

Spare Allocation and Reconfiguration in a Fault Tolerant Hypercube with Direct Connect Capability

Baback A. Izadi* and Füsün Özgüner
 Department of Electrical Engineering
 The Ohio State University
 Columbus, OH 43210

Abstract

This paper investigates hardware reconfiguration schemes to make the hypercube multicomputer fault tolerant. Two schemes are proposed; the Cluster Approach and the Enhanced Cluster Approach. The approaches are shown to be able to tolerate large number of failures without any performance degradation. It is further demonstrated that no modification to either the existing communication or computational algorithm is needed. Finally a gracefully degradable approach is presented to reconfigure when the number of faulty nodes are more than the available spares.

1 Introduction

For fault tolerance in hypercube multiprocessors, two fundamental approaches have been proposed. Some researchers have examined the software approach [1, 2, 3] with no added hardware and some performance degradation. Others have looked into hardware schemes where spare nodes and/or links are used to replace the faulty ones [4, 5, 6, 3, 7, 8, 9].

This paper investigates hardware schemes that allocate spares in a hypercube multiprocessor to make it fault tolerant. A hardware scheme was first proposed by Rennels [4]. Chau and Liestman [5] have employed a decoupling network to connect modules of active and spare processors. Several approaches have been proposed in [6, 3, 8] that perform reconfiguration using spare nodes and spare links. An augmented approach was illustrated by Witkowski and Lee [7]. Alam and Melhem [9] using hardware redundancy, developed augmented approaches to tolerate faulty nodes.

The schemes presented in this paper tolerate a large number of faults without any performance degradation and the resultant configuration does not affect either the communication or computational algorithms already developed for the hypercube multiprocessor. The allocation of spares is facilitated by using a routing element on each node similar in concept to the *Direct Connect Module* [10] of *Intel* hypercubes.

Each node of a d dimensional hypercube is represented

*Also with the Devry Institute of Technology, Columbus, Ohio 43209.




Figure 1: Hypercube of dimension 4 with clusters of dimension 2

by d -tuple $(b_{d-1} \cdots b_i \cdots b_0)$, where $b_i \in \{0, 1\}$. A subcube is defined by a unique d -tuple $\{0, 1, X\}^d$ where "0" and "1" are the *bound* coordinates, and "X" represents the *free* coordinates. A $(d-k)$ -dimensional subcube is represented by a d -tuple with k *bound* coordinates and $d-k$ *free* coordinates. Each spare is uniquely defined by d -tuple $\{0, 1, S\}^d$ where "0" and "1" represent the bound coordinates and S corresponds to a free coordinate of its assigned cluster. A link is specified uniquely by d -tuple $\{0, 1, -\}^d$ where "-" can be substituted by "0" and "1" to identify its connecting nodes. An intra-cluster spare link is defined by d -tuple $\{0, 1, S, -\}^d$ where S signifies spare and "-" can be substituted by "0" and "1" to identify its connecting spare nodes. An inter-cluster spare link is defined by $(d+1)$ -tuple $\{0, 1, S\}^{(d+1)}$ where S signifies spare and ("0", "1") are the *bound* coordinates of the regular node that its respective spare node is connected to.

2 Cluster Approach

An i dimensional cluster is merely a subcube of dimension i . Thus, a hypercube of dimension d can be divided into 2^j clusters where $d = i + j$. A single spare is assigned per cluster of nodes. Figure 1 depicts a hypercube of dimension 4 with clusters of dimension 2. In this scheme, the spare may logically replace the faulty node within its designated cluster. For example, in figure 1, spare 01SS may logically replace any one of nodes 0100, 0110, 0111, 0101 using spare links 01S00, 01S10, 01S11, 01S01 respectively.

An important factor influencing the selection of the cluster size is the mean time to failure of a node. Re-

liability studies [7] with the available components have shown that to attain 95 percent reliability in 10 years, approximately 25 percent spare overhead is needed. Therefore a cluster dimension of 2 should be acceptable.

It is assumed that a given faulty node retains its communication capability. This is a fair assumption since in hypercube multiprocessors such as the *iPSC/2*, the computation and the communication modules of a node are separate. Furthermore, since the complexity of the computation unit is much greater than that of the communication one, the probability of a failure in the computation unit is much higher.

