A Class of 2-Step Diagnosable Systems: Degree of Diagnosability and a Diagnosis Algorithm

Kaiyuan Huang
McGill University

Anindya Das
University of Montreal
Vinod K. Agarwal
McGill University

K. Thulasiraman
Concordia University

Abstract

A new diagnosability measure, $t/-1$-diagnosability, is introduced. A system is $t/-1$-diagnosable if 1) at least $f-1$ faulty units can be identified as long as the number of faulty units present, f, does not exceed t and 2) the system is 1-fault diagnosable. This class of diagnosable systems are fully characterized. In addition, an $O(n^{3.5})$ diagnosis algorithm is provided, which locates at least $f-1$ faulty units when $1 < f \leq t$ or the only faulty unit when $f = 1$. When there is no faulty unit, the algorithm certifies it. The algorithm is suited for any $t/-1$-diagnosable system.

1 Introduction

Research on system level diagnosis was pioneered by the work of Preparata, Metze and Chien [8]. They suggested that a system of interconnected computing units be diagnosed by first making the units of the system test each other and then analyzing the outcomes of these tests. Test outcomes are classified as fault-free or faulty. The set of test outcomes is called the *syndrome* of the system. All units are considered to be identical. They can test others or be tested by others. No postulate is to be made in the process of test outcome analysis either on the status (fault-free or faulty) of any of the units or on the correctness of any of the test outcomes produced by the testing units. It is assumed that test outcomes produced by fault-free testing units are always correct while those produced by faulty testing units can be anything (fault-free or faulty), irrespective of the status of the tested units. This kind of test outcome interpretation has since been known as the PMC model. They also introduced two diagnosability criteria, the *one-step t-diagnosability* and *sequential t-diagnosability*. A system is said to be *one-step t-diagnosable* if all faulty units can be identified from any syndrome produced by the system as long as the number of faulty units present does not exceed t. Similarly, a system is said to be ***sequentially t-diagnosable*** if at least one faulty unit can be identified from any syndrome produced by the system as long as the number of faulty units present does not exceed t. Friedman [5] introduced another measure of diagnosability, called t/s-diagnosability. A system is *t/s-diagnosable* if all faulty units can be located to within a set of no more than s units, as long as the number of faulty units present does not exceed t. Main contributions to this area of system level diagnosis include Hakimi and Amin's [6] characterization of t-diagnosable systems, Dahbura and Masson's [2] $O(n^{2.5})$ t-diagnosis algorithm, Sullivan's [11] and Raghavan's [9] diagnosability algorithms, characterization of sequentially t-diagnosable systems by Huang et al [7] and t/s-diagnosis algorithm by Das et al [3]. For the successful application of the t-diagnosability measure, a large number of tests between units are required. However, in existing multiprocessor systems such as the hypercube, connection is very limited. To overcome this limitation of classical diagnosability measure, Somani, Agarwal and Avis [10] and Das et al [4] proposed approaches which allow large numbers of faulty units even in sparsely connected systems. It is shown in [8] that even in a single loop architecture, a large number of faulty units, up to the square root of the number of units in the system can be sequentially diagnosed.

We will introduce in this paper a new diagnosing strategy — we locate all faulty units except for at most one which may be left unidentified, in which case a second step of diagnosis is needed to locate the remaining faulty unit. We call such a system a *$t/-1$-diagnosable system.* Just as t/t-diagnosability is on the boundary of t/s-diagnosability and one-step t-diagnosability, $t/-1$-diagnosability is on the boundary of multiple step diagnosability and one-step t-diagnosability. No more than two steps are needed to locate all faulty units.

0-7803-0593-0/92 $3.00 1992 IEEE

2 Preliminaries

A system is represented by its *test digraph* $D(V, A)$, where V is a set of vertices each of which corresponds to a unit of the system and A is a set of test arcs which correspond to inter-unit tests. A test arc (u, v) is in A if and only if u tests v. The *tester set* $\Gamma^{-1}(v)$ of unit v is the set of all the units which perform tests on unit v. Analogously, the tester set $\Gamma^{-1}(V')$ for a subset $V' \subseteq V$ is the set of those units which perform tests on some members of V' but are not themselves members of V'. The PMC model[8] is taken for interpreting test outcomes. A *fault set* is a set of units which are assumed to be faulty. Each fault set $F \in V$ stands for a unique system state. The collection of all possible fault sets or system states is denoted by U, which is the power set of V. As in the above-mentioned work, we will consider only those fault sets which contain no more than t faulty units. This universe is called the *collection of t-fault sets* and denoted by U_t. A *syndrome* of a system is the entire set of test outcomes. A fault set is said to be *consistent* with a syndrome if it can possibly produce this syndrome. Similarly, a family of fault sets $\mathcal{F} \subseteq U$ is said to be consistent with a syndrome if this syndrome is producible from every member of $\mathcal{F}$. We use the notation $\mathcal{F}(v)$ to associate with unit v the sets in $\mathcal{F}$ which contain unit v. There should not be any confusion from the context as to whether $\mathcal{F}$ represents a family of sets or a function. Let U_t be the collection of sets containing t or less vertices. This will be our scope of consideration for families of sets.

