

On diagnosability of large multiprocessor networks

R. Ahlswede and H. Aydinian¹

*Department of Mathematics, University of Bielefeld, POB 100131, D-33501
Bielefeld, Germany*

Keywords: system level diagnosis, PMC model, sequential diagnosis, diagnosability.

Abstract

We consider problems of fault diagnosis in multiprocessor systems. Preparata, Metze and Chien (1967) introduced a graph theoretical model for system-level diagnosis, in which processors perform tests on one another via links in the system. Fault-free processors correctly identify the status of tested processors, while the faulty processors can give arbitrary test results. The goal is to identify faulty processors based on the test results. A system is said to be t -diagnosable if faulty units can be identified, provided the number of faulty units present does not exceed t . We explore here diagnosis problems for n -cube systems and give bounds for diagnosability of the n -cube. We also describe a simple diagnosis algorithm A which is linear in time and which can be used for sequential diagnosis as well as for incomplete diagnosis in one step. In particular the algorithm applied to arbitrary topology based interconnection systems G with N processors improves previously known ones. It has sequential diagnosability $t_A(G) \geq \lceil 2N^{\frac{1}{2}} \rceil - 3$, which is optimal in the worst case.

1 Introduction

The concept of system-level diagnosis was introduced by Preparata, Metze and Chien [12] to perform automatic fault diagnosis in multiprocessor systems. In their graph theoretical model, called PMC model, a system S is composed of independent units u_1, \dots, u_n connected by communication links. The system is represented as an undirected graph $G = (V, E)$, where the vertices represent

¹ Supported by DFG-Schwerpunkt Nr. 1126 “Algorithmik großer und komplexer Netzwerke”.

units and edges represent interconnection links. In the PMC model diagnosis is based on a suitable set of tests between units. A unit u_i can test u_j iff the vertices corresponding to u_i and u_j in the graph $G = (V, E)$ of the system S are adjacent. The outcome of a test in which u_i tests u_j is denoted by a_{ij} , where $a_{ij} = 1$ if u_i finds u_j to be faulty and $a_{ij} = 0$ if u_i finds u_j to be fault-free.

The basic conditions of the PMC model are the following:

- The fault-free units give correct test outcomes.
- The answers of faulty units are unreliable.
- The number of faulty units t is bounded and all faults are permanent.

The set of tests for the purpose of diagnosis is represented by a set of directed edges where the presence of oriented edge (u_i, u_j) means that u_i tests u_j . Given a faulty set of units $F \subset V$ the set of all test outcomes $\{a_{ij}\}$ is called *syndrome*. The task is to identify the faulty units based on a syndrome produced by the system. In [12] two different kinds of strategies were introduced for implementing the diagnosis approach.

One-step diagnosis (or diagnosis without repair): a system is called t -fault diagnosable (or shortly t -diagnosable) in one step, if all faulty units can be uniquely identified from any syndrome, provided the number of faulty units does not exceed t .

Sequential diagnosis (or diagnosis with repair): a system is called sequentially t -diagnosable if it can identify at least one faulty unit from any syndrome, provided the number of faulty units does not exceed t . Under a sequential diagnosis strategy a system can locate a faulty unit, repair it and then repeat the process until all faulty units are repaired.

The *degree of diagnosability*, or simply diagnosability, of a system graph G is defined (for both kinds of strategies) as the maximum t such that the system is t -diagnosable.

The PMC model has been widely studied (see [4] for a good survey). It is known ([12]) that the maximum degree of diagnosability of a one-step diagnosis algorithm for any system is bounded from above by the minimum vertex degree of the interconnection graph. However, the real commercial multiprocessor systems are based on topologies of graphs with small average vertex degree (like grids, hypercubes, cube-connected cycles, trees etc).

Sequential diagnosis is a much more powerful strategy than one-step t -fault diagnosis. On the other hand the sequential diagnosis has the disadvantage of repeated execution of diagnosis and repair phases and may be time consuming for large systems.

That was the motivation for developing diagnosis algorithms (see [5]) which are able to diagnose in one step the status of a large fraction of the system units (i.e. if a “large” subset F' of the actual fault set F can be identified from any syndrome, provided $|F'| \leq t$). This approach is referred to as *incomplete*

diagnosis in one step.

Clearly, incomplete diagnosis in one step amounts to *one-step diagnosis* if $F' = F$, while sequential diagnosis is a strategy using multiple steps of incomplete diagnosis to identify and replace all faulty units. A common and usually implicit assumption is that no additional faults are introduced while the process of sequential diagnosis is going on.

Since the degree of diagnosability is different for different kinds of diagnosis strategies the notation $t_0(G)$ for one-step diagnosis and $t(G)$ for incomplete diagnosis in one step is appropriate. Actually the paper is concerned with $t(G)$ only except for Corollary 3, which provides an upper bound to the number of iterations in sequential diagnosis. Moreover, Section 4.3 deals with the diagnosability of a general graph G under algorithm A and here we use the notation $t_A(G)$.

The *diagnostic graph* DG of a system graph $G = (V, E)$, corresponding to a given syndrome, consists of bidirectional arcs, between every two neighbors of the original graph G , labelled by 0 or 1. Let $\{u, v\} \in E(G)$, then the presence of oriented edges (u, v) and (v, u) with $a_{uv} = 1$ and $a_{vu} = 0$ implies that v is faulty. In the sequel we assume that a diagnostic graph does not contain such “trivial” configurations. Thus the outcomes of any two neighbors coincide. Therefore we can represent a diagnostic graph as an undirected graph where each edge is labelled by a 0 or a 1.

Given a syndrome, a subset F of the vertex set V is called a *consistent fault set* if the assumption that the vertices in F are faulty and those in $V \setminus F$ are fault-free is consistent with the syndrome. The following simple facts are useful for obtaining upper and lower bounds for the diagnosability of a system graph.

Fact 1 Given a syndrome, let F_1, \dots, F_k be a collection of consistent fault sets with $|F_i| \leq t$; $i = 1, \dots, k$. Then G is not sequentially t -diagnosable if $\bigcap_{i=1}^k F_i = \emptyset$ and $\bigcup_{i=1}^k F_i = V$.

Given a diagnostic graph DG , define the subgraph G_0 consisting of edges labelled only by 0 (0-edges). The connected components of the graph G_0 will be called 0-components of DG .

