Correct Diagnosis of Almost All Faulty Units in a Multiprocessor System

K. Thulasiraman[1] Anindya Das[1] Kaiyuan Huang[2]

Vinod K. Agarwal[3]

Abstract

In a *t/t-diagnosable* system all faulty units can be located to within a set of no more than t units as long as the number of faulty units present does not exceed t. Furthermore, a unique doubtful unit can be identified; in other words, all faulty units, except possibly for one, can be correctly identified in a *t/t*-diagnosable system. An open question is "is *t/t*-diagnosability necessary for correctly identifying all but one faulty unit?" In this paper we address the above question and provide an answer. We establish necessary and sufficient conditions for correct diagnosis of all except possibly one faulty unit. In addition, we show that the faulty-free state is indistinguishable from a faulty state in a *t/t*-diagnosable system and propose a remedy. These considerations result in the definition and characterization of a new class of systems called *t/*-1 diagnosable systems.

1 Introduction

Research on system level diagnosis was pioneered by the work of Preparata, Metze and Chien [1]. 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. Each unit can test others or be tested by others. 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. Diagnosis is *correct* if every unit identified as faulty is indeed faulty and *complete* if every faulty unit is identified as so. Preparata et al. 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. One-step t-diagnosis is both correct and complete. Sequential t-diagnosis is correct but incomplete.

Since then, a lot of work has been done in this area. Among the contributors, Hakimi and Amin[2] presented the first full characterization of one-step t-diagnosability. An $O(n^{2.5})$ diagnosis algorithm was presented by Dahbura and Masson[3] for one-step t-diagnosable systems. Friedman [4] introduced a new measure of diagnosability, the *t/s-diagnosability,* allowing a certain number of fault-free units to be identified incorrectly as faulty while all faulty units are correctly identified. A system is said to be *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. *t/s*-diagnosis is complete but incorrect. The class of *t/s*-diagnosable systems has been studied by Das et al.[5,6] and Raghavan[7].

Kavianpour and Friedman[8] considered a very interesting special case of *t/s*-diagnosability, the *t/t-diagnosability.* In a *t/t*-diagnosable system, any set of not more than t faulty units can be located to within a set of no more than t units. They showed that with the same degree of connection, the degree of *t/t*-diagnosability might double the degree of t-diagnosability. Chwa and Hakimi[9] presented a characterization for *t/t*-diagnosable systems. Yang, Masson and Leonetti[10] provided a diagnosis algorithm for *t/t*-diagnosable systems. They further showed that in a *t/t*-diagnosable system there is at most one doubtful unit and their diagnosis algorithm can identify it if this is the case. In other words, all but one faulty unit can be correctly identified in a *t/t*-diagnosable system.

An interesting question is: Is *t/t*-diagnosability necessary for correctly identifying all but one faulty unit and if not what are the necessary and sufficient conditions? As will be shown in this paper, *t/t*-diagnosability is not necessary for this purpose. The necessary and sufficient conditions for correctly identifying all but one faulty unit will be provided.

A minor problem with *t/t*-diagnosis, as well as with *t/s*-diagnosis, is that the fault-free state is indistinguishable from a faulty state. As an example, consider a system which is one-step t'-fault diagnosable. Add an additional unit v to the system such that v tests some of the other units but is not tested by any other unit in the system. Let $t = t' + 1$. Obviously, the system is *t/t*-diagnosable and unit v is always identified as faulty even when the system is completely fault-free.

[1] School of Computer Science, Univ. of Oklahoma, Norman, Oklahoma, USA.
[2] Nortel, Ottawa, Canada.
[3] LV Software, San Jose, California, USA.

0-7803-5471-0/99/$10.00©1999 IEEE

This problem arises because *t/t*-diagnosability does not require that every unit be tested by at least one other unit.

