On Adaptive Fault Diagnosis for Multiprocessor Systems

Kumiko Nomura, Toshinori Yamada, and Shuichi Ueno

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

Abstract. We first consider adaptive serial diagnosis for multiprocessor systems. We present an adaptive diagnosis algorithm using $N+t-1$ tests, which is the smallest possible number, for an N-processor system modeled by a $(2t-1)$-connected graph with at most t faulty processors. We also present an adaptive diagnosis algorithm using minimum number of tests for a system modeled by cube-connected cycles. We consider adaptive parallel diagnosis as well. We show that for adaptive parallel diagnosis of an N-processor system modeled by a hypercube, three testing rounds are necessary and sufficient if the number of faulty processors is at most $\log N - \lceil \log(\log N - \lceil \log\log N \rceil + 4) \rceil + 2$. We also show that three testing rounds are necessary and sufficient for adaptive parallel diagnosis of a system modeled by cube-connected cycles of dimension greater than three.

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 [16]. 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 [16]. A complete characterization of t-diagnosable system was shown by Hakimi and Amin [9]. 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

P. Eades and T. Takaoka (Eds.): ISAAC 2001, LNCS 2223, pp. 86–98, 2001.

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 [15], tests can be determined dynamically depending on previous test results. Blecher [7] and Wu [17] 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 some practical systems modeled by sparse graphs has been considered in the literature [4,5,6,8,12,13,14]. Among others, Kranakis, Pelc, and Spatharis [14] showed adaptive diagnosis algorithms using minimum number of tests in the worst case for systems modeled by trees, cycles, and tori. Björklund [6] showed an adaptive diagnosis algorithm for an N-processor system modeled by a hypercube with at most t faulty processors. The algorithm uses $N+t-1$ tests if $t = \log N$, and $N + t$ tests if $t < \log N$.

This paper shows an adaptive diagnosis algorithm using minimum number of tests for systems modeled by cube-connected cycles. We also show an adaptive diagnosis algorithm using $N + t - 1$ tests for an N-processor system modeled by a $(2t-1)$-connected graph with at most t faulty processors. This is an extension of a previous result on systems modeled by complete graphs in the sense that an N-vertex complete graph K_N is $(2t-1)$-connected if $t < N/2$. Notice that our algorithm uses $N+t-1$ tests for an N-processor system modeled by a hypercube with at most t faulty processors if $t \leq (\log N+1)/2$, since an N-vertex hypercube is $\log N$-connected.

The adaptive parallel diagnosis has been considered as well in the literature [1,2,3,6,11,13]. 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 K_N with at most t faulty processors, 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]. 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. It is still open whether 3 testing rounds are sufficient for such systems, as mentioned in [6].

We partially answer the question above by showing 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 $\log N - \lceil \log(\log N - \lceil \log\log N \rceil + 4) \rceil + 2$. We also show that 3 testing rounds are necessary and sufficient for adaptive parallel diagnosis of systems modeled by cube-connected cycles of dimension greater than 3.

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 [16].

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, path and complete graph with N vertices by C_N, P_N, and K_N, respectively. C_N is called an even cycle if N is even, and odd cycle otherwise. The product of graphs G and H is a graph $G \times H$ with vertex set $V(G) \times V(H)$, in which (u, v) is adjacent to (u', v') if and only if either $u = u'$ and $(v, v') \in E(H)$ or $v = v'$ and $(u, u') \in E(G)$

An n-dimensional cube $Q(n)$ is recursively defined as follows: $Q(1) = P_2$; $Q(n) = Q(n-1) \times P_2$. It follows that $Q(n) = Q(p) \times Q(q)$ for any positive integers p and q such that $p + q = n$. $Q(n)$ has 2^n vertices, and the degree of a vertex is n.

The n-dimensional cube-connected cycles(CCC) is constructed from $Q(n)$ by replacing each vertex of $Q(n)$ with C_n in CCC. For any positive integer k, $[k]$ denotes $\{0, 1, \ldots, k-1\}$. For any positive integer n and $\boldsymbol{x} = x_{n-1}x_{n-2}\cdots x_0 \in [2]^n$ and $i \in [n]$, let $\chi_i(\boldsymbol{x}) = x_{n-1}\cdots x_{i+1}\overline{x_i}x_{i-1}\cdots x_0$, where $\overline{x_i} = 1 - x_i$, that is the complement of x_i. The n-dimensional CCC, denoted by $CCC(n)$, is the graph defined as follows:

$$V(CCC(n)) = [2]^n \times [n];$$
$$E(CCC(n)) = \{([\boldsymbol{x}, i], [\chi_i(\boldsymbol{x}), i]) : i \in [n]\} \cup \{([\boldsymbol{x}, i], [\boldsymbol{x}, j]) : j = (i \pm 1) \bmod n\}.$$

$CCC(n)$ has $n2^n$ vertices, and the degree of a vertex is 3.

