Theoretical Computer Science 452 (2012) 132–147

Contents lists available at SciVerse ScienceDirect

Theoretical Computer Science

journal homepage: www.elsevier.com/locate/tcs

A general technique to establish the asymptotic conditional diagnosability of interconnection networks

Iain A. Stewart *

School of Engineering and Computing Sciences, Durham University, Science Labs, South Road, Durham DH1 3LE, UK

ARTICLE INFO

Article history:
Received 14 February 2012
Received in revised form 2 May 2012
Accepted 12 May 2012
Communicated by R. Klasing

ABSTRACT

We develop a general and demonstrably widely applicable technique for determining the asymptotic conditional diagnosability of interconnection networks prevalent within parallel computing under the comparison diagnosis model. We apply our technique to replicate (yet extend) existing results for hypercubes and k-ary n-cubes before going on to obtain new results as regards folded hypercubes, pancake graphs and augmented cubes. In particular, we show that the asymptotic conditional diagnosability of: folded hypercubes $\{FQ_n\}$ is $3n-2$, pancake graphs $\{P_n\}$ is $3n-7$, and augmented cubes $\{AQ_n\}$ is $6n-17$. We demonstrate how our technique is independent of structural properties of the interconnection network G in question and essentially only dependent upon the minimal size of the neighbourhood of a path of length 2 in G, the number of neighbours any two distinct vertices of G have in common, and the minimal degree of any vertex in G.

1. Introduction

The design of interconnection networks is fundamental to parallel computing, for as to how one (directly) connects processors in some distributed-memory multiprocessor (along with accompanying design decisions relating to, for example, routing, flow control, switching and packaging) has a tremendous impact upon the resulting efficiency of the machine [5,6]. There is no one family of interconnection networks that is better than all of the others, for the quality of a family of interconnection networks depends upon the properties that happen to be of most relevance to a particular scenario. These properties include having low degree and high connectivity, being vertex- or edge-transitive, having simple and efficient routing and broadcast algorithms, being recursively decomposable, and possessing embedded Hamiltonian cycles or paths and cycles of a whole variety of lengths.

Not only should an interconnection network possess desirable properties such as those above but any distributed-memory multiprocessor should be able to tolerate a limited number of processor or link failures. This expectation has provoked much research on not just the sustainability of specific interconnection network properties in the presence of faults but also the detection of actual faults in a distributed-memory multiprocessor. It is with this latter research direction that we are concerned in this paper. Imagine the situation. A distributed multiprocessor system is known to possess some faulty processors but it is not known as to which processors are faulty. The problem is to detect the faulty processors, that is, to *diagnose* the set of faulty processors. Crucial to this diagnosis is the observation that we can use the processors of the system to do this, that is, we can undertake a self-diagnosis. As to how this is done depends upon the model adopted.

A popular model is the *comparison diagnosis model* (also called the *MM model*), advocated by Malek and Maeng [23,24]. In this model, a processor can send a message to any two of its neighbours who then send replies back to the processor. On receipt of these two replies, the processor compares them and proclaims that at least one of the two neighbours is faulty

* Tel.: +44 1913341720.
E-mail address: i.a.stewart@durham.ac.uk.

0304-3975/$ – see front matter © 2012 Elsevier B.V. All rights reserved.
doi:10.1016/j.tcs.2012.05.015

if the replies are different or that both neighbours are fault-free if the replies are identical. However, if the processor itself is faulty then no reliance can be placed on this proclamation. The goal is to use these tests made by various processors in order to deduce exactly which are the faulty processors. Obviously there are limits as to what can be done. For example, if all processors are faulty then there is no way that this can be detected (from any collection of tests undertaken). For a specific interconnection network (forming the underlying topology of some distributed-memory multiprocessor), there is a bound on the number of faulty processors that can necessarily be detected within this model and a considerable amount of research has been undertaken on determining this bound, or the *diagnosability*, for different interconnection networks (see, for example, [8–11,18–20,28,32] for a selection of results).

The diagnosability of an interconnection network is determined by the topology of the network; so, henceforth, we equate an interconnection network of processors in some distributed-memory multiprocessor with an undirected graph and we talk about faulty vertices as opposed to faulty processors. In [17], Lai et al. observed that the diagnosability of many interconnection networks increases if one rules out the possibility that a set of faulty vertices can contain all neighbours of some vertex, and they proposed a more refined notion of diagnosability, namely *conditional diagnosability*, where all of the above principles apply except that one has the *a priori* stipulation that a set of faulty vertices can never contain the set of all neighbours of some vertex (this observation is made in the context of the PMC model in [17] but is equally valid in the comparison diagnosis model). Alternatively, if one assumes that any processor in a multiprocessor system fails with equal independent probability then a simple statistical analysis shows that the likelihood that every neighbour of some given processor is faulty is extremely small in many interconnection networks (with this likelihood decreasing as the parameter n indexing the family increases). Results on conditional diagnosability in the comparison diagnosis model include those in, for example, [12–15,22,31,33–36] (we shall revisit some of these results later).

In this paper, we develop a general and demonstrably widely applicable technique for determining the asymptotic conditional diagnosability of interconnection networks prevalent within parallel computing (that is, the limiting behaviour of the conditional diagnosability of a family $\{X_n\}$ of interconnection networks as n increases). We apply our technique to replicate (yet extend) existing results for hypercubes and k-ary n-cubes before going on to obtain new results as regards folded hypercubes, pancake graphs and augmented cubes. In particular, we show that the asymptotic conditional diagnosability of: folded hypercubes $\{FQ_n\}$ is $3n-2$; pancake graphs $\{P_n\}$ is $3n-7$; and augmented cubes $\{AQ_n\}$ is $6n-17$. We demonstrate how our technique is independent of structural properties of the interconnection network G in question and only dependent upon (essentially) the minimal size of the neighbourhood of a path of length 2 in G, the number of neighbours any two distinct vertices of G have in common, and the minimal degree of any vertex in G. Whilst our technique is extremely powerful in that it reduces ascertaining the asymptotic conditional diagnosability to the elucidation of these three parameters, our application as regards augmented cubes shows that ascertaining these parameters is not always straightforward.

In the next section, we give basic definitions relating to interconnection networks (when viewed as undirected graphs) and diagnosability, before outlining related research on conditional diagnosability in Section 3. We detail our general technique in Section 4 before we apply this technique to hypercubes and k-ary n-cubes in Section 5 (we start with hypercubes and k-ary n-cubes as the application of our technique is particularly straightforward in these cases and we also have existing results to compare with; that said, we do establish new results for 3-ary n-cubes). In Section 6, we use our technique to establish new conditional diagnosability results for folded hypercubes, pancake graphs and augmented cubes, with the latter application being decidedly non-trivial. We present our conclusions and directions for further research in Section 7.

2. Basic definitions

In parallel computing, an interconnection network consists of a set of processors together with a set of bidirectional links involving certain pairs of distinct processors. Consequently, throughout we identify an *interconnection network* with an undirected graph $G = (V, E)$ with vertex set V and edge set E where there are no multiple edges or self-loops. The interconnection networks relevant to parallel computing come in families with each interconnection network of a family parameterised by some non-zero positive integer. For example, the family of hypercubes $\{Q_n\}$ are such that the vertex set of Q_n is $\{0, 1\}^n$ and there is an edge joining two vertices if, and only if, the corresponding bit-strings of length n differ in exactly one bit. When our domain is $\{0, 1\}$, we write $\bar{x}$ to denote 0 if x is 1 and 1 if x is 0; so, any vertex $x_1x_2 \ldots x_i \ldots x_n$ of Q_n is adjacent to $x_1x_2 \ldots \bar{x}_i \ldots x_n$, for every $i \in \{1, 2, \ldots, n\}$. We say that a vertex $x_1x_2 \ldots x_n$ has *weight* m if exactly m of the n bits are 1.

Let $G = (V, E)$ be an arbitrary graph. If $(u, v) \in E$ then we say that u (resp. v) is *adjacent* to v (resp. u) or that u (resp. v) is a *neighbour* of v (resp. u). We will be interested in certain aspects of a graph $G = (V, E)$. The *degree* of a vertex v is denoted $d_G(v)$ and defined as $|\{u \in V : (u, v) \in E\}|$, with $\Delta(G)$ being the degree of a vertex of minimum degree. Given a subset of vertices $U \subseteq V$, we define the *neighbourhood* of U, denoted $N_G(U)$, as the set of vertices each of which is adjacent to at least one vertex of U but which is not in U; that is, $N_G(U) = \{v \in V \setminus U : (u, v) \in E, \text{ for some } u \in U\}$. If H is a sub-graph of G involving the vertices of $U \subseteq V$ then we define $N_G(H)$ as $N_G(U)$ and $G \setminus U$ as the subgraph of G obtained by deleting all vertices of U and any edge that is adjacent with at least one vertex of U. A *path* ρ of length $m-1$ is a sequence of distinct vertices $(v_1, v_2, \ldots, v_m)$, for some $m \geq 1$, such that $(v_i, v_{i+1}) \in E$, for $i = 1, 2, \ldots, m-1$. A *connected component* of G is a maximal set of vertices with the property that there is a path in G from any vertex of this set to any other. We define $p_2(G)$ to be the minimum size of the neighbourhood of any path in G of length 2; that is,

$p_2(G) = \min\{|N_G(\rho)| : \rho$ is a path of length 2 in $G\}$. A *cycle* of length m is a path $(v_1, v_2, \ldots, v_m)$ of length $m \geq 2$ so that $(v_m, v_1) \in E$. The *girth* of G is the length of a shortest cycle in G. A *clique* of size k in a graph G is a subset of exactly k vertices each of which is adjacent to all the others. We define $c(G)$ to be the maximum number of vertices any pair of vertices are both adjacent to; that is, $c(G) = \max\{|N_G(u) \cap N_G(v)| : u, v \in V, u \neq v\}$. A graph G is *vertex-transitive* (resp. *edge-transitive*) if there is an automorphism of G mapping any chosen vertex (resp. edge) to any other chosen vertex (resp. edge). Additional details as regards the definitions above can be found in [16,29].

There are two basic models prevalent as regards fault diagnosis in interconnection networks: the *PMC model* (proposed by Preparata et al. [26]) and the *comparison diagnosis model* (also called the *MM model* and advocated by Malek and Maeng [23,24]). It is with the comparison diagnosis model that we are concerned in this paper (or, more precisely, a variant of it that we will detail in a moment). The comparison diagnosis model is as follows. Given a graph $G = (V, E)$ within which there may be *faulty* vertices, from some *fault set*, every vertex u of V tests every pair v and w of its neighbours by sending a test message to both neighbours and receiving replies. We assume that: all faults are permanent; and a faulty vertex always produces an incorrect response to any test message, so that two faulty vertices do not produce identical responses to any test messages. Suppose that u is a *healthy* vertex; that is, it is not faulty. If the replies from v and w are identical then the test result $s_u(v, w)$ is set at 0 (signalling that both v and w are healthy), otherwise $s_u(v, w)$ is set at 1 (signalling that at least one of v and w is faulty). However, if u is a faulty vertex then the test result $s_u(v, w)$ can be arbitrarily 0 or 1 with no reliance placed upon this result. The set of all test results for every vertex and its pairs of neighbours is called a *syndrome*. The general fault diagnosis problem is: given a graph $G = (V, E)$ and a syndrome, can we use the data therein to obtain exactly the set of faulty vertices and, if so, to find these faulty vertices?

