SYSTEM LEVEL DIAGNOSIS WITH LOCAL CONSTRAINTS

A. Das (i), K. Thulasiraman (i),
K.B. Lakshmanan (ii), V.K. Agarwal (iii)

(i) E.E. Dept., Concordia University, Montreal
(ii) Dept. of Math. and Comp. Sc., SUNY, Brockport, N.Y.
(iii) E.E. Dept., McGill University, Montreal

ABSTRACT

In this paper, the concept of a local fault constraint in system-level diagnosis for multiprocessor systems is introduced. Given local constraints, distributed and sequential algorithms for a ring of processors as well as other regular interconnected structures are presented. In all cases considered, the maximum number of faults that can be diagnosed in the local domain is determined. The number of faulty processors that can be diagnosed when information about local constraints is available is significantly larger than what is allowed by the classical t-diagnosis theory. In addition, the number of permissible fault patterns is also significantly higher. These advantages of diagnosis under local constraints over the classical t-diagnosis approach are striking when the connectivity of the system is much smaller than the total number of processors in the system.

I. INTRODUCTION

The concept of system-level fault diagnosis originated by Preparata, Metze and Chien [1] provides a potential framework for diagnosis of multiprocessor systems . Here each processor tests some of the other processors and produces test results, which are unreliable if the testing processor is itself faulty. The collection of all test results over the entire system is referred to as a syndrome.

The classical constraint used in the study of diagnosable systems is to assume that the number of faulty processors in the entire system is upper bounded by an integer t . The resulting t-diagnosis problem has been studied extensively in the literature [2,3,4]. The fault diagnosis of a given multiprocessor system clearly depends on the form of testing, the type of faults and the interpretation assigned to the test results [5,6,7]. We consider the comparison form of testing proposed by Chwa and Hakimi [8] and we assume that only permanent faults are present in the system.

In the classical t-diagnosability theory the value of t, and hence the largest number of processors that can be diagnosed, is limited by the connectivity of the processor interconnection graph of the system. This theory is of little value when applied to regular interconnected multiprocessor systems realised on a single chip. In these systems, it is expected that each processor can test only the neighboring processors in the interconnection. For example, in a rectangular grid-wrapped system, each processor is connected to exactly four processors, and hence the classical theory permits correct identification of all processors, only as long as the number of faulty processors in the entire system is at most four !

Instead of a single global constraint, in this paper, we consider local constraints on the number of faulty processors in the neighborhood of each processor in the multiprocessor system. The basic model and definitions are introduced in Section 2. In Section 3, we analyze the implication of local fault constraints to system-level diagnosis on a ring of processors. In Section 4, we study regular interconnected systems such as the closed rectangular grid, hexagonal grid and the octagonal grid when fault constraints are imposed on every local domain consisting of a processor and all its adjacent processors.

II. PRELIMINARIES

A multiprocessor system consists of n independent processors, or processors, $U=\{u_1,u_2,...,u_n\}$. In the comparison model of multiprocessor fault diagnosis, all processors in S are assigned to perform the same task. Upon completion, the outputs of some pairs of these processors are compared. The comparison assignment can be represented by an undirected graph $G=(U,E)$ where an edge e_{ij} belongs to E if and only if the outputs of u_i and u_j are compared.

An outcome a_{ij} is associated with each pair of processors whose outputs are compared, where $a_{ij}=0(1)$ if the outputs compared agree(disagree). Since only permanent faults are considered and we assume that the outputs of a fault-free and a faulty processor always disagree, it follows that $a_{ij}=0$ whenever both u_i and u_j are fault-free; $a_{ij}=1$ if one of u_i and u_j is fault-free and the other faulty; a_{ij} is unreliable if both u_i and u_j are faulty. $N(u_i)$ denotes the set of neighbors of u_i i.e. the set of all processors adjacent to u_i. An edge that has a 0(1) outcome associated with it is referred to as a 0-link(1-link). $N_0(u_i)$ and $N_1(u_i)$ denote the set of processors adjacent to u_i which are connected with u_i by a 0-link and a 1-link respectively.

