Diagnosability of Regular Systems

Antonio Caruso[1,2], Stefano Chessa[1,2], Piero Maestrini[1,2], Paolo Santi[1]

[1]*Istituto di Elaborazione dell'Informazione,*

Area della Ricerca del CNR di Pisa-S.Cataldo, 56100 Pisa, Italy.

[2]*Computer Science Department, University of Pisa,*

Corso Italia 40, 56125 Pisa, Italy.

e-mail: {caruso,ste}@di.unipi.it, {maestrini,santi}@iei.pi.cnr.it

Abstract

A novel approach aimed at evaluating the diagnosability of regular systems under the PMC model is introduced. The diagnosability is defined as the ability to provide a correct diagnosis, although possibly incomplete. This concept is somehow intermediate between one-step diagnosability and sequential diagnosability. A lower bound to diagnosability is determined by lower bounding the minimum of a "syndrome-dependent" bound t_σ over the set of all the admissible syndromes. In turn, t_σ is determined by evaluating the cardinality of the smallest consistent fault set containing an aggregate of maximum cardinality. The new approach, which applies to any regular system, relies on the "edge-isoperimetric inequalities" of connected components of units declaring each other non-faulty. This approach has been used to derive tight lower bounds to the diagnosability of toroidal grids and hypercubes, which improve the existing bounds for the same structures.

Keywords: System-level diagnosis, PMC model, sequential diagnosis, diagnosability, regular systems, isoperimetric inequalities.

1. Introduction

System-level diagnosis, which was introduced by Preparata, Metze and Chien [14], aims at diagnosing systems composed by units (usually processors) connected by point-to-point, bidirectional links. A system S is represented as a *system graph* $G = (\mathsf{V}, \mathsf{L})$, an undirected graph where nodes represent units and edges represent interconnection links. There exists edge[1] $\{u, v\} \in \mathsf{L}$ if and only if units u and v are interconnected. The cardinality $n = |\mathsf{V}|$ is called the *size* of the system. Units u and v are said to be *adjacent*, denoted $u \longleftrightarrow v$, if $\{u, v\} \in \mathsf{L}$.

In the *PMC model* [14], diagnosis is based on a suitable set of tests between adjacent units. The *testing* unit u provides a test sequence to the *tested* unit v, which returns an output sequence to u. The testing unit compares the actual and the expected output sequences and provides a binary test outcome, defined 0 if the actual and the expected results match and 1 otherwise.

The PMC model assumes that tests of faulty units performed by non-faulty units always return 1 (that is, the test has perfect coverage), while the test outcome of tests performed by faulty units is arbitrary. This *invalidation rule* is shown in Table 1. Alternate diagnostic models assume different invalidation rule [1], or are based on comparisons between units [12].

Testing unit	Tested unit	Test outcome
Faulty-Free	Faulty-Free	0
Faulty-Free	Faulty	1
Faulty	Faulty-Free	0 or 1
Faulty	Faulty	0 or 1

Table 1. Invalidation rule of the PMC model

The set of tests utilized for the purpose of diagnosis are represented by directed edges in the *diagnostic graph DG* $= (\mathsf{V}, \mathsf{E})$, where edge (u, v) from u to v exists if and only if unit u tests unit v. Edges in E are labeled with the binary test outcomes.

Given any set $\mathsf{V}_f \subseteq \mathsf{V}$ of faulty units (*actual fault set*), the set of all test outcomes is called *syndrome*, denoted σ. Given a syndrome σ, the task of determining the status of the units is known as *syndrome-decoding*. In the *centralized diagnosis* approach [14], the syndrome is decoded by an external, reliable computer, called *diagnoser*.

Preparata et al. also introduced the concepts of *one-step* and *sequential diagnosis* [14]. In the

[1] We denote $\{u, v\}$ a directed edge from unit u to v, and (u, v) an undirected edge.

former approach it is required that syndrome decoding identifies the status of all the units. The latter approach consists of several *diagnosis and repair* phases, the goal of each phase being the identification of at least one faulty unit. Once identified, faulty units are immediately repaired or replaced, thus reducing the number of faulty units. The process is iterated until all the faulty units have been removed.

A given system S is *one-step diagnosable*; that is, one-step diagnosis of S is always possible if the number of faults in the system is not above a parameter, the *one-step diagnosability*, tied to the structure of the diagnostic graph. A similar parameter for sequential diagnosis is called *sequential diagnosability*, and is usually far above the one-step diagnosability. One-step diagnosable systems have been characterized in [8], while sequentially diagnosable systems have been characterized in [9] and [13]. The problem of determining the one-step diagnosability of a given system was solved in [18], while the analogous problem for sequential diagnosis was shown to be co-NP complete in [15].

The one-step diagnosability of a system is limited above by the minimum number of the tests undergone by units in the system, that is, by the minimum of the node in-degrees in DG [8, 14]. For this reason, the one-step diagnosis approach is inadequate to the case of large systems based on regular or quasi-regular interconnection structures, such as hypercubes, tori and grids. This is the case of massive parallel systems and wafer-scale testing [16] which, on the other hand, appear the natural candidates for application of system-level diagnosis. In fact, one-step diagnosability of such structures is very small as compared to the number of units and, thus, to the potential number of faults.

Conversely, sequential diagnosis of regular or quasi-regular systems is feasible under more realistic fault situations. Lower bounds to sequential diagnosability of grids and hypercubes have been provided in [10]. However, the repeated execution of diagnosis and repair phases entailed by the sequential diagnosis approach may be time consuming and the underlying hypothesis, which excludes the occurrence of additional faults while the diagnosis is going on, may be unrealistic. Furthermore, the sequential diagnosis is not feasible in some environments, such as wafer-scale testing.

Given a system, its diagnostic graph $DG = (\mathsf{V}, \mathsf{E})$, set $\mathsf{V}_f \subseteq \mathsf{V}$ of units which are actually faulty in the system, and a syndrome σ resulting from V_f, we consider diagnosis algorithms which result is a partition of set V into subset F of units declared faulty, subset K of units declared non-faulty, and subset S of *suspect* units, i.e., the units whose status remains unidentified. The diagnosis is said to be correct if $\mathsf{F} \subseteq \mathsf{V}_f$ and $\mathsf{K} \subseteq \mathsf{V} - \mathsf{V}_f$. The diagnosis is said to be complete if $\mathsf{S} = \emptyset$. Using this viewpoint, one-step diagnosis is both *correct* and *complete*, while sequential diagnosis may leave unidentified, at every phase, the status of as many as $|\mathsf{V}| - 1$ units.

Diagnosis algorithms which are able to diagnose in one step the state of a large fraction of the system units, thus providing an almost complete diagnosis, have been reported in the literature [3][4][5][6][17]. To some extent, the algorithm of [10] displays the same behavior. The almost complete diagnosis provided by the preceding algorithms is guaranteed to be correct if the number of actual faults is not above an *algorithm-specific diagnosability*.

From this point on, we refer to *diagnosability* as to the capability of correctly diagnosing a large fraction of the system units by means of a suitable, unspecified algorithm. Using this definition, it is clear that the diagnosability is a lower bound to the sequential diagnosability.

