OPTIMAL MULTIPLE SYNDROME PROBABILISTIC DIAGNOSIS*

Sunggu Lee† and Kang G. Shin‡

† Department of Electrical Engineering
The University of Delaware
Newark, DE 19716.

‡ Real-Time Computing Laboratory
Department of Electrical Engineering and Computer Science
The University of Michigan
Ann Arbor, MI 48109-2122.

ABSTRACT

This paper addresses the distributed self-diagnosis of a multiprocessor/multicomputer system based on inter-processor tests with imperfect fault coverage (thus also permitting intermittently faulty processors). We show that by using multiple fault syndromes, it is possible to achieve significantly better diagnosis than by using a single fault syndrome, even when the amount of time devoted to testing is the same. We derive a multiple syndrome diagnosis algorithm which is *optimal* in the level of diagnostic accuracy achieved (among diagnosis algorithms of a certain type to be defined) and produces good results even with sparse interconnection networks and inter-processor tests with low fault coverage. Furthermore, we prove upper and lower bounds on the number of fault syndromes required to asymptotically produce 100% correct diagnosis as $N \rightarrow \infty$. Our solution and another multiple syndrome diagnosis solution by Fussell and Rangarajan [6] are evaluated both analytically and with simulations.

1. INTRODUCTION

This paper addresses the problem of the distributed diagnosis of faulty processors in a multiprocessor or multicomputer system. The diagnosis is done on the basis of a *fault syndrome* consisting of a collection of binary pass-fail inter-processor test results as in the PMC model [10]. However, unlike most of the system-level diagnosis methods based on the PMC model, 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. In addition, for ease of implementation, we assume that processors test one another by comparing the outputs of identical tasks (commonly referred to as *comparison-testing*). Since inter-processor tests with imperfect fault coverage and intermittently faulty processors can result in the same types of fault syndromes, we can handle intermittent as well as permanent faults.

Several authors [8,9,12] addressed the problem of diagnosing intermittent faults in t_i–diagnosable systems, in which if no more than t_i nodes are intermittently faulty, a non-faulty node will never be diagnosed as faulty [8]. However, because a node is identified as faulty only if there is sufficient evidence to definitely identify it as faulty given the upper bound t_i on the number of faulty nodes, these methods rarely achieve correct diagnosis (by which we mean that the diagnosed fault set is the same as the actual fault set).

Other authors [1,3,5,6] proposed probabilistic diagnosis algorithms which achieve correct diagnosis with high probability given intermittently faulty processors. Such probabilistic diagnosis algorithms offer the most general solutions with the highest level of *diagnostic accuracy*, defined as the percentage of diagnoses which are correct. Three arguments used to support probabilistic diagnosis algorithms are: (1) using analysis to show that high diagnostic accuracy 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 systems grows to infinity, diagnostic accuracy approaches 100% [1,6]. While argument (2), guaranteeing the most probable diagnosis, is the most appealing, it has been shown that finding the most probable diagnosis given the global syndrome information is an NP-hard problem [2,7]. From a practical perspective, argument (3) is insufficient as good diagnostic accuracy is desired for finite systems. However, since automated diagnosis is particularly important for large systems, asymptotically correct diagnosis is certainly a desirable property of any probabilistic diagnosis algorithm. In this paper, a diagnosis algorithm is presented which has the property of asymptotically correct diagnosis and is *optimal* (in diagnostic accuracy) among all diagnosis algorithms of a certain type to be defined. Our diagnosis algorithm is further supported by simulations using square mesh and TMR structures.

Blough *et. al.* [1] showed that they could asymptotically achieve 100% correct diagnosis in an N processor system as $N \rightarrow \infty$ provided that $\alpha(N) \log N$ tests were performed on each processor, where $\alpha(N) \rightarrow \infty$ arbitrarily slowly as $N \rightarrow \infty$. In their method, the number of tests on processor u_i is equivalent to the number of processors testing u_i. Fussell and Rangarajan [6] improved on [1] by showing that the same asymptotic result can be obtained for systems with lower connectivity (e.g., meshes or rings) if each pair of processors conducts multiple tests and the number of *these* tests on each processor grows faster than $\log N$. Fussell and Rangarajan's algorithm can be viewed as a *multiple syndrome diagnosis* algorithm, in which testing is conducted in stages and a fault syndrome is collected after each testing stage. Our analysis in Section 3 shows that the multiple syndrome diagnosis method [6] can achieve significantly higher diagnostic accuracy than diagnosis using a single syndrome, even when the time devoted to testing is the same.

In this paper, we improve upon Fussell and Rangarajan's (FR) algorithm [6] by deriving a multiple syndrome diagnosis algorithm which is *optimal* in the level of diagnostic accuracy achieved. Since multiple syndromes can be formed in many different ways and since many different types of syndrome information can be used in the diagnosis, we define a specific category of multiple syndrome diagnosis (of which the FR algorithm is a member) and restrict our analysis to this category. Our diagnosis algorithm is *provably optimal* among all multiple syndrome diagnosis algorithms which use the same type of syndrome information as the FR algorithm. In addition, our optimal multiple syndrome diagnosis algorithm has the same desirable asymptotic properties as the FR algorithm. Upper and lower bounds on the number of tests required for asymptotically correct diagnosis are derived.

The rest of this paper is organized as follows. Preliminaries and a description of the FR algorithm are given in Section 2. In Section 3, we use probability analysis to show that significantly higher diagnostic accuracy can be achieved by restricting the manner in which inter-processor comparison tests are conducted. Our optimal multiple syndrome diagnosis algorithm is derived in Section 4. Section 5 presents the asymptotic analysis for our diagnosis algorithm. In Section 6, we use simulations to provide a comparison of the FR algorithm, our optimal multiple syndrome diagnosis algorithm, and an optimal single syndrome diagnosis algorithm [7]. We conclude with Section 7.

2. BACKGROUND

2.1. Preliminaries

A system S is composed of N processing nodes, denoted by the set $V = \{u_0, \cdots, u_{N-1}\}$. Inter-processor testing is assumed to be

CH 2877-9/90/0000/0324/$01.00 – 1990 IEEE

done by comparison-testing, in which a test between two processors u_i and u_j is actually a comparison of their execution results/outputs for an identical task. The set of tests executed in this manner is represented by an undirected graph $G = (V, E)$, called the *testing graph*, where vertex $u_i \in V$ represents a processing node and undirected edge $e_{ij} \in E$ represents the fact that a comparison test is performed between u_i and u_j. The testing graph is assumed to be a subgraph of the graph representing the interconnection network of the system. A (fault) *syndrome SD* is a function from E to $\{0, 1\}$. $SD(e_{ij})$ is denoted by a_{ij} and is equal to 1 if and only if nodes u_i and u_j produce different execution results for the same task. A diagnosis is said to be *correct* if the set of nodes diagnosed to be faulty is the same as the actual fault set. For a given node u_i, let $\Gamma(u_i)$ denote the set of nodes that u_i tests (which are the same nodes that test u_i), and let $d(u_i) = |\{u_j \in \Gamma(u_i) : a_{ij} = 1\}|$. The fault status of u_i is denoted by δ_i for "u_i is faulty" and $\bar{\delta}_i$ for "u_i is non-faulty". The actual set of faulty nodes which are to be diagnosed is denoted by F'.