Fact 2 All vertices of a 0-component have the same status: “faulty” or “fault-free”.

This fact was used for lower estimates of diagnosabilities under sequential diagnosis (of certain regular graphs) and for the incomplete diagnosis approach in [9], [5]. For a given system let the number of faults be bounded by t . Suppose then that given a syndrome the diagnostic graph contains a 0-component K

of size $|K| > t$. Then clearly all vertices of K are fault-free. Thus, if the size of the largest 0-component is lower bounded by $t + 1$ (for all syndromes), then the degree of sequential diagnosability of the system $t(G) \geq t$.

Motivated by a certain model of parallel computing in unreliable networks Ahlswede and Koschnick [2] considered the following extremal problems for graphs.

Problem 1 Given a connected graph $G = (V, E)$, let $\lambda(G, c)$ denote the maximal number such that removal of any $\lambda(G, c)$ or less vertices results in a graph with a connected component of size at least c . Determine or estimate $\lambda(G, c)$.

Problem 2 Removing edges instead of vertices, define analogously the function $\mu(G, c)$ and determine or estimate $\mu(G, c)$.

Sometimes it is more convenient to deal with the function $\lambda^*(G, c)$ (resp. $\mu^*(G, c)$) = minimal number with the property that there exist $\lambda^*(G, c)$ vertices (resp. $\mu^*(G, c)$ edges) whose removal results in a graph with a maximal connected component of size $\leq c$. Observe that $\lambda^*(G, c) = \lambda(G, c + 1) + 1$ and $\mu^*(G, c) = \mu(G, c + 1) + 1$. We note that both functions $\lambda(G, c)$ and $\mu(G, c)$ are useful for diagnosis problems in multiprocessor systems. In fact the following derived quantity is essential. For a graph G define $m(G) = \max\{x : \lambda(G, x + 1) \geq x\}$.

Fact 3 If the number of faults $t \leq m(G)$, then for every syndrome at least $m(G) + 1$ units can be identified as fault-free. Thus the sequential diagnosability of a graph $t(G) \geq m(G)$.

Note, however, that in general $m(G)$ can be much smaller than the degree of sequential diagnosability. Consider for example a star graph G on $N = 2k + 1$ vertices. It is not hard to observe that the sequential diagnosability of this graph $t(G) = k$ while $m(G) = 0$.

A result in [14] implies that $\lceil 2N^{\frac{1}{2}} \rceil - 3 \leq t(G) \leq \lfloor (N - 1)/2 \rfloor$ holds for a connected graph G on N vertices. However, this result does not tell us how to identify a fault set present in a system for a given syndrome. Note also that the problem of determining the sequential diagnosability of a system is shown to be co-NP complete [13].

For short, let us use the notation $t(n)$, $\lambda(n, c)$ and $m(n)$ for n -cubes. Khanna and Fuchs [9] studied the function $m(G)$ in order to get estimates for diagnosability of the n -cube and other structures. It was shown (see also Caruso et al. [5]) that the degree of sequential diagnosability of the n -cube $t(n) = \Omega\left(\frac{2^n \log n}{n}\right)$. They also described a simple algorithm for sequential diagnosis (referred to as PARTITION) which actually reduces to finding a “largest”

0-component and declaring it to be fault-free. The algorithm gives good performance for graphs with large $m(G)$, as e.g. for n -cubes. The augmented version of the algorithm applied to arbitrary interconnection graphs on N vertices has diagnosability $\Omega(N^{\frac{1}{3}})$.

Recently Yamada et al. [16] improved the lower bound into $t(n) = \Omega(2^n/\sqrt{n})$, applying the result of Kleitman [10], which states that $\lambda(n, 2^{n-1}+1) \geq \binom{n}{n/2} - 1$ and hence $m(n) = \Omega(2^n/\sqrt{n})$. They also obtained the upper bound $t(n) = O\left(\frac{2^n \log n}{\sqrt{n}}\right)$ which greatly improves previously known ones. A sequential diagnosis algorithm (referred to as HYBRID) for arbitrary interconnection graphs G on N vertices was proposed in [16] with diagnosability $t_{HYBRID}(G) \geq \lceil \sqrt{N-1} \rceil - 1$.

In this paper we present new bounds for the function $\lambda(n, c)$ which imply bounds for incomplete diagnosis in one step. **We present a short proof for the upper bound in [16] mentioned above.** We also describe a simple diagnosis algorithm A which is linear in time and which can be used for sequential diagnosis as well as for incomplete diagnosis in one step. The algorithm applied to arbitrary topology based interconnection systems improves the $t_{HYBRID}(G)$ ([16]) to $t_A(G) \geq \lceil 2N^{\frac{1}{2}} \rceil - 3$, which is optimal in the worst case.

2 Bounds for $\lambda(n, c)$

Let $\mathcal{H}_n = \{0, 1\}^n$ denote the binary Hamming space and let $d(x, y)$ denote the Hamming distance between any two vectors $x, y \in \mathcal{H}_n$, defined as the number of coordinates in which they differ. We associate \mathcal{H}_n with the Hamming graph $G(\mathcal{H}_n)$ where two vertices $x, y \in \mathcal{H}_n$ are adjacent iff $d(x, y) = 1$. For $x \in \mathcal{H}_n$ let $B_r(x)$ denote the Hamming ball of radius r centered at x , that is, $B_r(x) = \{u \in \mathcal{H}_n : d(u, x) \leq r\}$. Define also a quasi-ball $B_r^\delta(x)$ consisting of the union of $B_r(x)$ and δ elements of \mathcal{H}_n whose distance from x is $r+1$. Thus $M = |B_r^\delta(x)| = b_r + \delta$, where $b_r := \sum_{i=0}^r \binom{n}{i}$ and $0 \leq \delta < \binom{n}{r+1}$.

For $A \subset \mathcal{H}_n$ the neighborhood of A is defined by $\sigma(A) = \{x : x \in \mathcal{H}_n, d(x, y) \leq 1 \text{ for some } y \in A\}$. The boundary of A is defined by $\Gamma A = \sigma(A) \setminus A$. For $n, M \in \mathbb{N}$, $1 \leq M < 2^n$, there is a unique representation (see [8])

$$M = \binom{n}{n} + \cdots + \binom{n}{k+1} + \binom{a_k}{k} + \cdots + \binom{a_s}{s} \quad (2.1)$$

for some $k, s \leq k \leq n$ and $n > a_k > a_{k-1} > \cdots > a_s \geq s \geq 1$.