This paper introduces an approach to the evaluation of the diagnosability of a wide class of regular systems. This approach enables derivation of tight lower bounds to the diagnosability of regular systems for which an *edge isoperimetric inequality* [11] is known. More specifically, lower bounds to the diagnosability of square toroidal grids and hypercubes are derived, which significantly improve the bounds of [10] and [4][5][6][17].

2. Diagnosability of Regular Graphs

In this section we study the problem of characterizing the diagnosability of a system S represented by the system graph G=(V,L). Let $DG = (\mathsf{V}, \mathsf{E})$ be the diagnostic graph associated with G. We make the following assumptions, which are kept throughout this paper:

- Adjacent units perform *mutual tests*: in other words, $\{u, v\} \in \mathsf{L}$ implies $(u, v) \in \mathsf{E}$ and $(v, u) \in \mathsf{E}$.
- The system graph G is k*-regular*, i.e. all vertices in V have the same degree k. Given the hypothesis of mutual tests, this implies that every unit tests, and is tested by, exactly k neighbors.

We denote Σ the set of all the possible syndromes which may result from executing all tests corresponding to the edges of DG, and σ any syndrome in Σ. A test of v performed by unit u

with outcome $\delta \in \{0,1\}$ is denoted by $u \xrightarrow{\delta} v$. The concise notation $u \xleftrightarrow{\gamma\ \delta} v$, with $\gamma, \delta \in \{0,1\}$, denotes both the test of u performed by v with outcome γ and the test of v performed by u with outcome δ.

The properties stated in the following lemma are immediate from the invalidation rule (Table 1) of the PMC model:

Lemma 2.1. *For any two adjacent units u, v in set V:*

(a) If $u \xleftrightarrow{1\ 0} v$, then u is faulty.

(b) If $u \xrightarrow{0} v$ and u is fault-free, then v is also fault-free.

(c) If $v \xrightarrow{0} u$ and u is faulty, then v is also faulty.

(d) If $u \xleftrightarrow{1\ 1} v$, then at least one unit between u and v is faulty.

For any given syndrome $\sigma \in \Sigma$, we define $DG_0 = (\mathsf{V}, \mathsf{E}_0)$ as the subgraph of DG of edge set $\mathsf{E}_0 \subseteq \mathsf{E}$, where E_0 is the set of edges labeled with test outcome 0; that is, $\mathsf{E}_0 = \{(u,v) \in \mathsf{E} \mid u \xrightarrow{0} v\}$.

Consider the strongly connected components of DG_0 and let $\mathsf{A}_1, \mathsf{A}_2, \ldots, \mathsf{A}_r$ $(r \geq 1)$ be the vertex sets of such components. Since every vertex belongs to exactly one (possibly trivial) component, the collection of subsets $\{\mathsf{A}_1, \mathsf{A}_2, \ldots, \mathsf{A}_r\}$ is a partition of V. The following properties are immediate from Lemma 2.1:

Lemma 2.2. *Given syndrome $\sigma \in \Sigma$, let $\mathsf{A}_1, \mathsf{A}_2, \ldots, \mathsf{A}_r$ be the vertex sets of the strongly connected components of $DG_0 = (\mathsf{V}, \mathsf{E}_0)$. For every i, $1 \leq i \leq r$:*

(A) All units in A_i are in the same (faulty or non-faulty) state.

(B) If there exist units $u, v \in \mathsf{A}_i$ with $u \xrightarrow{1} v$, then all units in A_i are faulty.

(C) If there exist units $u \in \mathsf{A}_i$, $v \in \mathsf{A}_j$ with $i \neq j$ and $u \xrightarrow{0} v$, then all units in A_i are faulty.

Proof.

Property (A) is immediate from properties (b) and (c) of Lemma 2.1. Property (B) is immediate from the combination of properties (a) or (d) and (c) of Lemma 2.1. Property (C) is immediate from properties (a) and (c) of Lemma 2.1, observing that must be $u \xleftrightarrow{1\ 0} v$. $\square$

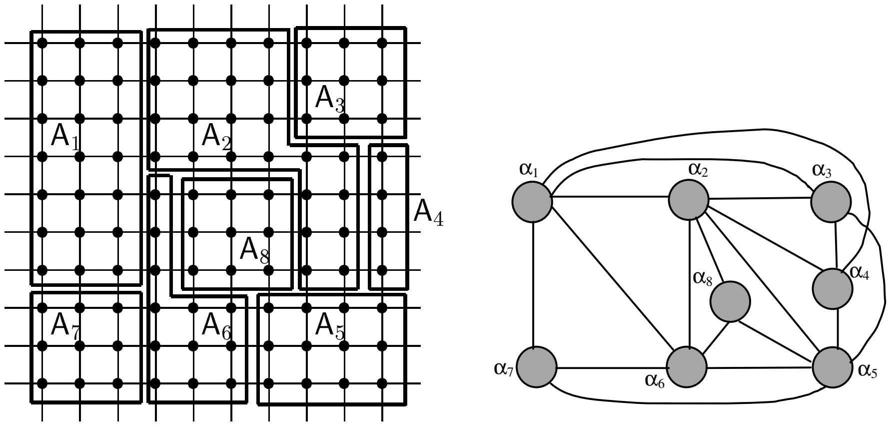

Figure 1. Z-Aggregates and corresponding contracted graph.

Consider all components satisfying conditions (B) or (C) of Lemma 2.2, and define set F as the union of all such components; it is immediate that F is completely faulty. If $\mathsf{F} \neq \emptyset$ an incomplete, unconditionally correct diagnosis is trivially available, and the goal of diagnosis is achieved. However, $\mathsf{F} = \emptyset$ in the worst case. In the following, we restrict consideration to the task of providing a diagnosis when $\mathsf{F} = \emptyset$. This also covers the more general problem of diagnosing the subsystem whose diagnostic graph is the subgraph of DG induced by vertex set $\mathsf{V} - \mathsf{F}$.

Given any $\sigma \in \Sigma$, assume $\mathsf{F} = \emptyset$. For every A_i, $i = 1, \ldots, r$ with $|\mathsf{A}_i| > 1$, this implies that adjacent units u, v in A_i must test each other with $u \overset{0\ 0}{\longleftrightarrow} v$ for every $u, v \in \mathsf{A}_i$. For this reason, components $\mathsf{A}_1, \mathsf{A}_2, \ldots, \mathsf{A}_r$ are also called *Z-aggregates*. Z-aggregates A_i, A_j, with $1 \leq i, j \leq r$, $i \neq j$, are said to be *adjacent* if there exist units $u \in \mathsf{A}_i$, $v \in \mathsf{A}_j$ such that $u \longleftrightarrow v$. It is immediate that must be $u \overset{1\ 1}{\longleftrightarrow} v$. From this point on, the system will be represented by a labeled version of the system graph $G = (\mathsf{V}, \mathsf{L})$, a undirected graph where $(u, v) \in \mathsf{L}$ is labeled 0 (or 1) if $u \overset{0\ 0}{\longleftrightarrow} v$ (or $u \overset{1\ 1}{\longleftrightarrow} v$).

We define the *contracted graph of G under syndrome σ*, denoted $CG_\sigma = (\mathsf{V}', \mathsf{L}')$ as the undirected, vertex-weighted graph obtained by contracting vertices in A_i into vertex $a_i \in \mathsf{V}'$. Vertices in V' are in a one-to-one correspondence with Z-aggregates. Edge $\{a_i, a_j\} \in \mathsf{L}'$ iff Z-aggregates A_i and A_j are adjacent. Vertex a_i is weighted with $\alpha_i = |\mathsf{A}_i|$. An example of diagnostic graph and the corresponding contracted graph are shown in Figure 1.

