Fault Diagnosis in Loop-Connected Systems*

KEWAL K. SALUJA

and

BRIAN D. O. ANDERSON

Department of Electrical Engineering, The University of Newcastle, New South Wales, 2308, Australia

ABSTRACT

The paper considers fault diagnosis in a large system comprising a collection of small subsystems or units which can test one another for the existence of a faulty condition. If subsystem α is not faulty and tests subsystem β, a correct indication of the status of β is obtained; if α is faulty, the test outcome contains meaningless information. A particular form of interconnection is examined. For a system with n units $u_0, u_1, \ldots, u_{n-1}$, for each i unit u_i tests $u_{i+1}, u_{i+2}, \ldots, u_{i+A}$ (modulo n arithmetic being understood), where A is a preselected integer. If t is the maximum number of faulty units, we show that when $t \leqslant A$, all faults are immediately diagnosable if $n \geqslant 2t+1$; we also show that when $t > A$, at least A faults can be diagnosed if and only if $n \geqslant s(t-As)+t+A+1$, where s is the integer which maximizes the quadratic function $f(x) = x(t-Ax)$ of the integer variable x.

1. INTRODUCTION

Many hardware and software systems comprise interconnections of smaller subsystems. In checking the reliability of large systems, it can often be that the individual subsystems are used to test one another, so that system fault diagnosis becomes the task of locating one or more subsystems which are faulty.

Based on this simple idea, a number of models have been developed and analyzed for what is now called system level diagnosis (see e.g. [1–5]). In this paper, we shall use one such model, a graph theoretic one developed by Preparata et al. [1]. This model and two of its simplest properties are described in the next section. Then in Sec. 3, the problem dealt with in this paper is

*Work supported by the Australian Research Grants Committee.

52 Vanderbilt Ave., New York, NY 10017 0020-0255/80/04075-18$01.75

formulated, and necessary notation developed. The problem is one of explaining what the relation between the total number of units and maximum number of faulty units must be, in a system of prescribed structure, in order that faults may be identifiable. Section 4 presents a necessary condition for diagnosability. Section 5 presents an algorithm with twofold significance. Firstly, it provides a technique for identifying faulty units. Secondly, it allows a proof in Sec. 6 that the necessity condition of Sec. 4 is also sufficient.

It has come to our notice that Karunanithi [6] has also studied the type of system structure discussed in this paper.

2. KNOWN RESULTS

The large system which is to be diagnosed for faults is composed of n units (e.g., computers, microprocessors, programs, etc.), denoted $u_0, u_1, \ldots, u_{n-1}$. These units need not be identical. Each unit is capable of testing one or more other units by applying a sequence of stimuli and observing the response. The system is represented by a directed graph in which nodes correspond to units and a directed arc b_{ij} is drawn from u_i to u_j if u_i tests u_j. The outcome of the test applied by u_i to u_j is denoted by a binary variable a_{ij}, where $a_{ij}=0$ (1) if u_i determines u_j to be fault free (faulty). Additionally, we assume that the information supplied by a faulty unit is unreliable. We can sum this up as follows:

$a_{ij}=0$ if u_i is fault free and u_j is fault free,
$a_{ij}=1$ if u_i is fault free and u_j is faulty,
$a_{ij}=d$ if u_i is faulty, where $d \in \{0,1\}$.

With this formulation Preparata et al. [1] then posed the problem of finding necessary and sufficient conditions for diagnosability of a system graph consisting of n units and having no more than t faults. It is quickly seen that such conditions are dependent on the structure of the graph and definition of diagnosability. Two types of diagnosability are defined below.

DEFINITION 1 [1]. A system of n units is one step (sequentially) t-fault-diagnosable if all faulty units (at least one faulty unit) can be identified without replacement or repair, provided the number of faulty units present does not exceed t.

Preparata et al. [1] in their work studied two special classes of systems. However, certain general conditions were also established and are stated below.

FACT 1 [1]. *Let a system S be t-fault-diagnosable (one step or sequentially). Then* $n \geqslant 2t+1$.

FACT 2 [1]. *In a one step k-fault-diagnosable system S, each unit is tested by at least k other units.*

3. PROBLEM FORMULATION AND NOTATION

In this section we generalize the definition of sequential diagnosability and define a generalized system structure [2].