We denote by $B_{n-k}^\delta((1, 1, \dots, 1))$ the quasi-ball of size M consisting of all

elements $x \in \mathcal{H}_n$ with Hamming weight $wt_H(x)$, the number of nonzero coordinates, greater than k and $\delta = \binom{a_k}{k} + \dots + \binom{a_s}{s}$ elements of weight k taken in colexicographic order. Thus $B_{n-k}^\delta((1, 1, \dots, 1))$ is a quasi-ball centered at $(1, \dots, 1)$. The following result is known as Harper's Vertex Isoperimetric Theorem.

Theorem H ([7]) *For every $A \subset \mathcal{H}_n$ with $|A| = M$ we have*

$$|\Gamma A| \geq |\Gamma B_{n-k}^\delta((1, 1, \dots, 1))|. \quad (2.2)$$

Equivalently with the parameters of (2.1)

$$|\sigma(A)| \geq \binom{n}{n} + \dots + \binom{n}{k} + \binom{a_k}{k-1} + \binom{a_{k-1}}{k-2} + \dots + \binom{a_s}{s-1}. \quad (2.3)$$

A weaker version of the theorem is as follows. Let $|A| = M = \binom{n}{n} + \dots + \binom{n}{k+1} + \binom{x}{k}$, where x is a real number and $\binom{x}{k} = x(x-1)\dots(x-k+1)/k!$. Then

$$|\sigma(A)| \geq \binom{n}{n} + \dots + \binom{n}{k} + \binom{x}{k-1}. \quad (2.4)$$

Remark 1 For the explicit version (2.3) of the theorem and the weaker version (2.4) we refer to [8] and [6] respectively. A dual form of Theorem H is given in [1]. Note also that the result holds for any translate of $B_{n-k}^\delta((1, 1, \dots, 1))$ in \mathcal{H}_n .

The next result gives a characterization of an optimal set $A \subset \mathcal{H}_n$ for which equality in (2.2) holds.

Theorem A ([3]) *Every optimal set $A \subset \mathcal{H}_n$ with $|A| = b_r + \delta$ contains a ball of radius r .*

As shown in [2], Theorem H can be used to estimate the quantity $\lambda(n, c)$ (and hence $m(\mathcal{H}_n)$). It was shown that for $N_{n,k+1} > 2/3 \cdot 2^n$, where $N_{n,k+1} := \binom{n}{n} + \dots + \binom{n}{k+1}$, one has $\lambda^*(n, N_{n,k+1}) = \binom{n}{k}$. It is also known ([10]) that $\lambda(n, 2^{n-1} + 1) \geq \binom{n}{n/2} - 1$ (the result mentioned in the Introduction). Here we extend these results to the following

Theorem 1 (i) *For $n \geq 2k$ we have*

$$\lambda^*(n, N_{n,k+1}) = \binom{n}{k}$$

(ii)

$$\lambda(n, N_{n,k+1}) = \begin{cases} \binom{n}{k} & , \text{if } n > 2k \\ \binom{n}{k} + 1 & , \text{if } n = 2k, k \geq 3. \end{cases}$$

Proof (i) We need the following simple observation.

Proposition 1 *Let $A \subset \mathcal{H}_n$ and $N_{n,i+1} < |A| < N_{n,i}$. Then*

$$|\Gamma A| > \begin{cases} \binom{n}{i-1} & , \text{if } n \geq 2i \\ \binom{n}{i} & , \text{if } n < 2i \end{cases}$$

Proof We have $|A| = N_{n,i+1} + \alpha$ for some integer $1 \leq \alpha < \binom{n}{i}$. Thus $\alpha = \binom{x}{i}$ for some real $x < n$ and in view of Theorem H

$$|\Gamma A| \geq \binom{x}{i-1} + \binom{n}{i} - \binom{x}{i}.$$

For $n \geq 2i$ we have $\binom{n}{i-1}(\frac{n-i+1}{i} - 1) > \binom{x}{i-1}(\frac{x-i+1}{i} - 1)$ and thus $\binom{x}{i-1} + \binom{n}{i} - \binom{x}{i} > \binom{n}{i}$. Also if $n < 2i$, then $\binom{x}{i-1} > \binom{x}{i}$ thus $|\Gamma A| > \binom{n}{i}$. \square

We note first that

$$\lambda^*(n, N_{n,k+1}) \leq \binom{n}{k}.$$

This is clear because the graph obtained by removing all vertices of weight k consists of two connected components of size $N_{n,k+1}$ and $N_{n,n-k+1}$. Thus it remains to show that

$$\lambda^*(n, N_{n,k+1}) \geq \binom{n}{k}. \quad (2.5)$$

We proceed by induction on n . For $n = 1, 2$ the statement is trivial.

Induction step: $n - 1 \rightarrow n$.

Let $H_0 \subset \mathcal{H}_n$ be the vectors with 0 in the last coordinate, and let $H_1 = \mathcal{H}_n \setminus H_0$.

Suppose, for a contradiction, the opposite relation holds in (2.5). That is, there exists an $S \subset \mathcal{H}_n$ with $|S| < \binom{n}{k}$ such that the largest connected component $A \subset \mathcal{H}_n \setminus S$ has size $|A| \leq N_{n,k+1}$.

Define also $S_0 = S \cap H_0$, $S_1 = S \cap H_1$ and $A_0 = A \cap H_0$, $A_1 = A \cap H_1$. W.l.o.g. suppose now that $|S_0| \leq \frac{1}{2}|S|$. Then $|S_0| < \frac{1}{2}\binom{n}{k} < \binom{n-1}{k}$, if $n \geq 2k + 1$ and $|S_0| < \binom{n-1}{k-1}$, if $n = 2k$.

In the case $n \geq 2k + 1$ by the induction hypothesis we have for $n - 1 \geq 2k$ $\lambda^*(n - 1, N_{n-1, k+1}) = \binom{n-1}{k}$ and since $|S_0| < \binom{n-1}{k}$ we infer that $|A_0| > N_{n-1, k+1}$.

In the case $n = 2k$ by the induction hypothesis we have for $n - 1 \geq 2(k - 1)$ $\lambda^*(n - 1, N_{n-1, k}) = \binom{n-1}{k-1}$ and since $|S_0| < \binom{n-1}{k-1}$ we infer that $|A_0| > N_{n-1, k}$.