3 Adaptive Diagnosis

In nonadaptive diagnosis, all tests are scheduled in advance. It is known that at least tN tests are necessary for nonadaptive diagnosis of an N-vertex graph with at most t faulty vertices [16].

In adaptive diagnosis, tests can be determined dynamically depending on previous test results. The following theorem shows a general lower bound for the number of tests necessary to adaptively diagnose a graph.

```
function Expand(G, H_0, F_0, t)
begin
    H ← H_0; F' ← F_0;
    while H ∪ F' ≠ V(G) and |F'| < t do
    begin
        Select any v ∈ V(G) − (H ∪ F') s.t. (u, v) ∈ E(G) for some u ∈ H;
        if outcome of test ⟨u, v⟩ is 0 then H ← H ∪ {v}
        else F' ← F' ∪ {v};
    end
    return(F')
end
```

Fig. 1. Function Expand

Theorem I [7] *If G is an N-vertex graph with at most t faulty vertices then $N + t - 1$ tests are necessary to adaptively diagnose G in the worst case.*

The following theorem shows upper bounds for the number of tests sufficient to adaptively diagnose hypercubes.

Theorem II [6] *$Q(n)$ is adaptively t-diagnosable using at most $N - t + 1$ tests if $t = n$, and using at most $N + t$ tests if $t < n$, where $N = 2^n$ is the number of vertices in $Q(n)$.*

3.1 $(2t - 1)$-Connected Graphs

In this section, we prove the following theorem.

Theorem 1. *Let G be an N-vertex graph and t be a positive integer. If G is $(2t - 1)$-connected and $t < N/2$ then G is adaptively t-diagnosable using at most $N + t - 1$ tests.*

Since K_N is $(N - 1)$-connected and $Q(n)$ is n-connected, we have the following corollaries:

Corollary I [7,17] *K_N is adaptively t-diagnosable using at most $N + t - 1$ tests if $t < N/2$.* □

Corollary 1. *$Q(n)$ is adaptively t-diagnosable using at most $N + t - 1$ tests if $t \le (n + 1)/2$ and $n \ge 2$, where $N = 2^n$ is the number of vertices in $Q(n)$.* □

3.1.1 Proof of Theorem 1

We need a preliminary result.

Lemma 1. *Let G be a t-connected graph, and F be a set of all faulty vertices with $|F| \le t$. If $H_0 \subseteq V(G) - F$, $H_0 \ne \emptyset$, and $F_0 \subseteq F$ then* Function Expand *shown in Fig. 1 identifies F using at most $|V(G)| - |H_0 \cup F_0|$ tests.*

Proof. We prove the lemma by a series of claims.

Claim 1. *If $H \cup F' \neq V(G)$, $H \neq \emptyset$, and $|F'| < t$ then there is a vertex $v \in V(G) - (H \cup F')$ such that $(u, v) \in E(G)$ for some vertex $u \in H$.*

Proof (of Claim 1). Since $|F'| < t$ and G is t-connected, $G - F'$ is connected. Since $V(G) - (H \cup F') \neq \emptyset$ and $H \neq \emptyset$, there is a vertex $v \in V(G) - (H \cup F')$ such that $(u, v) \in E(G)$ for some vertex $u \in H$. □

The following claim is obvious.

Claim 2. *$H \subseteq V(G) - F$, $H \neq \emptyset$, and $F' \subseteq F$.* □

Claim 3. *If $H \cup F' = V(G)$ or $|F'| = t$ then $F = F'$.*

Proof (of Claim 3). If $F \neq F'$ then we conclude by Claim 2 that $H \cup F' \neq V(G)$ and $|F'| < |F| \leq t$, which is a contradiction. Hence, $F = F'$. □

By Claims 1 and 3, Function Expand identifies F. Since each vertex of $V(G) - (H_0 \cup F_0)$ is tested at most once, Function Expand uses at most $|V(G)| - |H_0 \cup F_0|$ tests. □

Now we are ready to prove Theorem 1. Let G be a $(2t-1)$-connected graph, and F be a set of all faulty vertices with $|F| \leq t$. We prove the theorem by induction on t.

Since we can identify $F = \emptyset$ correctly with no test, the theorem holds for $t = 0$.

