

V. CONCLUDING REMARKS

Some directions of extension of the present results would be the following:

- 1) give the number of cascade-realizable multivalued functions with permuted input variable assignments;
- 2) give the number of symmetry types of multivalued functions realizable by cascades; and
- 3) extend the results to the case of disjunctive networks.

ACKNOWLEDGMENT

We would like to express our gratitude to Y. Kanada for his assistance at using the REDUCE/HLISP system, to H. Machida and the referees for their helpful comments on the manuscript, and to J. T. Butler for providing us a copy of [6].

REFERENCES

- [1] K. K. Maitra, "Cascaded switching networks of two-input flexible cells," *IRE Trans. Electron. Comput.*, vol. EC-11, pp. 136-143, Apr. 1962.
- [2] J. Sklansky, A. J. Korenjak, and H. S. Stone, "Canonical tributary networks," *IEEE Trans. Electron. Comput.*, vol. EC-14, pp. 961-963, Dec. 1965.
- [3] A. Maruoka and N. Honda, "Logical networks of flexible cells," *IEEE Trans. Comput.*, vol. C-22, pp. 347-358, Apr. 1973.
- [4] J. T. Butler and K. J. Breeding, "Some characteristics of universal cell nets," *IEEE Trans. Comput.*, vol. C-22, pp. 897-903, Oct. 1973.
- [5] J. T. Butler, "On the number of functions realized by cascades and disjunctive networks," *IEEE Trans. Comput.*, vol. C-24, pp. 681-690, July 1975.
- [6] —, "Fanout-free networks of multivalued gates," in *Proc. 7th Int. Symp. Multiple-Valued Logic*, Charlotte, NC, May 1977.
- [7] A. C. Hearn, *Reduce 2 User's Manual*, 2nd ed., Univ. Utah, Mar. 1973.
- [8] Y. Kanada, "Implementation of HLISP and algebraic manipulation language REDUCE-2," Inform. Sci. Lab., Univ. Tokyo, Tokyo, Japan, Jan. 1975.

Universal System Diagnosis Algorithms

JAMES E. SMITH

Abstract—A class of simple digital system diagnosis algorithms is presented, and two members of the class are examined in detail. The algorithms are based on the assumption that good units can be replaced during the diagnosis process. Information pertaining to the system testing structure is not used by the two principal algorithms, so they can be applied regardless of system structure. The efficiency of the algorithms in terms of good units replaced is analyzed, and they are shown to compare favorably with methods for special case systems that have been proposed by others.

Index Terms—Diagnosis algorithms, fault diagnosis, system diagnosis.

I. INTRODUCTION

One of the most important tools for locating failures in digital systems is the application of diagnostic test sets. Tests are usually derived for each of the subsystems (units), and the application of a test on some unit typically requires the use of some other unit(s) which may also be faulty. If a faulty unit is used in the application

of a test, the test result may be invalid. This test invalidation greatly increases the difficulty of system diagnosis and makes it a complex and interesting problem.

In [1] Preparata *et al.* proposed one of the first, and probably the best-known, models for system diagnosis. In this model, each unit has the capability of testing other units by itself, i.e., only one unit is required for the application of a test. The assumption is made that if a faulty unit performs a test, a fault-free unit could be judged faulty or a faulty unit could be judged fault-free. This type of invalidation of test results will be referred to as *symmetric invalidation*.

In [2], Barsi *et al.* proposed a different type of invalidation as being more likely when complex units are performing tests. This *asymmetric invalidation* assumes that all invalidation takes the form of a correct unit being judged faulty; i.e., a faulty unit always fails all of its tests.

While the models [1], [2] may work well if complex subsystems are used, in many digital systems a number of relatively simple units, e.g., adders and multiplexers, are employed to perform a test. To handle this more general case, a model is proposed in [3], [4] where multiple units may be used to perform a test, and the failure of any one of them causes the test to be invalid.