Given the contracted graph $CG_\sigma = (\mathsf{V}', \mathsf{E}')$, a *F-state assignment* assigns the state of `FAULTY` or `FAULT-FREE` to every $a_i \in \mathsf{V}'$. This is equivalent to assigning the state of `FAULTY` or `FAULT-`

$\mathtt{FREE}$ to every $v \in \mathsf{A}_i$, since, for every A_i, all units are in the same state. In order to be consistent with syndrome σ, a F-state assignment must satisfy the condition stated by the following theorem:

Theorem 2.3. *Given syndrome σ, assume that set F is empty. A F-state assignment of the contracted graph is consistent with σ iff the set of vertices labeled non-faulty is an independent set*[2].

Proof.
For any pair a_i, a_j of adjacent nodes of CG_σ, the outcomes of the mutual tests between units in A_i and units in A_j are 1, since set F has been assumed empty. If a_i is assumed $\mathtt{FAULT\text{-}FREE}$, every a_j adjacent to a_i in CG_σ must be faulty. This means that nodes labeled $\mathtt{FAULT\text{-}FREE}$ cannot be adjacent to nodes also labeled $\mathtt{FAULT\text{-}FREE}$. □

Given a F-state assignment consistent with σ, it is immediate that the set of units labeled $\mathtt{FAULTY}$ by this assignment, is a *Consistent Fault Set* (CFS for short) of syndrome σ [9].

In order to diagnose the state of at least one Z-aggregate A, we consider the indexes $G_0^\sigma(\mathsf{A})$, $G_1^\sigma(\mathsf{A})$, which generalize a concept introduced in [13]:

- $G_0^\sigma(\mathsf{A}_i)$: the minimum cardinality of the CFSs that do not include A_i.
- $G_1^\sigma(\mathsf{A}_i)$: the minimum cardinality of the CFSs that include A_i.

Given indexes $G_0^\sigma(\mathsf{A}_i)$ and $G_1^\sigma(\mathsf{A}_i)$, and assuming $G_0^\sigma(\mathsf{A}_i) \neq G_1^\sigma(\mathsf{A}_i)$, Z-aggregate A_i is diagnosed completely faulty if $G_0^\sigma(\mathsf{A}_i) > G_1^\sigma(\mathsf{A}_i)$ and the actual number of faults is less then $G_0^\sigma(\mathsf{A}_i)$. Observe that the number of actual faults must be at least $G_1^\sigma(\mathsf{A}_i)$. On the other hand, if $G_1^\sigma(\mathsf{A}_i) > G_0^\sigma(\mathsf{A}_i)$ and the actual number of faults is less then $G_1^\sigma(\mathsf{A}_i)$, then Z-aggregate A_i must be fault-free, and the number of actual faults must be at least $G_0^\sigma(\mathsf{A}_i)$.

Consequently, the *syndrome-dependent diagnosability* of A_i is defined as

$$t_\sigma(\mathsf{A}_i) = \max\big(G_0^\sigma(\mathsf{A}_i), G_1^\sigma(\mathsf{A}_i)\big) - 1. \tag{2.1}$$

In fact, Z-aggregate A_i can be diagnosed as either fault-free or completely faulty in the hypothesis that the number of actual faults is not above $t_\sigma(\mathsf{A}_i)$.

Given any A_i, the exact evaluation of index $G_0^\sigma(\mathsf{A}_i)$ is a difficult problem, which implies the identification of a *maximum, weighted independent vertex-set* of CG_σ with the constraint that A_i

[2] A set of vertices is said to be *independent* if no two elements of it are adjacent.

is in this set. Similarly, evaluating $G_1^{\sigma}(\mathsf{A}_i)$ implies the identification of a maximum, weighted independent vertex set of CG_σ with the constraint that A_i is not in this set. It is known that this problem is NP-complete [7].

The value of $t_\sigma(\mathsf{A}_i)$ is influenced by the choice of A_i. *The syndrome-dependent diagnosability* t_σ *of* G is defined as:

$$t_\sigma = \max_{i=1,\dots,r} t_\sigma(\mathsf{A}_i) \tag{2.2}$$

It is immediate that at least one Z-aggregate can be diagnosed, provided the number of actual faults is not above t_σ. Observe that the problem of determining t_σ has the same asymptotic complexity as determining $t_\sigma(\mathsf{A}_i)$.

The *diagnosability* of the given system, is defined as [13]:

$$t = \min_{\sigma \in \Sigma} t_\sigma \tag{2.3}$$

In fact, assuming that the cardinality of the actual fault set is not above t, at least one Z-aggregate can be correctly identified as fault-free or completely faulty for any admissible syndrome, thus fulfilling the requirements of diagnosis.

The problem of determining t has been proved to be co-NP complete by Raghavan and Tripathi [15]. Given the extreme difficulty of determining the exact diagnosability, recent research has been aimed at providing accurate approximations. This paper introduces a technique providing tight lower bounds to the diagnosability of a large class of regular graphs, i.e. those graphs for which an *edge-isoperimetric inequality* [11] is known. The new technique is presented in the next section, where bounds holding for symmetric grids and hypercubes are also derived.

3. A lower bound to diagnosability of regular systems

Consider set Σ of all syndromes, and the partition of this set into subsets $\Sigma_1, \dots, \Sigma_n$, where Σ_α is the subset of syndromes leading to Z-aggregates of cardinality at most α. In other words, denoting $\mathsf{A}^\sigma = \{\mathsf{A}_1, \mathsf{A}_2, \dots, \mathsf{A}_r\}$ the collection of Z-aggregates resulting from syndrome σ, subset Σ_α is defined as:

$$\Sigma_\alpha = \{\sigma \in \Sigma \mid \max_{\mathsf{A}_i \in \mathsf{A}^\sigma} |\mathsf{A}_i| = \alpha\}$$

The technique introduced in this paper is based on the property stated by the Theorem 2.3, i.e. independence of fault-free Z-aggregates. A lower bound to diagnosability is derived by taking $G_1^\sigma(\mathsf{A}_i) - 1$ as a lower bound to $t_\sigma(\mathsf{A}_i)$ for any $\mathsf{A}_i \in \mathsf{A}^\sigma$, and $G_1^\sigma(\mathsf{A}_r) - 1$ as a lower bound to t_σ for any $\sigma \in \Sigma$, where A_r is a Z-aggregate of maximum cardinality under σ. In fact, recalling 2.1, $t_\sigma(\mathsf{A}_i) = \max\big(G_0^\sigma(\mathsf{A}_i), G_1^\sigma(\mathsf{A}_i)\big) - 1 \geq G_1^\sigma(\mathsf{A}_i) - 1$. Further, from 2.2, $t_\sigma = \max_{\mathsf{A}_i \in \mathsf{A}^\sigma} t_\sigma(\mathsf{A}_i) \geq t_\sigma(\mathsf{A}_r)$

Recalling 2.3 and replacing t_σ by the preceding lower bound, the minimum of t_σ over Σ is evaluated in two-steps: the first step takes the minimum over Σ_α for each α, and the second step takes the minimum of such minima over α.

