

PROPERTIES OF c -CORRECTABILITY IN SELF-DIAGNOSING SYSTEMS

JON T. BUTLER*

Department of Electrical
and Computer Engineering
Naval Postgraduate School, Code 62-Bu
Monterey, CA 93943-5004

YONINA ROSEN

Department of Electrical
Engineering
University of Delaware
Newark, DE 19716

ABSTRACT

We propose a distributed disabling algorithm for a multiprocessing system in which each processor or unit is prevented from doing computation when it fails some number of tests by other units. The goal is to disable all faulty units and to enable all fault-free units. Specifically, a unit is disabled iff it fails d or more tests by enabled units (d -disabling rule). A multiprocessing system is c -correctable using the d -disabling rule iff all faulty units are permanently disabled and all fault-free units are permanently enabled after a finite number of applications of the disabling rule, provided there are no more than c faulty units. This models an unattended system where the removal of faulty units is done locally by simple and reliable circuitry. We give a sufficient condition for c -correctability in general systems and a necessary and sufficient condition in general systems where $c < d$. Then, we give necessary and sufficient conditions for c -correctability of two types of systems, (1) complete digraphs and (2) a new class of systems called segmented systems.

I. INTRODUCTION

In the systems diagnosis approach to reliability, testing is distributed. For example, in a multiprocessing system, processors test other processors producing pass or fail test results. The goal is to identify faulty units in the presence of incorrect information from such units. If there are too many faulty units, it may be impossible to uniquely identify them. For example, if *all* units are faulty, they may all produce pass test results, and it is impossible to distinguish between this and the case where all units are fault-free.

While testing is distributed, diagnosis may not be. Most papers on this subject have assumed a central diagnoser. In this case, system reliability depends critically on the reliability of the diagnoser. There has been a trend in recent years towards

*Research supported by NSF Grant #ECS-8203276 and a NAVELEX Chair Professorship tenured at the Naval Postgraduate School.

Report Documentation Page

Form Approved
OMB No. 0704-0188

Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.

1. REPORT DATE SEP 1988		2. REPORT TYPE		3. DATES COVERED	
4. TITLE AND SUBTITLE Properties of c-Correctability in Self-Diagnosing Systems				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S)				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Naval Postgraduate School, Department of Electrical and Computer Engineering, Monterey, CA, 93943				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited.					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT We propose a distributed disabling algorithm for a multiprocessing system in which each processor or unit is prevented from doing computation when it fails some number of tests by other units. The goal is to disable all faulty units and to enable all fault-free units. Specifically, a unit is disabled iff it fails d or more tests by enabled units (d-disabling rule). A multiprocessor system is c-correctable using the ddisabling ntle iff all faulty units are permanently disabled and all fault-free units are permanently enabled after a finite number of applications of the disabling rule, provided there are no more than c faulty units. This models an unattended system where the removal of faulty units is done locally by simple and reliable circuitry. We give a sufficient condition for c-correctability in general systems and a necessary and sufficient condition in general systems where c < d. Then, we give necessary and sufficient conditions for c-correctability of two types of systems, (1) complete digraphs and (2) a new class of systems called segmented systems.					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES 9	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

systems where the diagnosis is also distributed [2-8]. Meyer and Masson [7] propose a distributed diagnosis algorithm in which each unit has a "view" of the entire system based on tests it makes and on test results received by units that it finds to be fault-free. It is shown that, if there is an upper limit on the number of faulty units, the most common "view" is the correct one. This model is extended by Kuhl and Reddy [5,6] and Hosseini, Kuhl, and Reddy [3] to the case where links between units can also fail. Both are based on the Preparata, Metze, and Chien [9] model of systems diagnosis. However, there is then the problem of how the user identifies faulty units and removes them from the system. Kreuzer and Hakimi [4] address the first problem but not the second. A distributed diagnosis algorithm based on the Russell and Kime [10] model is shown by Holt and Smith [2]. Repair and graceful degradation models are proposed, using a message passing method in which fault-free units try to gain an accurate view of the status of various other units.

