $u_i \in V$, $d(u_i) = \{u_j \in \Gamma^{-1}(u_i) : a_{ji} = 1\}$. In discussing probabilities of events where the set of basic events is clear, we will use the notation $P(x)$ to denote the probability of the event x .

In our testing model, the result of a nonfaulty node testing another nonfaulty node is the only completely reliable test result. f_i will denote the prior fault probability of u_i . The probability parameters that describe the possible values of a_{ij} given different fault statuses of u_i and u_j are given in Table I. The fault status of u_k (fs_k) is denoted by δ_k (for u_k is faulty) and $\bar{\delta}_k$ (for u_k is nonfaulty). As an example, p_{ij} is the probability that a nonfaulty node u_i will correctly diagnose a faulty node u_j . Average values of probability parameters will be denoted by the corresponding letters without subscripts. Thus, for example, f and p will refer to the average f_i and p_{ij} values, respectively.

B. Categorization of Probabilistic Diagnosis Algorithms

We can classify methods for probabilistic diagnosis based on the amount of information used to identify faulty nodes. Looking at the diagnosis algorithm executed on each node, the maximum amount of information that can be used by each node is the entire syndrome. A probabilistic diagnosis method which uses the entire syndrome is defined as a *category 1 probabilistic diagnosis* method. The optimal diagnosis algorithm in this category is the algorithm which finds the most probable fault set given the syndrome. This diagnosis algorithm was used in [3]. This category of diagnosis is inherently inefficient because the entire syndrome must be reliably communicated to each node.

Suppose instead that each node is aware of only the part of the syndrome that directly implicates it as faulty or nonfaulty, i.e., each node u_i is only aware of the a_{ji} values such that $u_j \in \Gamma^{-1}(u_i)$. This is referred to as *local syndrome information*. It is possible to summarize this local syndrome information as the ordered pair $(d(u_i), |\Gamma^{-1}(u_i)|)$, where $d(u_i)$ is as defined above. In *category 2 probabilistic diagnosis*, faulty nodes are identified *one at a time*, using local syndrome information and the identity of the nodes diagnosed as faulty in previous steps. Thus, since the calculations of all of the nodes have to be examined to determine the order in which faulty nodes are identified, a summarized form of the global

TABLE II
CATEGORIES OF PROBABILISTIC DIAGNOSIS AND SYNDROME INFORMATION USED

Category	Syndrome Information Used
1	All
2	Local and nodes previously diagnosed as faulty
2A	Summarized local and nodes previously diagnosed as faulty
3	Local
3A	Summarized local

syndrome is effectively being used in the diagnosis. This method approximates the category 1 method by looking at individual nodes instead of subsets of nodes. *Category 2A probabilistic diagnosis* is similar to category 2 probabilistic diagnosis except that summarized local syndrome information is used. The diagnosis method of [1], [5] fall into this latter category. In *category 3 probabilistic diagnosis*, local syndrome information is used to identify each node as faulty or nonfaulty independently of the other nodes in the system. In *category 3A probabilistic diagnosis*, summarized local information is used to diagnose each node. The diagnosis method of [1] can be easily modified to fall into this category. Table II summarizes this categorization.

C. Probability Model and Previous Methods

To define a probability model, we need to define a probability space, which is a triple (Ω, Θ, P) , where Ω is the sample space, Θ is the event space, and P is a probability measure. However, for each category of probabilistic diagnosis, the type of syndrome information used in the diagnosis is different. Thus, for each category of probabilistic diagnosis, we use a different probability space in talking about the probability of certain types of syndromes and fault sets being present.

