Probabilistic Diagnosis of Multiprocessor Systems

SUNGGU LEE

POSTECH, Department of Electrical Engineering, P.O. Box 125, Pohang 790-600, Korea

KANG GEUN SHIN

The University of Michigan, Real-Time Computing Laboratory, Department of Electrical Engineering and Computer Science, Ann Arbor, MI 48109-2212

This paper critically surveys methods for the automated probabilistic diagnosis of large multiprocessor systems. In recent years, much of the work on system-level diagnosis has focused on probabilistic methods, which can diagnose intermittently faulty processing nodes and can be applied in *general* situations on *general* interconnection networks. The theory behind the probabilistic diagnosis methods is explained, and the various diagnosis algorithms are described in simple terms with the aid of a running example. The diagnosis methods are compared and analyzed, and a chart is produced, showing the comparative advantages of the various diagnosis algorithms on the basis of several factors important to probabilistic diagnosis.

Categories and Subject Descriptors: C.1.2 [**Processor Architectures**]: Multiple Data Stream Architectures—*MIMD; parallel processors*; D.4.5 [**Operating Systems**]: Reliability—*fault tolerance*; G.3 [**Mathematics of Computing**]: Probability and Statistics—*probabilistic algorithms (including Monte Carlo)*

General Terms: Algorithms, Performance

Additional Key Words and Phrases: Centralized and distributed self-diagnosis, comparison testing, fault-tolerant computing, probabilistic diagnosis, system-level diagnosis, system-level testing

INTRODUCTION

As large multiprocessing systems are increasingly being used in safety-critical applications, it is imperative that such multiprocessor systems be provided with good fault tolerance capabilities. Additionally, in order to maintain a highly reliable system, faulty PEs (processing elements) must be diagnosed and periodically removed (either physically or by reconfiguration) from the system. In large systems with more than about 1000 PEs, the fault diagnosis and reconfiguration tasks should be automated for efficient operation. However, the problem of identifying the faulty PEs in large systems is an extremely difficult task, especially since it is possible for faulty PEs to accuse nonfaulty PEs of being faulty. PEs can be *intermittently faulty* (a fault is defined to be intermittent if it is only occasionally present due to unstable hardware or varying hardware or soft-

The work reported in this survey was supported in part by NASA under grant NAG-1-296 and NAG-1-492. Any opinions, findings, and conclusions or recommendations expressed are those of the authors and do not necessarily reflect the views of the funding agencies.

CONTENTS

ware states [Siewiorek and Swarz 1982]), and tests to detect faulty PEs may fail to catch all possible faults.

This article presents a critical survey of the probabilistic approach to multiprocessor diagnosis. Due to the size of the systems being considered and the difficulty of the problem, we will only consider diagnosis at the level of a PE. This type of diagnosis is referred to as *system-level diagnosis*. A fundamental model for system-level diagnosis was developed some 25 years ago by Preparata et al. [1967]. In their model, referred to as the *PMC model*, it is assumed that PEs can test each other to arrive at separate conclusions about the fault status of other PEs. The system is modeled by a directed graph, called the *testing graph*, in which the vertices correspond to PEs, and the edges correspond to inter-PE-testing assignments. The (fault) *syndrome* is defined to be a binary labeling of the directed edges in which each label represents the result of the corresponding inter-PE test. As can be imagined, this can lead to an extremely large number of fault syndromes. The diagnosis subsystem is faced with the task of analyzing a syndrome to identify the set of faulty PEs. However, for every set of faulty PEs, there are a myriad of syndromes that can result because of the arbitrary manner in which faulty PEs and even nonfaulty PEs can evaluate other PEs. Nonfaulty PEs may exhibit arbitrary behavior due to other intermittently faulty PEs or because of the use of tests that are not able to catch all faults.

The currently available methods for system-level diagnosis can be broadly categorized into *deterministic* and *probabilistic* methods. *Deterministic* diagnosis methods are defined as those methods in which the entire fault set (or a well-defined subset of the fault set) can be uniquely identified from the syndrome provided that certain assumptions on the structure of the testing graph and the behavior of faulty and nonfaulty nodes are satisfied. By contrast, *probabilistic* diagnosis methods are defined as those methods that only attempt to correctly diagnose faulty nodes *with high probability* and require no restrictive assumptions on the structure of the testing graph.[1]

In the deterministic diagnosis methods, a restriction is imposed on the set of faulty nodes (such as an upper bound on the size of the fault set), and it is guaranteed that all faulty nodes (or a well-specified subset of the fault set) are caught by the diagnosis procedure. In these methods, the important issues are the *characterization* of testing graphs for which the diagnosis procedure is valid, procedures for determining the number of faulty nodes that can be diagnosed given a testing graph, and procedures for identifying the fault set.

In the probabilistic diagnosis methods, instead of placing restrictions on the set of faulty nodes, a probability model is

[1] Note that, using this definition, Maheswari and Hakimi's [1976] diagnosis method, which has been termed "probabilistic" in the past, belongs to the deterministic category since it uniquely identifies the entire fault set provided that (1) the sum of the prior fault probabilities of the faulty nodes is less than a prespecified bound and (2) the structure of the testing graph satisfies certain properties.

used to model the behavior of faulty and nonfaulty nodes, and based on this model, a procedure, which may be heuristic, is used to identify a set of nodes as faulty based on the syndrome observed. The important issues in probabilistic methods are the complexity of the diagnosis procedure and the "quality" of the diagnoses obtained. These methods frequently use probability parameters to describe the *probabilistic* behavior of nonfaulty and faulty nodes and to evaluate the quality of the diagnoses produced.

This article surveys and analyzes the currently available methods for probabilistic system-level diagnosis, since deterministic system-level diagnosis methods have been studied extensively, and several good surveys on this subject already exist [Dahbura 1988; Friedman and Simoncini 1980; Kime 1986]. For completeness, however, Section 1.4 includes a short discussion and analysis of the main results in deterministic diagnosis methods. The rest of Section 1 introduces the notation and definitions used in this survey and describes the proposed diagnosis classification scheme. Section 2 provides a preliminary discussion of probabilistic diagnosis, and Section 3 describes the probabilistic diagnosis algorithms in the literature, using examples to illustrate the algorithms. The issue of distributed self-diagnosis is discussed briefly in Section 4. Comparisons of the various methods are made in Section 5, and the survey concludes with Section 6.

1. PRELIMINARIES

1.1 Notation and Definitions

A system is composed of N *nodes* (processing elements), denoted by the set $V = \{u_1, \ldots, u_N\}$, where each node $u_i \in V$ is assigned a particular subset of the nodes in V to test. The set of testing assignments is represented by a directed graph $G = (V, E)$, called the *testing graph*, where (1) vertex $u_i \in V$ represents a node (processing element) and (2) edge $(u_i, u_j) \in E$ represents the fact that u_i tests u_j. The set of nodes that test a given node 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 $\Gamma_k^{-1}(u_i) = \{u_j \in \Gamma^{-1}(u_i) : a_{ji} = k\}$, $k = 0, 1$. $\gamma = \max_{u_i \in V}\{|\Gamma^{-1}(u_i)|\}$ is used to denote the maximum in-degree of the testing graph. The *fault coverage* of a test is defined as the probability that the test can detect a fault in the tested node given that there is a fault present. 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 test. a_{ij} is undefined if u_i does not test u_j. A (fault) *syndrome* S is a mapping from E to $\{0, 1\}$, defined such that for all $(u_i, u_j) \in E$, $S((u_i, u_j)) \equiv a_{ij}$.

Figure 1 shows an example of a testing graph and a syndrome. We assume that the testing graph is a subgraph of the graph representing the interconnection structure of the system. Although this assumption is not adopted by everyone, it makes the task of testing nodes substantially easier. Thus, if the system has a point-to-point interconnection structure, then a node can only test those nodes to which it is directly connected. The testing graph of Figure 1, for example, can map directly onto a 2D torus-wrapped mesh interconnection topology.

On the basis of a syndrome S, a *diagnosis* is performed when a set of nodes is identified as faulty. The diagnosis is said to be *correct* if there are no nonfaulty nodes mistakenly identified as faulty; otherwise, it is an *incorrect* diagnosis. Similarly, the diagnosis is said to be *complete* if all faulty nodes are identified as such; otherwise, the diagnosis is *incomplete*. If a diagnosis identifies the exact set of faulty nodes, then it is a correct and complete diagnosis; the *diagnostic accuracy* of a diagnosis algorithm refers to the percentage of diagnoses produced that are both correct and complete. A diagnosis algorithm is said to be *optimal* if it results in the highest possible level of diagnostic accuracy.