Note that the same syndrome could arise from different sets of faulty vertices; that is, there might be more than one set of faulty vertices *consistent* with the syndrome. Let $G = (V, E)$ be some graph and let $F_1, F_2 \subseteq V$ be two fault sets. We say that F_1 and F_2 and *distinguishable* if there is no syndrome consistent with both F_1 and F_2; otherwise, F_1 and F_2 are *indistinguishable*. A graph $G = (V, E)$ is said to be δ-*diagnosable* if given a syndrome s resulting from a set of at most δ faulty vertices, there is exactly one set of faulty vertices consistent with s. The maximum number δ for which a graph $G = (V, E)$ is δ-diagnosable is the *diagnosability* of G. Sengupta and Dahbura [27] were the first to provide structural conditions upon G for it to be δ-diagnosable. One remark we have is that the diagnosability of any graph $G = (V, E)$ is bounded above by $\Delta(G)$. To see this, suppose that u is some vertex of minimal degree in G and consider the following two sets of faulty vertices: the first fault set consists of all u's neighbours; and the second of all u's neighbours as well as u. It is not difficult to see that there is a syndrome that both of these sets of faults are consistent with.

A *conditional fault set* in $G = (V, E)$ is a set of faults with the property that for every vertex v of V, not all of v's neighbours in G are faults. If one assumes that all fault sets are always conditional and works within the framework above then the concept of *conditional diagnosability* arises. If there is a function $f(n)$ and an integer n_0 so that an interconnection network X_n from a family of interconnection networks $\{X_n\}$ has (resp. conditional) diagnosability $f(n)$, for every $n \geq n_0$, then we say that the family of interconnection networks has *asymptotic* (resp. *conditional*) *diagnosability* $f(n)$.

3. Related research

The conditional diagnosabilities of a number of families of interconnection networks have been considered, both within the PMC model (see, for example, [3,17,21,30]) and the comparison diagnosis model. The conditional diagnosabilities of the following interconnection networks under the comparison diagnosis model have previously been established: the conditional diagnosability of any BC-Network X_n (also called a hypercube-like network) is $3n - 5$ when $n \geq 5$ [14,15], and so this is true when X_n is an n-dimensional hypercube (see also [13,33]), an n-dimensional twisted cube (see also [35]), an n-dimensional crossed cube and an n-dimensional Möbius cube (see also [34]); the conditional diagnosability of any Cayley graph generated by transposition trees is $3n-8$ except for the case of the n-dimensional star graph when it is $3n-7$, under the proviso that $n \geq 4$ [22]; the conditional diagnosability of the alternating group network AN_n is $3n - 9$ when $n \geq 5$ [36]; the conditional diagnosability of the hypermesh $H_{n,k}$ is $3n(k - 1) - 2k - 1$ when $n \geq 3$ and $k \geq 4$ [31]; and the conditional diagnosability of the k-ary n-cube Q_n^k is $6n - 5$ when $n \geq 4$ and $k \geq 4$ [12]. The general technique used has been to assume that we have two conditional fault sets F_1 and F_2, of a certain size, in some graph G and to examine the structure of graphs such as $G \setminus (F_1 \cup F_2)$ and $G \setminus F_1$ under the assumption that F_1 and F_2 are indistinguishable. This analysis has been concerned with the existence of large connected components and tied to specific interconnection networks. As we see below, we can actually make this technique more generic by concentrating on the existence of connected components in the form of a K_2 (and not on large connected components) and by using some combinatorial arguments.

4. A general technique

In this section, we establish a general technique for ascertaining the conditional diagnosability of an arbitrary graph. Our technique is widely applicable, especially amongst graphs prevalent as interconnection networks as we subsequently demonstrate. Before detailing our technique, we establish some useful lemmas.

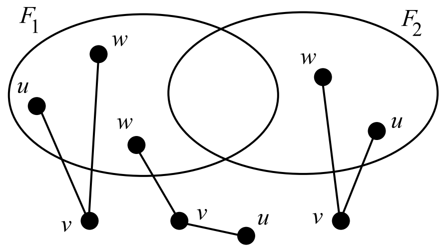

Fig. 1. Sengupta and Dahbura's classification.

4.1. Some useful lemmas

An extremely useful classification of when two fault sets are distinguishable has been established by Sengupta and Dahbura. We write $A \triangle B$ to denote the symmetric difference of two sets A and B, and we write $A \setminus B$ to denote the set $\{a \in A : a \notin B\}$.

Theorem 1 ([27, Theorem 1]). *Let $G = (V, E)$ be a graph and let $F_1, F_2 \subseteq V$ be fault sets where $F_1 \neq F_2$. The fault sets F_1 and F_2 are distinguishable if, and only if, at least one of the following conditions is satisfied in G:*

1. *there are $u, v \in V \setminus (F_1 \cup F_2)$ and $w \in F_1 \triangle F_2$ such that (u, v, w) is a path;*
2. *there are $u, w \in F_1 \setminus F_2$ and $v \in V \setminus (F_1 \cup F_2)$ such that (u, v, w) is a path;*
3. *there are $u, w \in F_2 \setminus F_1$ and $v \in V \setminus (F_1 \cup F_2)$ such that (u, v, w) is a path.*

We use Theorem 1 throughout. The three conditions in Theorem 1 can be visualised as in Fig. 1

The next lemma gives a simple upper bound on the conditional diagnosability of a graph.

Lemma 2. *Let G be a graph and let (v_1, u, v_2) be a path of length 2 in G such that $|N_G(\{v_1, u, v_2\})| = m$. The conditional diagnosability of G is at most m.*

Proof. Define $F_i = N_G(\{v_1, u, v_2\}) \cup \{v_i\}$, for $i = 1, 2$. By Theorem 1, F_1 and F_2 are indistinguishable. □

The next lemma provides useful information about the neighbourhoods of certain vertices lying outside two given indistinguishable conditional fault sets.

Lemma 3. *Let $G = (V, E)$ be a graph and let $F_1, F_2 \subseteq V$ be conditional fault sets, where $F_1 \neq F_2$. Suppose further that F_1 and F_2 are indistinguishable. If $x \in V \setminus (F_1 \cup F_2)$ is adjacent to some vertex of $F_1 \setminus F_2$ and $d_G(x) \geq 2$ then x is adjacent to: exactly one vertex of $F_1 \setminus F_2$; exactly one vertex of $F_2 \setminus F_1$; and $d_G(x) - 2$ vertices of $F_1 \cap F_2$.*

Proof. By Theorem 1: all neighbours of x lie in $F_1 \cup F_2$; x has exactly one neighbour in $F_1 \setminus F_2$; and x has at most one neighbour in $F_2 \setminus F_1$. As F_1 (resp. F_2) is a conditional fault set, x must have exactly one neighbour in $F_2 \setminus F_1$ (resp. $F_1 \setminus F_2$). □

Our final lemma deals with a trivial condition for two conditional fault sets to be distinguishable.

Lemma 4. *Let $F_1, F_2 \subseteq V$ be conditional fault sets, with $F_2 \subset F_1$. The sets F_1 and F_2 are distinguishable.*

Proof. Let $u \in F_1 \setminus F_2$. As F_1 is a conditional fault set, there exists a vertex $v \notin F_1$ such that u is adjacent to v. Again, as F_1 is a conditional fault set, there exists a vertex $w \notin F_1$ such that v is adjacent to w. Thus, by Theorem 1, F_1 and F_2 are distinguishable. □

4.2. Our general technique

We now describe our general technique to establish the conditional diagnosability of a graph $G = (V, E)$. In the next section, we demonstrate its efficacy with different classes of interconnection networks.

Let $F_1, F_2 \subseteq V$ be indistinguishable conditional fault sets so that $F_1 \neq F_2$ and both sets are of size at most $p_2(G)$. Our ultimate aim is to obtain a contradiction and thus, by Lemma 2, to show that the conditional diagnosability of G is exactly $p_2(G)$. By Lemma 4, it is not the case that $F_1 \subset F_2$ or $F_2 \subset F_1$. Consequently, there exists some $u \in F_1 \setminus F_2$ (and also some $u' \in F_2 \setminus F_1$; thus, $|F_1 \cap F_2| \leq p_2(G) - 1$). As F_2 is a conditional fault set, u has a neighbour v that is not in F_2 (and similarly u' has a neighbour v' that is not in F_1). We call our chosen vertex v u's corresponding partner vertex (and *vice versa*). The general situation can be visualised as in Fig. 2.

4.2.1. Establishing a K_2

The crux of our technique is to show that under certain circumstances and with the set-up as described in the previous paragraph, $G \setminus F_1$ and $G \setminus F_2$ both have connected components isomorphic to K_2 (we subsequently show how to use this fact to obtain a lower bound on the conditional diagnosability of G). To this end, suppose that $N_G(\{u, v\}) \not\subseteq F_2$; so, let

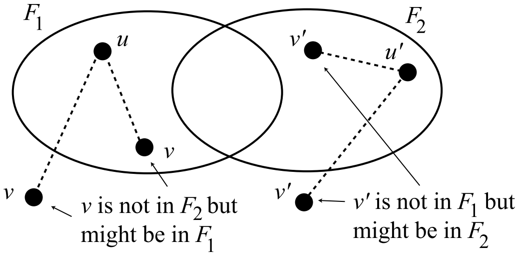

Fig. 2. The general situation.

$w \in N_G(\{u, v\}) \setminus F_2$ and let $T_3 = \{u, v, w\}$. The subgraph of G induced by the vertices of T_3 contains a path of length 2 of which one edge is (u, v).