We improve now the inequality $|A| \leq N_{n, k+1}$ to the inequality

$$|A| < N_{n, n-k+1}. \quad (2.6)$$

Observe first that for the case $n = 2k$ $N_{n, n-k+1} = N_{n, k+1}$ and so the bounds are equal. The identities $|A| = N_{n, n-k+1-j}$ with $0 \leq j \leq n - 2k$ are excluded, because $|S| > |\Gamma A|$, $|\Gamma A| \geq \binom{n}{n-k-j}$ by Theorem H, and $\binom{n}{n-k-j} \geq \binom{n}{k}$ for $n \geq 2k$ imply $|S| > \binom{n}{k}$, a contradiction. Suppose now $n \geq 2k + 1$ and let $N_{n, i+1} < |A| < N_{n, i}$, where $k + 1 < i < n - k$. Then (by Proposition 1) for $n \geq 2i$ we have $|S| > |\Gamma A| > \binom{n}{i-1} > \binom{n}{k}$, a contradiction. Similarly, for $n < 2i$ we have $|\Gamma A| > \binom{n}{i} > \binom{n}{n-k} = \binom{n}{k}$, again a contradiction. Therefore $|A|$ does not fall into the interval $[N_{n, n-k+1}, N_{n, k+1}]$ and since $|A| \leq N_{n, k+1}$ (2.6) follows (actually done with strict inequality).

Let us now bound $|S|$ from below. Denote by $\tilde{A}_0 \subset \mathcal{H}_{n-1}$ the set of vectors obtained from $|A_0|$ by deletion of the last coordinate. Clearly $|\tilde{A}_0| = |A_0|$. Observe then that

$$|\Gamma A_0| = |A_0| + |\Gamma \tilde{A}_0|. \quad (2.7)$$

It is also clear that

$$|\Gamma A| \geq |A_0| - |A_1| + |\Gamma \tilde{A}_0|. \quad (2.8)$$

By the observation above we have

$$|A| = |A_0| + |A_1| \leq N_{n, n-k+2} + \alpha, \quad (2.9)$$

where $\alpha < \binom{n}{n-k+1}$.

By the induction hypothesis we also have $|A_0| > N_{n-1, k+1}$. This together with (2.8) and (2.9) implies

$$\begin{aligned} |S| &\geq |\Gamma A| \geq 2|A_0| - |A| + |\Gamma \tilde{A}_0| > 2N_{n-1, k+1} + |\Gamma \tilde{A}_0| - N_{n, n-k+2} - \binom{n}{n-k+1} \\ &= N_{n, k+2} + \binom{n-1}{k+1} + |\Gamma \tilde{A}_0| - N_{n, n-k+1}. \end{aligned} \quad (2.10)$$

If now $k + 2 < n - k + 1$, that is, $n \geq 2k + 2$ then (2.10) implies

$$|S| > \binom{n}{k+2} + \binom{n-1}{k+1} > \binom{n}{k},$$

a contradiction. Thus, it remains to consider two cases. It is worth noticing that their treatment also can be used to cover the cases $n \geq 2k + 2$ which are done differently here.

Case $n = 2k + 1$: We have now $|A_0| > N_{2k,k+1}$. Suppose first $|A_0| \geq N_{2k,k}$. Then (2.10) implies

$$|\Gamma A| > 2N_{2k,k} - N_{2k+1,k+3} - \binom{2k+1}{k+2} = \binom{2k+1}{k+1} + \binom{2k}{k+1} > \binom{2k+1}{k},$$

a contradiction. Let now $|A_0| = N_{2k,k+1} + \beta$, where $1 \leq \beta < \binom{2k}{k}$. Then (2.10) gives

$$\begin{aligned} |\Gamma A| &> 2N_{2k,k+1} + 2\beta + |\Gamma \tilde{A}_0| - N_{2k+1,k+2} \\ &= N_{2k+1,k+2} + \binom{2k}{k+1} + 2\beta + |\Gamma \tilde{A}_0| - N_{2k+1,k+2} = \binom{2k}{k+1} + 2\beta + |\Gamma \tilde{A}_0|. \end{aligned} \quad (2.11)$$

Let us bound now $2\beta + |\Gamma \tilde{A}_0|$. By Theorem H we can write $\beta = \binom{x}{k}$, where $k \leq x < 2k$, and bound $|\Gamma \tilde{A}_0| \geq \binom{x}{k-1} + \binom{2k}{k} - \binom{x}{k}$. Hence $2\beta + |\Gamma \tilde{A}_0| \geq 2\binom{x}{k} + \binom{x}{k-1} + \binom{2k}{k} - \binom{x}{k} \geq \beta + \binom{2k}{k}$. This together with (2.11) implies

$$|\Gamma A| > \binom{2k}{k+1} + \beta + \binom{2k}{k} \geq \binom{2k+1}{k} + 2, \text{ a contradiction.}$$

Case $n = 2k$: We have $|A| \leq N_{2k,k+2} + \alpha$, where $\alpha < \binom{2k}{k+1}$, and by the induction hypothesis $|A_0| > N_{2k-1,k}$. Now (2.10) implies $|\Gamma A| \geq 2|A_0| + |\Gamma \tilde{A}_0| - N_{2k,k+2} - \alpha$. If $|A_0| \geq N_{2k-1,k-1}$, then $|\Gamma A| > N_{2k,k} + \binom{2k-1}{k-1} - N_{2k,k+1} = \binom{2k}{k} + \binom{2k-1}{k-1}$, a contradiction. Let now $|A_0| = N_{2k-1,k} + \beta$ with $1 \leq \beta < \binom{2k-1}{k-1}$. Then $|\Gamma A| > 2N_{2k-1,k} + 2\beta + |\Gamma \tilde{A}_0| - N_{2k,k+1} = \binom{2k-1}{k} + 2\beta + |\Gamma \tilde{A}_0|$.

Letting $\beta = \binom{x}{k-1}$, where $1 \leq x < 2k - 1$, we get

$$|\Gamma A| > \binom{2k-1}{k} + 2\binom{x}{k-1} + \binom{x}{k-2} + \binom{2k-1}{k-1} - \binom{x}{k-1} \geq \binom{2k-1}{k} + \beta + \binom{2k-1}{k-1} = \beta + \binom{2k}{k},$$

again a contradiction. This completes the proof of part (i). \square

Proof (ii) Consider first the case $n \geq 2k + 1$. As above, clearly we have $\lambda(n, N_{n,k+1}) \leq \binom{n}{k}$. The proof now goes along the same lines as for part (i). The only difference is that instead of the induction hypothesis we now use directly the result in (i).

