OPTIMAL ADAPTIVE PARALLEL DIAGNOSIS FOR ARRAYS

Toshinori YAMADA, Kumiko NOMURA, and Shuichi UENO

Department of Communications and Integrated Systems, Graduate School of Science and Engineering
Tokyo Institute of Technology, Tokyo 152–8552, Japan
E–mail: {yamada, nomura, ueno}@lab.ss.titech.ac.jp

ABSTRACT

We consider adaptive fault diagnosis for array multiprocessor systems. We show that three testing rounds are necessary and sufficient for adaptive parallel diagnosis of an N-processor system modeled by a d-dimensional square mesh [torus] if $N \geq (d+1)^{d/2}$ [$N \geq (2d+2)^d$].

1. INTRODUCTION

The system diagnosis has been extensively studied in the literature in connection with fault-tolerant multiprocessor systems. An original graph-theoretical model for system diagnosis was introduced in a classic paper by Preparata, Metze, and Chien [20]. In this model, each processor is either faulty or fault-free. The fault-status of a processor does not change during the diagnosis. The processors can test each other only along communication links. A testing processor evaluates a tested processor as either faulty or fault-free. The evaluation is accurate if the testing processor is fault-free, while the evaluation is unreliable if the testing processor is faulty. The system diagnosis is to identify all faulty processors based on test results.

A system is t-diagnosable if all faulty processors can always be identified provided that the number of faulty processors does not exceed t. It is well-known that a system with N processors is t-diagnosable only if $t < N/2$ and each processor is connected with at least t distinct other processors by communication links [20]. A complete characterization of t-diagnosable system was shown by Hakimi and Amin [10]. The original model is nonadaptive in the sense that all tests must be determined in advance. It can be shown that each processor must be tested by at least t distinct other processors in nonadaptive diagnosis if as many as t processors may be faulty. It follows that at least tN tests are necessary for nonadaptive diagnosis of an N-processor system with at most t faulty processors.

In adaptive diagnosis introduced by Nakajima [16], tests can be determined dynamically depending on previous test results. The adaptive diagnosis has been extensively studied in the literature [1–7, 9, 11–19, 21]. Among others, Blecher [7] and Wu [21] showed that $N+t-1$ tests are sufficient for adaptive diagnosis of an N-processor system with at most t faulty processors if the system is modeled by a complete graph and $t < N/2$. Moreover, Blecher [7] showed that $N+t-1$ is also the lower bound for the number of tests in the worst case. The adaptive diagnosis of practical systems modeled by sparse graphs has been also considered [4–6, 9, 13–15, 17–19].

The adaptive parallel diagnosis has been considered as well [1–3, 6, 12, 14, 17, 19]. In adaptive parallel diagnosis, each processor may participate in at most one test, either as a testing or tested processor, in each testing round. Beigrl, Hurwood, and Kahale [1] showed that for adaptive parallel diagnosis of an N-processor system modeled by a complete graph with at most t faulty processors, 3 testing rounds are necessary and sufficient if $2 \leq t \leq \sqrt{N/3}$, 4 testing rounds are necessary and sufficient if $2\sqrt{2N} \leq t \leq 0.03N$, 5 testing rounds are necessary if $t \geq 0.49N$, and 10 testing rounds are sufficient if $t < N/2$. Since at least $N+t-1$ tests are necessary for adaptive parallel diagnosis of an N-processor system with at most t faulty processors and there are at most $N/2$ tests in each testing round, $\lceil (N+t-1)/(N/2) \rceil$, which is 3 if $t \geq 2$, is a general lower bound for the number of testing rounds [2]. Okashita, Araki, and Shibata [19] considers adaptive parallel diagnosis of systems modeled by butterfly networks using small number of testing rounds. Björklund [6] showed that 4 testing rounds are sufficient for adaptive parallel diagnosis of an N-processor system modeled by a hypercube with at most $\log N$ faulty processors. Nomura, Yamada, and Ueno [17] showed that for adaptive parallel diagnosis of an N-processor system modeled by a hypercube, 3 testing rounds are necessary and sufficient if the number of faulty processors is at most $\log N - \lceil \log(\log N - \lceil \log\log N \rceil + 4) \rceil + 2$. They also showed that 3 testing rounds are necessary and sufficient for adaptive parallel diagnosis of a system modeled by cube-connected cycles of dimension greater than three [17].