In the past, theoretical work has been concerned with the *diagnosability* of systems. A system is *k-step t-fault diagnosable* if there exists a sequence of *k* applications of the test set and repairs of identified faulty units that allows all the faulty units originally present to be identified provided the number of original faulty units does not exceed *t*. An important special case occurs if only one application of the test set is required; that is, the system is *one-step t-fault diagnosable*. For *k-step* diagnosis, it is typically assumed that any repaired unit is fault-free and remains so until all units have been repaired.

One of the most commonly studied problems in system diagnosis is the determination of necessary and sufficient conditions under which a system is *t-fault diagnosable* [1]–[8]. Another problem often studied is the construction of diagnosable systems that are in some sense optimal [1], [2], [8], [9]. A third problem is the determination of diagnosis algorithms [2], [9]–[14].

In practice, "repair" can take the form of replacement or actual repair. For our purposes, it is convenient to speak of replacement, although this does not restrict the generality of any of the results. In most theoretical work it is assumed that no good units are replaced, but in [9], [10] this requirement is dropped. Such an assumption is less conservative and is somewhat closer to diagnosis techniques that are used in practice. Based on this assumption, Friedman [9] proposes a different measure of system diagnosability, *t-out-of-s (t/s) diagnosability*. A system is *t/s diagnosable* if a set of $f \leq t$ faulty units can be located and repaired by replacing at most *s* units. One-step and *k-step t/s* diagnosability are possible, and they are defined in the natural way.

In this paper, we propose a class of very simple diagnosis algorithms that allow the replacement of good units as in [9], [10]. The algorithms only rely on test results and are *independent of system structure*. Consequently, the algorithms can be applied to any system; it is in this sense that the algorithms are *universal*. The algorithms are also guaranteed to result in a correct system, provided that a fundamental condition on fault detectability is satisfied. Consequently, they are effective against the broadest class of diagnosable faults.

Two of the algorithms are studied in terms of the models [1], [2], [4]. This analysis leads to upper bounds on the number of

Manuscript received February 15, 1978; revised January 1, 1979.

The author is with the Department of Electrical and Computer Engineering, University of Wisconsin-Madison, Madison, WI 53706.

fault-free units that must be replaced, and the bounds can be evaluated for any system structure, whether regular or not. The algorithms are shown to compare favorably with other more complex algorithms developed for special-case systems, e.g., "single-loop" systems. Examples are given that show the algorithms can often perform much better than the bounds indicate.

II. NOTATION AND TERMINOLOGY

Models to be considered here are the ones presented in [1], [2], [4] as well as the model in [4] with asymmetric invalidation. It is assumed that the systems are *morphic* [4]. This will allow us to use graphical models. Many of the results to be presented here can be generalized so that semimorphic systems [4] and systems requiring more complex models [8], [15], [16] are included. Nevertheless, it is felt that the models to be considered here offer sufficient generality while allowing the insight that graphical models provide.

Since the models [1], [2] are special cases of the one given in [4] (with the appropriate type of invalidation), we present only the model for morphic systems given in [4]. This model can be formulated either in terms of "units" or "faults." In [4] faults are used; we use units.

A system S is composed of a set of n potentially faulty units $U = \{u_0, u_1, \dots, u_{n-1}\}$. There is a set of m tests, $\{t_0, t_1, \dots, t_{m-1}\}$, that can be applied to the units in S . A test t_j is a *complete test* for unit u_i if t_j always fails when u_i is the only faulty unit in S and t_j always passes when S contains no faulty units. Each test is assumed to be a complete test for one and only one unit, and at least one complete test exists for each unit. The set of tests that is complete for unit u_i is denoted as $t(u_i)$. For a set of units, $\{u_i, u_j, \dots, u_k\}$, $t(\{u_i, u_j, \dots, u_k\}) = t(u_i) \cup t(u_j) \cup \dots \cup t(u_k)$.

