Reducing the Number of Sequential Diagnosis Iterations in Hypercubes

Paolo Santi
Istituto di Informatica e Telematica
Via G. Moruzzi, 1
56124, Pisa – ITALY
paolo.santi@iit.cnr.it

Stefano Chessa
Dip. di Informatica
Via F. Buonarroti, 2
56127, Pisa – ITALY
and
ISTI – CNR
Via G.Moruzzi, 1
56124, Pisa – ITALY
ste@di.unipi.it

Abstract

In this note, we use a vertex-isoperimetric inequality to show that the number of test and repair iterations needed to perform sequential diagnosis of d-dimensional hypercubes is upper bounded by d-r, where $r \in \Theta(d)$. This result improves the best bound of d test and repair iterations previously known. Numerical evaluation has shown that the actual value of r ranges from 0.16d to 0.31d.

Index terms. Massively parallel systems, system-level diagnosis, sequential diagnosis, hypercubes.

1. Introduction

System-level diagnosis was introduced by Preparata, Metze and Chien [14] as a mean to diagnose systems composed by units (usually processors) connected by point-to-point, bidirectional links. A system is represented as an undirected graph G=(V,E), where node set V (with $|V|=n$) represents the system units and E represents the interconnection links. Edge (u,v) belongs to E iff u and v are interconnected, in which case u and v are said to be *adjacent*. The diagnosis in the PMC model in based on the outcomes of tests performed between adjacent units. Each test involves two units, called the *tester* and the *tested* unit, and is performed as follows. The tester unit u sends a test sequence to the tested unit v, which executes the test and returns the result to u. Unit u compares the received and the expected result and generates the test outcome, which is 0 if the test passes and 1 if it fails. If the tester unit is faulted, the test outcome is arbitrary. This *invalidation rule* is summarized in Table 1. The collection of the test results is represented by the labeled directed graph $G_S=(V,E_S)$, called the *syndrome graph*, where $(u,v)\in E_S$ if and only if unit u tests unit v. Every edge is labeled with the correspondent test outcome.

Table 1. Invalidation rule of the PMC model.

Testing unit	Tested unit	Test outcome
Fault-free	Fault-free	0
Fault-free	Faulty	1
Faulty	Fault-free	0 or 1
Faulty	Faulty	0 or 1

Given any set $V_f\subset V$ of faulty units (*actual fault set*), the set of all test outcomes is called *syndrome*, denoted σ. In the PMC approach the collection of the tests is determined prior to the execution of the diagnosis algorithm. Once all the tests have been executed, the syndrome is collected and decoded by an external, reliable computer, called *diagnoser*.

In [14], Preparata et al. introduced the concepts of *one-step* and *sequential diagnosis*. In the former approach it is required that syndrome decoding correctly identifies the status of all the units. In the sequential diagnosis approach, the diagnosis process is decomposed into several diagnosis and repair iterations. The goal of each iteration is to identify at least one faulty unit, which is immediately repaired or replaced, thus reducing the number of faulty units in the system. The process is iterated until all the faulty units have been repaired or replaced.

Given a system S, the *one-step diagnosability* is a parameter tied to the structure of the syndrome graph of S, which express the maximum number of faults in the system under which is always possible to produce a one-step diagnosis of the system. A similar parameter for sequential diagnosis is called *sequential diagnosability* (or, briefly, *diagnosability*) and is usually far above the one-step diagnosability. One-step diagnosable systems have been characterized in [5], while sequentially diagnosable systems have been characterized in [5] and [13].

The one-step diagnosability of a system is limited above by the minimum of the node in-degrees in G_S [5, 14]. For this reason, the one-step diagnosis approach is inadequate to the case of large systems based on regular interconnection structures, such as hypercubes. Conversely, sequential diagnosis of regular or quasi-regular systems is feasible under more realistic fault situations. Lower bounds to the sequential diagnosability of hypercubes have been provided in [2,8,9,10].

Another parameter usually considered in the evaluation of a sequential diagnosis strategy is the number of test and repair iterations needed to locate all the faulty units within the system. The number of iterations along with the complexity of the sequential diagnosis algorithm provide a measure of the time (i.e., of the computational overhead) needed to achieve a complete diagnosis of the system. Thus, reducing the complexity of the sequential diagnosis algorithm and the number of iterations is particularly important to implement a fast and efficient sequential diagnosis scheme. Furthermore, a short time-to-diagnosis is required to validate two common assumptions of sequential diagnosis, i.e., that faults are permanent and that no new fault occurs during the diagnosis process.