3 $t/-1$-Diagnosability and Characterization

In this section, we fully characterize the class of $t/-1$-diagnosable systems.

Definition 3.1 *A system is said to be* $t/-1$-diagnosable *if 1) it is 1-diagnosable and 2) when there are $f \leq t$ faulty units present at least $f-1$ of them can be identified.*

All proofs of the following results are omitted to save space.

Theorem 3.1 *A family $\mathcal{F}$ is consistent if and only if the following is satisfied:*

1. *If $\mathcal{F}$ does not contain the empty set, then for every test arc $(u, v) \in A$*

$$\mathcal{F}(v) \subseteq \mathcal{F}(u) \text{ or } \mathcal{F}(u) \cup \mathcal{F}(v) = \mathcal{F}.$$

2. *If $\mathcal{F}$ contains the empty set, then for every nonempty set $F \in \mathcal{F}$*

$$\Gamma^{-1}(F) = \emptyset.$$

Theorem 3.2 *A system is $t/-1$-diagnosable if and only if the following conditions are satisfied:*

1. *$\Gamma^{-1}(V')$ is not empty for every nonempty subset $V' \subseteq V$ of cardinality less than or equal to t. In addition, if $t = 1$, then $\Gamma^{-1}(V')$ must not be empty for every subset $V' \subseteq V$ of cardinality two.*
2. *If $\left|\bigcap_{F\in\mathcal{F}} F\right| < \max_{F\in\mathcal{F}}(|F|) - 1$ for some $\mathcal{F} \subseteq U_t$ with $\emptyset \notin \mathcal{F}$, then there exists a test arc $(u, v) \in A$ such that*

$$\mathcal{F}(v) \not\subseteq \mathcal{F}(u) \text{ and } \mathcal{F}(u) \cup \mathcal{F}(v) \neq \mathcal{F}.$$

Theorem 3.3 *A system is $t/-1$-diagnosable if and only if, for any subset $V' \subseteq V$, the following is satisfied:*

1. *$|\Gamma^{-1}(V')| > 0$ for $|V'| = 1$;*
2. *$|\Gamma^{-1}(V')| > \max(t-2, 0)$ for $|V'| = 2$;*
3. *$|\Gamma^{-1}(V')| > t - \lceil |V'|/2 \rceil$ for $|V'| \geq 3$.*

4 A Diagnosis Algorithm for $t/-1$-Diagnosable Systems

In this section, we present a polynomial time diagnosis algorithm for the class of $t/-1$-diagnosable systems. This algorithm will identify at least $f-1$ faulty vertices if $f > 1$ or the only faulty vertex if $f = 1$, where f is the number of faulty vertices, as long as the system is $t/-1$-diagnosable and f does not exceed t. In the following we will use some additional graph theory terms. Those readers who are not familiar with these terms are referred to [1]. For a given syndrome and a test digraph D(V,A), we construct the *implied fault graph* $G_L(V, E)$[2], which is an undirected graph and an edge (u, v) is in E if and only if u can be deduced to be faulty on the assumption that v is fault-free. Then we compute a minimum vertex cover of this graph and determine from it at least $f-1$ faulty vertices when $f > 1$ or the only faulty vertex when $f = 1$.

Suppose we can, by some means, find a minimum vertex cover of the implied fault graph G_L. The following algorithm can then be utilized to locate the faulty vertices. It identifies at least $f-1$ faulty vertices when $f > 1$ or the only faulty vertex when $f = 1$, as long as the number of faulty vertices present does not exceed t. Of course, the system should be $t/-1$-diagnosable.