Each test is performed by some set of units. Let $T(u_i)$ be the set of tests applied at least in part by unit u_i . If *symmetric invalidation* is assumed and u_i is faulty, the result of a test $t_j \in T(u_i)$ is unreliable in the sense that t_j may pass even though u_k is faulty and $t_j \in t(u_k)$, or it might fail if u_k is fault-free. If *asymmetric invalidation* is assumed, the result of $t_j \in T(u_i)$ is unreliable only in the sense that t_j might fail even though $t_j \in t(u_k)$ and u_k is fault-free. $T(\{u_1, u_j, \dots, u_k\})$ is defined to be $T(u_i) \cup T(u_j) \cup \dots \cup T(u_k)$; it is this property that makes the system morphic.

The diagnosis model just described can be used to represent a system as a directed graph. There is an internally labeled vertex in the graph for each unit. An edge is directed from the node labeled u_i to the node labeled u_j if $T(u_i) \cap t(u_j) \neq \emptyset$. The edges are labeled with the tests in the set $T(u_i) \cap t(u_j)$. In some cases it is convenient to assume that a test result is always valid, e.g., if the test is applied by some external unit that is assumed fault-free. If this occurs, the vertex internally labeled with the tested unit is externally labeled with the test.

The graphical models of [1], [2] result if each test is performed by only one unit. When this is the case, each edge has a unique label, and the label can be deleted without the loss of any invalidation information.

In system diagnosis, a *syndrome* is typically defined to be a binary vector representing the test outcomes. We define the *aggregate syndrome* to be an ordered set $\sigma = \langle \sigma_0, \sigma_1, \dots, \sigma_{n-1} \rangle$, where n is the number of units in S and σ_i is the total number of tests on u_i that fail for a given application of the test set.

We now present the diagnosis graphs of several example systems; the graphs help to explain the notation and are used in later discussion.

Example 1: This is a single-loop system; tests are performed






Fig. 1. Diagnostic graphs for three example systems.

by only one unit and each unit has only one test. Fig. 1(a) shows the diagnostic graph for a single-loop system.

Example 2: One can use the same interconnection structure suggested by single-loop systems but with a more complex testing procedure. Assume unit i applies inputs to unit $i + 1$, but unit $i + 2$ observes the outputs from unit $i + 1$ (addition is modulo n , where there are n units). Fig. 1(b) shows a diagnostic graph for a system with five units. Such a diagnosis scheme cannot be represented by the models [1], [2].

Example 3: A possible application of the models [1], [2] is to microcomputer arrays. Fig. 1(c) shows the model for a part of such a two-dimensional array. Test labels have been deleted since it is assumed that each test is performed by only one microcomputer.

In Fig. 1(b), $t(u_4) = t_4$; $T(u_4) = \{t_3, t_0\}$. In Fig. 1(c), if unit u_4 is faulty, a possible aggregate syndrome is $\sigma = \langle 0, 1, 0, 0, 4, 1, 0, 1, 0, \dots \rangle$ if symmetric invalidation is assumed.

III. A CLASS OF DIAGNOSIS ALGORITHMS

When diagnosing a system within the framework of the diagnosis model of Section II, the test outcome (pass/fail), the $t(u_i)$, and the $T(u_i)$ may all be considered. The test outcomes and the units they test [indicated by the $t(u_i)$] are clearly needed for any

nontrivial replacement strategy to work (a trivial strategy simply replaces all the units immediately). On the other hand, the $T(u_i)$ reflect the testing structure of the system, and we will discuss diagnosis methods that do not use this information. Consequently, the methods are independent of the structure of the system and can be applied to any system.

The fact that the $T(u_i)$ are unnecessary is also advantageous because this information may at times be difficult to determine. This is particularly true in systems where several fault-free subsystems are required to perform a test, and where there is some nonobvious dependency among the subsystems.

We recall that for unit u_i , σ_i is the member of the aggregate syndrome that indicates the number of tests on u_i that fail. Then the universal diagnosis algorithms belong to the following scheme.

STEP: Perform the tests t_0, t_1, \dots, t_{m-1} and compute σ ;
 Let $F = \{u_j | \sigma_j \neq 0 \text{ and } u_j \text{ has not been replaced}\}$;
 If $F = \emptyset$, then the system is assumed correct; exit;
 Choose F' , a nonempty subset of F ;
 Replace all units in F' ; go to STEP.