1.2 Testing Models

A *testing model* describes the test outcomes that are possible given that the testing and tested nodes are faulty

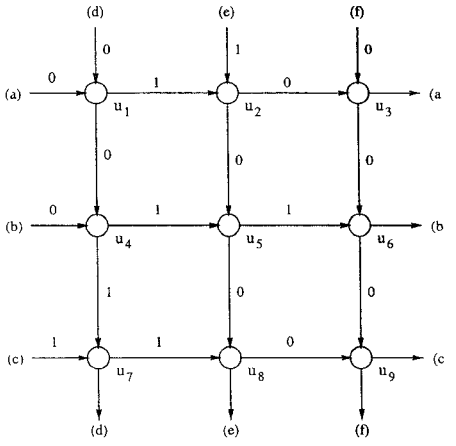

Figure 1. A testing graph and a syndrome.

and nonfaulty. Friedman and Simoncini [1980] present a complete tabulation of all of the nonequivalent testing models. The most general testing model, referred to as the *0-information tester* model, is one in which all test outcomes are possible regardless of the fault status of the testing and tested nodes. This testing model was actually used by Blount [1977] to produce a diagnosis method that guaranteed the most probable diagnosis but had exponential computational complexity.

Table 1 shows the probability parameters for the 0-information tester model. The probability parameter f_i is used to denote the prior fault probability of node u_i. The notation fs_m, $m = i$ or j, denotes the fault status of u_m. All possible combinations of fs_i, fs_j, and a_{ij} values are shown in Table 1 along with $P(a_{ij}|fs_i, fs_j)$, the probability of the test result a_{ij} given the fault status of the testing and tested nodes. Thus, q_{ij} is the probability that a nonfaulty node will incorrectly evaluate another nonfaulty node to be faulty. Since a_{ij} can only take on the values 0 and 1, the probability of $a_{ij} = 0$ under the same situation is $1 - q_{ij}$. The parameter q_{ij} was used by Blount [1977] to model the possibility of a faulty link between two nonfaulty nodes. p_{ij} is the probability that a nonfaulty node u_i will correctly evaluate a faulty node u_j. Hence, under a permanent fault model, p_{ij} is the fault coverage of the test applied by u_i on u_j. For ease of notation, p_{ij} will be referred to as fault coverage even when intermittent faults are permitted. r_{ij} and s_{ij} are the probabilities of a faulty node correctly diagnosing a nonfaulty and faulty node, respectively. As explained by Blount, r_{ij} and s_{ij} can model the extent to which a faulty node u_i can pass judgment on u_j. These two parameters are useful in modeling the behavior of faulty nodes.

Table 1. Probability Parameters for 0-Information Tester Model

fs_i	fs_j	a_{ij}	$P(a_{ij} \mid fs_i, fs_j)$
Good	Good	0	q_{ij}
Good	Good	1	$1 - q_{ij}$
Good	Faulty	1	p_{ij}
Good	Faulty	0	$1 - p_{ij}$
Faulty	Good	0	r_{ij}
Faulty	Good	1	$1 - r_{ij}$
Faulty	Faulty	1	s_{ij}
Faulty	Faulty	0	$1 - s_{ij}$

Most probabilistic diagnosis methods use a testing model that is more restrictive than the 0-information tester model. In the commonly used *partial-tester* model, the restriction $q_{ij} = 1$ is used. This implies that a nonfaulty node must always evaluate another nonfaulty node that it tests to be nonfaulty. Since the only way for $q_{ij} < 1$ is with a faulty testing link, the use of the partial-tester model is justified if the testing link (u_i, u_j) is assumed to be part of the node u_j. Another restriction that is sometimes used is $p_{ij} = 1$. This implies that faulty nodes can always be detected as such and thus requires tests with 100% fault coverage.

In many of the diagnosis algorithms, all nodes and all tests are treated identically. In this case, fixed values are assumed for the prior fault probability of a node, test fault coverage, and other parameters. Fixed probability parameter values 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.

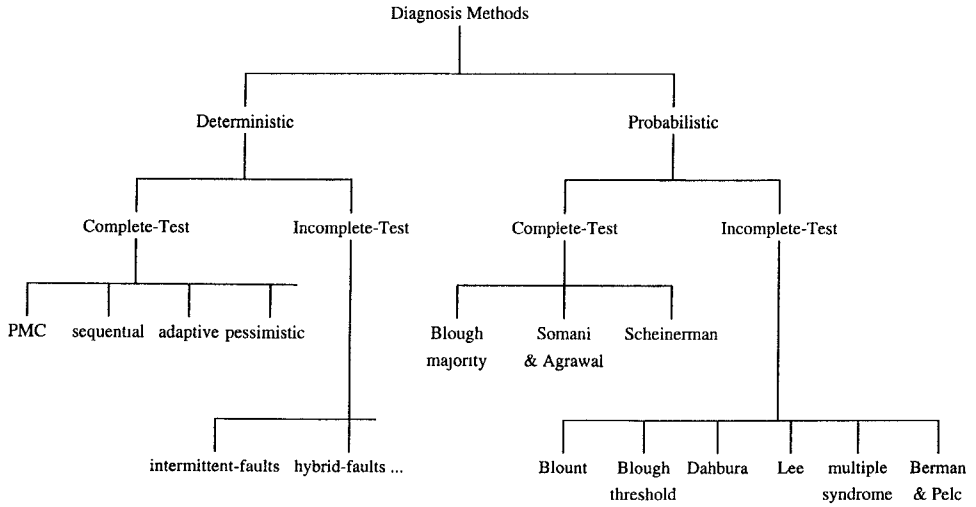

Figure 2. Classification of diagnosis methods.

1.3 Classification of System-Level Diagnosis Methods

Figure 2 shows the classification of diagnosis methods proposed in this survey. The categorization of deterministic and probabilistic methods has been described in the introduction. Diagnosis methods are further categorized into *complete-test* and *incomplete-test* methods. Complete-test methods assume $p_{ij} = 1$, i.e., that the system-level tests conducted by one node on another have complete (100%) fault coverage. Incomplete-test methods place no restrictions on p_{ij} values. Thus, complete-test methods must assume that all faulty nodes are permanently faulty and do not become faulty during the diagnosis. On the other hand, incomplete-test methods permit intermittently faulty nodes and nodes which become faulty during the diagnosis.

1.4 Overview of Deterministic Diagnosis Methods

This section provides a brief overview and analysis of the main results in deterministic system-level diagnosis methods. The interested reader is referred to Dahbura [1988], Friedman and Simoncini [1980], and Kime [1986] for more detailed surveys on deterministic diagnosis methods.

Complete-Test Methods

Most of the deterministic diagnosis methods that have appeared in the literature make the complete-test assumption ($p_{ij} = 1$). The PMC diagnosis model [Preparata et al. 1967] is representative of the efforts in this class of methods. A system is said to be *t-diagnosable* if all faulty nodes within the system can be identified without replacement provided the number of faulty nodes does not exceed t. Hakimi and Amin [1974] gave a complete characterization of t-diagnosable systems: specifically, a system with N nodes in which no two nodes test each other is t-diagnosable if and only if each node is tested by at least t other nodes. Figure 1 is an example of a t-diagnosable graph with $t = 2$.

An extremely large number of extensions and generalizations of the PMC model have been proposed in order to reduce the testing-graph requirements for diagnosability and expand the range of applicability of the diagnosis model. These include *sequential diagnosis* [Huang et al. 1989; Preparata

et al. 1967] (diagnosis with repair—diagnosis is done in stages with previously identified faulty nodes replaced at each stage), *adaptive diagnosis* [Hakimi and Nakajima 1984] (multistage diagnosis procedure in which a minimal number of tests are chosen at each stage based on the test results of the previous stage), *pessimistic diagnosis* [Friedman 1975; Kavianpour and Friedman 1978; Kavianpour and Kim 1991] (some nonfaulty nodes can be included in the diagnosed fault set), *p-t-diagnosability* [Maheswari and Hakimi 1976] (weighted diagnosis—prior node fault probabilities are used in the diagnosis), the *asymmetric invalidation model* [Barsi et al. 1976] (the restrictions $q_{ij} = 1$, $p_{ij} = 1$, *and* $s_{ij} = 1$ are used), Somani et al.'s [1987] model (diagnosis with respect to a *family* of allowable fault sets), and Russell and Kime's [1975] model. Russell and Kime's model is a generalization of the PMC model in which the relationships between faults and tests are formalized. Their model permits a test for a given fault to be invalidated by the presence of other faults.