Assume that there exists $b \geq 3$ and a connected subgraph of $G \setminus (F_1 \cap F_2)$ with vertex set T_b such that T_b contains b vertices including u, v and w and:

$$b + |N_G(T_b)| - (|F_2| + |F_1| - |F_1 \cap F_2|) > 0. \tag{1}$$

Denote $b + |N_G(T_b)| - (|F_2| + |F_1| - |F_1 \cap F_2|)$ by μ. We have that $|(T_b \cup N_G(T_b)) \setminus (F_1 \cup F_2)| \geq \mu > 0$. Suppose that $x \in (T_b \cup N_G(T_b)) \setminus (F_1 \cup F_2)$ and $y \notin F_1 \cup F_2$ are such that $(x, y) \in E$ and $d_G(x) \geq 2$. As the subgraph of G (actually, of $G \setminus (F_1 \cap F_2)$) induced by T_b is connected and $u \in T_b \cap (F_1 \setminus F_2)$, there is a path in $G \setminus (F_1 \cap F_2)$ for which the first two vertices are x and y (in some order) and for which the last vertex is u. By walking along this path we can find a path of length 3 so that the first two vertices lie outside $F_1 \cup F_2$ and the third vertex lies in $F_1 \setminus F_2$ or $F_2 \setminus F_1$. This yields a contradiction by Theorem 1. Hence, every vertex of $(T_b \cup N_G(T_b)) \setminus (F_1 \cup F_2)$ is adjacent only to vertices of $F_1 \cup F_2$ in G, with the consequence that every vertex of $(T_b \cup N_G(T_b)) \setminus (F_1 \cup F_2)$ is adjacent to some vertex of $F_1 \triangle F_2$ (recall that F_1 and F_2 are conditional fault sets). Thus, by Lemma 3, every vertex x of $(T_b \cup N_G(T_b)) \setminus (F_1 \cup F_2)$ is adjacent to: one vertex in $F_1 \setminus F_2$; one vertex in $F_2 \setminus F_1$; and $d_G(x) - 2$ vertices in $F_1 \cap F_2$. So, for every m such that $1 \leq m \leq \mu$, we must have that there are at least $m(\Delta(G) - 2) - c(G)\frac{m(m-1)}{2}$ distinct vertices in $F_1 \cap F_2$, and there is also at least 1 vertex in each of $F_1 \setminus F_2$ and $F_2 \setminus F_1$; that is, we must have that:

$$m(\Delta(G) - 2) - c(G)\frac{m(m-1)}{2} \leq |F_1 \cap F_2| \leq p_2(G) - 1. \tag{2}$$

If can obtain some T_b, as above, so that inequality (2) is violated, for some $m \in \{1, 2, \ldots, \mu\}$, then we obtain a contradiction and so must have that $N_G(\{u, v\}) \subseteq F_2$.

In order to obtain our contradiction (that is, to obtain T_b as required), it is feasible that we can iteratively build connected subgraphs with vertex sets $T_4, T_5, \ldots, T_b$ in $G \setminus (F_1 \cap F_2)$ so that for every $i \in \{4, 5, \ldots, b\}$, $T_{i-1} \subset T_i$. To this end, the following lemma provides a general lower bound on the size of the neighbourhood of T_i in terms of the size of the neighbourhood of T_{i-1}.

Lemma 5. *Fix $i \geq 4$. Let $T_{i-1} \subset T_i$ be subsets of vertices of some graph G so that T_{i-1} has size $i - 1$ and induces a connected subgraph of G, T_i has size i and induces a connected subgraph of G, and $T_{i-1} \cup \{x\} = T_i$. We have that*

$$|N_G(T_i)| \geq |N_G(T_{i-1})| + d_G(x) - (c(G) + 1)(i - 1).$$

Proof. Suppose that $z \in T_{i-1}$ and that $(x, z) \in E$. Let us count the neighbours of x in G. Each neighbour y of x in G has exactly one of 3 types:

- y lies in T_{i-1}, and there are at most $i - 1$ such neighbours (including the vertex z)
- $y \in N_G(T_{i-1})$ (and so $y \notin T_{i-1}$), and there are at most $c(G)(i - 1)$ such neighbours (as any vertex of T_{i-1} has at most $c(G)$ neighbours in common with x)
- $y \in N_G(T_i) \setminus N_G(T_{i-1})$ (trivially, $y \notin T_{i-1}$).

Thus, $|N_G(T_i)| - |N_G(T_{i-1})| \geq d_G(x) - (i - 1) - c(G)(i - 1) = d_G(x) - (c(G) + 1)(i - 1)$ and the result follows. □

We emphasise that our arguments above are intended to be as widely applicable as possible. For specific families of interconnection networks, the derived bounds and inequalities can be significantly tightened.

4.2.2. Having established a K_2

Suppose that we have proceeded as above and obtained that $N_G(\{u, v\}) \subseteq F_2$ (resp. $N_G(\{u', v'\}) \subseteq F_1$). Suppose also that our reasoning is such that our arguments apply equally well to any other vertex $u_1 \in F_1 \setminus (F_2 \cup \{u\})$ (resp. $u'_1 \in F_2 \setminus (F_1 \cup \{u'\})$) and its corresponding partner vertex $v_1 \notin F_2$ (resp. $v'_1 \notin F_1$). We are now in a position to possibly obtain an upper bound on $|F_1 \setminus F_2|$ (resp. $|F_2 \setminus F_1|$) and a lower bound on $|F_1 \cup F_2|$.

Let us assume that $|F_1 \setminus F_2| = \nu$. From above, for any $x \in F_1 \setminus F_2$, we have that exactly $d_G(x) - 1$ neighbours of x lie in F_2. Thus, if $1 \leq m \leq \nu$ then since $p_2(G) \geq |F_2|$, we have that

$$p_2(G) \geq m(\Delta(G) - 1) - c(G)\frac{m(m-1)}{2}. \tag{3}$$

Consequently, if some m for which $1 \leq m \leq \nu$ violates inequality (3) then we obtain a contradiction and we must have that $|F_1 \setminus F_2| < m$. An analogous statement can be made as regards $|F_2 \setminus F_1|$.

Let us assume that $|F_1 \triangle F_2| = \nu$. From above, for any $x \in F_1 \triangle F_2$, we have that at least $d_G(x) - 1$ neighbours of x lie in $F_1 \cup F_2$. Thus, if $1 \leq m \leq \nu$ then

$$|F_1 \cup F_2| \geq m(\Delta(G) - 1) - c(G)\frac{m(m-1)}{2}. \tag{4}$$

5. Applications

We now apply the methodology from the previous section. We begin with the hypercubes and the k-ary n-cubes, in order to illustrate how this methodology is applied and for which conditional diagnosability results have previously been obtained, before moving on to a range of other interconnection networks.

5.1. Hypercubes

Recall that it has already been shown independently in [13,33] that Q_n has conditional diagnosability $3n - 5$ when $n \geq 5$ (although this value was only established in [33] for $n \geq 7$).

It is easy to see that $p_2(Q_n) = 3n - 5$. Our basic assumption is that F_1 and F_2 are indistinguishable conditional fault sets in Q_n of size at most $p_2(Q_n) = 3n - 5$ such that $u \in F_1 \setminus F_2$ and $u' \in F_2 \setminus F_1$; so, in particular and with reference to the previous section, we have our vertex set $T_3 = \{u, v, w\}$. Assume further that $n \geq 29$ (we shall return to this assumption later). Note that $c(Q_n) = 2$ and that $\Delta(Q_n) = n$.

In the first phase of our reasoning, we apply the argument in Section 4.2.1. We have that $|F_1 \cap F_2| \leq 3n - 6$. So, $|N_{Q_n}(T_3)| > |F_1 \cap F_2|$ and we can build T_4 by augmenting T_3 with a vertex of $N_{Q_n}(T_3) \setminus (F_1 \cap F_2)$. By Lemma 5, $|N_{Q_n}(T_4)| \geq 4n - 14 > 3n - 6 \geq |F_1 \cap F_2|$ when $n > 8$. Build T_5 by augmenting T_4 with a vertex of $N_{Q_n}(T_4) \setminus (F_1 \cap F_2)$. By Lemma 5, $|N_{Q_n}(T_5)| \geq 5n - 26 > 3n - 6 \geq |F_1 \cap F_2|$ when $n > 10$. Continuing in this way yields T_{10} such that $|N_{Q_n}(T_{10})| \geq 10n - 131$ when $n > 17$. With reference to inequality (1), $\mu \geq 4n - 111 > 4$ when $n \geq 29$ (here, we use the fact that $|F_1 \cup F_2| \leq 2p_2(Q_n)$). Putting $m = 4$ in inequality (2) yields that $n \leq 14$ and so we obtain a contradiction. Thus, $N_{Q_n}(\{u, v\}) \subseteq F_2$.

In the second phase, we apply the argument in Section 4.2.2. There is nothing special about starting from the vertex u, above: if $u_1 \in F_1 \setminus (F_2 \cup \{u\})$ then we can proceed identically. Thus, if such a vertex u_1 exists then u_1 has some neighbour v_1 that is not in F_2 so that $N_{Q_n}(\{u_1, v_1\}) \subseteq F_2$. An analogous statement can be made as regards a vertex $u_1' \in F_2 \setminus (F_1 \cup \{u'\})$. Suppose that $|F_1 \setminus F_2| \geq 4$. Consequently, from inequality (3), $3n - 5 \geq 4(n-1) - 12 = 4n - 16$, which yields a contradiction when $n \geq 12$. Thus, we must have that $1 \leq |F_1 \setminus F_2| \leq 3$, and similarly that $1 \leq |F_2 \setminus F_1| \leq 3$; consequently, $|F_1 \cup F_2| \leq 3n - 2$. Further, if $|F_1 \triangle F_2| \geq 4$ then from inequality (4), $|F_1 \cup F_2| \geq 4n - 16$, which yields a contradiction when $n \geq 15$. Hence, $2 \leq |F_1 \triangle F_2| \leq 3$ with $|F_1 \cup F_2| \leq 3n - 4$.

In the third phase, we use the bound on $|F_1 \cup F_2|$ just established, in conjunction with some simple counting arguments, to obtain our contradiction. Suppose that $\{u, v\} \cap \{u', v'\} = \emptyset$. As $N_{Q_n}(\{u, v\}) \subseteq F_2$, $N_{Q_n}(\{u', v'\}) \subseteq F_1$ and $c(Q_n) = 2$, we must have that $4(n-1) - 8 \leq |F_1 \cup F_2| \leq 3n - 4$ (note that u and v have no neighbours in common, and nor do u' and v', as Q_n is bipartite). This yields a contradiction, and so we must have that $v = v'$. However, $N_{Q_n}(\{u, v, u'\}) \subseteq F_1 \cup F_2$ and in addition $u, u' \in F_1 \cup F_2$; thus, $|F_1 \cup F_2| \geq (3n - 5) + 2 = 3n - 3$, which yields a contradiction. Hence, if $n \geq 29$ then we have that Q_n has conditional diagnosability $3n - 5$; that is, the family of hypercubes $\{Q_n\}$ has asymptotic conditional diagnosability $3n - 5$.

Remark 6. Let us remark upon our initial assumption that n should be at least 29. We have chosen n to make it as small as possible yet so that the above arguments hold (essentially, with reference to above, we need μ to be at least 4 in order to obtain that $N_{Q_n}(\{u, v\}) \subseteq F_2$). We could have worked with T_7, for example, instead of T_{10}, but this would have required that $n \geq 46$. This can be calculated easily by hand but we actually employ a simple computer program to show that forcing n to be at least 29 is the best we can do (without employing a more detailed analysis than that in Sections 4.2.1 and 4.2.2 that is specific to hypercubes). We use this same computer program in the same way for the interconnection networks we consider below.