The following probability parameters are used. Given a node $u_i \in V$, f_i is the prior fault probability of u_i. Given a node $u_i \in V$ and a test task t_k, p_{ik} is the probability that u_i produces an incorrect result for task t_k given that u_i is faulty. p_{ik} values will be referred to as *fault coverage*. In part of the probability analysis, we will require the use of average probability parameter values. The average values of parameters will be denoted by the corresponding letters without subscripts. For example, f and p will refer to the average f_i and p_{ik} values, respectively.

The testing methods used in single and multiple syndrome diagnoses are referred to as *single syndrome testing* and *multiple syndrome testing*, respectively. Most of the previous work on diagnosis based on comparison-testing have assumed single syndrome testing. In single syndrome testing, it is assumed that the comparison test between a node u_i and another node $u_k \in \Gamma(u_i)$ is independent of any other comparison test. If two nodes u_i and u_j execute and compare more than one task, then $a_{ij} = a_{ji} = 1$ if any of the task outputs are different for the two nodes. The syndrome formed in this manner is the one used by a single syndrome diagnosis algorithm

In multiple syndrome testing, testing is done in stages, and in each testing stage, it is assumed that the same task is used in the comparison tests between a node u_i and nodes in $\Gamma(u_i)$. A "new" fault syndrome is formed after each testing stage using the same testing graph. In a single testing stage, each processing node is assigned at most one task to execute. Thus, all nodes in the same connected component of the testing graph must execute the same test task in a testing stage. The results of the comparison tests in different testing stages are assumed to be statistically independent of each other. The number of testing stages used in multiple syndrome testing is denoted as R.

There are several ways in which diagnosis can be done using the syndromes generated by the multiple syndrome testing method. In general, the optimal diagnosis method (in terms of diagnostic accuracy) is to produce the most probable diagnosis given all of the information contained in the multiple syndromes [7]. However, it has been shown that even with a single syndrome, finding the most probable diagnosis is an NP-hard problem [7]. In addition, an extremely high communication overhead is required to reliably distribute the syndrome information to all of the nodes. This problem is compounded when multiple syndromes are used.

Therefore, to obtain an efficient and practical diagnosis algorithm, we consider diagnosis algorithms in which each node u_i is only aware of the results of its tests with its immediate neighbors, referred to as *local* syndrome information. For a node $u_i \in V$, *summarized local* syndrome information is defined as $\{d^k(u_i) : 0 \le k \le R\}$, where $d^k(u_i) = d(u_i)$ for testing stage k (the number of nodes which test u_i to be faulty in stage k). There are two dimensions to the syndrome information: one dimension is $d^k(u_i)$ for a fixed k and the other dimension is the number of testing stages in which $d^k(u_i)$ is greater than a fixed threshold. Given an integer m, *m-threshold local* syndrome information is defined as $|\{k : 0 \le k \le R \text{ and } d^k(u_i) > m\}|$. *Category 3, 3A,* and *3AM* diagnosis are defined as diagnosis using local, summarized local, and m-threshold local (for any fixed m) syndrome information, respectively. This categorization is shown in Table 1. The diagnosis algorithm derived in this paper is the optimal category 3AM multiple syndrome diagnosis algorithm.

Category	Syndrome Info	Interpretation
3	Local	a_{ji} values for $u_j \in \Gamma(u_i)$
3A	Summarized local	$\{d^k(u_i) : 0 \le k \le R\}$
3AM	m-threshold local	$\lvert\{k : 0 \le k \le R \text{ and } d^k(u_i) > m\}\rvert$

Table 1: Categorization of diagnosis using local syndrome information.

2.2. Description of Fussell and Rangarajan's Algorithm

In this section, we describe Fussell and Rangarajan's (FR) algorithm [6] in detail since the FR algorithm is characteristic of category 3AM diagnosis. In the FR algorithm, testing is conducted in stages and two thresholds kv_i and sv_i are used. In testing stage i, it is assumed that all processors execute the same test task t_i. Let $T = \{t_1, \cdots, t_R\}$ be the set of R test tasks executed on all processors, where all tasks are treated identically. M is the total number of possible distinct incorrect results which a faulty processor can produce for a given test task. It is assumed that the probability distribution of producing incorrect results is uniform.

Algorithm FR:

0. Let $F \leftarrow \emptyset$ be the set of diagnosed faulty nodes;
1. **For** each $u_j \in V$ **do**
   $kv_j \leftarrow |\Gamma(u_j)| - 1$;
   $sv_j \leftarrow R - k_r R(1 - p(1 - \frac{p}{M})^{|\Gamma(u_j)|})$, where $1 \le k_r \le 2$;
2. **For** each $t_i \in T$ **do**
   for each $u_j \in V$ **do**
   if $d(u_j) > kv_j$
   then $L(i, j) = 1$;
   else $L(i, j) = 0$;
3. **For** each $u_j \in V$ **do**
   if $\sum_{t_i \in T} L(i, j) > sv_j$
   then $F \leftarrow F \cup \{u_j\}$;

In Step 1 of the description of the FR algorithm, kv_j is chosen to be $|\Gamma(u_j)| - 1$ and a range of values is indicated as being acceptable for the choice of sv_j. These thresholds were simply chosen in order for the the algorithm to satisfy desirable asymptotic properties. The authors proved that as $N \to \infty$, the diagnostic accuracy of the FR algorithm asymptotically approaches 100%. An earlier algorithm by the same authors [11] can be considered to be the same as the FR algorithm with $sv_j = 0$ for all $u_j \in V$.

3. ANALYSIS OF MULTIPLE SYNDROME TESTING

The two main differences between multiple and single syndrome testing are the use of multiple versus single syndromes and the constrained manner in which the syndromes are formed in multiple syndrome testing. In [11], examples are given to show that the additional information available when multiple syndromes are used permits correct diagnosis in some fault situations where an "accumulated" single syndrome results in incorrect diagnosis. In this section, we show that the way in which syndromes are formed in multiple and single syndrome testing also results in a significant difference in diagnostic capability.

Suppose we take all of the syndromes associated with the multiple testing stages in a multiple syndrome testing method and form an updated syndrome. The single syndrome generated using this process has the property that for a given node $u_i \in V$, all of the tests e_{ik} for $u_k \in \Gamma(u_i)$ use the same set of tasks in their testing. However, the syndrome used in single syndrome testing has no such restriction. In single syndrome testing, it is assumed that the test $e_{ik} \in E$ is independent of all of the other tests in E. As will be shown shortly, this "small" difference in testing method results in a significant difference in diagnostic capability.