Incomplete-Test Methods

Several incomplete-test deterministic diagnosis methods have been proposed. Mallela and Masson [1978] described a system as t_i-diagnosable when it is such that if no more than t_i nodes are intermittently faulty, then a nonfaulty node will never be diagnosed as faulty, and the diagnosis is at worst incomplete (i.e., some faulty nodes may not be identified). This model was generalized to combinations of permanently and intermittently faulty nodes in Mallela and Masson [1980] and Yang and Masson [1986]. Yang and Masson defined a syndrome to be *pf-compatible* if there is a set $U_f \subseteq V$, $|U_f| \leq t_i$, such that the syndrome can be produced under the assumption that only permanent faults exist and that U_f is the set of faulty nodes. They also gave an $O(|E|)$ fault diagnosis algorithm for t_i-diagnosable systems that finds all permanent fault sets U_f consistent with *pf*-compatible syndromes and identifies at least one faulty node over collections of syndromes significantly larger than the set of *pf*-compatible syndromes.

Analysis

Much progress has been made with deterministic diagnosis using the complete-test model, and there exist several methods which can be used effectively with sparse graphs and/or large numbers of faulty nodes (sequential diagnosis, adaptive diagnosis, pessimistic diagnosis, etc.). However, the basic limitation of the complete-test assumption ($p_{ij} = 1$) prevents the application of these methods in situations with intermittently faulty nodes, which is a severe limitation since intermittent faults are known to account for a large portion of the faults that occur in real systems. In fact, experimental studies have shown that more than 80% of the faults that occur in real systems are transient or intermittent faults [Horst et al. 1993; Siewiorek and Swarz 1982].

Deterministic incomplete-test methods tend to be too conservative and only identify nodes that are definitely faulty given the syndrome and fault set size limit. Thus, many intermittently faulty nodes can be missed. While deterministic diagnosis methods may be appropriate for the complete-test model, the probabilistic approach is more appropriate for the incomplete-test model.

2. PROBABILISTIC DIAGNOSIS—PRELIMINARY DISCUSSION

Probabilistic diagnosis methods do not make any prior assumptions about the set of faulty nodes and, in general, can be used with arbitrary testing graphs (it follows that the concept of diagnosability is not applicable to probabilistic diagnosis). As a consequence, probabilistic methods cannot guarantee that a correct and complete diagnosis is made. Thus, the quality of the diagnosis methods must be substantiated by other means.

Three arguments used to support probabilistic diagnosis algorithms are: (1) us-

ing analysis to show that high diagnostic accuracy is achieved in certain situations, (2) guaranteeing that the set of nodes most likely to be faulty given the syndrome is found, and (3) showing that as the number of nodes in the system grows to infinity, diagnostic accuracy approaches 100%. 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 [Blough 1988; Lee 1990]. From a practical perspective, argument (3) is insufficient since good diagnostic accuracy is desired for finite systems. However, since automated diagnosis is particularly important for large systems, asymptotically correct and complete diagnosis is certainly a desirable property of any probabilistic diagnosis algorithm. All three arguments have been used to support the probabilistic diagnosis methods surveyed in this article.

Probabilistic diagnosis methods have two serious limitations that must be understood before attempting to use the methods. Probabilistic methods usually require the use of probability parameters to model the behavior of faulty and nonfaulty nodes and to evaluate the quality of the diagnoses produced. The issue of obtaining these probability parameter values in real computer systems is an important unsolved problem which needs to be investigated. However, some of the diagnosis methods to be described can be implemented without knowing the probability parameter values or with imprecise, estimated values—in these methods, the probability parameters are used primarily for analyzing the diagnosis algorithms. A second limitation of all of the incomplete-test probabilistic diagnosis methods described in this survey is that they assume that the results of inter-PE tests performed by different nodes are statistically independent. However, since most of the diagnosis methods described permit arbitrary behavior of the faulty nodes, statistically dependent test results of *faulty* nodes do not negate the usefulness of the diagnosis methods, although they may lower the level of diagnosis accuracy produced.

Probabilistic methods differ in the types of probability parameters used and in the probability model used to define the probability of occurrence of particular syndromes. In the following, we discuss the different ways in which probability parameters and models are used in probabilistic diagnosis methods.

2.1 Probability Models

A probability model is characterized by defining 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. The probability model that is most often used, implicitly or explicitly, in most probabilistic methods is referred to as the *common* probability model. In this model, the sample space Ω is defined to be the set of all possible syndrome and fault set pairs given a testing graph G. Formally,

$$\Omega = \{(SD, F) : F \subseteq V \text{ and } SD \text{ is a function from } E \text{ to } \{0, 1\}\}.$$

Then, the event space Θ is taken to be the set of all possible subsets of Ω. The probability measure P is easily defined using the probability parameters defined above by assuming that all nodes in the set F are faulty and that all nodes in the set $V-F$ are nonfaulty. If the 0-information tester model is used, all possible syndrome and fault set pairs can have a finite probability value. However, if the partial-tester or complete-test model is used, then certain syndrome and fault set pairs will have zero probability value. As an example, under both testing models, a situation in which $u_i, u_j \in V - F$ and $SD((u_i, u_j)) \equiv a_{ij} = 1$ can never occur, and thus the probability of such a syndrome and fault set pair is 0.

Blough [1988] introduced a probability model that is more general than the common probability model. In his model, Blough modeled the behavior of faulty nodes, not by the parameters r_{ij} and s_{ij} defined in Section 2, but by assuming

that the faulty nodes behave in the manner most detrimental to the diagnosis algorithm. In Blough's model, the sample space Ω is the same as for the common model. However, the basic events of the model are defined to consist of all sets of syndrome and fault set pairs which have the same fault set and whose syndromes are identical except for the edges out of faulty nodes. Thus, a syndrome and fault set pair (SD', F') is contained in a basic event B defined as

$$B = \{(SD, F) : F = F' \text{ and } \forall (u_i, u_j) \in E \text{ with } u_i \in V - F, SD((u_i, u_j)) = SD'((u_i, u_j))\}.$$

B_{set} is defined as the set of all sets B such that B is a basic event of the testing graph G. The event space Θ is the set of all subsets B_{set}. In this model, the probability of correct and complete diagnosis by an algorithm A is defined to be the minimum of the probabilities of the syndrome and fault set pairs in the event B such that the fault set F is the set of nodes diagnosed to be faulty and B contains the actual syndrome observed.

Lee and Shin [1993] introduced another generalization of the common probability model. They noted that most probabilistic diagnosis methods use less than the global syndrome information in diagnosing the fault status of each node. They assumed a distributed self-diagnosis method in which each node diagnoses itself as faulty or nonfaulty based on a limited form of the global syndrome information, referred to as *partial* syndrome information. A different probability space was used for the diagnosis of each node. For a given node u_i, let the partial syndrome used in the diagnosis of u_i be denoted by SD_i, and let SD_i^{all} denote the set of all such partial syndromes. Then, for u_i, the sample space is defined as

$$\Omega_i = \{(SD_i, fs_i) : SD_i \in SD_i^{all}, fs_i \in \{\delta_i, \delta'_i\}\}.$$

The event space Θ_i is the set of all possible subsets of Ω_i. Finally, the definition of the probability measure P_i used is dependent on the partial syndrome information used.

Comparison of Probability Models

There are pros and cons to all of the probability models introduced. Using the common probability model, an exact evaluation can be made of the posterior fault probability of each fault set given a syndrome. Thus, it is possible to come up with a diagnosis algorithm which guarantees that the most probable diagnosis is made. Such a diagnosis algorithm can be shown to be *optimal* in diagnostic accuracy [Blough 1988]. However, since finding the most probable diagnosis in general testing graphs is NP-hard, this diagnosis algorithm has exponential computational complexity.[2]

Using Lee and Shin's probability model, it is possible to produce the most probable diagnosis given partial syndrome information using a polynomial-time algorithm. It can also be shown that such a diagnosis algorithm is *optimal* in diagnostic accuracy among all diagnosis algorithms that use the same type of partial syndrome information. This will be referred to as *locally optimal* to contrast with the definition of optimal diagnosis given previously. A serious limitation of both the common and Lee and Shin's probability models is the requirement that the behavior of both nonfaulty and faulty nodes must be known or estimated, in terms of the probability parameters, before any probability analysis can be done.