A fault set $F \subseteq U$ is a permissible fault set for a set of fault constraints if F satisfies the requirements of the fault constraints. Given a syndrome, F is an allowable fault set if and only if F is a permissible fault set, and the assumption that the processors in F are faulty and the processors in $U-F$ are fault-free is consistent with the given syndrome.

A processor u_j is said to belong to the local domain $L_k(u_i)$ if and if u_j lies within a distance k of u_i.

In this paper we study the fault diagnosis properties of regular interconnected systems such as the closed rectangular grid, the hexagonal grid and the closed octagonal grid when fault constraints are imposed on $L_k(u)$ for every processor u in the system. A system S is defined to be t-in-L_1 diagnosable if given a syndrome, all faulty processors can be uniquely identified provided that there are at most t faulty processors in $L_1(u)$ for every processor u in S.

III. DIAGNOSIS OF A RING OF PROCESSORS WITH LOCAL FAULT CONSTRAINTS

In this section, we analyze the implication of imposing local fault constraints on a ring of processors. Specifically, we wish to determine if, given a syndrome, we can uniquely determine the set of faulty processors as long as at most p out of any q consecutive processors are faulty. We also wish to develop diagnosis algorithms for these systems.

ISCAS '89

Theorem 1. Let S be a ring of N processors where N is even. Given that at most p processors are faulty out of any q consecutive processors the values for p and q which admit the maximum number of fault sets which can be uniquely diagnosed are p=2 and q=5.

Proof: Let $\{u_1, u_2, \ldots, u_n\}$ be the ring of processors.

case 1: $3 \leq p \leq n$.

We show that in this case, the set of permissible fault sets cannot be uniquely diagnosed. Consider the following syndrome: $a_{N1}=1$, $a_{12}=1$, $a_{23}=1$, $a_{34}=1$ and all other outcomes have value 0. Both $F_1=\{u_1, u_2, u_3\}$ and $F_2=\{u_1, u_3\}$ are allowable fault sets for this syndrome.

case 2: p=2, $q \geq 3$.

case 2.1: p=2, q=3 .

Consider the following syndrome: $a_{N1}=1$, $a_{12}=1$, $a_{23}=1$, $a_{34}=1$, $a_{45}=1$ and all other outcomes are 0. The fault sets $F_1=\{u_1, u_3, u_4\}$ and $F_2=\{u_1, u_2, u_4\}$, two permissible fault sets under the given fault constraint, are allowable faults for this syndrome.

case 2.2: p=2, q=4 .

Since N is even, let F be a fault set containing alternate processors in S. Then the fault sets F and F^c are allowable fault sets for the syndrome in which all outcomes are 1.

case 2.3: p=2, $q \geq 5$.

We note that if there are at most 2 faulty processors in any consecutive 5 processors then for any processor u, $L_2(u)$ which consists of 5 processors contains at most 2 faulty processors. Thus, given a permissible syndrome, the local diagnosis algorithm developed in [9] can be used to identify all processors correctly. Since the constraint p=2 and q=5 permits all fault sets which are valid when p=2 and $q \geq 5$, these values for p and q admit the maximum number of fault sets which can be uniquely diagnosed. //

If a fault constraint permits a fault set F and its complement U-F to be permissible fault sets, then given a valid syndrome, the faulty processors may not be correctly identified; the fault sets F and U-F generate a common syndrome. We note that if initially one processor v is correctly determined to be fault-free or less than half the total number of processors in the system are faulty then for any subset F of U, at most one of the subsets F and U-F can be an allowable fault set for a given syndrome.

Theorem 2. Let S be a ring of processors in which one of the following conditions is satisfied: (1) some processor is given to be fault-free
(2) less than half processors in the system are faulty
(3) N is odd.

Then the values for p and q which permit the maximum number of fault sets which can be uniquely diagnosed in S under the local constraint of at most p faulty processors in any q consecutive processors are p=2 and q=4 respectively.