To get a direct comparison, let us compare the difference in diagnostic capability between single and multiple syndrome testing when only

one syndrome is used in the multiple syndrome testing method. For simplicity of analysis, it is assumed that all nodes have identical parameter values and that the testing graph is regular with node-degree γ. The analysis can easily be individualized for each node (average parameter values are used for a given node's neighbors). Given a node u_i and any node $u_j \in \Gamma(u_i)$, let $A = P(a_{ji} = 1 \mid \delta_i) = p_i((1-f_j) + f_j(1 - \frac{p_j}{M}))$ $+ (1-p_i)(f_j p_j) = (1-f)p + fp(2 - p - \frac{p}{M})$ and $B = P(a_{ji} = 1 \mid \bar{\delta}_i) = f_j p_j = fp$. For single syndrome testing, the probability of having $z = d(u_i)$ one-links incident on u_i out of a maximum of γ links (denoted as z *one–links* : γ) given that u_i is faulty and non-faulty are

$$P(z \text{ one–links} : \gamma \mid \delta_i) = \binom{\gamma}{z} A^z (1-A)^{\gamma - z} , \quad (1a)$$

$$P(z \text{ one–links} : \gamma \mid \bar{\delta}_i) = \binom{\gamma}{z} B^z (1-B)^{\gamma - z} . \quad (1b)$$

For Eq. (1a), z of the γ neighbors of u_i are chosen such that those z neighbors test u_i to be faulty and the $\gamma - z$ other neighbors test u_i to be non-faulty. Eq. (1b) is obtained similarly. It follows that

$$P(\delta_i \mid z \text{ one–links} : \gamma)$$

$$= \frac{P(z \text{ one–links} : \gamma \mid \delta_i) f_i}{P(z \text{ one–links} : \gamma \mid \delta_i) f_i + P(z \text{ one–links} : \gamma \mid \bar{\delta}_i)(1-f_i)}$$

$$= \frac{1}{1 + \frac{1-f}{f} \left(\frac{B}{A}\right)^z \left(\frac{1-B}{1-A}\right)^{\gamma - z}} . \quad (1c)$$

For multiple syndrome testing with a single syndrome, the probability that there are z one-links incident on a node $u_i \in V$ given that u_i is non-faulty and faulty are:

$$P(z \text{ one–links} : \gamma \mid \bar{\delta}_i)$$

$$= \sum_{j=z}^{\gamma} \binom{\gamma}{j} f^j (1-f)^{\gamma - j} \binom{j}{z} p^z (1-p)^{j-z} = h(z) , \quad (2a)$$

$$P(z \text{ one–links} : \gamma \mid \delta_i)$$

$$= p \sum_{j=0}^{z} \binom{\gamma}{j} (1-f)^j f^{\gamma - j} \binom{\gamma - j}{z - j} (1 - \frac{p}{M})^{z-j} (\frac{p}{M})^{\gamma - z}$$

$$+ (1-p) h(z) \quad (2b)$$

$$= p\, g(z) + (1-p) h(z) .$$

For Eq. (2a), u_i must have z or more faulty neighbors, of which exactly z execute the test task incorrectly. For the first part of Eq. (2b), u_i produces an incorrect test result and has $j \le z$ non-faulty neighbors and $(z - j)$ faulty neighbors that either execute the test task correctly or produce a different incorrect result from u_i. It follows that the posterior fault probability of u_i given z one-links incident on u_i is

$$P(\delta_i \mid z \text{ one–links} : \gamma)$$

$$= \frac{f \left[p\, g(z) + (1-p) h(z)\right]}{f \left[p\, g(z) + (1-p) h(z)\right] + (1-f) h(z)} \quad (2c)$$

$$\simeq \begin{cases} 1 & \text{if } z = \gamma , \\ f(1-p) / (1 - fp) & \text{if } z \neq \gamma . \end{cases}$$

The approximation holds if M is large and $\gamma \le 4$. Then, Eq. (2c) is close to a delta function with a spike at $z = \gamma$ since f (prior fault probability) values are typically fairly small. If $\gamma > 4$, then Eq. (2c) becomes close to a step function. In general, Eq. (1c) is a much more smoothly increasing function of z than Eq. (2c).

Let us consider the testing time required for single versus multiple syndrome testing. Suppose that it takes τ units of time to execute each task and that each task has the same level of fault coverage p. Then, for multiple syndrome testing using a single syndrome, it takes τ units of time to obtain the syndrome since each node executes at most one task. In single syndrome testing, $\tau\,\gamma$ units of time are required for testing since each node executes γ tasks (in order to make γ comparison tests). Thus, using the same amount of testing time, we can execute γ test tasks (treated as one "large" test task) in one testing stage of the multiple syndrome testing method, thereby achieving an effective fault coverage of $1 - (1-p)^{\gamma}$. In Section 6, it is shown that good diagnosis results can be obtained when using multiple syndrome testing on a TMR structure. Thus, for concurrent diagnosis, "useful" tasks *can* be used for the test tasks required in single and multiple syndrome testing.

In Fig. 1, we have shown the distributions of $z = d(u_i)$ given u_i faulty and u_i non-faulty for both multiple syndrome testing (Eqs. (1a) and (1b)) and single syndrome testing (Eqs. (2a) and (2b)) using $M = 1000$, $\gamma = 4$, $p = 0.4$, and two different values of f. The fault coverage value p used for multiple syndrome testing is actually $1 - (1-p)^{\gamma} = 0.87$ since this level of fault coverage can be obtained in the same amount of testing time required to achieve fault coverage of p for single syndrome testing. A high level of diagnostic accuracy can be achieved if the syndrome information perceived when u_i is faulty is drastically different from the syndrome information perceived when u_i is non-faulty. From our analysis, it can be seen that the syndrome used in multiple syndrome testing fits this mold much more closely than the syndrome used in single syndrome testing.

4. OPTIMAL MULTIPLE SYNDROME DIAGNOSIS

Theorem 1: For any category of diagnosis, the algorithm which makes the most probable diagnosis for each node given the type of syndrome information available is the optimal (in terms of diagnostic accuracy) diagnosis algorithm.

Proof: Refer to [7].

To derive the optimal category 3AM multiple syndrome diagnosis algorithm, we must analyze the probability of a node $u_i \in V$ being faulty given the syndrome information for u_i. Category 3AM multiple syndrome diagnosis is based on m-threshold local syndrome information for any fixed m. Under the assumption that the test evaluation by a non-faulty processor is at least as good as the test evaluation by a faulty processor, the posterior fault probability of u_i is a non-decreasing function of $d^k(u_i)$ and also of $|\{k : 0 \le k \le R \text{ and } d^k(u_i) > m\}|$ for any m. Then, the optimal category 3AM multiple syndrome diagnosis algorithm is based on the selection of two thresholds. The first threshold z_{th_i} is for the number of one-links incident on u_i during a single testing stage. The second threshold H_{th_i} is for the number of testing stages in which a given node passes the first threshold. In the FR algorithm, $z_{th_i} = kv_i = \gamma - 1$ and $H_{th_i} = sv_i$. The thresholds kv_i and sv_i in the FR algorithm were simply chosen to prove desirable asymptotic properties of the algorithm.

Optimal threshold values can be obtained by calculating posterior fault probabilities. The optimal choice for z_{th_i} is obtained using Eq. (2c). The optimal z_{th_i} value, denoted as $z^*_{th_i}$, is equal to z such that $P(\delta_i \mid z \text{ one–links} : \gamma) \le 0.5$ and $P(\delta_i \mid z+1 \text{ one–links} : \gamma) > 0.5$, since u_i is more probably non-faulty (faulty) in the former (latter) case. Although calculation of Eq. (2c) for large values of γ is computationally expensive, since γ is at most the node-degree of the processor interconnection network, very large γ values will not be needed for most practical partially-connected systems. Also, from the analysis done in Section 3 (approximation for Eq. (2c)), it is evident that when $\gamma \le 4$, $z^*_{th_i} = \gamma - 1$. Since Eq. (2c) is a monotonically non-decreasing function of z, we can obtain $z^*_{th_i}$ for $\gamma > 4$ by evaluating Eq. (2c) for several values of z near γ.