[2] An interesting, but restrictive, result in Blough and Pelc [1992] shows that if complete fault coverage is assumed ($p_{ij} = 1$) and the asymmetric invalidation model [Barsi et al. 1976] is used ($s_{ij} = 1$, i.e., a faulty node is always able to correctly diagnose another faulty node that it tests), then most probable diagnosis can be achieved for bipartite graphs in time $O(|E|\sqrt{N})$. Under these same restrictions, an $O(N)$-time algorithm is also presented for ring graphs.

Table 2. Pros and Cons of Three Probability Models

Model	Pros	Cons
common	Exact analysis, optimal diagnosis possible	Requires estimates of how faulty nodes behave, exp. computational complexity for optimal diagnosis
Lee & Shin	Exact analysis, locally optimal diagnosis	Requires estimates of how faulty nodes behave, limits syndrome info used
Blough	No parameterization of behavior of faulty nodes	Exact analysis not possible, many indistinguishable algorithms, optimal diagnosis not possible

The main advantage of Blough's [1988] probability model is that the behavior of faulty nodes do not have to be parameterized before the model can be used. However, the result of this is that many diagnosis algorithms that produce different diagnosis results will all be evaluated as being equal under Blough's model. Thus, Blough's model will not be able to distinguish between two diagnosis algorithms, one of which may perform significantly better than the other in terms of diagnostic accuracy. It follows that exact probability analysis and optimal diagnosis are not possible using Blough's model. With his probability model, Blough only makes statements regarding the *asymptotic* or upper-bound behavior of various diagnosis algorithms. The pros and cons of the three probability models discussed are summarized in Table 2.

2.2 General Asymptotic Results

Blough [1988] presents several important results concerning the asymptotic behavior of diagnosis algorithms. Asymptotic diagnostic accuracy refers to the limiting value of diagnostic accuracy of an algorithm given that the number of nodes in the system is increased to infinity while retaining the original testing-graph structure. While Blough uses his own probability model, Blough's asymptotic results extend to the other probability models discussed above.

The first set of results address conditions under which no diagnosis algorithm is able to produce asymptotically correct and complete diagnosis. Blough [1988] proved that if the number of edges in the testing graph grows slower than N, then the diagnostic accuracy of all diagnosis algorithms approaches 0. This is intuitively obvious since isolated nodes must exist if the number of edges grows slower than N, the number of nodes. A *regular* graph is a graph in which the number of edges adjacent to a vertex is the same for all vertices in the graph. Blough proved that for regular testing graphs, the diagnostic accuracy of any diagnosis algorithm approaches 0 as $N \to \infty$ if the number of testing edges grows slower than N log N. Recently, Berman and Pelc [1990] were able to show that this same result holds for general testing graphs.

Several results have also been shown regarding conditions under which diagnosis algorithms can produce 100% accurate diagnosis. These results are given in the description of the various probabilistic diagnosis methods in the next section.

3. DESCRIPTION OF PROBABILISTIC DIAGNOSIS METHODS

In this section, we give an overview of the methods available in the literature for probabilistic system-level diagnosis. Unless otherwise stated, the probabilistic diagnosis methods described use the

common probability model. A running example based on the testing graph and syndrome shown in Figure 1 is used to illustrate some of the concepts and algorithms. Note that Figure 1 is a contrived example intended to illustrate some of the differences between the various diagnosis methods presented. Thus, many of the methods to be described produce an incorrect or incomplete diagnosis. However, with a larger testing graph and more example syndromes, it is expected that most of these diagnosis methods will perform reasonably well.

3.1 Complete-Test Probabilistic Methods

If system-level testing is assumed to have 100% fault coverage, then several good deterministic methods exist that can guarantee that the unique set of faulty nodes is identified provided that the number of faulty nodes is less than an upper bound t. The reason that probabilistic methods have been introduced for this problem is to permit diagnosis in situations with more than t faulty nodes and for general testing-graph structures.

Scheinerman [1987] presented a probabilistic diagnosis method for the complete-test model with desirable asymptotic properties. In this algorithm, a core group of nonfaulty nodes is identified by finding a strongly connected subgraph of G in which all links are labeled with a 0 and in which more than half of the total nodes are present. Every node with a path of 0-links from this core group is then added to this set of nonfaulty nodes. All other nodes are identified as faulty. Scheinerman showed that his algorithm produces asymptotically correct and complete diagnosis in random graphs in which a node is connected to another node with probability $(c \log N)/N$, where $c > 1/(1 - f)$. Scheinerman's algorithm does not work for the testing graph and syndrome shown in Figure 1. This is because if all 1-links are removed, the remaining graph does not contain a strongly connected subgraph of more than 4 nodes. Scheinerman's work is significant, however, since it provided the first proof of asymptotically correct and complete diagnosis.

Using Blough's probability model, Blough et al. [1992a] presented a complete-test probabilistic diagnosis method in which each node u_i simply diagnoses itself to be faulty or nonfaulty based on the majority opinion of its testers $\Gamma^{-1}(u_i)$. Thus, u_i diagnoses itself to be faulty if and only if $|\{u_j : a_{ji} = 1 \text{ and } u_j \in \Gamma^{-1}(u_i)\}| > |\Gamma^{-1}(u_i)|/2$. Blough et al. were able to show that asymptotically, as the size of the system grows to infinity, 100% accurate diagnoses can be obtained for testing graphs in the form of hypercubes if $f < 0.067$. Additionally, they showed that asymptotically correct and complete diagnosis can be obtained for a special class of testing graphs with $N \times \omega(N)$ testing links, where $\omega(N)$ is any function that approaches infinity, albeit arbitrarily slowly.

Example 1. Let us use Blough et al.'s [1992a] algorithm on the testing graph and syndrome shown in Figure 1. Since u_2 and u_7 are the only nodes tested to be faulty by more than $|\Gamma^{-1}(u_i)|/2 = 1$ other nodes, the set of nodes diagnosed to be faulty is $F = \{u_2, u_7\}$. It is noted that this diagnosis is incomplete since one of u_5 or u_6 must be faulty because u_5 accuses u_6 of being faulty.

Somani and Agrawal [1989; 1992] introduced three more complex probabilistic diagnosis algorithms for this problem. These algorithms are based on initially identifying all nodes as being *potentially* faulty or nonfaulty, using the potentially nonfaulty nodes to identify *definitely* nonfaulty nodes, and then using the definitely nonfaulty nodes to identify other definitely nonfaulty and faulty nodes. In the first algorithm, majority voting is used in the first step to identify each node as potentially faulty or nonfaulty. In the second step, an iteration is used in which majority voting among the potentially nonfaulty nodes is used to identify certain nodes as definitely nonfaulty; the test results of the nodes identified as definitely nonfaulty are then used di-

rectly to identify other nodes as definitely faulty or nonfaulty. In the third step, any remaining potentially faulty and nonfaulty nodes are identified as definitely faulty and nonfaulty, respectively. The second algorithm differs from the first algorithm only in that unanimous voting among the potentially nonfaulty nodes is used to identify definitely nonfaulty nodes. The third algorithm differs from the first algorithm in that unanimous voting is used in both the first and second steps. The second and third algorithms are meant to be successively simpler algorithms from the first algorithm. Somani and Agrawal use several lemmas and theorems to describe the conditions under which their algorithms can guarantee to produce correct or correct and complete diagnosis. They also use examples to show that good diagnostic accuracy is obtained for several $\sqrt{N} \times \sqrt{N}$ meshes.

Example 2. To illustrate Somani and Agrawal's [1989] algorithms, let us again use Figure 1 as an example. Using the first algorithm, majority voting in the first step results in identification of $\{u_2, u_7\}$ as potentially faulty and the rest of the nodes as potentially nonfaulty. Then, using majority voting among the nodes in $V - \{u_2, u_7\}$, we obtain the set of definitely nonfaulty nodes $NF = \{u_1, u_3, u_4, u_8, u_9\}$. Using the test results of the nodes in NF we obtain the set of definitely nonfaulty nodes $NF = NF \cup \{u_6\}$ and the set of definitely faulty nodes $F = \{u_2, u_5, u_7\}$. Using the second and third algorithms, we obtain the same diagnosis as in the first algorithm since majority voting among 1 or 2 incoming links is the same as unanimous voting among the incoming links.

3.2 Incomplete-Test Probabilistic Methods

Probabilistic diagnosis methods are most useful in those situations where nodes can be intermittently faulty and where system-level tests have significantly less than 100% fault coverage. However, even the best probabilistic diagnosis method is not able to produce accurate diagnosis when the percentage of faulty nodes is too large and/or the fault coverage of system-level tests is too low.