Sequential diagnosis of hypercubes has been addressed in several recent papers.

In [8], Kavianpour and Kim presented two strategies for the sequential diagnosis of hypercubes. The first strategy performs sequential diagnosis in at most *d* iterations, and has diagnosability $t \in \Theta\left(\sqrt{d \cdot n}\right)$. The second strategy differs slightly from the sequential diagnosis approach since it allows that the set of processors replaced at every iterations may possibly include, along with faulty units, at most one fault-free unit. Under this hypothesis the diagnosability is increased to *n*-2, and the number of iterations is still upper bounded by *d*.

In [10], Khanna and Fuchs introduced a cluster-based sequential diagnosis algorithm for hypercubes. The algorithm has diagnosability $t \in \Theta\left(\sqrt{d \cdot n}\right)$, and completes the diagnosis in $s \in O(d)$ test and repair iterations, where *s* can be larger than *d* in the worst-case. The same authors introduced in [9] the PARTITION algorithm for the sequential diagnosis of regular graphs. PARTITON is based on a graph partitioning, and has diagnosability $t \in \Omega\left(\frac{n \log d}{d}\right)$ when applied to hypercubes. The number of iterations needed by PARTITION to complete diagnosis is *d* in the worst-case.

The diagnosis of hypercubes has also been addressed using other approaches, such as *adaptive diagnosis*. In adaptive diagnosis, tests can be scheduled dynamically during the diagnosis process based on the outcomes of the tests performed so far, and the goal is to minimize the total number of tests needed to diagnose all the units in the system. Recent results along this trend of research can be found in [1,4,11].

In this paper, we exploit a known vertex-isoperimetric inequality for binary hypercubes to show that a straightforward modification of PARTITION allows diagnosis completion in at most *d*-*r* iterations, where $r \in \Theta(d)$. Our improved algorithm, called *i*-PARTITION, has the same diagnosability and time complexity of PARTITION, i.e., $t \in \Omega\left(\frac{n \log d}{d}\right)$ and O(*nd*), respectively. We remark that *i*-PARTITION improves the best upper bound *d* on the number of iterations known so far. The value of *r* for hypercubes of different dimensions has been determined through numerical evaluation, which has shown that the actual reduction on the number of iterations is about 30%.

The performances of the algorithms for sequential diagnosis of hypercubes proposed so far, as well as the performance of *i*-PARTITION, are summarized in Table 2.

Table 2. Performance of various sequential diagnosis algorithms for hypercubes.

Algorithm	Diagnosability	#iterations
KavianpourKim	$\Theta\left(\sqrt{d \cdot n}\right)$	$\leq d$
KhannaFuchs	$\Theta\left(\sqrt{d \cdot n}\right)$	O(d)
PARTITION	$\Omega\left(\frac{n \log d}{d}\right)$	$\leq d$
i-PARTITION	$\Omega\left(\frac{n \log d}{d}\right)$	$\leq d$-r

2. Vertex-isoperimetric inequalities

Let $G=(V,E)$ be a graph, with $|V|=n$. Given any $u,v\in V$, the *distance* between u and v is the length of the shortest path from u to v, and is denoted $d(u,v)$. Observe that if G is undirected, then $d(u,v)=d(v,u)$, and $d(u,v)$ can be used as a metric on G. Given any $A\subseteq V$, let $d(A,v)=\min\{d(u,v):u\in A\}$. Observe that $d(A,v)=0$ if and only if $v\in A$.

Def (vertex boundary):

Given any $A\subseteq V$, the vertex boundary of A, denoted ∂A, is defined as the set of vertices at distance at most 1 from A. Formally, $\partial A=\{v\in V: d(A,v)\leq 1\}$.

Observe that the vertex boundary of a set includes the set itself.

Def (vertex-isoperimetric inequality):

A vertex-isoperimetric inequality for a graph $G=(V,E)$ is defined by a function $g(m)$ such that $|\partial A|\geq g(m)$ for any $A\subseteq V$ with $|A|=m$.