We now use probability analysis to derive $H^*_{th_i}$, the optimal value of H_{th_i}. For simplicity of analysis, we will assume that $\gamma \leq 4$ so that $z^*_{th_i} = \gamma - 1$. (The changes required in the analysis when $\gamma > 4$ is discussed at the end of this section.) For a single syndrome,

$$P(\gamma \text{ one-links} : \gamma \mid \delta_i) = \begin{cases} A_1 \simeq p + (1-p)\,p^\gamma & \text{if } \Gamma(u_i) \subseteq F' \\ A_2 \simeq p & \text{otherwise} \end{cases}, \quad (3)$$

$$P(\gamma \text{ one-links} : \gamma \mid \bar{\delta}_i) = \begin{cases} B_1 = p^\gamma & \text{if } \Gamma(u_i) \subseteq F' \\ B_2 = 0 & \text{otherwise} \end{cases}. \quad (4)$$

The approximations in Eq. (3) hold since M is assumed to be large, that is,

$$A_1 = p_i(1 - \frac{p}{M})^\gamma + (1-p_i)\,p^\gamma \simeq p + (1-p)\,p^\gamma ,$$

$$p_i(1 - \frac{p}{M})^{\gamma-1} \leq A_2 \leq p_i \quad \Rightarrow \quad A_2 \simeq p .$$

If $\Gamma(u_i) \subseteq F'$, then A_1 and B_1 are the probabilities of having greater than z_{th_i} one-links incident on u_i given that u_i is faulty and non-faulty, respectively. A_2 and B_2 are the same probabilities when $\Gamma(u_i) \subseteq V - F'$. Next, the probabilities of having H syndromes in which $d(u_i) > z^*_{th_i}$ (denoted as H *passes*) given that u_i is faulty and non-faulty are

$$P(H \text{ passes} \mid \delta_i) = f^\gamma \binom{R}{H} A_1^H (1-A_1)^{R-H} + (1-f^\gamma) \binom{R}{H} A_2^H (1-A_2)^{R-H}, \quad (5)$$

$$P(H \text{ passes} \mid \bar{\delta}_i) = f^\gamma \binom{R}{H} B_1^H (1-B_1)^{R-H} + (1-f^\gamma) \binom{R}{H} B_2^H (1-B_2)^{R-H}. \quad (6)$$

For Eq. (5), the first additive term corresponds to the case when all of the neighbors of u_i are faulty and the second additive term corresponds to all other cases; for all cases, there must be H testing stages in which all of the neighbors of u_i test u_i to be faulty. Eq. (6) is obtained analogously. The posterior fault probability of u_i given that $d(u_i) > z^*_{th_i}$ for H syndromes is

$$P(\delta_i \mid H \text{ passes}) = \frac{P(H \text{ passes} \mid \delta_i)\,P(\delta_i)}{P(H \text{ passes} \mid \delta_i)\,P(\delta_i) + P(H \text{ passes} \mid \bar{\delta}_i)\,P(\bar{\delta}_i)} . \quad (7)$$

The value of H at which $P(\delta_i \mid H \text{ passes}) = 0.5$ is the optimal H_{th_i} value, $H^*_{th_i}$. It is difficult to determine $H^*_{th_i}$ directly because of the form of Eq. (7). However, since $H^*_{th_i}$ is to be used as a threshold for an integer quantity, we only need to determine $\overline{H}_{th_i} = \lfloor H^*_{th_i} \rfloor$. Eq. (7) is a monotonically non-decreasing function of H. Thus, $\overline{H}_{th_i}$ can be determined by calculating Eq. (7) for several values of H. This process is made simpler if a close upper bound for $H^*_{th_i}$ can be calculated. Denoting this upper bound by $\hat{H}_{th_i}$,

$$\hat{H}_{th_i} = \frac{\log\left[\frac{1-f_i}{f_i}\right]}{\log\left[\frac{A_1(1-B_1)}{(1-A_1)B_1}\right]} + R\,\frac{\log\left[\frac{1-B_1}{1-A_1}\right]}{\log\left[\frac{A_1(1-B_1)}{(1-A_1)B_1}\right]} . \quad (8)$$

Theorem 2: $\hat{H}_{th_i}$ is an upper bound for $H^*_{th_i}$.

Proof: Fig. 2 shows the distributions of three random variables Z_0, Z_1, and Z_2. Z_0 denotes the number of testing stages (out of R) that a faulty node u_i has $d(u_i) > z^*_{th_i}$. Thus, the distribution for Z_0 has Eq. (5) as its probability mass function. Z_1 is the binomial random variable with parameters R and A_1. Likewise, Z_2 is the binomial random variable with parameters R and A_2. Since $A_1 \geq A_2$, it is clear that the distribution for Z_1 is strictly to the right of the distribution for Z_2. Also, from the form of Eq. (5), we can see that the distribution for Z_0 must lie in between the distributions for Z_1 and Z_2. This relationship is shown graphically in Fig. 2. Since $B_1 \geq B_2$, the distribution for Eq. (6) can be shown to reside between two analogous random variable distributions. Thus, the distributions for Eqs. (5) and (6) are both shifted to the right when A_2 and B_2 are replaced by A_1 and B_1, respectively. When these replacements are made, $\hat{H}_{th_i}$ is the value of H at which $\left[P(H \text{ passes} \mid \delta_i)\,P(\delta_i)\right] = \left[P(H \text{ passes} \mid \bar{\delta}_i)\,P(\bar{\delta}_i)\right]$. Q.E.D.

To calculate $\overline{H}_{th_i}$, we need to execute the following procedure.

Procedure Calc_H:

1. Calculate $\hat{H}_{th_i}$ using Eq. (8);
2. **For** H from 0 to $\min\{R, \lfloor \hat{H}_{th_i} \rfloor + 1\}$ **do**
   calculate $P(\delta_i \mid H \text{ passes})$ using Eq. (7);
3. $\overline{H}_{th_i} \leftarrow H'$ such that $P(\delta_i \mid H' \text{ passes}) \leq 0.5$ and $P(\delta_i \mid H'+1 \text{ passes}) > 0.5$.

Algorithm OPTM, the optimal category 3AM multiple syndrome diagnosis algorithm, is essentially the same as the FR algorithm except that the thresholds are chosen differently. This, however, is a crucial difference since the performance of the algorithm hinges upon the choice of thresholds. Our algorithm is described below using the same assumptions as the FR algorithm with test tasks $T = \{t_1, \cdots, t_R\}$.

Algorithm OPTM:

1. **For each** $u_j \in V$ **doparallel**
   calculate $z^*_{th_j}$ using Eq. (2c);
   calculate $\overline{H}_{th_j}$ using procedure Calc_H;
2. **For each** $u_j \in V$ **doparallel**
   for each $t_i \in T$ **do**
   if $d(u_j) > z^*_{th_j}$
   then $L(i,j) = 1$;
   else $L(i,j) = 0$;
3. **For each** $u_j \in V$ **doparallel**
   if $\sum_{t_i \in T} L(i,j) > \overline{H}_{th_j}$
   then u_j is faulty;
   else u_j is non-faulty;

Although the OPTM algorithm is the optimal category 3AM diagnosis algorithm, it is not the optimal category 3 or even category 3A diagnosis algorithm. Since the effort required to obtain and use the syndrome information for category 3 or 3A is only slightly more than for category 3AM, it might seem worthwhile to derive the optimal category 3 and 3A multiple syndrome diagnosis algorithms. However, the diagnostic accuracy achieved by the OPTM algorithm is very close to the best possible even when category 3 or 3A syndrome information is used.