The problem of reliably disabling faulty units in systems diagnosis has received little attention. To the credit of Holt and Smith [2], "controllers" are proposed that disable units diagnosed as faulty. Unlike previous papers, we consider self-diagnosis in which the process of disabling faulty units is inherent. That is, the process of disabling a unit is built-in to the diagnosis algorithm. The reliable operation of the system depends on the reliability of a circuit which implements the rule. We choose to make the function of this circuit so simple that ultrareliability is achieved inexpensively (by redundancy, for example). Specifically, a unit is disabled iff it fails d tests by enabled units. This is the d -disabling rule. We assume an upper bound c on the number of faulty units, and we seek conditions which guarantee that all faulty units are disabled and all fault-free units are enabled after a finite number of applications of the d -disabling rule. A sufficient condition for c -correctability is given for general systems. The condition is expressed as a property of subsets of units and how they are interconnected by tests. A necessary and sufficient condition for c -correctability using the d -disabling rule is given for general systems in which $d < c$ holds. Next, we show necessary and sufficient conditions for two specific classes of systems

1. complete digraphs and
2. segmented systems.

The latter systems are new. They have a cyclical symmetry that extends over groups of units.

This paper is arranged as follows. Section III shows a sufficient condition for c -correctability in general systems. Section IV gives necessary and sufficient conditions for c -correctability in two specific systems.

II. BACKGROUND AND NOTATION

A *system* is a directed graph where nodes represent units or processors and arcs represent tests between units. Let $V = \{u_0, u_1, \dots, u_{n-1}\}$ be the set of units in the

system. Then, a directed arc exists from u_i to u_j iff u_i tests u_j . The test outcome is either pass or fail, depending on the status of the units involved in the test. Each unit is either *fault-free* or *faulty*. If the testing unit is fault-free, then the test outcome is a true representation of the status of the tested unit, pass if the tested unit is fault-free and fail if it is faulty. However, if the testing unit is faulty, the test outcome is arbitrarily pass or fail.

A complete set of test results is called a *syndrome*. The object of a diagnosis is to identify uniquely all faulty units given a syndrome. If the number of faulty units is small enough, then unique identification is possible for *all* possible arrangements of faulty units and *all* possible syndromes. Specifically, a system is *t-diagnosable* iff all faulty units can be uniquely identified provided there are no more than t of them. Preparata, Metze, and Chien [9] show necessary conditions for a system to be *t*-diagnosable and Hakimi and Amin [1] show necessary and sufficient conditions.

Each unit is either *enabled* or *disabled*. We assume initially that any unit can be arbitrarily enabled or disabled.

Definition: The *d-disabling rule* is as follows: a unit is disabled if it fails d or more tests by enabled units; otherwise, it is enabled.

The rule is applied continually to each unit without regard to order among units. We seek conditions that guarantee a faulty unit is eventually disabled and remains disabled at each application of the *d*-disabling rule and that a fault-free unit is similarly enabled.

Definition: A system is *c-correctable* using the *d*-disabling rule iff for

1. any arrangement of c or fewer faulty units,
2. any resulting set of test outcomes, and
3. any initial assignment of enable/disable to units,

the continual application of the *d*-disabling rule to each unit u permanently disables u if u is faulty and permanently enables u if u is fault-free.

Fig. 1 shows a system with six units, two of which are faulty. Assume that both produce fail test outcomes of all tests they apply and that both are initially enabled. The fault-free units produce a fail test outcome if the unit tested is faulty and pass if it is fault-free. Consider the application of the 1-disabling rule to this system. If the rule is applied first to the fault-free units, they will be disabled regardless of their initial status. The subsequent application of the 1-disabling rule to the faulty units will leave them enabled. Successive applications of the 1-disabling rule will produce no change, leaving faulty units permanently enabled and fault-free units permanently disabled. Thus, the system is *not* 2-correctable using the 1-disabling rule. However, it is 1-correctable, because the first application of the 1-disabling rule to a fault-free unit u testing the single faulty unit will enable u (it fails no tests). Then, a subsequent

application of the 1-disabling rule to the faulty unit disables it. Once the faulty unit is disabled, no fault-free unit is disabled. Thus, in the steady-state, all fault-free units are enabled, while the faulty unit is disabled.