Proof: Assume p=2 and q=4. We first show that if one fault-free processor v is given or found to be fault-free, all other processors can be identified correctly. We assume that the diagnosis procedure initiated at v proceeds clockwise. If a processor is fault-free then the adjacent processor can be correctly identified. If two consecutive processors are identified as faulty then the next processor can be correctly identified as fault-free. Thus only the situation shown below could pose a problem.

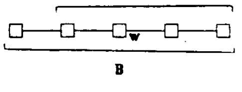

Since there are at most 2 faulty processors in any 4 consecutive processors, there are at most 2 faulty processors in A. Hence B contains at most 2 faulty processors. The processor w has at most 2 faulty processors in its local neighborhood $L_2(w)$. Thus a local diagnosis algorithm can be carried out with respect to processor w to determine its status. Thus if one processor is given to be fault-free or can be identified correctly to be fault-free then all other processors can be identified correctly.

We now show how one processor can be identified correctly if either (2) or (3) is true. We note that if a valid syndrome contains the sequence of consecutive outcomes 00, 011 or 110 then the processors adjacent to the 0-links are fault-free; for otherwise there is a sequence of 4 consecutive processors of which at least three are faulty.

We now claim that for a valid syndrome one of the following sequence of outcomes 00, 011 or 110 occurs. Assume the contrary. Then the following syndromes are the only syndromes which do not contain any of the sequences 00, 011 or 110: the syndrome s_1 in which all outcomes have value 1 and the syndrome s_2 in which 0 and 1 outcomes alternate.

case 1: Less than half the processors are faulty.

In this case, S contains two consecutive fault-free processors. Hence there exists at least one 0-link and the syndrome s_1 cannot occur. Since at most 2 of any 4 consecutive processors can be faulty, the syndrome s_2 corresponds to fault sets in which two faulty processors are followed by two fault-free processors and vice versa. But this contradicts the assumption that the number of faulty processors is less than the number of fault-free processors.

case 2: N is odd.

In this case, the syndrome s_2 cannot be present. Since at most 2 out of any 4 consecutive processors can be faulty, the syndrome s_1 corresponds to fault sets in which faulty and fault-free processors alternate; this is not possible since N is odd and S is a ring of processors.

This shows that one fault-free processor can be determined if either (2) or (3) is true. //

We observe that the diagnosis algorithm outlined in the proof of Theorem 2 can be designed to run sequentially on a host processor or in a distributed manner on the ring of processors.

IV. t-in-L_1 DIAGNOSIS OF REGULAR INTERCONNECTED SYSTEMS

The following Lemma is useful in proving the diagnosability results that appear in this section.

Lemma 1. Given a system S and a syndrome, let X_1 and X_2 be two distinct allowable fault sets for the given syndrome such that $X_1 \cup X_2 \neq U$ and for all processors $u \in U$, $L_1(u)$-X_1 and $L_1(u)$-X_2 are both non-empty. Then there exist processors $x, y \in U$ such that

(1) $x \in U - (X_1 \cup X_2)$
(2) $y \in (X_1 \cap X_2^c) \cup (X_2 \cap X_1^c)$
(3) $2 \leq d(x,y) \leq 3$

Proof: Since $X_1 \cup X_2 \neq U$ the set $U - (X_1 \cup X_2)$ is non-empty. Furthermore, since X_1 and X_2 are distinct, there exists at least one processor which belongs to one fault set and is not contained in the other. Thus there exist processors in U satisfying conditions (1) and (2). Now let x and y be processors in U satisfying conditions (1) and (2) respectively such that the distance $d(x,y)$ is minimum.

Assume $d(x,y) \geq 4$. Consider a processor w such that w is at a distance of at most $\lceil d(x,y)/2 \rceil$ from both x and y. Since $L_1(w)$-X_1 and $L_1(w)$-X_2 are both non-empty, there exists a processor $z \in L_1(w)$ satisfying condition (1) or (2). If z

satisfies condition (1), then $d(z,y) \leq d(w,y) + 1 < d(x,y)$; if z satisfies condition (2), then $d(x,z) \leq d(x,w) + 1 < d(x,y)$. In either case, the minimality of $d(x,y)$ is contradicted. Hence $d(x,y) \leq 3$.