In conclusion, the lower bound to t is defined by the following inequality:

$$t \geq \min_{1 \leq \alpha < n/2} \; \min_{\sigma \in \Sigma_\alpha} G_1^\sigma(\mathsf{A}_r) - 1 \tag{3.1}$$

where the admissible range $[1, n/2)$ for α derives from observation that the number of faults must be below $n/2$ [14], and from inequality $G_1^\sigma(\mathsf{A}_r) \geq \alpha$, which is immediate.

Determining a tight lower bound $T(\alpha)$ to $\min_{\sigma \in \Sigma_\alpha} G_1^\sigma(\mathsf{A}_r)$ is a difficult task for general graphs. However, the problem is solvable in the case of those regular graphs for which an edge-isoperimetric inequality is known.

Given $G = (\mathsf{V}, \mathsf{L})$ and any subset $\mathsf{X} \subseteq \mathsf{V}$, the set I of *internal edges* and the set B of *boundary edges*, are defined as $\mathsf{I}(\mathsf{X}) = \big\{\{u, v\} \in \mathsf{L} \mid u, v \in \mathsf{X}\big\}$ and $\mathsf{B}(\mathsf{X}) = \big\{\{u, v\} \in \mathsf{L} \mid u \in \mathsf{X}, v \notin \mathsf{X}\big\}$. The cardinalities of sets $\mathsf{I}(\mathsf{X})$ and $\mathsf{B}(\mathsf{X})$ are denoted $i(\mathsf{X})$ and $b(\mathsf{X})$. The problems of evaluating the maximum of $i(\mathsf{X})$ and the minimum of $b(\mathsf{X})$ over all subsets X of given cardinality x, are known as *edge isoperimetric problems* [11]. We denote:

$$\overline{i}(x) = \max_{\substack{\mathsf{X} \subset \mathsf{V} \\ |\mathsf{X}| = x}} i(\mathsf{X}) \qquad \underline{b}(x) = \min_{\substack{\mathsf{X} \subset \mathsf{V} \\ |\mathsf{X}| = x}} b(\mathsf{X}) \tag{3.2}$$

The preceding problems have been satisfactorily solved for several regular graphs [11] by providing an *edge isoperimetric inequality* for $\underline{b}(x)$, that is, a function $b^*(x)$ such that $b^*(x) \leq \underline{b}(x)$. From the k-regularity of G, an edge isoperimetric inequality for $\overline{i}(x)$ is easily derived from $b^*(x)$ observing that:

$$2 \cdot i(\mathsf{X}) + b(\mathsf{X}) = k \cdot |\mathsf{X}| \tag{3.3}$$

which implies $i^*(x) = (kx - b^*(x))/2 \geq \overline{i}(x)$.

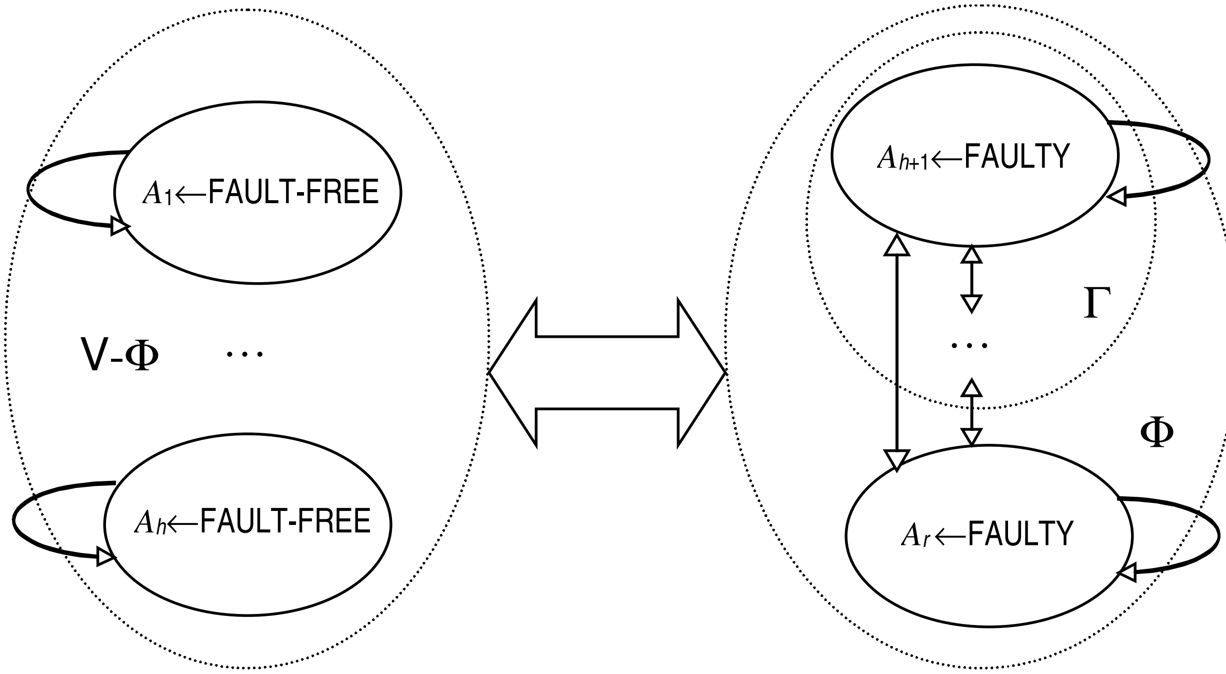

Figure 2. A Consistent Fault-Set of σ.

Let $G = (\mathsf{V}, \mathsf{L})$ be the k-regular graph of a given system, $\sigma \in \Sigma_\alpha$ a syndrome, and $\mathsf{A}^\sigma = \{\mathsf{A}_1, \mathsf{A}_2, \ldots, \mathsf{A}_r\}$ the resulting collection of Z-aggregates, with $\alpha_i = |\mathsf{A}_i| \leq \alpha$ for $i = 1, \ldots, r$. Without loss of generality, assume that $\alpha_r = \alpha$, i.e. A_r is one of the aggregates of maximum cardinality, and that $\{\mathsf{A}_1, \mathsf{A}_2, \ldots, \mathsf{A}_h\}$, $(1 \leq h < r)$ is an independent set. Defining set Φ as the union set of $\mathsf{A}_{h+1}, \ldots, \mathsf{A}_r$, it follows that Φ is a consistent fault set of σ of cardinality $\varphi = |\Phi|$. Note that the maximum aggregate A_r is considered `FAULTY`.

Denote by Γ the set $\Phi - \mathsf{A}_r = \bigcup_{i=h+1}^{r-1} \mathsf{A}_i$, and by $\gamma = \varphi - \alpha$ its cardinality. Let $\mathsf{L}(\mathsf{A}_r, \Gamma)$ be the set of edges connecting units in Γ to units in A_r (Figure 2).