Let t be a positive integer. For inductive step, assume that the theorem holds for any non-negative integer $t' < t$. Select any $v \in V(G)$. Let $u_1, u_2, \ldots, u_k$ be the vertices adjacent to v. We perform a sequence of tests $\langle u_1, v\rangle, \langle u_2, v\rangle, \ldots, \langle u_k, v\rangle$, and add u_i to T_j if the outcome of test $\langle u_i, v\rangle$ is j ($j = 0, 1$) until either of the following two events occurs: (i) $|T_0| = t$; (ii) $|T_1| = |T_0| + 1$. It should be noted that $k \geq 2t-1$ because G is $(2t-1)$-connected. Thus, either of (i) and (ii) always occurs. It is easy to see the following:

Claim 4. *$T_1 \subseteq F$ if v is fault-free, and $T_0 \cup \{v\} \subseteq F$ otherwise.* □

We distinguish two cases.

(i) $|T_0| = t$: Since $|T_0 \cup \{v\}| = t + 1$ and $|F| \leq t$, v is fault-free and $T_1 \subseteq F$ by Claim 4. Hence, by Lemma 1, Expand$(G, \{v\}, T_1, t)$ identifies F. The total number of tests performed is at most $|T_0| + |T_1| + (N - |T_1| - 1) = N + t - 1$.

(ii) $|T_1| = |T_0| + 1$: Let $s = |T_1|$ and $G' = G - (T_0 \cup T_1 \cup \{v\})$. By Claim 4, there exists at least s faulty vertices in $T_0 \cup T_1 \cup \{v\}$, and so G' has at most $t - s$ faulty vertices. It should be noted that $|T_0 \cup T_1 \cup \{v\}| = 2s$. Since G is $(2t-1)$-connected and $|V(G)| = N \geq 2t + 1$, G' is $(2(t-s)-1)$-connected and $|V(G')| = N - 2s \geq 2(t-s) + 1$. Thus, by inductive hypothesis, we can identify $F \cap V(G')$ using at most $(N - 2s) + (t - s) - 1 = N + t - 3s - 1$ tests. Let $H' = V(G') - F$. We further distinguish two cases.

(ii)-(a) $H' \cap \{u_1, u_2, \ldots, u_k\} \neq \emptyset$: Let $u \in H' \cap \{u_1, u_2, \ldots, u_k\}$. If the outcome of test $\langle u, v\rangle$ is 0 then v is fault-free, and so $T_1 \subseteq F$ by Claim 4. Thus, by Lemma 1, Expand$(G, H' \cup \{v\}, (F \cap V(G')) \cup T_1, t)$ identifies F. The total number

Algorithm 1 [14]
Step 1
Perform the first series of 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).

Fig. 2. Algorithm 1

of tests performed is at most $(2s-1)+(N+t-3s-1)+1+(s-1) = N+t-2$. If the outcome of test $\langle u, v\rangle$ is 1 then v is faulty, and so $T_0 \cup \{v\} \subseteq F$ by Claim 4. Thus, by Lemma 1, Expand$(G, H', (F \cap V(G')) \cup T_0 \cup \{v\}, t)$ identifies F. The total number of tests performed is at most $(2s-1)+(N+t-3s-1)+1+s = N+t-1$.

(ii)-(b) $H' \cap \{u_1, u_2, \ldots, u_k\} = \emptyset$: Since $2t-1 \leq |T_0| + |T_1| + |F \cap V(G')| \leq t+s-1$, we have $s \geq t$. On the other hand, $|T_0| \leq t-1$, and so we have $s = |T_1| = |T_0|+1 \leq t$. Thus, we conclude that $s = t$. Since $|T_0 \cup \{v\}| = |T_1| = t$, we have by Claim 4 that $F = T_0 \cup \{v\}$ or $F = T_1$. Notice that $H' = V(G')$ and $F \cap V(G') = \emptyset$. Thus, $T_0 \cup T_1 = \{u_1, u_2, \ldots, u_k\}$. Since G is $(2t-1)$-connected and $|T_0| = t-1 \leq 2(t-1)$, $G - T_0$ is connected, and so there exists some vertex $w \in T_1$ such that $(x, w) \in E(G)$ for some $x \in V(G') = H'$. If the outcome of test $\langle x, w\rangle$ is 0 then w is fault-free, and so we conclude that $F = T_0 \cup \{v\}$. If the outcome of test $\langle x, w\rangle$ is 1 then w is faulty, and so we conclude that $F = T_1$. Hence, we can identify F using at most $(2t-1)+(N-2t-1)+1 \leq N+t-1$ tests.

3.2 Cycles

We will use the following results on cycles proved in [14].

Theorem III [14] *Algorithm* 1 *shown in Fig.* 2 *adaptively diagnoses* C_N *using at most* $N+1$ *test if the number of faults is at most* 2 *and* $N \geq 5$.