Figure 1. Application of the 1-Disabling Rule.

There is no value of d for which the system is 2-correctable using the d -disabling rule (as can be demonstrated by an exhaustive enumeration of all possibilities). However, it is 2-diagnosable [9], and so all faulty units can be uniquely identified by a central diagnoser provided there are 2 or fewer of them. Thus, distributed diagnosis places a greater restriction on the number of faulty processors which can be tolerated. It is the penalty incurred for using only local information to identify the faulty/fault-free status of units.

III. GENERAL c -CORRECTABLE SYSTEMS

We begin by showing properties possessed by every c -correctable system using the d -disabling rule.

Definition: $\Gamma(u) = \{u_i \mid u \neq u_i \in V \text{ and } u_i \text{ tests } u\}$.

$\Gamma(u)$ is the set of units that test u .

Lemma 1: Every unit in a c -correctable system using the d -disabling rule is tested by at least $d + c - 1$ units.

Proof: On the contrary, suppose there exists a unit u in a c -correctable system that is tested by $d + c - 2$ or fewer other units. Consider a subset $C \subseteq \Gamma(u)$ such that $|C| = c - 1$. Let $F = C \cup \{u\}$ be the faulty units in the system. Assume all test results of u by units in C are pass. Then, the largest number of fault-free units testing u is $d - 1$ and u , having failed less than d tests, is permanently enabled.

Thus, the system is not c -correctable.

Q.E.D.

The condition of Lemma 1 becomes necessary and sufficient when c is strictly less than d , as shown in the next lemma.

Lemma 2: If $c < d$, then a system is c -correctable using the d -disabling rule iff every unit is tested by at least $d + c - 1$ other units.

Proof: (if) Since $c < d$, all fault-free units are permanently enabled. Since each unit is tested by at least $d + c - 1$ units, and each faulty unit is tested by no more than $c - 1$ other faulty units, there are at least d fault-free units testing each faulty unit. Since all fault-free units are enabled, each faulty unit fails d tests by enabled units, and so, by the d -disabling rule, is disabled.

(only if) On the contrary, assume there is a system that is c -correctable, but does not satisfy the condition. However, this is impossible since, by Lemma 1, all units in a c -correctable system using the d -disabling rule are tested by at least $d + c - 1$ units.

Q.E.D.

A limit on the number of units is given by:

Lemma 3: In a c -correctable system, $n \geq 2c + 1$, where n is the total number of units.

Proof: Since every faulty unit in a c -correctable system must be unambiguously identified as faulty, a c -correctable system is also c -diagnosable. From [9], a c -diagnosable system has the property $n \geq 2c + 1$.

Q.E.D.

We now show a sufficient condition for c -correctability using the d -disabling rule.

Definition: $\Gamma_{\geq d}^1(Z) = \{u_i \mid u_i \in Z \subseteq V \text{ and there are at least } d \text{ units in } Z \text{ that test } u_i\}$.

$\Gamma_{\geq d}^1(Z)$ is the set of units outside of Z tested by at least d units in Z .

Theorem 1: S is c -correctable using the d -disabling rule if for all $F \subseteq V$ with $|F| \leq c$, all subsets F' of F have the property, $F' \cap \Gamma_{\geq d}^1(Z) \neq \emptyset$, where $Z = V - F - \Gamma_{\geq d}^1(F')$.