The best possible probabilistic diagnosis method, in terms of diagnostic accuracy, is the one that guarantees that the most probable diagnosis given the syndrome is found. Blount [1977] described an early diagnosis method for solving this problem. Blount used the 0-information tester model and the common probability model to define a mapping from syndromes to fault patterns. Blount's algorithm finds the syndrome SD for which $P(SD, F)$ is the maximum for each possible fault set F. This information is encoded into a lookup table. Whenever a diagnosis needs to be made from an observed syndrome, the lookup table is accessed to find the most probable diagnosis F. Given $|E|$ edges and N nodes, there are $2^{|E|}$ possible syndromes and 2^N possible fault sets. Thus, to create the lookup table, $O(2^{N+|E|})$ calculations and $O(2^{|E|})$ memory locations are required. Lee [1990] used the partial-tester model and the common probability model to produce a more efficient method for finding the most probable diagnosis given a syndrome. In his method, he introduced several heuristics for making the search for the most probable diagnosis more efficient by bounding the search tree as early as possible. Although the average behavior of Lee's diagnosis algorithm is fairly good, the worst-case computational complexity of the method is $O(2^{|F|})$, where F is the set of faulty nodes found.

Realizing the limitations of finding the most probable diagnosis, Blough et al. [1992b] used Blough's probability model and presented an $O(|E|)$ algorithm which produces asymptotically correct and complete diagnosis provided that the number of testing links incident on each node is greater than log N. Examples were also given to show that this algorithm performs well for testing graphs with 100 and 1000 nodes and several sets of probability parameter values. In their algorithm, the number of 1-links directed toward a given node u_i is compared with a threshold value. Every node in which the

threshold value is exceeded is included into the fault set F. Next, all outgoing links from nodes in the set F are changed to be 1-links. Any nodes which then exceed the threshold value are included into F. This process is repeated until none of the nodes in $V - F$ exceed their respective threshold values. The property of asymptotically correct and complete diagnosis is proven with threshold values chosen as follows:

$$\forall u_i \in V, \quad \kappa_i = \frac{1}{2}|\Gamma^{-1}(u_i)|(f + p(1-f)). \tag{1}$$

Blough et al. [1992b] also describe a heuristic procedure by which better threshold values may be chosen. This heuristic procedure is based on the following lower-bound estimate of the probability of correct diagnosis:

$$\begin{aligned} &P(Correct\_Diag) \\ &\geq 1 - N(1-f)\sum_{J=k+1}^{\gamma}\binom{\gamma}{j}f^J(1-f)^{\gamma-J} \\ &\quad - Nf\sum_{J=0}^{\gamma}\left[\binom{\gamma}{j}(1-f)^J f^{\gamma-J}\right. \\ &\quad \left.\cdot\sum_{l=0}^{\min(J,k)}\binom{j}{l}p^l(1-p)^{J-l}\right], \end{aligned} \tag{2}$$

where a common threshold $k = \kappa_i$ $(u_i \in V)$ is used and $\gamma = \max_{u_i \in V}\{|\Gamma^{-1}(u_i)|\}$. Equation (2) is evaluated for all possible values of k, and the threshold value k resulting in the largest value for Eq. (2) is chosen. It is noted that since Eq. (2) only gives a pessimistic lower-bound value for $P(Correct\_Diag)$, it can produce *negative values*, and the threshold k chosen based on Eq. (2) is not necessarily optimal.

Example 3. To illustrate Blough et al.'s [1992b] algorithm, we need to add the parameters f_i and p_{ij} to the testing graph and syndrome shown in Figure 1. Let us use the fixed values $f = 0.01$ and $p = 0.9$. Then we obtain, for all $u_i \in V$, $k_i = 0.901$. Every node with 1 or more incoming 1-links exceeds this threshold. Thus, the set $F = \{u_2, u_5, u_7, u_8\}$ is the initial set of faulty nodes. However, after changing all of the outgoing links from F to be 1-links, *every* node has at least one incoming 1-link. Thus, the final fault set is $F = V$. If the heuristic procedure for selecting a better threshold k is used, the largest value for Eq. (2) is 0.98 with $k = 1$. Then, every node with 2 incoming 1-links exceeds this threshold, and the initial set of nodes diagnosed as faulty is $F = \{u_2, u_7\}$. In iteration (2), u_5 is added to F, and in iteration (3), u_8 is added to F. Thus, the final fault set in this case is $F = \{u_2, u_5, u_7, u_8\}$.

Dahbura et al. [1987] gave an $O(N^2)$ probabilistic diagnosis algorithm that is based on comparison testing. In their algorithm, they repeatedly select and remove from the testing graph a node which is incident on the largest number of 1-links until no 1-links remain in the testing graph. Using an assumed upper bound on the number of faulty nodes, Dahbura et al. show that for a completely connected testing graph, the probability of misdiagnosis is extremely small. Later, upon a reanalysis of Dahbura et al.'s algorithm, Lee [1990] was able to show that this algorithm had the same desirable asymptotic properties as Blough et al.'s algorithm [1992b]. Simulations also showed that Dahbura et al.'s algorithm performed significantly better than Blough et al.'s algorithm with the threshold in Eq. (1) for several different topologies and sets of probability parameter values. This demonstrates that asymptotic accuracy is not a *sufficient* measure of the goodness of a probabilistic diagnosis algorithm.

Example 4. Applying Dahbura et al.'s [1987] algorithm on the example of Figure 1 with $f = 0.01$ and $p = 0.9$, we find that nodes u_2 and u_7, with two incoming 1-links each, have the largest number of incoming 1-links. Arbitrarily choosing u_2 from among these nodes, we find that the initial fault set is $F = \{u_2\}$. u_2 and its incoming and outgoing links are removed from the graph G. Next, in $G - F$, u_7 has the largest number of incoming 1-

links—thus, $F = F \cup \{u_7\}$. After removing u_2 and u_7 from the graph G, u_5 and u_6 each have one incoming 1-link, and the rest of the nodes have no incoming 1-links. Arbitrarily choosing u_5 from the set $\{u_5, u_6\}$, we obtain $F = F \cup \{u_5\} = \{u_2, u_5, u_7\}$. Finally, in the graph $G - F$, there are no 1-links, and the algorithm terminates.

Lee and Shin [1993] presented another set of diagnosis algorithms. It was first noted that many of the previous probabilistic diagnosis algorithms used only partial syndrome information in diagnosing the fault status of each node. Several categories of diagnosis were defined based on the type of partial syndrome information used in the diagnosis of each node. Then, for each category of diagnosis, Lee's probability model was used to calculate posterior fault probability values for each node. Faulty nodes were identified based on these posterior fault probability calculations. All algorithms developed were shown to have the same desirable asymptotic properties as Blough's algorithm. The main advantage of Lee and Shin's diagnosis algorithms is that they produce the most probable diagnosis given a particular type of syndrome information (a locally optimal diagnosis). The main limitation of this type of method is that the diagnosis method is dependent on the use of probability parameter values.

Two of Lee and Shin's [1993] diagnosis algorithms are described here. The other algorithms are of similar character. In Algorithm OPT3A, each node u_i compares the number of 1-links incident on it with a threshold z_{th_i}. z_{th_i} is obtained as

$$z_{th_i} = \frac{\log\left(\frac{1-f_i}{f_i}\right)}{\log\left(\frac{A(1-B)}{(1-A)B}\right)} + |\Gamma^{-1}(u_i)| \frac{\log\left(\frac{1-B}{1-A}\right)}{\log\left(\frac{A(1-B)}{(1-A)B}\right)},$$

where $A = (1-f)p + fs$ and $B = f(1-r)$. All nodes in which the threshold is exceeded are diagnosed to be faulty. In Algorithm OPT2A, each node u_i calculates its posterior fault probability assuming partial syndrome information. The fault set F is initialized to $\emptyset$. The node u_j with the highest posterior fault probability is added to F. The posterior fault probabilities of all neighbors of u_j are updated given the knowledge that u_j is faulty. Then the node with the highest posterior fault probability in $V - F$ is again added to F. This process is repeated until all 1-links in G originate from or terminate on nodes in F.