DEFINITION 2 [2]. A graph of n nodes is said to be $D_{\delta A}$ if there exists an arc from u_i to u_j whenever $j-i=\delta k \pmod n$ for $k=1,2,\ldots,A$.

Figure 1 is an example of D_{12} for $n=9$.

It is easy to show that if δ and n are relatively prime, $D_{\delta A}$ graphs are isomorphic to D_{1A} graphs [1, 2]. Here we shall confine our attention to D_{1A} systems. With t the maximum number of faults, necessary and sufficient conditions for D_{1t} systems to be one step diagnosable [1] and sufficient conditions for D_{11} systems (single loop systems [1]) to be sequentially diagnosable are given in [1]. Necessary conditions for D_{11} systems to be sequentially diagnosable are developed in [3]. These conditions take the form of inequalities relating n and t.

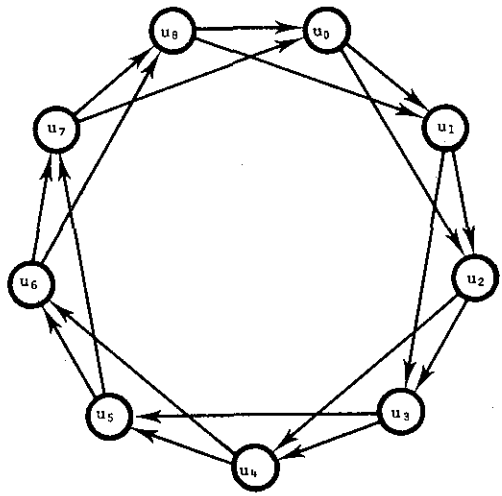

Fig. 1. A D_{12} system with $N=9$.

Therefore the following response pattern is possible:

$$
\begin{array}{l}
0\ldots000\ldots10\ldots00\ldots1 \\
\vdots\quad\vdots\quad\vdots\quad\vdots \\
0\ldots001\ldots10\ldots01\ldots1 \\
\underbrace{0\ldots011\ldots}_{t-A(k-1)}\underbrace{10\ldots11\ldots1}_{A+t-Ak} \\
\qquad\qquad\qquad\qquad\underbrace{\qquad\qquad}_{\text{repeats } k-1 \text{ times}}
\end{array}
$$

$$
\begin{array}{l}
\ldots0\ldots00\ldots10\ldots00\ldots1 \\
\quad\vdots\quad\vdots\quad\vdots\quad\vdots\quad\vdots \\
\ldots0\ldots01\ldots10\ldots01\ldots1 \\
\underbrace{\ldots0\ldots11\ldots}_{A+t-Ak}\underbrace{10\ldots11\ldots1}_{t-A(k-1)}
\end{array}
$$

In these repetitive strings it is impossible to locate the faulty units. For a conclusion which locates a faulty unit in any block of size $t-A(k-1)$ will imply the corresponding unit is faulty in every other block. That such a conclusion is not true is evident from the fault pattern.

Clearly in the above configuration

$$
\begin{aligned}
\text{total number of } N\text{'s} &= \sum_{i=1}^{k} n_i \\
&= (t-Ak)(k-1)+t-A(k-1) \\
&= (t-Ak)k+A,
\end{aligned}
$$

$$
\begin{aligned}
\text{total number of } F\text{'s} &= A(k-1)+t-A(k-1) \\
&= t,
\end{aligned}
$$

so that

$$n' = (t-Ak)k+A+t.$$

We now choose k, keeping A and t fixed, in order to maximize n'. Setting $dn'/dk = 0$ yields $k = t/2A$. However, k must be integer. Therefore n' will be maximized for $k = \lceil t/2A \rceil$ or $\lfloor t/2A \rfloor$. It can be checked that with $A \leqslant t$, these values of k satisfy $t-Ak \geqslant 0$.

Since an n' unit system is not diagnosable, it is necessary if an n unit system is to be diagnosable that

$$n \geqslant n'+1$$

or

$$n \geqslant \max_k \{(t-Ak)k+A+t\}+1$$

$$= \max\{m_1, m_2\}+1. \quad \blacksquare$$

We remark that since (as shown in the above proof)

$$\max\{m_1, m_2\} = \max_k \{(t-Ak)k+A+t\}$$

$$\geqslant \{(t-Ak)k+A+t\}_{k=1}$$

$$= 2t,$$

the necessity condition of Theorem 1 is automatically at least as demanding as the condition $n \geqslant 2t+1$ of Fact 1.