Let us consider category x probabilistic diagnosis, where x can be any of the categories defined above. Suppose the syndrome information used in diagnosing node $u_i \in V$ is denoted by SD_i . SD_i is a restricted form of the information present in the syndrome SD . For a given node u_i , its fault status set is defined as $Status_Set_i = \{\delta_i, \bar{\delta}_i\}$. Let us consider an arbitrary category x probabilistic diagnosis algorithm A . In order for the diagnosis by A to be correct, A 's diagnosis of each node $u_i \in V$ must be correct. Let A_i denote the part of A which diagnoses the fault status of node u_i , and define as basic events the pairs (SD_i, fs_i) , where SD_i is a "partial syndrome" and $fs_i \in Status_Set_i$. The set of all possible SD_i 's will be denoted by SD_i^{all} . The diagnosis of u_i by A when executed on a syndrome containing the partial syndrome SD_i is denoted by $Diag_{A_i}(SD_i) \in Status_Set_i$.

For category x probabilistic diagnosis on node u_i , the sample space $\Omega_i^x = \{(SD_i, fs_i) : SD_i \in SD_i^{all}, fs_i \in Status_Set_i\}$, the event space Θ_i^x is all possible subsets of Ω_i^x , and the probability measure P_i^x is defined for category x probabilistic diagnosis such that it is a legitimate probability measure. Although not explicit in our notation, the probability measure P_i^x is also dependent on the testing graph $G = (V, E)$. Given a testing graph G and a diagnosis algorithm A , let $Correct_G(A_i) = \{(SD_i, fs_i) : Diag_{A_i}(SD_i) = fs_i\}$. For a testing graph G , the probability of correct diagnosis of u_i by A is

$$P_i^x(Correct_G(A_i)) = \sum_{SD_i \in SD_i^{all}} P_i^x(SD_i, Diag_{A_i}(SD_i)). \quad (1)$$

There are a few diagnosis algorithms in the current literature which are closely related to the diagnosis algorithms that we present in this correspondence. The diagnosis algorithm that always finds the

most probable fault set given a syndrome is referred to as the most probable diagnosis (MPD) algorithm [3]. Since the MPD algorithm results in the highest achievable level of diagnostic accuracy, it is used to evaluate the performance of other diagnosis algorithms. Blough *et al.*'s (BSM) diagnosis algorithm [1] and Dahbura *et al.*'s (DSK) diagnosis algorithm [5] are category 2A probabilistic diagnosis algorithms. The DSK* algorithm is an improved version of the DSK algorithm presented in [6]. The MPD, BSM, DSK, DSK* diagnosis algorithms will be compared to the diagnosis algorithms presented in this correspondence and to each other. Although the algorithm presented in [9] is also an interesting probabilistic diagnosis algorithm, it is based on the use of "multiple syndromes" and is not directly comparable to the above algorithms.

The BSM algorithm [1] is an approximation algorithm for the method used to diagnose mixtures of permanent and intermittent faults in [4]. This algorithm compares $d(u_i)$ to a threshold value κ_i , which is the assumed maximum number of faulty nodes testing u_i . All nodes with $d(u_i) > \kappa_i$ are diagnosed to be faulty. Converting the formula for κ_i in [1] into our testing model and using our notation:

$$\kappa_i = \frac{1}{2} \sum_{u_j \in \Gamma^{-1}(u_i)} [f_j + p_{ji}(1 - f_j)]. \quad (2)$$

In the DSK algorithm [5], the one-condensation of the testing graph G is first copied into a graph G' . $d'(u_i)$ is set to the value of $d(u_i)$ in the new graph G' . The DSK algorithm repeatedly selects and removes from the graph G' a node with the greatest number of incident one-links ($d'(u_i)$ value) until no one-links remain in G' . Ties in $d'(u_i)$ values are broken arbitrarily.

The DSK* algorithm improves on the DSK algorithm by selecting the node u_k with the highest f_k (prior fault probability) value from among the nodes u_m with the maximum value of $d'(u_m)$. A total ordering is imposed on the nodes to break any ties that exist after comparing f_i values. In [6], we showed that DSK* performs significantly better than DSK when the relative magnitudes of f_i values are known. DSK is $O(|E|)$ and DSK* is $O(N^2)$.