Example 5. We again use the example of Figure 1 with the added parameters $f = 0.01$, $p = 0.9$, and $r = s = 0.5$. Using the equation shown above, we get $A = 0.896$, $B = 0.005$, and $z_{th_i} = 1.224$. Thus, using Algorithm OPT3A, we obtain the fault set $F = \{u_2, u_7\}$. For Algorithm OPT2A, we will use intuitive calculations rather than the complex equations that can be found in Lee and Shin [1993]. Initially, nodes u_2 and u_7 will have the highest posterior fault probabilities since u_2 and u_7 have the largest number of incident 1-links. Starting with the fault set $F = \emptyset$, first one, then the other of u_2 and u_7 are added to F. Next, we note that the nodes u_5, u_6, and u_8 each have one 1-link incident on them. However, the 1-link incident on u_8 comes from $u_7 \in F$. Thus, u_5 and u_6 are more likely to be faulty than u_8. Next, the 0-link incident on u_5 comes from the known faulty node $u_2 \in F$ while u_6's incident 0-link comes from $u_3 \in V - F$. Thus, u_5 is the node with the highest updated posterior fault probability. The final fault set is $F = \{u_2, u_5, u_7\}$ since all 1-links are accounted for by the nodes in F.

Fussell and Rangarajan [1989] introduced an entirely different type of diagnosis method. Given that the testing graph is a subgraph of the graph representing the interconnection structure, the previous incomplete-test diagnosis methods require each node to be tested by at

least log N other nodes for asymptotically correct and complete diagnosis. Fussell and Rangarajan improved on previous methods by showing that the same asymptotic result can be obtained for systems with lower connectivity (e.g., meshes or rings) if each pair of nodes conducts multiple tests and if the number of *these* tests on each node grows faster than log N. Fussell and Rangarajan's algorithm uses R stages of comparison testing. In testing stage i, all nodes are assumed to execute the same test task. After a testing stage, each node compares its results with the results of all adjacent nodes. The testing link between two nodes is labeled with a 1 for that testing stage if the two nodes have different results for the test task. In this manner, R independent syndromes are obtained. Two thresholds sv_t and kv_t are used. A node u_t is identified as faulty if and only if the number of testing stages in which it had greater than kv_t 1-links incident on it is greater than the second threshold sv_t. For all nodes $u_t \in V$, kv_t was chosen to be $|\Gamma^{-1}(u_t)| - 1$, and a range of values was indicated as being acceptable for sv_t. Rangarajan and Fussell [1992] derive better kv_t and sv_t threshold values and discuss a hierarchical version of the above algorithm in which (1) testing and diagnosis is conducted in *clusters* and (2) a third threshold is used for the number of clusters in which a node u_t is diagnosed to be faulty. For simplicity, we will only consider the algorithm by Fussell and Rangarajan [1989] since Rangarajan and Fussell [1992] is a hierarchical generalization of the first paper.

Example 6. To demonstrate Fussell and Rangarajan's [1989] algorithm, we need several syndromes of the form shown in Figure 1. Let us assume $R = 10$ testing stages with the first 2 syndromes identical to the syndrome shown in Figure 1. In syndromes 3 through 10, suppose that $a_{25} = 1$ and that all other link labels remain unchanged. Let us choose the threshold values $kv_t = 1$ and $sv_t = 7.5$. u_2 and u_7 have 2 ($> kv_t$) 1-links incident on them in 10 ($> sv_t$) syndromes. u_5 also has 2 ($> kv_t$) 1-links incident on it in 8 ($> sv_t$) syndromes. No other nodes have 2 1-links incident on them in greater than $sv_t = 7.5$ syndromes. Thus, the fault set is $F = \{u_2, u_5, u_7\}$.

Referring to the use of multiple testing stages and multiple syndromes as a *multiple-syndrome diagnosis* method, Lee and Shin [1990a] derived a locally optimal multiple-syndrome diagnosis algorithm using their own probability model. The only change introduced by Lee and Shin's multiple-syndrome diagnosis algorithm is the way in which the thresholds kv_t and sv_t are chosen. Lee and Shin's multiple-syndrome diagnosis algorithm is the same as Fussell and Rangarajan's algorithm except for the choice of the kv_t and sv_t thresholds. Posterior fault probability calculations are used to derive optimal values for kv_t and sv_t. This algorithm shares the same desirable asymptotic properties as Fussell and Rangarajan's algorithm.

Similar to a multiple-syndrome diagnosis strategy is a *sequential diagnosis* strategy for probabilistic diagnosis. In sequential diagnosis, diagnosis is conducted in stages, with nodes identified as faulty in the ith stage replaced with spares before commencing with the $(i + 1)$th stage of diagnosis. From a diagnosis viewpoint, the only difference between multiple-syndrome diagnosis and sequential diagnosis is the replacement of nodes identified as faulty in the sequential diagnosis strategy. Blough and Pelc [1993] present four algorithms for sequential diagnosis given the four possible combinations of complete and incomplete system-level tests and perfect and imperfect spares. Using their algorithms, the total number of tests required to produce asymptotically correct and complete diagnosis is $O(N)$ in the complete-test, perfect-spare model, $O(N \log N)$ in both of the intermediate models, and $O(N \log^2 N)$ in the incomplete-test imperfect-spare model. The basic diagnosis strategy takes place in two phases. A core set of non-

faulty nodes *NF* is identified in the first phase. This is followed by a second phase in which the nonfaulty nodes are used to determine the fault status of other nodes. The nodes found to be nonfaulty are added to *NF*, and the nodes found to be faulty are replaced by spares and tested. Multiple tests and multiple replacements of spares are used in the case of incomplete tests and imperfect spares, respectively. Blough and Pelc's sequential diagnosis algorithms permit asymptotically correct and complete diagnosis using testing graphs in the form of rings and meshes.

3.3 Distributed Self-Diagnosis

Distributed self-diagnosis can be executed with minimal message overhead if complete testing is assumed. In deterministic complete-test diagnosis methods such as in Hosseini et al. [1988] and Kuhl and Reddy [1980], test information is only transmitted once from a node u_i to another node u_m that it tests to be nonfaulty. By passing information along paths that are known to be fault free (on the basis of complete fault coverage tests), no redundant messages are transmitted. If a single test may be incomplete, but a collection of tests conducted by different testers of a common node u_i is assumed to be complete, as in Buskens and Bianchini [1993], then a similar method can be used to distribute testing information. Probabilistic complete-test methods can also use the same type of method to distribute testing information. However, if incomplete testing is assumed, then faulty nodes can never be identified with 100% certainty, and thus, the above methods cannot be used.

Three general approaches can be identified for performing distributed self-diagnosis using a probabilistic incomplete-test diagnosis model. In the first type of method, the communication of test results is mixed in with the actual testing and diagnosis. Thus, redundant copies of test results and/or partial diagnosis results must be communicated over multiple paths. The diagnosis is made taking into account the redundant copies received. In the second type of method, a reliable broadcast procedure is used to distribute every node's test results to every other node. After this is done, every nonfaulty node can then execute the appropriate diagnosis algorithm to arrive at its diagnosis of the overall system. Although the reliable broadcast operation incurs an extremely high message traffic overhead, methods such as in Lee and Shin [1990b] have shown that this operation can be implemented extremely efficiently on regular mesh and hypercube topologies. In the third type of method, each node only diagnoses the fault status of its immediate neighbors. This requires much less communication overhead than the alternative methods.

Berman and Pelc's [1990] diagnosis method is an example of the first approach to distributed self-diagnosis. Berman and Pelc's algorithm is designed for a special class of testing graphs with $O(N \log N)$ edges. These graphs are such that the set of nodes can be partitioned into subsets of completely connected nodes with $c \log N$ nodes each. Within a completely connected subset of nodes, a maximum clique of nodes connected by 0-links is found and labeled as nonfaulty; all other nodes in the subset are labeled as faulty. In the second stage, all nodes in each subset of completely connected nodes receive the diagnosis results of all other nodes. The final diagnosis for a given node u_i is the majority value of the received messages on u_i. The computational complexity of this algorithm is $O(N^c)$, where c is a constant that depends on the probability parameters f and p. Berman and Pelc are able to show that their algorithm has a diagnostic accuracy of at least $(1 - N^{-1}) \times 100\%$.

Pelc [1993] improved on Berman and Pelc [1990] by using a modified diagnosis algorithm with an $O(N)$ computational complexity. In his improved algorithm, the computationally expensive step of finding a maximum clique of nodes connected by 0-links is deleted. Instead, the set of nodes is partitioned into subsets of

nodes with $c \log N$ nodes each. The subsets are then connected in a cycle. All nodes within a subset test one another and test all nodes within an adjacent subset in the cycle. All test results are then communicated to the same set of nodes. This test-and-send procedure is repeated r times, for a fixed r. Then majority voting is used to determine the faulty nodes, and this diagnosis result is communicated to other nodes using a procedure similar to that in Berman and Pelc.