To see the reason for this, let us refer to the analysis of Section 3 and Fig. 1. Given a non-faulty node u_i, it is most likely to not have any one-links incident on it. Given a faulty node u_i, if it fails a test task t_j (in other words, t_j *covers* the fault), then it is most likely to have γ one-links incident on it; otherwise, if u_i passes t_j, then it is most likely to not have any one-links incident on it. Therefore, in all cases, u_i will most likely either have γ one-links or zero one-links incident on it. Given this observation, the syndrome information used in category 3AM contains almost all of the important category 3A syndrome information. In simulations using the experimental setup to be described in Section 6,

out of 80,000 syndromes generated for a 100-node torus-wrapped square mesh, there was exactly one syndrome in which a node u_i had neither $d(u_i) = \gamma$ nor $d(u_i) = 0$. Thus, the OPTM algorithm produced the optimal category 3A diagnosis over 99.998% of the time. Also, as explained in [7], the optimal category 3A diagnosis algorithm approximates the behavior of the optimal category 3 diagnosis algorithm since *average* probability parameter values are used in the analysis for category 3A. In summary, the OPTM algorithm is the optimal category 3AM multiple syndrome diagnosis algorithm and the "near-optimal" category 3A and category 3 multiple syndrome diagnosis algorithm.

In deriving the thresholds for the OPTM algorithm, we have used the assumption that $\gamma \leq 4$. If $\gamma > 4$, then it is possible that $z^*_{th_i} < \gamma - 1$. In that case, the A_1, A_2, B_1, and B_2 values used in Eqs. (3) — (8) must be changed. Let $\Delta_i \equiv |\Gamma(u_i) \cap F| = x$ (number of faulty neighbors of u_i). Then,

$$A_1' = P(d(u_i) > z^*_{th_i} : \gamma |\ \delta_i\ , \Delta_i = \gamma)\ , \cdots\ ,$$

$$A_{\gamma+1}' = P(d(u_i) > z^*_{th_i} : \gamma |\ \delta_i\ , \Delta_i = 0) \quad \text{and}$$

$$B_1' = P(d(u_i) > z^*_{th_i} : \gamma |\ \bar{\delta}_i\ , \Delta_i = \gamma)\ , \cdots\ ,$$

$$B_{\gamma+1}' = P(d(u_i) > z^*_{th_i} : \gamma |\ \bar{\delta}_i\ , \Delta_i = 0)$$

should be used instead of A_1, A_2, B_1, and B_2. A_1 and B_1 are replaced by A_1' and B_1', respectively. A_2 is replaced by A_2' through $A_{\gamma+1}'$ and B_2 is replaced by B_2' through $B_{\gamma+1}'$. Note that Eqs. (3) and (4) will now have more additive terms. It may be possible to combine some of the A_k' or B_k' values as was done in the analysis for the case of $z^*_{th_i} = \gamma - 1$.

5. ASYMPTOTIC ANALYSIS

One of the main desirable aspects of the FR algorithm was that it was shown to asymptotically achieve 100% correct diagnosis as $N \to \infty$ if $\gamma \geq 2$ and R grows faster than $\log N$. But, when $\bar{H}_{th_i}$ is calculated as in the previous section, we note that $\bar{H}_{th_i}$ is not necessarily one of the permitted values for sv_i in the FR algorithm (refer to Section 2.2). However, it is possible to directly prove that the OPTM algorithm also asymptotically achieves 100% correct diagnosis as $N \to \infty$. Let $\alpha(N)$ be any function of N such that $\lim_{N\to\infty} \alpha(N) = \infty$. In this section, we prove that the OPTM algorithm asymptotically achieves 100% correct diagnosis as $N \to \infty$ if $\gamma \geq 2$ and $R \geq \alpha(N) \log N$. We also prove that no category 3AM multiple syndrome diagnosis algorithm achieves 100% correct diagnosis as $N \to \infty$ if $R \leq \dfrac{\log N}{\alpha(N)}$.

For the asymptotic analysis of the OPTM algorithm, we need the following corollary [2] to a theorem proved by Chernoff [4].

Corollary 1: Let Z be a binomial random variable with parameters n and q. Then

$$P(Z \leq cnq) \leq e^{-(1-c)^2 nq/2}\ ,\quad 0 < c \leq 1\ ,$$

$$P(Z \geq cnq) \leq e^{-(c-1)^2 nq/3}\ ,\quad c \geq 1\ .$$

For the purposes of analysis, let us assume that $H^*_{th_i} = H^*_{th}$ for all $u_i \in V$ (the proofs also work when this does not hold). Since H^*_{th} is equivalent to $\bar{H}_{th}$ when used as a threshold, the analysis will be done assuming that H_{th} is the threshold used in the OPTM algorithm. Also, let a be the base of the logarithm unless otherwise specified. Let the random variable Y denote the number of faulty nodes u_i for which $d(u_i) > z^*_{th_i}$ for $\leq H_{th}$ syndromes. Let the random variable X denote the number of non-faulty nodes u_i for which $d(u_i) > z^*_{th_i}$ for $> H_{th}$ syndromes. For a multiple syndrome diagnosis situation, if there are no nodes which fit the requirements for random variables Y or X, then OPTM produces correct diagnosis. However, if there is any node that fits the requirement for random variable Y or X, then OPTM does not produce correct diagnosis. Thus, $1 - E[X] - E[Y] \leq$ P({OPTM produces correct diagnosis}) $\leq 1 - \max\{E[X], E[Y]\}$. The proofs of the lemmas and theorems in this section are based on showing limiting values for $E[X]$ and $E[Y]$.

Lemma 1: If $R \geq \alpha(N) \log N$, where $\lim_{N\to\infty} \alpha(N) = \infty$, and $RB_1 \leq H^*_{th} \leq RA_2$, then $\lim_{N\to\infty} E[X] = \lim_{N\to\infty} E[Y] = 0$.