We also remark that when $A=1$, the condition becomes

$$n \geqslant \left\lceil \frac{t}{2} \right\rceil \left\lfloor \frac{t}{2} \right\rfloor + t + 2,$$

and this can be easily checked to be equivalent to the condition of [1, 3], which has been shown to be necessary and sufficient for single loop systems to be sequentially diagnosable.

The case $A=2$ is studied in [4].

The case of $A > \lfloor t/2 \rfloor$ is studied in [2]. Actually, we can easily consider $A \geqslant \lceil t/2 \rceil$ to obtain a slightly sharper result for the case when t is even: a necessary, and as Sec. 6 shows, sufficient condition for two-step sequential diagnosability is

$$n \geqslant 2t+1.$$

5. SUFFICIENT CONDITION: PARTITIONING ALGORITHM

In this and the next section we shall show that the necessity condition of Sec. 4 is also a sufficient condition for a D_{1A} system to be sequentially diagnosable.

The first part of the proof involves a partitioning of the units of the system into a set of blocks, each block containing adjacent units. The blocks have certain key properties. For example, in each block either at least half the units are faulty, or, if it is known that the block contains a minimal number of faulty units consistent with the response vectors of the units within it, then one unit

of the block is identifiably normal (not faulty). The second part of the proof, contained in Sec. 6, involves showing that if the necessity condition of the previous section holds, at least one of the blocks must contain a minimal number of faulty units. This identifies one normal unit, and from this normal unit one may then clearly diagnose at least A faults in a D_{1A} system (assuming that the system contains at least A faults).

Before presenting the algorithm proper, an observation and some definitions are in order.

Consider k successive nodes and the associated subgraph and response vectors. From the response pattern within the subgraph, one can work out the various fault patterns or labelings of the nodes as normal or faulty, consistent with these response vectors. (A number of fault patterns, rather than just one, may in fact be consistent with the given response pattern in the subgraph.) By considering all such fault patterns we can then find a fault pattern which has the minimum number of nodes labeled as F. In general such a fault pattern may not be unique and may not correspond to the true distribution of faulty and normal units in the subgraph. However, it is evident that the true number of faults must be no smaller than the minimum number of faults determined as described above. It is important to note that because k is finite, *the determination of the minimum number of faults can proceed with a finite algorithm*, which is actually easy to make highly systematic.

DEFINITION 3. A block is a collection of successive nodes together with the associated subgraph and response vectors.

DEFINITION 4. A block is of type α if the response vectors of the block force the minimum number of faults in the block to be at least half the number of units in the block.

DEFINITION 5. A block is of type β if the response vectors of the block are such that when the minimum number of faults are present:

(a) the minimum number of faults is less than half the number of units in the block,

(b) the last A units are always faulty,

(c) the unit immediately prior to the last A units is normal (its response vector is therefore all 1's).

DEFINITION 6. The rightmost node of a block is the last node of a block.

In what follows we present the algorithm to partition the system into blocks. In each step of the algorithm the properties of the resulting blocks are stated in the form of lemmas. A flow chart describing the algorithm is given in Fig. 2.