Proof: Suppose the condition holds, but S is not c -correctable. Then, either (i) there is a set of faulty units $F \subseteq V$, where $|F| \leq c$, such that there is a nonempty subset $F' \subseteq F$ consisting entirely of permanently enabled units, (ii) there is a nonempty subset $G \subseteq V - F$ of fault-free units all of which are permanently disabled, or (iii)

there is a set of units $V_{NPDE} \subseteq V$ which are neither permanently disabled nor permanently enabled for some arbitrarily long sequence of applications of the d -disabling rule. In the case of (i), it must be that no unit in F' is tested by d or more permanently enabled fault-free units; that is, units in $Z = V - F - \Gamma_{\geq d}^{-1}(F')$. It follows that $F' \cap \Gamma_{\geq d}^{-1}(Z) = \phi$, a contradiction. In the case of (ii), it must be that $G \subseteq \Gamma_{\geq d}^{-1}(F')$, where F' is a set of enabled faulty units. If $Z = V - F - \Gamma_{\geq d}^{-1}(F')$, then $F' \cap \Gamma_{\geq d}^{-1}(Z) = \phi$, since units in F' are not permanently disabled, a contradiction. Consider (iii). Let V_E be the set of all permanently enabled units, and let V_D be the set of permanently disabled units. We can assume that $V_E = V - F - V_{NPDE}$ and $V_D = F - V_{NPDE}$, since otherwise there exists a permanently enabled faulty unit or a permanently disabled fault-free unit and this would fall under case (i) or (ii). Let $F' = V_{NPDE} \cap F$. It follows that $\Gamma_{\geq d}^{-1}(F') \supseteq (V - F) \cap V_{NPDE}$, since units in $(V - F) \cap V_{NPDE}$, which are all fault-free, can be disabled only by faulty units that are enabled at some time. Specifically, units in $V_D = F - F'$ cannot disable units in V_{NPDE} since they are permanently disabled. Thus, $Z = V - F - \Gamma_{\geq d}^{-1}(F') \subseteq V - F - V_{NPDE}$ consists of permanently enabled fault-free units exclusively. Since units in F' are not permanently disabled, $F' \cap \Gamma_{\geq d}^{-1}(Z) = \phi$, a contradiction.

Q.E.D.

IV. SPECIFIC c -CORRECTABLE SYSTEMS

A. COMPLETE DIGRAPHS

A *complete digraph* $G(V,E)$ is a digraph with node set V and edge set E such that for every ordered pair (u,v) where $u, v \in V, (u,v) \in E$. We have,

Lemma 4: A complete digraph on n units is c -correctable using the d -disabling rule iff

$$c < d \leq n - c. \quad (1)$$

Proof: (only if) Let S be a complete digraph that is c -correctable using the d -disabling rule, and, assume, on the contrary, that either $c \geq d$ or $d > n - c$. Suppose $c \geq d$. Let there be c faulty units that are initially enabled, and assume that each fails all fault-free units it tests and passes all faulty units. An application of the d -disabling rule to all fault-free units will cause them to be disabled. Since there are no enabled units which fail the faulty units, an application of the d -disabling rule to faulty units leaves them enabled. This situation is permanent. Thus, the system is not c -correctable. Now suppose $d > n - c$. Consider a faulty unit u . If there are c faulty units, u is tested by $n - c$ fault-free units. If all faulty units pass all faulty units, then u is permanently enabled because there are insufficiently many fail test outcomes. Thus, S is not c -correctable using the d -disabling rule.