A particular algorithm is defined by a deterministic method for choosing F' . We consider two particular algorithms where the choice of F' is independent of the $T(u_i)$.

Algorithm 1: Let $F' = F$.

Algorithm 2: Let $F' = \{u_i | u_i \in F \text{ and for all } u_j \in F \sigma_i \geq \sigma_j\}$.

These two universal algorithms are intuitively simple and are probably closer to methods used in practice than the more sophisticated algorithms that have been proposed. Algorithm 1 simply replaces all units that have not already been replaced and which have failed some test. Algorithm 2 replaces only those units that fail the greatest number of tests considering only those units that have not yet been replaced. It should be noted that we make the usual assumption that all replacement units are fault-free and remain fault-free at least until the algorithm terminates.

In order to characterize the sets of faulty units that the universal algorithms can correctly diagnose and replace, we first observe that in a system it is sometimes possible for a set of faulty units to be present such that invalidations occur which cause all tests to pass. Clearly, such a fault set can potentially cause any nontrivial diagnosis algorithm to fail. This is because if all the tests pass, there is nothing the algorithm can do but assume the system is fault-free. For this reason, we define a set of faulty units U^F to be *completely detectable* if some test will definitely fail if U^F or any nonempty subset of U^F is faulty. We observe that when asymmetric invalidation is assumed, any set of faulty units is completely detectable. For symmetric invalidations, fault detectability is discussed in [4]. Then:

Theorem 1: Assuming asymmetric invalidation, any algorithm belonging to the class we have defined will always terminate with a fault-free system regardless of what units are initially faulty.

Proof: At each step of the algorithm at least one unit is replaced. Since there is a finite number of units and each unit is only replaced once, the algorithm must eventually terminate. At the time it terminates, if a test fails, it must be performed on a replaced unit by a faulty nonreplaced unit. Any test on the faulty unit must pass, but this contradicts asymmetric invalidation. Hence, at termination, all tests must pass. Asymmetric invalidation implies all units must be fault-free. \square

Theorem 2: Assuming symmetric invalidation, if a completely detectable set of faulty units is initially present, the algorithms of the above class will always terminate with a fault-free system.

Proof: This is similar to Theorem 1; when the algorithm terminates and a test fails, it must be performed on a replaced (fault-free) unit by a faulty nonreplaced unit. Any such test results are invalid and the tests could have potentially passed. Consequently, any remaining set of faulty units is not completely detectable, so the original faulty units could not be completely detectable. The theorem follows immediately. \square

Theorems 1 and 2 are fundamental theorems. They essentially say that the algorithms are guaranteed to result in a correct system if any nontrivial algorithm guarantees a correct system. In addition, Theorem 2 makes the worst case assumption that a test result on a faulty unit is always invalid (i.e., it passes) whenever the test is performed by a faulty unit. In many cases when the worst case does not occur, correct repair results even though a set of faulty units is not completely detectable.

Another observation is that the only upper limit on the number of faulty units that can be diagnosed is the one that may be implicitly imposed by the detectability condition.

One can also make some observations regarding the computational complexity of the algorithms. With each step (an application of the tests and the choice of units for replacement), only $O(m)$ operations are performed where m is the number of tests. It seems that any nontrivial k -step diagnosis must be at least this complex for each step if the result of each test is even considered. The number of steps required by the algorithms is more difficult to analyze, but it is clearly bounded by the total number of units replaced since at least one unit is replaced at each step. Taking this into account, it follows from later results that for large classes of systems $O(f)$ steps are sufficient, where f units are actually faulty.

We observe that if Algorithm 1 is applied when there is asymmetric invalidation, then only one step is required, since all faulty elements must fail all their tests. Hence, when applied to such a system, Algorithm 1 is of complexity $O(m)$.