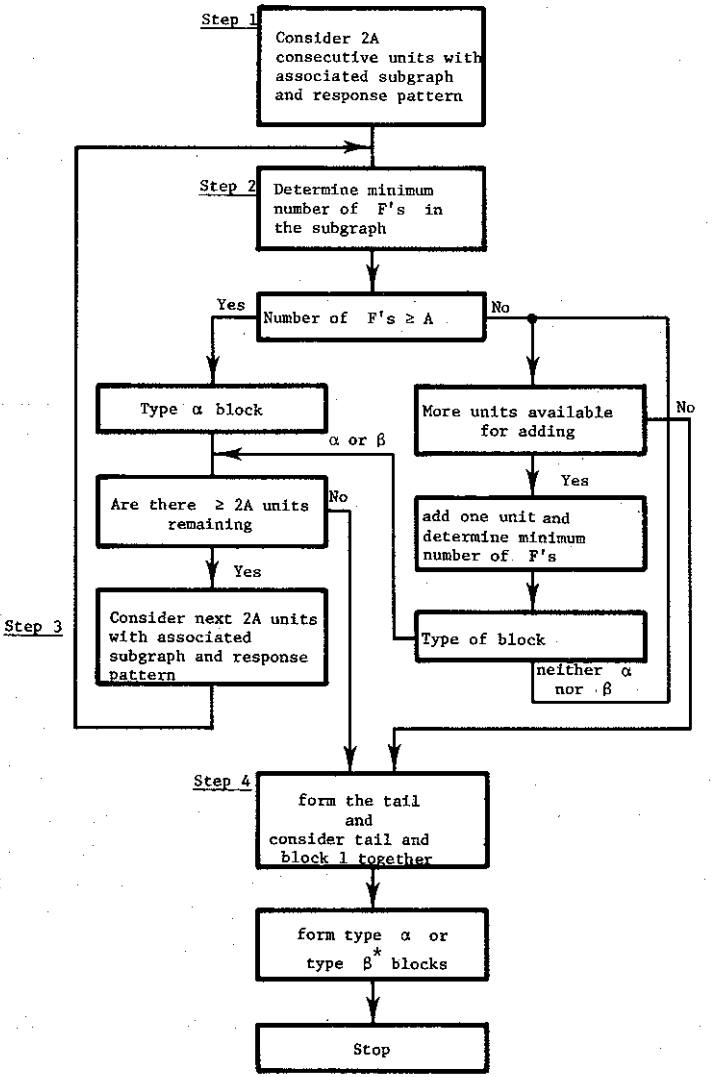

Fig. 2. Flow chart of the partitioning algorithm.

ALGORITHM AND LEMMAS

Step 1

Pick an arbitrary node u_i and form the block comprising $u_i, u_{i+1}, \ldots, u_{i+2A-1}$.

Step 2

Establish the minimum numbers of faults these $2A$ nodes must have. Two cases arise, according as the minimum is below A or not.

Case A: Minimum number of faults $\geq A$. The block is a type α block. Go to Step 3.

Case B: Minimum number of faults $< A$.

LEMMA 1. *Suppose that a block of length at least* $2A$ *can achieve a minimum number of faults* $f < A$. *Let* p *denote the number of normal units when there are* f *faults. Then if* f *faults are present, at least one unit is identifiable as normal. If this unit is faulty, at least* p *faults are present.*

Proof. Let $\gamma_1, \ldots, \gamma_p$ and $\delta_1, \ldots, \delta_p$ be two sets of normal units consistent with a minimum number of faults being present. Then $p > A$. Suppose $\gamma_p \neq \delta_p$. We shall deduce a contradiction.

Suppose γ_p is nearer the block end than δ_p. Because there are fewer than A faults, γ_i must test γ_{i+1} for all $i \in [1, p-1]$ and δ_p tests all units to the end, including γ_p. Since δ_p is the rightmost normal unit in one assignment with minimum number of faults, δ_p tests γ_p with a 1. Then since γ_i tests γ_{i+1} (with a zero) for $i \in [1, p-1]$, δ_p normal implies $\gamma_p, \gamma_{p-1}, \ldots, \gamma_1$ are faulty. But $p > A$ and the number of faults with $\delta_1, \ldots, \delta_p$ normal is less than A, providing the contradiction.

We now form a new block by adding one unit at the right, and redetermine the minimum number of faulty units for the new block.

LEMMA 2. *Suppose that a block of length at least* $2A$ *can achieve a minimum number of faults* $f < A$, *and let* γ_p *be the rightmost of the* p *normal units when this minimum is attained. Suppose one further unit is added to the right of this block. Then one of the following two possibilities holds*:

(i) *the extended block is of type* α,

(ii) *the extended block is not of type* α, *and* γ_p *is normal when the minimum number of faults is achieved.*

Proof. Suppose the extended block is not of type α, and γ_p is faulty when the minimum number of faults is attained. Then arguing as in the proof of Lemma 1, $\gamma_{p-1}, \ldots, \gamma_1$ are also faulty; with p faults in all, we should have $p = 2A - f > A$, or $p \geqslant A + 1$, and with $2A + 1$ units in all, the block must be type α, a contradiction. ■

In case (ii), it is clear that when the minimum number of faults is achieved, we can identify the rightmost unit guaranteed to be normal. It will be γ_p if all units tested by γ_p and within the block are tested with 1's, or some unit to the right of γ_p otherwise. There are at most $f+1$ units to the right of γ_p, and $f+1 \leqslant A$.