Since $\{\mathsf{A}_1, \ldots, \mathsf{A}_h\}$ is an independent set we have:

$$b(\Phi) = \sum_{i=1}^{h} b(\mathsf{A}_i) \tag{3.4}$$

Introducing Equation 3.3 in the left hand side of (3.4) yields:

$$\begin{aligned} b(\Phi) &= b(\mathsf{A}_r) + b(\Gamma) - 2|\mathsf{L}(\mathsf{A}_r, \Gamma)| \leq \\ &\leq b(\mathsf{A}_r) + b(\Gamma) = b(\mathsf{A}_r) + k\gamma - 2i(\Gamma) \\ &\leq b(\mathsf{A}_r) + k\gamma \end{aligned}$$

Assuming that $b^*(x) \le \underline{b}(x)$ is an edge isoperimetric inequality for the given graph, Equation 3.4 yields:

$$\begin{aligned} b(\mathsf{A}_r) + k\gamma &\ge \sum_{i=1}^{h} b(\mathsf{A}_i) \ge \sum_{i=1}^{h} \underline{b}(\alpha_i) \\ &\ge \sum_{i=1}^{h} b^*(\alpha_i) \end{aligned} \tag{3.5}$$

The following lemmas provide an upper bound to $b(A_r)$ and a lower bound to $\sum_{i=1}^{h} b^*(\alpha_i)$.

Lemma 3.1. *Let $G = (\mathsf{V}, \mathsf{L})$ be a k-regular graph and $\mathsf{A} \subseteq \mathsf{V}$ be the vertex set of a connected component of G with $|\mathsf{A}| = \alpha$. Then, $b(\mathsf{A}) \le \alpha(k-2) + 2$.*

Proof.

From the hypothesis that the given component is connected, there exists a tree which spans all the vertices in A. Hence the number of internal edges in A is at least the number of edges belonging to its spanning tree, that is, $i(\mathsf{A}) \ge \alpha - 1$ [2]. From Equation 3.3 we have $b(\mathsf{A}) = k\alpha - 2i(\mathsf{A}) \le k\alpha - 2\alpha + 2 \le \alpha(k-2) + 2$. □

Lemma 3.2. *Let $f : [0, \alpha] \to \mathcal{R}$ be concave, and $f(0) = 0$ then*

(1) $f(x_1) + f(x_2) \ge f(x_1 + x_2) \quad \forall\, x_1, x_2 \in [0, \alpha] : x_1 + x_2 \in [0, \alpha]$*;*

(2) the function $x \to f(x)/x$ is nonincreasing in $(0, \alpha]$;

(3) $f(x_1) + f(x_2) \ge f(\alpha) + f(x_1 + x_2 - \alpha) \quad \forall\, x_1, x_2 \in [0, \alpha] : x_1 + x_2 - \alpha \in [0, \alpha]$.

Proof.
Since f is concave, its second derivative satisfies $f''(x) \le 0$ for $x \in (0, \alpha]$. We obtain:

$$0 = f(0) \le f(x) + f'(x)(-x) \quad \Rightarrow$$

$$xf'(x) \le f(x) \tag{3.6}$$

Letting $h(x) = f(x)/x$ for $x \in (0, \alpha]$, by Equation (3.6) the first derivative h' of h satisfies

$$h'(x) = \frac{xf'(x) - f(x)}{x^2} \le 0$$

that is, $h(x)$ is nonincreasing in $(0, \alpha]$, which proves statement (2).

Let x_1, x_2 be as in statement (1) and assume without loss of generality $x_2 \leq x_1$. Then

$$\begin{aligned} f(x_1+x_2) &\leq f(x_1) + f'(x_1)x_2 && \text{(by concavity)} \\ &\leq f(x_1) + f(x_1)x_2/x_1 && \text{(by 3.6)} \\ &\leq f(x_1) + f(x_2) && \text{(by statement (2), and by } x_2 \leq x_1) \end{aligned}$$

which proves statement (1).

If we define $g : [0, \alpha] \to \mathcal{R}$ by $g(x) = f(\alpha - x) - f(\alpha)$, $x \in [0, \alpha]$ then g satisfies the hypothesis of the lemma. The application of statement (1) to g gives statement (3). □

Using Lemma 3.2 and 3.1 we obtain:

Lemma 3.3. *Given Z-aggregates $\mathsf{A}_1, \mathsf{A}_2, \ldots, \mathsf{A}_h$ with $\alpha_i = |\mathsf{A}_i|$ and $\alpha_i \leq \alpha$ $(1 \leq i \leq r)$, let $b^*(x) \leq \underline{b}(x)$ be an edge isoperimetric inequality for G satisfying the hypotheses of Lemma* 3.2. *Then*

$$\sum_{i=1}^{h} b^*(\alpha_i) \;\geq\; \frac{\sum_{i=1}^{h} \alpha_i}{\alpha} \cdot b^*(\alpha) \tag{3.7}$$

Proof.

Consider Z-aggregates $\mathsf{A}_1, \mathsf{A}_2, \ldots, \mathsf{A}_h$ of cardinalities $\alpha_1, \alpha_2, \ldots, \alpha_h$ with $\alpha_i < \alpha$ for $1 \leq i \leq h$. Let d be the number of terms equal to α among $\alpha_1, \alpha_2, \ldots, \alpha_h$, and w the number of the remaining terms. Let $B_0^* = \sum_{i=1}^{h} b^*(\alpha_i)$ and consider the iterative transformation of $\sum_{i=1}^{h} \alpha_i$ and $B_0^* = \sum_{i=1}^{h} b^*(\alpha_i)$ which, in a sequence of steps, rearranges the terms and their number as follows:

On step s $(s \geq 1)$:

- If there exist α_p, α_q with $p \neq q$ and $\alpha_p + \alpha_q < \alpha$, then replace terms α_p and α_q with the unique term $\alpha_p + \alpha_q$, and replace terms $b^*(\alpha_p)$, $b^*(\alpha_q)$ of B_{i-1}^* with $b^*(\alpha_p + \alpha_q)$, yielding sum B_i^*. The number w decreases by one, d remains unchanged and, by statement (1) of Lemma 3.2, $B_i^* \leq B_{i-1}^* \leq \sum_{i=1}^{h} b^*(\alpha_i)$.

- If there exist α_p, α_q with $p \neq q$ and $\alpha_p + \alpha_q = \alpha$, then replace terms α_p and α_q with α, and replace terms $b^*(\alpha_p)$, $b^*(\alpha_q)$ of B_{i-1}^* with $b^*(\alpha)$, yielding sum B_i^*. The number w decreases by 2, d increases by one and, by statement (1) of Lemma 3.2, $B_i^* \leq B_{i-1}^* \leq \sum_{i=1}^{h} b^*(\alpha_i)$.

- If there exist α_p, α_q with $p \neq q$ and $\alpha_p + \alpha_q > \alpha$, then replace terms α_p and α_q with α and $(\alpha_p + \alpha_q - \alpha)$, and replace terms $b^*(\alpha_p)$, $b^*(\alpha_q)$ of B_{i-1}^* with $b^*(\alpha)$ and $b^*(\alpha_p + \alpha_q - \alpha)$, yielding sum B_i^*. The number w decreases by one, d increases by one and, by statement (3) of Lemma 3.2, $B_i^* \leq B_{i-1}^* \leq \sum_{i=1}^{h} b^*(\alpha_i)$.

The iteration ends when $w \leq 1$, with $\sum_{i=1}^{h} \alpha_i$ expressed as $d \cdot \alpha + \beta$, where $0 \leq \beta < \alpha$, and $B_s^* = d \cdot b^*(\alpha) + b^*(\beta) \leq \sum_{i=1}^{h} b^*(\alpha_i)$.