Remark 7. Note that in applying our techniques so as to show that the family of hypercubes has asymptotic conditional diagnosability $3n - 5$, essentially the only structural properties of Q_n that we use are that $\Delta(Q_n) = n$, $c(Q_n) = 2$ and $p_2(Q_n) = 3n - 5$ (we also use the fact that Q_n is bipartite which, as it happens, we need not have used). In particular, if any other family of interconnection networks $\{X_n\}$ is such that $\Delta(X_n) = n$, $c(X_n) = 2$ and $p_2(X_n) = 3n - 5$ then we immediately obtain that $\{X_n\}$ has asymptotic conditional diagnosability $3n - 5$ too.

5.2. k-ary n-cubes

The k-ary n-cube Q_n^k, where $k \geq 3$ and $n \geq 2$, is defined as follows: it has vertex set $\{0, 1, \ldots, k-1\}^n$; and there is an edge $(u_1u_2\ldots u_n, v_1v_2\ldots v_n)$ if, and only if, there exists $i \in \{1, 2, \ldots, n\}$ such that $u_j = v_j$, for all $j \in \{1, 2, \ldots, n\} \setminus \{i\}$, and $u_i - v_i \in \{+1, -1\} (\text{mod } k)$. Recall that it has already been shown in [12] that Q_n^k has conditional diagnosability $6n - 5$ when $n \geq 4$ and $k \geq 4$.

Suppose that $k \geq 4$. It is easy to see that $p_2(Q_n^k) = 6n - 5$. As before, our basic assumption is that F_1 and F_2 are indistinguishable conditional fault sets in Q_n^k of size at most $p_2(Q_n^k) = 6n - 5$ such that $u \in F_1 \setminus F_2$ and $u' \in F_2 \setminus F_1$. Assume further that $n \geq 15$. Note that $c(Q_n^k) = 2$ and that $\Delta(Q_n^k) = 2n$.

Our first phase of reasoning proceeds similarly to as in the case of the hypercubes. We have that $|F_1 \cap F_2| \leq 6n - 6$; so, $|N_{Q_n^k}(T_3)| > |F_1 \cap F_2|$ and we can build T_4 by augmenting T_3 with a vertex not in $F_1 \cap F_2$. By Lemma 5, $|N_{Q_n^k}(T_4)| \geq 8n - 14 > 6n - 6 \geq |F_1 \cap F_2|$ when $n > 4$. Build T_5 by augmenting T_4 with a vertex not in $F_1 \cap F_2$. By Lemma 5, $|N_{Q_n^k}(T_5)| \geq 10n - 26 > 6n - 6 \geq |F_1 \cap F_2|$ when $n > 5$. Build T_6 by augmenting T_5 with a vertex not in $F_1 \cap F_2$. By Lemma 5, $|N_{Q_n^k}(T_6)| \geq 12n - 41 > 6n - 6 \geq |F_1 \cap F_2|$ when $n > 5$. Continuing in this way yields T_9 such that $|N_{Q_n^k}(T_9)| \geq 18n - 104$ when $n > 8$. With reference to inequality (1), $\mu \geq 6n - 85 > 4$ when $n \geq 15$. Putting $m = 4$ in inequality (2) yields that $8n - 20 \leq 6n - 6$ and so we obtain a contradiction. Thus, $N_{Q_n^k}(\{u, v\}) \subseteq F_2$.

In the second phase, we apply the argument in Section 4.2.2. There is nothing special about starting from the vertex u, above: if $u_1 \in F_1 \setminus (F_2 \cup \{u\})$ then we can proceed identically. Thus, if such a vertex u_1 exists then u_1 has some neighbour v_1 that is not in F_2 so that $N_{Q_n^k}(\{u_1, v_1\}) \subseteq F_2$. An analogous statement can be made as regards a vertex $u_1' \in F_2 \setminus (F_1 \cup \{u'\})$. Suppose that $|F_1 \setminus F_2| \geq 4$. Consequently, from inequality (3), $6n - 5 \geq 4(2n - 1) - 12 = 8n - 16$, which yields a contradiction when $n \geq 6$. Thus, we must have that $1 \leq |F_1 \setminus F_2| \leq 3$, and similarly that $1 \leq |F_2 \setminus F_1| \leq 3$; consequently, $|F_1 \cup F_2| \leq 6n - 2$. Further, if $|F_1 \triangle F_2| \geq 4$ then from inequality (4), $|F_1 \cup F_2| \geq 8n - 16$, which yields a contradiction when $n \geq 8$. Hence, $2 \leq |F_1 \triangle F_2| \leq 3$ with $|F_1 \cup F_2| \leq 6n - 4$.

In the third phase, we use the bound on $|F_1 \cup F_2|$ to obtain a contradiction. Suppose that $\{u, v\} \cap \{u', v'\} = \emptyset$. As $N_{Q_n^k}(\{u, v\}) \subseteq F_2$ and $N_{Q_n^k}(\{u', v'\}) \subseteq F_1$, we must have that $4(2n - 1) - 8 \leq |F_1 \cup F_2| \leq 6n - 4$ (note that u and v have no neighbours in common, and nor do u' and v', as Q_n^k has no cycles of length 3). This yields a contradiction, and so we must have that $v = v'$. However, $N_{Q_n^k}(\{u, v, u'\}) \subseteq F_1 \cup F_2$ and in addition $u, u' \in F_1 \cup F_2$; thus, $|F_1 \cup F_2| \geq (6n - 5) + 2 = 6n - 3$, which yields a contradiction. Hence, if $n \geq 15$ then we have that Q_n^k has conditional diagnosability $6n - 5$; that is, if $k \neq 3$ then the family of hypercubes $\{Q_n^k\}$ has asymptotic conditional diagnosability $6n - 5$.

Our approach as regards $\{Q_n^3\}$ follows the usual phases of reasoning. Suppose that $k = 3$.

Lemma 8. *When $n \geq 2$, $p_2(Q_n^3) = 6n - 7$.*

Proof. Let $\rho = (x, z, y)$ be a path of length 2. As Q_n^3 is edge-transitive [2], we may assume that $x = 00\ldots0$ and $z = 10\ldots0$. Consequently, w.l.o.g. we need look only at the cases when y is: $20\ldots0$; $110\ldots0$; and $120\ldots0$. The vertices: x and z have only $20\ldots0$ as a common neighbour; x and $20\ldots0$ have only z as a common neighbour; x and $110\ldots0$ have z and $010\ldots0$ as common neighbours; x and $120\ldots0$ have z and $020\ldots0$ as common neighbours; z and $20\ldots0$ have x as a common neighbour; z and $110\ldots0$ have $120\ldots0$ as a common neighbour; and z and $120\ldots0$ have $110\ldots0$ as a common neighbour. Consequently, $p_2(Q_n^3) = 6n - 7$ (witnessed by both $\rho = (x, z, 110\ldots0)$ and $\rho = (x, z, 120\ldots0)$). □

So, we have that $p_2(Q_n^3) = 6n - 7$, $c(Q_n^3) = 2$ and $\Delta(Q_n^3) = 2n$. Proceeding exactly as we did above but with these parameters and with $n \geq 15$, we obtain that T_9 is such that $|N_{Q_n^3}(T_9)| \geq 18n - 106$ (as $n > 8$), with the result that $\mu \geq 6n - 83 > 4$ (as $n \geq 15$). Putting $m = 4$ in inequality (2) yields that $8n - 20 \leq 6n - 8$ and so we obtain a contradiction. Thus, $N_{Q_n^3}(\{u, v\}) \subseteq F_2$.

Now we apply the argument in Section 4.2.2. There is nothing special about starting from the vertex u, above: if $u_1 \in F_1 \setminus (F_2 \cup \{u\})$ then we can proceed identically. Thus, if such a vertex u_1 exists then u_1 has some neighbour v_1 that is not in F_2 so that $N_{Q_n^3}(\{u_1, v_1\}) \subseteq F_2$. An analogous statement can be made as regards a vertex $u_1' \in F_2 \setminus (F_1 \cup \{u'\})$. Suppose that $|F_1 \setminus F_2| \geq 4$. Consequently, from inequality (3), $6n - 7 \geq 4(2n - 1) - 12 = 8n - 16$, which yields a contradiction. Thus, we must have that $1 \leq |F_1 \setminus F_2| \leq 3$, and similarly that $1 \leq |F_2 \setminus F_1| \leq 3$; consequently, $|F_1 \cup F_2| \leq 6n - 4$. Further, if $|F_1 \triangle F_2| \geq 4$ then from inequality (4), $6n - 4 \geq |F_1 \cup F_2| \geq 8n - 16$, which yields a contradiction. Hence, $2 \leq |F_1 \triangle F_2| \leq 3$ with $|F_1 \cup F_2| \leq 6n - 6$.

Suppose that $\{u, v\} \cap \{u', v'\} = \emptyset$. As $N_{Q_n^3}(\{u, v\}) \subseteq F_2$ and $N_{Q_n^3}(\{u', v'\}) \subseteq F_1$, we must have that $4(2n - 1) - 10 \leq |F_1 \cup F_2| \leq 6n - 6$ (note that u and v have only 1 common neighbour, as do u' and v'). This yields a contradiction, and so we must have that $v = v'$. However, $N_{Q_n^3}(\{u, v, u'\}) \subseteq F_1 \cup F_2$ and in addition $u, u' \in F_1 \cup F_2$; thus, $|F_1 \cup F_2| \geq (6n - 7) + 2 = 6n - 5$, which yields a contradiction. Hence, if $n \geq 15$ then we have that Q_n^3 has conditional diagnosability $6n - 7$; that is, the family of 3-ary n-cubes $\{Q_n^3\}$ has asymptotic conditional diagnosability $6n - 7$ (we remark that this result is new in that the results from [12] only apply to k-ary n-cubes when $k \geq 4$).

6. Some new results

We now use our methodology to establish conditional diagnosability results for some interconnection networks G for which hitherto no such results were known. We proceed as we did for the hypercubes and the k-ary n-cubes; namely, our basic assumption is that F_1 and F_2 are indistinguishable conditional fault sets in G of size at most $p_2(G)$ such that

$u \in F_1 \setminus F_2$ and $u' \in F_2 \setminus F_1$. So, in particular and with reference to the previous section, we have our graph T_3 with vertex set $\{u, v, w\}$. We make additional assumptions on n as appropriate. Our analysis is in three phases, as before: we first prove that $N_G(\{u, v\}) \subseteq F_2$; we then obtain a bound on $|F_1 \cup F_2|$; and we then establish a contradiction. Our applications are repetitive and so we only outline the essential numeric details within each phase.

6.1. Folded hypercubes

The *folded hypercube* FQ_n [7] is obtained by adding certain edges to Q_n: for every vertex $x_1x_2 \ldots, x_n$ of Q_n, we add the edge $(x_1x_2 \ldots x_n, \bar{x}_1\bar{x}_2 \ldots \bar{x}_n)$. Clearly, $\Delta(FQ_n) = n + 1$.