Testing required by the algorithms can be simplified on steps after the first because σ_i for a unit u_i that has already been replaced is not needed. If U^R is the set of units that have already been replaced, the tests in $t(U^R)$ do not need to be performed.

Finally, one could modify Algorithm 2 to select just one u_j with the maximum σ_j using some rule. This would tend to decrease the average number of good units replaced while increasing the number of steps required. In our later analysis we are only concerned with the worst case number of good units replaced; consequently, we will not complicate this analysis with some arbitrary selection rule.

IV. EFFICIENCY OF THE ALGORITHMS

Thus far we have only established that the proposed diagnosis algorithms always work given a detectability condition. However, if they require the replacement of large numbers of fault-free units, they are of little practical value. In this section, we analyze Algorithms 1 and 2 to determine bounds on the number of good units that may be replaced. The algorithms are shown to be efficient in terms of good units replaced for broad classes of systems, and the algorithms compare favorably with more complex special-case algorithms proposed by others.

For discussing the efficiency of the algorithms, some of the concepts introduced by Friedman [9], [10] are useful. From [9], a system S is k -step t/s *diagnosable* if by k applications of the diagnostic test set sequence any set of $f \leq t$ faulty units in S can be diagnosed and repaired by replacing at most s units. The parameter t reflects the maximum number of faulty units that are assumed to be possible. For the algorithms presented here, there is no such explicitly assumed upper bound. Hence, we choose to

discuss f/s diagnosability. That is, if f faulty units are actually present, at most s units need to be replaced in order to get a correct system. Consequently, s will typically be a function of f , while s is a function of t in [9], [10].

Before proceeding, we define two fundamental parameters that are used in the determination of worst case f/s diagnosability for a system composed of the set of units U .

$T^M = \text{maximum}_{u_i \in U} |T(u_i)|$; the maximum number of tests performed at least in part by any unit.

$t^m = \text{minimum}_{u_i \in U} |t(u_i)|$; the minimum number of tests performed on any unit.

Theorem 3: For the system S with either symmetric or asymmetric invalidation, the use of Algorithm 1 results in at worst $f/(T^M + 1)f$ diagnosability.

Proof: In order for a fault-free unit to be replaced, it must fail some test, i.e., some test result on the unit must be invalid. The maximum number of invalid test results is fT^M . Consequently, fT^M is a bound on the number of fault-free units replaced. Then with the f faulty units, at most $(T^M + 1)f$ units are replaced. \square

The bound on diagnosability of Theorem 3 can be reached if each faulty unit invalidates a test on T^M different good units. In situations where $(T^M + 1)f > n$, where n is the total number of units in the system, then clearly one should take f/n as a bound on diagnosability. This same comment holds for all bounds on diagnosability given in this paper.

For single-loop systems, $|T(u_i)| = 1$ for any unit u_i . Theorem 1 implies that such systems are $f/2f$ diagnosable when Algorithm 1 is used. In [9] diagnosis algorithms are given that are specifically for single-loop systems. It is interesting to compare Algorithm 1 with these other methods.

In [9], six different strategies are given. Strategy 2 is a t/t diagnosis approach that assumes only faulty units are replaced. Of the five strategies that allow replacement of fault-free units, Strategy 1 is a one-step approach and Strategies 5 and 6 are two-step methods. As a consequence, Strategies 1, 5, and 6 are rather inefficient, and it is not fair to compare general k -step procedures with them. On the other hand, Strategies 3 and 4 are k -step methods and do offer a fair comparison. Strategy 3 is a $t/3t - 1$ diagnosis method while Strategy 4 gives $t/2t - 1$ diagnosis. Hence, our $f/2f$ procedure compares very favorably in general, especially when one considers the difference between f/s and t/s diagnosis.

As in [9], the average efficiency of the algorithms is usually much better than the bounds. To demonstrate this, we present an example that is taken from [9].