III. OPTIMAL PROBABILISTIC DIAGNOSIS ALGORITHMS

In this section we use probability analysis to derive optimal diagnosis algorithms for category 2, 2A, 3, and 3A probabilistic diagnosis. Each category of probabilistic diagnosis is defined by the type of syndrome information used in the diagnosis. Given a certain type of syndrome information, the optimal diagnosis algorithm is to make the most probable diagnosis at each node. This was shown formally for category 1 probabilistic diagnosis by Blough [2]. We now show this for general categories of probabilistic diagnosis. Suppose we are working with category x probabilistic diagnosis. Let OPT_x denote the algorithm which makes the most probable diagnosis for each node u_i given SD_i .

Theorem 1: For any category x probabilistic diagnosis algorithm A , $P(\{A \text{ produces correct diagnosis}\}) \leq P(\{OPT_x \text{ produces correct diagnosis}\})$.

Proof: Consider a testing graph G and an arbitrary category x probabilistic diagnosis algorithm A . A produces correct diagnosis if and only if the diagnosis of each node is correct. Let A_i denote the

diagnosis of node u_i by A . For an arbitrary node $u_i \in V$,

$$\begin{aligned} P_i^x(\text{Correct}_G(A_i)) &= \sum_{SD_i \in SD_i^{all}} P_i^x(SD_i, \text{Diag}_{A_i}(SD_i)) \\ &= \sum_{SD_i \in SD_i^{all}} P_i^x(\text{Diag}_{A_i}(SD_i) | SD_i) P_i^x(SD_i) \\ &\leq \sum_{SD_i \in SD_i^{all}} P_i^x(\text{Diag}_{OPT_x}(SD_i) | SD_i) P_i^x(SD_i) \\ &= P_i^x(\text{Correct}_G(OPT_x)). \end{aligned}$$

Since $P_i^x(\text{Correct}_G(A_i)) \leq P_i^x(\text{Correct}_G(OPT_x))$ for all $u_i \in V$, the theorem follows. Q.E.D.

The strategy used for deriving an optimal diagnosis algorithm is based on the calculation of the posterior fault probability for each node given the syndrome information for that category of probabilistic diagnosis. For diagnosis category x , each node u_i is diagnosed as faulty or nonfaulty depending on the fault status that results in a higher probability measure P_i^x . As shown in the proof of Theorem 1, the probability of correct diagnosis of node u_i is the sum of $P_i^x(SD_i, \text{Diag}_{A_i}(SD_i))$, where the sum is over all possible partial syndromes SD_i and $\text{Diag}_{A_i}(SD_i)$ is the diagnosis of u_i by an algorithm A (one of $\bar{\delta}_i$ or δ_i). Thus, if algorithm A 's diagnosis for partial syndrome SD_i is the one that results in the highest value of P_i^x for all syndromes SD_i , then the probability of correct diagnosis of node u_i is maximized.

Several assumptions are made in our probability analysis. It is assumed that the probability parameter values of different nodes are independent. f_i, p_{ij}, r_{ij} , and s_{ij} probability parameter values are assumed to be greater than 0 and less than 1. The probability of a faulty node having a one-link incident on it is assumed to be greater than the probability of a nonfaulty node having a one-link incident on it. This last assumption is required for any probabilistic diagnosis method to work.