Lemma 9. *When $n \geq 4$, $c(FQ_n) = 2$ and $p_2(FQ_n) = 3n - 2$.*

Proof. It is easy to show that the folded hypercube FQ_n is vertex-transitive (see, for example, [29]). Hence, w.l.o.g. in order to find $c(FQ_n)$ and $p_2(FQ_n)$ it suffices to look at the paths (x, z, y) where: $x = 0 \ldots 0$, $z = 10 \ldots 0$ and $y = 110 \ldots 0$; and $x = 0 \ldots 0$, $z = 10 \ldots 0$ and $y = 01 \ldots 1$. The only other path of length 2 from x to y is: $(x, w = 010 \ldots 0, y)$ in the first case; and $(x, w = 1 \ldots 1, y)$ in the second case. In both cases: (x, z, y, w) is a cycle of length 4; x and y have z and w as their only common neighbours; z and w have x and y as their only common neighbours; x and z have no common neighbours; and z and y have no common neighbours. Hence, $c(FQ_n) = 2$ and $N_{FQ_n}(\{x, y, z\}) = 3n - 2$. $\square$

We remark that the conditional diagnosability of a folded hypercube has been studied but only under the PMC model when it was shown to be $4n - 3$ when $n = 5$ or $n \geq 8$ [37].

Assume that $n \geq 28$. In the first phase, we build T_{10} so that $N_{FQ_n}(T_{10}) \geq 10n-121$ and hence so that $\mu \geq 4n-107$. Thus, as $n \geq 28$, we must have that $\mu > 4$. Putting $\mu = 4$ in inequality (2) yields that $4n - 16 \leq 3n - 3$, which yields a contradiction. Thus, $N_{FQ_n}(\{u, v\}) \subseteq F_2$. In the second phase, suppose that $|F_1 \setminus F_2| \geq 4$. Consequently, from inequality (3), $3n - 2 \geq 4n - 12$, which yields a contradiction. Thus, we must have that $1 \leq |F_1 \setminus F_2| \leq 3$, and similarly that $1 \leq |F_2 \setminus F_1| \leq 3$; consequently, $|F_1 \cup F_2| \leq 3n + 1$. Further, if $|F_1 \triangle F_2| \geq 4$ then from inequality (4), $|F_1 \cup F_2| \geq 4n - 12$, which yields a contradiction. Hence, $2 \leq |F_1 \triangle F_2| \leq 3$ with $|F_1 \cup F_2| \leq 3n - 1$. In the third phase, suppose that $\{u, v\} \cap \{u', v'\} = \emptyset$. As $N_{FQ_n}(\{u, v\}) \subseteq F_2$ and $N_{FQ_n}(\{u', v'\}) \subseteq F_1$, we must have that $4n - 8 \leq |F_1 \cup F_2| \leq 3n - 2$ (note that u and v have no neighbours in common, and nor do u' and v'). This yields a contradiction, and so we must have that $v = v'$. However, $N_{FQ_n}(\{u, v, u'\}) \subseteq F_1 \cup F_2$ and in addition $u, u' \in F_1 \cup F_2$; thus, $|F_1 \cup F_2| \geq 3n - 2 + 2 = 3n$, which yields a contradiction. Hence, if $n \geq 28$ then we have that FQ_n has conditional diagnosability $3n - 2$; that is, the family of folded hypercubes $\{FQ_n\}$ has asymptotic conditional diagnosability $3n - 2$.

6.2. Pancake graphs

The *pancake graph* P_n [1] has vertex set S_n consisting of all permutations of $\{1, 2, \ldots, n\}$ and there is an edge joining $u_1u_2 \ldots u_n$ and $v_1v_2 \ldots v_n$ if, and only if, there exists some $i \in \{2, 3, \ldots, n\}$ such that $v_1v_2 \ldots v_n = u_iu_{i-1} \ldots u_1u_{i+1}u_{i+2} \ldots u_n$; that is, $v_1v_2 \ldots v_n$ is obtained from $u_1u_2 \ldots u_n$ by 'reversing' a prefix of $u_1u_2 \ldots u_n$. Trivially, P_n is regular of degree $n - 1$. It is not difficult to prove that when $n \geq 3$, the pancake graph P_n has girth 6 (an explicit proof is given in [25]); consequently, when $n \geq 3$ we have that $p_2(P_n) = 3n - 7$ and $c(G) = 1$.

Assume that $n \geq 20$. In the first phase, we build T_9 so that $N_{P_n}(T_9) \geq 9n - 79$ and hence so that $\mu \geq 3n - 56$. Thus, we must have that $\mu \geq 4$. Putting $\mu = 4$ in inequality (2) yields that $4n - 18 \leq 3n - 8$, which yields a contradiction. Thus, $N_{P_n}(\{u, v\}) \subseteq F_2$. In the second phase, suppose that $|F_1 \setminus F_2| \geq 4$. Consequently, from inequality (3), $3n - 7 \geq 4n - 14$, which yields a contradiction. Thus, we must have that $1 \leq |F_1 \setminus F_2| \leq 3$, and similarly that $1 \leq |F_2 \setminus F_1| \leq 3$; consequently, $|F_1 \cup F_2| \leq 3n - 4$. Further, if $|F_1 \triangle F_2| \geq 4$ then from inequality (4), $|F_1 \cup F_2| \geq 4n - 14$, which yields a contradiction. Hence, $2 \leq |F_1 \triangle F_2| \leq 3$ with $|F_1 \cup F_2| \leq 3n - 6$. In the third phase, suppose that $\{u, v\} \cap \{u', v'\} = \emptyset$. As $N_{P_n}(\{u, v\}) \subseteq F_2$ and $N_{P_n}(\{u', v'\}) \subseteq F_1$, we must have that $4n - 8 \leq |F_1 \cup F_2| \leq 3n - 6$ (note that u and v have no neighbours in common, and nor do u' and v'). This yields a contradiction, and so we must have that $v = v'$. However, $N_{P_n}(\{u, v, u'\}) \subseteq F_1 \cup F_2$ and in addition $u, u' \in F_1 \cup F_2$; thus, $|F_1 \cup F_2| \geq 3n - 7 + 2 = 3n - 5$, which yields a contradiction. Hence, if $n \geq 20$ then we have that P_n has conditional diagnosability $3n - 7$; that is, the family of pancake graphs $\{P_n\}$ has asymptotic conditional diagnosability $3n - 7$.

6.3. Augmented cubes

The *augmented cube* AQ_n [4] is obtained by adding certain edges to Q_n. We call the edges of Q_n, within AQ_n, the *b-edges*, to denote that they result from flipping one bit, and we call the additional edges the *s-edges*, to denote that they result from flipping a suffix of bits. In more detail, for every vertex $x = x_1x_2 \ldots x_n$ of AQ_n and for every $s \in \{1, 2, \ldots, n - 1\}$, there is an s-edge $(x_1x_2 \ldots x_n, x_1x_2 \ldots x_{s-1}\bar{x}_s \ldots \bar{x}_n)$. In particular, AQ_n is regular of degree $2n - 1$.

In order to apply the techniques of the previous section, we need to ascertain $p_2(AQ_n)$ and $c(AQ_n)$. As we shall see below, doing so is not always as straightforward as it has been hitherto. In order to obtain these values, we need to examine the different types of paths of length 2 that can arise within AQ_n. (We remark that the conditional diagnosability of an augmented cube has been studied under the PMC model and shown to be $8n - 27$ when $n \geq 5$ [3].)

Theorem 10. *For the augmented cube AQ_n, where $n \geq 5$, we have that $c(AQ_n) = 4$ and $p_2(AQ_n) = 6n - 17$.*

Proof. Let $\rho = (x, z, y)$ be an arbitrary path. As the augmented cube AQ_n is vertex-transitive [4], w.l.o.g. we may assume that the vertex x is $00\ldots0$. Our path ρ has one of four types depending upon the types of the two edges involved. We consider these paths according to their types. For every such path ρ, what we do below is examine this path and see whether there is also an edge (x, y) and whether there are any other paths of length 2 from x to y (we call such paths 2-*paths*).

Before we begin, we note that every path $\rho = (x, z, y)$ has a *dual* path, namely the path obtained by 'reversing the operations' corresponding to the edges (x, z) and (z, y). So, for example, if the operation corresponding to the b-edge (x, z) is to flip the bth bit and the operation corresponding to the s-edge (z, y) is to flip the bank of bits from the sth up to the nth then the dual path of ρ is obtained by starting from x and first flipping the bank of bits from the sth up to the nth to get the vertex z' and then flipping the bth bit of z' to get y. The dual path is always a different 2-path from the original path.

In what follows, we use subscripts to denote specific bits of vertices; for example, $0\ldots01_s1\ldots1$ denotes the vertex where the first $s-1$ bits are 0 and the last $n-(s-1)$ bits are 1.

Case (a): Suppose that we have a path $\rho = (x, z, y)$ so that (x, z) is an s-edge and (z, y) is a b-edge; so, $z = 0\ldots01_s1\ldots1$, with $1 \leq s \leq n-1$ (if (x, z) is a b-edge and (z, y) is an s-edge then we simply interchange the roles of x and y). Suppose that y is obtained from z by flipping bit b.

Sub-case (i): $b < s$, and so the weight of z is $(n+1)-s+1 = n+2-s \geq 3$. Every path of length 2 from x to y must contain exactly one s-edge (if it consists of two s-edges then we obtain a contradiction as we would have $1 = y_n = x_n = 0$, and if it contains no s-edge then y would have weight 2).

1. If $b = s-1$ then we have that $y = 0\ldots01_{s-1}1\ldots1$ and there is an edge (x, y) (which is an s-edge). Apart from ρ and its dual path $(x, 0\ldots01_{s-1}0\ldots0, y)$, there are also 2-paths $(x, 0\ldots01_{s-2}0\ldots0, y)$ and $(x, 0\ldots01_{s-2}1\ldots1, y)$ (assuming that $s \geq 3$).
2. If $b \leq s-2$ then $y = 0\ldots01_b0\ldots01_s1\ldots1$. There is no edge (x, y). Apart from ρ and its dual path $(x, 0\ldots01_b0\ldots0, y)$, in the case that $b = s-2$ only there are also 2-paths $(x, 0\ldots01_{s-1}0\ldots0, y)$ and $(x, 0\ldots01_{s-2}1\ldots1, y)$.

Sub-case (ii): $b = s$, and so $y = 0\ldots01_{s+1}1\ldots1$. There is an edge (x, y), which is an s-edge if $s \leq n-2$ and a b-edge if $s = n-1$.

1. If $s \leq n-2$ then apart from the path ρ and its dual path $(x, 0\ldots01_s0\ldots0, y)$, there are also 2-paths $(x, 0\ldots01_{s+1}0\ldots0, y)$ and $(x, 0\ldots01_{s+2}1\ldots1, y)$.
2. If $s = n-1$ then apart from the dual path $(x, 0\ldots010, y)$, there are no other 2-paths.