Proof: In the following, let $m = \lfloor H^*_{th} \rfloor$. The comments refer to Fig. 2.

$$E[X] = \sum_{u_i \in V} P(> H^*_{th}\ passes \mid \bar{\delta}_i)\, P(\bar{\delta}_i)$$

$$\leq \sum_{u_i \in V} \left[\sum_{k=m+1}^{R} \binom{R}{k} B_1^k (1 - B_1)^{R-k} \right] (1 - f_i)$$

{distribution shifted right}

$$= (1 - f)\, N \sum_{k=m+1}^{R} \binom{R}{k} B_1^k (1 - B_1)^{R-k}$$

$$\leq (1 - f)\, N \sum_{k=m}^{R} \binom{R}{k} B_1^k (1 - B_1)^{R-k}$$

$$= (1 - f)\, a^{\log N}\, P(X \geq c_1 RB_1)$$

$$\leq (1 - f)\, a^{\log N}\, e^{-(c_1 - 1)^2 RB_1/3} \quad \text{if } c_1 = \frac{m}{RB_1} \geq 1.$$

Since $H^*_{th} \geq m \geq RB_1$ and $R \geq \alpha(N) \log N$, $\lim_{N\to\infty} E[X] = 0$. Similarly, for $E[Y]$,

$$E[Y] = \sum_{u_i \in V} P(\leq H^*_{th}\ passes \mid \delta_i)\, P(\delta_i)$$

$$\leq \sum_{u_i \in V} \left[\sum_{k=0}^{m} \binom{R}{k} A_2^k (1 - A_2)^{R-k} \right] f_i$$

{distribution shifted left}

$$= f\, N \sum_{k=0}^{m} \binom{R}{k} A_2^k (1 - A_2)^{R-k}$$

$$= f\, a^{\log N}\, P(X \leq c_2 RA_2)$$

$$\leq f\, a^{\log N}\, e^{-(1-c_2)^2 RA_2/2} \quad \text{if } c_2 = \frac{m}{RA_2} \leq 1 \text{ and } c_2 > 0.$$

$(H^*_{th} \leq RA_2) \Rightarrow (m \leq RA_2)$ and $R \geq \alpha(N) \log N$. Thus, it follows that $\lim_{N\to\infty} E[Y] = 0$. **Q.E.D.**

Theorem 3: If $R \geq \alpha(N) \log N$, where $\lim_{N\to\infty} \alpha(N) = \infty$, then P({OPTM produces correct diagnosis}) $\to 1$ as $N \to \infty$.

Proof: Let the random variables Z_0, Z_1 and Z_2 be as defined in the proof of Theorem 2. As shown in Fig. 2, Z_0 is sandwiched in between Z_2 and Z_1. Since Z_1 and Z_2 are binomial random variables, $E[Z_1] = RA_1$ and $E[Z_2] = RA_2$. Thus, $RA_2 \leq E[Z_0] \leq RA_1$. Likewise, if we let the random variable W_0 denote the number of syndromes for which a non-faulty node u_i has $d(u_i) > z^*_{th_i}$, then $RB_2 \leq E[W_0] \leq RB_1$. Then, from Eqs. (3) and (4),

$$B_1 = p^\gamma \leq p\,(1 - \frac{p}{M})^{\gamma - 1} \leq A_2$$

since M has been assumed to be a large number.

Now, let D be a category 3AM multiple syndrome diagnosis algorithm with $z_{th_i} = z^*_{th_i}$ and $H_{th} = 0.5\,R(B_1 + A_2)$. Then, since $RB_1 \le H_{th} \le RA_2$, D produces correct diagnosis as $N \to \infty$ provided the conditions of the theorem are satisfied. Thus, for any $\varepsilon > 0$, there exists an $N' > 0$ such that P({D produces correct diagnosis}) $> 1 - \varepsilon$. H^*_{th} is the optimal H_{th} threshold value. Therefore, P({OPTM produces correct diagnosis}) > P({D produces correct diagnosis}) $> 1 - \varepsilon$. Since this holds for any $\varepsilon > 0$, the theorem follows. **Q.E.D.**

Theorem 3 could also have been proven by using the fact that the FR algorithm is in category 3AM since the OPTM algorithm is the optimal category 3AM multiple syndrome diagnosis algorithm and Theorem 3 was proven for the FR algorithm [6]. However, we have provided a more direct proof of Theorem 3.

We now prove lower bounds on the number of testing stages required for asymptotically correct category 3AM multiple syndrome diagnosis. In [2], Blough essentially proved that if $\le \log N/\alpha(N)$ tests are performed on each processor, where $\lim_{N\to\infty} \alpha(N) = \infty$, then no category 3AM diagnosis algorithm which uses a single syndrome can achieve asymptotically correct diagnosis as $N \to \infty$. It follows that if γ is constant and $R \le \log N/\alpha(N)$, where $\lim_{N\to\infty} \alpha(N) = \infty$, then no category 3AM multiple syndrome diagnosis algorithm can achieve asymptotically correct diagnosis as $N \to \infty$. However, what happens when $R = \log N$? The following theorem answers this question.

Theorem 4: If $A_1 < 1 - \frac{1}{a}$ (recall that a is the base of the logarithm) and $R = \log N$, then for any category 3AM multiple syndrome diagnosis algorithm D, P({D produces correct diagnosis}) $\to 0$ as $N \to 0$.

Proof: Suppose that $A_1 < 1 - \frac{1}{a}$ and $R = \log N$. Then,

$$\lim_{N\to\infty} E[Y] \ge \lim_{N\to\infty} Nf\ (1 - A_1)^R = f \lim_{N\to\infty} \left[(1 - A_1)\, a\right]^{\log N} = \infty .$$

Thus, as $N \to \infty$, P({OPTM produces correct diagnosis}) $\to 0$ and likewise for any other category 3AM multiple syndrome diagnosis algorithm. **Q.E.D.**

As an example of the use of Theorem 4, if $a = 2$ (as in a hypercube structure), Theorem 4 tells us that we must have $A_1 \ge 0.5$, which implies that $p > 0.4$ (with $\gamma \ge 2$). It is unlikely that we can get such a high p (fault coverage) value using a single test task of short duration. A slightly higher upper bound for A_1 can be obtained by using a closer lower bound approximation for $E[Y]$. In summary, it appears unlikely that any category 3AM multiple syndrome diagnosis algorithm can achieve asymptotically correct diagnosis when $R = \log N$ unless a is very small, implying a quickly growing logarithm function, or very long test tasks are used.

6. SIMULATIONS

Simulations were conducted to evaluate the performance of the diagnosis algorithms studied. In assigning prior fault probability values, failure arrivals were assumed to follow a Poisson process. For each node u_i, a time value τ_i, corresponding to the length of time u_i has been in the system, was generated from a uniform distribution over the interval $[0, T]$ for some T. Then $f_i = 1 - e^{-\lambda\tau_i}$ was assigned to u_i, where $\lambda = \text{MTTF}^{-1}$ is the mean failure arrival rate.

The simulated experiments were conducted on a Sun 4/280 for a 100-node torus-wrapped square mesh and a 300-node TMR structure. A TMR structure is a 2-regular graph in which nodes are clustered into completely connected components of size 3 each. We used $T = 1K$ hours and MTTF values of $50K$ and $100K$ hours, resulting in λT values of 0.1 and 0.01 respectively. Clearly, the same results can be obtained by decreasing T and increasing MTTF by the same factor. Thus, this can model components which have been in the system for different lengths of times and components which have different MTTF values (their τ_i values can be adjusted). Given MTTF values of $50K$ and $100K$, $E[f_i] = 0.0099$ and 0.0050, respectively.