This paper shows that 3 testing rounds are necessary and sufficient for adaptive parallel diagnosis of an N-processor system modeled by a d-dimensional square mesh [torus] if $N \geq (d+1)^{d/2}$ [$N \geq (2d+2)^d$].

2. PRELIMINARIES

A multiprocessor system is modeled by a graph in which the vertices represent processors and edges represent communication links. Each vertex is either faulty or fault-free. A pair of adjacent vertices can test each other. A test performed by u on v is represented by an ordered pair $\langle u, v \rangle$. The outcome of a test $\langle u, v \rangle$ is $1(0)$ if u evaluates v as faulty(fault-free). The outcome is accurate if u is fault-free, while the outcome is unreliable if u is faulty. A graph is t-diagnosable if all faulty vertices can always be identified from test results provided that the number of faulty vertices is not more than t. If an N-vertex graph G is t-diagnosable then $t < N/2$ and the minimum degree of a vertex is at least t [20].

We denote the vertex set and edge set of a graph G by $V(G)$ and $E(G)$, respectively. For $S \subseteq V(G)$, $G - S$ is the graph obtained from G by deleting the vertices in S. For a positive integer k, a graph G is said to be k-connected if $G - S$ is connected for any $S \subseteq V(G)$ with $|S| \leq k - 1$. A graph is said to be k'-connected for any integer $k' \leq 0$ for convenience. We denote a cycle and path with N vertices by C_N and P_N, respectively. C_N is called an even cycle if N is even, and odd cycle otherwise. The product of graphs $G_1, G_2, \ldots, G_k$ is a graph $G = G_1 \times G_2 \times \cdots \times G_k$ with vertex set $V(G) = V(G_1) \times V(G_2) \times \cdots V(G_k)$, in which $(u_1, u_2, \ldots, u_k)$ is adjacent to $(v_1, v_2, \ldots, v_k)$ if and only if there exists an integer j such that $(u_j, v_j) \in E(G_j)$ and $u_i = v_i$ for every $i \neq j$.

The d-dimensional $m_1 \times m_2 \times \cdots \times m_d$ mesh, denoted by $M(m_1, m_2, \ldots, m_d)$, is the graph deifned as

$$M(m_1, m_2, \ldots, m_d) = P_{m_1} \times P_{m_2} \times \cdots \times P_{m_d}.$$

It follows that

$$M(m_1, m_2 \ldots, m_d) = M(m_1, \ldots, m_p) \times M(m_{p+1}, \ldots, m_d)$$

for any positive integer $p < d$. $M(m_1, m_2, \ldots, m_d)$ has $m_1 m_2 \cdots m_d$ vertices, and the minimum degree of a vertex is d. $M(m_1, m_2, \ldots, m_d)$ is called the d-dimensional m-sided mesh and denoted by $M_d(m)$ if $m_1 = m_2 = \cdots = m_d = m$.

The n-dimensional cube $Q(n)$ is defined as the n-dimensional 2-sided mesh $M_n(2)$. $Q(n)$ has 2^n vertices, and the degree of a vertex is n. $Q(n)$ can be represented as $Q(p) \times Q(q)$ for any positive integers p and q with $p + q = n$.

The d-dimensional $m_1 \times m_2 \times \cdots \times m_d$ torus, denoted by $D(m_1, m_2, \ldots, m_d)$, is the graph deifned as

$$D(m_1, m_2, \ldots, m_d) = C_{m_1} \times C_{m_2} \times \cdots \times C_{m_d}.$$

It follows that

$$D(m_1, m_2, \ldots, m_d) = D(m_1, \ldots, m_p) \times D(m_{p+1}, \ldots, m_d)$$

for any positive integer $p < d$. $D(m_1, m_2, \ldots, m_d)$ has $m_1 m_2 \cdots m_d$ vertices, and the degree of a vertex is $2d$. $D(m_1, m_2, \ldots, m_d)$ is called the d-dimensional m-sided torus and denoted by $D_d(m)$ if $m_1 = m_2 = \cdots = m_d = m$.