(if) Let $c < d \leq n - c$, and assume the system is not c -correctable using the d -disabling rule. Either there is (i) a fault-free unit that is permanently disabled, (ii) a faulty unit that is permanently enabled, or (iii) a unit that is neither permanently disabled nor permanently enabled for some arbitrarily long sequence of applications of the d -disabling rule. For a fault-free unit to be permanently disabled as in (i), it must be tested by d or more enabled faulty units. But from $c < d$, this is impossible. Thus, all fault-free units are permanently enabled. For a faulty unit to be permanently enabled as in (ii), it must be tested by no more than $d - 1$ enabled fault-free units. However, this is impossible since there are at least $n - c$ permanently enabled fault-free units, and from $d \leq n - c$, it follows that each faulty unit is tested by at least d enabled fault-free units. Let V_{ED} be a nonempty set of units which are neither permanently enabled nor permanently disabled. Let V_E be the set of all permanently enabled units and V_D the set of all permanently disabled units. We can assume that $V_E \subseteq V - F$ and $V_D \subseteq F$, since otherwise this would fall under case (i) or (ii). It follows that $|V_E| \leq d - 1$; otherwise all faulty units are disabled, and thus all fault-free units are enabled, which implies $V_{NPDE} = \phi$. We now show that V_{NPDE} contains no fault-free unit. On the contrary, such a unit must fail at least d tests by faulty units, which implies $d \leq c$, contradicting the condition $c < d$. Thus, V_{NPDE} contains only faulty units, and $|V_{NPDE} \cup V_D| \leq c$. However, $n = |V_{NPDE} \cup V_D| + |V_E| \leq c + d - 1$, contradicting the condition $d \leq n - c$.

Q.E.D.

B. SEGMENTED SYSTEMS

Definition: $G_{s,m}(V,E)$ is a *segmented system* if

1. $V = A_0 \cup A_1 \cup \dots \cup A_{s-1}$,
2. $|A_i| = m, 0 \leq i \leq s-1$,
3. $A_i \cap A_j = \phi, i \neq j, 0 \leq i, j \leq s-1$, and
4. $E = \{(u,v) \mid u \in A_i \text{ and } v \in A_{i+1}, \text{ where index addition is modulo } s\}$.

A segmented system consists of $s - 1$ groups of m units each. The only tests that exist are between adjacent groups, in which case, all possible tests exist.

Lemma 5: A segmented system $G_{s,m}(V,E)$ is c -correctable using the d -disabling rule iff

$$\frac{c(2-s \bmod 2)}{s} < d \leq m - c + 1. \quad (2)$$

Proof: (if) On the contrary, suppose there is a segmented system where the condition holds yet is not c -correctable using the d -disabling rule. Then, either there is (i) at

least one permanently enabled faulty unit, (ii) at least one permanently disabled fault-free unit, or (iii) at least one unit that is neither permanently enabled nor permanently disabled. Assume (i) holds. Let $u \in A_i$ be a permanently enabled faulty unit. Thus, there exists an integer a such that after a applications of the d -disabling rule u is always enabled. For this steady-state condition, we observe the following. *There can be at most $d - 1$ enabled fault-free units in A_{i-1} . Either there are no disabled fault-free units in A_{i-1} or there is at least one. **If there are none, A_{i-1} has at least $m - d + 1$ faulty units, and among A_{i-1} and A_i there are at least $m - d + 2$ faulty units. Thus, $m - d + 2 \leq c$. However, this contradicts the rightmost inequality of (2). If A_{i-1} has at least one disabled fault-free unit, there are at least d enabled faulty units in A_{i-2} . Since A_{i-2} has at least one enabled faulty unit, a similar argument yields a contradiction to the rightmost inequality of (2) or the conclusion that A_{i-4} has d faulty units, etc.. If s is even, every other A_j can have at least d faulty units, for a total of at least $ds/2$ faulty units. If s is odd, every A_j can have at least d faulty units, for a total of at least ds faulty units. Thus, $\frac{ds}{2-s \bmod 2} > c$, contradicting the leftmost inequality on (2). The proof for the case of (ii) is included in the above (beginning at **). Now consider case (iii). Let V_E be the set of permanently enabled units and V_D the set of permanently disabled units. We can assume $V_E = V - F - V_{NPDE}$ and $V_D = F - V_{NPDE}$, where V_{NPDE} is the nonempty set of units which are neither permanently disabled nor disabled; otherwise we have case (i) or (ii). Further, $V_{NPDE} \cap F \neq \phi$; otherwise it follows that there are no fault-free units in V_{NPDE} , since such units can be disabled only by faulty units in V_{NPDE} , which implies that $V_{NPDE} = \phi$. Let $u \in V_{NPDE} \cap F$, where $u \in A_i$. The proof that this leads to a contradiction is included in the above (beginning at *).