Example 4: In a single-loop system S with 32 units, the units $u_{11}, u_{12}, u_{13}, u_{21}, u_{23}, u_{24}$, and u_{25} are faulty. We denote the aggregate syndrome after the i th test application as σ^i . Then

$$\sigma^1 = \langle 0000000000010000000001110100000 \rangle.$$

Note that because of notational differences our σ^i and R_i in [9] differ by a cyclic shift of one position (recall that σ_0 reflects the test results for u_0).

According to Algorithm 1, units $u_{11}, u_{21}, u_{22}, u_{23}, u_{24}$, and u_{26} are replaced. Then, at step 2,

$$\sigma^2 = \langle 00000000000100000000001100000 \rangle.$$

Hence, units u_{12} and u_{25} are replaced. Unit u_{26} is not replaced as it was replaced at step 1. Then,

$$\sigma^3 = \langle 00000000000010000000000000000 \rangle.$$

Now, unit u_{13} is replaced, and σ^4 is all 0's implying a fault-free system.

In this example, seven units were faulty and nine were replaced. In [9] Strategies 3 and 4 resulted in eleven and nine units being replaced, respectively. \square

Of course, Algorithm 1 can be applied to other than single-loop systems. For example, its application to the systems of Fig. 1(b) results in at worst $f/3f$ diagnosability and for the arrays of Fig. 1(c) $f/5f$ diagnosability is the worst case.

We now turn to Algorithm 2 which is generally more efficient than Algorithm 1 in terms of replacement of fault-free units.

Theorem 4: For the system S with symmetric invalidation, the use of Algorithm 2 results in at worst $f/(T^M - t^m + 2)f$ diagnosability provided that $t^m \leq T^M$.

Proof: (By induction on f) $f = 0 \Rightarrow$ no units are replaced, and the theorem holds trivially.

Let U^F be the set of faulty units initially present, and assume the theorem holds for all $|U^F| < f$. At the end of some step $j, j \geq 1$, let U_j^R be the set of replaced units and U_j^F the set of replaced faulty units. We consider two cases.

Case 1: There is some $j \geq 1$ where $|U_j^R| \leq 2|U_j^F|$. At this point, the system still contains $|U^F - U_j^F|$ faulty units. We could conceptually begin Algorithm 2 at this point with the partially repaired system as input. Since $j \geq 1$, some units have been replaced; furthermore, because $|U_j^R| \leq 2|U_j^F|$, some faulty units have been replaced and $|U^F - U_j^F| < f$. By the induction hypothesis, at most $|U^F - U_j^F|(T^M - t^m + 1)$ additional fault-free units are replaced as Algorithm 2 continues. Considering the $|U_j^R - U_j^F|$ fault-free units that have already been replaced in the first j steps, a total of at most $|U^F - U_j^F|(T^M - t^m + 1) + |U_j^R - U_j^F|$ fault-free units are replaced. Now, $|U^F - U_j^F| = |U^F| - |U_j^F|$, $|U_j^R - U_j^F| = |U_j^R| - |U_j^F|$, and

$$|U_j^R| \leq 2|U_j^F|$$

from the condition for Case 1. Substitution then implies that at most $|U^F|(T^M - t^m + 1) - |U_j^F|(T^M - t^m)$ fault-free units are replaced. Finally, since $t^m \leq T^M$, $|U^F|(T^M - t^m + 1)$ is a bound on the fault-free units replaced. Since f faulty units are replaced and $|U^F| = f/(T^M - t^m + 2)f$ diagnosability results.

Case 2: At all steps, $|U_j^R| > 2|U_j^F|$. Then it is possible to match each faulty unit with a fault-free unit that is replaced during a step no later than the step during which the faulty unit is replaced. Further, each faulty unit can be paired with a different fault-free unit.

Lemma: The sum of the initially invalid test results on any pair is at least t^m .

Proof: If the fault-free unit was replaced when it failed more than t^m tests, then clearly the claim holds. If the fault-free unit was replaced when it failed $i < t^m$ tests, then at least i test results on the fault-free unit were initially invalid, and at least $t^m - i$ test results on the faulty unit in the pair were invalid, or the faulty unit would have been replaced at a step prior to the replacement of the fault-free unit. Hence, at least $t^m - i + i = t^m$ test results were invalid for the pair. This proves the lemma.