For all diagnosis algorithms evaluated, 1000 fault situations were produced and diagnosed assuming p values of 0.1 to 0.5 (in 0.1 increments) and the two MTTF values given above. For each fault situation, $R = 16$ testing stages were used, resulting in 16 syndromes. The OPTM algorithm was compared with the FR algorithm and the OPT3A algorithm, a category 3A single syndrome diagnosis algorithm. In [7], it is shown that the algorithms presented in Blough *et. al.* [1] and Dahbura *et. al.* [5] are category 3A single syndrome diagnosis algorithms, and that the OPT3A algorithm is the optimal category 3A single syndrome diagnosis algorithm. Assuming that it takes τ units of time to execute a single test task, multiple syndrome testing requires $R\tau$ time units while single syndrome testing requires $\gamma\tau$ time units. Thus, for single syndrome diagnosis, if the same amount of time is devoted to testing, it is possible to use test tasks which are R/γ times as long as those used in multiple syndrome diagnosis. Therefore, in the simulations for the OPT3A algorithm, $p' = 1 - (1 - p)^{R/\gamma}$ was used as the fault coverage value.

In the simulations for the FR algorithm, we must choose values for kv_i and sv_i. The FR algorithm specifies that $kv_i = \gamma - 1$ but indicates that a range of values is acceptable for sv_i (refer to Section 2.2). If the equation for sv_i in Section 2.2 is used, it is possible to get a negative value for sv_i. Since a negative sv_i threshold value implies that all nodes in V will be diagnosed to be faulty, this possibility is discounted. Then the modified equation for sv_i is $\max\{0\,,\ R - 2R(1 - p(1 - \frac{p}{M})^\gamma)\}$ $\le sv_i \le R - R(1 - p(1 - \frac{p}{M})^\gamma)$. In our simulations, sv_i was chosen to be the value halfway between the lower and upper bounds for sv_i. M was chosen to be 1000. Tables 2 and 3 show the values of $\overline{H}_{th}$ and sv_i for the torus-wrapped square mesh and TMR structure, respectively. $\overline{H}_{th}$ values shown in Tables 2 and 3 are for both $f = 0.0099$ and $f = 0.0050$ unless otherwise specified. sv_i and sv_i^{mid} (used in the simulations) are independent of f. In Table 3, threshold values for $p = 0.7$ and $f = 0.0099$ are shown to demonstrate that $\overline{H}_i$ and sv_i values do diverge.

	H_{th}	sv_i	sv_i^{mid}
$p = 0.1$	0	0 — 1.60	0.80
$p = 0.2$	0	0 — 3.20	1.60
$p = 0.3$	0	0 — 4.80	2.40
$p = 0.4$	0	0 — 6.40	3.20
$p = 0.5$	0	0 — 8.00	4.00

Table 2: Threshold values for torus-wrapped square mesh ($\gamma = 4$).

	H_{th}	sv_i	sv_i^{mid}
$p = 0.1$	0	0 — 1.60	0.80
$p = 0.2$	0	0 — 3.20	1.60
$p = 0.3$	0	0 — 4.80	2.40
$p = 0.4$	0	0 — 6.40	3.20
$p = 0.5$	1	0 — 8.00	4.00
$p = 0.7$	4 ($f = 0.0099$)	6.40 — 11.20	8.80

Table 3: Threshold values for TMR structure ($\gamma = 2$).

Figs. 3 and 4 show the results of the simulations for the torus-wrapped square mesh and TMR structure, respectively. In all cases, the OPTM algorithm performs significantly better than the OPT3A algorithm, with the difference more acute when p is small. The FR algorithm performs the same as the OPTM algorithm for $p = 0.1$, but then quickly falls off in accuracy as different $\lfloor sv_i^{mid} \rfloor$ threshold values are used. Similar results were obtained for all simulations attempted.

7. CONCLUSION

In this paper, we have derived an optimal category 3AM (and near-optimal category 3A and category 3) multiple syndrome diagnosis algorithm. Using probability analysis, multiple syndrome testing is shown to be more effective than single syndrome testing. Our simulation results support the probability analysis. It is proven that Algorithm OPTM, the optimal category 3AM multiple syndrome diagnosis algo-

rithm, achieves 100% correct diagnosis in an N processor system as $N \to \infty$ provided that $R \geq \alpha(N) \log N$ testing stages are used, where $\alpha(N) \to \infty$ arbitrarily slowly as $N \to \infty$. It is also shown that no category 3AM can achieve asymptotically correct diagnosis as $N \to \infty$ if γ is constant and $R \leq \log N/\alpha(N)$. If γ is constant, the computational complexity of the OPTM algorithm is $O(R)$, which is the minimum possible for any multiple syndrome diagnosis algorithm since R testing stages are required.

The OPTM algorithm requires each processor to execute identical tasks with its neighbors, send and receive the results of the tasks from its neighbors, compare the results received with its own results, and execute a diagnosis procedure to determine whether it should diagnose itself to be faulty or non-faulty. This distributed algorithm must be executed by a diagnostic component which operates in a fail-safe mode or is part of the *hard-core* of the processor. There must also be reliable communication between neighboring nodes. This can be accomplished by sending identical messages along multiple disjoint paths [7]. The thresholds used by the OPTM algorithm can be precomputed. Thus, the OPTM algorithm simply requires the diagnostic component to accumulate integer quantities and compare them against precomputed threshold values. The diagnostic component can therefore be a very simple digital circuit.

REFERENCES

[1] D. M. Blough, G. F. Sullivan, and G. M. Masson, "Almost certain diagnosis for intermittently faulty systems," *Digest of Papers, FTCS-18*, pp. 260-265, 1988.

[2] D. M. Blough, *Fault Detection and Diagnosis in Multiprocessor Systems*, Ph. D. Dissertation, The Johns Hopkins University, Baltimore, Maryland, 1988.

[3] M. L. Blount, "Probabilistic treatment of diagnosis in digital systems," *Digest of Papers, FTCS-7*, pp. 72-77, 1977.

[4] H. Chernoff, "A measure of asymptotic efficiency for tests of a hypothesis based on the sum of observations," *Annals Math. Stat.*, vol. 23, pp. 493-507, 1952.

[5] A. T. Dahbura, K. K. Sabnani, and L. L. King, "The comparison approach to multiprocessor fault diagnosis," *IEEE Trans. Comp.*, vol. C-36, no. 3, pp. 373-378, March 1987.

[6] D. Fussell and S. Rangarajan, "Probabilistic diagnosis of multiprocessor systems with arbitrary connectivity," *Digest of Papers, FTCS-19*, pp. 560-565, 1989.

[7] S. Lee, *Probabilistic Multiprocessor and Multicomputer Diagnosis*, Ph. D. Dissertation, The University of Michigan, Ann Arbor, Michigan, 1990.

[8] S. Mallela and G. M. Masson, "Diagnosable systems for intermittent faults," *IEEE Trans. Comp.*, vol. C-27, no. 6, pp. 560-566, June 1978.

[9] S. Mallela and G. M. Masson, "Diagnosis without repair for hybrid fault situations," *IEEE Trans. Comp.*, vol. C-29, no. 6, pp. 461-470, June 1980.

[10] F. P. Preparata, G. Metze, and R. T. Chien, "On the connection assignment problem of diagnosable systems," *IEEE Trans. Electron. Comp.*, vol. EC-16, no. 6, pp. 848-854, December 1967.

[11] S. Rangarajan and D. Fussell, "A probabilistic method for fault diagnosis of multiprocessor systems," *Digest of Papers, FTCS-18*, pp. 278-283, 1988.

[12] C. L. Yang and G. M. Masson, "A fault identification algorithm for t_i-diagnosable systems," *IEEE Trans. Comp.,*, vol. C-35, no. 6, pp. 503-510, June 1986.