Replacing $d = (\sum_{i=1}^{h} \alpha_i - \beta)/\alpha$ in the latter inequality yields:

$$\sum_{i=1}^{h} b^*(\alpha_i) \geq \left(\frac{\sum_{i=1}^{h} \alpha_i}{\alpha} \right) \cdot b^*(\alpha) \ \text{ if } \beta = 0$$

and

$$\frac{\sum_{i=1}^{h} \alpha_i}{\alpha} \, b^*(\alpha) + \beta \left(\frac{b^*(\beta)}{\beta} - \frac{b^*(\alpha)}{\alpha} \right)$$

otherwise.

Being $\beta < \alpha$, it is immediate from statement (2) of Lemma 3.2 that $b^*(\beta)/\beta - b^*(\alpha)/\alpha > 0$. Thus:

$$\sum_{i=1}^{h} b^*(\alpha_i) > \frac{\sum_{i=1}^{h} \alpha_i}{\alpha} b^*(\alpha)$$

□

Combining the preceding results, a lower bound to the diagnosability is given by the following theorem:

Theorem 3.4. *Let $G = (\mathsf{V}, \mathsf{L})$ be a k-regular system graph, with $n = |\mathsf{V}|$, and $b^*(x)$ be an edge-isoperimetric inequality for $\underline{b}(x)$ satisfying the hypotheses of Lemma* 3.2*. The diagnosability t of G satisfies*

$$t \geq \min_{1 \leq \alpha < n/2} \left(\max \left(\frac{n b^*(\alpha) + 2\alpha(\alpha - 1)}{k\alpha + b^*(\alpha)}, \alpha \right) - 1 \right)$$

Proof.

Given syndrome σ, from Equation 3.1 diagnosability t is bounded by

$$t \geq \min_{1 \leq \alpha < n/2} \ \min_{\sigma \in \Sigma_\alpha} G_1^\sigma(\mathsf{A}_r) - 1$$

Introducing the bounds provided by Lemma 3.1 and Lemma 3.3 in the left and in the right hand side of 3.5, and recalling that $\sum_{i=1}^{h} \alpha_i = n - \varphi$, we obtain:

$$\alpha(k-2) + 2 + k\gamma \geq \frac{n - \varphi}{\alpha} b^*(\alpha)$$

and from $\gamma = \varphi - \alpha$:

$$\alpha(k-2) + 2 + k(\varphi - \alpha) \geq \frac{n-\varphi}{\alpha} b^*(\alpha)$$

that is:

$$\varphi \geq \frac{nb^*(\alpha) + 2\alpha(\alpha - 1)}{k\alpha + b^*(\alpha)} \tag{3.8}$$

On the other hand, from $\mathsf{A}_r \subseteq \Phi$ must also be:

$$\varphi \geq \alpha \tag{3.9}$$

The preceding inequalities hold for any collection of Z-aggregates whose cardinality is at most α, that is for any $\sigma \in \Sigma_\alpha$. Recalling that the range of interest for α is $[1, n/2)$, a lower bound to $\min_{\sigma\in\Sigma_\alpha} G_1^\sigma(\mathsf{A}_r)$ is given by $T(\alpha) = \max\big((nb^*(\alpha)+2\alpha(\alpha-1))/(k\alpha+b^*(\alpha)), \alpha\big)$. In turn, a lower bound to the diagnosability t is provided by taking the minimum of $T(\alpha)$ for $1 \leq \alpha < n/2$. □

3.1. Toroidal Grids

A *square wrap-around grid* (also called *toroidal square grid*) of size $n = l \times l$, where l is a positive integer, is represented by the product graph $G_4 = C_l \times C_l$: where C_l is a cycle of length l. Toroidal grids are 4-regular structures, due to the wrap-around links crossing the border.

Grid NG_4 of size n, also called *simple square grid*, is derived from G_4 of the same size by removing the wrap-around links. In NG_4 the degrees of units lying on the border are smaller than the degree of internal units: hence, simple grids are quasi-regular structures.

A lower bound to the diagnosability of toroidal square grids is stated in the following theorem.

Theorem 3.5. *The diagnosability t of square toroidal grids of size $n \geq 136$ satisfies:*

$$t \geq \frac{16n - 4\sqrt{3 + x_1 + x_2} + (3 + x_1 + x_2)^{\frac{3}{2}})}{16 + 8\sqrt{3 + x_1 + x_2}} - 1$$

where $x_1 = (32n^2 + 8\sqrt{2n}(2n+1)^{\frac{3}{2}} - 40n - 1)^{\frac{3}{2}}$ and $x_2 = (1 - 16n)/x_1$.

Proof.

By Theorem 3.4, $t \geq \min_{1\leq\alpha<n/2} T(\alpha) - 1$, where $T(\alpha) = \max(g(\alpha), \alpha)$ and $g(\alpha) = \frac{nb^*(\alpha)+2\alpha(\alpha-1)}{4\alpha+b^*(\alpha)}$.

Assuming $\alpha < n/2$, an edge isoperimetric inequality for toroidal grids [11] is given by $b^*(\alpha) = \min(4\sqrt{\alpha}, 2\sqrt{n})$. Observe that $b^*(\alpha) = 4\sqrt{\alpha}$ for $\alpha \leq n/4$ and $b^*(\alpha) = 2\sqrt{n}$ for $n/4 < \alpha < n/2$.

Consider the following cases:

- $n/4 < \alpha < n/2$: it is immediate that

$$\min_{n/4<\alpha<n/2} \max\{g(\alpha), \alpha\} \geq \min_{n/4<\alpha<n/2} \alpha = \frac{n}{4}$$

- $1 \leq \alpha \leq n/4$: observe that $b^*(\alpha) = 4\sqrt{\alpha}$ satisfies the hypotheses of Theorem 3.4, which thus provides a bound to diagnosability. Introducing $b^*(\alpha)$ in $g(\alpha)$, yields:

$$g(\alpha) = \frac{2n + \alpha^{\frac{3}{2}} - \sqrt{\alpha}}{2\sqrt{\alpha} + 2} \tag{3.10}$$

  Since $g(\alpha) \geq \alpha$ iff $\alpha \leq \alpha^*$, where α^* has been determined using Mathematica™ as $\alpha^* = (4-2^{4/3}(24n-1)/x_3+2^{2/3}x_3)/6$ and $x_3 = (108n^2+12\sqrt{3n(27n+2)}(n+1)-144n-2)^{1/3}$, we obtain

$$\begin{aligned} \min_{1\leq\alpha\leq n/4} T(\alpha) &= \min\left(\min_{1\leq\alpha\leq\alpha^*} g(\alpha), \min_{\alpha^*<\alpha\leq n/4} \alpha\right) \\ &= \min\left(\min_{1\leq\alpha\leq\alpha^*} g(\alpha), \alpha^*\right) \end{aligned} \tag{3.11}$$

  The minimum of $g(\alpha)$ in $[1, \alpha^*]$ occurs when α equals $\bar{\alpha} = (3 + x_1 + x_2)/4$, where the expressions of x_1 and x_2 are given in the statement of the theorem. Substituting $\bar{\alpha}$ in Equation 3.10 we obtain:

$$g(\overline{\alpha}) = \frac{16n - 4\sqrt{3 + x_1 + x_2} + (3 + x_1 + x_2)^{\frac{3}{2}}}{16 + 8\sqrt{3 + x_1 + x_2}} \tag{3.12}$$