Sub-case (iii): $b > s$, and so the weight of z is $(n+1)-s-1 = n-s$. We have that $y = 0\ldots01_s\ldots10_b1\ldots1$.

1. If $s \leq n-3$ then (as above) every path of length 2 from x to y must contain exactly one s-edge. Apart from ρ and its dual path $(x, 0\ldots01_b0\ldots0, y)$, in the case that $b = s+1$ only (when $y = 0\ldots01_s01\ldots1$) are there 2-paths $(x, 0\ldots01_s0\ldots0, y)$ and $(x, 0\ldots01_{s+2}1\ldots1, y)$.
2. If $s = n-2$ then $y = 0\ldots0101$ or $y = 0\ldots0110$. Apart from the path ρ and its dual path $(x, 0\ldots0010, y)$ or $(x, 0\ldots0001, y)$, respectively, when $y = 0\ldots0101$ there are 2-paths $(x, 0\ldots0100, y)$ and $(x, 0\ldots0001, y)$, and when $y = 0\ldots0110$ there are 2-paths $(x, 0\ldots0100, y)$ and $(x, 0\ldots0010, y)$.
3. If $s = n-1$ then $y = 0\ldots0010$. Apart from ρ and its dual path $(x, 0\ldots0011, y)$, there are no other 2-paths although there is an edge (x, y).

Case (b): Suppose that we have a path $\rho = (x, z, y)$ so that (x, z) and (z, y) are both s-edges where $z = 0\ldots01_s1\ldots1$ and $y = 0\ldots01_s1\ldots10_t0\ldots0$ (if $s > t$ then we simply interchange the roles of x and y).

1. Suppose that $t-s \geq 3$. There are no 2-paths apart from ρ and its dual path $(x, 0\ldots01_t\ldots1, y)$ (if there were a 2-path then both edges would need to be s-edges and this is impossible).
2. Suppose we have that $t = s+2$. Apart from ρ and its dual path $(x, 0\ldots01_{s+2}1\ldots1, y)$, there are 2-paths $(x, 0\ldots01_s0\ldots0, y)$ and $(x, 0\ldots01_{s+1}0\ldots0, y)$.
3. Suppose that $t = s+1$. Apart from ρ and its dual path $(x, 0\ldots01_{s+1}1\ldots1, y)$, there are no other 2-paths although there is an edge (x, y).

Case (c): Suppose that we have a path $\rho = (x, z, y)$ so that (x, z) and (z, y) are both b-edges where $z = 0\ldots01_b0\ldots0$ and $y = 0\ldots01_b0\ldots01_{b'}0\ldots0$ (if $b' < b$ then we simply interchange the roles of x and y).

1. Suppose that $b'-b \geq 2$. There are no 2-paths apart from ρ and its dual path $(x, 0\ldots01_{b'}0\ldots0, y)$, unless $b = n-2$ and $b' = n$ when $y = 0\ldots0101$ and there are 2-paths $(x, 0\ldots0010, y)$ and $(x, 0\ldots0111, y)$.
2. Suppose that $b' = b+1$ and $b' \leq n-1$. Apart from ρ and its dual path $(x, 0\ldots01_{b+1}0\ldots0, y)$, there are 2-paths $(x, 0\ldots01_b1\ldots1, y)$ and $(x, 0\ldots01_{b+2}1\ldots1, y)$.
3. Suppose that $b' = b+1$ and $b' = n$, when $y = 0\ldots011$. Apart from ρ and its dual path $(x, 0\ldots01, y)$, there are 2-paths $(x, 0\ldots0100, y)$ and $(x, 0\ldots0111, y)$. There is also an edge (x, y).

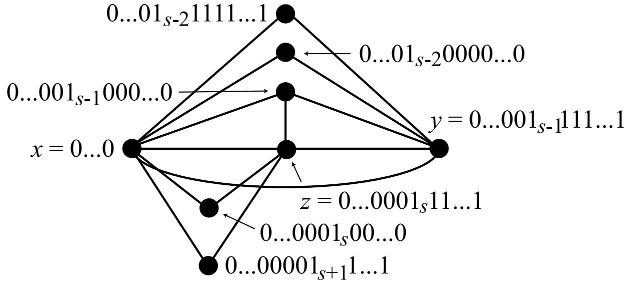

Fig. 3. Case (a)(i)(1) (where $b = s - 1$).

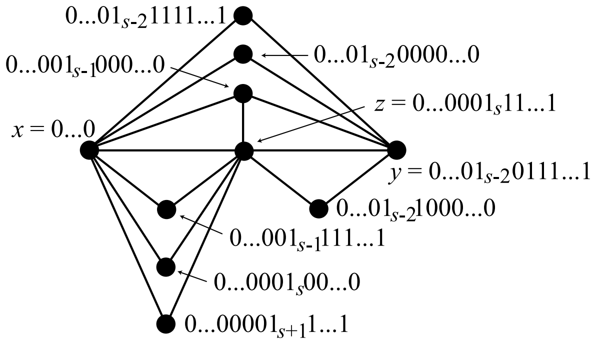

Fig. 4. The first sub-case of Case (a)(i)(2) where $b = s - 2$.

By inspecting the different cases above, we see that $c(AQ_n) = 4$.

We can now use the above classification to obtain $p_2(AQ_n)$. For each path $\rho = (x, z, y)$ in AQ_n of length 2, of one of the types above, we need to consider $N_{AQ_n}(\{x, y, z\})$. So, not only do we have to consider the neighbours common to x and y, which are readily available from above, we also need to consider the neighbours common to x and z and also to z and y (bearing in mind that some vertex might be a neighbour of each of x, y and z). As it turns out, this means splitting some of the cases above into sub-cases. Recall that we are looking for a path $\rho = (x, z, y)$ which minimises $|N_{AQ_n}(\{x, y, z\})|$.

Consider Case (a)(i)(1) where $b = s - 1$. The common neighbours can be listed as follows, ignoring that $x = 0 \ldots 0$ (resp. $y = 0 \ldots 01_{s-1}1 \ldots 1$, $z = 0 \ldots 01_s 1 \ldots 1$) is a common neighbour of y and z (resp. x and z, x and y):

- x, y: $0 \ldots 01_{s-1}0 \ldots 0$; $0 \ldots 01_{s-2}0 \ldots 0$; $0 \ldots 01_{s-2}1 \ldots 1$
- x, z: $0 \ldots 01_{s-1}0 \ldots 0$; $0 \ldots 01_s 0 \ldots 0$; $0 \ldots 01_{s+1}1 \ldots 1$
- y, z: $0 \ldots 01_{s-1}0 \ldots 0$.

Of course, we are assuming that $s \geq 3$ above (if not then there is a reduction in common neighbours and the size of $N_{AQ_n}(\{x, y, z\})$ increases). Note that we are using our classification above, of the neighbours common to x and y, to determine the neighbours common to x and z and y and z too. There are repetitions above. We can picture the edges involving the vertices of $\{x, y, z\}$ and any common neighbours as in Fig. 3. Consequently, we have that $|N_{AQ_n}(\{x, y, z\})| = 6n - 15$.

Consider Case (a)(i)(2) where $b \leq s-2$. We need to split this case into two sub-cases. The first sub-case is when $b = s-2$. The common neighbours can be listed as follows, ignoring that $z = 0 \ldots 01_s 1 \ldots 1$ is a common neighbour of $x = 0 \ldots 0$ and $y = 0 \ldots 01_{s-2}01_s 1 \ldots 1$:

- x, y: $0 \ldots 01_{s-1}0 \ldots 0$; $0 \ldots 01_{s-2}0 \ldots 0$; $0 \ldots 01_{s-2}1 \ldots 1$
- x, z: $0 \ldots 01_{s-1}1 \ldots 1$; $0 \ldots 01_{s-1}0 \ldots 0$; $0 \ldots 01_s 0 \ldots 0$; $0 \ldots 01_{s+1}1 \ldots 1$
- y, z: $0 \ldots 01_{s-1}0 \ldots 0$; $0 \ldots 01_{s-2}10 \ldots 0$.

We can picture the edges involving the vertices of $\{x, y, z\}$ and any common neighbours as in Fig. 4. So, we have that $|N_{AQ_n}(\{x, y, z\})| = 6n - 15$.

The second sub-case of Case $(a)(i)(2)$ is when $b \leq s - 3$. The common neighbours can be listed as follows, ignoring that $z = 0 \ldots 01_s 1 \ldots 1$ is a common neighbour of $x = 0 \ldots 0$ and $y = 0 \ldots 01_b 0 \ldots 01_s 1 \ldots 1$:

- x, y: $0 \ldots 01_b 0 \ldots 0$
- x, z: $0 \ldots 01_{s-1}1 \ldots 1$; $0 \ldots 01_{s-1}0 \ldots 0$; $0 \ldots 01_s 0 \ldots 0$; $0 \ldots 01_{s+1}1 \ldots 1$
- y, z: $0 \ldots 01_b 1 \ldots 10_s 0 \ldots 0$; $0 \ldots 01_{b+1}1 \ldots 10_s 0 \ldots 0$.

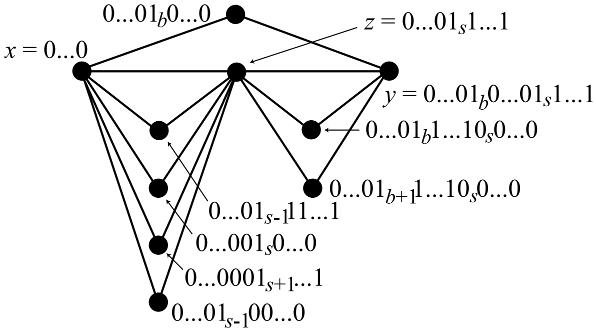

Fig. 5. The second sub-case of Case (a)(i)(2) where $b \leq s - 3$.

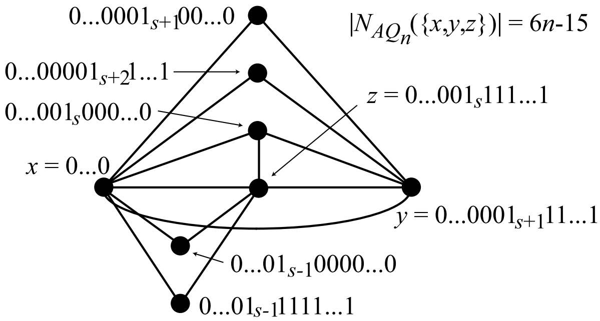

Fig. 6. Case (a)(ii)(1) where $s \leq n - 2$.

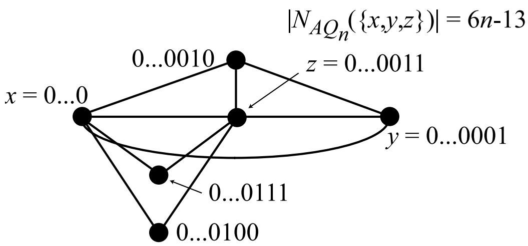

Fig. 7. Case (a)(ii)(2) where $s = n - 1$.