More generally, we have the following result:

LEMMA 3. *Suppose that a block is not of type α, the rightmost normal unit is known when the minimum number of faults is achieved, and there are no more than $A-1$ units to the right of this particular unit. Suppose the block is extended on the right by one unit. Then either*

(i) *the extended block is of type α,*

(ii) *the extended block is not of type α, and γ_p is the rightmost normal unit when a minimum number of faults are present, with no more than A units of the extended block to its right, or*

(iii) *the extended block is not of type α, and there is a rightmost normal unit to the right of γ_p when a minimum number of faults are present, and therefore with no more than $A-1$ units of the extended block to its right.*

The proof is a trivial modification of that used for Lemma 2.

The process of extending the block by adding one unit at a time can clearly be continued until one of the following possibilities is attained:

(a) the extended block is of type α,

(b) the extended block is not of type α, and when the minimum number of faults is attained, there is a rightmost unit which is guaranteed to be normal, and this unit is followed by A faulty units (i.e., the block is of type β), or

(c) no further units are available for addition.

With either of the first two possibilities we go to Step 3. For possibility (c), we go to Step 4.

Step 3

Let the number of units not so far assigned to a block be y.

Case A. If $y \geqslant 2A$, consider the first $2A$ nodes following those assigned to blocks, and go to Step 2.

Case B. If $y<2A$, go to Step 4.

Step 4

When Step 4 is encountered, all units have been assigned to blocks of type α or β, save a consecutive sequence from which, for one reason or another, no further type α or type β block can be formed. We shall call this remaining group of units the tail. Of course, the tail may evanesce in certain situations. In this case, the algorithm ends.

Number the blocks sequentially, so that block 1 follows the tail. Let n_i,f_i denote the number of normal and faulty units in block i when the minimum number of faults is attained in that block, and let n_t,f_t denote the corresponding quantities for the tail. Let k be the number or blocks excluding the tail. For a type α block $f_i \geqslant n_i$, and for a type β block $n_i>f_i$.

LEMMA 4. *Either* $n_t>f_t$ *or* $n_i>f_i$ *for some* $i\in[1,k]$.

This is an immediate consequence of Fact 1.

LEMMA 5. *If* $n_t>f_t$ *and the tail contains a minimum number of faults, a normal unit can be identified in the tail. If this unit is actually faulty, there exist at least* n_t *faulty units in the tail.*

Proof. If $n_t>f_t$, the tail arises either from Step 2, Case B (an attempt to form a type β block fails through insufficiency of units) or Step 3, Case B (fewer than $2A$ units in the tail). Minor variation of the proofs of Lemma 2 and 3 gives the result. ■

It is possible that the only block is the tail, in which case the algorithm terminates. This case is labeled case A in the collection given below, and Lemma 5 is crucial to its analysis. The range of cases is as follows:

Case A: The only block is the tail block.
Case B: $n_t+n_1\leqslant f_t+f_1$.
Case C: $n_1\leqslant f_1$, $n_t>f_t$, $n_t+n_1>f_t+f_1$.
Case D: $n_1>f_1$, $n_t\leqslant f_t$, $n_t+n_1>f_t+f_1$.
Case E: $n_1>f_1$, $n_t>f_t$.

Case B. Combine the tail block and block 1 into a single new block 1. At least half the units in the composite block must be faulty, and it is a type α block.

The algorithm terminates with the total number of blocks in the partition being k.

Case C.