The case $n = 2k$ requires an additional argument. Note first that $\lambda(2k, N_{2k,k+1}) \leq \binom{2k}{k} + 1$. This can be seen by removing the vertices of weight k and any vertex

from each of the two connected components. Suppose now the largest connected component A , after removal of $\binom{2k}{k} + 1$ vertices, has size $|A| < N_{2k,k+1}$. Define also A_0 and A_1 as above. Let us show then that $|A| \leq N_{2k,k+2} + \binom{2k-1}{k+1}$. Assuming the opposite and using (2.3) observe that for $k \geq 3$ and for some $k \leq x < 2k - 1$ we have

$$|\Gamma A| \geq \binom{2k-1}{k} + \binom{x}{k-1} + \binom{2k}{k+1} - \binom{2k-1}{k+1} - \binom{x}{k} > \binom{2k}{k} + 1,$$

a contradiction.

Case 1: $|A| = N_{2k,k+2} + \binom{2k-1}{k+1}$.

Denote $F = \mathcal{H}_n \setminus (A \cup S)$ and observe that $|F| = N_{2k,k+1} + \binom{2k-1}{k} - 1$. Then by Theorem H

$$|\Gamma F| \geq \binom{2k}{k} + 1 = |S|.$$

This means that F is optimal and by Theorem A it contains a ball of size $N_{2k,k+1}$. The latter clearly implies that the largest connected component has size at least $N_{2k,k+1}$, a contradiction.

Case 2:

$$|A| = N_{2k,k+2} + \binom{2k-1}{k+1} - 1. \quad (2.12)$$

Then it is easy to see that $|\Gamma A| \geq \binom{2k}{k} + 1 = |S|$. Hence A is an optimal set and it contains a ball of size $N_{2k,k+2}$ (w.l.o.g. we may assume that A contains the ball centered at $(1, \dots, 1)$). This together with (2.12) implies that

$$|A_0| < N_{2k-1,k+2} + \binom{2k-1}{k+1} < N_{2k-1,k+1}, \quad |A_1| < N_{2k-1,k+1} + \binom{2k-1}{k+1} < N_{2k-1,k}.$$

Suppose now that $|S_0| \leq \binom{2k-1}{k-1}$ (resp. $|S_1| \leq \binom{2k-1}{k-1}$). This (by our previous observation) implies $|A_0| \geq N_{2k-1,k}$ (resp. $|A_1| \geq N_{2k-1,k}$), a contradiction.

Case 3: $|A| = |A_0| + |A_1| \leq N_{2k,k+2} + \binom{2k-1}{k+1} - 2$.

Assuming (w.l.o.g.) that $|S_0| \leq \binom{2k-1}{k-1}$ and hence $|A_0| \geq N_{2k-1,k}$ we get

$$|S| \geq |\Gamma A| \geq |A_0| - |A_1| \geq 2N_{2k-1,k} - N_{2k,k+2} - \binom{2k-1}{k+1} + 2 \geq \binom{2k}{k} + 2,$$

a contradiction which completes the proof of part (ii). \square

Next we give upper bounds for the function $\lambda(n, c)$.

We describe a regular separation of the vertices of the Hamming graph $G(\mathcal{H}_n)$. For convenience of the description, we identify \mathcal{H}_n with the set of vectors $\mathcal{H}_n^* := \{-1, 1\}^n \subset \mathbb{R}^n$ using $1 \rightarrow -1$ and $0 \rightarrow 1$. This is an exchange of the coordinates in elements of \mathcal{H}_n . Thus we can speak about an identical graph $G(\mathcal{H}_n^*)$. Note that the Hamming distance between any $x, y \in \mathcal{H}_n^*$ can be evaluated by their inner product $\langle x, y \rangle$, that is, $d(x, y) = \frac{1}{2}(n - \langle x, y \rangle)$.

The idea is to separate the elements of \mathcal{H}_n^* into equal sized parts by mutually orthogonal hyperplanes of \mathbb{R}^n . It is known that for any $n = 2^k$ there exist Hadamard matrices of order n . Recall that a $(+1, -1)$ -matrix H of size $n \times n$ is called a Hadamard matrix of order n , if $HH^T = nI_n$. Hadamard matrices H_n of order $n = 2^k$ can be constructed as k -th Kronecker power of matrix

$$H_2 = \begin{bmatrix} 1 & 1 \\ 1 & -1 \end{bmatrix}. \text{ Note that the corresponding } (0,1)\text{-matrix without all-zero}$$

column can be viewed as the simplex code of length $2^k - 1$ (well known in Coding Theory [11]) with a generator matrix of size $k \times 2^k - 1$ consisting of all-nonzero column vectors.

Let $\langle v_1 \rangle, \dots, \langle v_n \rangle$ be the hyperplanes defined by $\langle v_i \rangle = \{x : \langle v_i, x \rangle = 0\}$, $i = 1, \dots, n$. Given an integer $1 \leq r \leq n$ let us define the set of *sign sequences* $\Sigma := \{+, -\}^r$. Let $x \in \mathcal{H}_n^*$ and let $(\sigma_1, \dots, \sigma_r) \in \Sigma$. We say that $Sign(x) = (\sigma_1, \dots, \sigma_r)$ if $Sign\langle x, v_i \rangle = \sigma_i$, $i = 1, \dots, r$, (where for a real number a , like $\langle x, v_i \rangle$, $Sign a$ is defined in the natural way). Let $\Sigma_1, \dots, \Sigma_{2^r}$ be the elements of Σ in some fixed order. Define the sets $B_i = \{x \in \mathcal{H}_n^* : Sign(x) = \Sigma_i\}$; $i = 1, \dots, 2^r$. Clearly these sets are disjoint. Denote the set of remaining elements of \mathcal{H}_n^* by S_r , that is, $S_r = \{x \in \langle v_i \rangle \cap \mathcal{H}_n^* : 1 \leq i \leq r\}$. The hyperplanes $\langle v_1 \rangle, \dots, \langle v_r \rangle$ separate the points of \mathbb{R}^n into classes which have different signs. Therefore we have the following.

