

Diagnosis in Distributed Systems

A Presentation by
Gunjan Khanna
Padma Varadharajan

Diagnosis

- Important to be able to figure out which component has failed
- In order to initiate recovery
- Multiple processes running and thousands of nodes.
 - Can a centralized algorithm work

System Diagnosis

- To figure out which component has failed if any.
- Based on testing by each node in most algorithms.
 - The test results are called *Syndromes*.
 - A centralized tester figures out the results of the syndromes using a graph.
 - $N > (2t+1)$ where t is the number of faulty receivers needed to be diagnosed.
- Probabilistic diagnosis and distributed diagnosis.

Group Membership Problem

- Trying to figure out whether your neighbor is functioning properly or not.
 - Keeping track of which processes have failed and which are functioning properly
- Varies from loose synchrony to virtual synchrony
- Broadcast based coordinator based and token ring based.

Poles Apart or R They ?

Membership Algorithm	System Diagnosis
Fail-Stop and crash failures	Failures detected by the test given
Testing might not be active	Testing is always active
They can tolerate any number of failures	They usually tolerate only a fixed number of failures like the PMC model
False alarms	They provide a guarantee with identification
Integrated with the application protocol	Not necessarily but run independently

NEW_SELF System Diagnosis algorithm

- Node P_i tests its neighbors and keeps it in list $TESTED_BY(P_i)$.
- It receives the diagnostic information from all the fault-free nodes stores it.
- It retests all the nodes and verifies the information and validates it
- This information is forwarded to all nodes that test P_i .



Distributed Online Diagnosis

- Algorithm which guarantees 'correct' diagnosis.
- Under limited model (consistent liars), diagnosis is also complete.
- General case of unrestricted arbitrary faults also discussed.
- Technique shall identify node failures that occur during diagnosis algorithm execution
- Faults are manifested as corrupted diagnostic information maintained at node or exchanged between nodes.

Setup


- Collection of distributed nodes, $V(S)$ interconnected by bidirectional edges $E(S)$
- Each node assigned fault state s –
 - 0 – Fault free
 - 1 – Faulty
- Set of performed tests – Testing assignment
- Collection of corresponding test results- Syndrome
- Effect a distributed diagnosis.

Node Validation Mechanism




Features of Algorithm

- If N_x validates N_a , represented as



$a.Syndromes[0] = 1$
 $a.Syndromes[1] = 2$
 $a.Syndromes[2] = 4$
 $a.Syndromes[3] = -$
 $a.Syndromes[4] = 5$
 $a.Syndromes[5] = 0$



(a) Example Syndromes Array (b) Validation Graph

Algorithm for 1 consistent Liar

- Nodes are listed in sequential order.
- Node N_x requests syndrome information from next two nodes in ordered node list.
- If N_y does not validate with N_z , N_x requests information from next node, N_w .
- N_x identifies first node that validates with N_w as fault free
- Diagnostic information received from validated node is used to update local Syndrome array.

Example

n_1	n_2	n_3					
0 <table border="1"><tr><td>1</td></tr></table>	1	0 <table border="1"><tr><td>1</td></tr></table>	1	0 <table border="1"><tr><td>1</td></tr></table>	1	0 <table border="1"><tr><td>2</td></tr></table>	2
1							
1							
1							
2							
1 <table border="1"><tr><td>2</td></tr></table>	2	1 <table border="1"><tr><td>2</td></tr></table>	2	1 <table border="1"><tr><td>2</td></tr></table>	2	1 <table border="1"><tr><td>2</td></tr></table>	2
2							
2							
2							
2							
2 <table border="1"><tr><td>3</td></tr></table>	3	2 <table border="1"><tr><td>3</td></tr></table>	3	2 <table border="1"><tr><td>3</td></tr></table>	3	2 <table border="1"><tr><td>3</td></tr></table>	3
3							
3							
3							
3							
3 <table border="1"><tr><td>4</td></tr></table>	4	3 <table border="1"><tr><td>4</td></tr></table>	4	3 <table border="1"><tr><td>4</td></tr></table>	4	3 <table border="1"><tr><td>4</td></tr></table>	4
4							
4							
4							
4							
4 <table border="1"><tr><td>5</td></tr></table>	5	4 <table border="1"><tr><td>5</td></tr></table>	5	4 <table border="1"><tr><td>5</td></tr></table>	5	4 <table border="1"><tr><td>5</td></tr></table>	5
5							
5							
5							
5							
5 <table border="1"><tr><td>0</td></tr></table>	0	5 <table border="1"><tr><td>0</td></tr></table>	0	5 <table border="1"><tr><td>0</td></tr></table>	0	5 <table border="1"><tr><td>0</td></tr></table>	0
0							
0							
0							
0							

(a) Syndromes arrays received by n_0

(b) n_0 's final Syndromes array

Assumptions:


- 1) Every message transmitted by fault-free node is received correctly.
- 2) Receiver of message can identify sender of message.
- 3) Absence of message can be detected.
- 4) Every node can communicate with every other node
- 5) Faulty node distributes incorrect information to all nodes that requests information

Working of Algorithm


- Every fault-free node correctly identifies the nearest fault free node in the ordered node list.
- After one round of the algorithm, a validation path exists between any pair of fault free nodes.
- Syndrome entries corresponding to fault free nodes are identical in all fault free nodes after a fixed no. of rounds.

Extensions

- For t consistent liars where common-mode failure is possible, a validation path of length t containing $t+1$ nodes is found, to ensure that the first node in the path is fault free.
- For inconsistent liars, algorithm provides correct, but not complete diagnosis.



One Inconsistent Liar



Pros and Cons of Algorithm

- Absence of centralized supervisor-controlled diagnosis.
- No assumption made about fault free node being able to accurately project state of node it is testing.
- Increased overhead owing to syndromes being passed irrespective of whether changes have been effected.

Fault Identification using Finite State Machine Model