Algorithm 4.1

Input: *Given a minimum vertex cover $\hat{K}$ and implied fault graph G_L.*

Output: *Nonempty fault set $\hat{F} \subseteq F$ with $|\hat{F}| \geq |F| - 1$ when $|F| > 1$ and $|\hat{F}| = 1$ when $|F| = 1$, where F is the set of faulty vertices present.*

Step 1: *If $|\hat{K}| = 1$ or $|\hat{K}| < t$, set $\hat{F} = \hat{K}$ and go to End.*

Step 2: *If $|\hat{K}| = t$, determine if there is a vertex $v \in \hat{K}$ such that $|N(v) - \hat{K}| = 1$. If "yes", set $\hat{F} = \hat{K} - \{v\}$. Go to End.*

Step 3: *Set $\hat{F} = \hat{K}$.*

End

All proofs of the following results are omitted to save space.

Theorem 4.1 *For any given syndrome produced by a nonempty set F of not more than t faulty vertices, a nonempty subset of faulty vertices $\hat{F} \subseteq F$ with $|\hat{F}| \geq |F| - 1$ can be identified by an application of Algorithm 4.1 on the implied fault graph G_L and any minimum vertex cover of G_L, as long as the system is $t/-1$-diagnosable. The complexity of the algorithm is $O(e)$, where e is the number of edges in the implied fault graph G_L.*

The implied fault subgraph with respect to vertex $v \in V$, denoted by G_L^v, is a subgraph of G_L formed by removing from G_L vertex v and all edges incident on v. The neighbor set of a vertex u in G_L^v is denoted by $N^v(u)$ and the neighbor set of u in G_L^v with respect to F (the set of neighbors which are not in F) is denoted by $N_F^v(u)$.

The following labeling procedure, similar to that used by Dahbura and Masson in their diagnosis algorithm for t-diagnosable systems[2], partitions the set of vertices into two subsets, the set of vertices supposed to be fault-free and the set of those vertices supposed to be faulty. The procedure may fail for some vertices. It reports whether or not it is successful when it terminates.

Procedure $LABEL(v)$

Input: Implied fault subgraph G_L^v with respect to v, neighbor set $N(v)$ of v and maximum matching $\hat{M}^v$ of G_L^v.

Output: Return status (successful or failed), labels (fault-free or faulty) on the vertices in G_L^v and vertex cover K^v of G_L^v.

Step 1: Initialize stack S.

Step 2: Mark all $\hat{M}^v$-unsaturated vertices "fault-free" and place them on S. Set the stack top pointer Top equal to the number of these vertices. If $Top = 0$, set the return status as "failed" and go to End.

Step 3: If $Top = 0$, set the return status as "successful" and go to Step 6. Remove from the stack the vertex pointed to by Top. Let it be v_i. Set $Top = Top - 1$. If v_i is marked "faulty", go to Step 5.

Step 4: Let D be the set of vertices in $N^v(v_i)$ already marked. If any vertex in D is marked "fault-free", set the return status as "failed" and go to End; otherwise, mark all vertices in $N^v(v_i) - D$ "faulty", place them on S and set $Top = Top + |N^v(v_i) - D|$. Go to Step 3.

Step 5: Let v_j be the vertex matched by v_i under $\hat{M}^v$. If v_j is marked "faulty", set the return status as "failed" and go to End. If v_j is unmarked, mark it "fault-free", place it on S and set $Top = Top + 1$. Go to Step 3.

Step 6: Let U be the set of vertices in G_L^v left unmarked. If $|U| = 0$, go to Step 7. Otherwise, U contains exactly two vertices matched by each other under $\hat{M}^v$. Let them be u_1 and u_2. Mark one vertex adjacent to v in G_L "faulty" if there is any; otherwise, mark any one of them "faulty". Mark the other vertex "fault-free".

Step 7: Place all vertices marked "faulty" into K^v.

End

Now we are ready to present our algorithm for finding a minimum vertex cover of the implied fault graph.

Algorithm 4.2 (Minimum Vertex Cover)

Input: *Implied fault graph G_L.*

Output: *Minimum vertex cover $\hat{K}$ of G_L.*

Step 1: *Set $s = n$.*

Step 2: *For each vertex v in G_L do:*

1. *Form subgraph $G_L^v = G_L - v$ by removing from G_L vertex v and all edges incident on v.*
2. *Compute a maximum matching $\hat{M}^v$ of G_L^v.*
3. *Apply $LABEL(v)$ to partition $V - \{v\}$ into two subsets K_v and $V - \{v\} - K^v$.*

4. *If the status returned from LABEL(v) is "successful", then do the following:*
   (a) *If $N(v) - K^v \neq \emptyset$ then set $K' = K^v \cup \{v\}$; otherwise, set $K' = K^v$.*
   (b) *Set $s_v = |K'|$.*
   (c) *If $s_v < s$, then set $\hat{K} = K'$ and $s = s_v$.*

End

Theorem 4.2 *Algorithm 4.2 can correctly compute a minimum vertex cover of G_L, as long as the system is $t/-1$-diagnosable and the number of faulty vertices present does not exceed t. The execution time of Algorithm 4.2 is upper bounded by $O(n^{3.5})$.*

The whole diagnosis algorithm for $t/-1$-diagnosable systems is formally described below.