Def (binary hypercube):

A binary hypercube of dimension d is composed by $n=2^d$ units. Every unit u is labeled with a d-digits binary number denoted lab(u). Units are connected based on the Hamming distance of their labels, denoted d_H: edge (u,v) exists if and only if $d_H(lab(u),lab(v))=1$.

In the following, we use the term hypercube to refer to a binary hypercube, and we use only logarithms to the base 2.

From the definition above, it is immediate that hypercubes are d-regular structures, and have diameter d.

In general many vertex-isoperimetric inequalities for a given graph G can be defined. A vertex-isoperimetric inequality for hypercubes has been derived in [7,12], and is stated in the following Theorem.

Theorem 1.

Let G be the d-dimensional hypercube, and let $A \subseteq V$, with $|A|=m$. Then $|\partial A| \geq \sum_{i=0}^{r+1} \binom{d}{i}$, *where*

$$r = \max\left\{ k \in \{1,\dots,d\} : \sum_{i=0}^{k} \binom{d}{i} \leq m \right\}.$$

3. The improved PARTITION algorithm

The PARTITION algorithm introduced in [9] is composed of two phases, which we report for completeness:

– *Phase 1: Fault-Free Subset Identification*

The goal of this phase is to identify a subset of fault-free units. Each unit tests each one of its neighbors. The outcomes of these tests are used to construct the syndrome graph G_S. Let G_P be the subgraph of G_S induced by the edge set $E_{00}=\{(u,v)\in E_S$: (u,v) and (v,u) are labeled 0$\}$. Locate a connected component of size at least t+1 in G_P, where t is properly defined. All the units in this component are diagnosed fault-free.

– *Phase 2: Iterative Diagnosis and Repair*

The goal of this phase is to iteratively diagnose and repair faulty units. Select an arbitrary unit, say u, diagnosed as fault-free in the previous step of the algorithm, and construct a breadth-first search tree of G rooted at u. Let h denote the height of the tree, and let L_i, $0\leq i\leq h$, be the set of units at distance i from u. Starting from the top of the tree, units in L_i are used to diagnose units in L_{i+1}. At each step, units diagnosed as faulty are repaired. At step h, all the faulty units have been repaired.

In [9] it is shown that if t is the largest integer satisfying inequality $\phi_G(x+1)\geq x$, where $\phi_G(x)$ is the x-partition number of G,[1] then a connected component of G_P of size at least t+1 must exist, and it must be composed by fault-free units. Hence, diagnosis performed by PARTITION is correct. The diagnosability of the algorithm is t, which is proved to be $\Omega\left(\frac{n \log d}{d}\right)$. Furthermore, since $h\leq d$, the number of iterations needed to complete diagnosis is d in the worst-case.

As suggested by the authors themselves, the number of iterations needed to perform diagnosis with PARTITION can be reduced by using a collection of breadth-first search trees, instead of a single tree, to diagnose and repair faulty units. However, this solution would leave the worst-case upper bound of d iterations of sequential diagnosis unchanged. In the following, we propose a different implementation of the second phase of PARTITION, which requires at most d-r iterations to complete diagnosis, where $r\in\Theta(d)$. We call this improved implementation of the algorithm *i*-PARTITION.

Algorithm *i*-PARTITION

- *Phase 1:* as in the PARTITION algorithm
- *Phase 2-improved: Iterative Diagnosis and Repair*

  Let A_1 be the node set of the connected component of G_P identified in the Phase 1. In the first step, all the units in $\partial A_1 - A_1$ are diagnosed by units in A_1, and the units identified as faulty are repaired. At step j, units in $A_j=\partial A_{j-1}$ diagnose units in $\partial A_j - A_j$. The process is iterated until $A_j=V$.

Theorem 2.

Algorithm i-PARTITION:

a) *has diagnosability* $t\in\Omega\left(\frac{n \log d}{d}\right)$;

b) *has time complexity O(nd);*

c) *requires at most the same number of test and repair iterations as algorithm PARTITION to complete diagnosis.*

Proof.

[1] See [9] for the definition of x-partition number.

Properties a) and b) follow immediately by the definition of *i*-PARTITION and by [9]. Property c) is proved by observing, that, for every i, set L_i as defined in algorithm PARTITION is contained in set ∂A_i. ❑

4. Upper bound on the number of iterations of *i*-PARTITION

In this section, we derive an upper bound to the number of iterations needed by *i*-PARTITION to complete diagnosis using the vertex-isoperimetric inequality for hypercubes of Theorem 1.

