A general algorithm for detecting faults under the Comparison Diagnosis Model

Iain Stewart

Durham University, U.K.

Contents

Basic issues relating to diagnosability

Given: an interconnection network (graph) in which there are some faulty processors (nodes) that we must find.

We can perform tests:

- a node u tests 2 of its neighbours v and w by sending a message to both neighbours
- if u is healthy and v and w respond identically then the syndrome table entry $s_u(v,w) = 0$
- if u is healthy and v and w respond differently then $s_u(v,w) = 1$
- if u is faulty then $s_u(v,w) \in \{0,1\}$ arbitrarily

Assumptions :

- all faults are permanent
- a faulty node always produces an incorrect (message) response
- no 2 faulty nodes ever produce identical incorrect (message) responses.

This model is known as the Comparison Diagnosis Model [MM81].
There are other models, notably the PMC model [PMC67], where individual neighbours are tested.

(86 citations on Google Scholar)

[MM81] J. Maeng, M. Malek, Proc. 7th Ann. Symp. Computer Architecture
[PMC67] F. Preparata, G. Metze, R. Chien, IEEE Trans. Electron. Comput.

A fault distribution $\rightsquigarrow$ a set of syndromes.

A syndrome $\rightsquigarrow$ a set of fault distributions.

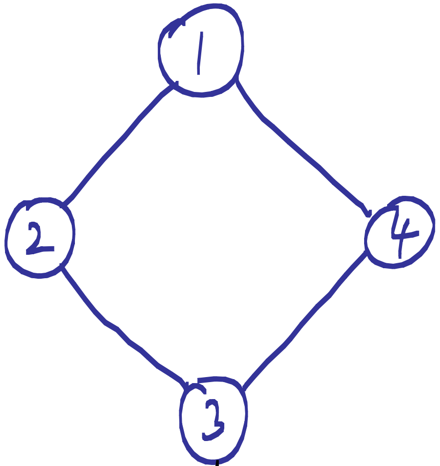

graph

s	(1,3)	(2,4)
1		0
2	0	
3		1
4	1	

s	(1,3)	(2,4)
1		1
2	0	
3		1
4	0	

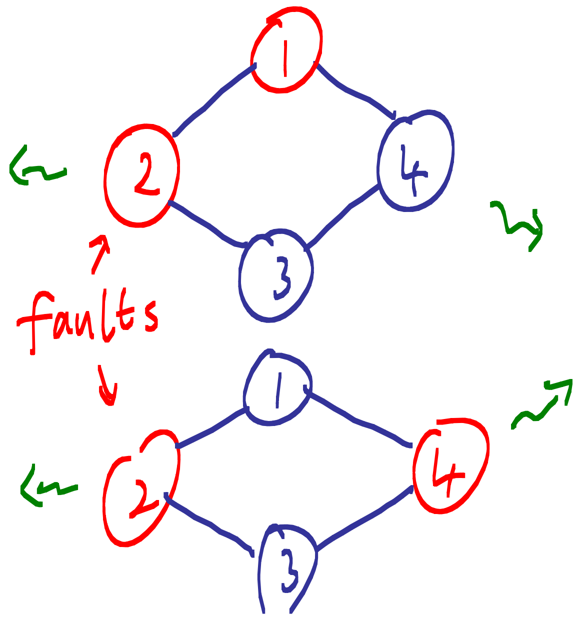

s	(1,3)	(2,4)
1		1
2	1	
3		1
4	1	

Fundamental Question 1

Given a graph G, what is the maximal number t so that given any syndrome resulting from a fault distribution of $\leq t$ faults, there is exactly one such fault distribution corresponding to the syndrome?

The number t is the diagnosability of G.

Fundamental Question 2

If G has diagnosability t, what is the complexity of finding the fault distribution corresponding to a syndrome resulting from a fault distribution of $\leq t$ faults?

Some basic results relating to diagnosability

There is a simple upper bound on the diagnosability of G.

Let u be a node of G of degree d.