Algorithm 4.3 (Fault Identification)

Input: *Test graph $D(V,A)$ and syndrome Syn.*

Output: *Fault set $\hat{F} \subseteq F$ such that $|\hat{F}| \geq |F| - 1$ when $|F| > 1$ and $\hat{F} = F$ when $|F| = 1$, where F is the set of faulty vertices present.*

Step 1: *If the test outcomes are all "0", set $\hat{F} = \emptyset$ and go to End. Construct the implied fault graph $G_L(V,E)$ from $D(V,A)$ and Syn.*

Step 2: *Find a minimum vertex cover $\hat{K}$ of G_L (using Algorithm 4.2).*

Step 3: *Determine a fault set $\hat{F}$ (using Algorithm 4.1).*

End

Theorem 4.3 *Algorithm 4.3 Identifies at least $f-1$ faulty vertices when $f > 1$ or the only faulty vertex when $f = 1$, provided the system is $t/-1$-diagnosable and the number of faulty vertices f is no larger than t. The algorithm can be executed in $O(n^{3.5})$ time.*

5 Conclusions

We have introduced a new diagnosability measure, the $t/-1$-diagnosability, and presented a complete characterization of $t/-1$-diagnosable systems. We have also given an $O(n^{3.5})$ $t/-1$ diagnosis algorithm. The $t/-1$-diagnosable systems have the following characteristics: 1) It is two-step diagnosable; 2) the degree of diagnosability may double that of one-step t-diagnosability and 3) if a system is both $t/-1$-diagnosable and 1-fault tolerant, one phase of diagnosis and repair can bring a failed system back into operation. This new diagnosability is a compromise between one-step t-diagnosability and sequential t-diagnosability.

References

[1] J. A. Bondy and U. S. R. Murty. *Graph Theory with Applications.* Elsevier North Holland, New York, 1976.

[2] A. T. Dahbura and G. M. Masson. An $O(n^{2.5})$ fault identification algorithm for diagnosable systems. *IEEE Trans. Comput.*, c-33(6):486–492, Jun. 1984.

[3] A. Das, K. Thulasiraman, V. K. Agarwal, and K. B. Lakshmanan. t/s-diagnosable systems: A characterization and diagnosis algorithm. In *Proc. 15th Int. Conf. Graph-Theoretic Concepts in Computer Science*, Rolduc, Holland, 1989.

[4] A. Das, K. Thulasiraman, K. B. Lakshmanan, and V. K. Agarwal. System level diagnosis with local constraints. In *Proc. IEEE Int. Symp. Circ. and Syst.*, pages 1891–1894, Portland, Oregon, 1989.

[5] A. Friedman. A new measure of digital system diagnosis. In *Digest FTCS-5*, pages 167 – 170, Jun. 1975.

[6] S. L. Hakimi and A. T. Amin. Characterization of connection assignment of diagnosable systems. *IEEE Trans. Comput.*, c-23(1):86–88, Jan. 1974.

[7] S. Huang, J. Xu, and T. Chen. Characterization and design of sequentially t-diagnosable systems. In *Digest FTCS-19*, pages 554–559, Jun. 1989.

[8] F. P. Preparata, G. Metze, and R. T. Chien. On the connection assignment problem of diagnosable systems. *IEEE Trans. Electron. Comput.*, ec-16(12):848–854, Dec. 1967.

[9] V. Raghavan. *Diagnosability Issues in Multiprocessor Systems.* PhD thesis, University of Minnesota, 1989.

[10] A. Somani, V. Agarwal, and D. Avis. A generalized theory for system-level diagnosis. *IEEE Trans. Comput.*, c-36(5):538–546, May 1987.

[11] G. F. Sullivan. A polynomial time algorithm for fault diagnosability. In *Proc. 25th Symp. Found. of Comput. Sci.*, pages 148–156, 1984.