We will introduce a new diagnosis strategy - locating all but one faulty unit and certifying a fault-free system when every unit is fault-free. We call a system with such characteristics *t/*-1-*diagnosable*. More precisely, a system is *t/* - 1-*diagnosable* if 1) at most one faulty unit may evade identification, provided the number of faulty units does not exceed *t* and 2) the system is 1-fault diagnosable. With such a system, we can identify all but one faulty unit. Those units identified as faulty are repaired or replaced. A second step of diagnosis is then carried out. If the replacement parts are good, the possibly remaining faulty unit is identified. Therefore, the system is two-step diagnosable. 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

Full characterizations of *t/* - 1-diagnosable systems will be presented in Section III.

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 terms unit and vertex will be used interchangeably, as well as the terms test and test arc. 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'. For example, consider the test digraph shown in Figure 1. There are five vertices and each of them is tested by two adjacent vertices. The tester set $\Gamma^{-1}(\{v_0,v_1\})$ is identical to $\{v_3,v_4\}$. Faults are assumed to be permanent and classified as fault-free or faulty, which are represented by the binary values 0 and 1 respectively. The PMC model [1] is taken for interpreting test outcomes. That is, a fault-free unit evaluates a unit to be faulty if and only if the tested unit is faulty while a test outcome produced by a faulty testing unit can be arbitrary, irrespective of the status of the tested unit. A *fault set* is a set of units which are assumed to be faulty. Each fault set $F \subseteq V$ stands for a unique system state. In particular, the empty set ϕ represents the normal fault-free 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. As a test outcome produced by a faulty unit is arbitrary, a lot of different syndromes may be producible from the same fault set. Again, consider Figure 1. There are two faulty units, v_1 and v_3. The test outcomes are noted on the arcs. For a faulty testing unit, the outcome it produces may be either 0 or 1. There are 16 possible syndromes corresponding to this fault situation. On the other hand, different fault sets may produce the same syndrome. A fault set is said to be *consistent* with a syndrome if it can possibly produce this syndrome. It is not difficult to see that a fault set F is consistent with a syndrome if and only if every test into F has a 1 outcome and every test of which neither end is in F has a 0 outcome. 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. For instance, consider a family of sets $\mathcal{F} = \{\{v_0, v_1\}, \{v_1, v_2\}, \{v_1,v_3\}\}$. In this case, we have $\mathcal{F}(v_1) = \{\{v_0, v_1\}, \{v_1, v_2\}, \{v_1,v_3\}\}$and $\mathcal{F}(v_2) = \{\{v_1, v_2\}\}$. There should not be any confusion from the context as to whether $\mathcal{F}$ represents a family of sets or a function.

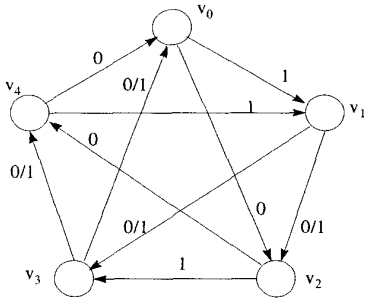

Figure 1: An example of a test digraph

3 t/ - 1-Diagnosability and Characterization

In this section we will give the necessary and sufficient conditions for correctly identifying all but one faulty unit and certifying the fault-free state. We will then introduce a new diagnosability measure which captures these characteristics. But first, we give necessary and sufficient conditions for a family of fault sets to be consistent.

Lemma 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 non-empty set* $F \in \mathcal{F}$

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

Proof:

(Necessity) Assume that $\mathcal{F}$ is consistent. If $\mathcal{F}$ does not contain the empty set and there is a test arc $(u, v) \in A$ such that $\mathcal{F}(v) \not\subset \mathcal{F}(u)$ and $\mathcal{F}(u) \cup \mathcal{F}(v) \neq \mathcal{F}$, then there is a set F' in $\mathcal{F}(v)$ which is not in $\mathcal{F}(u)$ and a non-empty set F'' in $\mathcal{F}$ which is in neither of $\mathcal{F}(u)$ and $\mathcal{F}(v)$ as shown in Figure 2. In order for F' to be a consistent fault set, this test arc must assume the outcome of "1". On the other hand, the test outcome can only assume the value of "0" for F'' to be a consistent fault set. So, whatever value this test takes as its outcome, either F' or F'' *is* not consistent. Hence $\mathcal{F}$ is not consistent, a contradiction. On the other hand, if $\mathcal{F}$ contains the empty set and there is a set $F \in \mathcal{F}$ with $\Gamma^{-1}(F) \neq \phi$, F and ϕ cannot be both consistent with a common syndrome, again contradicting the hypothesis that $\mathcal{F}$ is consistent. (Note: If ϕ is a consistent fault set then all test outcomes must be 0.)