To prove that $d(x,y) \geq 2$, we show that the assumption $d(x,y) = 1$ leads to a contradiction. Assume $d(x,y) = 1$. Then the link between x and y is a 0-link with respect to one fault set and a 1-link with respect to the other, contradicting the assumption that X_1 and X_2 share a common syndrome. //

Theorem 3. The maximum value of t which permits a closed rectangular grid S to be t-in-L_1 diagnosable given that less than half the processors in S are faulty, is 3.

Proof: The theorem is proved by contradiction. Assume there exist two permissible fault sets X_1 and X_2 sharing a common syndrome s, such that there are at most 3 faulty processors in $L_1(u)$ for every processor u in S. Then there exist processors x and y satisfying the conditions of Lemma 1 such that the distance $d(x,y)$ is minimum. We arrive at a contradiction by showing that $d(x,y)$ must have a value other than 2 or 3.

We observe that the status of all processors in $L_1(x)$ remain unchanged with respect to both X_1 and X_2 since x is fault-free in the presence of either fault set. We also note that there cannot be a path of fault-free processors between x and y with respect to either fault set, otherwise X_1 and X_2 cannot share a common syndrome.

case 1: $d(x,y) = 2$.

We divide this case into two subcases, all others being symmetric to one of these subcases.

case 1.1: (See Fig. 1.(a))

Without loss of generality, we assume y to be fault-free with respect to X_1 and faulty with respect to X_2. From the arguments given earlier, both w_1 and w_2 are faulty with respect to both X_1 and X_2. Since y is fault-free and $L_1(y)$ contains at least one other fault-free processor in the presence of X_1, these processors must be faulty in the presence of X_2. Thus $L_1(y)$ contains at least 4 faulty processors with respect to X_2, contradicting the assumption that there are at most 3 faulty processors in $L_1(u)$ for every processor u in S.

case 1.2: (See Fig. 1.(b))

The processor z must be faulty with respect to both fault sets, otherwise there is a path of fault-free processors between x and y. If either w_1 or w_2 is fault-free with respect to X_1 or X_2, then case 1.1. occurs and we arrive at a contradiction. If both w_1 and w_2 are faulty with respect to X_1 and X_2, then since y is faulty in the presence of one of these fault sets, $L_1(z)$ contains more than 3 faulty processors with respect to X_1 or X_2; a contradiction.

case 2: $d(x,y) = 3$.

We consider two subcases, all others being symmetric to one of these subcases.

case 2.1: (See Fig. 1.(c))

The processors w_1, w_2 and w_3 are all faulty with respect to both X_1 and X_2, otherwise the minimality of $d(x,y)$ is contradicted. Since y is faulty with respect to either X_1 or X_2, $L_1(w_2)$ contains more than 3 faulty processors in the presence of one of these fault sets, a contradiction.

case 2.2: (See Fig. 1.(d))

If any of the processors w_1, w_2 and w_3 is fault-free with respect to either X_1 or X_2, then one of the cases discussed earlier occurs. Hence we assume w_1, w_2 and w_3 to be faulty with respect to X_1 and X_2. Since y is faulty in the presence of one of these fault sets, $L_1(w_2)$ contains at least 4 faulty processors with respect to either X_1 or X_2, a contradiction.

It follows that the system S is 3-in-L_1 diagnosable given that less than half the processors in S are faulty. Fig. 2 shows a closed rectangular grid in which two fault sets having at most 4 faulty processors in $L_1(u)$ for every processor u, share a common syndrome. This proves that the maximum value for t is 3. //

The proof of the following result is similar to the proof of Theorem 3.

Theorem 4. The maximum values of t which permit a closed hexagonal grid and a closed octagonal grid to be t-in-L_1 diagnosable given that less than half the processors in S are faulty, are 4 and 5 respectively.

The following Lemma provides a method for developing diagnosis algorithms for regular interconnected structures with local constraints.

Lemma 2. Given a system S and a syndrome, let u be a processor in S such that $L_1(u)$ has at most k + 1 faulty processors, $|L_1(u)| = 2k + 1$ and at least 2 processors in $N(u)$ have been correctly identified. Then u can be correctly identified.