In hypercubes with *Direct Connect Capability* [10, 11], the cost of communication is nearly constant between any two given nodes. To facilitate reconfiguration, the routing elements described below are used at each node. Figure 2 depicts the block diagram of the router for the regular and the spare nodes. The regular node router is an enhanced version of *Intel's Direct Connect Module (DCM)* with an additional channel routing element provided to connect a given node to its respective spare. The spare node router is a simple switch which is used to either link two of the regular nodes within a cluster together or to connect the spare to one of its cluster nodes. The former is used to bypass the faulty link by connecting the two appropriate channel routing elements together. For example in figure 2, connecting the routing elements of nodes 00 and 01 (linking *Node 00 In* to *Node 01 out* and *Node 01 In* to *Node 00 out*) can provide an alternative path to 010—in subcube 01XX of figure 1. The *Spare In*, *Spare Out* of each router of the nodes 0100 and 0101 are then connected to the *Node 00 Out*, *Node 00 In* and *Node 01 Out*, *Node 01 In* of the spare router. The latter, connecting the spare to one of its cluster nodes, is done to tolerate a node failure. This is accomplished by connecting the *Spare Routing Element* of the faulty node (figure 2(a)) to the appropriate *Node Routing Unit* of the spare node. Three cases exist which may involve routing data through the faulty node. The first one is when the message originates from the faulty node. Here the spare sends its data via the serializer (figure 2) out to the appropriate *Node XX Out* channel. The message is then routed to the *DCM* of the faulty node via the *Spare In*. Depending on the final destination node, any of *Channel Outs* may be selected. The second case is when the destination of the message is the faulty node. Once the data reaches the *DCM* of the faulty node, it is automatically routed to the channel *Spare Out* instead of channel *To Node*. The spare's *DCM* would consequently receive the message using the appropriate channel (*Node XX In*). It then deserializes the message and sends it to the spare. The third case is when the faulty node is neither the source nor the destination of the message but is used merely as an intermediate hop. In this case, the spare is not involved at all, and routing is done normally. Note that the connection of the spare router to the active node(s) is established after the reconfiguration and remains intact thereafter.

Upon detecting a node failure, the spare within the re-




Figure 2: Communication Module Architecture for (a) Regular Node (b) Spare Node




Figure 3: Reconfiguration of a faulty hypercube

spective cluster logically replaces the faulty node. The spare consequently activates its link to the faulty node and disable the rest of its links. The spare’s communication module then forwards/receives its data to/from other nodes via the *DCM* of the faulty node as discussed above.

One intra-cluster link failure per cluster can be tolerated. Upon detection of a link failure, the router of the spare node is used to establish a parallel path to the faulty link. The *Processing Elements* at the two ends of the faulty link would logically replace their *channel routing elements*, which connects them to the faulty link (figure 2), with the *spare channel routing element*. An inter-cluster link failure is fatal.

Example: Figure 3 illustrates the reconfiguration of hypercube upon detection of faults in nodes 0111, 1000 and links 111–,000–. The *Ch0 routing element* of nodes 0000, 0001, 1110, and 1111 are replaced with the *spare channel routing element*. The faulty nodes are replaced by their designated spare as discussed before. If there exists a need to establish a link between node 1111 and 1000, the path (1111 → 11SS → 1110 → 1100 → 1000 → 10SS) is set up. Note that no modification of the communication algorithm is needed.

3 Enhanced Cluster Approach

The *Cluster Approach* has two shortcomings. It does not allow multiple faults within a cluster and it fails to reconfigure upon an inter-cluster link failure. To overcome these deficiencies, in this scheme, the spares are connected in such a way to form a $d - 2$ cube. Figure 4 depicts an *Enhanced Cluster Hypercube (ECH)* of dimension 4. The same figure also demonstrates the scalability of this topology. Each regular and spare node is connected to $d + 1$ and $4 + (d - 2) = d + 2$ nodes respectively. Intra-cluster and inter-cluster spare links are uniquely defined by $(d + 1)$ -tuple and d -tuple respectively. For example, the spare link connecting the spare node 00SS and the node 0001 is identified as 00S01. Also the inter-cluster spare link connecting the spare 01SS and 00SS is labeled 0 – SS. The block diagram of the router for the regular node is the same as before (figure 2(a)). The router for the spare node needs $d + 2$ routing elements instead of 4 elements as shown in figure 2(b). The spares are provided with the *simultaneous multi-channel communication capability*. More specifically, each spare at the same time can connect $\lfloor \frac{d}{2} \rfloor$ pairs of its links together.




Figure 4: ECH of dimension 4




Figure 5: Reconfiguration of ECH

The *ECH* can tolerate up to 3 faulty nodes per cluster. Reconfiguration is accomplished using the following algorithm. The algorithm first checks whether the number of spares is sufficient for reconfiguration and if there exists a cluster with all four nodes being faulty. Next, a single faulty cell per faulty cluster of dimension 2 is selected at random and its task is assigned to the respective spare. If all faults are covered, then reconfiguration is accomplished. If there are still faulty nodes which have no spare assignment, the free spares in the next higher dimension are considered. The process continues until all faulty nodes are covered by the spares.

The *Enhanced Cluster Hypercube* can tolerate both intra-cluster and inter-cluster link failures. Parallel path(s) to the faulty link(s) are established. All other unnecessary spare links are disabled. Reconfiguration fails if any of the two nodes at the end of a faulty link is faulty too. It also fails if two faulty links have a common node.

Example: Figure 5 demonstrates the reconfiguration of the *Enhanced Cluster Hypercube* upon detection of faults in links 0-00, 100- and nodes 1011, 1100, 1110. *Simultaneous multi-channel communication capability* of the spare node 10SS is used to logically replace node 1011 and at the same time establish a path between nodes 1001 and 1000. Inter-cluster link failure, 0-00, requires the use of two spare nodes 01SS and 00SS to be replaced with 01S00 → 0 – SS → 00S00. The spare node 01SS logically replaces the faulty node 1100 using the path –1SS → 11S00. Similarly 11SS replaces 1110 using the link 11S10.

4 Gracefully Degradable Approach

As part of a comprehensive solution to a fault-tolerant hypercube, the possibility of more faulty nodes than available spares needs to be considered. In such a case,