Algorithm 1
Step 1
Perform in 2 testing rounds all tests along all edges of C_N in the clockwise direction.
Step 2
If there is a sequence $a \xrightarrow{1} b \xrightarrow{1} c \xrightarrow{1} d \xrightarrow{0} e$ in test outcomes of Step 1
then perform one additional test (e, d);
If there is a sequence $a \xrightarrow{1} b \xrightarrow{1} c \xrightarrow{0} d$ and there are only two 1's in test outcomes of Step 1
then perform one additional test (d, c);
If there is a sequence $a \xrightarrow{1} b \xrightarrow{0} c \xrightarrow{1} d \xrightarrow{0} e$ and there are only two 1's in test outcomes of Step 1
then perform one additional test (e, d);
If there is a sequence $a \xrightarrow{1} b \xrightarrow{0} c \xrightarrow{0} d$ and there is only one 1 in test outcomes of Step 1
then perform one additional test (d, c).

Figure 1: Algorithm 1

3. DIAGNOSIS FOR GRAPH PRODUCTS

In this section, we show an optimal adaptive parallel diagnosis algorithm for products of t-connected graphs and even cycles.

We need a couple of known results. The following theorem is an easy corollary of a classic theorem by Dirac [8].

Theorem I *Let t be a positive integer. If G is a t-connected graph with at least $2t$ vertices and S is a set of vertices of G with $|S| \leq t$ then every vertex v of S has a distinct vertex in $V(G) - S$ adjacent to v.* □

The following theorem is proved in [17].

Theorem II [17] *Algorithm 1 shown in Figure 1 adaptively diagnoses an even cycle C_N in 3 testing rounds if the number of faults is not more than 2 and $N \geq 6$.* □

Now, we are ready to prove our main theorem.

Theorem 1 *Let t be a positive integer, and let G be a t-connected graph with at least $2t$ vertices. For any even number $n \geq \max\{t + 3, 6\}$, Algorithm 2 shown in Figure 2 adaptively diagnoses $G \times C_n$ in 3 testing rounds if the number of faults is not more than $t + 2$.*

Proof : Algorithm 2 works in two steps. In the first step, we perform in two testing rounds all tests along all copies of C_n in the clockwise direction. A copy of C_n is said to

Algorithm 2
Step 1
Perform in 2 testing rounds all tests along all edges in all copies of C_n in the clockwise direction. Let $\mathcal{F}$ be the set of all faulty copies of C_n.
Step 2
If $t+1 \leq |\mathcal{F}| \leq t+2$ **then**
perform tests in one more testing round according to Step 2 of Algorithm 1, and identify the faults;
If $1 \leq |\mathcal{F}| \leq t$ **then**
diagnose all vertices in all faulty copies of C_n by corresponding vertices in distinct fault-free copies of C_n in one more testing round;
If $|\mathcal{F}| = 0$ **then**
identify the faults as empty.

Figure 2: Algorithm 2

be fault-free if it has no faulty vertex, and faulty otherwise. Since $n \geq t+3$, a copy of C_n is faulty if and only if it has a test outcome of 1.

Let $\mathcal{F}$ be the set of all faulty copies of C_n. The second step of our algorithm is distinguished in four cases depending on $|\mathcal{F}|$.

If $t+1 \leq |\mathcal{F}| \leq t+2$ then each faulty copy of C_n has at most two faulty vertices, which can be identified in one more testing round as can be seen by Theorem II since $n \geq 6$.

If $1 \leq |\mathcal{F}| \leq t$ then every faulty copy F_n of C_n has a distinct fault-free copy H_n of C_n in which each vertex v_H of H_n is adjacent to the corresponding vertex v_F of F_n by Theorem I since G is t-connected. By performing the tests $\langle v_H, v_F \rangle$ for all faulty copies of C_n in one testing round, we can identify all the faults.

If $|\mathcal{F}| = 0$ then we know from the test results in the first step that there is no fault. □

4. COROLLARIES

We show in this section some corollaries of Theorem 1 for tori, meshes, and hypercubes.

4.1. Tori

Theorem 2 *Let d be a positive integer, and m be an even number at least $\max\{2d+2, 6\}$. Then, $D_d(m)$ can be adaptively diagnosed in 3 testing rounds if the number of faults is not more than $2d$.*

Proof : If $d = 1$ then $D_d(m) = C_m$, and so the theorem holds by Theorem II since $m \geq 6$.