A. Diagnosis Using Purely Local Syndrome Information

Given an arbitrary node $u_i \in V$ and a node $u_j \in \Gamma^{-1}(u_i)$, let $A_{ji} = P(a_{ji} = 1 | \delta_i)$ and $B_{ji} = P(a_{ji} = 1 | \bar{\delta}_i)$. It is assumed that $A_{ji} > B_{ji}$. Using the parameters of our testing model

$$\begin{aligned} A_{ji} &= (1 - f_j)p_{ji} + f_j s_{ji} \\ B_{ji} &= f_j(1 - r_{ji}). \end{aligned}$$

Let the local syndrome information used by node $u_i \in V$ be denoted by LS_i . Then we can define our probability measure P_i^3 as

$$\begin{aligned} P_i^3(LS_i, \delta_i) &= f_i \prod_{u_j \in \Gamma_1^{-1}(u_i)} A_{ji} \prod_{u_j \in \Gamma_0^{-1}(u_i)} (1 - A_{ji}) \\ P_i^3(LS_i, \bar{\delta}_i) &= (1 - f_i) \prod_{u_j \in \Gamma_1^{-1}(u_i)} B_{ji} \prod_{u_j \in \Gamma_0^{-1}(u_i)} (1 - B_{ji}). \end{aligned}$$

It can be checked that P_i^3 is a legitimate probability measure. It follows that

$$\begin{aligned} P_i^3(\delta_i | LS_i) &= \frac{P_i^3(LS_i, \delta_i)}{P_i^3(LS_i, \delta_i) + P_i^3(LS_i, \bar{\delta}_i)} \\ &= \frac{1}{1 + \frac{1-f_i}{f_i} \prod_{u_j \in \Gamma_1^{-1}(u_i)} \left(\frac{B_{ji}}{A_{ji}}\right) \prod_{u_j \in \Gamma_0^{-1}(u_i)} \left(\frac{1-B_{ji}}{1-A_{ji}}\right)}. \quad (3) \end{aligned}$$

To make the most probable diagnosis for u_i based on u_i 's local syndrome information, u_i must be diagnosed to be faulty if and only if (3) > 0.5 . This follows from the fact that $P_i^3(\bar{\delta}_i|LS_i) = 1 - P_i^3(\delta_i|LS_i)$. The most probable diagnosis for u_i is the diagnosis that results in a higher value of P_i^3 . Thus, the optimal category 3 probabilistic diagnosis algorithm is given as follows.

Algorithm OPT3:

For all nodes $u_i \in V$, do

1. calculate posterior fault probability for u_i using (3);
2. if $P(\delta_i|LS_i) > 0.5$, then label u_i as faulty;
otherwise, label u_i as nonfaulty;

In category 3A probabilistic diagnosis, the strategy is to determine the fault status of each node based on summarized local syndrome information, i.e., the number of neighbors of the node that test it to be faulty. The main advantages of using summarized local syndrome information (instead of local syndrome information) are that the resulting diagnosis algorithms are simpler, less dependent on the accuracy of probability parameter values, and implementable as constant-time distributed algorithms. In the probability analysis for the category 3A probabilistic diagnosis on an arbitrary node $u_i \in V$, we shall assume that the testing graph is regular (has constant node-degree) and that the average probability parameter values of nodes in $\Gamma^{-1}(u_i)$ are being used. Then, going through a similar process as the previous analysis, we can write the posterior fault probability of u_i given $z = d(u_i)$ one-links directed into u_i out of a maximum of $\gamma = |\Gamma^{-1}(u_i)|$ (denoted by z one-links: γ) as

$$P_i^{3A}(\delta_i|z \text{ one-links: } \gamma) = \frac{1}{1 + \frac{1-f_i}{f_i} \left(\frac{B}{A}\right)^z \left(\frac{1-B}{1-A}\right)^{\gamma-z}}, \quad (4)$$

where $A = (1-f)p + fs$ and $B = f(1-r)$.

Implicit in the above analysis is the fact that the partial syndrome information used in category 3A probabilistic diagnosis is denoted by z one-links: γ (note that $z = d(u_i)$). Also, although we have not yet formally defined the probability measure P_i^{3A} it is easy to see what $P_i^{3A}(z \text{ one-links: } \gamma, \delta_i)$ and $P_i^{3A}(z \text{ one-links: } \gamma, \bar{\delta}_i)$ must be by referring to the analysis for P_i^3 and (4). A similar procedure will be used in describing category 2 and 2A probabilistic diagnosis (the notation for the partial syndrome information used and the definition of the formal probability measure will be obvious from the discussion).