LEMMA 6. *Under Case C, if f_t+f_1 is the actual number of faults in a composite block comprising the tail and block 1, there exists an identifiable normal unit in the tail. If this unit is actually faulty, the number of faults in the composite block is at least n_t+n_1.*

Proof. The first conclusion follows from Lemma 5. If the unit referred to in the lemma statement is faulty, there exist at least n_t faulty units in the tail, while the minimum number of faults in block 1 is $f_1 \geqslant n_1$. ■

Combine the tail block and block 1 into a single new block 1. Call this block type β^*, with the following definition:

DEFINITION 7. A block is of type β^* if the response vectors of the block are such that

(a) there exist integers n_i, f_i with n_i+f_i the total number of units, and with $n_i > f_i \geqslant A$,
(b) the minimum number of faults is at least f_i,
(c) there exists a certain unit with the property that if it is faulty, there are at least n_i faulty units in the block.

Any type of β block is also a type β^* block. The algorithm terminates with the total number of blocks in the partition equal to k.

Case D.

LEMMA 7. *Under Case D, if f_t+f_1 is the actual number of faults in a composite block comprising the tail and block 1, there exists an identifiable normal unit in block 1. If this unit is actually faulty, the number of faults in the composite block is at least n_t+n_1.*

Proof. Block 1 must be a type β block, since $n_1 > f_1$. If the composite block has f_t+f_1 faults, block 1 has only f_1. Therefore, by Lemma 3, it has an identifiable normal unit. If the unit is faulty, there must exist n_1 faults in block 1 and at least $f_t \geqslant n_t$ faults in the tail block, i.e. a total of at least n_t+n_1 faults. ■

Combine the tail block and block 1 into a single new block 1, which again will be a type β^* block. The algorithm terminates with the total number of blocks in the partition equal to k.

Case E. Notice that the conditions force block 1 to be of type β. The general idea in case E is to transfer, one unit at a time, units from the left hand

end (start) of block 1 to the right hand end of the tail block. This process is continued until one or more of the following happens:

(a) The new blocks have $n_t + n_1 \leqslant f_t + f_1$.
(b) The new block 1 is a type α block.
(c) The new tail block is a type α block.
(d) The new tail block has $f_t \geqslant A$.

In case $f_t \geqslant A$ before any transfers, the tail block is type β^* and the algorithm ends. We now explain the transfer idea in more detail, indicating why one of these possibilities must in due course be encountered.

LEMMA 8. *Let block* 1 *be a type* β *block, and transfer one unit from the left hand end of block* 1 *to the right hand end of the tail block. With* n_t, f_t, n_1, f_1 *having their usual meanings for pretransfer quantities and* $n_t^*, f_t^*, n_1^*, f_1^*$ *for posttransfer quantities, the following hold*:

$$n_1^* = n_1, \quad f_1^* = f_1 - 1 \quad \text{or} \quad n_1^* = n_1 - 1, \quad f_1^* = f_1$$

and

$$n_t^* \leqslant n_t, \quad f_t^* \geqslant f_t + 1 \quad \text{or} \quad n_t^* = n_t + 1, \quad f_t^* = f_t.$$

Moreover, the new block 1 *is either type* α *or type* β.

Proof. If the leftmost unit (before transfer) in block 1 is labeled faulty in any labeling consistent with minimum fault count, it is clear that $n_1^* = n_1$, $f_1^* = f_1 - 1$. Suppose this unit is labeled normal in all minimum fault count labelings. Clearly $n_1^* \geqslant n_1 - 1$. Suppose $n_1^* = n_1$. Then apply labels to the shortened block 1 consistent with a minimum fault count, and then add a faulty unit at the left end. This gives the original block 1 with a minimum fault count and with leftmost unit faulty, a contradiction. Hence $n_1^* = n_1 - 1$, and so $f_1^* = f_1$.

In case $n_1^* > f_1^*$, a very minor extension of this argument shows that the new block 1 must be type β. Otherwise, $n_1^* = f_1^*$, since $n_1 > f_1$ prevents $n_1^* < f_1^*$. In this case, the block is of type α.

If there exists a labeling of normal and faulty nodes in the original tail with minimum fault count such that any node labeled normal and pointing to the first unit of the original block 1 points with a zero and none points with a one, then clearly $n_t^* = n_t + 1$, $f_t^* = f_t$. One easily argues that otherwise, $n_t^* \leqslant n_t$, $f_t^* \geqslant f_t + 1$. ∎

The unit by unit transfer is, as noted above, repeated until one or more of conditions (a) through (d) is first fulfilled. If (a), (c), or (d) is never fulfilled, (b) must be eventually. This is because block 1 is continually shrinking, and cannot maintain its type β character when it contains $2A$ units.