There are $|U^F|$ pairs, so a sum of at least $|U^F|t^m$ tests are invalid for all the pairs. At most $T^M|U^F|$ total invalidations are possible, so any nonpaired fault-free units must have at least one test result invalidated by the remaining $T^M|U^F| - t^m|U^F|$ invalidations. Hence, at most $T^M|U^F| - t^m|U^F|$ nonpaired fault-free units are replaced. With the $|U^F|$ paired ones, a total of $|U^F|(T^M - t^m + 1)$ fault-free are replaced. With the $|U^F| = f$ faulty units, $f/(T^M - t^m + 2)f$ diagnosability results. \square

By considering different system parameters, other bounds may be derived that are at times better than the one in Theorem 4. However, the one given was chosen because it is easy to state and simple to apply. Further, for a given T^M and $t^m \leq T^M$ a system can be constructed that reaches the bound. In such a system, each faulty unit has only one valid test, and each fault-free unit has one invalid test; Algorithm 2 takes two steps; all the tests pass in the second step.

If $t^m > T^M$, at worst $f/2f$ diagnosis is possible because the excess in t^m can only make diagnosis more precise than when $t^m = T^M$.

Better bounds are also possible if system structure is restricted. We examine one such class of systems that occur frequently in the literature. This class of systems can be modeled as in [1], [2] and have $|t(u_i)| = |T(u_i)|$ for all u_i . Members of this class include single-loop systems, D_{δ} systems [1], and systems where u_i tests u_j implies u_j tests u_i . The array of Fig. 1(c) is of the latter type. We call such systems *balanced*. If all the units in a balanced system have the same $|t(u_i)|$, then clearly $T^M = t^m$ and $f/2f$ diagnosis follows immediately from Theorem 4. However, the same result holds for balanced systems where the $|t(u_i)|$ are unequal.

Theorem 5: In a balanced system S with symmetric invalidation, Algorithm 2 results in at worst $f/2f$ diagnosis.

Proof: The proof parallels the proof of Theorem 4 and is omitted due to its length. \square

Example 5: As a consequence of Theorem 5, the arrays of Fig. 1(c) are $f/2f$ diagnosable by Algorithm 2. A more careful analysis shows that much better diagnosability results from the use of Algorithm 2, at least for small f . In particular, $1/1$, $2/2$, $3/3$, $4/5$ diagnosability result. \square

We now consider Algorithm 2 when asymmetric invalidation is assumed.

Theorem 6: For the system S with asymmetric invalidation, the use of Algorithm 2 always results in at worst $f/(\lfloor T^M/t^m \rfloor + 1)f$ diagnosability.

Proof: Due to asymmetric invalidation, any faulty unit must fail at least t^m tests. In order for a fault-free unit to be replaced, it must fail at least as many tests as some faulty unit. Since $T^M f$ is an upper bound on the total number of tests on fault-free units, at most $\lfloor T^M f/t^m \rfloor$ fault-free units can fail t^m or more tests. \square

V. CONCLUSIONS

From a practical standpoint, the algorithms presented here have the following advantages:

- 1) They are easy to describe and easy to understand; they follow intuition in that only units that fail tests are replaced.
- 2) They are computationally simple.
- 3) They do not depend on system structure; the $T(u_i)$ are not needed. This is important because in some cases the $T(u_i)$ may be difficult to obtain.
- 4) The number of good units replaced is usually low; a notable exception may be systems for which the $|t(u_i)|$ are relatively small and where the $|T(u_i)|$ are relatively large. For these systems, some algorithm depending on system structure may be necessary for efficient replacement.

- 5) The number of units that can be faulty is bounded only by a fundamental detectability condition that also appears to bound any nontrivial diagnosis algorithm.

From a more theoretical point of view, we observe that there is a tradeoff between the complexity of a diagnosis algorithm and its efficiency in terms of good units replaced. At one end of the scale is the trivial algorithm that ignores both test results and system structure and simply replaces all units every time. The algorithms given here are another critical point on this scale; here only test results are considered.