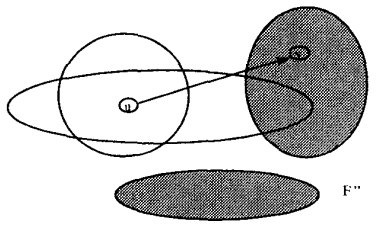

Figure 2: An inconsistent family of fault sets

(Sufficiency) If the conditions are satisfied, we can construct a syndrome which is producible from every member of $\mathcal{F}$. We consider the following cases.

Case 1: $\mathcal{F}$ does not contain the empty set. By assumption, condition 1 is satisfied. For every test arc $(u, v) \in A$, if $F(v) \subseteq F(u)$, we assign to it the outcome of "0"; otherwise, we have $F(u) \cup F(v) = F$ and we assign to the arc the outcome of "1". It is easy to see that this assignment of test outcomes contributes to a syndrome consistent with all the members of $\mathcal{F}$.

Case 2: $\mathcal{F}$ contains the empty set and every member of $\mathcal{F}$ has an empty tester set. Assign to all test arcs the outcome of "0". Thus every member of $\mathcal{F}$, when considered to be the fault set, can produce this syndrome. This completes the proof of the lemma.

For every syndrome, there is a family of fault sets which are consistent with it. The actual fault set is obviously a member of this family. To ensure that all but one faulty vertex is correctly identified from a consistent family of more than one fault set, the number of vertices common to all members of this family must not be less than the number of vertices in the largest member minus one.

Theorem 3.1

1. *The fault-free state is distinguishable from every faulty state of no more than t faulty vertices if and only if* $\Gamma^{-1}(V')$ *is not empty for any nonempty subset* $V' \subseteq V$ *of cardinality less than or equal to t.*

2. *All except one faulty vertex can be correctly identified in the presence of no more than t* faulty *vertices if and only if the following condition is satisfied:*
   If $|\cap_{F \in \mathcal{F}} F| < \max_{F \in \mathcal{F}}(|F|) - 1$ *for some* $\mathcal{F} \subseteq U_t$ *with* $\phi \notin \mathcal{F}$ *then there exists a test arc* $(u, v) \in A$ *such that*
   $F(v) \not\subseteq F(u)$ *and* $F(v) \cup F(u) \neq F$.

Proof: The correctness proof follows directly from Lemma 3.1 and the above discussions, and hence omitted.

Furthermore, the above characterization can be put in a way similar to the one-step t-diagnosable system characterization of Allan et. al [11].

Theorem 3.2 *The fault-free state can be certified and all but one faulty vertex can be correctly identified in the presence of no more than t faulty vertices if and only if the following is satisfied for any subset* $V' \subseteq V$:

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

Proof: **(Necessity)**

1. By Theorem 3.1 Condition 1 is necessary for the fault-free state to be distinguishable.
2. Assume that $|\Gamma^{-1}(V')| \leq t - 2$ for some $V' \subseteq V$ with $|V| = 2$. Let $F' = \Gamma^{-1}(V')$, $F'' = V' \cup \Gamma^{-1}(V')$and $\mathcal{F} = \{F', F''\}$. Now we have the following:
   (a) $|F|, |F''| \leq t$,
   (b) $|F' \cap F''| = |F| < \max(|F|, |F''|) - 1 = |F| + 1$ and
   (c) Condition 2 of Theorem 3.1 is not satisfied, for a test arc (u,v), $u \in F'$, $v \in V'$.
3. If $|\Gamma^{-1}(V')| + \lceil |V|/2 \rceil \leq t$ for some subset $V' \subseteq V$ of cardinality larger than 2, then we can partition V' into two non-empty subsets V_1 and V_2 such that $|V_1 \cup \Gamma^{-1}(V')| \leq t$ and $|V_2 \cup \Gamma^{-1}(V')| \leq t$. Let $F' = V_1 \cup \Gamma^{-1}(V')$, F'' = and $V_2 \cup \Gamma^{-1}(V')$ and $\mathcal{F} = \{F', F''\}$. Note that $|F' \cap F''| = |\Gamma^{-1}(V')| < \max(|F|, |F''|) - 1$ as V' contains more than 2 vertices. We can see that, for every test $(u, v) \in A$, either $F(v) \subseteq F(u)$ or $F(u) \cup F(v) = F$. Condition 2 of Theorem 3.1 is not satisfied.