We can picture the edges involving the vertices of $\{x, y, z\}$ and any common neighbours as in Fig. 5. So, we have that $|N_{AQ_n}(\{x, y, z\})| = 6n - 14$.

Henceforth, for brevity, we only give the figure corresponding to each of the cases in our classification (or sub-case if necessary) together with the size of the corresponding $N_{AQ_n}(\{x, y, z\})$ (Figs. 6–21).

In consequence, we have that $p_2(AQ_n) = 6n - 17$ and the result follows. $\square$

Assume that $n \geq 25$. In the first phase, we build T_9 so that $N_{AQ_n}(T_9) \geq 18n - 188$ and hence so that $\mu \geq 6n - 145$. Thus, as $n \geq 25$, we must have that $\mu > 4$. Putting $\mu = 4$ in inequality (2) yields that $8n - 36 \leq 6n - 18$, which yields a contradiction. Thus, $N_{AQ_n}(\{u, v\}) \subseteq F_2$. In the second phase, suppose that $|F_1 \setminus F_2| \geq 4$. Consequently, from inequality (3), $6n - 17 \geq 8n - 32$, which yields a contradiction. Thus, we must have that $1 \leq |F_1 \setminus F_2| \leq 3$, and similarly that $1 \leq |F_2 \setminus F_1| \leq 3$; consequently, $|F_1 \cup F_2| \leq 6n - 14$. Further, if $|F_1 \triangle F_2| \geq 4$ then from inequality (4), $|F_1 \cup F_2| \geq 8n - 32$, which yields a contradiction. Hence, $2 \leq |F_1 \triangle F_2| \leq 3$ with $|F_1 \cup F_2| \leq 6n - 16$. In the third phase, suppose that $\{u, v\} \cap \{u', v'\} = \emptyset$. As $N_{AQ_n}(\{u, v\}) \subseteq F_2$ and $N_{AQ_n}(\{u', v'\}) \subseteq F_1$, we must have that $8n - 24 \leq |F_1 \cup F_2| \leq 6n - 16$. This yields a contradiction, and so we must have that $v = v'$. However, $N_{AQ_n}(\{u, v, u'\}) \subseteq F_1 \cup F_2$ and in addition $u, u' \in F_1 \cup F_2$; thus, $|F_1 \cup F_2| \geq 6n - 17 + 2 = 6n - 15$, which yields a contradiction. Hence, if $n \geq 25$ then we have that AQ_n has conditional diagnosability $6n - 17$; that is, the family of augmented cubes $\{AQ_n\}$ has asymptotic conditional diagnosability $6n - 17$.

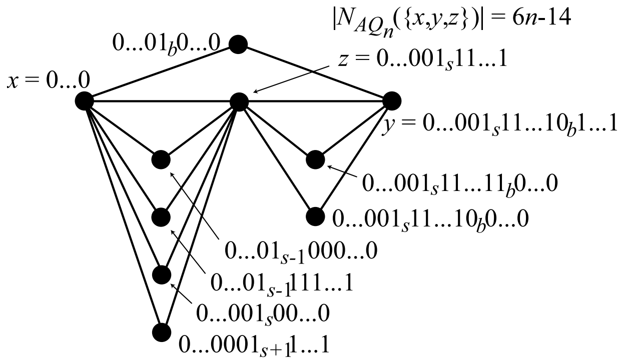

Fig. 8. The first sub-case of Case (a)(iii)(1) where $s \leq n-3$ and $s+2 \leq b \leq n-1$.

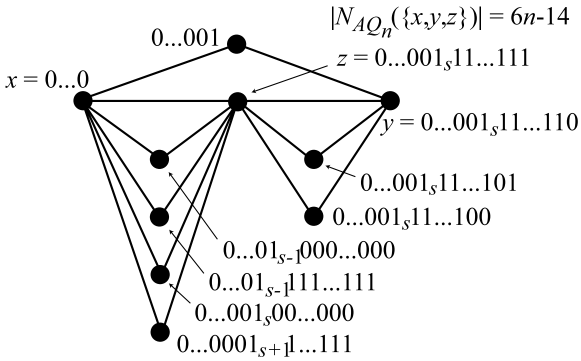

Fig. 9. The second sub-case of Case (a)(iii)(1) where $s \leq n-3$ and $b = n$.

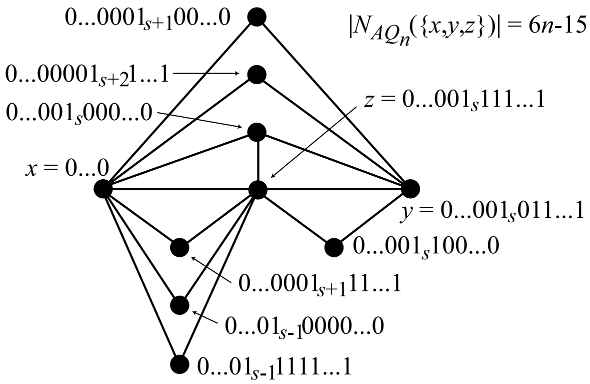

Fig. 10. The third sub-case of Case (a)(iii)(1) where $s \leq n-3$ and $b = s+1$.

7. Conclusions

In this paper we have developed and applied a powerful method for ascertaining the (asymptotic) conditional diagnosability of interconnection networks under the comparison diagnosis model. Our method only relies upon the combinatorial content of certain parameters associated with an interconnection network and, to some extent, is independent of the internal structure of the interconnection network.

We have a number of comments. We have expressly developed and applied our technique *so as to make our technique as widely applicable as possible*. As such, the value of n, with regard to some family of interconnection networks $\{X_n\}$, at which a conditional diagnosability result applies can be relatively large (for example, with the hypercubes our method yields that Q_n has conditional diagnosability $3n-5$ when $n \geq 29$ whereas it is known from [13] that Q_n has conditional diagnosability $3n-5$ when $n \geq 5$). If we were to apply our method, and in particular the results from Section 4, *specifically to hypercubes*, so as to utilise the internal structure of hypercubes, then we could get this value of n down considerably (probably even

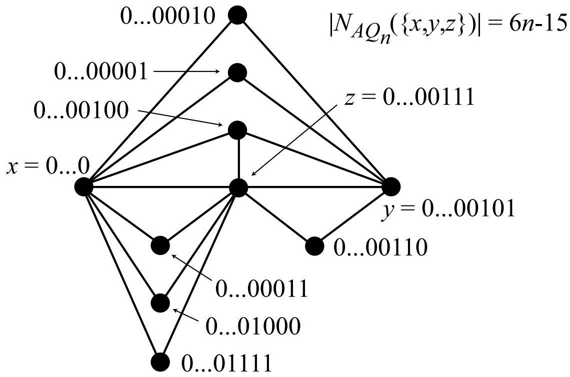

Fig. 11. The first sub-case of Case (a)(iii)(2) where $s = n - 2$ and $b = n - 1$.

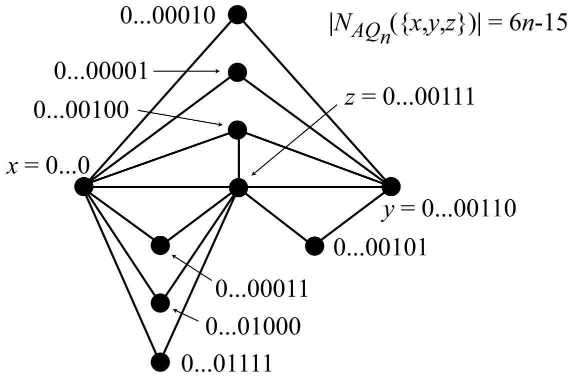

Fig. 12. The second sub-case of Case (a)(iii)(2) where $s = n - 2$ and $b = n$.

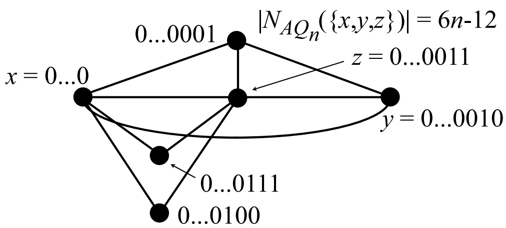

Fig. 13. Case (a)(iii)(3) where $s = n - 1$ and $b = n$.

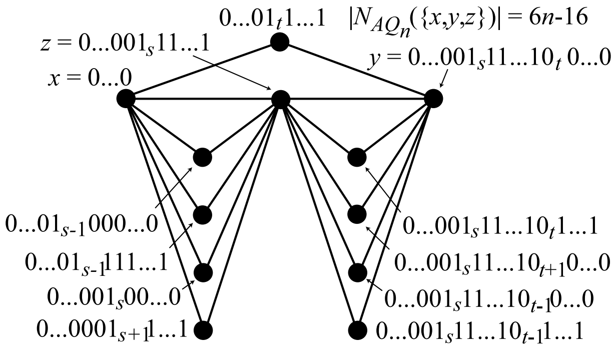

Fig. 14. Case (b)(1) where $t - s \geq 3$.

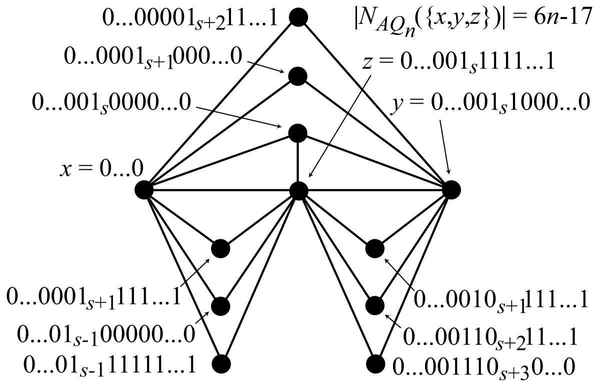

Fig. 15. Case (b)(2) where $t - s = 2$.

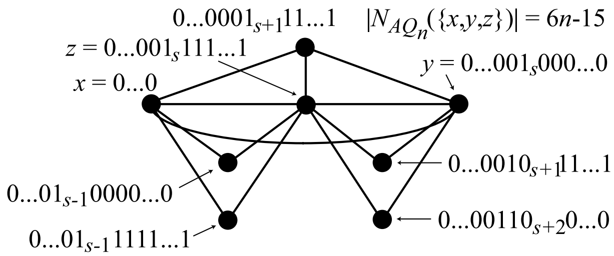

Fig. 16. Case (b)(3) where $t - s = 1$.