Theorem 3.

Algorithm i-PARTITION completes diagnosis in at most d-r steps, where

$$r = \max\left\{ k \in \{1,..,d\} : \sum_{i=0}^{k} \binom{d}{i} \le t+1 \right\}.$$

Proof.

Let A_1 be the connected component identified in step 1 of the *i*-PARTITION algorithm; component A_1 is correctly diagnosed as fault-free in the first phase of *i*-PARTITION, and $|A_1| \geq t+1$. By Theorem 1, it follows that $|\partial A_1| \ge \sum_{i=0}^{r+1} \binom{d}{i}$, where $r = \max\left\{ k \in \{1,..,d\} : \sum_{i=0}^{k} \binom{d}{i} \le t+1 \right\}$.

Hence, at least $\sum_{i=0}^{r+1} \binom{d}{i}$ units are correctly diagnosed and, if necessary, repaired after the first iteration. At iteration j, the number of units diagnosed is at least $\sum_{i=0}^{r+j} \binom{d}{i}$. The proof follows by observing that, for $j=d-r$, we have $\sum_{i=0}^{r+j} \binom{d}{i} = \sum_{i=0}^{d} \binom{d}{i} = n$. ❑

Theorem 4.

Let r be defined as in the statement of Theorem 3. Then r $\in \Theta(d)$.

Proof.

Let $r = \max\left\{ k \in \{1,..,d\} : \sum_{i=0}^{k} \binom{d}{i} \le t+1 \right\}$. If $r \geq d/2$, the Theorem is trivially proved. Hence, in the following we assume $r<d/2$.

We can upper and lower bound t+1 as follows:

$$\sum_{i=0}^{r}\binom{d}{i} \le t+1 \le \sum_{i=0}^{r+1}\binom{d}{i}. \qquad (1)$$

Since $r<n/2$, and observing that $\left(\frac{d}{i}\right)^i \le \binom{d}{i}$ (see, for instance, [3] pg. 102), we have:

$$\sum_{i=0}^{r}\binom{d}{i} \ge \binom{d}{r} \ge \left(\frac{d}{r}\right)^r.$$

Similarly, observing that $\binom{d}{i} \le \frac{d^i}{i!}$, we have:

$$\sum_{i=0}^{r+1}\binom{d}{i} \le (r+1)\binom{d}{r+1} \le (r+1)\frac{d^{r+1}}{(r+1)!} = \frac{d^{r+1}}{r!}.$$

Thus, inequality (1) can be rewritten as:

$$\left(\frac{d}{r}\right)^r \le t+1 \le \frac{d^{r+1}}{r!}.$$

Taking the logarithm, we have:

$$r\log\left(\frac{d}{r}\right) \le \log(t+1) \le (r+1)\log d - \log r!.$$

Thus,

$$\log(t+1) \in \Theta(\log t) \in \Omega\left(r\log\left(\frac{d}{r}\right)\right)$$

Furthermore, since $\log r! \in \Theta(r\log r)$, we have:

$$\log(t+1) \in \Theta(\log t) \in O((r+1)\log d - \log r!) = O\left(r\log\left(\frac{d}{r}\right)\right).$$

It follows that:

$$\log t \in \Theta\left(r\log\left(\frac{d}{r}\right)\right). \qquad (2)$$

Equation (2) states that functions log t and $f(r) = r\log\left(\frac{d}{r}\right)$ have the same magnitude, where function $f(r)$ is defined in the interval $\left(0, \frac{d}{2}\right)$. In principle, the magnitude of r with respect to d could be arbitrarily small. However, this magnitude is related to the magnitude of log t (and, consequently, of t), which cannot be too small. In fact, by Theorem 3 of [9] we have $t \in \Omega\left(\frac{n \log d}{d}\right)$, from which it follows $\log t \in \Omega(d)$. Since $f(r)$ is strictly increasing in the interval $\left(0, \frac{d}{2}\right)$, it follows that the magnitude of r must be at least such that $\log t \in \Theta(f(r)) \in \Omega(d)$. The thesis follows by observing that for $f(r) \in \Theta(d)$ we have $r \in \Theta(d)$.