(Sufficiency) Condition 1 of Theorem 3.1 is clearly satisfied if we have the conditions of this theorem. For any family of sets $\mathcal{F} \subseteq U_t$ with $\phi \notin \mathcal{F}$, we consider the following cases.

1. Assume that there are two sets $F', F'' \in \mathcal{F}$ such that $|F' - F''| \geq 2$. Without loss of generality, assume that $|F' - F''| \geq |F'' - F|$. Note that $|F' \cap F''| \leq t - |F' - F''|$ and $|F' \oplus F''| \leq 2|F' - F''|$. By conditions 2 and 3, we have $\Gamma^{-1}(F' \oplus F'') - (F' \cup F'') \neq \phi$. This means that there is a test $(u, v) \in A$ with F' (or F'') in $F(v)$, F' (or F'') not in $F(u)$, and F'' (or F') not in $F(u) \cup F(v)$. That is, condition 2 of Theorem 3.1 is satisfied.
2. On the other hand, assume that there are no two members F' and F'' of $\mathcal{F}$ such that $|F' - F''| \geq 2$. We have $|F' - F''| \leq 1$ for any $F', F'' \in \mathcal{F}$. Let f_{max} be the cardinality of the largest sets in $\mathcal{F}$, F_{max} be one of the largest members of $\mathcal{F}$ and f_{min} be the cardinality of the smallest members of $\mathcal{F}$. It is obvious that $f_{max} - f_{min} \leq 1$.
   (a) Assume $f_{max} - f_{min} = 1$. Let F'_{min} be one of the smallest sets. One can easily see that $F'_{min} \subset F_{max}$. If there is no other vertex set in $\mathcal{F}$, then F'_{min} is a desired fault set of size $f_{min} = f_{max} - 1$when $\mathcal{F}$ happens to be the consistent family of fault sets. If there is another vertex set, say F, in $\mathcal{F}$, it can be of size either f_{min} or f_{max} If it is of size f_{max}, it must contain all the vertices in F'_{min}. Therefore, we still have a desired fault set when $\mathcal{F}$ happens to be the consistent family of fault sets. This argument can be repeated for all sets of cardinality f_{max}. Suppose there is a set $F'' \in \mathcal{F}$ of cardinality f_{min}. It must be included in F_{max} and contains the only vertex in F_{max} which is not in F'_{min}. Let this vertex be u. There is also a vertex in F'_{min} which is not in F''. Let this one be v. From condition 2, we know that there is a test either on vertex u or on vertex v performed by a vertex outside of $F'_{min} \cup F''$, which means that condition 2 of Theorem 3.1 is satisfied.
   (b) Assume, otherwise, $f_{max} - f_{min} = 0$. If there are two sets, say, F' and F'' in $\mathcal{F}$, then $|F' \cap F''| = |F'| - 1$. $F' \cap F''$ is a desired fault set if $\mathcal{F}$ contains no more sets of vertices. If there is a

third member in F, say $F \in$ F, it contains exactly | F' | - 1 vertices of F' and | F' | - 1 vertices of F'' and it is of the same cardinality as F' and F''. If F contains $F' \cap F''$, we have a desired fault set when F happens to be the family of consistent fault sets. This argument can be repeated for all other members of F which contain $F' \cap F''$. Let $\bar{F}$ be a member of F $F' \cap F''$. Let {u F'' - F' = { }. $F$1 - 1 members of F' and must contain both vertex u v and | | - 2 members of F' F''.