3.3 CCC's

Theorem 2. *$CCC(n)$ is adaptively 3-diagnosable using at most $N+2$ tests if $n \geq 4$, where $N = n2^n$ is the number of vertices in $CCC(n)$.*

Proof. Suppose $n \geq 4$ and $F \subseteq V(CCC(n))$ is a set of all faulty vertices with $|F| \leq 3$.

Let $p = \lfloor n/2 \rfloor$ and $q = n - p$ $(= \lceil n/2 \rceil)$. For any $k \in [2^{n-1}]$, set $m_k = 4p$ if $k < 2^{n-2}$, and $m_k = 4q$ otherwise. Notice that $m_k \geq 8$ since $n \geq 4$. For any $k \in [2^{n-1}]$ and any $i \in [m_k]$, define $v_{k,i}$ as follows: If $k < 2^{n-2}$ then

$$v_{k,i} = \begin{cases} [\boldsymbol{b}_1 \cdot 0 \cdot \boldsymbol{b}_0 \cdot 0, i] & \text{if } i < p, \\ [\boldsymbol{b}_1 \cdot 1 \cdot \boldsymbol{b}_0 \cdot 0, 2p - 1 - i] & \text{if } p \leq i < 2p, \\ [\boldsymbol{b}_1 \cdot 1 \cdot \boldsymbol{b}_0 \cdot 1, i - 2p] & \text{if } 2p \leq i < 3p, \\ [\boldsymbol{b}_1 \cdot 0 \cdot \boldsymbol{b}_0 \cdot 1, 4p - 1 - i] & \text{if } 3p \leq i, \end{cases}$$

where $\boldsymbol{b}_1 \in [2]^q$ and $\boldsymbol{b}_0 \in [2]^{p-2}$ are the q most and $p-2$ least significant bits of the $(n-2)$-bit binary representation of k, respectively, and $\boldsymbol{a} \cdot \boldsymbol{b}$ denotes the concatenation of $\boldsymbol{a}$ and $\boldsymbol{b}$; If $k \geq 2^{n-2}$ then

$$v_{k,i} = \begin{cases} [0 \cdot \boldsymbol{b}'_1 \cdot 0 \cdot \boldsymbol{b}'_0, i + p] & \text{if } i < q, \\ [1 \cdot \boldsymbol{b}'_1 \cdot 0 \cdot \boldsymbol{b}'_0, n - 1 + q - i] & \text{if } q \leq i < 2q, \\ [1 \cdot \boldsymbol{b}'_1 \cdot 1 \cdot \boldsymbol{b}'_0, i + p - 2q] & \text{if } 2q \leq i < 3q, \\ [0 \cdot \boldsymbol{b}'_1 \cdot 1 \cdot \boldsymbol{b}'_0, n - 1 + 3q - i] & \text{if } 3q \leq i, \end{cases}$$

where $\boldsymbol{b}'_1 \in [2]^{q-2}$ and $\boldsymbol{b}'_0 \in [2]^p$ are the $q-2$ most and p least significant bits of the $(n-2)$-bit binary representation of $k - 2^{n-2}$, respectively. Define that if $k < 2^{n-2}$ then

$$\overline{v_{k,i}} = \begin{cases} [\boldsymbol{x}, n-1] & \text{if } j = 0, \\ [\boldsymbol{x}, p] & \text{if } j = p - 1, \\ [\chi_j(\boldsymbol{x}), j] & \text{otherwise,} \end{cases}$$

and if $k \geq 2^{n-2}$ then

$$\overline{v_{k,i}} = \begin{cases} [\boldsymbol{x}, p-1] & \text{if } j = p, \\ [\boldsymbol{x}, 0] & \text{if } j = n - 1, \\ [\chi_j(\boldsymbol{x}), j] & \text{otherwise,} \end{cases}$$

where $v_{k,i} = [\boldsymbol{x}, j]$. For any $k \in [2^{n-1}]$, let $V_k = \{v_{k,i} : i \in [m_k]\}$. It is easy to see the following claims:

Claim 5. *$(V_0, \ldots, V_{2^{n-1}-1})$ is a partition of $V(CCC(n))$.* □

Claim 6. *$(v_{k,i}, \overline{v_{k,i}}) \in E(CCC(n))$ and $\overline{v_{k,i}} \in V(CCC(n)) - V_k$ for any $k \in [2^{n-1}]$.* □

Claim 7. *The subgraph of $CCC(n)$ induced by V_k is isomorphic to a cycle C_{m_k} for any $k \in [2^{n-1}]$. In particular, $(v_{k,i}, v_{k,(i \pm 1) \bmod m_k}) \in E(CCC(n))$.* □