Fault distribution A

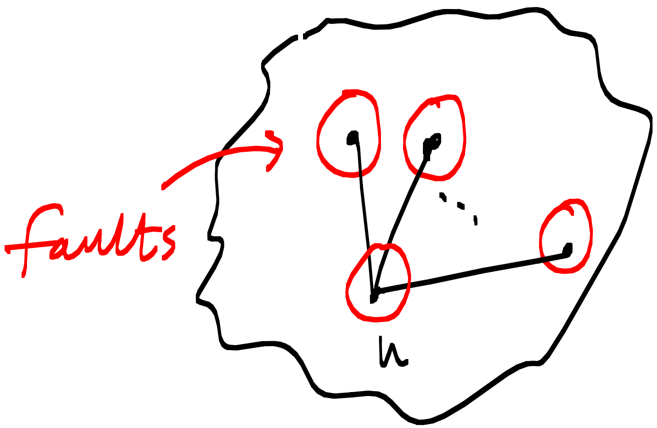

these two fault distributions clearly have a common syndrome: so

Fault distribution B

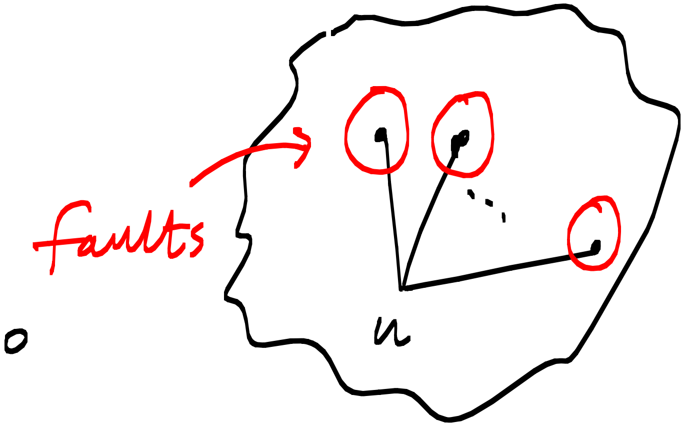

$$\text{diagnosability}(G) \leq \text{min. node degree in } G$$

Theorem [SD92]

A graph $G=(V,E)$ with $|V|=n$ has diagnosability at least t iff

(a) $n \geq 2t+1$

(b) each node has order at least t

(c) for each $X \subseteq V$ s.t. $|X| = n-2t+p$, where $0 \leq p \leq t-1$, the graph G_X does not have a spanning subgraph of $K_{t-p,t-p}$ as a node-induced subgraph.

[SD92] A. Sengupta, A.T. Dahbura, IEEE Trans. Computers

Some more workable sufficient conditions:

Theorem [CLTH04]

If $G=(V,E)$ is a regular graph of degree d, has connectivity d, and $|V| \geq 2d+3$, then G has diagnosability d.

Theorem [CT09]

If $G=(V,E)$ is a graph where every node is the root of an extended star with t branches then G has diagnosability $\geq t$.

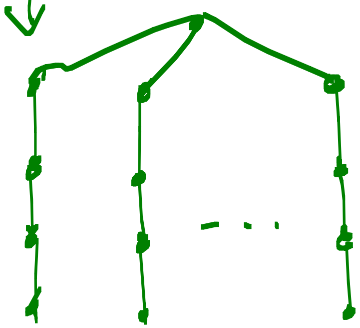

[CLTH04] C.-P. Chang, L.-P. Lai, J.J.-M. Tan, L.-H. Hsu, IEEE Trans. Comput.
[CT09] C.-F. Chang, J.J.-M. Tan, IEEE Trans. Computers

Algorithms solving the fault diagnosis problem

Theorem [SD92]

Suppose that $G=(V,E)$ is a graph of diagnosability t and $|V|=n$. There is an algorithm of time complexity $O(n^5)$ that given a syndrome corresponding to a fault distribution of $\leq t$ faults outputs the unique corresponding fault distribution.

Theorem [YT07]

Further, if G has minimal (resp. maximal) degree δ (resp. Δ) then the time complexity can be improved to $O(\delta \cdot \Delta^3 n)$.