$\bar{F}$ is included in $F'U F''$. We also note that there is a vertex in $F' \cap F''$ which is not contained in $\bar{F}$ Let this vertex be w From Condition 3 of the theorem, we can see that there is a test on one of the three vertices , v w performed by a vertex outside of F', F'' and $\bar{F}$. This means that condition 2 of Theorem F.

Theorem 3.1 are satisfied as long as the conditions of this theorem are satisfied. Hence the sufficiency of this theorem is justified.

As we have seen, Conditions 2 and 3 of Theorem 3.2 are necessary and sufficient for identifying all but one faulty unit in the system. This means that *t/t* diagnosability is not necessary although sufficient for correct diagnosis of all except possibly one faulty unit. *tl-*

conditions to those required for *t/t* diagnosability. The necessary and sufficient conditions for -diagnosability given by Hakimi [9] can be restated as follows (they gave it in a different but equivalent form): A system is -diagnosable if and only if | $^{-1}$( ) | > -⌈ $V$1/2 for any subset V' V with | | ≥ the only difference is in the case when | V' *t/t-* Γ^{-1} V') | > - 1 while in our case | $^{-1}$( )| > t - 2. In addition, Condition 1 of Theorem 3.2 makes the fault distinguishable. Furthermore, if the system is 1-diagnosable, the but one faulty unit and repairing them. These characteristics lead to a new

Definition 3.1 *A system is said to be t* - 1-diagnosable it *is -diagnosable and 2) when there are f ≤ t least f* - 1

Roughly, Theorem 3.2 can be viewed as a characterization of *t* - 1-diagnosability. The only exception is that when t = 1 the *t/* By Condition 2 of the theorem, Γ (V' $V$1 = 2. -diagnosability [2], a system is -diagnosable if and only if |Γ (V' $V' \subseteq$, | V' = 1, 2. The following result follows directly:

Theorem 3.3 - 1 *subset* $V' \subseteq$

1. |Γ (V' $V$1 = 1;
   |Γ (V' - 2, 0) for | V'
3. |Γ (V' -⌈ $V$1/2 for | V' ≥ 3.

Conclusion

In this paper we presented the necessary and sufficient conditions

Theorem 3.3 and the characterization of *t t*-diagnosable systems *t/t*-diagnosable system is also /-1-diagnosable provided each unit is tested by at least one other unit.

[1] F. P. Metze, and R. T. Chien, "On the connection *IEEE Trans. Electron.* ., vol. EC -854, Dec. 1967.

[2] S. L. Amin, "Characterization of connection assignment of diagnosable systems," vol. C-23, pp. 86

[3] A. T. Dahbura and G. M. $O(n^{2.5}$ fault identification algorithm for diagnosable systems," *on Computing,* vol. C -492, Jun. 1984.

[4] *Digest FTCS-5,* - 170, Jun. 1975.

[5] A. Das, K. Agarwal, and K. B. Lakshmanan, "Diagnosis of -diagnosable systems," Journal of Circuits, Systems and Computers, vol. 1, no. 4, pp. 353-371,

[6] A. Das, K. Thulasiraman, and V. K. "Diagnosis of t/t+1-Diagnosable Systems", *SIAM on Computing*, vol. 23, Oct. 1994, pp. 895-905.

Raghavan, *Diagnosability Issues in Multiprocessor Systems.* thesis, University of Minnesota, 1989.

[8] Kavianpour and A. D. Friedman, "Efficient design of easily *Proc. 3rd USA Comput.* ., pp. 251

[9] K. Chwa and S. systems," *IEEE Trans. on Computing,* -3O, pp. 414-422,

[10] C. Yang, G. M. Masson, and R. and identification in t1/tl-diagnosable systems," *Computing,* vol. C -643, Jul. 1986.

[11] F. J. Allan, T. Kameda, and S. diagnosability analysis of a system," *IEEE Trans. on Computing,* -24, pp. 1040-1042, Oct. 1975.