(only if) Assume there is a c -correctable segmented system using the d -disabling rule in which the condition does not hold. Thus, either (i) $d > m - c + 1$ or (ii) $c(2-s \bmod 2)/s \geq d$. Suppose (i) holds. Let A_{i-1} contain at least $\min(m, c-1)$ faulty units, and let A_i contain at least one faulty unit. It follows that there are at most $m - c + 1$ fault-free units in A_{i-1} . From (i), there are fewer than d fault-free units in A_{i-1} , allowing the faulty unit in A_i to be permanently enabled. This contradicts the assumption that the system is c -correctable. Suppose (ii) holds. If s is even, there can be at least $ds/2$ faulty units in the system or at least d for every other A_j . Let all faulty units be initially enabled, and let each fail all tested fault-free units. Then, all units in the A_j 's consisting of fault-free units exclusively are disabled, and there are no enabled fault-free units to disable the successor faulty units. Thus, all faulty units are permanently enabled, contradicting the assumption that the system is c -correctable. If s is odd, there are at least ds faulty units in the system or at least d faulty units for every A_j . In a similar manner, it follows that all fault-free units can be permanently disabled, contradicting the assumption that the system is c -correctable.

Q.E.D.

V. CONCLUDING REMARKS

We propose a new process of self-diagnosis where the disabling of faulty units is an integral part of the diagnosis. We show conditions under which correct diagnosis is achieved; i.e. fault-free units are enabled and faulty units are disabled. The disabling mechanism, which must be done ultrareliably, is simple, so that it is constructed at reasonable cost. The approach is practical and narrows the gap between the *theory* of systems diagnosis and the practical application of that theory.

REFERENCES

- [1] S. L. Hakimi and A. T. Amin, "Characterization of the connection assignment problem of diagnosable systems," *IEEE Trans. on Comput.*, vol. C-23, Jan. 1974, pp. 86-88.
- [2] C. S. Holt and J. E. Smith, "Self-diagnosis in distributed systems," *IEEE Trans. on Comp.*, vol. C-34, pp. 19-32, Jan. 1985, pp. 19-32.
- [3] S. H. Hosseini, J. G. Kuhl, and S. Reddy, "On self-fault diagnosis of the distributed systems," *IEEE Trans. on Comp.*, vol. C-7, Feb. 1988, pp. 248-251.
- [4] S. E. Kreutzer and S. L. Hakimi, "Distributed diagnosis and the system user," *IEEE Trans. on Comp.*, vol. C-37, Jan. 1988, pp. 71-79.
- [5] J. Kuhl and S. Reddy, "Distributed fault tolerance for large multiprocessor systems," *Proc. of the 7th Inter. Symp. on Comp. Architecture*, June 1980, pp. 23-30.
- [6] J. Kuhl and S. Reddy, "Fault diagnosis in fully distributed systems," in *Proc. of the 11th Inter. Conf. on Fault Tolerant Computing*, June 1981, 100-105.
- [7] G. Meyer and G. M. Masson, "An efficient fault diagnosis algorithm for symmetric multiple processor architectures," *IEEE Trans. on Comput.*, vol. C-27, Nov. 1978, pp. 1059-1063.
- [8] R. Nair, "Diagnosis, self-diagnosis, and roving diagnosis in distributed digital systems," Coord. Sci. Lab., Univ. of Illinois, Urbana, IL Report R-823, Sept. 1978.
- [9] F. Preparata, G. Metze, and R. Chien, "On the connection assignment problem of diagnosable system," *IEEE Trans. Electron. Comput.*, vol. EC-16, Dec. 1967, pp. 848-854.
- [10] J. Russell and C. Kime, "System fault diagnosis: closure and diagnosability with repair," *IEEE Trans. on Comput.*, vol C-24, Dec. 1975, pp. 1078-1088.