[YT07] X. Tang, Y. Y. Tang, IEEE Trans. Computers

Theorem [Y03]*

The fault diagnosis problem in a hypercube Q_n, where the diagnosability is n, can be solved in $O(n^2 \cdot 2^n)$ time.

Theorem [CT09]*

If $G=(V,E)$ is a graph where every node is the root of an extended star with t branches then G has diagnosability $\geq t$ and the fault diagnosis problem can be solved in $O(\Delta N)$ time, where $|V|=N$ and G has maximal degree Δ.

[Y03] X. Yang, Proc. 12th Asian TCA Symp.

* Motivation for this work

Chiang and Tan's algorithm

Consider some node u:

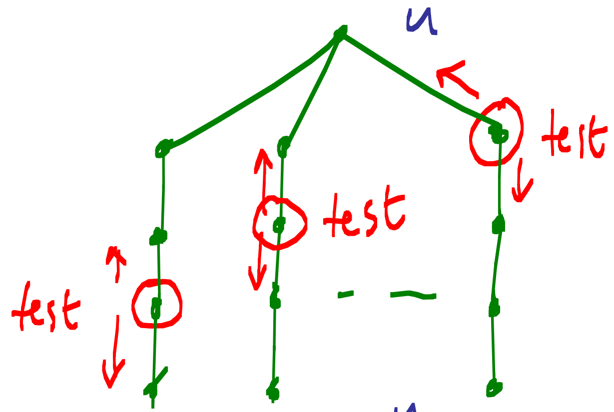

look at the test results induced by the nodes of the extended star.

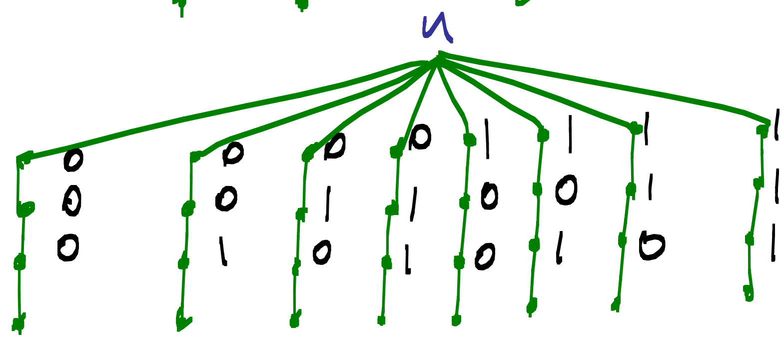

Each branch can be in one of 8 states.

Let R_0 = #branches in state

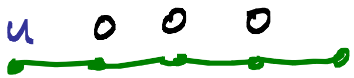

Let R_1 = #branches in state

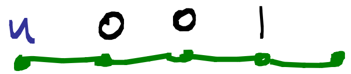

and so on.

So, $\sum_{i=0}^{7} |R_i| = t$. <u>Claim</u> u is healthy iff $|R_0| \geq |R_4|$.

$\Rightarrow$ "local" algorithm for diagnosability

Chiang and Tan's approach yields an $O(\Delta N)$ time algorithm.

E.g., every node of a hypercube Q_n is the root of an extended star of n branches $\Rightarrow$ $O(n \cdot 2^n)$ time algorithm

(better than Yang's $O(n^2 \cdot 2^n)$ time algorithm).

E.g., every node of a star graph is the root of an extended star of n branches $\Rightarrow$ $O(n \cdot n!)$ time algorithm.

Drawbacks of Chiang and Tan's approach:

- existence of extended stars need to be verified
- the algorithm uses an explicit representation of extended stars.

+ other (minor) drawbacks.