Proof: Let F denote the set of processors in $N(u)$ which have been correctly identified. If any member of F is fault-free then the status of u can be determined correctly. We now consider the case when all processors in F have been identified to be faulty. Let F_0 and F_1 represent the set of processors in $N(u)$ - F which have 0-links and 1-links repectively with u. We observe that $|F_0 \cup F_1| = 2k - |F|$. If

$$|F| + |F_1| > k + 1 \qquad \text{...(i)}$$

then u can be declared faulty; u can be declared fault-free if

$$|F| + |F_0| + 1 > K + 1 \qquad \text{...(ii)}$$

Both (i) and (ii) cannot be satisfied simultaneously; for otherwise the assumption that there are atmost k + 1 faults in $L_1(u)$ is violated or the processors in F have been identified incorrectly. At least one of the conditions (i) and (ii) is satisfied if we ensure that

$$|F| + \max\{F_0 + 1, F_1\} > k + 1 \qquad \text{...(iii)}$$

Since $|F| \geq 2$ and $\max\{F_0 + 1, F_1\} \geq \lfloor (2k - |F|)/2 \rfloor + 1$, condition (iii) is satisfied for all permissible values of F_0 and F_1. //

Theorem 5. Let S be a rectangular grid in which there are at most 3 faulty processors in $L_1(u)$ for every processor u in S. Given a syndrome and a fault-free processor v in S, all processors in S can be correctly identified.

Proof: Given a fault-free processor v, all processors in $L_1(v)$ can be correctly identified. Consider the processors x_1, x_2, x_3, and x_4 in Fig. 3.(a). Since each of these is adjacent to 2 processors in $L_1(u)$ and there are at most 3 faulty processors in $L_1(u)$ for every processor u in S, the processors x_1, x_2, x_3, and x_4 can be correctly identified. Now assuming all processors within a rectangle with sides containing at least 3 processors have been correctly identified, we show that the processors in the enclosing rectangle can be correctly identified.

Let R be the rectangle in which all processors have been correctly identified (Fig. 3.(b)). Let x_1 be a processor in R which is adjacent to the unidentified processor x_2 on A such that x_1 does not lie on a corner of R. If x_1 or one of its adjacent processors which lie on the same column as x_1 is fault-free then one of the processors on A can be identified correctly. If all three are faulty then x_2 can be identified as fault-free since there are already three faulty processors in $L_1(x_1)$. Once one processor on A has been correctly identified, the status of an adjacent processor on A can be determined since two of its neighboring processors have been correctly identified: one of the identified neighboring processors is in R and the other lies on A. Applying this technique repeatedly, all processors lying on A can be identified. The processors lying on B, C, and D

can be identified similarly. The status of the processors lying on the corners of the enclosing rectangle can be determined correctly since each of these is adjacent to two processors which have been identified correctly.

By induction, all processors in the system can be correctly identified. //

The procedure used in the proof of the above theorem can be used to develop a diagnosis algorithm for a 3-in-L_1 closed rectangular grid where less than half the processors are faulty. To begin with, we assume a processor v to be fault-free and apply the procedure outlined above with the following modification. Each time a processor is labeled faulty or fault-free, it selects one of its adjacent processors which has been labeled earlier to be its parent. When all processors have been labeled, a backward phase is initiated in which a leaf processor sends a value -1(1) to its parent if it is labeled faulty(fault-free). An internal processor waits until it has the values from all its children, adds these values, and sends the total to its parent after incrementing or decrementing the total by 1 depending on its own fault-free or faulty label. If at any stage of this backward phase, a processor finds it has more than 3 faulty processors within its immediate neighborhood, a failure message is propagated up the tree. If v receives a failure message or the sum of all the values received is negative then v is identified correctly to be faulty and the same procedure is initiated at a processor w at a distance 2 from v. If both processors are found to be faulty then a fault-free processor can easily be identified within the immediate neighborhood of the processor lying between x and y. One more sweep of the basic procedure identifies all the processors. If either v or subsequently w does not receive a failure message or the sum of all the values received is non-negative then all processors have been correctly identified during the current sweep.