Let $E_k = \{(v_{k,i}, v_{k,(i+1) \bmod m_k}) : i \in [m_k]\}$ for any $k \in [2^{n-1}]$ and any $i \in [m_k]$. For each $k \in [2^{n-1}]$, perform test $\langle v_{k,i}, v_{k,(i+1) \bmod m_k} \rangle$ in order of $i = 0, 1, \ldots, m_k - 1$ until the outcome of test $\langle v_{k,i}, v_{k,(i+1) \bmod m_k} \rangle$ is 1 for some i or we have m_k tests. Let $X = \{(v_{k,i}, v_{k,(i+1) \bmod m_k}) :$ the outcome of test $\langle v_{k,i}, v_{k,(i+1) \bmod m_k} \rangle$ is $1\}$. Then, it is easy to see the followings:

Claim 8. *$|E_k \cap X| \le 1$ for any $k \in [2^{n-1}]$.* □

Claim 9. *If $E_k \cap X = \emptyset$ then every vertex of V_k is fault-free.* □

Claim 10. *If $(v_{k,i}, v_{k,(i+1) \bmod m_k}) \in X$, at least one of $v_{k,i}$ and $v_{k,(i+1) \bmod m_k}$ is faulty.* □

We have $|X| \le |F| \le 3$ by Claims 8 and 10. There are four cases.

(i) $|X| = 3$: Let Y denote the set of vertices incident with an edge in X. Every vertex of $V(CCC(n)) - Y$ is fault-free since Y has three faulty vertices by Claims 8 and 10. If $(v_{k,i}, v_{k,(i+1) \bmod m_k}) \in X$ then one of $v_{k,i}$ and $v_{k,(i+1) \bmod m_k}$ is faulty and the other is fault-free by Claim 10. If $i \ge 1$ then $v_{k,i}$ is fault-free, for otherwise $|F| \ge |\{v_{k,0}, v_{k,1}, \ldots, v_{k,i}\}| + |X - (v_{k,i}, v_{k,(i+1) \bmod m_k})| \ge 4$, which is a contradiction. Since $v_{k,i}$ is fault-free, $v_{k,(i+1) \bmod m_k}$ is faulty. If $i = 0$ then test $v_{k,0}$ by $v_{k,n-1} \in V(CCC(n)) - Y$. If the outcome of test $\langle v_{k,n-1}, v_{k,0}\rangle$ is 1 then $v_{k,0}$ is faulty, and otherwise $v_{k,1}$ is faulty. Hence, we can identify F using at most $N = n \times 2^n$ tests.

(ii) $|X| = 2$: If $E_k \cap X = \emptyset$ then every vertex of V_k is fault-free by Claim 9. Thus, we can diagnose V_k with m_k tests. If $E_k \cap X \ne \emptyset$ then $|E_k \cap X| = 1$ by Claim 8, and so $|X - E_k| = 1$. It follows that V_k has at most two faulty vertices by Claim 10. Thus, from Claim 7 and the fact that $m_k \ge 8$, we can diagnose all vertices of V_k by applying Algorithm 1 for C_{m_k}. Notice that if $(v_{k,i}, v_{k,(i+1) \bmod m_k}) \in E_k \cap X$ then it suffices for Algorithm 1 to perform at most $(m_k - i)$ additional tests in order to diagnose V_k, since the outcome of $i+1$ tests $\langle v_{k,i}, v_{k,(i+1) \bmod m_k}\rangle (j \in [i+1])$ can be used to diagnose V_k. Thus, we can diagnose V_k with at most $m_k + 1$ tests. Since $|\{k : E_k \cap X\}| = |X| = 2$, we can identify F with at most $\sum_{i \in [2^{n-1}]} m_k + |X| = N + 2$ tests.

(iii) $|X| = 1$: Let $u_{k,i} \in X$ for some $k \in [2^{n-1}]$ and $i \in [m_k]$. Then, every vertex of $V(CCC(n)) - V_k$ is fault-free. We further distinguish three cases.

(iii)-(a) $i = 0$ or $i = 1$: We can identify F by testing $v_{k,j}$ by $\overline{v_{k,j}}$ for every $j \in [m_k]$, since $v_{k,j} \in V(CCC(n)) - V_k$ by Claim 6. The total number of tests performed is at most $N - m_k + (i+1) + m_k \le N + 2$.