Equation (4) is an increasing function of z since $A > B$. By setting (4) equal to $\frac{1}{2}$ and solving for z , we can determine a threshold value z_{th_i} . When $z > z_{th_i}$ ($z < z_{th_i}$), node u_i is likely to be faulty (fault-free) since $P(\delta_i|z \text{ one-links: } \gamma) > 0.5$ ($P(\bar{\delta}_i|z \text{ one-links: } \gamma) > 0.5$). Solving for z_{th_i} ,

$$z_{th_i} = \frac{\log\left(\frac{1-f_i}{f_i}\right)}{\log\left(\frac{A(1-B)}{(1-A)B}\right)} + \gamma \frac{\log\left(\frac{1-B}{1-A}\right)}{\log\left(\frac{A(1-B)}{(1-A)B}\right)}. \quad (5)$$

This results in the following optimal category 3A probabilistic diagnosis algorithm. The optimality of this algorithm follows from the monotonically increasing property of (4) and Theorem 1.

Algorithm OPT3A:

For all nodes $u_i \in V$, do

1. calculate z_{th_i} using (5);
2. if $d(u_i) > z_{th_i}$, then label u_i as faulty;
otherwise, label u_i as nonfaulty.

B. Category 2 and 2A Probabilistic Diagnosis

The first node that is identified to be faulty should be the one with the highest posterior fault probability. Once the first node is identified to be faulty on the basis of (3), we must identify the node with the next highest posterior fault probability. However, we can use the fact that one node has already been identified to be faulty to update the posterior fault probability of adjacent nodes. Suppose u_k has previously been identified to be faulty. Let $u_i \in \Gamma(u_k)$ be an arbitrary as-yet undiagnosed node. Then, from Table I, we know that $A_{ki} = P(a_{ki} = 1|\delta_i) = s_{ki}$ and $B_{ki} = P(a_{ki} = 1|\bar{\delta}_i) = 1 - r_{ki}$. In general, if the nodes in $H_{i1} \subseteq \Gamma_1^{-1}(u_i)$ and $H_{i0} \subseteq \Gamma_0^{-1}(u_i)$ have previously been identified to be faulty, then the result is (6), which is shown at the bottom of the page.

Relying purely on probabilistic information, the process of identifying faulty nodes should stop when there does not exist any node with a posterior fault probability greater than 0.5. This can result in a fault set F which is not a vertex cover of G_1 , the one-condensation of the testing graph G . A fault set which is not a vertex cover of G_1 can not have produced the syndrome for which the diagnosis is made. Thus, a better stopping condition is to continue identifying faulty nodes until the resulting fault set is a vertex cover of G_1 . (In effect, we are using a little bit more syndrome information.) By careful analysis [6], it can be shown that except under extremely extraordinary circumstances, the most probable fault set is a subset of V_1 , the node set of G_1 . The fault set found by using (6) and this stopping condition is also typically a subset of V_1 . The following is the optimal category 2 probabilistic diagnosis algorithm.

Algorithm OPT2:

0. Let $F \leftarrow \emptyset$ be the set of diagnosed faulty nodes;
let $G_2 = (V_2, E_2)$ be a copy of G_1 ;
1. **For all $u_i \in V$ do**
—use (6) with $H_{i1} = H_{i0} = \emptyset$ to calculate posterior fault probabilities;
2. **While $E_2 \neq \emptyset$, do**
 - 2a. Let u_k be the node with highest posterior fault probability;
 - 2b. $F \leftarrow F \cup \{u_k\}$;
 - 2c. **For all nodes $u_j \in \Gamma(u_k)$ do**
—update posterior fault probability of u_j using (6);
 - 2d. Update E_2 by removing all links to and from u_k ;