The second approach to distributed self-diagnosis is very general but requires high communication overhead. Each node must initiate a reliable broadcast procedure to broadcast its test results. This reliable broadcast procedure is described in general terms in Yang and Masson [1988]. Algorithms designed for specific architectures such as hypercubes [Ramanathan and Shin 1988] are also available. After testing, suppose that the test evaluation results of each node are combined into a single message. Up to t faulty nodes can be tolerated by having each node send $2t + 1$ copies of its message to every other node along node-disjoint paths [Yang and Masson 1988]. Since there are $N(N - 1)$ possible sender–receiver pairs, there are $(2t + 1)N(N - 1)$ such message transmissions. Lee and Shin [1990b] describe a method for implementing such an "all-to-all reliable broadcast" operation within a few milliseconds on large systems with over 10K nodes by using wormhole routing. After this reliable broadcast procedure, each node can use majority voting to determine the test evaluation results of all other nodes.

The third approach to distributed self-diagnosis applies to those diagnosis methods that identify each node as faulty or nonfaulty based only on the test evaluations of nodes directly connected to it. In these diagnosis methods, each node u_i must reliably receive the tests results of nodes $u_j \in \Gamma^{-1}(u_i)$. If comparison testing is being used, then no extra communication is required for diagnosis since each node knows the results of its comparison tests with its neighbors after the testing phase. With other inter-PE-testing methods, it may be necessary to execute a reliable multicast procedure in which each node sends a message reliably to all of its adjacent nodes. Although this method has extremely low communication overhead when compared to the previous methods, the diagnosis result is dispersed throughout the system. To obtain a global diagnosis, the results of the individual diagnoses must be combined.

4. COMPARISON

A large number of probabilistic diagnosis methods have been introduced in this survey. Table 3 shows a comparison of the various diagnosis methods on the basis of several factors important to probabilistic diagnosis. Acronyms based on the authors' last names are used to refer to the various diagnosis algorithms or sets of diagnosis algorithms. The computational complexities of the FR [Fussell and Rangarajan 1989] and LS2 [Lee and Shin 1990a] algorithms are dependent on R, the number of testing stages used, and γ, the maximum number of nodes testing any given node $u_i \in V$. The computational complexity of the BEP [Berman and Pelc 1990] algorithm is $O(N^c)$, where c is a constant dependent on the probability parameters f and p. The $O(N \log N)$ complexity of the PEL [Pelc 1993] algorithm is for sequential diagnosis with perfect spares. If imperfect spares are assumed, the computational complexity increases to $O(N \log^2 N)$.

All of the diagnosis algorithms shown except for the SA [Somani et al. 1987] algorithm have been proven to produce asymptotically correct and complete diagnosis provided certain prespecified conditions are met. An "optimal-diagnosis" diagnosis algorithm is one which identifies all of the faulty nodes correctly with the maximum probability. The notation "local" under the "Optimal Diagnosis" column in Table 3 refers to the fact that those algorithms only produce *locally* optimal diagnoses (i.e., with the restriction that only local syndrome in-

Table 3. Comparison of Probabilistic Diagnosis Methods

	Comput. Complexity	Optimal Diagnosis	Asymptotic Guarantee	# Links Required	# Prob. Parameters
Complete-Test:					
SCH [Scheinerman 1987]	$O(N^2)$	No	Yes	$O(N \log N)$	None
BSM1 [Blough et al. 1992a]	$O(N)$	No	Yes	$O(N\,\omega(N))$	None
SA [Somani and Agrawal 1989]	$O(N^2)$	No	No	N.A.	None
Incomplete-Test:					
BLO [Blount 1977]	$O(2^{N+\|E\|})$	Yes	Yes	$O(N \log N)$	All
LEE [Lee 1990]	$O(2^{\|F\|})$	Yes	Yes	$O(N \log N)$	4
BSM2 [Blough et al. 1992b]	$O(\|E\|)$	No	Yes	$O(N \log N)$	2
DSK [Dahbura et al. 1987]	$O(N^2)$	No	Yes	$O(N \log N)$	None
LS1 [Lee and Shin 1993]	$O(N^2)$	Local	Yes	$O(N \log N)$	4
FR [Fussell and Rangarajan 1989]	$O(R\,\gamma)$	No	Yes	$O(N)$	1
LS2 [Lee and Shin 1990a]	$O(R\,\gamma)$	Local	Yes	$O(N)$	2
BLP [Blough and Pelc 1993]	$O(N \log N)$	No	Yes	$O(N)$	2
BEP [Berman and Pelc 1990]	$O(N^c)$	No	Yes	$O(N \log N)$	None
PEL [Pelc 1993]	$O(N)$	No	Yes	$O(N \log N)$	None

formation is available). The number of testing links necessary for the asymptotic guarantee are shown in the fourth column. For the BSM [Blough et al. 1992a] algorithm, $\omega(N)$ is any function that approaches infinity as $N \to \infty$, albeit arbitrarily slowly. The proofs for the figures in this column can be found in the respective references and in Blough [1988] and Lee [1990]. Finally, the fifth column shows the number of probability parameters used in the respective diagnosis algorithms. While probability parameters are necessary in analyzing the probabilistic diagnosis algorithms, the diagnosis algorithm itself need not necessarily use all or even any of the parameters.

From Table 3, several patterns are apparent concerning the probabilistic diagnosis algorithms. The algorithms that guarantee the optimal diagnosis all have exponential computational complexity. Polynomial-time algorithms can only guarantee locally optimal diagnosis. At most $O(N \log N)$ testing links are required to guarantee asymptotically correct and complete diagnosis. The incomplete-test algorithms requiring $O(N)$ testing links are either multiple-syndrome or sequential diagnosis algorithms. Although the BSM1 [Blough et al. 1992a] algorithm requires only $O(N\ \omega(N))$ testing links, it is a complete-test algorithm that uses a special type of testing graph not normally used for computation purposes. Several of the probabilistic diagnosis algorithms do not require any probability parameters. However, these algorithms are not necessarily better than the algorithms that require 1 or more probability parameters since the latter algorithms with rough estimates of the necessary parameters may produce better diagnosis results than the former algorithms.

Besides the factors listed in Table 3, another factor that is important in evaluating a probabilistic diagnosis algorithm is diagnostic accuracy. However, it is difficult to quantify and compare this parameter since different probability and diagnosis models are used by the probabilistic diagnosis algorithms. Analyses of diagnostic accuracy can be found in some of the respective references, and comparisons of the diagnostic accuracies of several of the algorithms shown can be found in Lee [1990].

Since most practical computer systems will not be operational with large numbers of faulty nodes, it can be argued that the simplest probabilistic algorithm that works well with small numbers of faulty nodes should be used for diagnosis. However, since there is still the question of the validity of the complete-test assumption, the complete-test determinis-

tic and probabilistic diagnosis methods may not be desirable. The deterministic incomplete-test methods are also questionable because they tend to produce incomplete diagnoses in many cases. Thus, the polynomial-time algorithms in the incomplete-test section of Table 3 are all appropriate for fast diagnosis with small numbers of possibly intermittently faulty nodes. If a multiple-syndrome or sequential diagnosis algorithm is acceptable, the FR [Rangarajan and Fussell 1992], LS2 [Lee and Shin 1990a], and BLP [Blough and Pelc 1993] algorithms are the most suitable. For one-step diagnosis, the BSM2 [Blough et al. 1992b], DSK [Dahbura et al. 1987], and LS1 [Lee and Shin 1993] algorithms have similar execution times. Although the PEL [Pelc 1993] algorithm has lower computational complexity ($O(N)$), it has a fairly high constant factor and requires a special type of testing graph.

5. CONCLUSION

This article has surveyed probabilistic diagnosis methods, which in recent years have shown much promise of bridging the gap between the theory and practical application of system-level diagnosis ideas. All of the components necessary for the automated probabilistic diagnosis of large multiprocessor systems have been discussed. Additionally, the theory behind the probabilistic diagnosis methods, including testing models, probability models, and asymptotic diagnostic accuracy, has been presented. The various probabilistic diagnosis algorithms in the literature have been described in simple terms, and the algorithms presented have been analyzed and compared on the basis of several factors important to probabilistic diagnosis.

While significant progress has been made toward the development of diagnosis algorithms that are suitable for the diagnosis of large multiprocessor systems, their practicality and usefulness have yet to be demonstrated with physical experimental systems. Toward this end, Bianchini et al. [1990; 1992] recently presented the first application and implementation of theoretical diagnosability results to a real distributed network environment. Their results are limited to the original PMC model of diagnosis [Preparata et al. 1967]. More general experimental work of this nature needs to be conducted in the future to determine whether the newly proposed probabilistic diagnosis methods can be practical and useful.