(iii)-(b) $2 \le i \le m_k - 2$: Perform test $\langle \overline{v_{k,j}}, v_{k,j}\rangle$ in order of $j = 0, 1, \ldots$ until the outcome of test $\langle \overline{v_{k,j}}, v_{k,j}\rangle$ is 0 for some $j = l$. Notice that $\overline{v_{k,i}}$ is fault-free since $\overline{v_{k,i}} \in V(CCC(n)) - V_k$ by Claim 6. Thus,

$$\begin{cases} v_{k,0}, v_{k,1}, v_{k,2} \in F & \text{if } l = 3, \\ v_{k,0}, v_{k,1}, v_{k,i+1} \in F & \text{if } l = 2, \\ v_{k,0}, v_{k,i+1} \in F & \text{if } l = 1, \text{ and} \\ v_{k,i+1} \in F & \text{if } l = 0. \end{cases}$$

If $l \le 1$ then we test $v_{k,j}$ by $\overline{v_{k,j}}$ for every integer j, $i+2 \le j \le m_k - 1$. If the outcome of test $\langle \overline{v_{k,j}}, v_{k,j}\rangle$ is 1 then $v_{k,j}$ is faulty. Hence, we can identify F using at most $N - m_k + (i+1) + 2 + (m_k - i - 2) \le N + 1$ tests.

(iii)-(c) $i = m_k - 1$: In this case, $v_{k,0}$ is faulty and $v_{k,j}$ is fault-free for any integer j, $3 \le j \le m_k - 1$. Thus, if the outcome of test $\langle \overline{v_{k,1}}, v_{k,1}\rangle$ is 0, then $F = \{v_{k,0}\}$; If the outcome of test $\langle \overline{v_{k,1}}, v_{k,1}\rangle$ is 1 and the outcome of test

$\langle \overline{v_{k,2}}, v_{k,2} \rangle$ is 0, then $F = \{v_{k,0}, v_{k,1}\}$; If the outcome of test $\langle \overline{v_{k,1}}, v_{k,1} \rangle$ is 1 and the outcome of test $\langle \overline{v_{k,2}}, v_{k,2} \rangle$ is 1, then $F = \{v_{k,0}, v_{k,1}, v_{k,2}\}$. Hence, we can identify F using at most $N + 2$ tests.

(iv) $|X| = 0$: By Claim 9, we can identify $F = \emptyset$ using N tests.

By (i), (ii), (iii), and (iv), we can diagnose $CCC(n)$ using at most $N + 2$ tests. □

4 Adaptive Parallel Diagnosis

In adaptive parallel diagnosis, several tests may be performed simultaneously in a testing round, but each vertex can participate in at most one test. That is, the tests in a testing round are a directed matching on the vertices. Since at least $N+t-1$ tests are necessary for adaptive parallel diagnosis of an N-vertex graph with at most t faulty vertices and there are at most $N/2$ tests in each testing round, $\lceil (N + t - 1)/(N/2) \rceil$ is a general lower bound for the number of testing rounds. Thus we have the following.

Theorem IV [2] *If G is a graph with at most t faulty vertices then* 3 *testing rounds are necessary to adaptively diagnose G provided that $t \geq 2$.*

4.1 Even Cycles

The following theorem will be used in the next section.

Theorem 3. *An even cycle C_N can be adaptively diagnosed in* 3 *testing rounds if the number of faults is not more than* 2 *and $N \geq 6$.*

Proof. In **Step** 1 of Algorithm 1 shown in Figure 2, all tests can be performed in two rounds, since N is even. In **Step** 2, just one test is performed, and this can be done in a testing round. Thus we have the theorem.

4.2 Hypercubes

The following theorem is shown in [6].

Theorem V [6] *$Q(n)$ can be adaptively diagnosed in* 4 *testing rounds if the number of faults is not more than n and $n \geq 3$.*

We prove the following theorem.

Theorem 4. *$Q(n)$ 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 $n \geq 4$.*

4.2.1 Proof of Theorem 4 Let $t = n - \lceil \log(n - \lceil \log n \rceil + 4) \rceil + 2$. $Q(n)$ is represented as $Q(n-t+2) \times Q(t-2)$. Notice that $t \geq 3$ since $n \geq 4$. We need a few technical lemmas.

Lemma 2. $|V(Q(n-t+2))| > t$.

Proof. $|V(Q(n-t+2))| = 2^{n-t+2} = 2^{\lceil \log(n - \lceil \log n \rceil + 4) \rceil} \geq n - \lceil \log n \rceil + 4 > t$. □

Lemma 3. *For any $S \subseteq V(Q(n))$ with $|S| \leq n$, each vertex in S has a distinct adjacent vertex in $V(Q(n)) - S$.*

Proof. We prove the lemma by induction on n. The case when $n = 1$ is trivial. Assume that the lemma holds if $n = k$. Let S be a set of vertices of $Q(k+1)$ with $|S| \leq k+1$. Since $Q(k+1) = Q(k) \times P_2$, $Q(k+1)$ can be decomposed into two disjoint copies of $Q(k)$, say $Q_1(k)$ and $Q_2(k)$. We distinguish two cases.