The main advantage in going from a category 2 to a category 2A probabilistic diagnosis method is that the resulting diagnosis is less susceptible to inaccurate probability parameter values. In category 2A probabilistic diagnosis, the first node to be identified as faulty is the node with the highest posterior fault probability as calculated using (4). But then, the equation for updating the posterior fault probability of nodes adjacent to previously identified faulty becomes similar to (6). If the nodes in $H_{i1} \subseteq \Gamma_1^{-1}(u_i)$ and $H_{i0} \subseteq \Gamma_0^{-1}(u_i)$ have

$$P_i^2(\delta_i|LS_i, H_{i1}, H_{i0}) = \frac{1}{1 + \frac{1-f_i}{f_i} \prod_{u_j \in \Gamma_1^{-1}(u_i) - H_{i1}} \left(\frac{B_{ji}}{A_{ji}}\right) \prod_{u_j \in \Gamma_0^{-1}(u_i) - H_{i0}} \left(\frac{1-B_{ji}}{1-A_{ji}}\right) \prod_{u_j \in H_{i1}} \left(\frac{1-r_{ji}}{s_{ji}}\right) \prod_{u_j \in H_{i0}} \left(\frac{r_{ji}}{1-s_{ji}}\right)}. \quad (6)$$

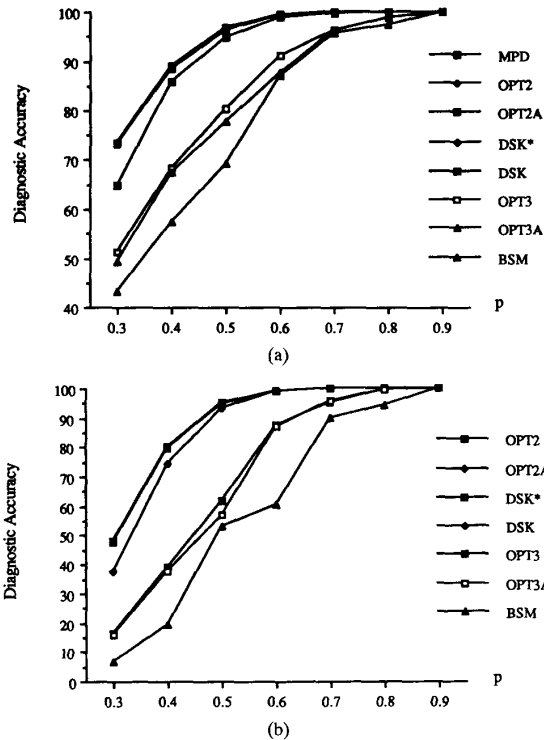


Fig. 1. Diagnostic accuracy on (a) Q_8 and (b) Q_{10} with MTTF = 50K hours.

previously been identified to be faulty, then the result is (7), shown at the bottom of the page. Thus, Algorithm OPT2A is the same as Algorithm OPT2 with (6) replaced by (7).

IV. SIMULATION RESULTS

Simulations were conducted to evaluate the performance of the diagnosis algorithms studied. However, since the diagnosis algorithms introduced in this correspondence were designed to be *optimal* (in diagnostic accuracy) for each category of probabilistic diagnosis studied, the results of simulations are not as important as they would be for heuristic diagnosis algorithms. Detailed analysis presented in [6] also shows that the diagnosis algorithms presented in this correspondence perform well even with highly inaccurate probability parameter value estimates.

Numerous simulations were conducted on a Sun 3/280 for hypercubes of dimension six through ten using various sets of probability parameter values. Each node in the hypercube was assigned to test each of its immediate neighbors. Let Q_r denote a hypercube of dimension r . The diagnosis algorithm MPD was only executed on hypercubes Q_6 through Q_8 with MTTF = 50 000 hours because of its high computational cost. Fig. 1(a) and (b) shows the results of the simulations on a Q_8 and Q_{10} with MTTF = 50 000 hours.