* This work was supported in part by NASA under Grants NAG-1-492 and NAG-1-296 and the Office of Naval Research under Contract N00014-85-k-0531. Any opinions, findings, and conclusions or recommendations expressed in this paper are those of the authors and do not necessarily reflect the views of the funding agencies.

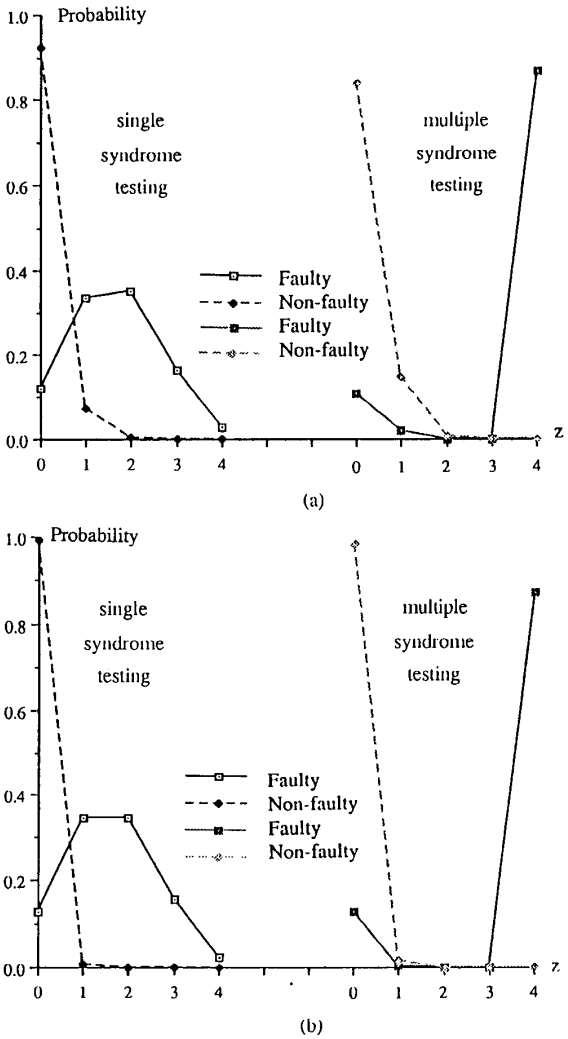

Figure 1: Probability distributions with MTTF of (a) 10K and (b) 100K.

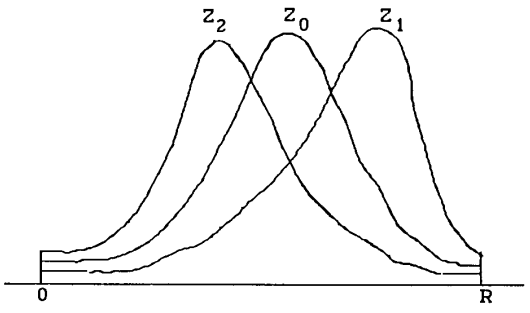

Figure 2: Distributions of random variables Z_0, Z_1, and Z_2.

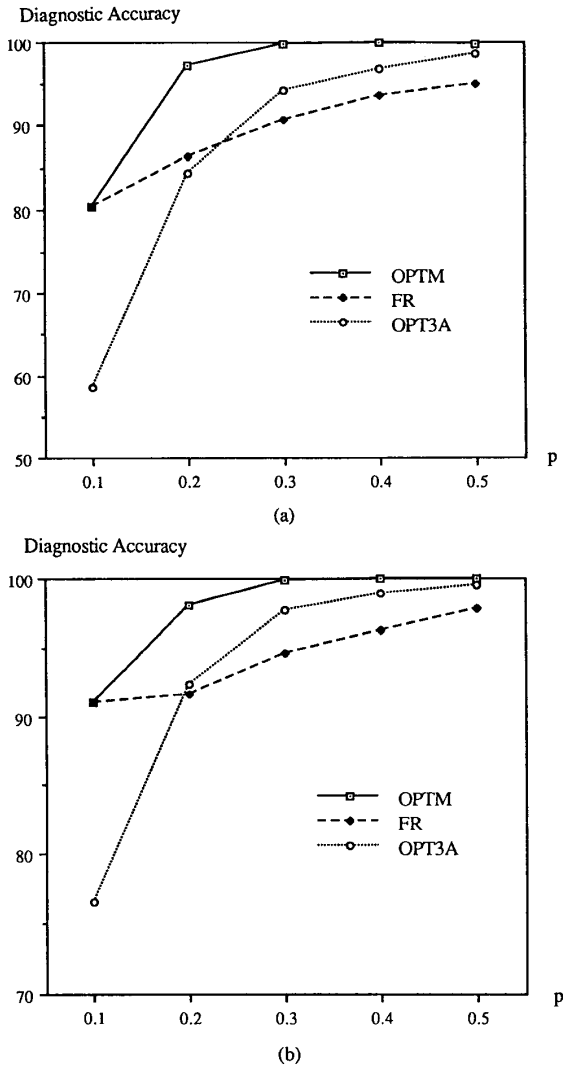

Figure 3: Accuracy with sq. mesh and MTTF of (a) 50K and (b) 100K.

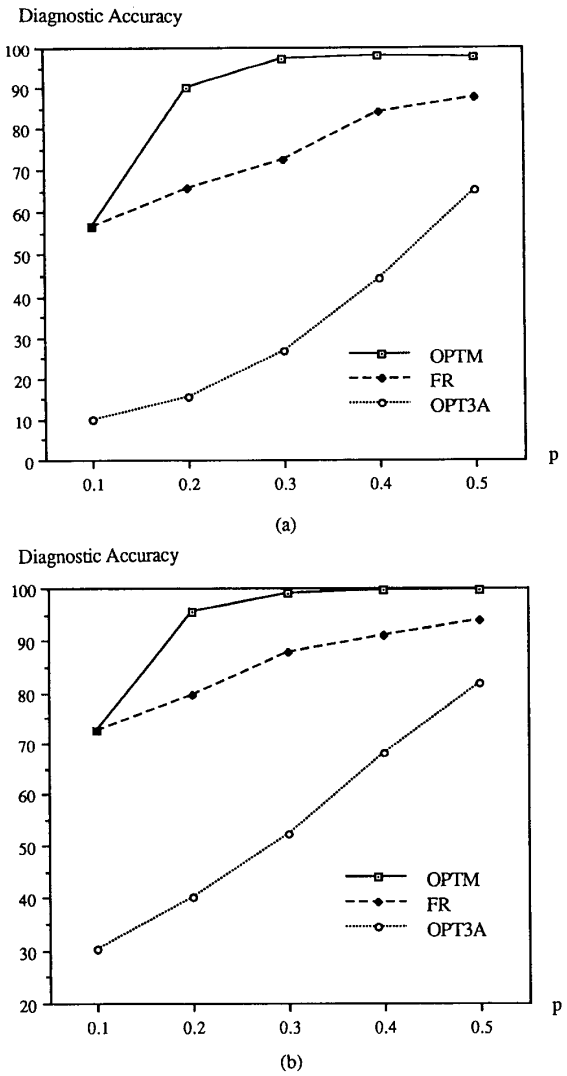

Figure 4: Accuracy with TMR structure and MTTF of

(a) 50K and (b) 100K.