  Since $\bar{\alpha} < \alpha^*$ and $g(\bar{\alpha}) < \alpha^*$ for $n > 2$, we have that:

$$\min\left(\min_{1\leq\alpha\leq\alpha^*} g(\alpha), \alpha^*\right) - 1 \geq g(\bar{\alpha}) - 1$$

Observing that $g(\bar{\alpha}) < n/4$ and $\alpha^* < n/4$ for $n > 136$, if the latter inequality holds, $t \geq g(\bar{\alpha})-1$ in both the preceding cases. □

The preceding lower bound to the diagnosability, which is $\Theta\left(n^{2/3}\right)$, has been evaluated numerically. A listing for selected values of the size n of the grids is reported in Table 2 as entry LT. This table also reports numerical evaluation of Khanna and Fuchs's [10] and Chessa and Maestrini's [6] lower bounds to diagnosability of simple square grids (entries K and C1, respectively), and of Chessa's lower bound to diagnosability of square toroidal grids [5] (entry C2). It is seen that the lower bound provided by Theorem 3.5 is above the bounds of [5] and [6], the best previously known results, and far above the bound of [10].

Table 2. Lower bounds to the diagnosability of square grids.

n	K	C1	C2	LT
256	26	34	47	50
1024	75	101	128	136
4096	206	281	340	359
16384	547	762	885	929
65536	1435	2022	2273	2377

3.2. Hypercubes

The binary d-cube $\mathcal{H}_n$, also called *hypercube*, is a graph with $n = 2^d$ nodes labeled with d-digits binary numbers. Nodes are connected based on the *Hamming distance* of their labels: edge $\{u, v\}$ exists iff the Hamming distance of the labels of u and v is 1. Hypercubes are n-regular structures, and the degree is an increasing function of the size. The diagnosability of hypercubes is given by the following theorem.

Theorem 3.6. *The diagnosability t of the hypercube of size n satisfies:*

$$t \geq \min\left(\min_{1 \leq \alpha \leq \frac{n}{e}} \max\bigl(g(\alpha), \alpha\bigr), \frac{n}{e}\right) - 1$$

where e is the base of the natural logarithms[3]*, and*

$$g(\alpha) = \frac{n \log n - n \log \alpha + \alpha \log 4 - \log 4}{2 \log n - \log \alpha}$$

Proof.
By Theorem 3.4, we have that:

$$t \geq \min_{1 \leq \alpha < n/2} T(\alpha) - 1 = \min_{1 \leq \alpha < n/2} \left(\max(g(\alpha), \alpha)\right) - 1$$

where $g(\alpha) = \frac{n b^*(\alpha) + 2\alpha(\alpha - 1)}{\alpha \log n + b^*(\alpha)}$. An edge isoperimetric inequality for hypercubes [11] is

$$b^*(\alpha) = \begin{cases} \frac{\alpha(\log n - \log \alpha)}{\log 2} & \text{if } \alpha > 0 \\ 0 & \text{if } \alpha = 0 \end{cases}$$

Function $b^*(\alpha)$ satisfies the hypotheses of Theorem 3.4 in $[0, n/e]$. In this interval $g(\alpha)$ has been determined by means of Mathematica™, yielding:

[3] In this paper, all the logarithms are to base e.

Table 3. Comparison of the Diagnosability bounds for Hypercubes structures.

n	K	LT
64	15	18
256	54	62
1024	196	220
4096	711	786
16384	2607	2846
65536	9651	10432

$$g(\alpha) = \frac{n \log n - n \log \alpha + \alpha \log 4 - \log 4}{2 \log n - \log \alpha}$$

Observing that $\min_{n/e<\alpha<n/2}\big(\max(g(\alpha), \alpha)\big) > n/e$, the expression of $\min_{1\leq\alpha<n/2} T(\alpha)$ is expressed as:

$$\min_{1\leq\alpha<\frac{n}{2}} T(\alpha) = \min\left(\min_{1\leq\alpha\leq\frac{n}{e}} \max\big(g(\alpha), \alpha\big), \frac{n}{e}\right),$$

which proves the theorem. □

So far, the only known bound to the hypercube diagnosability have been provided by Khanna and Fuchs in [10]. Similarly to [10], it is seen that t is $\Omega\left(\frac{n \log\log n}{\log n}\right)$.

The lower bound provided by the preceding theorem has been evaluated numerically. A listing for selected values of n (the size of the hypercubes) is reported as entry LT in Table 3, along with the numerical evaluation of [10] (entry K in the table). It is seen that the lower bound of Theorem 3.6 improves the previous result.

4. Comparison with previous results

The lower bound to diagnosability provided in this paper is based on equations 2.1, 2.2 and 2.3. This appears to be the natural approach to the problem, although it was discarded in the previous research due to the extreme difficulty of the involved computational problems, as explained in Section 2. The difficulty has been now circumvented by devising approximations which rely on isoperimetric inequalities of regular graphs. Although approximations may appear relatively coarse, they enabled derivation of a bound which is tighter than those previously known [4, 5, 6, 10, 17].

Although seemingly different, the approaches reported in the previous literature are somehow related to the one introduced in this paper. The extend to which they are related, as well as the reason of relative inaccuracy, are informally explained in the following.

In [10], Khanna and Fuchs also consider a diagnosis algorithm based on the aggregation of units which declare each other non-faulty, and derive a lower bound to diagnosability based on a parameter of the system graph, called the *k-partition number*. Given a connected graph $G = (\mathsf{V}, \mathsf{L})$, the k-partition number of G, denoted $\Phi_G(k)$, is the largest integer p such that, for every $\mathsf{A} \subseteq \mathsf{V}$ with $|\mathsf{A}| = p$, the subgraph of G induced by the vertices in $\mathsf{V} - \mathsf{A}$ has a connected component of size k or larger. A lower bound to the diagnosability of G is derived by evaluating the largest integer k such that $\Phi_G(k+1) \geq k$. Function $\Phi_G(x)$ is derived from an upper bound to the *vertex degree sum*[4] of subsets $\mathsf{A} \subseteq \mathsf{V}$ with $|\mathsf{A}| = x$.

The approach of [10] is related to the "natural" approach used in this paper for the following reasons. Consider a labeled system graph $G = (\mathsf{V}, \mathsf{L})$, and assume $\Phi_G(k+1) \geq k$. Given any $\mathsf{V}_1 \subseteq \mathsf{V}$ with $|\mathsf{V}_1| \leq k$, consider the connected components $\mathsf{A}_1, \mathsf{A}_2, \ldots, \mathsf{A}_r$ of the subgraph of G induced by node set $\mathsf{V} - \mathsf{V}_1$. Components $\mathsf{A}_1, \mathsf{A}_2, \ldots, \mathsf{A}_r$ are independent aggregates, being separated by units in V_1. Given inequality $\Phi_G(k+1) \geq k$, there exists at least one aggregate, say A_1, with $|\mathsf{A}_1| \geq k+1$. It follows that $G_1^\sigma(\mathsf{A}_1) \geq k+1$ and $G_0^\sigma(\mathsf{A}_1) \leq k$. This means that inequality $\Phi_G(k+1) \geq k$ is sufficient to ensure the existence of at least one aggregate A_1 such that $\max\{G_0^\sigma(\mathsf{A}_1), G_1^\sigma(\mathsf{A}_1)\} \geq k+1$. This aggregate is diagnosed as non-faulty in the assumption that the number of faults is at most k.