TABLE III
DIAGNOSTIC ACCURACY WITH RESPECT TO
MPD FOR Q_8 WITH MTTF = 50 000 HOURS

p	0.3	0.4	0.5	0.6	0.7	0.8	0.9
DSK	-11.4%	-3.6%	-2.1%	-0.6%	-0.2%	0.0%	0.0%
DSK*	-0.1%	-0.6%	-0.2%	0.0%	0.0%	0.0%	0.0%
OPT2A	-0.2%	-0.4%	-0.5%	0.0%	0.0%	0.0%	0.0%
OPT2	+0.1%	+0.1%	0.0%	0.0%	0.0%	0.0%	0.0%

Similar results were obtained for all hypercube dimensions and MTTF values used. Table III shows the relative performance of the DSK, DSK*, OPT2A, and OPT2 algorithms with respect to the globally optimal algorithm MPD for a Q_8 with MTTF = 50 000 hours.

V. CONCLUSION

The main contributions of this correspondence are the categorization of probabilistic diagnosis algorithms according to the type of syndrome information used in the diagnosis of each node and the development of *optimal* diagnosis algorithms for each category. This categorization is significant because the communication requirements for distributed self diagnosis are different for each category and because existing diagnosis algorithms fit well into the categories defined. For all categories except category 1 (which uses global syndrome information), these optimal diagnosis algorithms are computationally efficient algorithms with $O(N^2)$ or lower computational complexity.

REFERENCES

- [1] D. M. Bough, G. F. Sullivan, and G. M. Masson, "Almost certain diagnosis for intermittently faulty systems," in *Dig. Papers, FTCS-18*, 1988, pp. 260-265.
- [2] D. M. Blough, "Fault detection and diagnosis in multiprocessor systems," Ph.D. dissertation, Johns Hopkins Univ., Baltimore, MD, 1988.
- [3] M. L. Blount, "probabilistic treatment of diagnosis in digital systems," in *Dig. Papers, FTCS-7*, 1977, pp. 72-77.
- [4] A. T. Dahbura and G. M. Masson, "Greedy diagnosis of hybrid fault situations," *IEEE Trans. Comput.*, vol. C-32, no. 8, pp. 777-782, Aug. 1983.
- [5] A. T. Dahbura, K. K. Sabnani, and L. L. King, "The comparison approach to multiprocessor fault diagnosis," *IEEE Trans. Comput.*, vol. C-36, no. 3, pp. 373-378, Mar. 1987.
- [6] S. Lee, "Probabilistic multiprocessor and multicomputer diagnosis," Ph.D. dissertation, Univ. Michigan, Ann Arbor, 1990.
- [7] S. Mallela and G. M. Masson, "Diagnosable systems for intermittent faults," *IEEE Trans. Comput.*, vol. C-27, no. 6, pp. 560-566, June 1978.
- [8] F. P. Preparata, G. Metzger, and R. T. Chien, "On the connection assignment problem of diagnosable systems," *IEEE Trans. Electron. Comput.*, vol. EC-16, no. 6, pp. 848-854, Dec. 1967.
- [9] S. Rangarajan and D. Fussell, "Diagnosing arbitrarily connected parallel computers with high probability," *IEEE Trans. Comput.*, vol. 41, no. 5, pp. 606-615, May 1992.

$$P_i^{2A}(\delta_i | z \text{ one-links} : \gamma, H_{i1}, H_{i0}) = \frac{1}{1 + \frac{1-f_i}{f_i} \left(\frac{B}{A}\right)^{z-|H_{i1}|} \left(\frac{1-B}{1-A}\right)^{\gamma-z-|H_{i0}|}} \prod_{u_j \in H_{i1}} \left(\frac{1-r_{ji}}{s_{ji}}\right) \prod_{u_j \in H_{i0}} \left(\frac{r_{ji}}{1-s_{ji}}\right) \quad (7)$$