Yang's approach in hypercubes

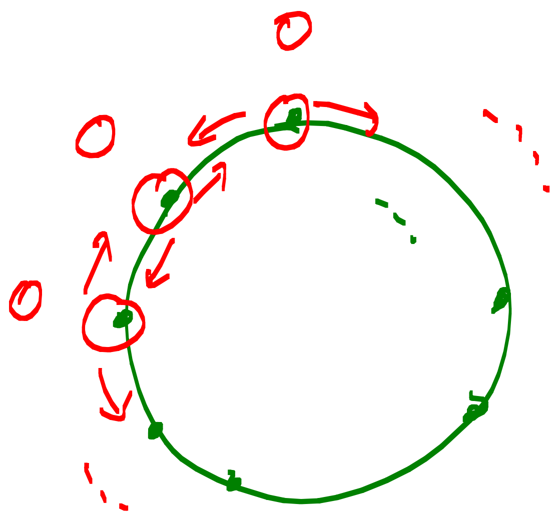

Suppose we have a "long cycle of 0's".
If the length of the cycle is > diagnosability then all nodes are healthy.

- "Decompose" Q_n into smaller hypercubes "joined" in the "shape" of a hypercube.
- Smaller hypercubes are Hamiltonian.
- There must exist some "healthy" sub-cubes.
- Use these healthy sub-cubes to test the rest of Q_n.

Our approach

- There is nothing special about Hamiltonian cycles: any sufficiently large connected component will do.
- So long as our graph can be partitioned into "sufficiently many" "sufficiently large" connected components, this will do.
- We start at a node, assume it is healthy and build up a component of nodes that are healthy, assuming our original node is healthy.

all test results 0

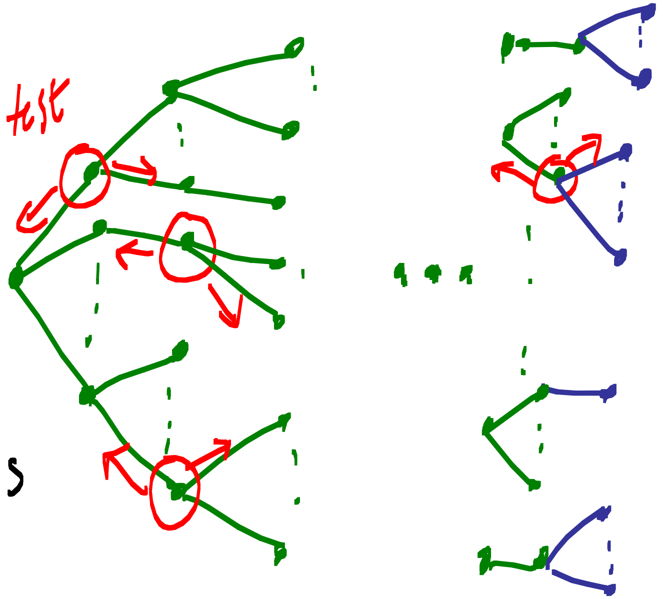

Start from u and build tree of "healthy" nodes

border of "faulty" nodes

The border is an articulation set or the remaining nodes of G

If the diagnosability of $G \leq$ connectivity of G then the border consists of all faulty nodes.

All we have to ensure is that we start at an appropriate node u.

Theorem

Let $G=(V,E)$ be a graph where the diagnosability of G is at most the connectivity of G and where G can be partitioned into "enough connected components of large enough size" then the fault diagnosability problem can be solved for G in $O(\Delta|V|)$ time, where Δ is the maximal degree of G.

Immediately applicable to:

hypercubes, crossed cubes, twisted cubes, folded hypercubes, enhanced hypercubes, augmented cubes, shuffle cubes, twisted N-cubes, k-ary n-cubes, (n,k)-stars, pancake graphs, ...

Conclusions

- Our time complexity matches that of Chiang and Tan's algorithm but our conditions for application are
  - less severe
  - do not impact upon the algorithm construction
- Our algorithm is conceptually simpler.
- Our algorithm can yield a sub-tree spanning all healthy nodes.
- Unlike Chiang and Tan's algorithm, we don't need to build the whole syndrome table or re-use existing entries.

Future research

So far, the focus of fault diagnosis has been on centralized algorithms.

No attention has been paid to distributed algorithms (even though the syndrome table is constructed in a distributed fashion).

Preliminary results show that a distributed version of our algorithm will be a dramatic improvement on distributed versions of existing algorithms.