When condition (a) is first fulfilled, go to Case B.

When condition (b) is first fulfilled and condition (a) is not, go to Case C.

When condition (c) is first fulfilled and condition (a) is not, go to Case D.

When condition (d) is first fulfilled and conditions (a), (b), and (c) are not, Lemma 5 is applicable to the new tail, which is type β^*, and the algorithm terminates with the total number of blocks in the partition equal to $k+1$.

The end result is that either every block is a type β or type β^* block, or that the tail block is the only block, in which case Lemma 5 applies.

6. SUFFICIENCY CONDITION: PROOF BASED ON PARTITION

We first dispense with a simple case. Suppose the only block resulting from the algorithm is the tail block. This means that the partitioning algorithm gets to Step 2, Case B, and in attempting to form a type β block, does not succeed before all units are exhausted.

LEMMA 9. *Suppose that the algorithm of Sec. 5 leads to a partition with only one block, a tail block. Then $n \geq 2t+1$ is a sufficient condition to identify all or at least A faults, whichever is the lesser.*

Proof. The condition $n \geq 2t+1$ ensures that $n_t > f_t$, and by Lemma 5 a unit is identifiable as normal; for if the unit referred to in the statement of Lemma 5 were faulty, there would be n_t faults at least, and the condition $n \geq 2t+1$ would be violated, since $n = n_t + f_t$, $n_t > f_t$, implying $t \geq n_t$.

Now trace paths of arcs emanating from this unit with the test outcomes on each arc in the path comprising all zeros save for the outcome of the last arc of the path, which is one. This last arc points to a faulty unit, and in this way, one can identify all faults or at least A faults. ■

As noted in Sec. 4, with

$$m_1 = \left\lceil \frac{t}{2A} \right\rceil \left(t - A \left\lceil \frac{t}{2A} \right\rceil \right) + t + A, \tag{6.1a}$$

$$m_2 = \left\lfloor \frac{t}{2A} \right\rfloor \left(t - A \left\lfloor \frac{t}{2A} \right\rfloor \right) + t + A, \tag{6.1b}$$

the condition $n \geq \max(m_1, m_2) + 1$ certainly ensures $n \geq 2t+1$. The former condition is the key one, as we now show.

THEOREM 2. *With $t \geq A$, a D_{1A} system is sequentially $\lceil t/A \rceil$ step t-fault-diagnosable if*

$$n \geq \max(m_1, m_2) + 1. \tag{6.2}$$

Proof. Lemma 9 establishes the result if the tail block is the only block resulting from the algorithm. So suppose every block is type α or type β^*. Since at least half the units of a type α are faulty, there must be at least one type β^* block present whenever $n \geqslant 2t+1$, a condition ensured in this case.

Define variables r_i and x_i associated with block i as follows:

Type α:

$$x_i = \left\lceil \tfrac{1}{2} \text{ (number of units in block)} \right\rceil,$$

$$r_i = 0.$$

Type β^*:

$$x_i = f_i - A \geqslant 0,$$

$$r_i = n_i - x_i = n_i - f_i + A > A.$$

In a type α block, there are at least x_i faults. In a type β^* block, the minimum number of faults is at least $x_i + A$. If this number is achieved, a certain unit must be identifiably normal. If the unit is not normal, the number of faults is at least $n_i = r_i + x_i > A + x_i$. This gives rise to the following crucial observation. Let $r_{\max} = \max_i r_i$, and suppose there are $q > 0$ blocks of type β^* and s blocks in all. Suppose

$$r_{\max} + \sum_{i=1}^{s} x_i + (q-1)A \geqslant t+1. \tag{6.3}$$

Then a normal unit can be identified in the block (or blocks) associated with $r_{\max}$. For the left side of (6.3) represents the sum of the minimum number of faults which arise when in a block associated with one $r_{\max}$ a certain unit is faulty, and all other blocks have the minimum number of faults.

Let $x = \sum_{i=1}^{s} x_i$, so that (6.3) becomes

$$r_{\max} + x + (q-1)A \geqslant t+1. \tag{6.4}$$

The remainder of the proof is concerned with showing that (6.2) implies (6.4). (This in turns allows identification of a normal unit and then at least A faults.) There are two cases.