There is little doubt that including structural information in the determination of F' can reduce the number of fault-free units that are replaced. However, in many situations the extra structural information can complicate the algorithm. Consequently, research is being directed at designing f/s diagnosis algorithms that take structure into account and which provide a good balance between algorithmic complexity and efficiency in terms of good units replaced.

ACKNOWLEDGMENT

The author would like to thank Prof. C. R. Kime for many interesting discussions on system diagnosis, and Mr. J. A. McPherson for his careful reading of the manuscript and many helpful comments.

REFERENCES

- [1] F. P. Preparata, G. Metze, and R. T. Chien, "On the connection assignment problem of diagnosable systems," *IEEE Trans. Electron. Comput.*, vol. EC-16, pp. 848-854, Dec. 1967.
- [2] F. Barsi, F. Grandoni, and P. Maestrini, "A theory of diagnosability of digital systems," *IEEE Trans. Comput.*, vol. C-25, pp. 585-593, June 1976.
- [3] C. R. Kime, "An analysis model for digital system diagnosis," *IEEE Trans. Comput.*, vol. C-19, pp. 1063-1073, Nov. 1970.
- [4] J. D. Russell and C. R. Kime, "System fault diagnosis: Closure and diagnosability with repair," *IEEE Trans. Comput.*, vol. C-24, pp. 1078-1089, Nov. 1975.
- [5] —, "System fault diagnosis: Masking, exposure, and diagnosability without repair," *IEEE Trans. Comput.*, vol. C-24, pp. 1155-1161, Dec. 1975.
- [6] S. L. Hakimi and A. T. Amin, "Characterization of connection assignment of diagnosable systems," *IEEE Trans. Comput.*, vol. C-23, pp. 86-88, Jan. 1974.
- [7] F. J. Allan, T. Kameda, and S. Toida, "An approach to the diagnosability analysis of a system," *IEEE Trans. Comput.*, vol. C-24, pp. 1040-1042, Oct. 1975.
- [8] S. N. Maheshwari and S. L. Hakimi, "On models for diagnosable systems and probabilistic fault diagnosis," *IEEE Trans. Comput.*, vol. C-25, pp. 228-236, Mar. 1976.
- [9] S. Karunanithi and A. D. Friedman, "System diagnosis with t/s diagnosability," in *Proc. Seventh Annual Int. Conf. Fault-Tolerant Computing*, June 1977, pp. 65-71.
- [10] A. D. Friedman, "A new measure of digital system diagnosis," in *Digest of Papers, 1975 Int. Symp. Fault-Tolerant Computing*, June 1975, pp. 167-170.
- [11] A. M. Corluhan and S. L. Hakimi, "On an algorithm for identifying faults in a T -diagnosable system," in *Proc. 1976 Johns Hopkins Conf. Inform. Sci. Syst.*, Mar. 1976, pp. 370-375.
- [12] G. G. L. Meyer and G. M. Masson, "An efficient fault diagnosis algorithm for symmetric multiple processor architectures," *IEEE Trans. Comput.*, vol. C-27, pp. 1059-1063, Nov. 1978.
- [13] G. G. L. Meyer, "Fault diagnosis of modular networks with a small number of faults," in *Proc. Fifteenth Ann. Allerton Conf. Commun., Contr. and Comput.*, Sept. 1977.
- [14] T. Kameda, S. Toida, and F. J. Allan, "A diagnosing algorithm for networks," *Inform. Contr.*, vol. 29, pp. 141-148, 1975.
- [15] M. Adham and A. D. Friedman, "Digital system fault diagnosis," *J. Design Automation Fault-Tolerant Computing*, vol. 1, pp. 115-132, Feb. 1977.
- [16] C. R. Kime, "A theory of digital system fault diagnosis," *Dep. Elec. Comput. Eng., Univ. of Wisconsin-Madison, Tech. Rep. ECE-77-2*, Feb. 1977.