(i) $S \subseteq V(Q_1(k))$: The vertices in $Q_2(k)$ corresponding to the vertices in S are the desired vertices.

(ii) $S \cap V(Q_1(k)) \neq \phi$ and $S \cap V(Q_2(k)) \neq \phi$: Let $S_i = S \cap V(Q_i(k))$ $(i = 1, 2)$. Since $|S_i| \leq k (i = 1, 2)$, S_i has a desired set of vertices in $Q_i(k)$ by the inductive hypothesis. □

Now we are ready to describe our algorithm. Our algorithm works in two steps. It is well-known that $Q(n)$ has a Hamilton cycle. In the first step, we perform in two testing rounds all tests along a Hamilton cycle in all copies of $Q(n-t+2)$ in the clockwise direction. A copy of $Q(n-t+2)$ is said to be fault-free if it has no faulty vertex, and faulty otherwise. The following is immediate from Lemma 2.

Lemma 4. *A copy of $Q(n-t+2)$ is faulty if and only if the tests along a Hamilton cycle have an outcome of* 1. □

Let $\mathcal{F}$ be the set of all faulty copies of $Q(n-t+2)$.

The second step of our algorithm is distinguished in three cases depending on $|\mathcal{F}|$.

If $|\mathcal{F}| = t$ then each faulty copy of $Q(n-t+2)$ has just one faulty vertex, which we can identify from the test results in the first step.

If $|\mathcal{F}| = t-1$ then each faulty copy of $Q(n-t+2)$ has at most two faulty vertices, which we can identify in one more testing round by Theorem 3.

If $|\mathcal{F}| \leq t-2$ then for each faulty copy Q_F of $Q(n-t+2)$, there is a distinct fault-free copy Q_H of $Q(n-t+2)$ in which each vertex v_H is adjacent to the corresponding vertex v_F in Q_F by Lemma 3. By performing the tests $\langle v_H, v_F \rangle$ for all faulty copies of $Q(n-t+2)$ in one testing round, we can identify all the faults.

Our algorithm is summarized in Fig. 3.

4.3 CCC's

The following theorem is proved based on adaptive serial diagnosis for CCC's in Section 3.3.

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

Fig. 3. Algorithm 2

Theorem 5. *$CCC(n)$ can be adaptively diagnosed in 3 testing rounds if the number of faults is not more than 3 and $n \geq 4$.*

Proof. Let $(V_0, V_1, \ldots, V_{2^{n-1}-1})$ be a partition of $V(CCC(n))$ defined in the proof of Theorem 2. Our algorithm works in two steps. By Claim 7, every block $V_k (k \in [2^{n-1}])$ is isomorphic to C_{m_k}. In the first step, we perform in two testing rounds all tests along a cycle C_{m_k} in all block V_k in the clockwise direction. A block V_k is said to be fault-free if it has no faulty vertex, and faulty otherwise. Since every block V_k has $4\lceil n/2 \rceil \geq 4$ vertices, we have the following.

Lemma 5. *V_k is faulty if and only if the tests along a cycle C_{m_k} have an outcome of 1.* □

Let $\mathcal{F}$ be the set of all faulty blocks. $|\mathcal{F}| \leq 3$ by the assumption. The second step of our algorithm is distinguished in four cases depending on $|\mathcal{F}|$.

(i) $|\mathcal{F}| = 3$: Each block $V_k \in \mathcal{F}$ has only one faulty vertex since there are at most three faulty vertices. Thus faulty vertices can be identified from the test results in the first step.

(ii) $|\mathcal{F}| = 2$: Each block $V_k \in \mathcal{F}$ has at most 2 faulty vertices, which we can identify in one more testing round by Theorem 3.

(iii) $|\mathcal{F}| = 1$: It is easy to see from Claim 6 that each vertex v_F in the block $V_F \in \mathcal{F}$, there exists a distinct vertex $\overline{v_F} \in V(CCC(n)) - V_F$ adjacent with v_F. We perform tests $\langle \overline{v_F}, v_F \rangle$ for all vertices v_F in V_F in one testing round.

(iv) $|\mathcal{F}| = 0$: From the test results in the first step, we know that there is no fault.

5 Concluding Remarks

1. We can prove that 4 testing rounds are necessary and sufficient to adaptively diagnose an odd cycle C_N if the number of faulty vertices is at most 2 and $N \geq 5$.