If $d \geq 2$ then $D_d(m) = D_{d-1}(m) \times C_m$. Since $D_{d-1}(m)$ is a $2(d-1)$-connected graph with $m^{d-1} \geq 4(d-1)$ vertices and m is an even number at least $\max\{2d+2, 6\}$, the theorem holds for $d \geq 2$ by Theorem 1. □

4.2. Meshes

Theorem 3 *Let d be an integer at least 2, and m be an even number at least $\max\{\sqrt{d+1}, 4\}$. Then, $M_d(m)$ can be adaptively diagnosed in 3 testing rounds if the number of faults is not more than d.*

Proof : The proof is similar to that of Theorem 2. Since m is an even number at least 4, $M_2(m)$ has a Hamilton cycle, and m^2 is an even number at least 16. Therefore, the theorem holds for $d = 2$ by Theorem II.

If $d \geq 3$ then $M_d(m) = M_{d-2}(m) \times M_2(m)$. Since $M_{d-2}(m)$ is a $(d-2)$-connected graph with $m^{d-2} \geq 2(d-2)$ vertices, $M_2(m)$ has a Hamilton cycle, and $|V(M_2(m))| = m^2$ is an even number at least $\max\{d+1, 16\}$, the theorem holds for $d \geq 3$ by Theorem 1. □

4.3. Hypercubes

The following theorem proved in [17] can be derived from Theorem 1.

Theorem III [17] *$Q(n)$ can be adaptively diagnosed in 3 testing rounds if the number of faults is not more than $n - \lceil \log(n - \lceil \log n \rceil + 4) \rceil + 2$ and $n \geq 4$.*

Proof : Let $t = n - \lceil \log(n - \lceil \log n \rceil + 4) \rceil$. $Q(n)$ is represented as $Q(n-t) \times Q(t)$. Notice that $t \geq 1$ since $n \geq 4$. Since

$$\begin{aligned} |V(Q(n-t))| &= 2^{n-t} \\ &= 2^{\lceil \log(n - \lceil \log n \rceil + 4) \rceil} \\ &\geq n - \lceil \log n \rceil + 4 \\ &\geq \max\{t+3, 6\}, \end{aligned}$$

$Q(n-t)$ has a Hamilton cycle, and $Q(t)$ is t-connected, we have the theorem by Theorem 1. □

5. CONCLUDING REMARKS

1. $Q(3)$ can be adaptively diagnosed in 3 testing rounds if the number of faults is at most 3, as mentioned in [14]. Notice that $Q(2)$ is just C_4. We can prove that $Q(n)$ can be adaptively diagnosed in 3 testing rounds if the number of faults is not more than $n - \lceil \log(n - \lceil \log n \rceil + 3) \rceil + 2$ and $n \geq 3$. The proof is similar to that of Theorem III but more complicated. It is still open whether 3 testing rounds are sufficient to adaptively diagnose $Q(n)$ with at most t faulty vertices even if $n - \lceil \log(n - \lceil \log n \rceil + 3) \rceil + 3 \leq t \leq n$.

2. We can prove the following more general results.

Theorem 4 *Let d be an integer at least 2, and let $m_1, \ldots, m_d$ be integers at least 3 such that $\prod_{i \in I} m_i$ is an even number at least $\max\{2(d - |I|) + 3, 6\}$ and $\prod_{j \in \{1,\ldots,d\}-I} m_j \neq 3$ for some $I \subseteq \{1, \ldots, d\}$. Then, $D(m_1, \ldots, m_d)$ can be adaptively diagnosed in 3 testing rounds if the number of faults is not more than $2(d - |I| + 1)$.* □

Theorem 5 *Let d be an integer at least 3, and let $m_1, \ldots, m_d$ be integers at least 2 such that $\prod_{i \in I} m_i$ is an even number at least $\max\{d - |I| + 3, 6\}$. Then, $M(m_1, \ldots, m_d)$ can be adaptively diagnosed in 3 testing rounds if the number of faults is not more than $d - |I| + 2$.* □

The proofs are similar to those of Theorems 2 and 3, but more complicated. Theorems 4 and 5 assume that one of $m_1, \ldots, m_d$ is even. It is open whether 3 testing rounds are sufficient to adaptively diagnose these graphs even if all of $m_1, \ldots, m_d$ are odd.