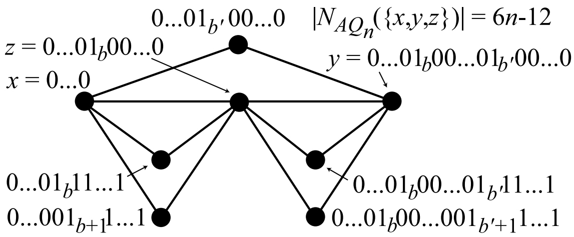

Fig. 17. The first sub-case of Case (c)(1) where $b' - b \geq 2$ and $b' \neq n$.

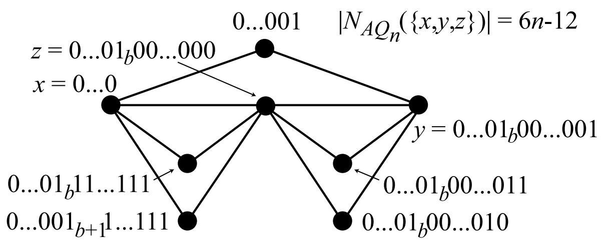

Fig. 18. The second sub-case of Case (c)(1) where $b' - b \geq 3$ and $b' = n$.

to 5). This same comment can be made as regards other interconnection networks to which we apply our methods, and, naturally, we would like to reduce the values of n for which our conditional diagnosability results apply in the cases of folded hypercubes, pancake graphs and augmented cubes. We envisage that we will quite easily be able to do this but leave this to the future, given that the focus in this paper is on establishing our general technique and its efficacy.

We feel that we have just touched the tip of the iceberg as regards the application of our technique, in that we conjecture that it is much more widely applicable than we have shown here (future research will verify this claim). However, as the situation with the augmented cubes denotes, the application of our technique is not always straightforward. What the results in this paper have shown is that there are combinatorial properties of interconnection networks G that are worthy of more study, notably $p_2(G)$.

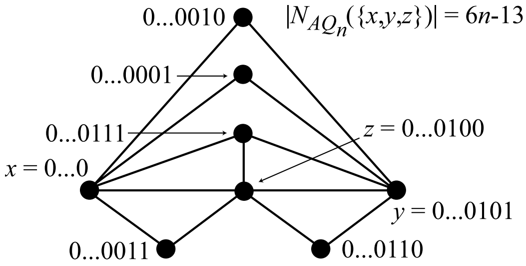

Fig. 19. The third sub-case of Case (c)(1) where $b' - b = 2$ and $b' = n$.

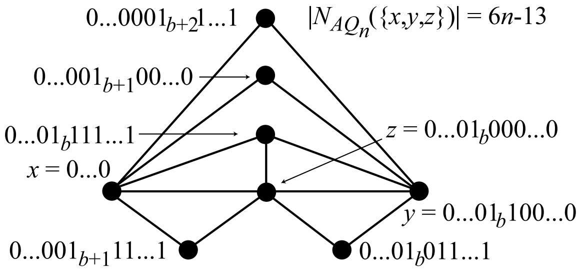

Fig. 20. Case (c)(2) where $b' - b = 1$ and $b' \leq n - 1$.

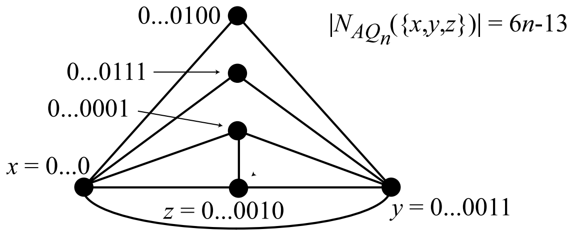

Fig. 21. Case (c)(2) where $b' - b = 1$ and $b' = n$.

Finally, we also feel that a general method, analogous to that here, can be developed in other diagnostic scenarios, notably as regards conditional diagnosis in the PMC model and also (non-conditional) diagnosis in both the PMC and comparison diagnosis models. Again, this claim will be studied in future.

References

[1] S.B. Akers, B. Krishnamurthy, A group-theoretic model for symmetric interconnection networks, IEEE Transactions on Computers 38 (4) (1989) 555–566.
[2] Y.A. Ashir, I.A. Stewart, On embedding cycles in k-ary n-cubes, Parallel Processing Letters 7 (1) (1997) 49–55.
[3] N.-W. Chang, S.-Y. Hsieh, Conditional diagnosability of augmented cubes under the PMC model, IEEE Transactions on Dependable and Secure Computing 9 (1) (2012) 46–60.
[4] S.A. Choudum, V. Sunitha, Augmented cubes, Networks 40 (2) (2002) 71–84.
[5] W.J. Dally, Principles and Practices of Interconnection Networks, Morgan Kaufmann, 2004.
[6] J. Duato, S. Yalamanchili, L.M. Ni, Interconnection Networks: An Engineering Approach, IEEE Computer Society Press, 1997.
[7] A. El-Amawy, S. Latifi, Properties and performance of folded hypercubes, IEEE Transactions on Parallel and Distributed Systems 2 (1) (1991) 31–42.
[8] J. Fan, Diagnosability of crossed cubes under the comparison diagnosis model, IEEE Transactions on Parallel and Distributed Systems 13 (10) (2002) 1099–1104.
[9] S.-Y. Hsieh, Y.-S. Chen, Strongly diagnosable product networks under the comparison diagnosis model, IEEE Transactions on Computers 57 (6) (2008) 721–732.
[10] S.-Y. Hsieh, Y.-S. Chen, Strongly diagnosable systems under the comparison diagnosis model, IEEE Transactions on Computers 57 (12) (2008) 1720–1725.
[11] S.-Y. Hsieh, T.-Y. Chuang, The strong diagnosability of regular networks and product networks under the PMC model, IEEE Transactions on Parallel and Distributed Systems 20 (3) (2009) 367–378.
[12] S.-Y. Hsieh, C.-Y. Kao, Determining the conditional diagnosability of k-ary n-cubes under the MM* model, in: Proceedings of 18th International Symposium on Structural Information and Communication Complexity SIROCCO'11, in: Lecture Notes in Computer Science, vol. 6796, Springer, 2011, pp. 78–88.
[13] G.-H. Hsu, C.-F. Chiang, L.-M. Shih, L.-H. Hsu, J.J.M. Tan, Conditional diagnosability of hypercubes under the comparison diagnosis model, Journal of Systems Architecture 55 (2) (2009) 140–146.

[14] G.-H. Hsu, G.-F. Chiang, J.J.M. Tan, Comparison-based conditional diagnosability on the class of hypercube-like networks, Journal of Interconnection Networks 11 (3–4) (2010) 143–156.
[15] G.-H. Hsu, J.J.M. Tan, Conditional diagnosability of BC networks under the comparison diagnosis model, Processings of International Computer Symposium ICS'08 1 (2008) 269–274.
[16] L.-H. Hsu, C.-K. Lin, Graph Theory and Interconnection Networks, CRC Press, 2008.
[17] P.-L. Lai, J.J.M. Tan, C.-P. Chang, L.-H. Hsu, Conditional diagnosability measures for large multiprocessor systems, IEEE Transactions on Computers 54 (2) (2005) 165–175.
[18] P.-L. Lai, J.J.M. Tan, C.-H. Tsai, L.-H. Hsu, The diagnosability of the matching composition network under the comparison diagnosis model, IEEE Transactions on Computers 53 (8) (2004) 1064–1069.
[19] C.-W. Lee, S.-Y. Hsieh, Diagnosability of two-matching composition networks under the MM^* model, IEEE Transactions on Dependable and Secure Computing 8 (2) (2011) 246–255.
[20] C.-W. Lee, S.-Y. Hsieh, Determining the diagnosability of $(1, 2)$-matching composition networks and its applications, IEEE Transactions on Dependable and Secure Computing 8 (3) (2011) 353–362.
[21] C.-K. Lin, T.-L. Kung, J.J.M. Tan, Conditional-fault diagnosability of multiprocessor systems with an efficient local diagnosis algorithm under the PMC model, IEEE Transactions on Parallel and Distributed Systems 22 (10) (2011) 1669–1680.
[22] C.-K. Lin, J.J.M. Tan, L.-H. Hsu, E. Cheng, L. Liptak, Conditional diagnosability of cayley graphs generated by transposition trees under the comparison diagnosis model, Journal of Interconnection Networks 9 (2008) 83–97.
[23] J. Maeng, M. Malek, A comparison connection assignment for self-diagnosis of multiprocessor systems, in: Proceedings of 11th International Symposium on Fault-Tolerant Computing, 1981, pp. 173–175.
[24] M. Malek, A comparsion connection assignment for diagnosis of multiprocessor systems, in: Proceedings of 7th Annual Symposium on Computer Architecture, ACM Press, 1980, pp. 31–35.
[25] Q.T. Nguyen, S. Bettayeb, The upper bound and lower bound of the genus of pancake graphs, in: Proceedings of 14th IEEE Symposium on Computers and Communications ISCC'09, IEEE Press, 2009, pp. 800–804.
[26] F.P. Preparata, G. Metze, R.T. Chien, On the connection assignment problem of diagnosable systems, IEEE Transactions on Computers EC-16 (1967) 448–454.
[27] A. Sengupta, A.T. Dahbura, On self-diagnosis multiprocessor systems: diagnosis by the comparison approach, IEEE Transactions on Computers 41 (11) (1992) 1386–1396.
[28] D. Wang, Diagnosability of hypercubes and enhanced hypercubes under the comparison diagnosis model, IEEE Transactions on Computers 48 (12) (1999) 1369–1374.
[29] J. Xu, Topological Structure and Analysis of Interconnection Networks, Springer, 2001.
[30] M. Xu, K. Thulasiraman, X.-D. Hu, Conditional diagnosability of matching composition networks under the PMC model, IEEE Transactions on Circuits and Systems II: Express Briefs 56 (11) (2009) 875–879.
[31] E. Yang, X. Yang, Q. Dong, J. Li, Conditional diagnosability of hypermeshes under the comparison model, Information Processing Letters 111 (4) (2011) 188–193.
[32] J. Zheng, S. Latifi, E. Regentova, K. Luo, X. Wu, Diagnosability of star graphs under the comparison diagnosis model, Information Processing Letters 93 (1) (2005) 29–36.
[33] S. Zhou, The conditional diagnosability of hypercubes under the comparison model, in: Proceedings of WASE Conference on Information Engineering ICIE'09, IEEE Press, 2009, pp. 459–462.
[34] S. Zhou, The conditional diagnosability of Möbius cubes under the comparison model, in: Proceedings of International Conference on Information and Automation, IEEE Press, 2009, pp. 96–100.
[35] S. Zhou, The conditional diagnosability of twisted cubes under the comparison model, in: Proceedings of International Symposium on Parallel and Distributed Processing with Applications ISPA'09, IEEE Press, 2009, pp. 696–701.
[36] S. Zhou, W. Xiao, Conditional diagnosability of alternating group networks, Information Processing Letters 110 (10) (2010) 403–409.
[37] Q. Zhu, S.Y. Liu, M. Xu, On conditional diagnosability of the folded hypercubes, Information Sciences 178 (4) (2008) 1069–1077.