Lemma 1 S_r is a vertex separating set for B_1, \dots, B_{2^r} , that is, any path between the vertices of two distinct classes B_i and B_j contains a vertex of S_r .

Theorem 2 Given integers $n = 2^k$ and $1 \leq r \leq k$, we have

$$\lambda(n, 2^{n-r} - |S_r|/2^r) \leq |S_r|. \quad (2.13)$$

Proof Let v_1, \dots, v_k be the row vectors of a submatrix G of H_n such that all columns of G are distinct.

Lemma 2 For every B_i and B_j ($1 \leq i, j \leq 2^r$) there exists a permutation π of the coordinates such that $\pi(B_i) = B_j$.

Proof Let us write the vectors of \mathcal{H}_n^* as concatenation $(x|y)$ of vectors $x, y \in \mathcal{H}_{n/2}^*$. We use the following property of G .

Claim: For any given v_t ($1 \leq t \leq k$) there exists a permutation π of coordinates (columns) such that for the resulting vectors $\pi(v_1), \dots, \pi(v_k)$ we have $\pi(v_t) = (1^{n/2} | -1^{n/2})$ and all other vectors $\pi(v_i)$ ($i \neq t; 1 \leq i \leq k$) are of the form $(u_i | u_i)$ (where $1^{n/2}$ and $-1^{n/2}$ are all-one and all-minus one vector).

This fact is obvious. Imagine replacing v_t to the last position and then put the column vectors in colexicographic order.

Suppose now Σ_i and Σ_j are two sign sequences which differ only in the t -th coordinate. In view of the claim, we may assume that $v_t = (1^{n/2} | -1^{n/2})$ and all other vectors v_l ($1 \leq l \leq k$) are of the form $(u_i | u_i)$. It is easy to see that, if $(x|y) \in B_i$, then $(y|x) \in B_j$. The latter (with the claim above) implies that there exists a permutation π such that $\pi(B_i) = B_j$. This argument can be successively applied for all sets B_i . \square

Lemma 2 implies that $|B_i| = |B_j|$ for $1 \leq i, j \leq 2^r$ and hence $|B_i| = (2^n - |S_r|)/2^r$. This clearly completes the proof of Theorem 2. \square

Corollary 1 For $n = 2^k$ and $1 \leq r \leq k$

$$\lambda(n, 2^{n-r}) < r \binom{n}{n/2} = \Theta(r2^n / \sqrt{n}). \quad (2.14)$$

Proof By definition, S_r consists of those vectors of \mathcal{H}_n^* which are orthogonal to one of v_1, \dots, v_r . Since there are exactly $\binom{n}{n/2}$ vectors orthogonal to any vector $v \in \mathcal{H}_n^*$, we conclude that $|S_r| < r \binom{n}{n/2}$. \square

Remark 2 Corollary 1 can be extended to arbitrary n . Partition the coordinate set $[n]$ into two parts of size $n_1 = 2^{\lfloor \log n \rfloor}$ and $n_2 = n - n_1$. Observe now that defining all sets (in separation of \mathcal{H}_{n_1}) as above for part $[n_1]$ and taking all their 2^{n_2} extensions in part $[n_2]$ we get that $\lambda(n, 2^{n-r}) = O(r2^n / \sqrt{n})$, where $1 \leq r \leq \lfloor \log n \rfloor$.

Conjecture For $n = 2^k$ and $1 \leq r \leq k$

$$\lambda(n, 2^{n-r} - |S_r|/2^r) = |S_r|. \quad (2.15)$$

Note that (in view of Theorem 1) the conjecture holds for $r = 1$.

3 Diagnosability of the n -cube

Theorem 1 implies: if the number of faulty vertices $\binom{n}{k-1} < t \leq \binom{n}{k}$, ($k \leq n/2$), then there exists a set of vertices $A \subset \mathcal{H}_n$ with $|A| \geq N_{n,k+1}$ that can be identified as “fault-free” and the vertices ΓA can be identified as “faulty”.

Thus the status of at least $|\sigma(A)|=|A \cup \Gamma A|$ elements can be identified in one step. This, together with Theorem H, implies

Corollary 2 (i) *Let t be the number of faulty vertices and let $\binom{n}{k-1} < t \leq \binom{n}{k}$, $k \leq n/2$. Then the status of at least $N_{n,k}$ vertices can be identified in one step. In particular, for $k = n/2$, the status of at least $N_{n,n/2} = 2^{n-1} + \binom{n-1}{\frac{n}{2}-1}$ vertices can be identified.*

(ii) (Yamada et al. [16]) *Given integer $n \geq 3$ we have $m(\mathcal{H}_n) \geq \binom{n}{\lfloor \frac{n}{2} \rfloor}$ and hence the degree of sequential diagnosability of the n -cube $t(\mathcal{H}_n) > \binom{n}{\lfloor \frac{n}{2} \rfloor}$*

An important parameter in sequential diagnosis is the number of tests and repair iterations needed to locate all the faulty units within the system (see [5], [15]). Thus, reducing the number of iterations is an important task in implementation of a diagnosis scheme. It was shown in [15] that this number for n -cubes is upper bounded by $\Theta(n)$. As a direct consequence of Theorem 1 we get

Corollary 3 *Let $\binom{n}{k-1} < t \leq \binom{n}{k}$, $k \leq n/2$, then the number of iterations needed for sequential diagnosis is at most k .*

The next upper bound for the sequential diagnosability of the n -cube is obtained by Yamada et al. [16]. Let us give a short proof for this result using Theorem 2.

Theorem 3 [16] $t(\mathcal{H}_n) = O(2^n \log n / \sqrt{n})$.

Proof As noticed in [16], it is sufficient to prove the theorem for the case when n is a power of 2. Let $n = 2^k$ and let B_1, \dots, B_{2^r} and S_r be the sets described in Theorem 2. Define the sets $S_r^* = S_r \cup \Gamma S_r$ and $B_i^* = B_i \setminus \Gamma S_r$ ($i = 1, \dots, r$). Thus B_i^* is the set of inner points of B_i (that is $B_i^* \cup \Gamma B_i^* = B_i$). Note that Lemma 2 immediately implies that $|B_i^*| = |B_j^*|$ for all $1 \leq i, j \leq 2^r$ and hence $|B_i^*| = 2^{n-r} - |S_r^*|/2^r$.