REFERENCES

BARSI, F., GRADONI, F., AND MAESTRINI, P. 1976. A theory of diagnosability of digital systems. *IEEE Trans. Comput. C-25*, 6 (June), 585–593.

BERMAN, P., AND PELC, A. 1990. Distributed probabilistic fault diagnosis in multiprocessor systems. *Dig. Papers FTCS-20* (June), 340–346.

BIANCHINI, R. P., JR., AND BUSKENS, R. W. 1992. Implementation of on-line distributed system-level diagnosis theory. *IEEE Trans. Comput. 41*, 5 (May), 616–626.

BIANCHINI, R., JR., GOODWIN, K., AND NYDICK, D. S. 1990. Practical application and implementation of distributed system-level diagnosis theory. *Dig. Papers FTCS-20* (June), 332–339.

BLOUGH, D. M. 1988. Fault detection and diagnosis in multiprocessor systems. Ph.D. dissertation, The Johns Hopkins University, Baltimore, Md.

BLOUGH, D. M., AND PELC, A. 1993. Diagnosis and repair in multiprocessor systems. *IEEE Trans. Comput. 42*, 2 (Feb.), 205–217.

BLOUGH, D. M., AND PELC, A. 1992. Complexity of fault diagnosis in comparison models. *IEEE Trans. Comput. 41*, 3 (Mar.), 318–324.

BLOUGH, D. M., SULLIVAN, G. F., AND MASSON, G. M. 1992a. Efficient diagnosis of multiprocessor systems under probabilistic models. *IEEE Trans. Comput. 41*, 9 (Sept.), 1126–1136.

BLOUGH, D. M., SULLIVAN, G. F., AND MASSON, G. M. 1992b Intermittent fault diagnosis in multiprocessor systems. *IEEE Trans. Comput. 41*, 11 (Nov.), 1430–1441

BLOUNT, M. L. 1977. Probabilistic treatment of diagnosis in digital systems *Dig. Papers FTCS-7*, 72–77.

BUSKENS, R. W., AND BIANCHINI, R. P., JR. 1993. Distributed on-line diagnosis in the presence of arbitrary faults. *Dig. Papers FTCS-21* (June), 470–479.

DAHBURA, A. T. 1988. System-level diagnosis: A perspective for the third decade. In *Concurrent Computation: Algorithms, Architectures, Technologies* Plenum, New York.

DAHBURA, A. T., SABNANI, K. K., AND KING, L. L. 1987. The comparison approach to multipro-

cessor fault diagnosis. *IEEE Trans. Comput. C-36*, 3 (Mar.), 373–378.

FRIEDMAN, A. D. 1975. A new measure of digital system diagnosis. *Dig. Papers FTCS-5*, 167–170.

FRIEDMAN, A. D., AND SIMONCINI, L. 1980. System-level fault diagnosis. *Computer 13*, 3 (Mar.), 47–53.

FUSSELL, D., AND RANGARAJAN, S. 1989. Probabilistic diagnosis of multiprocessor systems with arbitrary connectivity. *Dig. Papers FTCS-19*, 560–565.

HAKIMI, S. L., AND AMIN, A. T. 1974. Characterization of connection assignment of diagnosable systems. *IEEE Trans. Comput. C-23*, 1 (Jan.), 86–88.

HAKIMI, S. L., AND NAKAJIMA, K. 1984. On adaptive system diagnosis. *IEEE Trans. Comput. C-33*, 3 (Mar.), 234–240.

HORST, R., JEWETT, D., AND LENOSKI, D. 1993. The risk of data corruption in microprocessor-based systems. *Dig. Papers FTCS-23* (June), 576–585.

HOSSEINI, S. H., KUHL, J. G., AND REDDY, S. M. 1988. On self-fault diagnosis of the distributed systems. *IEEE Trans. Comput. 37*, 2 (Feb.), 248–251.

HUANG, S., XU, J., AND CHEN, T. 1989. Characterization and design of sequentially t-diagnosable systems. *Dig. Papers FTCS-19* (June), 554–559.

KAVIANPOUR, A., AND FRIEDMAN, A. D. 1978. Efficient design of easily diagnosable systems. In the *3rd USA-JAPAN Computer Conference*. 251–257.

KAVIANPOUR, A., AND KIM, K. H. 1991. Diagnosabilities of hypercubes under the pessimistic one-step diagnosis strategy. *IEEE Trans. Comput. 40*, 2 (Feb.), 232–237.

KIME, C. 1986. System diagnosis. In *Fault-Tolerant Computing Theory and Techniques*. Vol. 2. Prentice-Hall, Englewood Cliffs, N.J.

KUHL, J. G., AND REDDY, S. M. 1980. Distributed fault-tolerance for large multiprocessor systems. In the *7th International Symposium on Computer Architecture*. ACM, New York, 23–30.

LEE, S. 1990. Probabilistic multiprocessor and multicomputer diagnosis. Ph.D. dissertation, The Univ. of Michigan, Ann Arbor, Mich.

LEE, S., AND SHIN, K. G. 1993. Optimal and efficient probabilistic distributed diagnosis schemes. *IEEE Trans. Comput. 42*, 7 (July), 882–886.

LEE, S., AND SHIN, K. G. 1990a. Optimal multiple syndrome probabilistic diagnosis. *Dig. Papers FTCS-20* (June), 324–331.

LEE, S., AND SHIN, K. G. 1990b. Interleaved all-to-all reliable broadcast on meshes and hypercubes. In the *1990 International Conference on Parallel Processing*. Vol. III. Pennsylvania State University, University Park, Pa., 110–113.

MAHESWARI, S. H., AND HAKIMI, S. L. 1976. On models for diagnosable systems and probabilistic fault diagnosis. *IEEE Trans. Comput. C-25*, 3 (Mar.), 228–236.

MALLELA, S., AND MASSON, G. M. 1980. Diagnosis without repair for hybrid fault situations. *IEEE Trans. Comput. C-29*, 6 (June), 461–470.

MALLELA, S., AND MASSON, G. M. 1978. Diagnosable systems for intermittent faults. *IEEE Trans. Comput. C-27*, 6 (June), 560–566.

PELC, A. 1993. Efficient distributed diagnosis in the presence of random faults. *Dig. Papers FTCS-23* (June), 462–469.

PREPARATA, F. P., METZE, G., AND CHIEN, R. T. 1967. On the connection assignment problem of diagnosable systems. *IEEE Trans. Elec. Comput. EC-16*, 6 (Dec.), 848–854.

RAMANATHAN, P., AND SHIN, K. G. 1988. Reliable broadcast in hypercube multicomputers. *IEEE Trans. Comput. 37*, 12 (Dec.), 1654–1657.

RANGARAJAN, S., AND FUSSELL, D. 1992. Diagnosing arbitrarily connected parallel computers with high probability. *IEEE Trans. Comput. 41*, 5 (May), 606–615.

RUSSELL, J. D., AND KIME, C. R. 1975. System fault diagnosis: Closure and diagnosability with repair. *IEEE Trans. Comput. C-24*, 11 (Nov.), 1078–1088.

SCHEINERMAN, E. 1987. Almost sure fault tolerance in random graphs. *SIAM J. Comput. 16*, 12 (Dec.), 1124–1134.

SIEWIOREK, D. P., AND SWARZ, R. S. 1982. *The Theory and Practice of Reliable System Design*. Digital Equipment Corporation, Bedford, Mass.

SOMANI, A. K., AND AGRAWAL, V. K. 1992. Distributed diagnosis algorithms for regular interconnected structures. *IEEE Trans. Comput. 41*, 7 (July), 899–906.

SOMANI, A. K., AND AGRAWAL, V. K. 1989. Distributed syndrome decoding for regular interconnected structures. *Dig. Papers FTCS-19* (June), 70–77.

SOMANI, A. K., AGRAWAL, V. K., AND AVIS, D. 1987. A generalized theory for system level diagnosis. *IEEE Trans. Comput. C-36*, 5 (May), 538–546.

YANG, C. L., AND MASSON, G. M. 1988. A distributed algorithm for fault diagnosis in systems with soft failures. *IEEE Trans. Comput. 37*, 11 (Nov.), 1476–1479.

YANG, C. L., AND MASSON, G. M. 1986. A fault identification algorithm for t_1-diagnosable systems. *IEEE Trans. Comput. C-35*, 6 (June), 503–510.

Received August 1991; final revision accepted October 1993.