❑

Observe that the bound reported in Theorem 4 is asymptotic, and that the derivation of a more accurate bound for r would be very important for practical purposes. Unfortunately, the derivation of r requires the knowledge of t, and only asymptotic lower bounds on t are known. Thus, the first step in deriving a more accurate bound for r would be the derivation of an explicit formula for t, which appears to be a very difficult task.

Given the analytical difficulties explained above, we have evaluated the value of r numerically. The results of this evaluation for hypercubes of different dimensions, where the values of t are taken from [2], are reported in Table 3. As it is seen, the percentage of reduction in the number of iterations ranges from 16.7% to 31.3%, and it seems to show an increasing behavior with d.

Table 3. Numerical evaluation of r for hypercubes of different dimension.

n	*t*	*d*	*r*	*%improvement*
64	15	6	1	16.7%
256	54	8	2	25%
1024	196	10	3	30%
4096	711	12	3	25%
16384	2607	14	4	28.6%
65536	9651	16	5	31.3%

Before ending this note, we want to remark that, to the best of our knowledge, the problem of determining a non-trivial lower bound to the minimum number of iterations needed by any algorithm to complete diagnosis in hypercubes is still open. Thus, it is not clear how close the performance of our algorithm is to that of an optimal algorithm.

References

[1] A. Bjorklund, "Optimal Adaptive Fault Diagnosis of Hypercubes", Proc. Scandinavian Workshop on Algorithm Theory, SWAT 2000, LNCS 1851, pp. 527-534, 2000.

[2] A. Caruso, S. Chessa, P. Santi, and P. Maestrini, "Diagnosability of Regular Structures", Journal of Algorithms, vol.45, no.2, pp.126-143, Nov. 2002.

[3] T. H. Cormen, C. E. Leiserson, and R. L. Rivest, "Introduction to Algorithms", *The MIT Press*, 1990.

[4] C. Feng, L. N. Bhuyan, and F. Lombardi, "Adaptive system-level diagnosis for hypercube multiprocessors", IEEE Transactions on Computers, vol. 45, n. 10, pp. 1157 –1170, Oct 1996.

[5] S. L. Hakimi and A. T. Amin, "Characterization of Connection Assignment of Diagnosable Systems", IEEE Transactions on Computers, vol. C-23, n.1, pp. 86-88, 1974.

[6] S. Huang and J. Xu and T. Chen, "Characterization and Design of Sequentially *t*-Diagnosable Systems", IEEE Proc. 19th Int. Symp. on Fault-Tolerant Computing, Washington, D.C., USA, pp.554-559, 1989.

[7] G.O.H. Katona, "The Hamming-Sphere has Minimum Boundary", Studia Scientarum Mathematicarum Hungarica, vol. 10, pp. 131 – 140, 1975.

[8] A. Kavianpour and K. H. Kim, "A Comparative Evaluation of Four Basic System-Level Diagnosis Strategies for Hypercubes", IEEE Trans. Reliability, vol. 41, no 1 , pp. 26-37, March 1992.

[9] S. Khanna and W.K. Fuchs, “A Graph Partitioning Approach to Sequential Diagnosis”, IEEE Trans. Comp., vol. 46, no 1, pp. 39 – 47, Jan. 1997.

[10] S. Khanna and W. K. Fuchs, “A Linear Time Algorithm for Sequential Diagnosis in Hypercubes”, Journal of Parallel and Distrib. Comput., vol. 26, pp-48-53, April 1995.

[11] E. Kranakis, A. Pelc, “Better Adaptive Diagnosis of Hypercubes”, IEEE Trans. Comput., Vol. 49, no. 10, pp. 1013-1020, Oct. 2000.

[12] I. Leader, “Discrete Isoperimetric Inequalities”, AMS Proc. of Symposia in Applied Mathematics, vol. 44, pp.57-80, 1991.

[13] P. Maestrini and C. L. Liu, “On the sequential diagnosability of a class of digital systems”, IEEE Proc. 11th Int. Symp. on Fault-Tolerant Computing, pp.112-115, 1981.

[14] F. P. Preparata, G. Metze and R. T. Chien, “On the Connection Assignment Problem of Diagnosable Systems”, IEEE Trans. Comput., Vol. EC-16, pp. 848-854, Dec. 1967.