For $v \in \mathcal{H}_n^*$ let $S(v)$ denote the set of vertices (of \mathcal{H}_n^*) orthogonal to v (equivalently at Hamming distance $n/2$ from v). We know that $|S(v)| = \binom{n}{n/2}$ and the size of the set of the boundary vertices $\Gamma S(v)$ does not depend on the choice of v (the vertices of any $S(u)$ can be obtained from $S(v)$ by a translate in the Hamming space). Thus, taking say $v = 1^n$, we get $|\Gamma S(v)| = \#\{\text{boundary of the vertices of weight } n/2\} = 2\binom{n}{\frac{n}{2}-1}$. Hence $|S_r^*| = |S_r| + |\Gamma S(v)| < 3r\binom{n}{n/2}$.

Define now the following syndrome: all inner edges of $B_1^*, \dots, B_{2^r}^*$ (i.e. the edges of the corresponding induced subgraph) are labelled by 0's, all remaining edges are labelled by 1's. For each $x \in \mathcal{H}_n$ we define two fault sets $F(x)$ and $F(\bar{x})$ so that $x \in F(x)$ and $x \notin F(\bar{x})$.

If $x \in S_r$, then $F(x) := S_r^*$ and $F(\bar{x}) := S_r^* \setminus \{x\}$. The vertices of S_r have no neighbors from B_i^* , therefore $F(x)$ and $F(\bar{x})$ are consistent fault sets.

If $x \in (\Gamma S_r \cap B_i)$, then $F(x) := S_r^*$ and $F(\bar{x}) := (S_r^* \setminus \{x\}) \cup B_i^*$, $i = 1, \dots, 2^r$. By definition of S_r^* , no vertex of S_r^* has neighbors from two distinct classes B_i^* and B_j^* . Therefore $F(x)$ and $F(\bar{x})$ are consistent fault sets.

Finally, if $x \in B_i^*$, then $F(x) := S_r^* \cup B_i^*$, $F(\bar{x}) := S_r^*$ and obviously $F(x)$ and $F(\bar{x})$ are consistent fault sets.

Observe now that each fault set has size not greater than $t := |S_r^*| + |B_i^*| < 3r \binom{n}{\frac{n}{2}} + |B_i^*|$. This together with Fact 1 implies that \mathcal{H}_n is not sequentially t -diagnosable. If now $r = k$, then $|B_i^*| = (2^n - |S_r^*|)/n$ and hence $t < 3 \log n \binom{n}{\frac{n}{2}} + 2^n/n = O(\frac{2^n \log n}{\sqrt{n}})$. \square

4 Diagnosis Algorithm

Given a connected graph $G = (V, E)$ and a syndrome, that is, a diagnostic graph $DG = (V, E')$, where each edge of E is labelled by a 0 or 1.

Step 1. Partition the vertices of DG into 0-components K_1, \dots, K_ℓ . Denote $\mathcal{K} = \{K_1, \dots, K_\ell\}$.

Step 2. Construct the contracted graph $G_c = (V_c, E_c)$ as follows.

Each component K_i contracts to vertex $a_i \in V_c$ and $\{a_i, a_j\} \in E_c$ if and only if there is an edge $\{u, v\}$ (labelled with 1) in E with $u \in K_i$ and $v \in K_j$. To each vertex a_i of V_c assign the weight $wt(a_i) = |K_i|$. Thus G_c is an undirected graph with weights on vertices. Clearly $\sum_{a \in V_c} wt(a) = |V|$. The weight of a subgraph $G' \subset G_c$ is defined by $wt(G') = \sum_{b \in V'} wt(b)$, where V' is the vertex set of G' .

Step 3. Find a spanning tree TG_c of G_c .

Step 4. Partition the vertex set of TG_c into subsets T_1, \dots, T_p , each containing at least two vertices, such that the induced subgraph of each subset forms a star. Let z_i denote the center of the star S_i , $i = 1, \dots, p$. Define also

$$w_i = \min\{wt(z_i), wt(S_i \setminus \{z_i\})\}, \quad \alpha_i = \max\{wt(z_i), wt(S_i \setminus \{z_i\})\},$$

$$\bar{w} = w_1 + \dots + w_p, \quad \bar{\alpha} = \alpha_1 + \dots + \alpha_p.$$

Step 5. Determine $\Delta = \max_{1 \leq i \leq p} \{\alpha_i + \bar{w} - w_i\}$. Suppose $\Delta = \alpha_r + \bar{w} - w_r$. Suppose also the number of actual faults $t \leq \Delta - 1$.

Step 6. If $wt(z_r) = w_r$, then the vertex z_r is labelled as “faulty”. The component $K_{i_r} \in \mathcal{K}$ corresponding to z_r is diagnosed as faulty set.

If $wt(z_r) = \alpha_r$, then z_r is labelled as “non-faulty” and the remaining vertices of S_r are labelled as “faulty”. The components corresponding to the vertices $S \setminus \{z_i\}$ are diagnosed as faulty sets.

4.1 The algorithm works

Only the Step 6 requires to argue about. Observe that the faulty vertices (vertices of the original graph) corresponding to each star S_i have size at least w_i . Therefore z_r is faulty, if $wt(z_r) = w_r$, otherwise the remaining vertices of S_r are faulty and hence the total number of faults, corresponding to the given syndrome, is at least $\bar{w} - w_r + \alpha_r = \Delta > t$ (# actual faults), a contradiction. Similarly z_r is non-faulty, if $wt(z_r) = \alpha_r$, and consequently all other vertices of S_r are faulty. Thus the algorithm allows to identify the status of at least one vertex. The status of the remaining vertices can be identified iteratively applying the “diagnosis and repair” procedure ([14]). Furthermore this can be done applying at most $d(G)$ (= diameter of G) iterations.

4.2 Complexity

Let a connected graph $G = (V, E)$ and a diagnostic graph DG be given. By using depth first search or breadth first search one can verify that the complexity of each of the following problems is $O(|E|)$: Location of 0-components, Construction of the contracted graph G_c , Finding a spanning tree of a connected graph G_c , Partition of a tree into stars. Thus the overall complexity of the algorithm is $O(|E|)$.