The relative inaccuracy of the bound provided in [10] resides in the fact that index G_1^σ which is evaluated for an aggregate of cardinality at least $k+1$, does not account for the additional faults needed to separate the remaining aggregates, which must be independent.

The approach used in [4] [5] [6] [17] is based on a diagnosis algorithm (called DAGS) which diagnoses a large fraction of the system units in a single phase. Given syndrome σ, DAGS constructs first a subset $\mathsf{S} \subseteq \mathsf{V}$ of "suspect" units, by matching the units which test each other as faulty, and then partitions $\mathsf{V} - \mathsf{S}$ into subsets $\mathsf{A}_1, \mathsf{A}_2, \ldots, \mathsf{A}_h$, called Z-aggregates. Z-aggregates are defined as the connected components of the subgraph of $G = (\mathsf{V}, \mathsf{L})$ induced by vertex set $\mathsf{V} - \mathsf{S}$. The set of Z-aggregates $\mathsf{A}_1, \ldots, \mathsf{A}_h$ is independent, due to the separating set S. It should be observed

[4]Given graph $G = (\mathsf{V}, \mathsf{L})$, the *vertex degree sum* of the vertex subset $\mathsf{A} \subseteq \mathsf{V}$ is defined as $vds(\mathsf{A}) = \sum_{x \in \mathsf{A}} |\{\ \{x, y\} \in \mathsf{L}$ such that $x, y \in \mathsf{A}\}|$.

that aggregates constructed by DAGS are proper subsets of aggregates considered in this paper and in [10]. Letting A be a Z-aggregate of maximum cardinality α, $G_1^\sigma(\mathsf{A})$ is lower bounded by $t_\sigma = \alpha + |\mathsf{S}|/2$. This bound relies on property that at least $|\mathsf{S}|/2$ units in S must be faulty, although they remain unidentified.

A lower bound to the diagnosability is derived by taking the minimum of t_σ over the admissible range of α. Although quite accurate, the bounds of [4][5][6][17] are below the bound derived in this paper, due to the fact that aggregates defined by DAGS are proper subsets of the aggregates considered in this paper.

5. Final Remarks

A novel approach aimed at evaluating the diagnosability of regular systems under the diagnostic model of [14] has been introduced. The proposed approach applies to any regular system graph for which an edge-isoperimetric inequality is known. This is the case of toroidal grids and hypercubes, and lower bounds to the diagnosability of these structures have been derived in this paper. The new bounds are tighter than those previously known [4][5][6][10][17].

The issue of defining a diagnosis algorithm which provides correct diagnosis in the occurrence of at most t faults is not explicitly addressed in this paper. A diagnosis algorithm would simply consist in declaring fault-free the Z-aggregate A_r of maximum cardinality, and faulty all the Z-aggregates adjacent to A_r. If there exist multiple Z-aggregates of the same, maximum cardinality (which cannot be pairwise adjacent), all of them are diagnosed fault-free, and every Z-aggregate adjacent to at least one of them is diagnosed faulty. As in the case of DAGS, the diagnosis returned by this algorithm is generally incomplete.

Although it could be viewed as a sequential diagnosis algorithm (and, consequently, diagnosability t may be taken as a lower bound to the sequential diagnosability), the above outlined algorithm does much more than just sequential diagnosis, since it may be expected to identify the status of a large fraction of the system units.

The latter claim is supported by the simulative analysis of DAGS [3][5][6][17], showing that the diagnoses provided by this algorithm are almost complete. Considering that Z-aggregates constructed by DAGS are proper subsets of those defined in this paper, and that, with high probability, the maximum cardinality Z-aggregate constructed by DAGS is unique and the remaining aggregates are much smaller, it may be expected that DAGS and the above described algorithm return the same diagnoses in most cases. This matter is the object of continuing investigation.

References

[1] F. Barsi, F. Grandoni, and P. Maestrini. A theory of diagnosability of digital systems. *IEEE Transactions on Computers*, C-25(6):585–593, June 1976.

[2] B. Bollobas. *Modern Graph Theory*. Springer, 1998.

[3] A. Caruso, S. Chessa, P. Maestrini, and P. Santi. Diagnosis of regular structures. In *IEEE Proc. Dependable Systems and Networks*, pages 213–222, New York, NY, June 2000.

[4] A. Caruso, S. Chessa, P. Maestrini, and P. Santi. Fault–diagnosis of grid structures. *Theoretical Computer Science, Elsevier*, to appear.

[5] S. Chessa. *Self-Diagnosis of Grid-Interconnected Systems with Application to Self-Test of VLSI Wafers*. PhD thesis, TD-2/99, University of Pisa, Italy, Department of Computer Science, March 1999.

[6] S. Chessa and P. Maestrini. Correct and almost complete diagnosis of processors grids. Technical Report B4.13-5, Istituto di Elaborazione dell'Informazione del CNR, Pisa, Italy, May 1999.

[7] M. Garey and D. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman and Company, New York, 1979.

[8] S. L. Hakimi and A. T. Amin. Characterization of connection assignment of diagnosable systems. *IEEE Transactions on Computers*, C-23(1):86–88, January 1974.

[9] S. Huang, J. Xu, and T. Chen. Characterization and design of sequentially t-diagnosable systems. In *IEEE Proc.* 19^{th} *Int. Symp. on Fault-Tolerant Computing*, pages 554–559, Washington, D.C., USA, June 1989.

[10] S. Khanna and W. K. Fuchs. A graph partitioning approch to sequential diagnosis. *IEEE Transactions on Computers*, 46(1):39–47, January 1996.

[11] I. Leader. Discrete isoperimetric inequalities. *AMS Proc. of Symposia in Applied Mathematics*, 44:57–80, 1991.

[12] J. Maeng and M. Malek. A comparison connection assignment for self-diagnosis of multicomputer systems. In *IEEE Proc.* 11^{th} *Int. Symp. on Fault-Tolerant Computing*, pages 173–175, June 1981.

[13] P. Maestrini and C. L. Liu. On the sequential diagnosability of a class of digital systems. In *IEEE Proc.* 11^{th} *Int. Symp. on Fault-Tolerant Computing*, pages 112–115, June 1981.

[14] F. P. Preparata, G. Metze, and R. T. Chien. On the connection assignment problem of diagnosable systems. *IEEE Transactions on Computers*, EC-16(12):848–854, December 1967.

[15] V. Raghavan and A. Tripathi. Sequential diagnosability is co-NP complete. *IEEE Transactions on Computers*, 40(5):584–595, May 1991.

[16] S. Rangarajan, D. Fussel, and M. Malek. Built-in testing of integrated circuit wafers. *IEEE Transactions on Computers*, 39(2):195–205, February 1990.

[17] P. Santi. *Evaluation of a Self-Diagnosis Algorithm for Regular Structures*. PhD thesis, TD–6/00, University of Pisa, Italy, Department of Computer Science, March 2000.

[18] G. Sullivan. A polynomial time algorithm for fault diagnosability. In *Proceedings of the* 25^{th} *Int Symp. on Foundations of Computer Science*, pages 148–156, Los Angeles, Ca., USA, October 1984.