2. We can prove that $CCC(3)$ is also adaptively 3-diagnosable using at most $N+2$ tests. The proof is similar to that of Theorem 2. Notice that $CCC(2)$ is just C_8. We can show that $CCC(3)$ can be adaptively diagnosed in 4 testing rounds if the number of faults is at most 3. It is open if 3 testing rounds are sufficient for $CCC(3)$.
3. $Q(3)$ can be adaptively diagnosed in 3 testing rounds if the number of faults is at most 3, as mentioned in [13]. 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 4 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 $n - \lceil \log(n - \lceil \log n \rceil + 3) \rceil + 3 \leq t \leq n$. A similar approach based on the decomposition of $Q(n)$ into subcubes can be found in [13], in which it is shown that $Q(n)$ is adaptively n-diagnosable using $N + 3n/2$ tests if $n \geq 3$, and $Q(n)$ is adaptively diagnosable in 11 testing rounds if the number of faulty vertices is not more than n and $n \geq 3$.
4. We can prove that a d-dimensional torus can be adaptively diagnosed in 3 testing rounds if the number of faulty vertices is at most $2d$ and the number of vertices in the side is even. We can also show that a d-dimentional mesh can be adaptively diagnosed in 3 testing rounds if the number of faulty vertices is at most d. The details will appear in the forthcoming full version of the paper.

Acknowledgments

The adaptive diagnosis algorithms in Section 3 are a joint work with Ken Hayase. The research is a part of CAD21 Project at TIT.

References

1. R. Beigel, W. Hurwood, and N. Kahale. Fault diagnosis in a flash. *Proc. of 36th FOCS*, pages 571–580, 1995. 87
2. R. Beigel, S.R. Kosaraju, and G.F. Sullivan. Locating faults in a constant number of testing rounds. *Proc. of 1st SPAA*, pages 189–198, 1989. 87, 94
3. R. Beigel, G. Margulis, and D.A. Spielman. Fault diagnosis in a small constant number of parallel rounds. *Proc. of 5th SPAA*, pages 21–29, 1993. 87
4. R.P. Bianchini, Jr., and R. Buskens. An adaptive distributed system level-diagnosis algorithm and its implementation. *Proc. of 21st Int. Symp. on Fault Tolerant Computing*, pages 222–229, 1991. 87
5. R.P. Bianchini, Jr., K. Goodwin, and D.S. Nydick. Practical application and implementation of distributed system-level diagnosis thoery. *Proc. of 20st Int. Symp. on Fault Tolerant Computing*, pages 332–339, 1984. 87
6. A. Bjöklund. Optimal adaptive fault diagnosis of hypercubes. *Proc. of SWAT,* LNCS 1851, pages 527–534, 2000. 87, 89, 94
7. P.M. Blecher. On a logical problem. *Discrete Mathematics*, 43:107–110, 1983. 87, 89

8. C. Feng, L.N. Bhuyan, and F. Lombardi. Adaptive system-level diagnosis for hypercube multiprocessors. *IEEE Trans. Computers*, 45:1157–1170, 1996. 87
9. S.L. Hakimi and A.T. Amin. Characterization of connection assignment of diagnossable systems. *IEEE Trans. Computers*, 23:86–88, 1974. 86
10. S.L. Hakimi and K. Nakajima. On adaptive system daignosis. *IEEE Trans. Computers*, 33:234–240, 1984.
11. S.L. Hakimi, M. Otsuka, E.F. Schmeichel, and G.F. Sullivan. A parallel fault identification algorithm. *Journal of Algorithm*, 11:231–241, 1990. 87
12. S.L. Hakimi and E.F. Schmeichel. An adaptive algorithm for system level daignosis. *Journal of Algorithms*, 5:524–530, 1984. 87
13. E. Kranakis and A. Pelc. Better adaptive daignosis of hypercubes. *IEEE Trans. Computers*, 49:1013–1020, 2000. 87, 97
14. E. Kranakis, A. Pelc, and A. Spatharis. Optimal adaptive fault diagnosis for simple multiprocessor systems. *Networks*, 34:206–214, 1999. 87, 91
15. K. Nakajima. A new approach to system daignosis. *Proc. 19th allerton Conf. Commun. Contr. and Computing*, pages 697–706, 1981. 87
16. F.P. Preparata, G. Metze, and R.T. Chien. On the connection assignment problem of daiganosable systems. *IEEE Trans. Electronic Computers*, 16:848–854, 1967. 86, 88
17. Pei Yuan Wu. Partial solution to problem 81-6. *Journal of Algorithms*, 12:379–380, 1982. 87, 89