6. REFERENCES

[1] R. Beigel, W. Hurwood, and N. Kahale, "Fault diagnosis in a flash," in *Proc. 36th FOCS*, pp. 571–580, 1995.

[2] R. Beigel, S. Kosaraju, and G. Sullivan, "Locating faults in a constant number of testing rounds," in *Proc. 1st SPAA*, pp. 189–198, 1989.

[3] R. Beigel, G. Margulis, and D. Spielman, "Fault diagnosis in a small constant number of parallel rounds," in *Proc. 5th SPAA*, pp. 21–29, 1993.

[4] R. Bianchini, Jr. and R. Buskens, "An adaptive distributed system-level diagnosis algorithm and its implementation," in *Proc. 21st International Symposium on Fault Tolerant Computing*, pp. 222–229, 1991.

[5] R. Bianchini, Jr., K. Goodwin, and D. Nydick, "Practical application and implimentation of distributed system-level diagnosis theory," in *Proc. 20th International Symposium on Fault Tolerant Computing*, pp. 332–339, 1984.

[6] A. Björklund, "Optimal adaptive fault diagnosis of hypercubes," *Lecture Notes in Computer Science*, vol. 1851, pp. 527–534, 2000.

[7] P. Blecher, "On a logical problem," *Discrete Mathematics*, vol. 43, pp. 107–110, 1983.

[8] G. Dirac, "Généralizations du théorème de menger," *C.R. Acad. Sci. Paris*, vol. 250, pp. 4252–4253, 1960.

[9] C. Feng, L. Bhuyan, and F. Lombardi, "Adaptive system-level diagnosis for hypercube multiprocessors," *IEEE Transactions on Computers*, vol. 45, pp. 1157–1170, 1996.

[10] S. Hakimi and A. Amin, "Characterization of connection assignment of diagnosable systems," *IEEE Transactions on Computers*, vol. 23, pp. 86–88, 1974.

[11] S. Hakimi and K. Nakajima, "On adaptive system diagnosis," *IEEE Transactions on Computers*, vol. 33, pp. 234–240, 1984.

[12] S. Hakimi, M. Otsuka, and E. Schmeichel, "A parallel fault identification algorithm," *Journal of Algorithms*, vol. 11, pp. 231–241, 1990.

[13] S. Hakimi and E. Schmeichel, "An adaptive algorithm for system level diagnosis," *Journal of Algorithms*, vol. 5, pp. 524–530, 1984.

[14] E. Kranakis and A. Pelc, "Better adaptive diagnosis of hypercubes," *IEEE Transactions on Computers*, vol. 49, no. 10, pp. 1013–1020, 2000.

[15] E. Kranakis, A. Pelc, and A. Spatharis, "Optimal adaptive fault diagnosis for simple multiprocessor systems," *Networks*, vol. 34, pp. 206–214, 1999.

[16] K. Nakajima, "A new approach to system diagnosis," in *Proc. 19th Allerton Conf. Commun. Contr. and Computing*, pp. 697–706, 1981.

[17] K. Nomura, T. Yamada, and S. Ueno, "On adaptive fault diagnosis for multiprocessor systems," *Lecture Notes in Computer Science*, vol. 2223, pp. 86–98, 2001.

[18] A. Okashita, T. Araki, and Y. Shibata, "Adaptive diagnosis of butterflies with optimal number of tests," *Technical Report of IEICE*, vol. 101, no. 488, COMP2001-70, pp. 55–62, 2001. (in Japanese).

[19] A. Okashita, T. Araki, and Y. Shibata, "Adaptive diagnosis of butterfly networks," *IPSJ SIG Notes*, vol. 2001, no. 79, 2001-AL-79, pp. 29–36, 2001. (in Japanese).

[20] F. Preparata, G. Metze, and R. Chien, "On the connection assignment problem of diagnosable systems," *IEEE Transactions on Electronic Computers*, vol. 16, pp. 848–854, 1967.

[21] P. Wu, "Partial solution to problem 81-6," *Journal of Algorithms*, vol. 3, pp. 379–380, 1982.