4.3 Diagnosability of the algorithm

We have seen before that the algorithm works if the actual number of faults $t < \Delta$. Let us denote $\Delta^*(t, G) = \min \Delta(t, G)$, where the minimum is taken over all syndromes produced by all faulty sets F with $|F| \leq t$. Define also $\delta(G) = \max\{t : \Delta^*(t, G) > t\}$. Then, clearly, the degree of diagnosability $t(G) \geq \delta(G) - 1$. Thus the degree of diagnosability of our algorithm $t_A(G)$ can be estimated by giving lower bounds for $\delta(G)$. Note that $\delta(G) > m(G)$. This follows directly from the definition of Δ . Let us bound now $\Delta = \max_{1 \leq i \leq p} \{\alpha_i + \bar{w} - w_i\}$ from below. It is easy to see that the following is true.

Claim *The quantity Δ attains the minimum if $wt(z_i) = 1$, for all $i = 1, \dots, p$, that is, when the centers of all stars have weight one.*

Let now $wt(z_i) = 1$, $i = 1, \dots, p$ and let $\beta = \max\{\alpha_i\}$. Then we have $\Delta \geq$

$\beta + p - 1 \geq \beta + N/(\beta + 1) - 1$. Hence

$$\delta(G) \geq \min_{1 \leq \beta \leq N} \left(\beta + \frac{N}{\beta + 1} - 1 \right) = 2(\sqrt{N} - 1). \quad (4.1)$$

Thus we get $\delta(G) \geq \max\{m(G) + 1, 2\sqrt{N} - 2\}$. Therefore given a lower bound m^* for $m(G)$, the diagnosability of the algorithm

$$t_A(G) \geq \max\{m^*, 2\sqrt{N} - 3\}. \quad (4.2)$$

In particular this implies the following

Corollary 4 *For an arbitrary interconnection graph on N vertices the diagnosability of our algorithm $t_A(G) \geq \lceil 2N^{\frac{1}{2}} \rceil - 3$.*

In fact, the algorithm is optimal for “bad graphs”: there exist connected graphs on N vertices with sequential diagnosability $\lceil 2N^{\frac{1}{2}} \rceil - 3$.

Example Let $k = N^{\frac{1}{2}}$ be an integer and let DG be a diagnostic graph on N vertices shown in Figure 1.

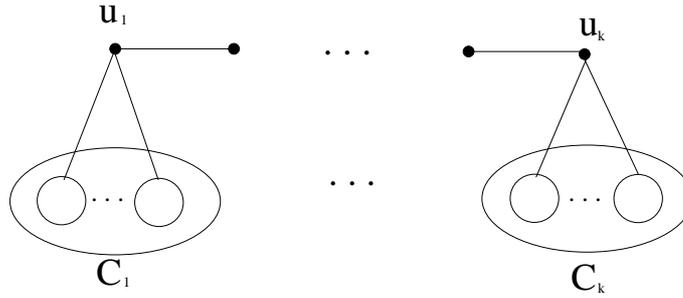


Fig. 1.

DG: each set of vertices C_i with $|C_i| = k - 1$, ($i = 1, \dots, k$) represents a union of some 0-components (denoted by circles), where the edges incident with vertices u_1, \dots, u_k are labelled by 1's. Let us denote $U = \{u_1, \dots, u_k\}$. Define then the faulty sets F_1, \dots, F_k as $F_i = (U \setminus \{u_i\}) \cup C_i$, $i = 1, \dots, k$. Note that $|F_i| = 2k - 2$. All these sets are consistent fault sets (their intersection is empty and the union is the vertex set of G). Therefore, G is not sequentially $(2N^{\frac{1}{2}} - 2)$ -diagnosable. More generally, one can show that for any integer N there are connected graphs on N vertices and maximal degree $k \leq N^{\frac{1}{2}}$ with sequential diagnosability $\lceil 2N^{\frac{1}{2}} \rceil - 3$. In particular there are such k -trees.

References

- [1] R. Ahlswede and G.O.H. Katona, Contributions to the geometry of Hamming spaces, Discrete Math. 17, 1-22, 1977.

- [2] R. Ahlswede and K.U. Koschnick, Note on an extremal problem arising for unreliable networks in parallel computing, *Discrete Math.* 47, 137-152, 1983.
- [3] L.A. Aslanian, An isoperimetric problem and related extremal problems for discrete spaces (in Russian), *Problemy Kibernetiki* 36, 85-127, 1979.
- [4] M. Barborak, M. Malek, and A. Dahbura, The consensus problem in fault-tolerant computing, *ACM Computing Surveys*, Vol. 25, No. 2, 171-220, 1993.
- [5] A. Caruso, S. Chessa, P. Maestrini, and P. Santi, Diagnosability of regular systems, *J. Algorithms* 45, 126-143, 2002.
- [6] P. Frankl, A lower bound on the size of a complex generated by an antichain, *Discrete Math.* 76 (1), 51-56, 1989.
- [7] L.H. Harper, Optimal numberings and isoperimetric problems on graphs, *J. Combin. Theory* 1, 385-395, 1966.
- [8] G.O.H. Katona, The Hamming sphere has minimum boundary, *Studia Scient. Math. Hungarica* 10, 131-140, 1975.
- [9] S. Khanna and W.K. Fuchs, A graph partitioning approach to sequential diagnosis, *IEEE Trans. Comput.* 46 (1), 39-47, 1996.
- [10] D.J. Kleitman, On a problem of Yuzvinsky on separating the n -cube, *Discrete Math.* 60, 207-213, 1986.
- [11] F.J. MacWilliams and N.J.A. Sloane, *The Theory of Error Correcting Codes*, North Holland Mathematical Library, 1977.
- [12] F.P. Preparata, G. Metze, and R.T. Chien, On the connection assignment problem of diagnosable systems, *IEEE Trans. Comput.* EC 16 (12), 848-854, 1967.
- [13] V. Raghavan and A. Tripathi, Sequential diagnosability is co-NP complete, *IEEE Trans. Comput.* 40 (5), 584-595, 1991.
- [14] J.D. Russell and C.R. Kime, System fault diagnosis: closure and diagnosability with repair, *IEEE Trans. Comput.* C-24, 1078-1089, 1975.
- [15] P. Santi and S. Chessa, Reducing the number of sequential diagnosis iterations in hypercubes, *IEEE Trans. Comput.* 53, 89-92, 2004.
- [16] T. Yamada, T. Otsuka, A. Watanabe, and S. Ueno, On sequential diagnosis of multiprocessor systems, *Discrete Appl. Math.* 146, 311-342, 2005.