Case A: $q = 1$. As noted above, (6.2) implies $n \geqslant 2t+1$, and since $r_{\max} > A$ or $r_{\max} - A \geqslant 1$, (6.2) therefore implies

$$n + r_{\max} - A \geqslant 2t+2. \tag{6.5}$$

Also, by simply counting up the units in the various blocks, we have

$$n \leqslant 2x + r_{\max} + A. \tag{6.6}$$

(The inequality can only arise if there are type α blocks with an odd number of units.) Now (6.6) and (6.5) imply

$$2x + 2r_{\max} \geqslant 2t + 2,$$

or (6.4) with $q = 1$.

Case B: $q \geqslant 2$. By counting up the units in the various blocks, we see that

$$n \leqslant (r_{\max} + A)q + 2x. \tag{6.7}$$

Now (6.4) holds provided[3]

$$qr_{\max} + qx + q(q-1)A \geqslant qt + 1,$$

and *a fortiori,* in the light of (6.7), if

$$n - Aq - 2x + qx + q(q-1)A \geqslant qt + 1,$$

i.e.,

$$n + (q-2)x \geqslant -q^2A + q(t+2A) + 1. \tag{6.8}$$

Because $q \geqslant 2$, (6.8) is guaranteed, and thus (6.4) also, if

$$n \geqslant -q^2A + q(t+2A) + 1$$

or if

$$n \geqslant \max_q \left[-q^2A + q(t+2A) \right] + 1. \tag{6.9}$$

The function $-q^2A + q(t+2A)$ regarded as a quadratic in a real variable q takes its maximum value at $q = (t/2A) + 1$. Accordingly, as a function of an integer variable q, it takes its maximum value at one or both of

$$q_1 = \left\lceil \frac{t}{2A} \right\rceil + 1, \qquad q_2 = \left\lfloor \frac{t}{2A} \right\rfloor + 1.$$

[3]The fact that all variables are integer valued is crucial.

Substituting these values into (6.9) yields, as required, the condition

$$n \geqslant \max(m_1, m_2) + 1. \qquad \blacksquare \quad (6.2)$$

7. CONCLUSIONS

The paper has given a necessary and sufficient condition for the sequential diagnosability of a system comprising a collection of interconnected subsystems with a certain pattern for the interconnections. The condition is readily checked, involving comparison of the number of subsystems with a readily calculated integer function of two other integer variables: the maximum number of faults and a parameter associated with the pattern. The result has embedded within it a number of other results obtained earlier.

A number of problems remain. It might, for example, be interesting to study in detail the computational complexity of the fault diagnosis algorithm presented in the paper. Again, more general structures than the so-called $D_{1,A}$ structure of the paper could be studied. With n the number of units and with (δ, n) coprime, it is known that a $D_{\delta,A}$ structure is isomorphic to the $D_{1,A}$ structure, and so this case is readily dealt with. But there is as yet no tool for the case of δ, n not coprime. Greater deviations again from the $D_{1,A}$ structure could also be contemplated.

REFERENCES

1. F. P. Preparata, G. Metze, and R. T. Chien, On the connection assignment problem of diagnosable systems, *IEEE Trans. Computers* EC-16(12):848–854 (Dec. 1967).
2. S. Karunanithi and A. D. Friedman, Analysis of digital systems using a new measure of system diagnosis, *IEEE Trans. Computers* C-28(2):121–133 (Feb. 1979).
3. F. P. Preparata, Some results on sequential diagnosable systems, in *Proceedings of the Hawaii International Conference on Systems Science*, Jan. 1968, pp. 623–626.
4. K. K. Saluja and B. D. O. Anderson, t-fault $t/2$-step sequentially diagnosable systems, *Proceedings of IEEE*, Vol 67 (12): 1678–79 (Dec. 1979)..
5. K. K. Saluja, System level diagnosis: A Survey, in *Proceedings of the IREE International Electronics Convention*,, Sydney, Aug. 1979.
6. S. Karunanithi, System level self-diagnosis of digital systems, Ph. D. dissertation, Department of Electrical Engineering Univ. of Southern California, Los Angeles, Calif. 1978.

Received July 1979