The algorithm can be run sequentially on a host processor where the syndrome has been collected or in a distributed manner on the rectangular grid itself. In the latter case it is assumed that each processor has a non-faulty core on which the algorithm can be run. The complexity of the sequential algorithm is $O(n+E)$ and the distributed algorithm has time complexity $O(n)$ and message complexity $O(E)$.

Similar algorithms can be developed to diagnose 4-in-L_1 hexagonal grid systems and 5-in-L_1 octagonal grid systems where either one fault-free processor is given or less than half the processors in the system are faulty.

REFERENCES.

[1] Preparata, F.P., Metze, G., Chien, R.T.: On the Connection Assignment Problem of Diagnosable Systems. IEEE Trans. Electr. Compt., vol EC-16(1967), pp. 848-854.

[2] Dahbura, A.T., Masson, G.M.: An $O(n^{2.5})$ Fault Identification Algorithm for Diagnosable Systems. IEEE Trans. Compt., vol. c-33 (1984), pp. 486-492 .

[3] Hakimi, S.L., Amin, A.: Characterization of the Connection Assignment of Diagnosable Systems. IEEE Trans. Comp., vol c-23 (1974), pp. 86-88.

[4] Sullivan, G.: The Complexity of System-Level Fault Diagnosis and Diagnosability. Ph.D. Dissertation, Yale University (1986).

[5] Yang, C., Masson, G.M.: A Fault Identification Algorithm for t_i-Diagnosable Systems. IEEE Trans. Comp., vol c-35 (1986), pp. 503-510.

[6] Dahbura, A., Sabnani, K., King, L.: The Comparison Approach to Multiprocessor Fault Diagnosis. IEEE Trans. Comp., vol c-36 (1987) pp. 373-378.

[7] Blough, D.M., Sullivan, G.F., Masson, G.M.: Almost Certain Diagnosis for Intermittently Faulty Systems. Dig. 18th Int. Symp. on Fault-Tolerant Comp., IEEE Comp. Soc. Press (1988), pp. 260-265.

[8] Chwa, K.Y., Hakimi, S.L.: Schemes for Fault-Tolerant Computing: A Comparison of Modularly Redundant and t-Diagnosable Systems. Inform. Contr., vol. 49 (1976), pp. 585-393 .

[9] Das, A., Lakshmanan, K.B., Thulasiraman, K., Agarwal, V.K.: Distributed Diagnosis for Regular Interconnected Systems. Proc. Second International Conf. on Supercomputing, vol 3 (1987), pp. 277-283 .

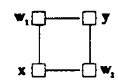

Fig. 1(a)

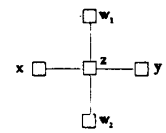

Fig. 1(b)

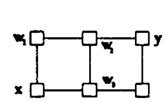

Fig. 1(c)

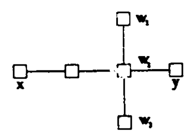

Fig. 1(d)

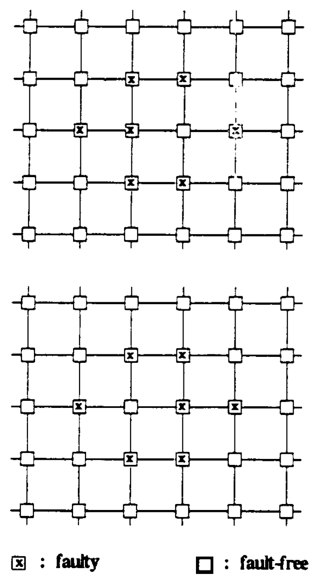

Fig. 2

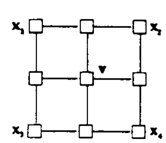

Fig. 3(a)

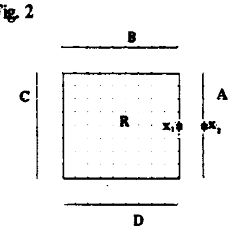

Fig. 3(b)